

Using Adaptive Mesh Refinement to Study Grid Resolution Effects for Shock/Boundary-Layer Interactions

Michael E. Olsen*

NASA Ames Research Center, Moffett Field, CA 94035

Randolph P. Lillard†

NASA Johnson Space Center, Houston, TX 77058

Adaptive Mesh Refinement (AMR) promises a much more computationally efficient means to obtain a discrete approximation to a continuous boundary value problem of a specified accuracy than classic isotropic grid refinement. The AMR capability of OVERFLOW is utilized to provide estimates of the exact analytical solutions to problems of interest for turbulence modeling. Predictions of surface pressure and skin friction, essentially the state of stress at the surface, shows little difference with grids believed to be “grid resolved.” Velocity profiles, on the other hand, show marked differences in flows with shocks. The AMR method, as implemented in OVERFLOW 2.2k, appears to provide the ability to produce arbitrarily accurate solutions at a predictable cost much smaller than classic uniform mesh refinement.

I. Nomenclature

r	=	radial position
r_0	=	model surface radius
C_p	=	pressure coefficient $\frac{p-p_\infty}{\frac{1}{2}\rho_\infty u_\infty^2}$
c_{f_x}	=	skin friction coefficient $\frac{\tau_w}{\frac{1}{2}\rho_\infty u_\infty^2}$
c	=	Bachalo-Johnson bump chord
L_r	=	refinement level (=NBREFINE OVERFLOW parameter)
N_0	=	number of points in baseline grid
N_p	=	number of points in AMR refined grid
N_c	=	number of points in uniformly refined grid
M_∞	=	free stream Mach number $\frac{u_\infty}{a_\infty}$
u_∞	=	free stream velocity
a_∞	=	free stream sound speed
ρ	=	mass density
p	=	pressure
e	=	internal energy $= \int_0^T c_v dT$
u	=	axial velocity
x	=	axial distance
r_2	=	solution residual
p_t	=	total pressure
τ_w	=	wall shear stress

II. Introduction

Reynolds Averaged Navier-Stokes (RANS) techniques will continue to be indispensable for the foreseeable future to solve complex aerospace design and analysis problems, but an implicit requirement is that the RANS simulations are “grid resolved.” The addition of Adaptive Mesh Refinement (AMR) to OVERFLOW provided, in principle, a means to

*Research Scientist, N230-2, Moffett Field CA 94035-1000, Associate Fellow AIAA

†Research Scientist, Johnson Space Center, Houston TX 77058

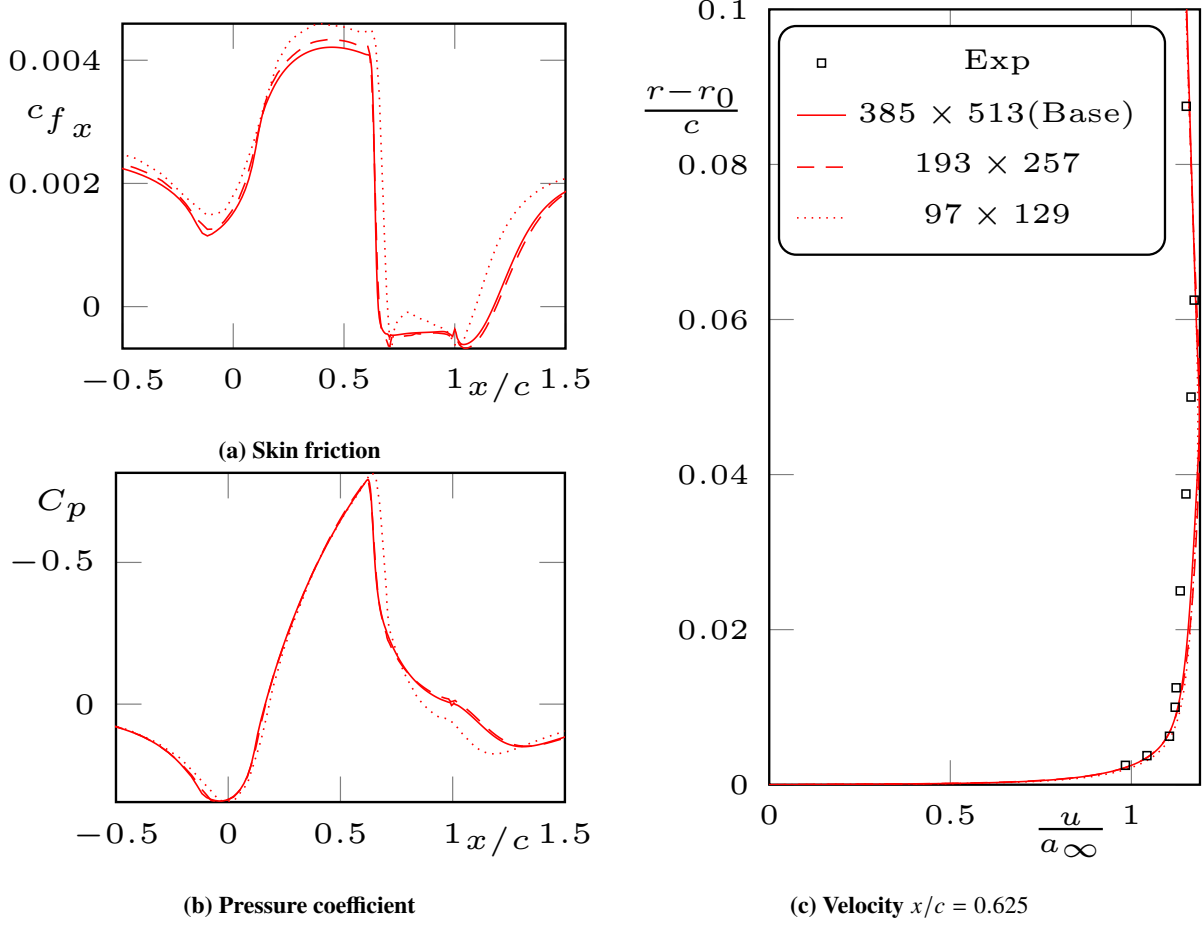


Fig. 1 Bachalo-Johnson bump apparent grid convergence—Lag- ν_T turbulence model

obtain an arbitrarily accurate solution at much lower cost than classic mesh refinement. While most uses of CFD may not justify ensuring every solution is grid resolved, when creating turbulence models, having this ability is a necessity.

For a two-dimensional flow, each successive level of uniform grid refinement quadruples the number of grid points, and simultaneously stiffens the numerical system to be solved. Multigrid is, in principle, useful for ameliorating the stiffness issue, but at least a sixteenfold increase in the computational work is required to obtain a solution with two levels of uniform grid refinement. For three-dimensional flow, this would be a sixty-four fold increase in workload to drop the grid spacing by a factor of 4. AMR promises to provide the ability to obtain solutions on grids that are, in principle, arbitrarily fine. The question addressed here is whether, in practice, it can provide these solutions at a reasonable computational cost.

