Seventh International Conference on Computational Fluid Dynamics (ICCFD7), Big Island, Hawaii, July 9-13, 2012

ICCFD7-1003

Improved Finite-Volume Method for Radiative Hydrodynamics

Alan Wray*
Corresponding author: alan.a.wray@nasa.gov
* NASA Ames Research Center, USA

Abstract: Fully coupled simulations of hydrodynamics and radiative transfer are essential to a number of fields ranging from astrophysics to engineering applications. Of particular interest in this work are hypersonic atmospheric entries and associated experimental apparatus, e.g., shock tubes and high enthalpy testing facilities. The radiative transfer calculations must supply to the CFD a heating term in the energy equation in the form of the divergence of the radiative heat flux and the radiative heat fluxes to bounding surfaces. It is most efficient to solve the radiative transfer equation on the same grid as the CFD solution, and this work presents an algorithm with improved accuracy compared to more conventional approaches on both structured and unstructured grids. Results are shown for a number of idealized test problems and for shock radiation during hypersonic entry. Issues of parallelization within a radiation sweep are also discussed.

Keywords: Numerical Algorithms, Computational Fluid Dynamics, Radiative Transfer

1 Introduction

There are several areas of application of CFD in which radiative heat transfer processes are significant relative to other energy-transfer mechanisms such as convection and conduction. Two such areas of particular interest to NASA are spacecraft reentry and solar and stellar structure and dynamics. For the former, simulations both of actual reentry events and of the earthbound experiments intended to gather data about reentry physics often require radiative heat transfer as a component. In stars radiative processes are significant both as a local energy transfer mechanism in the deep layers and as the principal energy-loss process in the outermost part of the body of the star, as well as the mechanism which creates most of the observable data about stars, namely their electromagnetic spectra.

It is therefore necessary to have efficient and effective algorithms for simulating coupled radiation and fluid (including plasma) systems. Efficiency is especially significant because radiation in general introduces *three* more dimensions into the problem: two giving the direction of propagation and one for the frequency or wavelength. In addition, the physical properties of radiation are qualitatively distinct from those of fluids in that radiation typically travels in essentially straight lines independent of the fluid motion and is generated and absorbed by mechanisms that depend in a complex way on the local fluid properties. A large number of algorithms have been developed to simulate radiative processes, even when only considering methods applicable to multidimensional, combined fluid-radiative systems. To name a few, there are spherical harmonic methods, discrete ordinate methods, and long- and short-characteristic schemes. We will concentrate here on methods that have the most general applicability and that are most easily combined with CFD, namely the finite volume methods. For simplicity of presentation of the mathematical development, we will restrict ourselves to non-scattering media and will not explicitly display wavelength variation in the equations. The techniques developed here are, however,

fully applicable with full spectral variation and with scattering. The results of section 3.4 utilize a full spectral analysis.

Finite volume methods have the significant advantage that they can be carried out using the same grid as the fluid mechanics, whether structured or unstructured in organization. This grid commonality greatly simplifies computation and bookkeeping of radiation data, especially in parallel implementations. But, of course, questions of sufficiency of the resolution arise since the CFD grid is presumably tailored for fluid physics rather than for radiation. So to avoid excessive, radiation-driven grid refinement or expensive and diffusive interpolations onto secondary grids, it is important to try to maximize the accuracy obtainable on a given grid for the radiation simulation. This is the goal of the present effort.

The radiative transfer equation (RTE) for the radiative intensity I in direction $\mathbf{\Omega}$ in a non-scattering medium is given by

$$\Omega \cdot \nabla I(\mathbf{x},\Omega) = \chi(\mathbf{x})\Big(S(\mathbf{x}) - I(\mathbf{x},\Omega)\Big) \tag{1.1}$$

where S is the radiative source function and χ is the opacity. Frequency dependence has been suppressed for simplicity of notation and presentation. Upon integration of (1.1) over a computational cell we obtain

$$\int_{\text{cell surface}} I(\mathbf{x},\Omega)\Omega \cdot d\mathbf{S} = \int_{\text{cell volume}} \chi(\mathbf{x})\Big(S(\mathbf{x}) - I(\mathbf{x},\Omega)\Big)dV \tag{1.2}$$

Discretizing (1.2) by approximating the volume integral with the value of the integrand at the center times the volume and the surface integral by a summation over the faces of face-centered values times the corresponding areas, one obtains

$$\sum_{\text{faces } k} I(\mathbf{x}_k,\Omega)\Omega \cdot \Delta\mathbf{S}_k = \chi(\mathbf{x}_\text{c})\Big(S(\mathbf{x}_\text{c}) - I(\mathbf{x}_\text{c},\Omega)\Big)V \tag{1.3}$$

where $\mathbf{x}_\text{c}$ is the position of the cell center, $\mathbf{x}_k$ is the center of face k, V is the cell volume, and $\Delta\mathbf{S}_k$ is the outward-pointing surface area vector of face k of the cell. Splitting the surface summation into incoming and outgoing parts and using cell-center values for the outgoing facial intensities in each cell, one gets the standard finite-volume method for radiative transfer [1]:

$$I(\mathbf{x}_\text{c},\Omega) = \frac{\displaystyle\sum_{k,\Omega\cdot\Delta\mathbf{S}_k<0} I(\mathbf{x}_k,\Omega)\left|\Omega \cdot \Delta\mathbf{S}_k\right| + \chi_\text{c} S_\text{c} V}{\displaystyle\sum_{k,\Omega\cdot\Delta\mathbf{S}_k>0} \Omega \cdot \Delta\mathbf{S}_k + \chi_\text{c} V} \tag{1.4}$$

where

$$S_\text{c} \equiv S(\mathbf{x}_\text{c}), \quad \chi_\text{c} \equiv \chi(\mathbf{x}_\text{c}) \tag{1.5}$$

2 Improved formulation

Three approximations are made in passing from (1.2) to (1.4): (a) the surface integrals over each face are approximated by multiplying the face-center value by the corresponding area; (b) the volume integral is approximated by multiplying the cell-center values by the volume; and (c) the outgoing intensity values

are approximated by the cell-center value. Approximation (a) is difficult to improve upon since a better formulation would require carrying or computing information about the spatial distribution of the radiation over a face, resulting in a substantial increase in complexity; we therefore continue to make this approximation. If this approximation turns out to be a poor one in a given simulation, then greater spatial resolution is likely to be the best solution. Approximations (b) and (c) become weak when the source S is rapidly varying or when the opacity χ is large (i.e., in an optically thick medium). These are both common situations, so improvement in this area may be valuable. To cite a particular example, the CFD grid may be designed to resolve temperature variations so that the temperature T does not vary too rapidly across a cell. But the radiative source S frequently varies as a power of T, e.g. as T^4 for grey media, for which substantial change across a cell may still be present.

In order to improve on approximations (b) and (c), we note that the radiation crossing a cell starts by entering at those faces with inward normals in the radiation direction $\mathbf{\Omega}$ and then travels across the cell to the faces with outward normals. It is this outward-flowing radiation that must be calculated for input to the neighboring cells that share those outgoing faces, or to the body surface for those cells with outward faces coincident with a body surface. This cross-cell flow of radiation is essentially a one-dimensional process along direction $\mathbf{\Omega}$. The sense in which it is not exactly one-dimensional arises from the fact that the incoming and outgoing faces are not generally all equidistant from each other in direction $\mathbf{\Omega}$.

The volume-integral approximation (b) in (1.3) and the approximation (c) of using cell-center values of I at the outgoing faces in (1.4) can both be eliminated by using the 1-d exact radiative transfer solution for constant χ and S along direction Ω as a function of the distance s from the incoming faces:

$$\bar{I}(s,\Omega) = \bar{I}(0,\Omega)e^{-\chi s} + (1-e^{-\chi s})S \tag{2.1}$$

$\bar{I}(s,\Omega)$ is the cross-sectional average intensity at distance s along direction $\mathbf{\Omega}$. The average incoming intensity $\bar{I}(0,\Omega)$ is approximated as the area-weighted average intensity over all faces with incoming-directed radiation for direction $\mathbf{\Omega}$.

To determine the effective cross-cell propagation distance d in direction $\mathbf{\Omega}$, we divide the volume V by the total area of the incoming (or, equivalently, outgoing) faces projected on that direction:

$$d \approx \frac{V}{A}, \text{ where } A = \sum_{i,\Omega\cdot\Delta\mathbf{S}_i<0} \left|\Omega\cdot\Delta\mathbf{S}_i\right| \tag{2.2}$$

The intensity on each outgoing face $\mathbf{x}_o$ is then given by

$$I(\mathbf{x}_o,\Omega) = \frac{\sum_{i,\Omega\cdot\Delta\mathbf{S}_i<0} I(\mathbf{x}_i,\Omega)\left|\Omega\cdot\Delta\mathbf{S}_i\right|}{\sum_{i,\Omega\cdot\Delta\mathbf{S}_i<0} \left|\Omega\cdot\Delta\mathbf{S}_i\right|} e^{-\chi_c d} + (1-e^{-\chi_c d})S_c \tag{2.3}$$

where $\mathbf{x}_i$ indicates an incoming face center. This method will be referred to in what follows as the exponential-constant method since the source is assumed constant in the cell. Formula (2.3) is self-consistently obtained either by using (2.1) directly or by integrating over the volume on the rhs of (1.2), using (2.1) as the intensity in the integrand, and then solving for the outgoing intensity as was done for

(1.4). This method accounts more accurately for emission and absorption acting between the incoming and outgoing faces of a computational cell, essentially by treating those processes with sub-cell accuracy using an analytic form. Eqs. (2.3) and (1.4) are the same to first order in the optical thickness of the cell, $\chi_c d$. A second advantage of this formulation is that values of I are only required at the cell faces but, using (2.1), can be computed at any point in the cell.

