A Research Note on Progress in Calculation Verification for Code Project A

Trevor Tippetts X-1
Francis X. Timmes X-2
James R. Kamm X-1
Jerry S. Brock X-1

Los Alamos National Laboratory
Los Alamos, NM 87545 USA

September 29, 2006

LA-UR-06-7225

List of Figures

1. Summary

What's New:

- A new optimization technique for conducting a calculation verification analysis on complex multi-physics problems has been developed and applied to a few test problems. The new technique uses a minimization algorithm to produce an estimated exact solution and a self-consistent convergence rate. Code verification is the term used when an analytic exact solution is available. Calculation verification is the term used when an analytic solution is unavailable.

Results:

- The new method was trained and assessed on the Sedov problem with the Eulerian code RAGE. Since the Sedov problem has an exact solution, this allowed a comparison between the estimated solution with the analytical solution. In general, the estimated solution on 1D uniform grids produces a shock front of nearly the same strength and location as the analytic solution while obtaining similar rates of spatial convergence as the exact solution. Our results on problems that do not admit an exact solution are reported elsewhere.

Recommended Directions:

- Continue applying such calculation verification methods to an increasing number of multi-physics problems in 1D and 2D where an exact solution is not known, using both uniform and adaptive meshes. This is a key growth direction for verification efforts to bridge the gap between analytical test problems and complex applications.

2.0 Introduction

An important element of physics-based code verification activities is the quantitative analysis of the error between computed and exact solutions for well-defined problems. This approach is based on two assumptions that constrain the applicability of the approach. The first assumption is that an exact solution is available, while the second assumption is the adoption of a model for how the spatial-temporal discretization errors behave. When the first assumption is violated, an estimated exact solution will often be used as a stand-in for the exact solution. This approach is sometimes referred to as calculation verification (Celik et al. 2005 and references therein). Normally, a highly refined numerical simulation is used as the estimated exact solution.

Recently, Smitherman, Kamm & Brock (2005) suggested a new technique to calculate a point-wise estimation of the exact solution that is mathematically consistent with the convergence properties of the numerical simulations. That is, the estimated exact solution is produced as part of the calculation verification. Their new technique allows for cases when the convergence is non-monotonic (i.e., oscillatory convergence) and frees the error model from being a simple power-law function of grid spacing.

In essence, Smitherman et al. (2005) investigated setting the absolute value of the difference between the numerical solution and an estimated exact solution equal to the error model. For a 1D uniform grid and a power-law function of the grid spacing their idea takes the form

$$|\hat{y} - y_n| = A\,(\Delta x)^{\alpha} \quad , \tag{1}$$

where $\hat{y}$ is the estimated exact solution of the field of interest (e.g., the density) , y_n is the numerical solution computed on a given grid, Δx is the grid spacing of reference grid, A is the magnitude of the error in the field on the reference grid, and α is the spatial convergence rate. It must be understood that equation (1) is computed at each grid point, so if there are 100 cells in the numerical solution then there are 100 equations. At each grid point, equation (1) has three unknowns; the estimated exact solution $\hat{y}$, the pre-factor A and the convergence rate α. Three grids are then the minimum number of grids necessary to specify the three unknowns at each grid point. Let the numerical solution be computed on coarse c, medium m, and fine f grids. It is essential to recognize that the absolute value in equation (1) ensures that the three equations for the three unknowns cannot be solved for analytically. unless $\Delta x_c/\Delta x_m = \Delta x_m/\Delta x_f$. Either way, equation (1) has multiple solutions. Smitherman et al. (2005) thus sought the roots of the equation

$$\mathbf{F} = \begin{bmatrix} F_1 \\ F_2 \\ F_3 \end{bmatrix} = \begin{bmatrix} |\hat{y} - y_c| - A\,(\Delta x_c)^{\alpha} \\ |\hat{y} - y_m| - A\,(\Delta x_m)^{\alpha} \\ |\hat{y} - y_f| - A\,(\Delta x_f)^{\alpha} \end{bmatrix} = 0 \ . \tag{2}$$

All the quantities in equation (2) are evaluated on the same grid. Either the medium and fine computational solutions are spatially averaged up to the coarse grid level, or coarse and medium grids are interpolated down to the fine grid level. Which direction is more robust may be problem dependent and is discussed in more detail below for the Sedov problem.