As separation prediction is the current focus of work, three representative separated flowfields are investigated—one supersonic, one transonic, and one low speed. Eddy viscosity turbulence models with extensive history of use and good prediction capabilities are used to provide a representative sample so that the behavior of the AMR can be assessed in a way that is independent of any single models quirks, but relevant to current practice. The results indicate that AMR can produce answers to the posed boundary value problems to arbitrarily high accuracy.

III. Method

The flow solver used was the 2.2k version of OVERFLOW. Two versions of this code were employed, one the stock version as distributed, and the other a modified version with the Lag- ν_T [1] turbulence model implemented. The results shown in this paper utilized matrix dissipation [2, 3]. As the grids are refined, the details of the numerical smoothing become less and less relevant for any numerically consistent scheme, but the lower dissipation inherent in

matrix dissipation should allow grid independence with less grid use than more dissipative schemes. One detail to note is that for the transonic and supersonic cases, the eigenvalue limiter values were nonzero [2], whereas for the subsonic case, the eigenvalue limiter values were zero [3], which should allow grid independence earlier for skin friction for this subsonic case.

The adaptation used is described in detail in [4, 5], with all the grids here using only the near-body adaption described in [5]. The sensor used is the second undivided difference function [4],

$$S = \max_{i=j,k,l} \left\{ \max_{Q \text{ components}} \left[\left(\frac{q_{i-1} - 2q_i + q_{i+1}}{2q_{\text{ref}}} \right)^2 \right] \right\} \quad (1)$$

where $Q = [\rho, \rho u, \rho v, \rho w, \rho e_0]$, the conservative mean variables at each grid point. The reference divisors for these variables are $q_{\text{ref}} = [\rho_\infty, \rho_\infty u_\infty, \rho_\infty v_\infty, \rho_\infty w_\infty, \rho_\infty e_{t\infty}]$. This function is a measure of the linearity of the solution. It is the ratio of the 2nd derivative of the conservative variables (without the associated grid metrics) in computational space to their freestream magnitude. This goes to zero as the grid is refined and linear interpolation is a better and better approximation of the solution between grid points.

The refinement is controlled by two coarsening/refinement parameters which choose which regions of the grid will be either refined (EREFINE), if the sensor function exceeds this value, or coarsened if it is below (ECOARSEN). In this paper, the coarsening parameter was set to a value small enough (ECOARSEN = 10^{-28}) that the grids are never coarsened. The refinement parameter was set at a value small enough (EREFINE = 10^{-7}) to continue to refine the grid so that, at least to two levels of refinement, the results mimicked classical grid refinement on the Bachalo-Johnson bump flow.

The overset capability of OVERFLOW is utilized to accomplish the near-body grid refinement. Regions that require refinement are replaced by locally doubled grids. This refinement is done in computational space, and refines even the grid lines defining the walls using cubic spline interpolation with limiters to eliminate oscillations at grid corners with slope discontinuities.

The baseline reference solutions are obtained using classic relaxation techniques [2]. Specifically, multigrid is utilized and the solution residuals are brought down over eight orders of magnitude. Once the baseline solution is created, then this is used as the starting solution for grid refinement. An adaption cycle is started with the refinement level NBREFINE increased by one (the baseline grid is NBREFINE = 0), and the grid is adapted every 1,000 steps ten times. Then, the solution is adapted twice more with 5,000 steps between adaptations. Finally the solution is converged without adaption for an additional 10,000 steps to completely converge the solution on the final grid at this level of refinement. This algorithm was used for both the transonic and supersonic solutions, but the subsonic solution used 2,000 iterations between adaptations in the first ten cycles, as the lower Mach number solutions required more iterations. Low Mach preconditioning was used on the subsonic case, but not on the other transonic and supersonic cases.

All adaption cycles did *not* use the multigrid functionality within OVERFLOW (MULTIG = `.false.`), as this could cause the solution to diverge. In all these cases, time to solution was not optimized; rather, a procedure that would reliably converge the solution at increasing levels of refinement was chosen. The method described above allowed convergence with grid refinements 6 levels higher than the baseline grid on all flowfields reported. As the starting solution for the next level of grid refinement was a fully converged solution at the previous level of refinement, this can be thought of as a form of multigrid, “full multigrid” in OVERFLOW parlance, or sometimes referred to as grid sequencing. Essentially, the lower frequency error has been eliminated for the most part before the next finer grid solution is attempted.

The SA (more precisely SA-noft2[6]) and SST (the 1994 version considered “standard”[6]) turbulence models in OVERFLOW have been verified to be coded consistently [7] using the turbulence modeling resource website [6]. The Lag- ν_T model is verified to be consistently coded as the code is modified in the course of turbulence model development. Verification of the implementation was redone at the beginning of this work, as described in the start of the results section.

IV. Results

A. Bachalo-Johnson Bump

On some iconic flowfields, grid spacing used in previous work was arrived at by classic grid refinement, to the point where further refinement produces changes that are small enough to discourage further refinement. An example of this would be the Bachalo-Johnson [8] transonic bump flow. With judicious grid creation, a solution that would be




Fig. 2 Insensitivity of surface pressure (C_p) to mesh refinement on Bachalo-Johnson bump





Fig. 3 Insensitivity of skin friction (c_{f_x}) to mesh refinement on Bachalo-Johnson bump

considered grid converged can easily be obtained with 358×513 grid points. This grid system was devised as a baseline for the TTR [3, 9, 10] turbulence model, a model which can require more grid to obtain truly grid converged solutions, and was believed to be essentially grid converged for eddy viscosity models.


A comparison of solutions for surface pressure, skin friction and a velocity profile just upstream of the shock/boundary-layer interaction region is shown in Fig. 1. This solution is for the eddy viscosity Lag- ν_T model, and the results are in line with previous work. More intricate models could potentially have more stringent grid requirements, and investigating the actual grid requirements for such models was the genesis of this paper.

In the process of preparing solutions for a Reynolds-stress model, this grid was utilized as the baseline, and as an excuse to bring AMR to its extreme conclusion, the solution refinement was carried out to eight levels. It was expected that diminishing alterations would set in and by at most four levels of refinement, the process would have concluded, at least to plotting accuracy. This expectation was not met, and fear that a bug had been introduced in the current implementation of the Reynolds-stress model replaced the initial complacency. However, model implementation error was ruled out when the model, as currently coded, reproduced previous results on a variety of flowfields. The next possibility investigated was that this was a previously unseen emergent feature of the Reynolds-stress model. Investigating this possibility, eddy viscosity model solutions were obtained using the same AMR methodology. It was expected that after a few grid refinement levels, boredom would ensue, but again—surprisingly large changes persisted in the velocity profiles near the shock/boundary-layer interaction region with the simpler eddy viscosity models.


