NASA Contractor Report 181848
ICASE Report No. 89-35

ICASE

ZONAL MULTIGRID SOLUTION OF COMPRESSIBLE FLOW PROBLEMS ON UNSTRUCTURED AND ADAPTIVE MESHES

Dimitri J. Mavriplis

Contract No. NAS1-18107, NAS1-18605
May 1989

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, Virginia 23665-5225

Operated by the Universities Space Research Association

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23665-5225

ZONAL MULTIGRID SOLUTION OF COMPRESSIBLE FLOW PROBLEMS ON UNSTRUCTURED AND ADAPTIVE MESHES

D. J. Mavriplis

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA

ABSTRACT

This work is concerned with the simultaneous use of adaptive meshing techniques with a multigrid strategy for solving the two-dimensional Euler equations in the context of unstructured meshes. To obtain optimal efficiency, methods capable of computing locally improved solutions without recourse to global recalculations are pursued. A method for locally refining an existing unstructured mesh, without regenerating a new global mesh is employed, and the domain is automatically partitioned into refined and unrefined regions. Two multigrid strategies are developed. In the first, time-stepping is performed only in the locally refined regions of the domain on the fine mesh levels, and throughout the entire domain on the coarsest mesh levels. In the second method, time-stepping is performed on a global fine mesh covering the entire domain, and convergence acceleration is achieved through the use of zonal coarse grid accelerator meshes, which lie under the adaptively refined regions of the global fine mesh. Both schemes are shown to produce similar convergence rates to each other, and also with respect to a previously developed global multigrid algorithm, which performs time-stepping throughout the entire domain, on each mesh level. However, the present schemes exhibit higher computational efficiency due to the smaller number of operations on each level.

This research was supported under the National Aeronautics and Space Administration under NASA Contract Nos. NAS1-18107 and NAS1-18605 while the author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23665.

1. INTRODUCTION

The use of unstructured triangular or tetrahedral meshes in two or three dimensions respectively for solving compressible flow problems about complicated geometries has become more widespread in recent years [1,2,3]. The advantages of unstructured meshes lie in their ability to deal with completely arbitrary geometries, while providing a natural setting for the use of adaptive meshing techniques. On the other hand, the implementation of well known efficient algorithms on unstructured meshes is more difficult or may not even be possible in certain cases. On structured meshes, multigrid methods have proven to be among the most efficient algorithms for computing steady-state solutions of compressible flow problems. A multigrid strategy achieves accelerated convergence of a set of fine grid equations by repeatedly computing corrections to the fine grid solution, obtained by time-stepping on coarser grids. In a previous paper [4], it was shown how a multigrid algorithm can be implemented on unstructured meshes. The underlying idea consists of treating the various coarse and fine meshes of the multigrid sequence to be completely independent from one another, and determining the patterns for transferring variables and residuals back and forth between the various meshes in a preprocessing operation, where an efficient tree-search algorithm is employed. Since this multigrid algorithm assumes no dependency between the various meshes, it can easily be combined with an adaptive meshing technique [5], where the finest meshes of the multigrid sequence are generated by local adaptive refinement, as determined by the solution on the previous coarser mesh. However, this simple strategy for combining adaptive meshing with a multigrid algorithm does not represent the optimum in solution efficiency. For cases where the adaptive meshing results in refinement of the mesh in several small localized areas of the domain, the new fine mesh and the previous coarse mesh exhibit the same resolution throughout most of the domain, and thus time-stepping on these areas on both mesh levels is unnecessary.

A major requirement in the development of an optimal adaptive multigrid method is that localized increases in resolution need not require the reconsideration of the entire global problem. This requirement has implications for both the unstructured grid generation procedure, as well as the flow solver. Thus, local adaptive mesh refinement should be achieved by a local restructuring of an existing mesh, rather than through a complete regeneration of a new global mesh. Similarly, the improved solution due to this local increase in mesh resolution should be computed by modifying an existing global solution through mostly local operations.

