



**AIAA 2002-0838**

**TRANSONIC DRAG PREDICTION  
USING AN UNSTRUCTURED  
MULTIGRID SOLVER**

**D. J. Mavriplis**

ICASE

MS 132C, NASA Langley Research Center

Hampton, VA

and

**David W. Levy**

Cessna Aircraft Co.

Wichita, KS

**40th AIAA Aerospace Sciences Meeting**  
**January 14-17 2002, Reno NV**



# TRANSONIC DRAG PREDICTION USING AN UNSTRUCTURED MULTIGRID SOLVER

D. J. Mavriplis

ICASE

MS 132C, NASA Langley Research Center

Hampton, VA

and

David W. Levy

Cessna Aircraft Co.

Wichita, KS

## **Abstract**

This paper summarizes the results obtained with the NSU3D unstructured multigrid solver for the AIAA Drag Prediction Workshop held in Anaheim, CA, June 2001. The test case for the workshop consists of a wing-body configuration at transonic flow conditions. Flow analyses for a complete test matrix of lift coefficient values and Mach numbers at a constant Reynolds number are performed, thus producing a set of drag polars and drag rise curves which are compared with experimental data. Results were obtained independently by both authors using an identical baseline grid, and different refined grids. Most cases were run in parallel on commodity cluster-type machines while the largest cases were run on an SGI Origin machine using 128 processors. The objective of this paper is to study the accuracy of the subject unstructured grid solver for predicting drag in the transonic cruise regime, to assess the efficiency of the method in terms of convergence, cpu time and memory, and to determine the effects of grid resolution on this predictive ability and its computational efficiency. A good predictive ability is demonstrated over a wide range of conditions, although accuracy was found to degrade for cases at higher Mach numbers and lift values where increasing amounts of flow separation occur. The ability to rapidly compute large numbers of cases at varying flow conditions using an unstructured solver on inexpensive clusters of commodity computers is also demonstrated.

## **Introduction**

Computational fluid dynamics has progressed to the point where Reynolds-averaged Navier-Stokes solvers have become standard simulation tools for predicting aircraft aerodynamics. These solvers are routinely

used to predict aircraft force coefficients such as lift, drag and moments, as well as the changes in these values with design changes. In order to be useful to an aircraft designer, it is generally acknowledged that the computational method should be capable of predicting drag to within several counts. While Reynolds-averaged Navier-Stokes solvers have made great strides in accuracy and affordability over the last decade, the stringent accuracy requirements of the drag prediction task have proved difficult to achieve. This difficulty is compounded by the multitude of Navier-Stokes solver formulations available, as well as by the effects on accuracy of turbulence modeling and grid resolution. Therefore, a particular Navier-Stokes solver must undergo extensive validation including the determination of adequate grid resolution distribution, prior to being trusted as a useful predictive tool. With these issues in mind, the AIAA Applied Aerodynamics technical committee organized a Drag Prediction Workshop, held in Anaheim CA, June 2001,<sup>1</sup> in order to assess the predictive capabilities of a number of state-of-the-art computational fluid dynamics methods. The chosen configuration, denoted as DLR-F4<sup>2</sup> and depicted in Figure 1, consists of a wing-body geometry, which is representative of a modern supercritical swept wing transport aircraft. Participants included Reynolds-averaged Navier-Stokes formulations based on block-structured grids, overset grids, and unstructured grids, thus affording an opportunity to compare these methods on an equal basis in terms of accuracy and efficiency. A standard mesh was supplied for each type of methodology, with participants encouraged to produce results on additionally refined meshes, in order to assess the effects of grid resolution. A Mach number versus lift coefficient ( $C_L$ ) matrix of test cases was defined, which included mandatory and optional cases. The calculations were initially run by the participants without knowledge of the experimental data, and a

Copyright © 2002 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved.

[illegible]

This paper describes the results obtained for this workshop with the unstructured mesh Navier-Stokes solver NSU3D.<sup>4-6</sup> This solver has been well validated and is currently in use in both a research setting and an industrial production environment. Results were obtained independently by both authors on the baseline workshop grid, and on two refined grids generated independently by both authors. All required and optional cases were run using the baseline grid and one refined grid, while the most highly refined grid was only run on the mandatory cases. The runs were performed on three different types of parallel machines at two different locations.