The fundamental assumption of the error model in equation (1) is that error decreases as zone size decreases. For any cell, numerical values of a physical variable output by the simulation, y_c, y_m, and y_f , may be in one of six

possible sequences and three possible states. The ratio

$$R = \frac{y_m - y_f}{y_c - y_m} \tag{3}$$

has been used to categorize the behavior of y when the grid size is decreased as (i) monotonically convergent when $0 < R < 1$, (ii) oscillatory convergent when $R < 0$, and (iii) monotonically divergent when $R > 1$. Although oscillatory convergence is allowed by the left hand side of equation (1) , the error model in the right-hand side does not allow it. Most sets of points in the majority of simulations do not exhibit oscillatory convergence, so calculating monotonic convergence has been an acceptable method for most problems. Oscillatory convergence has been seen to occur in non-smooth problems, e.g., at shock fronts and discontinuity boundaries. This will have ramifications for the Sedov problem discussed below. Either the technique discussed by Smitherman et al. (2005) or the technique discussed here have the potential to solve for fields that display either monotonic or oscillatory convergence behavior.

As noted by Smitherman et al. (2005), choosing to solve equation (2) by a Newton-like root-find has its idiosyncrasies. It may even be the best approach. An alternative strategy, which we study here, is optimization. We collapse the multidimensional root finding into one dimension by summing the squares of the individual functions F_i in equation (2) to obtain new function, $f = \mathbf{F} \cdot \mathbf{F}$, which is positive definite and has a global minimum of zero exactly at all solutions of the original set of nonlinear equations (Press et al. 1992). An advantage of a minimization approach is that the test for ``downhill" on a single surface is intrinsically one-dimensional. There is no analogous conceptual procedure for finding a multidimensional root, where ``downhill" must mean downhill simultaneously in N separate function spaces, thus allowing a multitude of trade-offs, as to how much progress in one dimension is worth compared with progress in another (Press et al. 1992). Minimization of a scalar composite error also allows for the use of more grids than parameters in the error model, which could reduce the sensitivity of the parameters to any individual grid. A potential disadvantage of the minimization approach is that the new function f may have has a great number of local minima, not all of which are a zero of the the original set of nonlinear equations. Thus, while such methods can still occasionally fail by coming to rest on a local minimum of f, they might succeed when a direct attack on $\mathbf{F}$ via a Newton-like method may stumble. We have chosen to use the TNBC optimization package (Nash 1984ab), which uses a constrained, truncated-newton algorithm to find a local minimum of f. This package does not assume that the function f is convex and, so, cannot guarantee a global solution but it does assume that the function is bounded below.

3.0 Results

We trained and assessed our new optimization method on the venerable Sedov test problem before attempting more complicated problems that do not possess an analytic solution.

We used RAGE 20060331.0240 on the Linux cluster QSC to perform the numerical simulations. RAGE is an adaptive mesh refinement hydrodynamics code that uses a high-order Godunov (direct Eulerian, piecewise linear) method with an approximate Riemann solver to obtain the solution of the gas dynamics equations. The calculations considered here were all done on 1D uniform grids in order to be consistent with the error model (equation 1) that governs our analysis. The RAGE input deck for the Sedov problem were taken from the SourceForge archive. After the problems were run, John Grove's AMHCTOOLS (2005a, 2005b) was used to extract the solution data on the

native grid from the binary dump files. If one requests the simulation data directly from RAGE, the data on the native mesh are interpolated onto a uniform mesh. Extracting the solution data on the native mesh is important for proper verification analysis, particularly on adaptive meshes.

For the Sedov problem, a finite amount of energy is deposited at the origin at an initial time. The problem of finding self-similar, one-dimensional solutions for compressible hydrodynamics was considered by Sedov (1959), Taylor (1950), and von Neumann (1947). Sedov provided the most general closed-form solution, which we employ in the forms considered by Kamm (2000b). A succinct description of the solution to the Sedov problem, including regularization of the singularities at the lower limits of integration and fortran code for generating solutions, is given by Timmes, Gisler & Hrbek (2005), spatial-temporal convergence properties are discussed in Timmes, Fryxell & Hrbek (2006a), and a verification analysis on multi-dimensional versions of the problem are discussed in Timmes, Fryxell & Hrbek (2006b).