In the case of structured meshes, efficient strategies for combining adaptive mesh refinement and multigrid techniques, which permit local increases in resolution without global recalculations have been developed. For example, Berger and Jameson [6] proposed an adaptive multigrid method where time-stepping on the adaptively generated fine grids is only performed in the refined regions. Usab [7] and later Dannenhoffer [8] have developed a multigrid method for compressible flow problems based on Ni's scheme [9], which also operates on a sequence of adaptively generated quadrilateral meshes. A new finer mesh level, which is comprised of various arbitrarily shaped zonal mesh patches, is created by adaptively refining specific zones of the previous mesh level. The solution is cycled back and forth between these zonal fine meshes and the global coarse meshes until convergence is achieved. The final converged solution over the entire domain is then obtained by considering the composite mesh formed by the union of the regions of the global coarse grid which do not lie under any refined zones, and the local fine meshes. These methods are also similar to the so-called "Fast Adaptive Composite" (FAC) method described in [10,11]. In this more theoretical approach, a mul-

tilevel method is used to solve a problem with widely disparate length scales. An initial coarse regular cartesian mesh is employed to resolve the large scales. In regions where smaller scales need to be resolved, additional levels are created by constructing zonal refined cartesian meshes in these regions. The underlying assumption is that each mesh level operates on a well defined set of length scales.

The objective of this work is to investigate methods for efficiently combining adaptive meshing and multigrid techniques in the context of unstructured meshes. While much use may be made of the experience developed for structured meshes, several important differences and additional requirements must be considered. For example, the special treatment required at refined/unrefined inner boundaries for structured meshes is no longer needed for unstructured meshes. These interfaces represent a breakdown in the structure of the mesh in the former case, and are dealt with automatically in the context of unstructured meshes. The partitioning of the domain into refined and unrefined regions, which is trivial for the structured mesh case, is complicated in the unstructured mesh case by the fact that the boundaries of these domains are a function of the connectivity of the mesh, which is altered during the adaptive refinement procedure. On the other hand, due to the use of pointers and indirect addressing on unstructured meshes, the size, location, shape and number of refined regions may be completely arbitrary, thus permitting the use of general and efficient adaptive refinement procedures.

In the present work, we also require the possibility of dealing with non-nested refinements. Thus, if upon the generation of a new mesh level, a region of the domain is not refined, it is not precluded from being refined later on, upon the generation of subsequent mesh levels. Non-nested refinements of this type require the possibility of direct communication between non-adjacent mesh levels. Such refinement strategies often arise in the solution of compressible flow problems where, for example, the resolution of a downstream flow feature can only proceed once adequate resolution of some particular upstream flow feature has been achieved. Finally, we also require the possibility of solving different equations on the coarse and fine mesh levels. In this manner, we may concentrate on the solution accuracy when time-stepping on the fine grid, and on economy and speed of convergence when visiting the coarser grids.

The above requirements represent an effort to develop the most general possible adaptive multigrid strategy. Unstructured meshes are employed to facilitate adaptivity and deal with complex geometries. The unstructured multigrid algorithm assumes coarse and fine grid levels to be completely independent from one another (except at boundaries). Mesh refinement is permitted on any mesh level, in any region of the domain, and provisions for solving different equations on different mesh levels are included.

In this work, two alternate adaptive multigrid strategies have been developed. In the first approach, which is similar to that previously employed for structured meshes, new mesh levels are formed by generating local fine grid patches through adaptive refinement of the previous coarser grid, and the final solution is obtained on a composite mesh which contains the zonal fine grids and the unrefined regions from the coarse grids. The second method obviates the need for defining a composite mesh to obtain the final solution. The finest mesh level consists of a global mesh which covers the entire domain, generated by successive adaptive refinements of the previous coarser meshes. To maintain efficiency and locality, the underlying coarse meshes are taken as zonal meshes which only cover the region of the domain which lies under a refined area.

2. UNSTRUCTURED MULTIGRID SOLVER

The governing equations to be solved are the Euler equations. These describe the conservation of mass momentum and energy of a fluid in the absence of viscous forces and heat conduction. In integral form, these equations read

$$\frac{\partial}{\partial t}\iint_{\Omega} w\, dxdy + \int_{\partial\Omega} (fdy - gdx) = 0$$

where Ω is a fixed volume with boundary $\partial\Omega$, x and y are cartesian coordinates, and

$$w = \begin{bmatrix} \rho \\ \rho u \\ \rho v \\ \rho E \end{bmatrix} \quad f = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho uH \end{bmatrix} \quad g = \begin{bmatrix} \rho v \\ \rho vu \\ \rho v^2 + p \\ \rho vH \end{bmatrix}$$