The NSU3D code solves the Reynolds averaged Navier-Stokes equations on unstructured meshes of mixed element types which may include tetrahedra, pyramids, prisms, and hexahedra. All elements of the grid are handled by a single unifying edge-based data-structure in the flow solver.<sup>1</sup>

Flow variables are stored at the vertices of the mesh, and the governing equations are discretized using a central difference finite-volume technique with added artificial dissipation. The matrix formulation of the artificial dissipation is employed, which corresponds

The basic time-stepping scheme is a three-stage explicit multistage scheme with stage coefficients optimized for high frequency damping properties,<sup>7</sup> and a CFL number of 1.8. Convergence is accelerated by a local block-Jacobi preconditioner in regions of isotropic grid cells, which involves inverting a  $5 \times 5$  matrix for each vertex at each stage.<sup>8-10</sup> In boundary layer regions, where the grid is highly stretched, a line smoother is employed, which involves inverting a block tridiagonal along lines constructed in the unstructured mesh by grouping together edges normal to the grid stretching direction. The line smoothing technique has been shown to relieve the numerical stiffness associated with high grid anisotropy.<sup>11</sup>

The single equation turbulence model of Spalart and Allmaras<sup>13</sup> is utilized to account for turbulence effects. This equation is discretized and solved in a manner completely analogous to the flow equations, with the exception that the convective terms are only discretized to first-order accuracy.

AMERICAN INSTITUTE OF AERONAUTICS AND ASTRONAUTICS PAPER 2002-0838

clusters and an SGI Origin 2000, using the MPI programming model exclusively.

### Grid Generation

The baseline grid supplied for the workshop was generated using the VGRIDns package.<sup>17</sup> This approach produces fully tetrahedral meshes, although it is capable of generating highly stretched semi-structured tetrahedral elements near the wall in the boundary-layer region, and employs moderate spanwise stretching in order to reduce the total number of points. A semi-span geometry was modeled, with the far-field boundary located 50 chords away from the origin, resulting in a total of 1.65 million grid points, 9.7 million tetrahedra, and 36,000 wing-body surface points. The chordwise grid spacing at the leading edge was prescribed as 0.250 mm and 0.500 mm at the trailing edge, using a dimensional mean chord of 142.1 mm. The trailing edge is blunt, with a base thickness of 0.5 % chord, and the baseline mesh contained 5 grid points across the trailing edge. The normal spacing at the wall is 0.001 mm, which is designed to produce a grid spacing corresponding to  $y^+ = 1$  for a Reynolds number of 3 million. A stretching rate of 1.2 was prescribed for the growth of cells in the normal direction near the wall, in order to obtain a minimum of 20 points in the boundary layer.

Because the NSU3D solver is optimized to run on mixed element meshes, the fully tetrahedral baseline mesh is subsequently converted to a mixed element mesh by merging the semi-structured tetrahedral layers in the boundary layer region into prismatic elements. This is done in a pre-processing phase where triplets of tetrahedral layers are identified and merged into a single prismatic element, using information identifying these elements as belonging to the stretched viscous layer region as opposed to the isotropic inviscid tetrahedral region. The merging operation results in a total of 2 million created prismatic elements, while the number of tetrahedral cells is reduced to 3.6 million, and a total of 10090 pyramidal elements are created to merge prismatic elements to tetrahedral elements in

regions where quadrilateral faces from prismatic elements are adjacent to tetrahedral elements.

A higher resolution mesh was generated by the second author using VGRIDns with smaller spacings in the vicinity of the wing root, tip, and trailing edge, resulting in a total of 3 million grid points, and 73,000 wing-body surface points. One of the features of this refined grid is the use of a total of 17 points across the wing trailing edge versus 5 for the baseline grid. After the merging operation, this grid contained a total of 3.7 million prisms and 6.6 million tetrahedra.