If the stopping criteria were surface pressure and skin friction, or velocity profiles outside of the vicinity of the shock, the refinement process would have indeed stopped with the baseline grid declared “grid converged.” For example, the surface pressure predictions for three eddy viscosity models, the Spalart-Allmaras [11] model, the Menter SST [12] model and the Lag- ν_T [1, 13] model, are shown in Fig. 2. Note that in this figure, the *coarsest* grid solutions (the “baseline grid”), are the *finest* grid solutions in figure Fig. 1, and there is no discernible difference among grid resolutions, even when the AMR is allowed to refine the grid resolution up to 64 times finer than the baseline solution.




(a) Baseline grid (every 4th grid line)



(b) AMR grid system, refinement level 3



(c) Solution (ρ) refinement level 3



(d) Solution (ρ) near bump, refinement level 3

Fig. 4 Bachalo-Johnson bump grid systems and solution

The same indifference to grid resolution is shown by the surface skin friction for these models as can be seen in Fig. 3. One might infer from this that the baseline grid, Fig. 4a, is indeed sufficiently fine for these models.




Fig. 5 Velocity profiles showing the effect of AMR

Looking at velocity profiles upstream of the interaction region (Fig. 5a) would lead to the same conclusion. Note that this figure includes results from three different turbulence models, as well as from effective grid resolutions differing by a factor of eight, and the solutions agree extremely well. Looking at a profile downstream of the shock/boundary-layer interaction region (Fig. 5c) shows more difference in the predictions of the three turbulence models (as one might expect), but there are relatively small differences due to this sizeable grid resolution change. However, in the region of the shock/boundary-layer interaction, there is a significant effect of increased grid resolution, and both Lag- ν_T and Menter-SST show what appears to be a discontinuity in velocity beginning to appear.

The refined grid system, at three refinement levels, is shown in Fig. 4b, and the density profiles in the converged solution at this level of refinement are shown in Fig. 4c. The density field along with the grid system close to the bump is shown in Fig. 4d. The various levels of the refinement grids are clearly visible in the closeup picture, with the finest grid refinement highlighting the shock structure. It is also interesting in this picture to note the areas that have not been refined to as high a degree. Recall that the AMR strategy chosen should never coarsen regions, so the unrefined regions here have never required further refinement during the entire AMR solution procedure to this level of refinement.

What is happening is shown in Fig. 6, where the Mach contours of the supersonic bubble of the Bump flow are compared for the three eddy viscosity models over 4 different effective grid densities. The shock structure evident in the lower row (which should be equivalent to a grid dimension of 3073×4097) is suggested by the Mach contours of the top row, which is the already fine baseline grid (385×513). As a point of reference, in the paper introducing the Lag- ν_T model, the finest grid run for this same test case was 358×161 , smaller in both directions.

All three eddy viscosity model solutions contain a λ shock system, with two very weak shocks coalescing above (for the Lag- ν_T and SST models) or at (for the SA model) the edge of the incoming boundary layer. The boundary layer edge is visible in Fig. 6 as the knee in the Mach number contours, where the boundary layer dissipation has changed the total enthalpy, affecting the sound speed below the boundary layer edge. Moving from the wall to the freestream, contours of constant Mach number run upstream below the boundary layer edge, and downstream above the boundary layer edge. This creates a very visible “knee” in the Mach number contours for all three turbulence models. For the