ρ is the fluid density, u and v the cartesian velocity components, and p is the pressure. The total energy and enthalpy are respectively given by

$$E = \frac{P}{(\gamma-1)\rho} + \frac{u^2+v^2}{2}, \qquad \text{and} \qquad H = E + \frac{P}{\rho}$$

These equations are discretized in space using a finite-volume vertex scheme, where the flow variables w are stored at the vertices of the triangles. The control volume for each node i is taken as the union of all triangles with a vertex at i, as shown in Figure 1, and the f and g fluxes are integrated along the boundary of the control volume. Explicit artificial dissipative terms are added to the equations to maintain stability. The resulting spatially discretized equations are integrated in time to obtain the steady-state solution using a five-stage Runge-Kutta-type time-stepping scheme. The full multigrid algorithm begins by computing the solution to the problem at hand on a coarse mesh. When convergence has been obtained, a new finer mesh is generated. Initially, (usually for the first two levels), this is performed by globally regenerating a new mesh exhibiting finer resolution in all regions of the domain. However, for subsequent levels, new meshes are generated by local adaptive refinement of the previous coarser grid. The patterns for transferring flow variables and residuals to this new finer grid are then determined. Since the elements of the new finer mesh are in general not nested with respect to the previous mesh, this is a non-trivial task, which is performed using an efficient tree-search algorithm to rapidly identify elements of one mesh which overlap particular elements on another mesh. Once this preprocessing operation has been performed, and the transfer rules determined, the flow variables are transferred to the new fine grid, thus serving as the initial conditions for time-stepping on this mesh. A multigrid saw-tooth cycle is then employed to solve the equations on this new finer mesh, using the previous mesh as a background coarse grid. When convergence is obtained, a third finer mesh is generated, transfer patterns are determined, and the flow variables are transferred to this new mesh. Time-stepping resumes on this mesh using all three meshes as a sequence in the multigrid saw-tooth cycle. This procedure may be repeated as many times as necessary, each time adding a new finer mesh to the multigrid sequence, as shown in Figure 2. A more detailed description of the discretization, time-stepping and application of multigrid to the Euler equations on unstructured meshes can be found in the literature [4,5], and will not be repeated here. However, special characteristics of this algorithm which affect the implementation of a zonal multigrid strategy will be discussed.

The first of these is the use of different artificial dissipation models on coarse and fine meshes. On the finest mesh of the sequence, the main concern is the accuracy of the solution. Thus, a second-order accurate dissipation model is employed. This is constructed as a finite-

volume approximation to an undivided biharmonic operator in the flow variables. In the vicinity of a shock, stronger dissipative terms are required, and thus, a finite-volume approximation to an undivided Laplacian operator is employed. This results in a scheme which is second-order accurate throughout the flow-field, except in the vicinity of a shock, where it becomes locally first-order accurate. On coarser mesh levels, accuracy is no longer a concern, but speed of convergence is now important. Thus, a strongly dissipative first-order accurate formulation is employed, where the artificial dissipation terms are constructed as a finite-volume approximation to an undivided Laplacian operator. Hence, in the interest of optimizing accuracy on the fine mesh and speed of convergence on the coarse mesh, different equations are solved on the coarse and fine meshes. The coarse grid equations also differ from the fine grid equations due to the addition of the defect correction on coarse grids. This term is constructed at each point on the coarse grids as

$$P_{2h} = I_h^{2h} R_h(w_h) - R_{2h}(I'^{2h}_h w_h)$$

i.e. the difference between the restricted fine-grid residuals and the residuals of the restricted fine-grid flow variables computed on the coarse grid. The h and 2h subscripts denote fine grid and coarse grid values respectively. R represents residuals, and w the flow variables. I_h^{2h} and I'^{2h}_h denote the restriction operators employed for transferring residuals and flow variables respectively, to coarser grids, and in general are not identical. The defect correction assures that the evolution of the coarse grid equations are driven by the fine grid residuals. At convergence, the residuals restricted from the fine grid vanish, but the residuals of the restricted flow variables do not. Hence the defect correction does not vanish, and must be equal and opposite to the residuals of the discretized governing equations on the coarse grid. Thus, although the governing equations are discretized conservatively on all grids, at convergence, conservation is only obtained on the finest grid of the sequence.