A further improvement in formula (2.3) can be easily attained by treating the source function S for direction $\mathbf{\Omega}$ as a function of a scalar variable x, $0 \le x \le d$, and approximating it with two linear functions, one from the incoming faces to the cell center and one from the cell center to the outgoing faces:

$$S(x) \approx \begin{cases} S_\mathrm{i} + (S_\mathrm{c} - S_\mathrm{i})x/(d/2) & 0 \le x \le d/2 \\ S_\mathrm{c} + (S_\mathrm{o} - S_\mathrm{c})(x - d/2)/(d/2) & d/2 \le x \le d \end{cases} \tag{2.4}$$

For the incoming and outgoing S values, S_i and S_o, we use the average of the central value S_c and the area-weighted averages of the neighboring central values:

$$S_\mathrm{i} \equiv \frac{1}{2}\left(S_\mathrm{c} + \frac{\sum_{i,\Omega\cdot\Delta\mathbf{S}_i<0} S(\mathbf{x}_i,\Omega)\left|\Omega\cdot\Delta\mathbf{S}_i\right|}{\sum_{i,\Omega\cdot\Delta\mathbf{S}_i<0}\left|\Omega\cdot\Delta\mathbf{S}_i\right|} \right) \quad \text{and} \quad S_\mathrm{o} \equiv \frac{1}{2}\left(S_\mathrm{c} + \frac{\sum_{i,\Omega\cdot\Delta\mathbf{S}_i>0} S(\mathbf{x}_i,\Omega)\left|\Omega\cdot\Delta\mathbf{S}_i\right|}{\sum_{i,\Omega\cdot\Delta\mathbf{S}_i>0}\left|\Omega\cdot\Delta\mathbf{S}_i\right|} \right) \tag{2.5}$$

With these approximations to S at the outgoing and incoming faces, the intensities I at the cell center and at the outgoing faces become

$$I(\mathbf{x}_\mathrm{c},\Omega) = \frac{\sum_{i,\Omega\cdot\Delta\mathbf{S}_i<0} I(\mathbf{x}_i,\Omega)\left|\Omega\cdot\Delta\mathbf{S}_i\right|}{\sum_{i,\Omega\cdot\Delta\mathbf{S}_i<0}\left|\Omega\cdot\Delta\mathbf{S}_i\right|} e^{-\frac{1}{2}\chi_\mathrm{c} d} + \left(1 - \frac{1-e^{-\frac{1}{2}\chi_\mathrm{c} d}}{\frac{1}{2}\chi_\mathrm{c} d}\right) S_\mathrm{c} + \left(\frac{1-e^{-\frac{1}{2}\chi_\mathrm{c} d}}{\frac{1}{2}\chi_\mathrm{c} d} - e^{-\frac{1}{2}\chi_\mathrm{c} d}\right) S_\mathrm{i} \tag{2.6}$$

and

$$I(\mathbf{x}_\mathrm{o},\Omega) = I(\mathbf{x}_\mathrm{c},\Omega) e^{-\frac{1}{2}\chi_\mathrm{c} d} + \left(1 - \frac{1-e^{-\frac{1}{2}\chi_\mathrm{c} d}}{\frac{1}{2}\chi_\mathrm{c} d}\right) S_\mathrm{o} + \left(\frac{1-e^{-\frac{1}{2}\chi_\mathrm{c} d}}{\frac{1}{2}\chi_\mathrm{c} d} - e^{-\frac{1}{2}\chi_\mathrm{c} d}\right) S_\mathrm{c} \tag{2.7}$$

This method will be referred to as the exponential-linear method since the source is assumed piecewise-linear in a cell. It is clear that (2.6) and (2.7) require somewhat more computation that the constant-source approximation (2.3); however, there are many common factors in these expressions and only one exponential evaluation is still required. The improved accuracy is often worth the extra arithmetic, especially, as will be shown, when there are large and/or rapid variations in the source function S. A similar argument could be applied to the opacity χ, but the extra computation required to implement, say,

a linear variation of χ in a cell, is substantial, requiring, for one thing, error function evaluations; in addition, the improved accuracy, even for rapidly varying χ, is not found to be significant. Finally, it is elementary to show that the rhs's of (2.6) and (2.7) are non-negative, as required for a radiative intensity.

The total radiative heating in a cell, needed for the CFD energy equation, is easily computed from (1.4), (2.3), or (2.7) as

$$\int_{\text{cell}} \nabla \cdot \mathbf{q}\, dV = \int_{\text{surface}} \mathbf{q} \cdot d\mathbf{S} \approx \int_{\Omega} \sum_{\text{faces } k} I(\mathbf{x}_k, \Omega) \mathbf{\Omega} \cdot \Delta \mathbf{S}_k \, d\Omega \tag{2.8}$$

No differences are needed to compute this quantity, only a summation over the cell faces. The other quantity typically needed in a rad-hydro calculation is the heat flux at a body surface. Assuming that face k of a cell lies on this surface, the heat flux through that surface is simply:

$$\mathbf{q}_k \approx \int_{\Omega} \mathbf{\Omega}\, I(\mathbf{x}_k, \Omega)\, d\Omega \tag{2.9}$$

3 Verification and Comparison

In this section we will be plotting two basic functions of the radiative intensity I: the radiative energy E and the radiative heat flux $\mathbf{q}$:

$$E(\mathbf{x}) = \frac{1}{c} \int_{4\pi} I(\mathbf{x}, \mathbf{\Omega}) d\Omega\,; \quad \mathbf{q}(\mathbf{x}) = \int_{4\pi} \mathbf{\Omega} I(\mathbf{x}, \mathbf{\Omega}) d\Omega \tag{3.1}$$

3.1 One-dimensional problems

In this sub-section we will compare results from the original (eq. (1.4)) and modified (eq. (2.3) or (2.6) and (2.7)) finite-volume RTE methods with exact analytic solutions in one spatial dimension. For such problems the angular integrations in (3.1) reduce to a simple sum or difference respectively of the intensities for radiation travelling in the forward and reverse directions; hence there is no error associated with angular quadrature.

First we consider an isolated Gaussian source that emits into a weakly absorbing medium bordered by a more strongly absorbing layer. Such a geometry is analogous to that of a reentry-vehicle's radiating shock layer and wall boundary layer. The source function S and opacity χ are shown in figure 1. For this case, $S(x) = 10 \exp\left(-\left((x - 0.25)/0.02\right)^2\right)$.

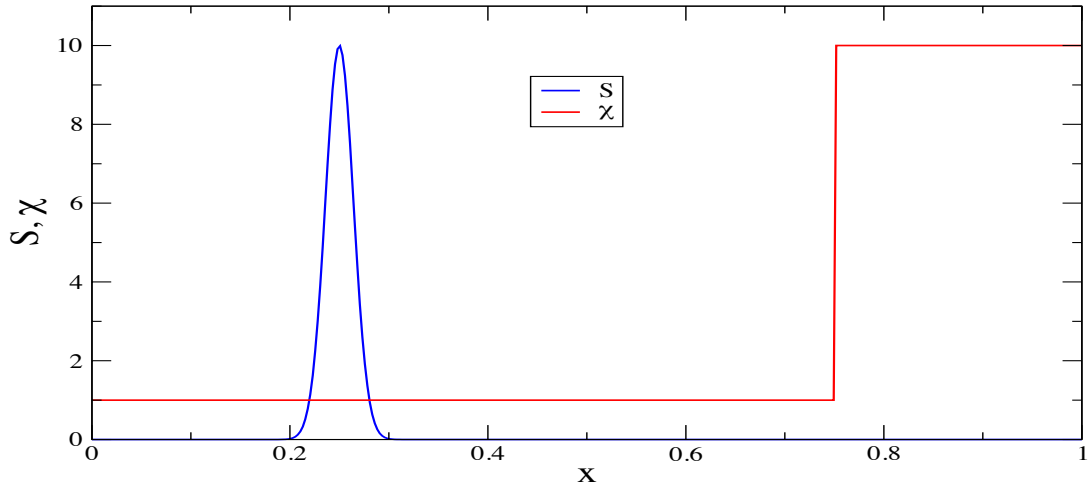


Figure 1. Source and opacity functions for comparison case 1: Gaussian source.

The RTE can be solved analytically for these source and opacity functions. First we compare solutions obtained using the original finite-volume method for various grid resolutions; this is shown in figure 2. As can be seen, convergence is nicely monotonic with increasing number of grid points n_x, but a rather fine mesh is required to converge the peak region near $x = 0.25$; even $n_x = 400$ is not very satisfactory. The situation is dramatically improved by using the exponential-constant finite volume method, as shown in figure 3 on the same scale; a zoom-in on the peak region for this method is shown in figure 4.