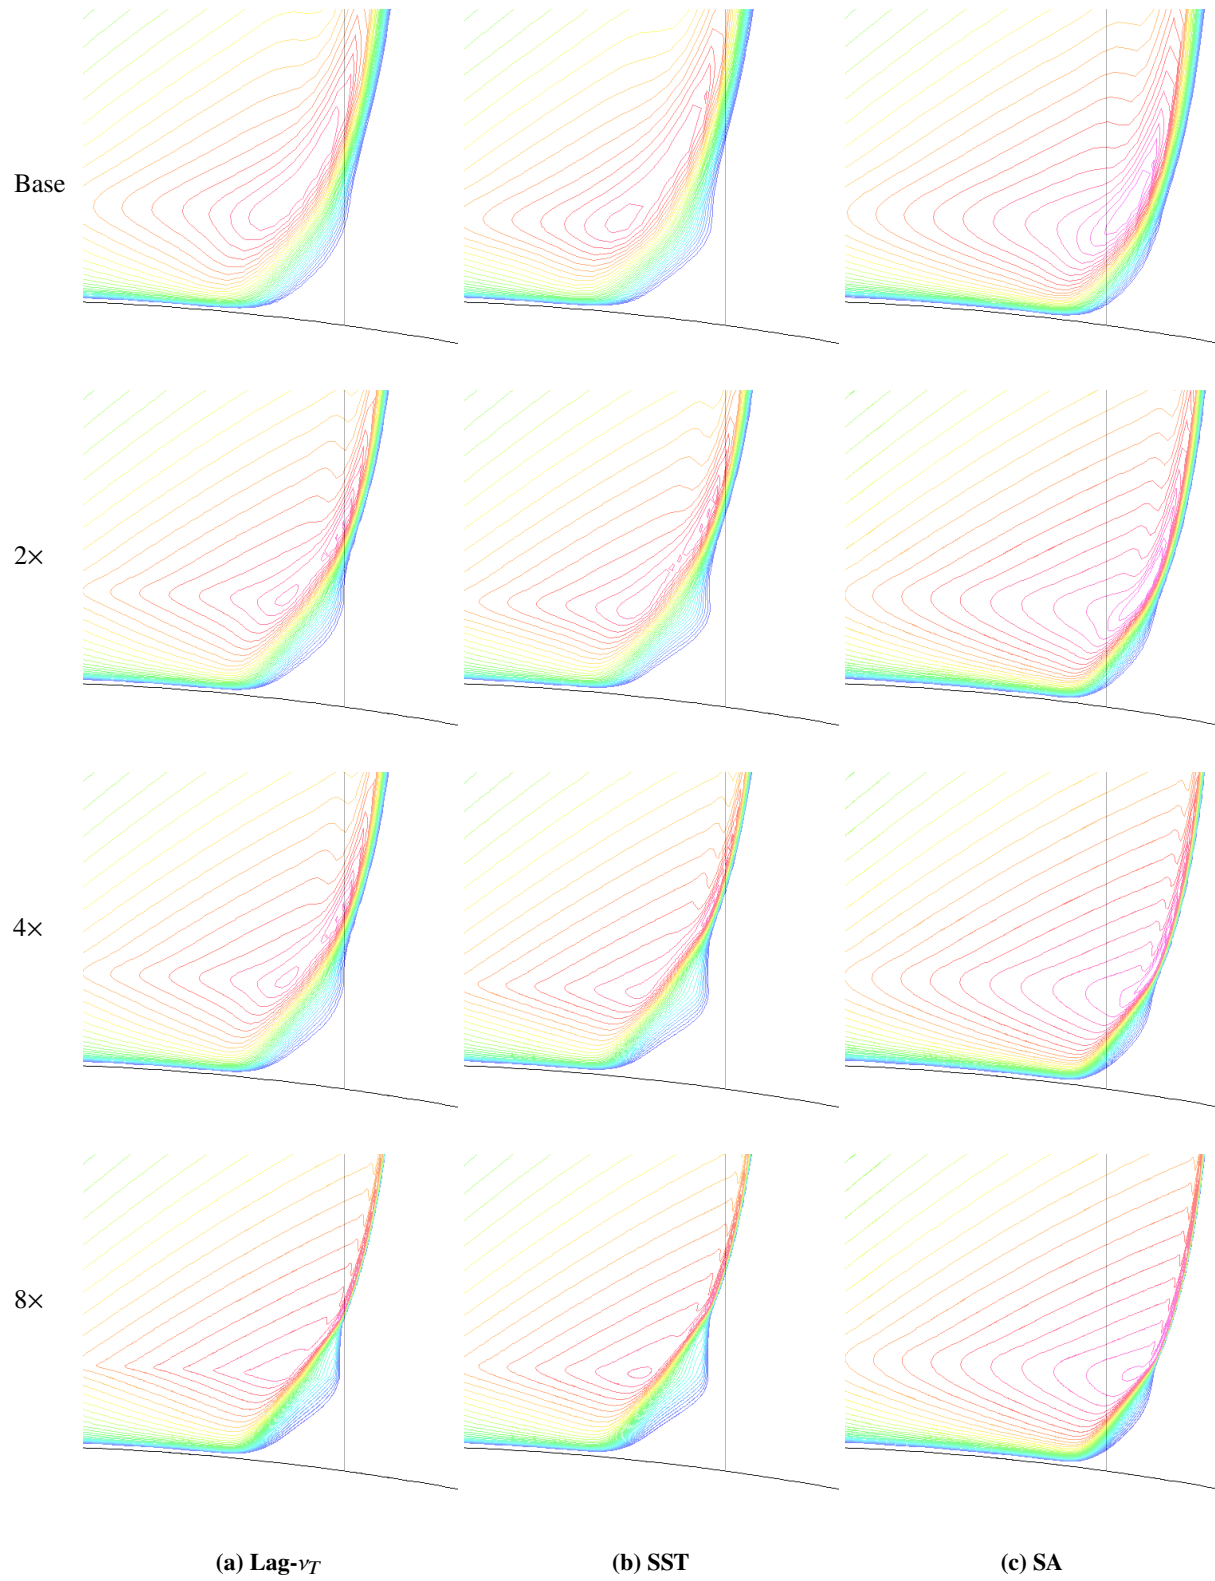


Fig. 6 Increasing mesh refinement on Bachalo-Johnson bump—supersonic bubble Mach number contours (vertical line shows $x/c = 0.688$ profile location)




Fig. 7 Comparison of AMR to uniform grid refinement as function of refinement level

models where the shocks join above the boundary layer edge, the result is an inviscid shock system with two very weak shocks and supersonic flow upstream between the two legs of the λ .




Fig. 8 Baseline grid residual

Another way of looking at this is presented in Fig. 7b, which is a plot of the ratio of the AMR grid size to the grid size from uniformly doubling the entire grid. At six levels of refinement, the AMR grid is less than 5% of what uniform refinement would entail. This implies that the sixth level AMR refinement grid has only slightly more grid points than the number resulting from two levels of uniform refinement.

When the solutions are compared away from the lambda shock region, at least for the eddy viscosity models studied in this paper, the differences diminish rapidly. For this flowfield, the region of largest change with grid refinement is the outer boundary layer edge and the inviscid flow immediately beyond. This makes the influence on the solution smaller than it could be if the changes were closer to the wall. By the same token, the only way to assess how resolving these features will affect the solution is by obtaining solutions like these where the features are resolved.

While for these eddy viscosity models, the change in the solution is dramatic only in a limited region near the shock system, a larger question is the relative cost. There is no way to know how far the solution is from the exact answer without performing a grid convergence study. Using uniform refinement to reach the solution resolution that is attainable with AMR is cost prohibitive even in 2D, where the cost (in grid size) of uniform grid refinement goes as 4^{L_r} with refinement level L_r . For this case, AMR is a dramatic improvement over

classic grid refinement. Figure 7a shows the grid size growth as a function of the level of refinement, giving the ratio of the converged grid size at each refinement level to the initial grid size. The solid black line on this graph is for classic grid refinement. At 6 levels of refinement the grid size with AMR is over 100 times the initial grid size, but it is far below what would be required with global grid refinement which is literally off the chart in this figure. Furthermore, the dashed gray line (adjusted to match up with the trend of the higher grid levels for SA and Lag- ν_T) of linear, as opposed to quadratic, growth implies that the improvement will continue to be dramatic. This linear (in grid spacing) growth could be because the structure that is continuing to be refined is the shock, which is not actually a 2D region, but rather a 1D “line.” For 3D solutions, this would then imply that, for cases where a shock is the limiting structure being refined, the growth would be 4^{L_r} , as opposed to 8^{L_r} , again a major improvement in efficiency, at least in regards to grid size.

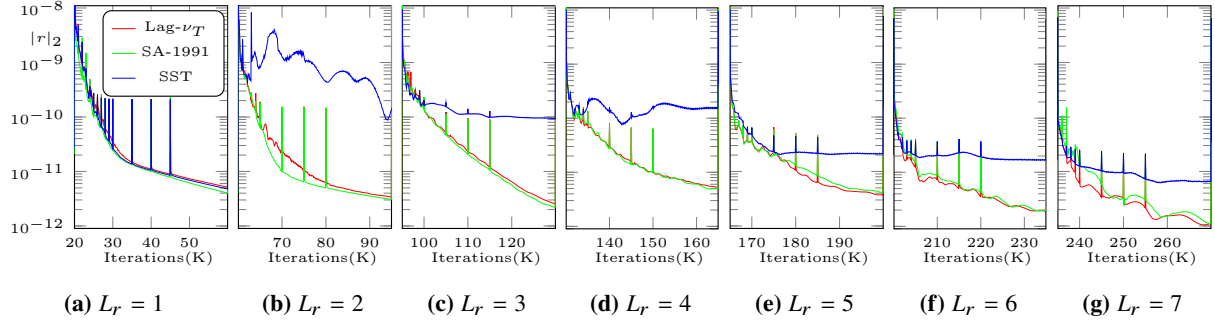


Fig. 9 Residual history as grid refinement level L_r is increased

The actual cost is not just the grid size, but also the iterations required to obtain the grid refinement and the additional iterations to converge the solution. There is concern that increasing grid refinement will slow convergence [4, 14]. At least for the relaxation choices used in this study, this concern was not borne out. Granted, the cost of obtaining the finer grid solutions does require both the iterations with adaptation cycles to obtain the refined grid itself and then more iterations to completely converge the solution on that new grid system. However, the total number of iterations required is comparable to the number of iterations used to obtain the baseline grid solution, and independent of the grid level (Figs. 8 and 9).