3. MESH ADAPTIVITY AND DOMAIN DECOMPOSITION

The generation and adaptive refinement of unstructured triangular meshes has previously been described in detail [4,5]. The basis for this construction is the Delaunay triangulation algorithm. Given a set of points in a domain, the Delaunay triangulation of these points represents a unique way of joining them together to form a set of non-overlapping triangles which completely cover the domain. One of the properties of this construction is that no vertex from any triangle may be contained within the circumcircle from any other triangle. This property may be employed to construct a triangulation, or modify an existing triangulation in a sequential manner. Assuming an initial triangulation exists, and a distribution of additional points to be triangulated is at hand, the new points are introduced one at a time into the existing structure. Each time a new point is introduced, the triangles whose circumcircles are intersected by the new point are located. The union of these intersected triangles forms a convex polygonal region, which contains the new point. The connectivity inside this region is then restructured as shown in Figure 3, by replacing the existing structure with that obtained by joining the newly inserted point to all the vertices of the polygonal region.

This construction thus forms an ideal method for adaptive local refinement of an existing unstructured mesh, without recourse to global mesh regeneration. In this case, the initial triangulation consists of the coarse grid to be refined, and the distribution of new points to be inserted is obtained adaptively, by examining the flow solution on the coarse grid. In this work, the first undivided difference of fluid density, computed along edges of the coarse mesh are examined. When the difference along a particular edge is larger than some fraction of the aver-

age difference along all edges of the mesh, a new point is added midway along that edge. When an unstructured mesh is adaptively refined in this manner, the boundaries delimiting refined and unrefined regions are not known at the outset, since they depend on the newly generated mesh connectivity. However, an automatic partitioning of the domain into refined and unrefined regions may be achieved by monitoring the restructuring of the mesh which occurs as new mesh points are introduced. When a new point is triangulated into the mesh, the region which is affected consists of the polygonal domain formed by the union of all triangles whose circumcircles are intersected by the new point (c.f. Figure 3). Thus, all newly formed triangles in this region are flagged as belonging to a refined region. After all the new points have been introduced, the refined regions simply consist of the union of all the flagged triangles. The size, shape, location and number of refined regions may be completely arbitrary. To locate the boundaries delimiting each individual region merely requires identifying mesh edges (and their associated nodes) which simultaneously border on a refined triangle and an unrefined triangle. In addition, information from neighboring boundary points is often required by the flow solution algorithm. These may be determined by performing a tree-search out along all edges emanating from the boundaries of the refined regions. If multiple layers of neighboring boundary points are required, the tree-search may simply be repeated, each time beginning from the newly determined boundary of neighboring points, as shown in Figure 4.

Since the mesh connectivity in the adaptively refined regions is determined by the Delaunay triangulation algorithm, rather than by equal subdivisions of existing coarse grid triangles, the newly formed triangular elements in these regions will not, in general, be nested with the previous coarse grid triangles. Furthermore, to obtain a smooth distribution of elements, the mesh in the newly refined regions is smoothed by slightly redistributing the mesh points in these regions according to a Laplacian filtering operation [4]. Thus, in general, none of the mesh points in the newly refined regions will coincide with the points from the original coarse mesh. Hence, the unstructured multigrid algorithm of [4], which operates on a sequence of unrelated coarse and fine meshes, must still be employed when considering adaptively generated fine meshes. The boundaries (and associated neighboring boundary points) which delimit the refined regions must, however, coincide exactly with mesh edges and nodes from the previous coarser mesh.

4. THE ZONAL FINE-GRID SCHEME

The first zonal multigrid strategy considered is identified as a zonal fine-grid scheme. This scheme follows the approaches previously employed for structured meshes [6,7,10]. To illustrate the differences between this strategy and the global unstructured multigrid algorithm of [5], it is useful to consider a sequence of linear one-dimensional grids as shown in Figures 5a, 5b and 5c. The grids on the bottom of each part of the figure represent coarse meshes, with higher levels corresponding to successive adaptive refinements of this initial mesh. The global multigrid strategy of [5] is depicted in part a) of the figure. Each level consists of a grid covering the entire domain. The first level contains a grid of uniform spacing. The second level is obtained by reproducing the previous level mesh, with added refinement in the right-hand side of the domain. The third mesh level is obtained by further refining the level 2 mesh in the extreme left and right-hand sides of the domain. The final fine mesh is a global mesh of unequal spacings. A fine-grid solution covering the entire domain is obtained by solving the discretized equations on this mesh, using the previous global coarse meshes as multigrid accelerator meshes. However, as can be seen from the figure, there exists regions such as in the center of the domain, where no change in resolution occurs between the various mesh