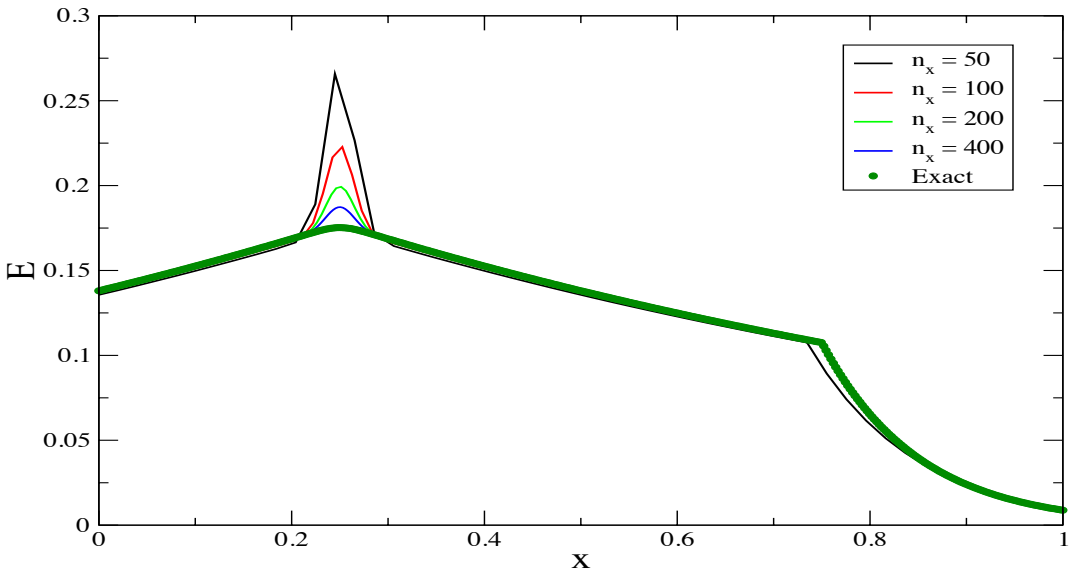


Figure 2. Radiative energy for comparison case 1: original finite volume method.

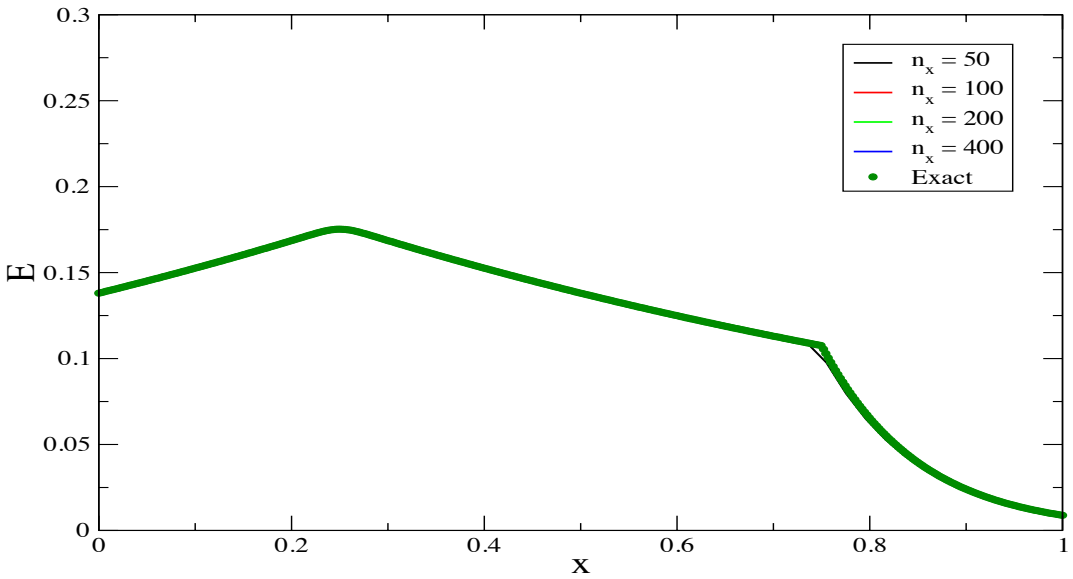


Figure 3. Radiative energy for comparison case 1: exponential-constant finite volume method.

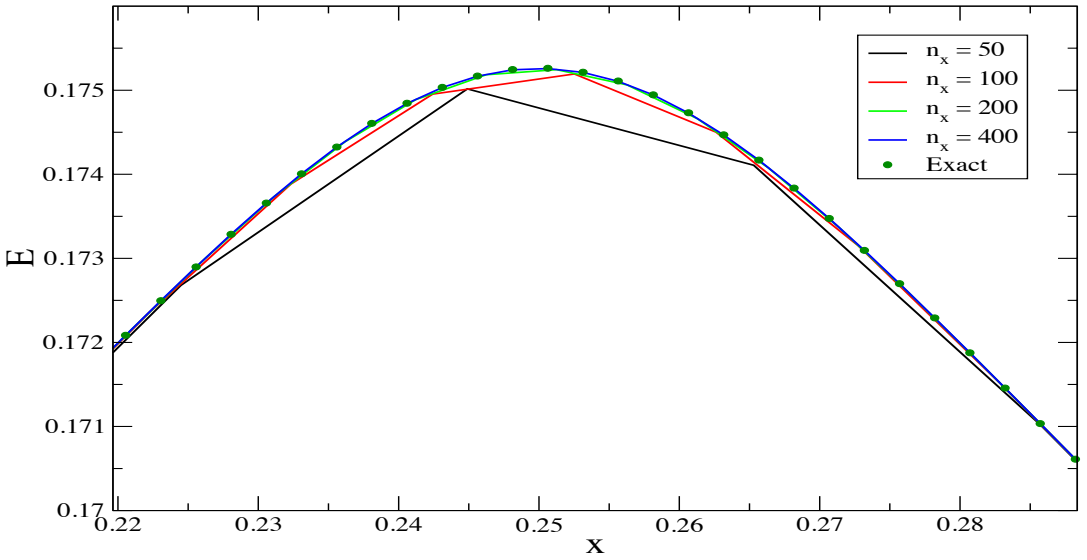


Figure 4. Radiative energy for comparison case 1: zoom-in on the peak region for the exponential-constant finite volume method

We now investigate the spatial-resolution properties of the various finite-volume schemes by using a sinusoidally oscillating source as shown, along with the opacity, in figure 5. This source function is given by $S(x) = 10 + 9\sin(200x)$.

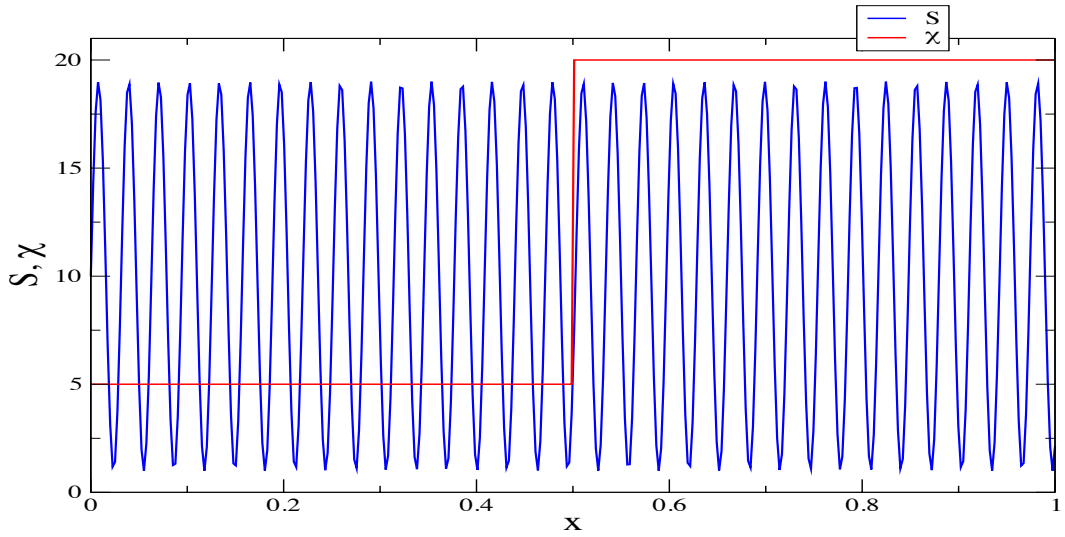


Figure 5. Source and opacity for comparison case 2: sinusoidal source.

As for the Gaussian source, this problem admits an analytic solution to which comparisons will be made. In the Gaussian source case just discussed, the two improved finite-volume methods described in section 2 give essentially the same results, and only the exponential-constant method (eq. (2.3)) results were shown. In the current sinusoidal-source case, this is not true, and in fact the exponential-linear method (eq. (2.7)) gives somewhat better results. In figure 6 we show a comparison, for a quite coarse resolution (n_x = 50), of the results from all three methods. It is clear that, for this marginally adequate grid, the exponential-linear scheme is best; it is also clear that either exponential scheme is considerably superior to the original scheme, showing, for instance, much less spurious oscillation.