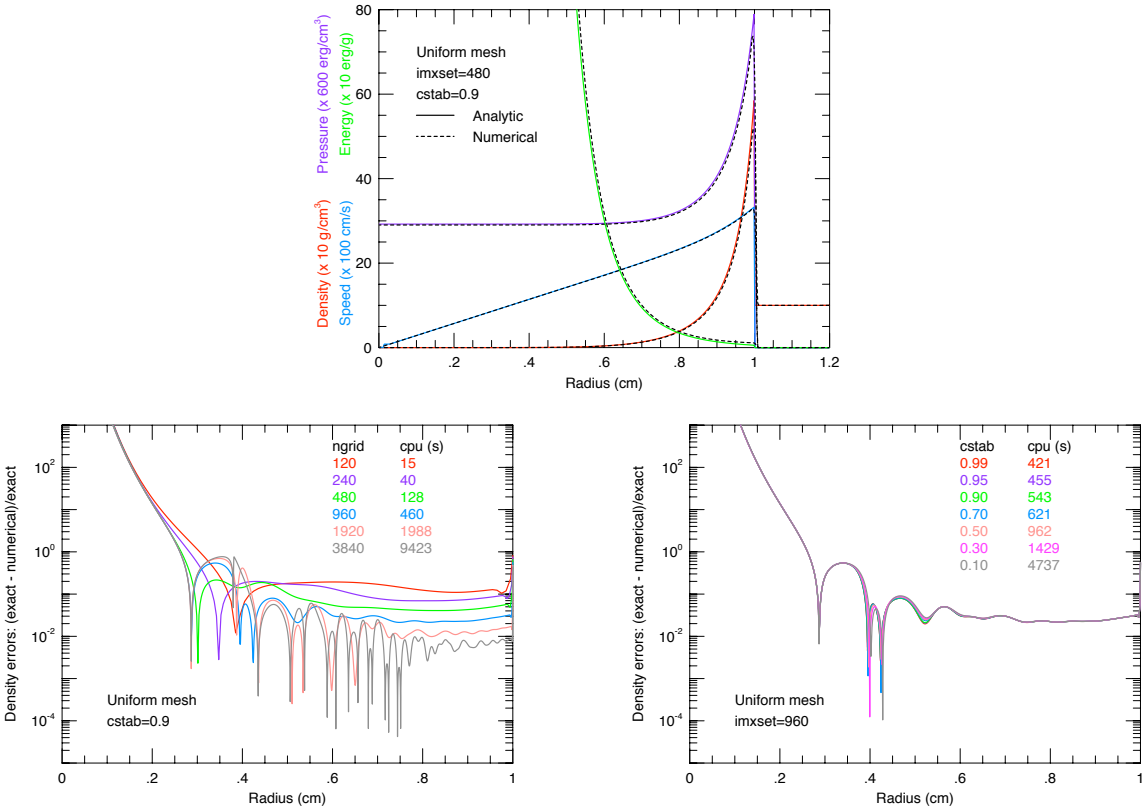

Figure 1. - Summary of the Sedov problem in 1D. Shown are numerical and exact solutions (upper right), relative errors in the density for a variety of uniform grids at a fixed time-step control value (lower left), and the relative error in the density for a variety of time-step control values on a fixed uniform grids.

The upper panel of Figure 1 displays a representative solution on a 480 cell uniform mesh with a time-step controller of cstab=0.9 at the final time of t=1.0 s. The parameter cstab sets the time-step based on the local sound speed and the material velocity, Δt = cstab $\cdot$ Δx / (c + max($|v_x| + |v_y| + |v_z|$)), and determines the time-step in

the numerical solution of the Sedov problem. Initialization of the Sedov problem never fails to generate a spirited discussion whose antagonists are divided between depositing all the energy into a single central zone or depositing the energy in a small fixed size region. While the one-cell case is perhaps a more authentic way of initializing the problem, it is rarely seen in the refereed literature (Reile & Gehren 1991; Buchler et al 1997; Fryxell et al. 2000; Swesty & Myra 2006). For the 1D results discussed in this report, the single-zone initialization was used. Results for both initialization procedures on 1D grids are discussed in Timmes, Fryxell, & Hrbek (2006). Exact and numerical solutions are overlayed in Figure 1 for the mass density, pressure, specific internal energy, and material speed. The analytic and numerical solutions generally agree at this view-graph norm level of comparison.