As the relaxation parameters used are very much what would be chosen normally (with the proviso that local multigrid has been dispensed with, as it could cause solution divergence above two levels of grid refinement), this is a heartening development. It bears repeating that obtaining the fastest convergence possible was not the guiding principle in choosing the relaxation strategy. It is certainly possible that coarser grid systems can be converged more rapidly than the techniques used in this work. The finer AMR grid systems may not allow the more aggressive relaxation parameter choices feasible on the initial coarser grids. However, in practice, having obtain the same efficiency used on the initial coarser grids. However, in practice, having on a fixed problem, is a more likely use environment for these methods, so the results could be more universal than the pessimism would dictate. A final point to note here is that these solutions were driven to extreme levels of convergence at all grid refinement levels, where the turbulence models would allow it. This is standard practice in our turbulence model development, but is not standard practice for many other uses of CFD.

The last factor of the cost, however, is the actual CPU time per grid point. This could conceivably rise with all the extra machinations required as more and more grids are created and the total number of blanked points increases with higher grid refinement. Such is not the case, as this number proves remarkably constant at least up to six levels of refinement. The cases were run without increasing the number of compute nodes, and a fixed number of MPI processes (from 4 to 7) and OpenMP threads (generally 4). On the Pleiades cluster, the CPU time per point for these cases ranged from 50ns (for up to 5 levels of refinement) to 80ns (at eight levels of refinement). For these simulations, the percentage of blanked points was around 4% at one level of refinement and reached 8% with 7 levels of refinement. Overall parallel efficiency (measured) was roughly 90%, at all grid refinement levels, but the solution size increased by over two orders of magnitude from the baseline to the level 7 grids, and the cases still ran on a single node for all grid sizes, which might explain the doubling in CPU time per grid point at the higher refinement levels. The ratio of cache size to total memory used will drop as memory usage increases with refinement level.

The bottom line for this case was that AMR enabled the creation of a grid resolved solution to this boundary value problem that would have been computationally prohibitive in its absence. The cost went roughly as the number of points in the refined grids, and for these turbulence models the problem did not require progressively more fine grid iterations

to obtain each successive grid refinement. Each further level required roughly 35K more iterations to converge it four orders of magnitude for the models that would allow that level of convergence. Finally, the cost per point was roughly constant, rising slightly above 5 levels of refinement. The number of points grew not by a factor of 4 at each grid level, but actually by a factor of 2, a significant savings as the refinement progressed to the next level. In short, this is a viable means of obtaining the “exact” solution, a very powerful tool when crafting turbulence models.

B. Brown, Brown & Kussoy Supersonic Cylinder-Flare

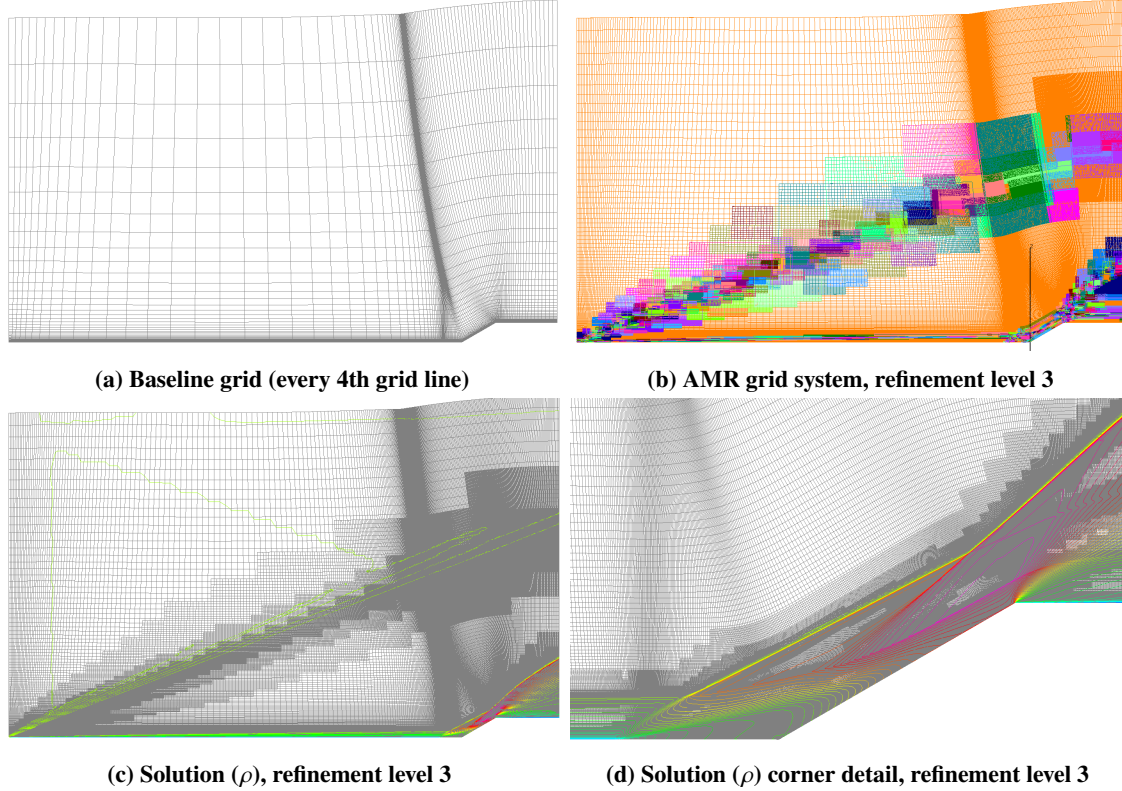


Fig. 10 Cylinder-flare grid systems and solution

The second case considered here is a $M_\infty = 2.89$ cylinder-flare experiment [15] which includes a wealth of measurements, including surface pressure, and LDV measurements. A 2.54 cm cylinder aligned with the freestream flow delivers a well defined flat plate boundary layer to a 30° flare. This case has been investigated by others [16, 17], and was suggested as a test case to the Turbulence Model Benchmarking Working Group. The case is well documented, having been the focus of a long campaign with multiple experimentalists, but led by James Brown over its entirety.

The case is a good test case for turbulence model comparisons, as it has very different solutions between the two industry standard turbulence models [16, 17]. For this study, it is used as a test for the AMR methodology, as it has a relatively rich shock structure and extensive separated flow regions. The baseline grid for this case is 449 (streamwise) \times 257 (wall normal).

The shock system can be seen in Fig. 10c. This grid system does not include the nose of the geometry, but instead starts at 0.7m ahead of the start of the flare, chosen to provide the correct inflow boundary layer height at the first measurement station 0.45m upstream (Fig. 13a). Both the SA and Lag- ν_T models match the velocity profile reasonably well, though neither predicts the nearest measurement point at this station, consistent with [16, 17]. Not consistent with those results, the SST model is separated. In this paper, the strain based production is used, where both other groups chose the vorticity based production.