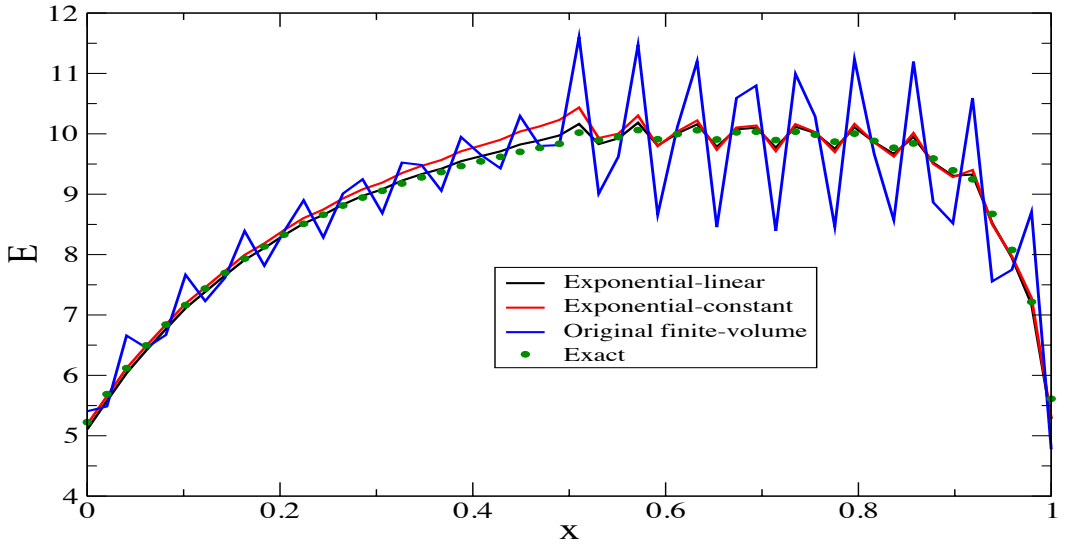


Figure 6. Radiative energy for comparison case 2: three finite-volume methods for a grid of n_x = 50.

3.2 Three-dimensional problems

We now move on to multidimensional applications and compare the results of the finite volume methods to those of other numerical methods, since analytical solutions are not generally available for multidimensional problems. In figures 7(a,b,c,d) below, a Gaussian radiative source and opacity are

present which are distributed on a hyperboloidal surface, emulating the reentry shock around a zero angle-of-attack axisymmetric bluff body. The source and opacity have the following forms:

$$S(x,y,z) = \begin{cases} \exp\left(-\left((x - x_0(y,z))/L_x\right)^2\right) & \text{for } x \geq x_0(y,z) \\ 0 & \text{otherwise} \end{cases}$$

$$x_0(y,z) = h_x + h_a\left(\left((y-h_y)/h_b\right)^2 + \left((z-h_z)/h_c\right)^2 + 1\right)^{\frac{1}{2}} \tag{3.2}$$

$$\chi(x,y,z) = c_1 + c_2 S(x,y,z)$$

In figures 7-9, contours of cell-average radiative heating, div(**q**), are shown in the z-midplane; x is streamwise, and y and z are cross-stream. In figure 7, we compare [1] a Cartesian-mesh, reference solution (figure 7(a)) on a fine grid (400×400×401) obtained using a short characteristic RTE solver; [2] a solution using the original finite-volume method on a 20×20×21 shock-fitted mesh (figure 7(b)); and [3] a solution using the exponential-constant method on the same fitted mesh (figure 7(c)). The reference solution uses 54 quadrature points in $\mathbf{\Omega}$; the finite-volume solutions shown use the S6 quadrature method which has 48 points in $\mathbf{\Omega}$; use of the same quadrature method as for the short-characteristic method (not shown) has negligible effect on the results. The shock-fitted mesh allows the use of coarse grids while still retaining reasonable accuracy. Compared to the fine-grid, characteristic method result (7(a)), we see that the original finite-volume method (7(b)) shows a greater difference in peak heating (weaker) than the exponential-constant method (7(c)), especially off the centerline. The difference is admittedly small, and smaller than the difference between the finite-volume results compared to the reference. On refining the grid to 50×50×51 (not shown), the difference between the two finite-volume methods almost disappears, and the difference from the reference is much reduced. Figure 7(d) shows the results from the exponential-constant method on a 100×100×101 grid, for which effective convergence to the reference solution has been obtained.

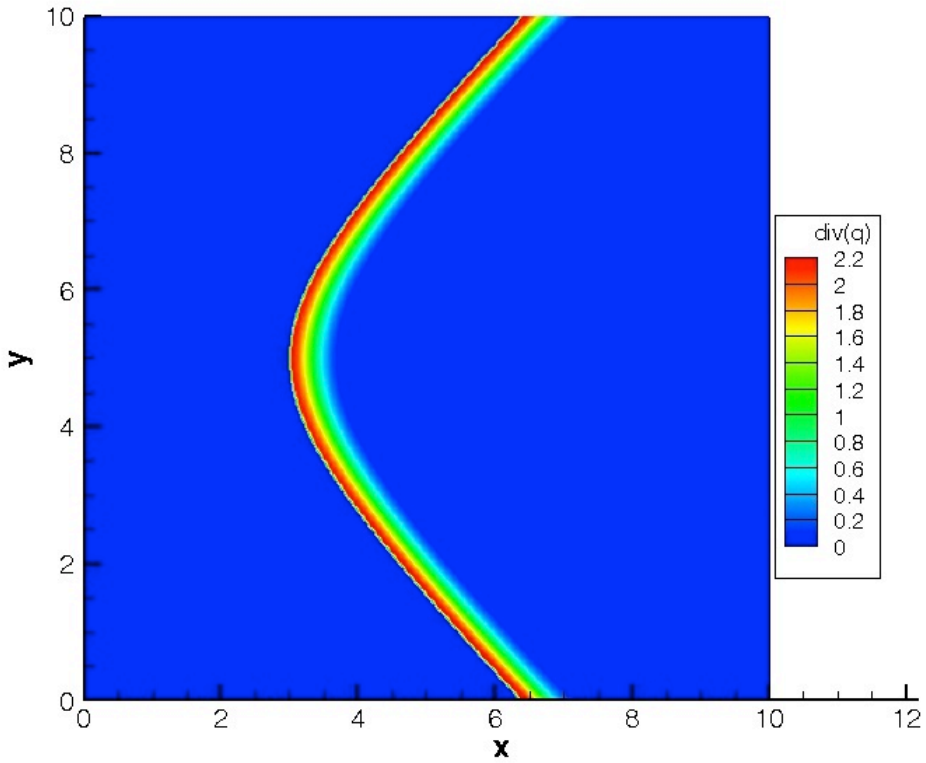


Figure 7(a): div(**q**) using a 400×400×401 Cartesian-mesh, short-characteristic-based method.

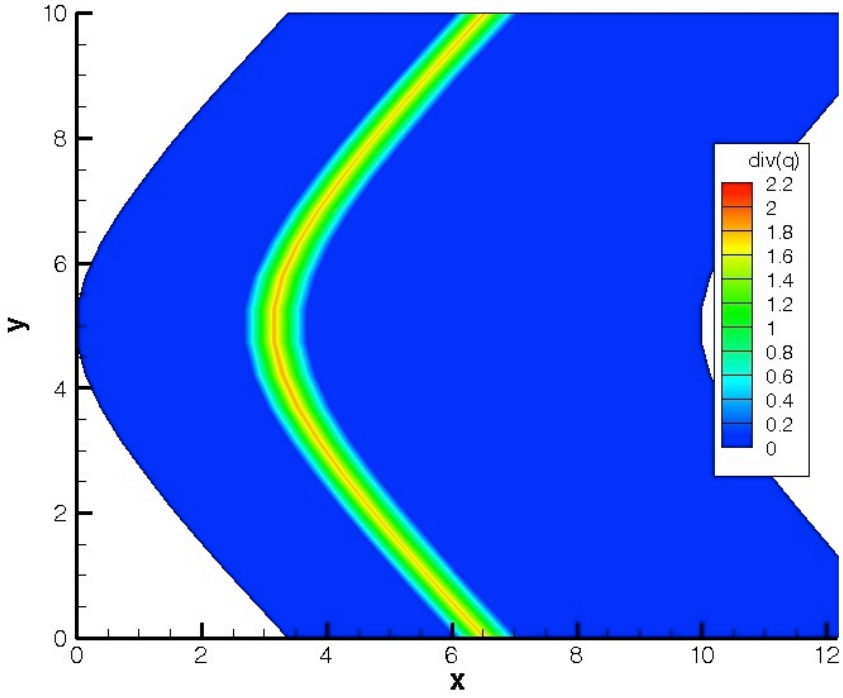


Figure 7(b): div(**q**) using the original finite-volume method on a shock-fitted 20×20×21mesh.