An additional fine mesh was obtained by the first author through global refinement of the baseline workshop mesh. This strategy operates directly on the mixed prismatic-tetrahedral mesh, and consists of subdividing each element into 8 smaller self-similar elements, thus producing an 8:1 refinement of the original mesh.<sup>18</sup> The final mesh obtained in this manner contained a total of 13.1 million points with 16 million prismatic elements and 28.8 million tetrahedral elements, and 9 points across the blunt trailing edge of the wing. This approach can rapidly generate very large meshes which would otherwise be very time consuming to construct using the original mesh generation software. One drawback of the current approach is that newly generated surface points do not lie exactly on the original surface description of the model geometry, but rather along a linear interpolation between previously existing surface coarse grid points. For a single level of refinement, this drawback is not expected to have a noticeable effect on the results. An interface for re-projecting new surface points onto the original surface geometry is currently under consideration.

The baseline grid was found to be sufficient to resolve all major flow features. The computed surface pressure coefficient on the baseline grid for a Mach number of 0.75, Reynolds number of 3 million, and  $C_L = 0.6$  is shown in Figure 2, illustrating good resolution of the upper surface shock. A small region of separation is also resolved in the wing root area, as shown by the surface streamlines for the same flow conditions, in Figure 3.


**Table 1 Grids and Corresponding Run Times**

Grid	No. Points	No. Tets	No. Prisms	Memory	Run Time	Hardware
Grid 1	$1.65 \times 10^6$	$2 \times 10^6$	$3.6 \times 10^6$	2.8 Gbytes	2.6 hours	16 Pentium IV 1.7GHz
Grid 1	$1.65 \times 10^6$	$2 \times 10^6$	$3.6 \times 10^6$	2.4 Gbytes	8 hours	4 DEC Alpha 21264 (667MHz)
Grid 1	$1.65 \times 10^6$	$2 \times 10^6$	$3.6 \times 10^6$	3.0 Gbytes	45 min.	64 SGI Origin 2000 (400MHz)
Grid 2	$3.0 \times 10^6$	$3.7 \times 10^6$	$6.6 \times 10^6$	4.2 Gbytes	8 hours	8 DEC Alpha 21264 (667MHz)
Grid 3	$13 \times 10^6$	$16 \times 10^6$	$28.8 \times 10^6$	27 Gbytes	4 hours	128 SGI O2000 (400MHz)


Figure 4 depicts the computed  $y^+$  values at the break section for the same flow conditions, indicating values well below unity over the entire lower surface and a majority of the upper surface. The convergence history for this case is shown in Figure 5. The flow is initialized as a uniform flow at freestream conditions, and ten single grid cycles (no multigrid) are employed to smooth the initialization prior to the initiation of the multigrid iteration procedure. A total residual reduction of approximately 5 orders of magnitude is achieved over 500 multigrid cycles. Convergence in the lift coefficient is obtained in as little as 200 multigrid cycles for this case, although all cases are run a minimum of 500 multigrid cycles as a conservative convergence criterion. This convergence behavior is representative of the majority of cases run, with some of the high Mach number and high  $C_L$  cases involving larger regions of separation requiring up to 800 to 1000 multigrid cycles. A flow solution on the baseline grid requires 2.8 Gbytes of memory and a total of 2.6 hours of wall clock time (for 500 multigrid cycles) on a cluster of commodity components using 16 Pentium IV 1.7 GHz processors communicating through 100 Mbit Ethernet. This case was also run on 4 DEC Alpha processors, requiring 2.4 Gbytes of memory and 8 hours of wall clock time. This case was also benchmarked on 64 processors (400MHz) of an SGI Origin 2000, requiring 3 Gbytes of memory and 45 minutes of wall clock time. The memory requirements are independent of the specific hardware and are only a function of the number of partitions used in the calculations.

The cases using the 3 million point grid were run on a cluster of 8 DEC Alpha processors communicating through 100 Mbit Ethernet and required approximately 8 hours of wall clock time and 1.2 Gbytes of memory.


The 13 million point grid cases were run on an SGI Origin 2000, using 128 processors and required 4 hours of wall clock time and 27 Gbytes of memory. A description of the three grids employed and the associated computational requirements on various hardware platforms is given in Table 1.