The image in the lower left of Figure 1 displays the absolute value of the relative errors in the density for a variety of uniform grids at the final time of 1.0 s. The relative cpu cost on a single processor of increasing the spatial resolution is given. The singularity at the origin means T($r \rightarrow 0$) $\rightarrow \infty$, implying large errors. For radii $\gtrsim$ 0.4 cm, there is a steady decline in the magnitude of the errors between the origin and the shock front as the spatial resolution is increased. Figure 1 suggests that for $r \gtrsim 0.4$ cm, the density (and other quantities) have roughly linear convergence rates ($\alpha \sim 1$). Including the $r \lesssim 0.4$ cm region leads to smaller convergence rates because of the persistent errors near the origin.

The image in the lower right of Figure 1 displays the absolute value of the relative errors in the density for a variety of time-step control values on a uniform grid of 960 cells. The relative cpu cost on a single processor of increasing the temporal resolution is shown. Figure 1 implies that the density (and other quantities) have a convergence rate near zero, which suggests that the spatial discretization may dominate the error budget.

In order to use equation (2) all the solutions must be put on the same mesh. One choice is to spatially average the medium and fine gridded solutions onto the coarse mesh. Another is to prolong the medium and coarse mesh solutions to the same uniform fine mesh. Both directions were tried. First we used the coarse to fine interpolation. This tended to produce unphysical jumps in the estimated exact solution as the minimization procedure switched from being on a monotonic convergent solution branch to being on an oscillatory convergent branch. We then used the fine to coarse averaging. This gave smoother estimated exact solutions, but still displayed some jumps in the solution in the post-shock region. These were identifed as being regions where there was a switch from being monotonically convergent to oscillatory convergent. While three grids are then the minimum number of grids necessary to specify the three unknowns at each grid point, one could use more grids to over-constrain the problem, which may weight the solution such that the algorith, stays on a convergent monotonic branch. In the interest of using the minimum number of grids, we experimented by only allowing monotonic convergence. This figures below show the results of this numerical experiment.

Figure 2 provides some results of the optimization procedure with the enforced monoticity constraint for the Sedov problem using 1D uniform grids of 48 (coarse), 96 (medium) and 192 (fine) cells at the final time of 1.0 s. From these three grids we seek an estimated exact solution, the magnitude of the error term, and the convergence rate at each grid point.

The plots on the left display the quality of the estimated exact solution for the density field, while the plots on the right show the same quantities for the pressure field. The upper panel details the convergence rate. The density generally shows second-order convergence behind the shock while the pressure has $\alpha \sim 0.5$. Both the density and

pressure have linear converge rates near the shock front. Zeroth order convergence is attained by both fields ahead of the shock because the initial conditions and the analytical solutions precisely agree. A Cusp appears in the point-wise convergence rate for the density field because the estimated exact solution first smaller than, and then greater than, the numerical solution on the reference grid.

The second panel from the top gives the magnitude of the pre-factor A in the error model (see equation 1). The third panel from the top shows the relative error of the fit (see equation 2), which should be small since there are three unknowns and three equations for each grid point. The fourth panel from the top shows the density (upper left) and pressure (upper right) fields on the three grids, the estimated exact solutions for the fields, the analytical solution for the fields. An analysis of this panel, the desired product, will be discussed in more detail below. The bottom panel displays the error between the estimated exact solution and the numerical solution on the three grids.

The plots in the lower half of Figure 2 compare the estimated exact solution for the density (left) and pressure (right) with the analytical solution in the vicinity of the shock front. In general, the estimated solutions produce a shock front of nearly the same strength at nearly the same location as the analytic solution. The estimated exact solution tends to be larger in the immediate post-shock region and to exhibit spikes. Of course, the most notable feature of the estimated exact solutions is the negative density, negative pressure region ahead of the shock front. This non-physical feature appears because each grid resolution gives a different physical width to the shock front; the coarsest grid spreads the shock front over a greater distance than the finer grids. This region just ahead of the shock violates our enforced monotonic convergence constraint, forcing the minimization technique to yield a dip in the density ahead of the shock front.