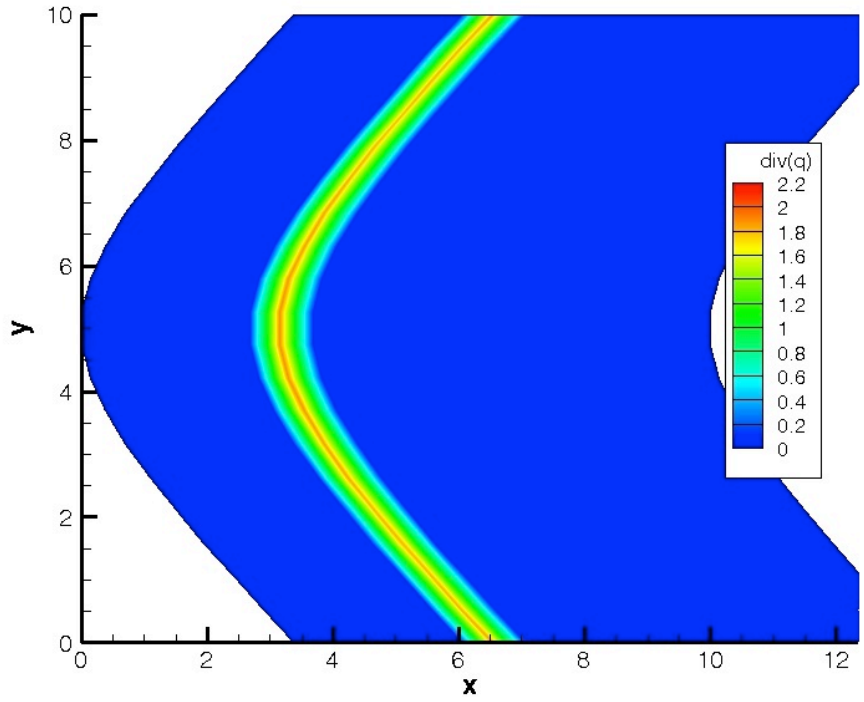


Figure 7(c): div(**q**) using the exponential-constant finite-volume method on a shock-fitted 20×20×21mesh.

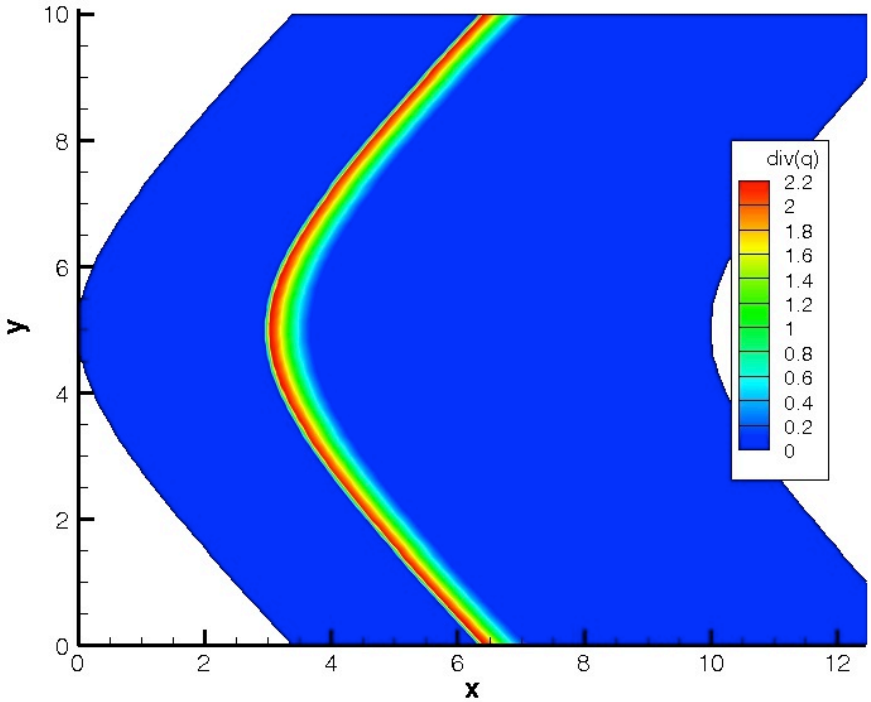


Figure 7(d): div(**q**) using the exponential-constant finite-volume method on a shock-fitted 100×100×101 mesh.

Figures 8(a,b,c) show a similar set of results, in this case for a medium having a 10× higher opacity in the emitting region ($c_2 \to 10c_2$ in eq. (3.2)) and a 10× smaller opacity ($c_1 \to c_1/10$ in eq. (3.2)) in the post-shock region. Grid convergence for all three finite volume methods is slower for this opacity field, most likely because the stronger absorption in the emitting region requires finer resolution; midplane results in the figures are shown for a 50×50×51 grid. Similar conclusions hold as for the previous case in that the exponential-constant method is closer to the reference than the original finite volume method. Figure 9 shows a close-up along the centerline in *y* of the region near the hyperboloidal source for this case. It is clear that for both meshes used, 20×20×21 and 50×50×51, the exponential-constant method gives results somewhat closer to the reference solution. It is also clear that the finer grid gives clearly superior results; in fact, a grid of 100×100×101 for this problem (not shown) gives results essentially converged to the reference. The exponential-linear method gives results that differ only inconsequentially from the exponential-constant scheme. A point to note is that the small maximum and minimum in the reference solution just to the left of *x*=3 are almost certainly nonphysical artifacts of the angular quadrature, a so-called "ray" effect; finite volume methods are less susceptible to this effect than are characteristic methods, and the finite-volume results obtained here do not show this small anomaly.

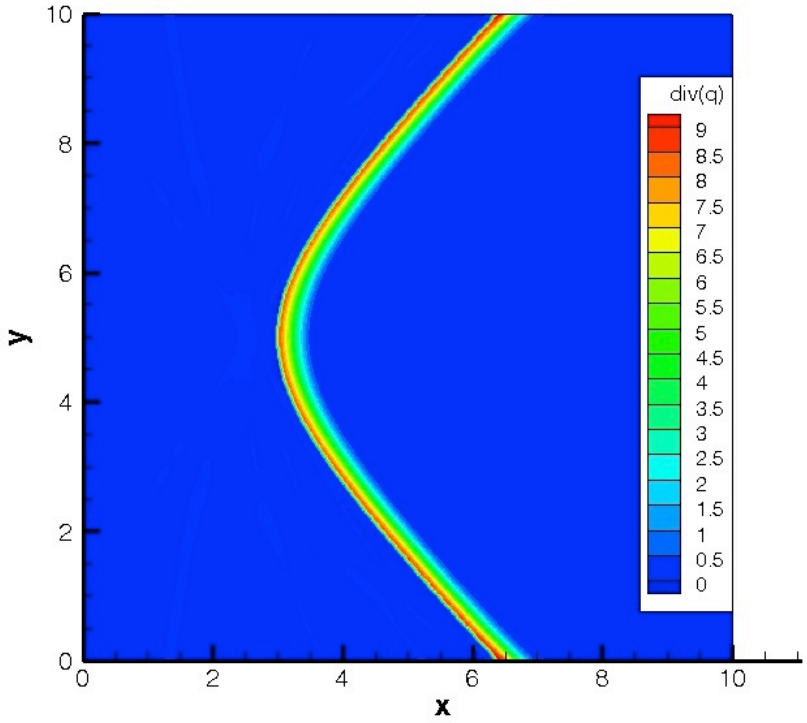


Figure 8(a): div(**q**) using a 400×400×401 Cartesian-mesh, short-characteristic method.

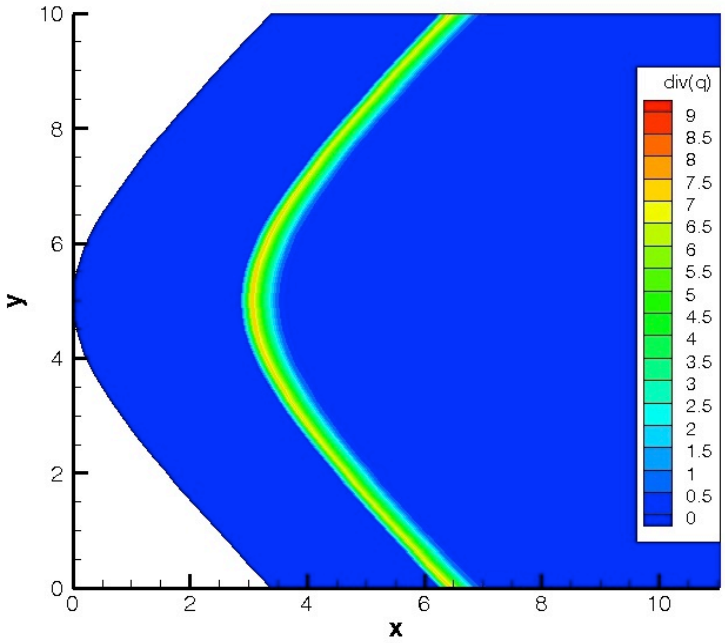


Figure 8(b): div(**q**) using the original finite-volume method on a shock-fitted 50×50×51mesh.