**Fig. 2 Baseline Grid and Computed Pressure Contours at Mach=0.75,  $C_L = 0.6$ , Re = 3 million**



**Fig. 3** Computed Surface Oil Flow Pattern in Wing Root Area on Baseline Grid for Mach=0.75,  $C_L = 0.6$ , Re = 3 million



**Fig. 4** Computed  $y^+$  on wing surface at span break section on baseline grid for Mach=0.75,  $C_L = 0.6$ , Re = 3 million



**Fig. 5** Density Residual and Lift Coefficient Convergence History as a Function of Multigrid Cycles on Baseline Grid for Mach=0.75,  $C_L = 0.6$ , Re = 3 million

## Results

The workshop test cases comprised two required cases and two optional cases. These cases are described in Table 2. For all cases the Reynolds number is 3 million. The first test case is a single point at Mach = 0.75 and  $C_L = 0.5$ . The second test case involves the computation of the drag polar at Mach=0.75 using incidences from -3.0 to +2.0 degrees in increments of 1 degree. The optional Cases 3 and 4 involve a matrix of Mach and  $C_L$  values in order to compute drag rise curves. Since an automated approach for computing fixed  $C_L$  cases has not been implemented, a complete drag polar for each Mach number was computed for Cases 3 and 4. For the baseline grid, the incidence for the prescribed lift value was then interpolated from the drag polar using a cubic spline fit, and the flow was recomputed at this prescribed incidence. The final force coefficients were then interpolated from the values computed in this case to the prescribed lift values, which are very close to the last computed case. For the 3 million point grid, the force coefficient values at the prescribed lift conditions were interpolated directly from the 6 integer degree cases within each drag polar.

**Table 2 Definition of Required and Optional Cases for Drag Prediction Workshop**

Case	Description
Case 1 (Required) Single Point	Mach = 0.75, $C_L = 0.500$
Case 2 (Required) Drag Polar	Mach = 0.75 $\alpha = -3^\circ, -2^\circ, -1^\circ, 0^\circ, 1^\circ, 2^\circ$
Case 3 (Optional) Constant $C_L$ Mach Sweep	Mach = .50..60..70..75..76..77..78..80 $C_L = 0.500$
Case 4 (Optional) Drag Rise Curves	Mach = .50..60..70..75..76..77..78..80 $C_L = 0.400, 0.500, 0.600,$

All cases were computed using the baseline grid (1.6 million points), and the medium grid (3 million points). Only the required cases were computed using the finest grid (13 million points) due to time constraints. Table 3 depicts the results obtained for Case 1 with the three different grids. The drag is seen to be computed accurately by all three grids, although there is a 10.6 count variation between the 3 grids. However, the incidence at which the prescribed  $C_L = 0.5$  is achieved is up to 0.6 degrees lower than that observed experimentally. This effect is more evident in the  $C_L$  versus incidence plot of Figure 6, where the computed lift values are consistently higher than the experimental values. Since this discrepancy increases with the higher resolution grids, it cannot be attributed to a lack of grid resolution. The slope of the computed lift curve is about 5% higher than the experimentally determined slope, and is largely unaffected by grid resolution.

Figure 7 provides a comparison of computed surface pressure coefficients with experimental values at the experimentally prescribed  $C_L$  of 0.6 (where the effects are more dramatic than at  $C_L = 0.5$ ) as well as at the experimentally prescribed incidence of 0.93 degrees, at the 40.9 % span location. When the experimental incidence value is matched, the computed shock location is aft of the experimental values, and the computed lift is higher than the experimental value, while at the prescribed lift condition, the shock is further forward and the suction peak is lower than the experimental values.

This bias in lift versus incidence was observed for a majority of the numerical solutions submitted to the workshop,<sup>3</sup> and thus might be attributed to a model

geometry effect or a wind tunnel correction effect, although an exact cause has not been determined. When plotted as a drag polar,  $C_L$  versus  $C_D$  as shown in Figure 8, the results compare favorably with experimental data. Although the drag polar was computed independently by both authors using the baseline grid, the results of both sets of computations were identical (as expected) and thus only one set of computations is shown for the baseline grid. The computational results on this grid compare very well with experiment in the mid-range (near  $C_L = 0.5$ ), while a slight over-prediction of drag is observed for low lift values, which decreases as the grid is refined.