Figure 3 explores the grid dependence of the estimated exact solution by duplicating the previous calculations on a more spatially refined set of three grids, 192 (coarse), 384 (medium), and 768 (fine) cells. The plots on the left depict the quality of the estimated exact solution for the density field, and the plots on the right show the same information for the pressure field. As in the (48, 96, 192) grid triplet, the density generally shows second-order convergence behind the shock while the pressure has $\alpha \sim 0.5$. There is some ``chatter'' in the pressure field behind the shock which is perplexing; the estimated exact solution is oscillating erratically about a good value (relative error $\sim 10^{-5}$). The cliff in convergence rate and dip in the solution around error around 0.4 cm may be due to a density in the numerical solution hitting a some small fixed value. The plots in the lower half of Figure 3 compare the estimated exact solution for the density (left) and pressure (right) with the analytical solution in the vicinity of the shock front. In general, these estimated exact solutions are much smoother than those for the (48, 96, 192) grid triplet in Figure 2. Unfortunately, the unphysical dip in the density and pressure ahead of the shock front is still present. This is probably an unavoidable consequence of enforcing montonic convergence for a discontinuity.

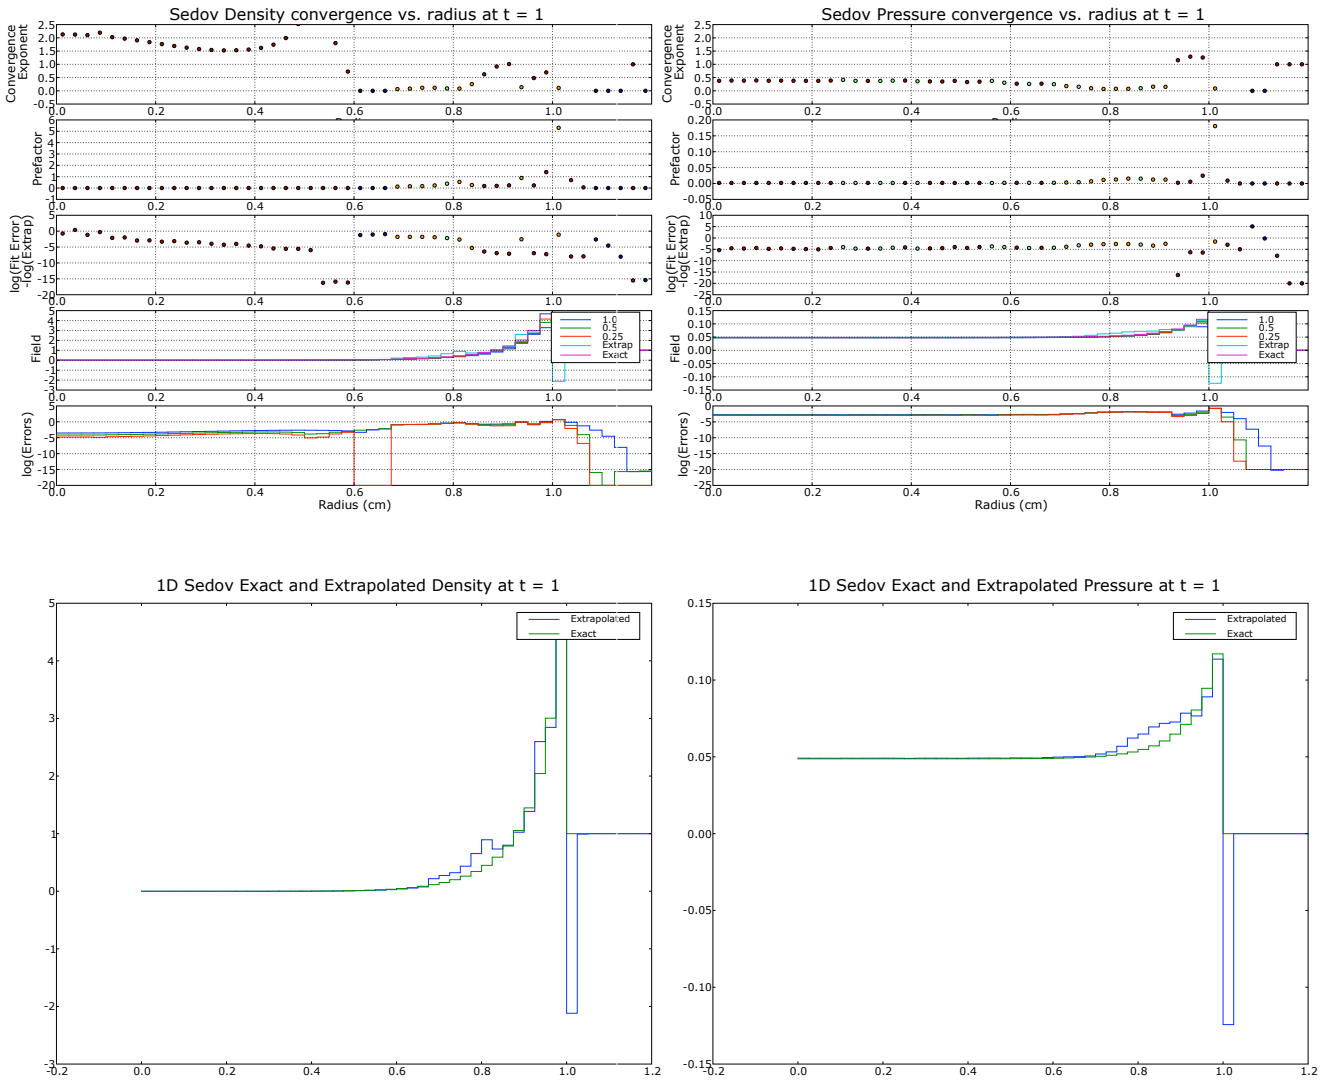

Figure 2. - Summary of the estimated exact solution on grids with 48, 96, and 192 cells at the final time of t=1.0 s. The plot in the upper left pertains to the quality of the estimated exact solution and convergence rate for the density field, while the plot in the lower left compares the estimated exact solution for the density with the analytical solution in the vicinity of the shock front. The plots on the right display the same information for the pressure field. The negative density and pressure ahead of the shock may be a consequence of enforcing montonic convergence for a discontinuity.