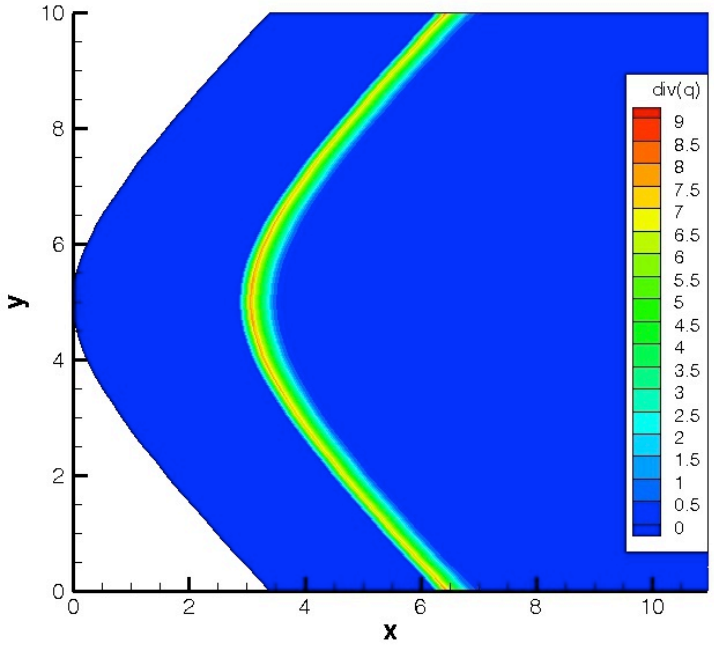


Figure 8(c): div(**q**) using the exponential-constant finite-volume method on a shock-fitted 50×50×51mesh.

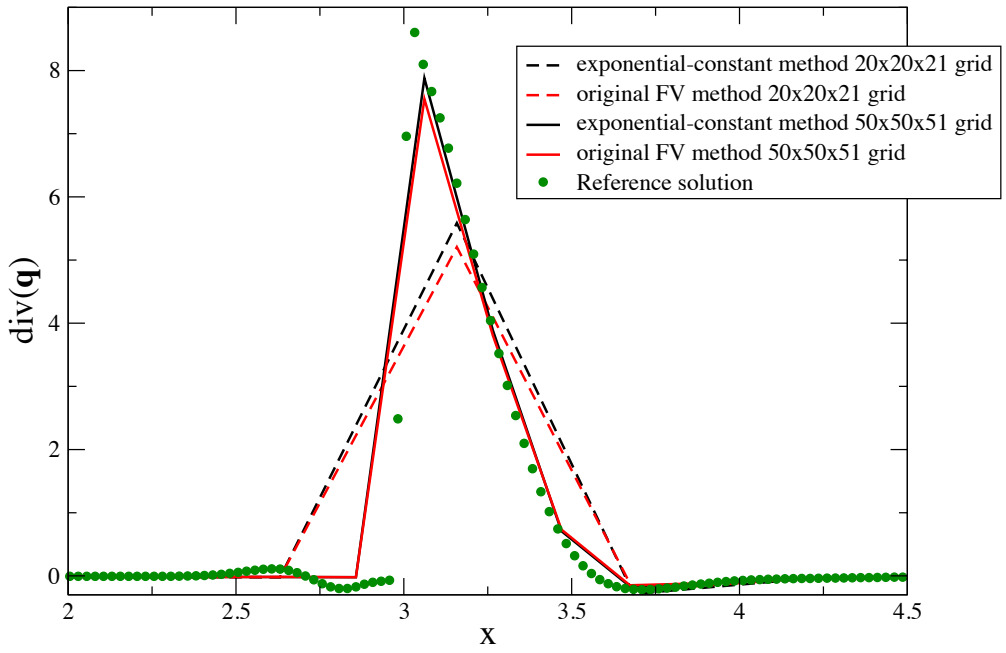


Figure 9. Centerline values of radiative heating rate (div(**q**)) zoomed in to show the peak values near the hyperboloidal source. Green circles denote the reference fine-grid solution; black curves denote the exponential-constant method results; red curves denote the original finite-volume method results. The grids used to obtain the black and red curves are 20×20×21 for the dashed lines and 50×50×51 for the solid lines. The reference solution used a 400×400×401 grid.

3.3 Plane-parallel problems

Plane-parallel radiation problems are a special case of three-dimensional geometry in which the domain is bounded on two sides by infinite parallel walls and in which the source and opacity depend only on the Cartesian coordinate normal to those walls. The radiation still travels through three dimensions as usual, but the resultant radiative heat transfer and other properties of the radiation field necessarily depend on only that single coordinate. In this case the angular quadrature shown in equation (3.1) can be done analytically in terms of exponential integrals; there remains the integration over one spatial variable (eq. (1.1)) which must be done numerically except in special cases. This semi-analytic solution is used to construct the reference plane-parallel data for the source and opacity shown in figure 10(a). For improved accuracy in the reference solution, the numerical quadrature in x is performed on a grid 100 times as fine as for the finite-volume methods, but for plotting clarity only every 100th point is shown in figures 10(b-e).

Plane-parallel problems provide a convenient check on the angular quadrature and spatial transfer accuracy of numerical methods that integrate the RTE in three spatial dimensions, such as the finite-volume methods being considered here. As always in solving such problems fully numerically, there are errors introduced by the finiteness of the numerical domain in the transverse directions and by the discrete angular quadrature. In the solutions shown we use a domain which is 50 wall-separations in length and width and which is discretized into 51×51 cells in those directions, which we take to be y and z; 100 points span the wall-to-wall direction x. The S8 angular quadrature method is used; it has 80 quadrature points in solid-angle space.

We simulate a problem containing a compact radiative source imbedded in an absorbing medium that is bordered on one side by a transparent medium (figure 10(a)). Plots of the radiative energy and the radiant heat flux are shown in figures 10(b-f). Note that in these figures the black curve for the exponential-linear case is often nearly coincident with and underneath the red curve for the exponential-constant method.

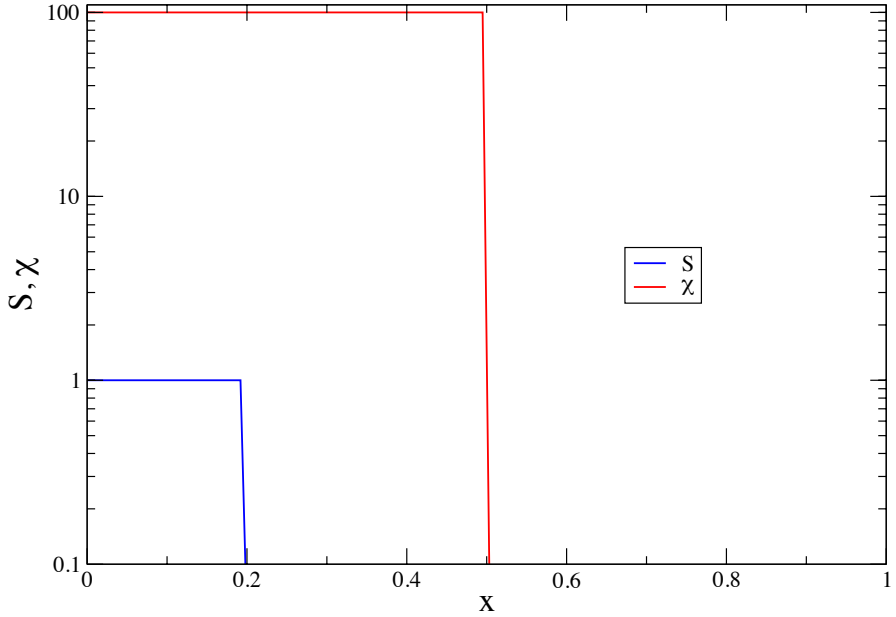


Figure 10(a). Source and opacity functions for the plane-parallel test case.

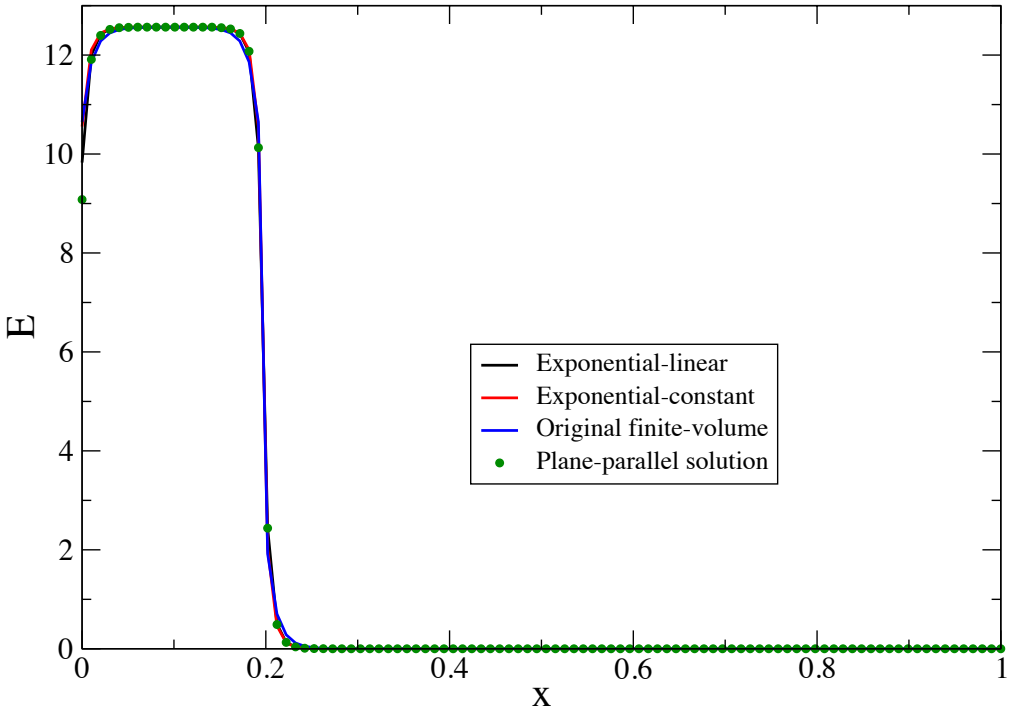


Figure 10(b). The radiative energy for the plane-parallel test case over the entire domain.