The grid system with refinement (Figs. 10b and 10c) highlights the shocks in this flowfield, including the one that emanates from the leading edge of the cylinder. This weak shock is actually not consistent with the experiment, but it does not interact with the shock system associated with the flare. The actual nose is an ogive, but given the agreement

with the experimental velocity field at the furthest upstream measurement point, it was not considered necessary to model the exact nose geometry when the baseline grid was constructed. The upstream boundary layer has a slightly different pressure history than that of the experiment, but the differences are believed to be smaller than plotting accuracy at the upstream location.

Boundary conditions are no-slip, adiabatic walls, with a uniform freestream upstream of all the shocks. The outflow boundary condition is simple—zero order extrapolation. The freestream turbulence intensity is what would be consistent with flight conditions in the computations, and has not been matched to the wind tunnel conditions. For these turbulence models, there is not expected to be any sensitivity to freestream turbulence levels.




Fig. 11 Sensitivity of surface pressure ($\frac{p}{p_t}$) to mesh refinement on Brown, Brown, Kusoy flare




Fig. 12 Sensitivity of skin friction (c_{f_x}) to mesh refinement on Brown, Brown, Kusoy flare

For this paper, the flowfield was chosen because it is a well defined, representative separated supersonic flow. There are a large number of shocks and expansion fans that interact (Fig. 10d) along with what should be a relatively large separated region. The shock waves are much stronger than those of the Bachalo-Johnson bump and expansion fans are an important part of the overall structure. How AMR interacts with these features is the focus of this work. The matching of experimental details or the sensitivity of any specific turbulence model to those details is not. The current work is a vetting of the tool of AMR rather than its use for turbulence model development. However, being able to eliminate computational uncertainty will be a powerful help in assessing and quantifying the ability of a given turbulence model to predict this flowfield.

The surface pressure, Fig. 11, and skin friction, Fig. 12, are reasonably well predicted on the baseline grid by the Lag- ν_T models, but the point of pressure rise does move slightly forward with increasing grid refinement. The remainder of the pressure and skin friction are quite well predicted by the baseline grid. Even the skin friction in the separated region ahead of and behind the start of the flare does not change a great deal—including the prediction of the point of reattachment (where c_{f_x} crosses 0 for $x > 0$) as the grid is doubled and doubled again.

However, the velocity profiles paint a different picture. The furthest upstream velocity profile (Fig. 13a) is insensitive for the SA and Lag- ν_T models, as this is an attached supersonic essentially flat plate velocity profile. The SST results for




Fig. 13 Velocity profiles showing the effect of AMR

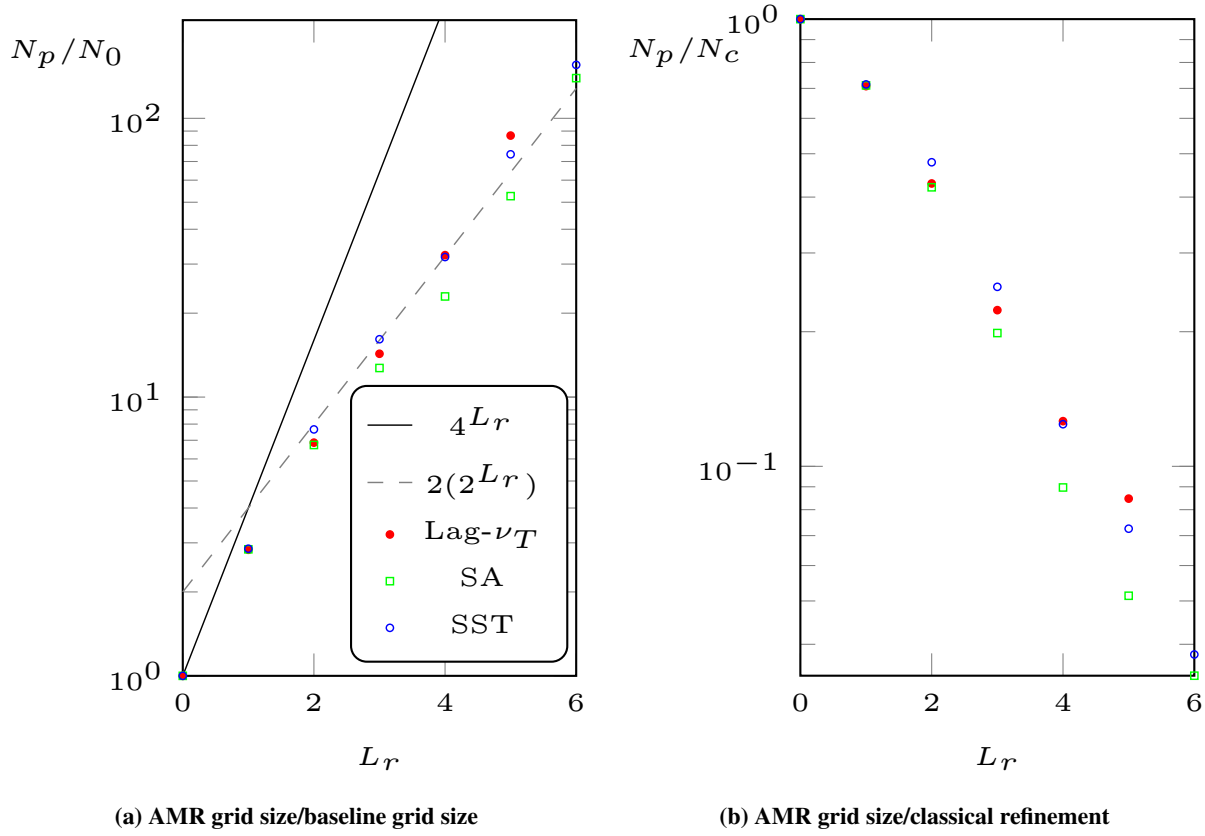


Fig. 14 Comparison of AMR to classic grid refinement for the supersonic flare

this case have a separation that extends significantly further upstream than the experiment, causing the upstream profile not to match the experiment or the other turbulence models. In this figure the range of grid refinement is larger. The baseline (solid lines) is compared with three levels of grid refinement ($\Delta x|_3 = \frac{1}{8}\Delta x|_0$) as dashed lines, and with five levels ($\Delta x|_5 = \frac{1}{32}\Delta x|_0$) as dotted lines. For the attached profiles in the upstream location, all three are indistinguishable, but in the profiles starting at the beginning of the flare ($x = 0$) there is daylight between the baseline grid profiles and those with three levels of refinement. The three and five level of refinement profiles are on top of each other throughout

in this plot. The same description applies to the next profile plotted (Fig. 13c). At the final station with experimental data, there is less difference between the various profiles. Also, the differences between the turbulence model solutions are smaller. The shock is well resolved with the refined grids, where the baseline grid has a smeared shock, but other than near the shock the effect of grid refinement is much less at this far downstream location.

The computational efficiency for this case is the critical question, and thankfully, the experience here mirrors what was found for the transonic flowfield with a single shock embedded in a subsonic flowfield. Again, the grid system grew by a factor a little over 100 (Fig. 14a) when the grid was refined 6 levels, very much in line with the transonic results. This leads to a similar picture for the grid point efficiency (Fig. 14b), with the grid system 3% of what a classically refined grid system would be for this level of refinement.