This behavior suggests an under-prediction of induced drag, possibly due to inadequate grid resolution in the tip region or elsewhere. The absolute drag levels have been found to be sensitive to the degree of grid refinement at the blunt trailing edge of the wing. The drag level is reduced by 4 counts when going from the 1.6 million point grid, which has 5 points on the trailing edge, to the 3 million point grid, which has 17 points on the trailing edge. Internal studies by the second author using structured grids have shown that up to 33 points on the blunt trailing edge are required before the drag does not decrease any further. In the current grid generation environment, and without the aid of adaptive meshing techniques, the generation of highly refined trailing edge unstructured meshes has been found to be problematic, thus limiting our study in this area.

Figure 9 provides an estimate of the induced drag factor, determined experimentally and computationally on the three meshes.

**Table 3 Results for Case 1; Experimental Values 1:ONERA, 2:NLR, 3:DRA; Grid1\*: Performed by first author, Grid1+: Performed by second author. Experimental data and 3 M point grid results are interpolated to specified  $C_L$  condition along drag polar.**

Case	$C_L$	$\alpha$	$C_D$	$C_M$
Experiment <sup>1</sup>	0.5000	+ .192°	0.02896	-.1301
Experiment <sup>2</sup>	0.5000	+ .153°	0.02889	-.1260
Experiment <sup>3</sup>	0.5000	+ .179°	0.02793	-.1371
Grid1(1.6Mpts)*	0.5004	-.241°	0.02921	-.1549
Grid1(1.6Mpts)+	0.4995	-.248°	0.02899	-.1548
Grid2(3.0Mpts)	0.5000	-.417°	0.02857	-.1643
Grid3(13Mpts)	0.5003	-.367°	0.02815	-.1657






Fig. 6 Comparison of Computed Lift as a function of Incidence for Three Different Grids versus Experimental Results




Fig. 8 Comparison of Computed versus Experimental Drag Polar for Mach=0.75 using Three Different Grids




Fig. 7 Comparison of Computed Surface Pressure Coefficients at Prescribed Lift and Prescribed Incidence versus Experimental Values for Baseline Grid at 40.9 % span location




Fig. 9 Comparison of Computed versus Experimental Induced Drag Factor for Mach=0.75 using Three Different Grids




Fig. 10 Comparison of Computed versus Experimental Idealized Profile Drag at Mach=0.75 using Three Different Grids




Fig. 11 Comparison of Computed versus Experimental Pitching Moment for Mach=0.75 using Three Different Grids

For  $C_L^2$  up to about 0.36, when the flow is mostly attached, induced drag is underpredicted by approximately 10 %, as determined by comparing the slopes of the computational and experimental curves (using a linear curve fit) in this region. Grid refinement appears to have little effect on the induced drag in this region. At the higher lift values, the 3 million point grid yields higher  $C_L$  and lower  $C_D$  values, which is attributed to a slight delay in the amount of predicted

flow separation. Results for the 13 million point grid are not shown, since a fully converged solution could not be obtained at the highest incidence. On the other hand, it should be noted that the wind tunnel experiments used a boundary layer trip at 15% and 25% chord on the upper and lower surfaces, while all calculations were performed in a fully turbulent mode. Examination of the generated eddy viscosity levels in the calculations reveals appreciable levels beginning between 5% to 7% chord. The exact influence of transition location on overall computed force coefficients has not been quantified and requires further study.

Figure 10 shows the idealized profile drag<sup>19</sup> which is defined by the formula:

$$C_{DP} = C_D - C_L^2 / (\pi AR) \quad (1)$$

where  $AR$  is the aspect ratio. Plotting  $C_{DP}$  generally results in a more compact representation of the data, allowing more expanded scales. It also highlights the characteristics at higher  $C_L$ , where the drag polar becomes non-parabolic due to wave drag and separation. In the non-parabolic region, the error in drag is relatively large at a constant  $C_L$ .