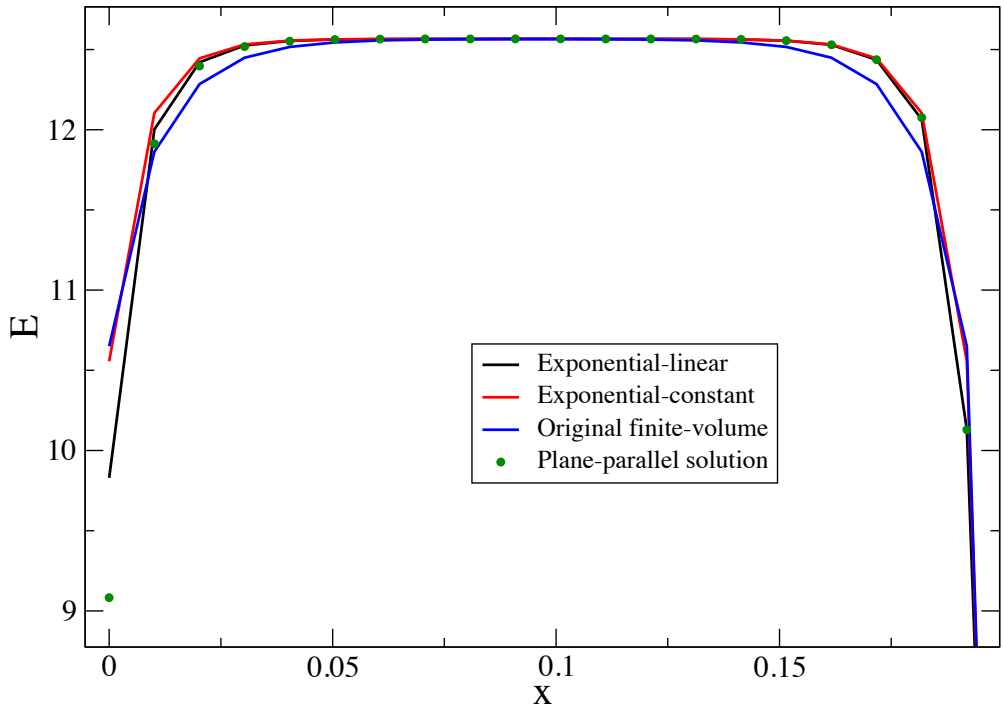


Figure 10(c). The radiative energy zoomed-in on the non-zero source region for the plane-parallel test case.

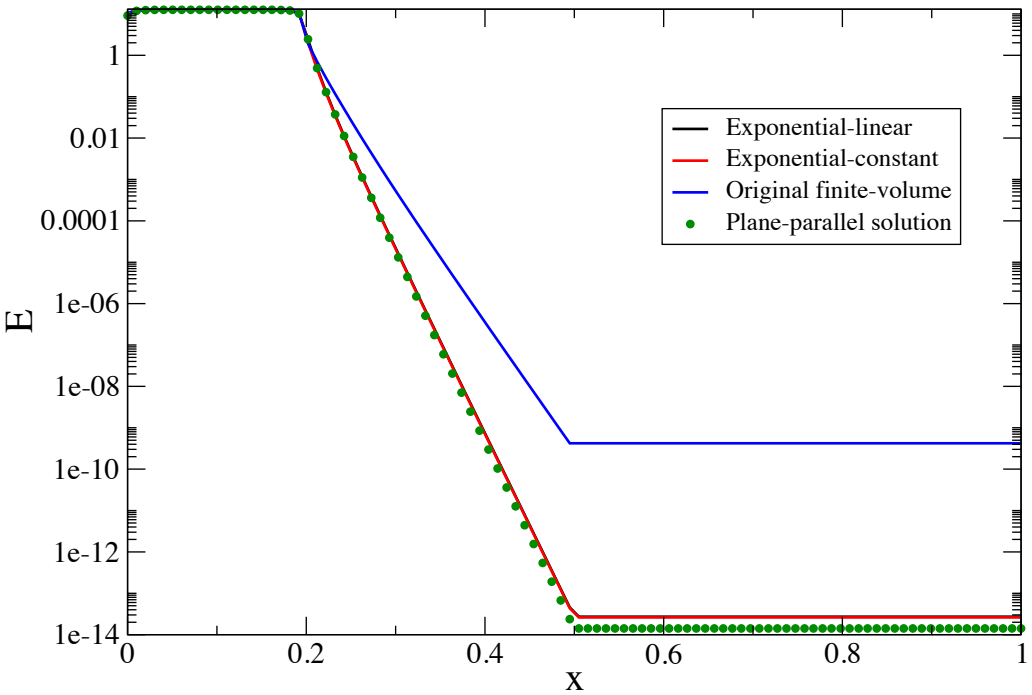


Figure 10(d). The radiative energy in log coordinates for the plane-parallel test case.

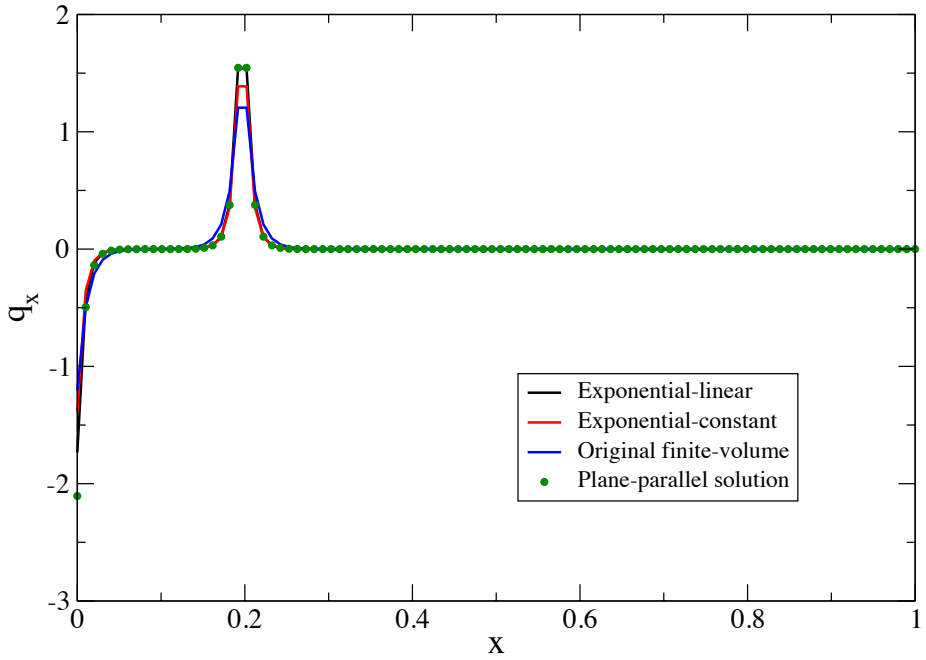


Figure 10(e). Radiative heat flux for the plane-parallel test case.

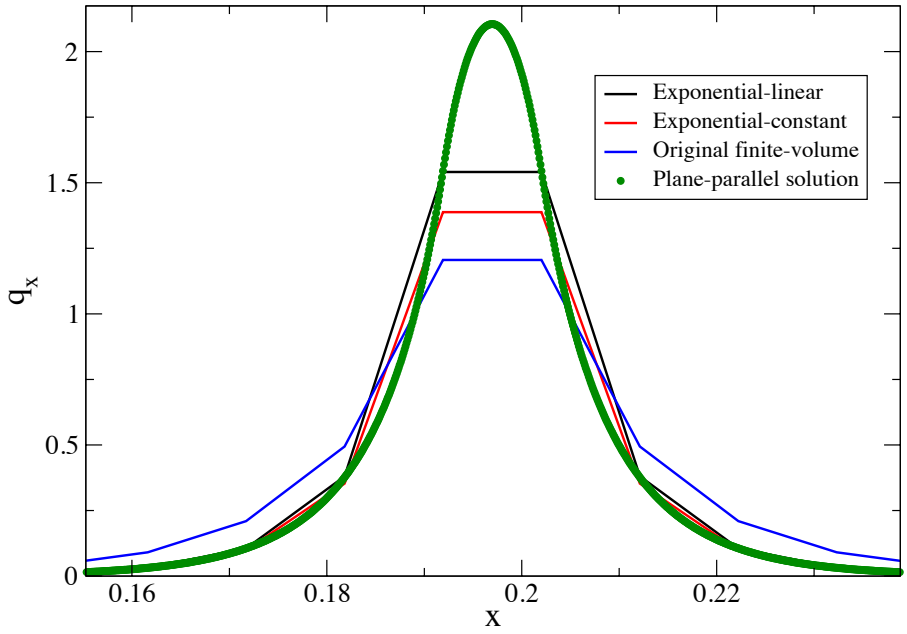


Figure 10(f). A zoom-in on the radiative heat flux for the plane-parallel test case showing the semi-analytic solution at full resolution.