Again, even though there are more shocks, the structure refined is a curved line—one dimensional, not a two-dimensional region. The equally powerful expansion fans, which are not localized, are apparently not requiring a refinement of the entire domain, which could conceivably have added to the grid point count required in a much more substantial manner. Looking at the detailed picture of the flowfield with the associated grid system in light gray (Fig. 10d), the region above the convex corner at the top has large domains where the density field (and apparently the other conservative variables) do not have large amounts of variation, and thus are not refined in the same manner as the shocks attract. The additional grid is required at each refinement level to improve the solution at the shock discontinuity. In contrast, the regions without shocks reach a level of grid density to match the exact solution derivatives, and no more refinement is required.

In terms of the additional pieces of the efficiency puzzle, the higher levels of grid refinement do not require more and more iterations as higher refinement levels are reached. More overall iterations are needed to obtain the solution for each grid level, but the number of iterations needed on the sixth refinement was essentially the same as needed on the first. Where 35K iterations sufficed for the transonic case, 50K was sufficient (with 2K steps between the first adaptation steps) to handle the higher Mach number case.

C. Driver CS0




Fig. 15 CS0 grid system and solution - 3 levels of refinement




Fig. 16 Baseline grid residuals

The final case studied was the Driver CS0 flowfield [18], a low speed case that has no shocks and is another standard test case for turbulence models. The baseline grid used is the finest grid from previously published work [1]. This case utilizes low-Mach preconditioning, but the eigenvalue limiters are set to zero as mentioned in the method section, since this is a completely subsonic case. This helps the grid convergence of the skin friction by removing unnecessary numerical dissipation.

Not as much detail will be shown for this case, as the grid is converged both in surface data and velocity profiles by the first grid refinement. The solution domain is long and thin (Fig. 15) and as is evident in the hodgepodge of patches distributed over the entire domain (Fig. 15a), there are no isolated regions that are favored for refinement. Instead it has the look of a quilting party familiar with and fond of Leroy Neiman's art. The solution (Fig. 15b shows the distribution of axial velocity) is extremely smooth, and while the separated region has a bit more fine structure in terms of the grid distribution, there are large regions that are single grid patches.




Fig. 17 Residual history as grid refinement level L_r is increased




Fig. 18 Insensitivity of surface pressure (C_p) to mesh refinement on Driver CS0 flow




Fig. 19 Insensitivity of skin friction (c_{f_x}) to mesh refinement on Driver CS0 flow

The residuals have the same pattern as the transonic and supersonic cases (Figs. 16 and 17), with the convergence taking no longer as the grid is refined. Each new level of refinement takes the same number of iterations, on the same order as those required to obtain the baseline solution, with the work to refine the grid accomplished as the overall residual is reduced throughout the entire domain. The strategy used in choosing where to add grid points as the refinement starts, focusing on the regions with the highest deviation from linearity, appears to work extremely well. A minimal number of extra iterations is required to converge the overall flowfield to this level of grid refinement.

Both surface pressure and skin friction (Figs. 18 and 19) are largely converged on the baseline grid, though the pressure field has some changes visible from the baseline to the first refinement. There is essentially no discernible difference between the first and second refinement with any of these turbulence models. The skin friction is nearly

indistinguishable from the baseline for both refinements.

The velocity profiles (Fig. 20) are similarly insensitive, showing little variation after the first refinement, with the largest differences occurring at the boundary layer edge and further downstream. The differences become small enough that comparison with higher grid resolution now requires more and more care since smaller differences become more significant in terms of the solution differences.




Fig. 20 Velocity profiles showing the effect of AMR

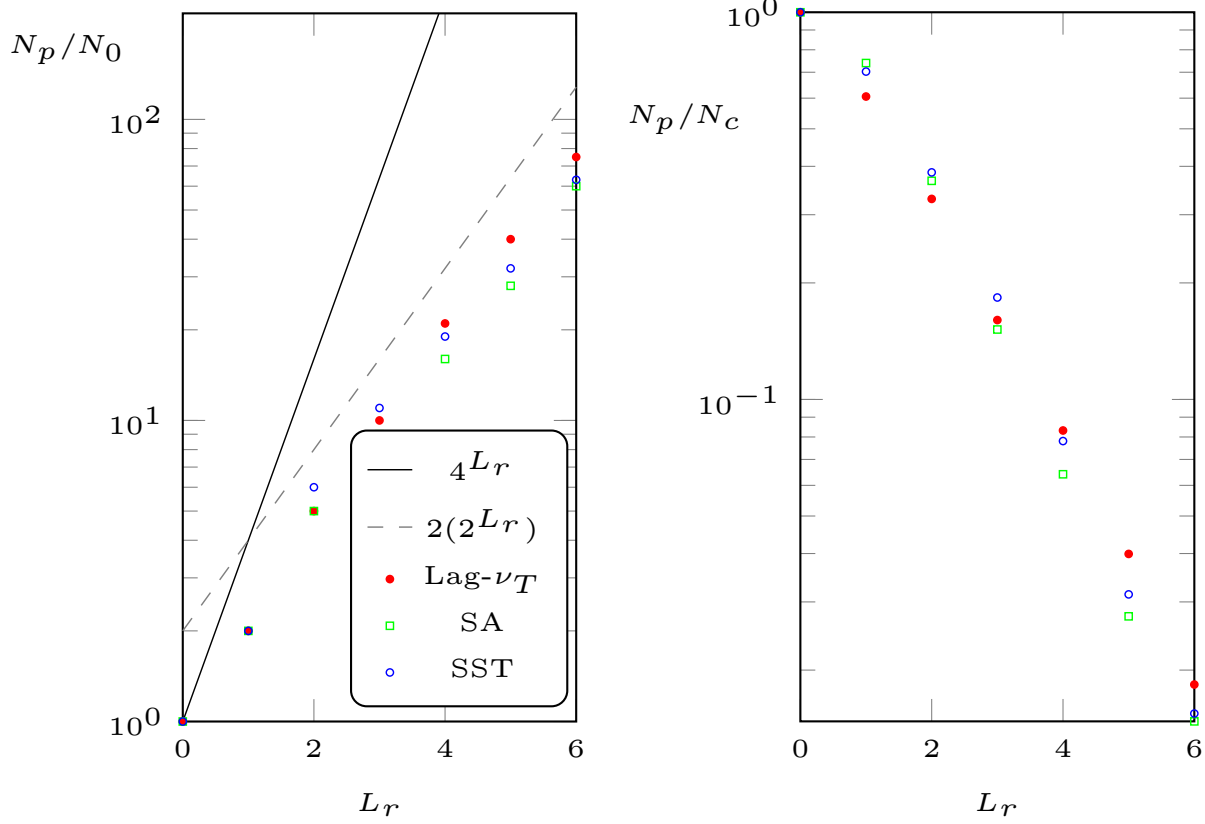
The final question to be addressed is how much grid is required for this smooth solution—does the smoothness make the overall grid growth more pronounced than the case with the flow discontinuities? If anything, AMR is an even bigger win for this case (Fig. 21). On one hand, this might be expected, as the solution should have a bounded derivative field, and there is no discontinuity that will always require more grid resolution. However, the grid addition seemed to be added over the entire domain rather than in a region that has a dimension one less than the solution domain. At least for this case, which does contain a separation that is a good fraction of the domain, the shock dominated flow fields required more grid to drive the solution to the same level of “grid linearity.”