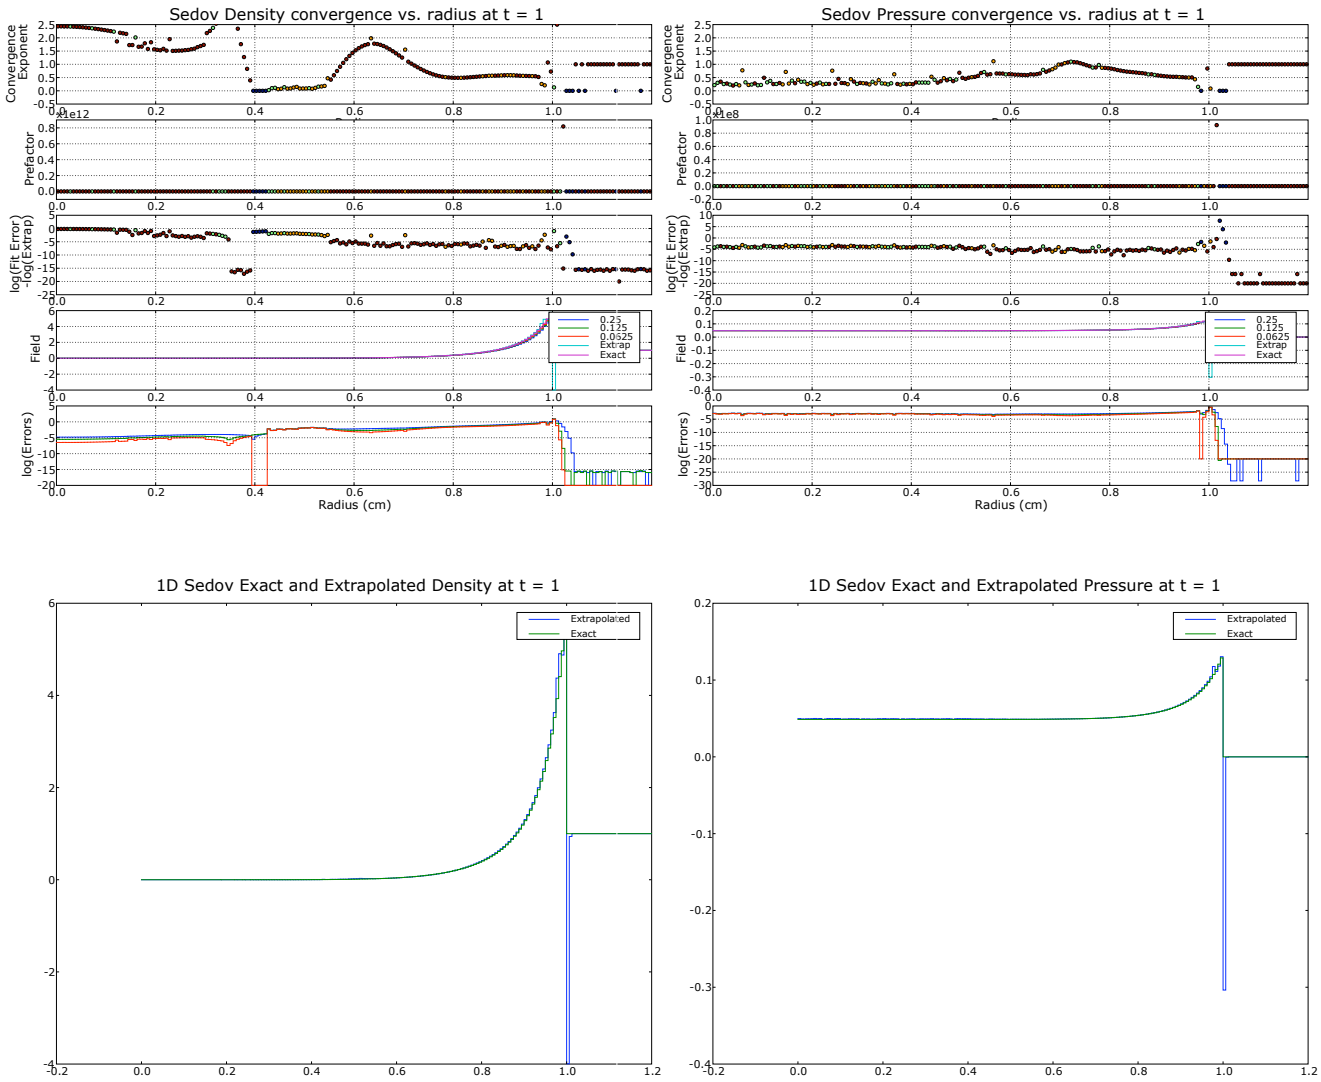

Figure 3. - Summary of the estimated exact solution on grids with 384, 768, and 1536 cells at the final time of t=1.0 s. The plot in the upper left pertains to the quality of the estimated exact solution and convergence rate for the density field, while the plot in the lower left compares the estimated exact solution for the density with the analytical solution in th vicinity of the shock front. The plots on the right display the same information for the pressure field. The negative density and pressure ahead of the shock may be a consequence of enforcing montonic convergence for a discontinuity or a limitation of the technique for this type of problem.

3. Discussion and Future Directions

This research note has described a new optimization technique for conducting a verification analysis on complex multi-physics problems that do not have an exact solution. Previous efforts along these lines considered a Newton-like root-find technique to calculate a point-wise estimation of the exact solution that is mathematically consistent with the convergence properties of the numerical simulations. The new minimization technique was applied to the venerable Sedov problem to derive an estimated exact solution and a self-consistent convergence rate on 1D uniform grids. For a triplet of coarse grids, the estimated exact solution were shown to produce unphysical answers ahead of the shock front, suggesting the grids were too coarse for this calculation verification technique, and/or that the numerical simulations were under-resolved. For a triplet of finer grids (three times the resolution), the estimated exact solution displays a shock front of very nearly the same strength and spatial location as the analytical. A discussion of other problems is reported elsewhere.