As can be seen in figures 10(b) and (e), the overall accuracy from all three finite volume methods for the radiative energy E and heat flux q_x is good, but, zooming in on the region where the source is non-zero, figure 10(c), the exponential methods show greater accuracy in the shoulder regions of E, and the exponential-linear method shows somewhat better accuracy than the exponential-constant method.

Furthermore, looking at the radiative energy in log coordinates, figure 10(d), clearly shows greater accuracy of the exponential methods in regions away from the source where E is small. Finally, figure 10(f) shows that the exponential methods, and especially the exponential-linear method, yield a more accurate heat flux compared to the original finite-volume method, particularly in the important peak region. In this figure all grid points in x are shown for the semi-analytic reference solution.

3.4 Reentry body heat shield

Our final example will be a solution for the radiative heat flux onto the heat shield of a CEV-like body reentering the Earth's atmosphere. The conditions of the flow are a speed of 9 km/sec at an altitude of 48 km at an angle of attack of 27.5° and zero yaw. These conditions were chosen because they are not too far from being in thermochemical equilibrium while still having a significant amount of radiative heating. The fluid immediately after the bow shock is to some degree in a nonequilibrium thermochemical state, but, for simplicity of calculation of the excited-state populations, the fluid is assumed to be in a state of local thermodynamic equilibrium for radiation purposes. The flowfield temperature and species densities have been taken from a CFD solution obtained with the DPLR hypersonic CFD code [2]. The grid used for the radiation solution is the same as that used for the CFD, which is a multi-block, structured mesh.

The NASA Ames hypersonic radiation code HyperRad was used to obtain the radiative source and opacity throughout the flowfield as a function of wavelength on the CFD grid, and the finite-volume RTE solver in HyperRad was used to solve for the three-dimensional radiation field, including the heat flux onto the body surface. The computation of the source and opacity is implemented in a parallel algorithm using spatial domain decomposition. The RTE solver implements a parallel algorithm for the developments of sections 1 and 2 of this paper. Parallelism for the RTE solution is attained by treating different spectral data points in parallel and simply summing over wavelength to form the total radiative properties such as the heat flux vector and its divergence. For the particular problem shown here, 128 processors were used on the NAS Pleiades computer with essentially 100% parallel efficiency since spectral points are completely independent and only the final integrations over wavelength require communication. Further parallelism over the angular quadrature directions can also be invoked for non-scattering problems but that was not done here.

The radiative solution to which we will compare results for this problem is obtained from the legacy NEQAIR code [3], which has been used for many years to predict radiative heat loads on reentering bodies and to predict the radiation from shock-tube experiments. This code does not have three-dimensional radiative transport, but instead uses an approximate plane-parallel method ("tangent slab"), under the assumption that shock-layer heating can be reasonably approximated in this manner at least for the forebody in regions where the shock layer is near the surface and locally nearly uniform. It is clear that such a method will in general overestimate the peak radiative heating since the radiative source strength will decrease away from the peak following the shock curvature whereas the plane-parallel assumption is that it is constant on local tangent planes. In parts of the flow where the body curvature is strong and where the shock layer is far from parallel to the body, results from the NEQAIR and HyperRad algorithms may differ significantly. In addition, the spectral data used by NEQAIR has been updated in HyperRad to the current state of knowledge for the line strengths and positions of the radiating species, thus the emissivities and opacities differ between the two codes. Cases using from 10,000 to 30,000 spectral sample points in wavelength space were used by HyperRad with essentially no resulting difference in the computed radiation field.

The HyperRad solutions on the given DPLR grid (2.35 million points) turn out to be nearly identical between the exponential-constant and original finite volume methods. This fact can be used to conclude that this grid leads to a well resolved RTE solution, since it is in cases of low grid resolution that the exponential finite-volume methods would give different, and, based on the results of this paper, presumably more accurate, answers.

Figure 11 shows an oblique view of the surface-normal component of the radiative heat flux; (a) is the NEQAIR solution and (b) is the HyperRad finite-volume solution on the same color scale. The peak value on the forebody differs by a factor of about 3 between them, but the distribution across the forebody is similar. This difference in peak level most likely results at least partly from a difference between the plane-parallel approximation and the full 3D algorithm, as discussed previously; differences in the spectral data could also contribute. The differences in the shoulder regions are more interesting. Assuming they are due to three-dimensional effects, this result shows that the radiative heating on this part of the spacecraft is even larger than the peak heating on the forebody. That is not to say, of course, that the total heat load, including convective heating, is necessarily larger there: as discussed above, the trajectory point being considered was chosen to be near equilibrium conditions and is not at peak radiative or total heating.

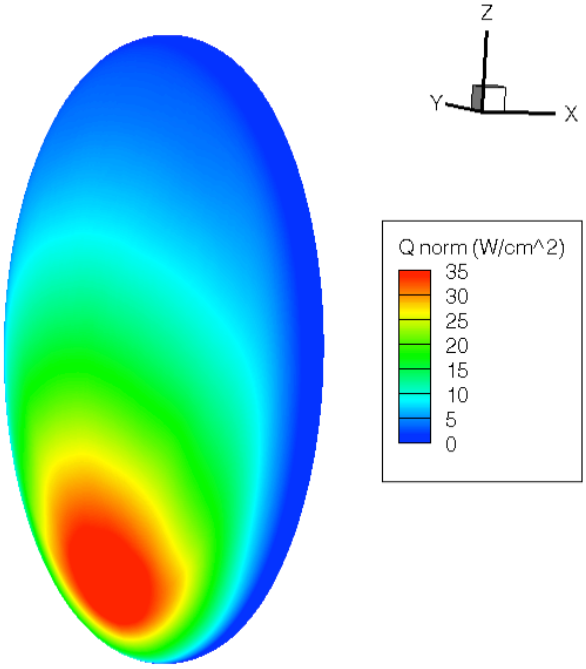


Figure 11(a). NEQAIR solution for a CEV-like forebody in Earth reentry.

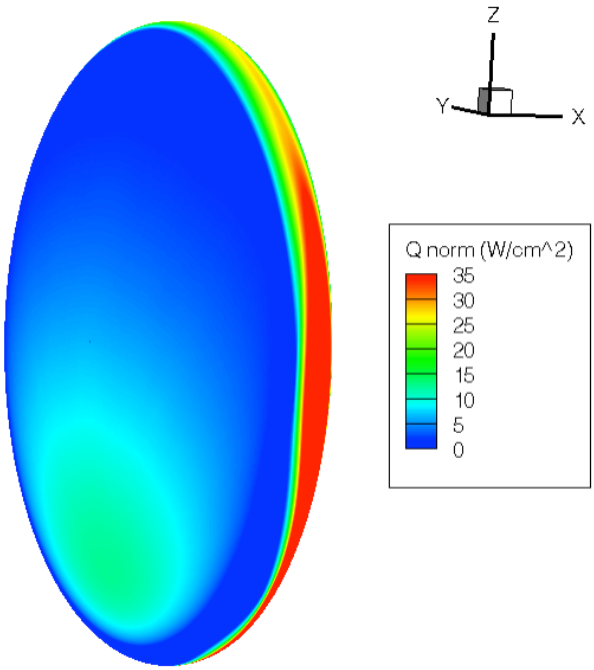


Figure 11(b). HyperRad solution for a CEV-like forebody in Earth reentry.

4 Conclusions

We have found clear evidence that the new finite-volume algorithms give improved results as compared to the classical one in many cases. The differences are modest in most instances, but computational radiative CFD is typically an expensive undertaking, in terms of computational cost, because of the increased dimensionality of the problem (6 vs. 3). In fully coupled simulations, in which the radiative contributions to heating are calculated simultaneously with the fluid solution, the radiative portion of the work can easily dominate and may force a reduction in solution accuracy in order to make the problem tractable with given resources. Therefore, any improvement in accuracy that does not entail significantly increased computation is worth adopting, and the methods developed here, we believe, fit in that category. The exponential-constant solution for the CEV-like forebody problem cost approximately 6% more than the classical FV method.

5 References

[1] Modest, M., *Radiative Heat Transfer, 2nd ed.*, Academic Press, p. 525, 2003.

[2] Wright, M. J., Candler, G. V., and Bose, D., "Data-Parallel Line Relaxation Method for the Navier–Stokes Equations," AIAA Journal, Vol. 36, No. 9, 1998, pp. 1603–1609.

[3] Whiting, E. E., Park, C., Liu, Y., Arnold, J. O., and Paterson, J. A., "NEQAIR96, Nonequilibrium and Equilibrium Radiative Transport and Spectra Program: User's Manual," NASA Rept. 1389, Dec. 1996.