The pitching moment is plotted as a function of  $C_L$  in Figure 11 for all three grids versus experimental values. The pitching moment is substantially underpredicted with larger discrepancies observed for the refined grids. This is likely a result of the over-prediction of lift as a function of incidence, as mentioned earlier and illustrated in Figure 6. Because the computed shock location and suction peaks do not line up with experimental values, the predicted pitching moments can not be expected to be in good agreement with experimental values.

Figure 12 depicts the drag rise curves obtained for Cases 3 and 4 on the baseline grid and the first refined grid (3 million points). Drag values are obtained at four different constant  $C_L$  values for a range of Mach numbers. Drag values are predicted reasonably well except at the highest lift and Mach number conditions. There appears to be no improvement in this area with increased grid resolution, which suggests issues such as transition and turbulence modeling may account for these discrepancies. However, since the two grids have comparable resolution in various areas of the domain, grid resolution issues still cannot be ruled out at this stage.

The results obtained for Cases 3 and 4 can also be plotted at constant Mach number, as shown in the drag polar plots of Figure 13. The plots show similar trends, with the drag being slightly overpredicted at low lift values on the coarser grid and with the refined grid achieving better agreement in these regions. For the higher Mach numbers, the drag is substantially underpredicted at the higher lift values. These discrepancies at the higher Mach numbers and lift conditions

point to an under-prediction of the extent of the separated regions of flow in the numerical simulations. The comparison of idealized profile drag in Figure 10 also suggests that the drag due to flow separation is not predicted accurately at the higher lift conditions. However, the character of the curves also suggest that the error may be due as well to the  $C_L$  offset (shown in Figure 6). Additional information concerning the regions of flow separation found in the wind tunnel would be needed to more accurately quantify the nature of the error.

The above results indicate that the current unstructured mesh Navier-Stokes solver achieves a reasonably good predictive ability for the force coefficients on the baseline grid over the majority of the flow conditions considered. The overall agreement, particularly at the low lift values, is improved with added grid resolution, while the more extreme flow conditions which incur larger amounts of separation are more difficult to predict accurately. On the other hand, the observed bias between computation and experiment in the lift versus incidence values has an adverse affect on the prediction of pitching moment. While the source of this bias is not fully understood, it was observed for a majority of independent numerical simulations undertaken as part of the subject workshop<sup>3</sup> and can likely be attributed to geometrical differences or wind tunnel corrections.

The results presented in this paper involve a large number of individual steady-state cases. For example, on the baseline grid, a total of 72 individual cases were computed, as shown in Figure 14, to enable the construction of Figures 8, 12, and 13. The majority of these cases were run from freestream initial conditions for 500 multigrid cycles, while several cases particularly in the high Mach number and high lift regions were run 800 to 1000 cycles to obtain fully converged results. The baseline cases (500 multigrid cycles) required approximately 2.6 hours of wall clock time on a cluster of 16 commodity PC processors. This enabled the entire set of 72 cases to be completed within a period of one week. This exercise illustrates the possibility of performing a large number of parameter runs, as is typically required in a design exercise, with a state-of-the-art unstructured solver on relatively inexpensive parallel hardware.




Fig. 12 Comparison of Computed versus Experimental Drag Rise Curves Three Different  $C_L$  values on Two Different Grids





Fig. 13 Comparison of Computed versus Experimental Drag Polars for Mach=0.6 and Mach=0.8 on Two Different Grids



**Fig. 14 Depiction of All 72 Individual Cases run on Baseline Grid Plotted in Drag Polar Format**

## Conclusions

A state-of-the-art unstructured multigrid Navier-Stokes solver has demonstrated good drag predictive ability for a wing-body configuration in the transonic regime. Acceptable accuracy has been achieved on relatively coarse meshes of the order of several million grid points, while improved accuracy has been demonstrated with increased grid resolution. Grid resolution remains an important issue, and considerable expertise is required in specifying the distribution of grid resolution in order to achieve a good predictive ability without resorting to extremely large mesh sizes. These issues can be resolved to some degree by the use of automatic grid adaptation procedures, which are planned for future work. The predictive ability of the numerical scheme was found to degrade for flow conditions involving larger amounts of flow separation. Slight convergence degradation was observed on two of the grids for the cases involving increased flow separation, while a fully converged result could not be obtained on the finest grid (13 million points) for the highest lift case at a Mach number of 0.75. The current results utilized the Spalart-Allmaras turbulence model exclusively, and the effect of other turbulence models in this regime deserves additional consideration. The rapid convergence of the multigrid scheme coupled with the parallel implementation on commodity networked computer clusters has been shown to produce a useful design tool with quick turnaround time.