levels, and thus, time-stepping on all levels in such regions is not necessary. Part b) of Figure 5 depicts the strategy employed for the zonal fine-grid scheme. The level 1 mesh consists of the same global coarse mesh employed in the previous scheme. However, the level 2 mesh is now a zonal fine grid which only covers the region in which the previous coarse grid has been refined. Hence, the portion of the level 2 mesh in Figure 5 a) which exhibits the same resolution as the mesh from the previous level is omitted in the present scheme. Similarly, the level 3 mesh contains only the zonal regions which have been further refined with respect to the previous levels. Note that, because of the non-nested refinement occurring in this case, the left-hand side zonal mesh on level 3 must communicate directly to the level 1 mesh. For structured meshes, it has been proposed that this difficulty may be avoided, and more equally spaced meshes on each level may be preserved if the zonal mesh which appears on the left-hand side of level 3 is repositioned to level 2. However, for unstructured meshes, the repositioning of a group of points on a previous mesh level will alter the connectivity of that mesh, and hence that of all subsequent finer mesh levels. This in turn, will require the recomputation of the intergrid transfer coefficients, as well as altering the boundaries of the zonal refined regions on the various mesh levels. Such a procedure is therefore not considered viable for unstructured meshes.

The zonal fine-grid scheme must be constructed such that it converges to a final global solution which is identical to that obtained on the finest grid using the global multigrid strategy of [5]. This final solution must be obtained on a composite mesh, which is formed by the zonal meshes on the highest mesh level, and portions of the domain on lower mesh levels which do not lie under a refined region from a higher level. From Figure 5 b), the final composite mesh is seen to consist of two zones from level 3, one zone from level 2, and one zone from level 1. Recalling that different equations are solved for on the coarse and fine meshes in [4,5], the equivalent procedure in this case consists of employing the second-order accurate dissipation model in all zones of the various mesh levels, which compose the final composite mesh, and the first-order accurate dissipation model in all other regions. Furthermore, the zones on the various mesh levels which do not belong to the final composite mesh, must lie under a refined region from a higher mesh level, and thus, the equations in these regions must be augmented by the defect correction from the higher mesh levels. Hence, it can be seen that, on any given mesh level, different equations may be solved for in various zones of the domain. Conservation, therefore, cannot be guaranteed across boundaries delimiting two such zones, and is also not preserved in the zones where the defect correction is added. However, conservation and second-order accuracy must be preserved on the final composite mesh. Due to the different stencils employed in the first and second-order accurate dissipation models, this requires the use of extra boundary information on the zonal fine meshes. The discretization of the convective terms in the Euler equations, as well as the first-order dissipation model, only require information about the immediate neighbors of a point, which corresponds to a three-point stencil on the linear one dimensional grids of Figure 5. Conversely, the second-order dissipation model, which requires the discretization of a biharmonic operator, employs information about the neighbors of a point, as well as about the neighbors of these neighbors, thus corresponding to a five-point stencil on these one-dimensional grids. A simplified two-level mesh system is shown in Figure 6. Part a) of the figure depicts the meshes employed in the global multigrid strategy, while part b) illustrates the two meshes employed in the zonal fine-grid scheme. From level 1 in part b), the points i = 1, 2, 3, 4 are seen to be part of the final composite mesh, whereas points i = -3, -2, -1, 0 lie under a refined region from level 2, and thus form a coarse grid accelerator mesh. Point i = 0 delimits the boundary between these two zones on level 1.