V. Conclusions

AMR in OVERFLOW was executed on subsonic, transonic, and supersonic cases with separated turbulent flows. Although standard grid techniques provided converged solutions for surface pressure and skin friction, AMR brought out more detailed features inside the solution domain where shock waves interact with separation zones. AMR provided the ability to obtain these features accurately at a cost orders of magnitude lower than possible with uniform grid refinement. In addition, AMR, as executed in OVERFLOW, achieved numerical convergence similar to standard techniques and provided better grid efficiency than expected.

The ability to obtain an arbitrarily accurate solution is a vital capability when developing turbulence models. While classic grid refinement promised this capability, the cost (even in two-dimensional flows) limited the ability to assess new turbulence models. It has been demonstrated that the AMR capability in OVERFLOW provides a much more robust ability to assess model behavior. This capability is, of course, important for any use of CFD when assessing solution uncertainty. It also allows for more rapid solution turnaround, since grid refinement can be used to obtain a sufficiently refined grid system without having to create initial grids which are overly refined. For flows with complex shock systems, this capability is likely to become a necessity even for prosaic CFD calculations.

The capability will become more important as our ability to predict fluid flow improves. While the situation looks



(a) AMR grid size/baseline grid size

(b) AMR grid size/classical refinement

Fig. 21 Efficiency of AMR compared to classic grid refinement for the CS0 subsonic flow

very good for steady flows, unsteady flows add a further level of complexity. In addition to the obvious issue of the shock moving, which will require the refinement to follow it, allowing the solution to refine the time step locally would also have great utility. Even without a time dependent grid refinement strategy, simply having a refinement technique that monitors the progress of an unsteady solution, then restarts with a grid refined in space (at every point that would require refinement for a steady solution), but fixed in time, would be most useful.

The present work demonstrates that for steady, two dimensional flowfields, AMR provides a reliable and efficient means of obtaining an arbitrarily accurate approximation to the analytical solution. Further, it implies that for three-dimensional flowfields, the same will be possible, albeit at greater cost.

Acknowledgements

This research was sponsored by NASA's Transformational Tools and Technologies (TTT) Project of the Transformative Aeronautics Concepts Program under the Aeronautics Research Mission Directorate.

References

- [1] Olsen, M., and Coakley, T., “The lag model, a turbulence model for nonequilibrium flows,” AIAA Paper 2001-2564, Jun 2001. URL <https://doi.org/10.2514/6.2001-2564>.
- [2] Olsen, M., and Prabhu, D., “Application of OVERFLOW to hypersonic perfect gas flowfields,” AIAA Paper 2001-2664, Jun 2001. URL <http://dx.doi.org/10.2514/6.2001-2664>.
- [3] Olsen, M. E., and Lillard, R. P., “Revised Reynolds-Stress and Triple Product Models,” AIAA Paper 2017-3954, Jun 2017. URL <https://doi.org/10.2514/6.2017-3954>.
- [4] Buning, P. G., and Pulliam, T. H., “Near-Body Grid Adaption for Overset Grids,” AIAA Paper 2016-3326, Jun 2016. URL <https://doi.org/10.2514/6.2016-3326>.
- [5] Buning, P., and Pulliam, T., “Cartesian Off-Body Grid Adaption for Viscous Time-Accurate Flow Simulations,” AIAA Paper 2011-3693, Jun 2011. URL <https://doi.org/10.2514/6.2011-3693>.
- [6] Rumsey, C., “Langley Research Center Turbulence Modeling Resource,” , May 2018. [Online; accessed 9-May-2018], <http://turbmodels.larc.nasa.gov/>.
- [7] Childs, Marissa L., Pulliam, Thomas H., and Jespersen, Dennis C., “OVERFLOW Turbulence Model Resource Verification Results,” NAS Technical Report 2014-03, NASA, June 2014. NAS Technical Report NAS-2014-03.
- [8] Bachalo, W.D., and Johnson, D., “Transonic Turbulent Boundary Layer Separation Generated on an Axisymmetric Flow Model,” *AIAA Journal*, Vol. 24, 1986, pp. 437–443.
- [9] Olsen, M. E., “Prediction of Separation with a Third-Order-Moment Model,” AIAA Paper 2015-1968, <http://dx.doi.org/10.2514/6.2015-1968>, January 2015.
- [10] Olsen, M. E., “Reynolds-stress and triple-product models applied to flows with rotation and curvature,” AIAA Paper 2016-3942, Jun 2016. doi:10.2514/6.2016-3942, URL <https://doi.org/10.2514/6.2016-3942>.
- [11] Spalart, P.R., and Allmaras, S., “A One Equation Turbulence Model for Aerodynamic Flows,” *La Recherche Aerospatiale*, Vol. 1, 1994, pp. 5–21.
- [12] Menter, F.R., “Two Equation Eddy Viscosity Model for Engineering Applications,” *AIAA Journal*, Vol. 32, 1994, pp. 1299–1310.
- [13] Olsen, M., Coakley, T., and Lillard, R., “The Lag Model Applied to High Speed Flows,” AIAA Paper 2005-0101, Jan 2005. URL <https://doi.org/10.2514/6.2005-101>.
- [14] Su, X., “Accurate and robust adaptive mesh refinement for aerodynamic simulation with multi-block structured curvilinear mesh,” *International Journal for Numerical Methods in Fluids*, Vol. 77, No. 12, 2015, pp. 747–766. doi:10.1002/flid.4004, URL <https://onlinelibrary.wiley.com/doi/abs/10.1002/flid.4004>.
- [15] Brown, J.D., Brown, J.L., and Kussoy, M.I., “A documentation of two-and three-dimensional shock-separated turbulent boundary layers,” Tech. rep., NASA, Jul 1988. NASA TM 101008.
- [16] DeBonis, J.R., “Evaluation of Industry Standard Turbulence Models on an Axisymmetric Supersonic Compression Corner,” Tech. rep., NASA, 2015. NASA TM 2015-2853.
- [17] Dudek, J., and Carlson, J.-R., “Evaluation of Full Reynolds Stress Turbulence Models in Fun3D,” AIAA Paper 2017-0541, Jan 2017. URL <https://doi.org/10.2514/6.2017-0541>.
- [18] DRIVER, D., “Reynolds shear stress measurements in a separated boundary layer flow,” AIAA Paper 1991-1787, Jun 1991. URL <https://doi.org/10.2514/6.1991-1787>.