## References

<sup>1</sup>AIAA Drag Prediction Workshop, Anaheim, CA, <http://www.aiaa.org/tc/apa/dragpredworkshop/dpw.html>, June 2001.

<sup>2</sup>G. Redeker, DLR-F4 wing body configuration, Technical report, August 1994, AGARD Report AR-303, Vol II.

<sup>3</sup>M. Hemsch, Statistical analysis of CFD solutions from the Drag Prediction Workshop, AIAA Paper 2002-0812, January 2002.

<sup>4</sup>D. J. Mavriplis and V. Venkatakrishnan, A unified multigrid solver for the Navier-Stokes equations on mixed element meshes, *International Journal for Computational Fluid Dynamics*, 8:247-263, 1997.

<sup>5</sup>D. J. Mavriplis and S. Pirzadeh, Large scale parallel unstructured mesh computations for 3D high-lift analysis, AIAA paper 99-0537, presented at the 37th AIAA Aerospace Sciences Meeting, Reno NV, January 1999.

<sup>6</sup>D. J. Mavriplis, Parallel performance investigations of an unstructured mesh Navier-Stokes solver, ICASE Report No. 2000-13, NASA CR 2000-210088, March 2000.

<sup>7</sup>B. van Leer, C. H. Tai, and K. G. Powell, Design of optimally smoothing multi stage schemes for the Euler equations, AIAA Paper 89-1933, June 1989.

<sup>8</sup>K. Rienslagh and E. Dick, A multigrid method for steady Euler equations on unstructured adaptive grids, In *6th Copper Mountain Conf. on Multigrid Methods*, NASA conference publication 3224, pages 527-542, 1993.

<sup>9</sup>E. Morano and A. Dervieux, Looking for O(N) Navier-Stokes solutions on non-structured meshes, In *6th Copper Mountain Conf. on Multigrid Methods*, pages 449-464, 1993, NASA Conference Publication 3224.

<sup>10</sup>N. Pierce and M. Giles, Preconditioning on stretched meshes, AIAA paper 96-0889, January 1996.

<sup>11</sup>D. J. Mavriplis, Directional agglomeration multigrid techniques for high-Reynolds number viscous flows, AIAA paper 98-0612, January 1998.

<sup>12</sup>M. Lallemand, H. Steve, and A. Dervieux, Unstructured multigriding by volume agglomeration: Current status, *Computers and Fluids*, 21(3):397-433, 1992.

<sup>13</sup>P. R. Spalart and S. R. Allmaras, A one-equation turbulence model for aerodynamic flows, *La Recherche Aéronautique*, 1:5-21, 1994.

<sup>14</sup>G. Karypis and V. Kumar, A fast and high quality multi-level scheme for partitioning irregular graphs, Technical Report Technical Report 95-035, University of Minnesota, 1995. A short version appears in Intl. Conf. on Parallel Processing 1995.

<sup>15</sup>W. Gropp, E. Lusk, and A. Skjellum, *Using MPI: Portable Parallel Programming with the Message Passing Interface*, MIT Press, Cambridge, MA, 1994.

<sup>16</sup>OpenMP: Simple, portable, scalable SMP programming, <http://www.openmp.org>, 1999.

<sup>17</sup>S. Pirzadeh, Three-dimensional unstructured viscous grids by the advancing-layers method, *AIAA Journal*, 34(1):43-49, 1996.

<sup>18</sup>D. J. Mavriplis, Adaptive meshing techniques for viscous flow calculations on mixed-element unstructured meshes, AIAA paper 97-0857, January 1997.

<sup>19</sup>E. N. Tinoco, An assessment of CFD prediction of drag and other longitudinal characteristics, AIAA Paper 2001-1002, January 2001.