In order to illustrate the use of extra boundary information, points i = 1, 2 have been reproduced in dashed lines on level 2. Considering only three-point stencils, in part b) of the figure, it is seen that the stencil at point 0 on level 1 differs from the corresponding three-point stencil at point 0 on level 2. In fact, the stencil at this point on level 2 is identical to that which is employed at this point on the finest mesh of the global multigrid strategy (level 2, part a) of Figure 6). Recalling that the purpose of this construction is to obtain a discretization on the fine composite mesh which is identical to that employed on the finest mesh of the global multigrid strategy, we conclude that the values at point 0 on level 2 in the zonal fine-grid scheme must be recomputed using information from point 1. Point 1 itself need not be recomputed, since the same three-point stencil applies to this point on both levels. However, when considering a five-point stencil, a similar argument indicates that points 0 and 1 must both be recomputed on level 2, making use of information from points 2 and 3. Thus, in general, time-stepping must be performed on one extra layer of boundary points for the zonal fine meshes, employing information from three extra layers of boundary points. In practice, due to non-essential details in the present solver, time-stepping was performed on two extra layers of boundary points, making use of information from a total of four layers of boundary points. In extending this to unstructured meshes in two dimensions, the success of such a strategy requires that such boundary points and their neighbors, as well as the connectivity of the mesh in these regions, be identical on the two adjacent mesh levels, and that the multigrid transfer of variables and residuals in these regions be effected by simple injection.

To illustrate the effectiveness of the zonal fine-grid scheme, the Euler equations were solved in the transonic regime over a tandem airfoil configuration on a sequence of unstructured and adaptive meshes, first by means of the global unstructured multigrid approach, and then using the zonal fine-grid strategy. The finest mesh of the sequence for the global multigrid scheme, which also corresponds to the final composite mesh of the zonal scheme, is depicted in Figure 7d. The mesh contains 5632 nodes and represents the eighth level in the global multigrid sequence. The first two levels of the sequence were generated globally, and the subsequent six levels were obtained by adaptive mesh refinement. Parts a), b), and c) of Figure 7 also illustrate the 4th, 5th, and 6th mesh levels (out of a total of eight levels) employed in the zonal fine-grid scheme. As expected, these levels are comprised of various zonal fine meshes, corresponding to the regions of the domain which have been adaptively refined from the previous level. From these figures, non-nested refinement is seen to occur in some regions of the domain between levels 4 and 5. On the higher mesh levels, the refined regions are seen to occur mainly in the vicinity of the shocks and the stagnation points of the flow. The eighth and finest mesh level for the zonal scheme is similar to that shown in Figure 7c, with the exception that the zonal meshes are much smaller and more densely packed. In Table 1, the total number of mesh points used in the zonal meshes on each mesh level, including the required extra boundary points, are compared with the number of mesh points on each mesh used in the global multigrid strategy. Starting from identical global coarse meshes, the zonal fine-grid scheme results in fewer points on higher mesh levels, thus offering the potential for increased solution efficiency. The computed Mach contours in the flow-field for the converged solution on the mesh of Figure 7d are shown in Figure 8. The freestream Mach number is 0.7, and the incidence of the flow is 3°. A strong shock is formed on the upper surface of the forward airfoil, and a weaker shock is present on the upper surface of the rear airfoil. The computed surface pressure distribution for both airfoils is shown in Figure 9. The shocks and other regions of large gradients in the flow-field are well resolved by the adaptive mesh of Figure 7d.

The convergence of the two schemes is depicted in Figure 10. For the global multigrid scheme, convergence to steady-state is monitored by plotting the RMS average of the residuals of the discretized continuity equation (i.e. $\frac{\partial \rho}{\partial t}$) on the finest grid, versus the number of multigrid cycles. For the zonal fine-grid scheme, convergence is measured by the RMS average of the density residuals, as previously, but taken over the final composite mesh. It is important to realize that the criterion for convergence in this case must be the reduction of the residuals computed on the composite mesh, for it is the steady-state solution to the equations discretized on this mesh which is ultimately being solved for, and which corresponds identically to the converged solution obtained with the global multigrid algorithm. From Figure 10, the rate of convergence per multigrid cycle is seen to be nearly identical for both schemes, indicating that the omission of time-stepping in regions of the domain where the mesh is not refined has little effect on the overall convergence rate. The simultaneous solution of different equations in various zones on a given mesh level, and the non-conservative formulation of these equations in particular zones of the domain are also seen to have no adverse effect on the rate of convergence. Although reduction of the residuals far below the levels of truncation error is not useful in practice, the residuals on the final composite mesh for the zonal fine-grid scheme were driven all the way to machine zero, to ensure the resulting composite mesh discretization was fully consistent with the multilevel zonal mesh discretizations, which are used to drive the convergence. Figure 11 depicts the convergence rate of the two schemes in terms of work units, where a work unit corresponds to the time required for one time-step on the finest grid of the global multigrid sequence (c.f. Figure 7d). For this case, the zonal fine-grid scheme is seen to produce a savings of roughly 50% over the global multigrid strategy, due to the lower number of operations required in a single multigrid cycle. It should however be noted that the comparison of Figure 11 only represents an estimate of the relative costs of the two schemes, which was obtained by comparing the total number of grid points where time-stepping is performed, on all mesh levels, for both schemes. This was necessary, since the zonal code is at present a non-optimized experimental code.