In the near future the root-find and minimization techniques should be compared on a series of test problems, with and without analytical solutions. Experiments that over-constrain the problem with more grids should be explored, although there are serious cpu and social reasons to stick with the minimum number of grids necessary. Experiments that explore how both techniques performs on problems that display pure oscillatory convergence are needed to better assess how these methods perform for non-smooth problems (i.e., shock fronts and material interfaces). We anticipate continuing to apply such calculation verification methods to an increasing number of complex multi-physics problems on 1D adaptive and 2D axisymmetric uniform meshes. This technology represents a key growth direction for verification efforts to bridge the gap between analytical test problems and complex applications.

4. Acknowledgments

This work was supported by Jerry Brock, Kim New, and Scott Doebling. The reasoned technical input of Palmer Smitherman, Bill Rider, and Francois Hemez was helpful in working through the details of this investigation. Los Alamos National Laboratory is operated by the Los Alamos National Security, LLC for the National Nuclear Security Administration of the U.S. Department of Energy under contract DE-AC52-06NA25396.

5. References

Brock, J., ``Isolating Temporal-Discretization Errors for Separate-Verification Analysis", AIAA Aerospace Sciences Conference, January 2004, AIAA-2004-0741, LA-UR-03-9160.

Buchler, J.R., Kollath, Z., & Marom, A., ``An Adaptive Code for Radial Stellar Model Pulsations" , Astrophys. Space Sci. 253, 139, 1997.

Celik, I., Li, J., Gusheng, H., Shaffer, C., ``Limitations of Richardson Extrapolations and Some Remedies" July 2005, Journal of Fluids Engineering, 127, 795.

Fryxell, B., Olson, K.,Ricker, P., Timmes, F.X., Zingale, M., Lamb, D.A., MacNeice, P., Rosner, R., Truran, J.W., & Tufo, K., ``FLASH: An Adaptive Mesh Hydrodynamics Code for Modeling Astrophysical Thermonuclear Flashes", Astrophys. J. Suppl. 131, 273, 2000.

Grove, J., AMHCTOOLS, LA-CC-05-052, 2005a.

Grove, J., AMHCTOOLS, LA-UR-05-7425, 2005b.

Kamm, J.R., ``Evaluation of the Sedov-von Neumann-Taylor Blast Wave Solution", LA-UR-00-6055, 2000b.

Nash, S.G., ``Newton-like minimization via the Lanczos method", SIAM J. Numer. Anal. 21, 770, 1984.

Nash, S.G., ``User's guide for TN/TNBC", Technical Report 397, Department of Mathematical Sciences, The Johns Hopkins University, Baltimore.

Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P., ``Numerical Recipes in FORTRAN: The Art of Scientific Computing", Second Edition, Cambridge University Press, New York, NY, p. 376, 1992.

Reile, C., & Gehren, T., ``Numerical simulation of photospheric convection in solar-type stars I. Hydrodynamical test calculations", Astron. Astrophys. 242, 142, 1991.

Smitherman, D.P., Kamm, J.R., Brock, J.S., ``Calculation Verification: Pointwise Estimation of Solutions and Their Method-associated Numerical Error", LA-UR-05-8002, 2005.

Sedov, L.I., ``Similarity and Dimensional Methods in Mechanics", Academic Press, New York, 1959.

Swesty, F.D., & Myra, E.S., ``A Numerical Algorithm for Modeling Multigroup Neutrino-Radiation Hydrodynamics in Two Spatial Dimensions", Astrophys. J., submitted, 2006.

Taylor, G.I., ``The formation of a blast wave by a very intense explosion", Proc. Roy. Soc. London A 201, pp. 159-174 (1950).

Timmes, F.X., Gisler, G., Hrbek, G.M., Automated Analyses of the Tri-Lab Verification Test Suite on Uniform and Adaptive Grids for Code Project A, LA-UR-05-6865, 2005.

Timmes, F.X., Fryxell, B., Hrbek, G.M., ``Spatial and temporal convergence properties of the Tri-Lab Verification Test suite for code project A", LA-UR-06-6444, 2006a.

Timmes, F.X., Fryxell, B., Hrbek, G.M., ``Two- and Three-dimensional properties of the Tri-Lab Verification Test Suite for Code Project A", LA-UR-06-6697, 2006b.