5. THE ZONAL COARSE-GRID SCHEME

The use of a composite mesh and the requirement of solving different equations in various zones on a given mesh level in the previous scheme leads to considerable coding complications. Thus, a simpler method of combining adaptive mesh refinement with multigrid has been pursued. This new method, identified as the zonal coarse-grid scheme, operates on a global fine mesh and employs a series of zonal coarse-grid accelerator meshes to enhance convergence. On the finest mesh, a fully conservative discretization of the governing equations with a second-order accurate dissipation formula is employed. On the zonal coarse meshes, the discretized governing equations are augmented by the defect correction from higher mesh levels, and a first-order accurate dissipation model is employed. Thus, on any given mesh level, similar equations are solved for in all zones of the domain. Because the final solution is obtained on a global fine mesh, as opposed to a composite mesh, the tasks of convergence monitoring, adaptive mesh refinement, and postprocessing of the final solution are greatly simplified.

To illustrate the zonal coarse-grid strategy, consider the linear one-dimensional grids of Figure 5. Part a) of the figure depicts the equivalent global multigrid strategy of [5], and part c) the strategy of the present method. As can be seen from the figure, the zonal coarse-grid method operates on the same global fine mesh as the global multigrid method, thus ensuring convergence to identical final solutions for both schemes. On the coarser mesh levels, regions

of the domain which exhibit the same resolution as the next finer mesh level are omitted in the present scheme, thus yielding a series of zonal coarse meshes. The relationship between the present zonal coarse-grid scheme and the zonal fine-grid scheme of the previous section can be examined by observing Figures 5 and 6. For the simplified case of one-dimensional grids, assuming for the purposes of this argument that identical equations are solved for on all mesh levels, the two schemes are seen to differ only by the order in which time-stepping is performed on the unrefined regions of the domain. Thus, for example, in the two-level case of Figure 6, the right-hand side of the domain, where refinement does not occur, is assigned to level 1 for the zonal fine grid scheme, whereas, for the zonal coarse-grid scheme it would be assigned to level 2.

As in the previous scheme, additional boundary points surrounding the zonal meshes are required. However, in this case, these are no longer required to ensure the accuracy and consistency of the final solution, since they only apply to the coarse grid accelerator meshes, but are essential in preventing a deterioration in the rate of convergence of the overall scheme. The success of any multigrid algorithm depends on the ability of the driving scheme to rapidly damp out all high-frequency error components. The present method achieves this by the use of a five-stage Runge-Kutta-type time-stepping scheme. Assuming a compact stencil (as is the case for the coarse grid discretizations), a five-stage scheme requires information from neighboring points up to five levels deep. Thus, for points immediately bordering on a zonal coarse mesh (and to a lesser extent, for interior neighbors), information from up to five extra layers of boundary points is required to ensure the same high-frequency damping characteristics as for interior points. In practice, good convergence rates may be obtained without using a full five layers of boundary points, and significant convergence deterioration is only observed when less than two extra layers are employed. These practical findings may however be case dependent, since they may be easily affected by the relative shape and size of the zonal coarse grids.

The generation of the various mesh levels in the zonal coarse-grid strategy may be integrated into the solution process as follows. A global coarse mesh is first generated, and a converged solution is obtained on this mesh. A duplicate copy of this mesh is then stored in memory. The new global fine mesh is then obtained by adaptively refining the current version of this global coarse mesh, while the automatic domain decomposition method of Section 3 is employed to monitor the regions of the domain where mesh restructuring due to adaptive refinement occurs. The multigrid transfer patterns between the previous coarser mesh which has been stored in memory, and the new current fine mesh, are then determined. Using these transfer patterns, it is then possible to determine the regions of the coarse grid which lie under a refined region on the new fine mesh. All regions of the coarse mesh which do not lie under a refined region are then omitted, thus forming a series of zonal coarse meshes on this level. This procedure may be repeated each time a new finer mesh is generated. Note that this strategy also allows for non-nested refinements, and may require communication between non-adjacent mesh levels in various zones of the domain (c.f. Figure 5c).

The zonal coarse-grid strategy was employed to solve the same test case as described in the previous section. A total of eight mesh levels were employed, with the finest mesh being identical to that employed by the global multigrid strategy, which is depicted in Figure 7d. The zonal coarse meshes on levels 3, 4, and 5 cover the same regions of the domain as the level 4, 5, and 6 level meshes of the previous scheme, respectively, as depicted in Figure 7. In Table 1, the number of mesh points employed on each mesh level with the zonal coarse-grid scheme is compared to the number of mesh points employed on the various mesh levels with

the previous two schemes. While the finest mesh level in the present scheme is identical to that employed in the global multigrid scheme, savings result from a reduction of the number of points on the coarser mesh levels. Although the relative sizes of the meshes on each level differ between the zonal coarse and fine grid schemes, the total number on all levels for both schemes is identical, further illustrating the notion that the present scheme may be thought of as a repositioning of the unrefined mesh regions from the zonal fine-grid scheme to higher mesh levels (c.f. Figure 6). The final solution obtained at convergence is identical to that obtained with the previous two methods, and is illustrated in Figures 8 and 9. In Figure 10, the multigrid convergence rate of the present scheme, as measured by the RMS average of the density residuals on the global fine mesh versus the number of multigrid cycles, is compared with the convergence rates of the other two schemes. As expected, these convergence rates are approximately equal, indicating that the omission of time-stepping in the regions of the coarse grids which are duplicated without refinement on the next finer mesh levels, does not significantly affect convergence. In Figure 11, the relative cost of this scheme is compared with that of the previous two schemes, by plotting convergence rate versus the number of work units, where a work unit corresponds to the amount of time required for a single time-step on the finest mesh of the of the sequence. As previously, an estimate of the amount of work required in a multigrid cycle for each individual scheme was obtained by summing the number of points where time-stepping was performed on each mesh level. A straight comparison of the amount of CPU time required by the various methods was not possible, since the zonal coarse-grid code is a non-optimized experimental code. From Figure 11, it is seen that the zonal coarse-grid scheme offers a potential savings of 50% over the global multigrid strategy for this case, which is roughly the same as that produced by the zonal fine grid scheme.

6. CONCLUSION

Two methods for efficiently combining adaptive meshing techniques and a multigrid strategy in the context of unstructured meshes have been described. The first method makes use of local fine meshes and obtains the final solution on a composite mesh which consists of a union of zonal meshes of varying resolution from different multigrid mesh levels. The second method operates on a global fine mesh and accelerates convergence with a sequence of local coarse meshes. The two methods were shown to yield comparable solution efficiencies for the computation of inviscid transonic flow, while showing substantial gains as compared to a previously employed global multigrid strategy. The exact amount of savings achieved with these methods will in general be very case dependent. The computation of flows with numerous features such as shocks and high gradients, which requires high mesh resolution in many regions of the domain, will yield only minor improvements. However, flows with a small number of highly localized features offer the possibility of truly large savings.

As both schemes exhibit similar performance, the relative benefits of one scheme over the other are mainly concerned with such issues as ease of implementation and simplicity of coding. In this respect, the zonal coarse-grid scheme would appear to constitute the preferred choice, as it may more easily be implemented as a modification to an existing global multigrid strategy. At present, the codes developed for both schemes in this work are only experimental programs. At each mesh level, an entire global grid is stored in memory, but time-stepping and residual evaluations are omitted in the appropriate zonal regions of the domain. Future work will concentrate on further optimizing these codes, both in terms of memory and CPU time, by storing only the zonal meshes in memory, and by increased vectorization. Finally, these schemes need not be restricted to the Euler equations in two-dimensions, and may be employed

for the solution of alternate equations on unstructured meshes in two and three dimensions, provided an efficient global multigrid strategy has previously been developed. As the present global unstructured multigrid strategy has been successfully applied to the Navier-Stokes equations [12], it appears these schemes should easily carry over to the solution of viscous flows.

REFERENCES

1. Jameson, A., Baker, T. J., and N. P. Weatherill, "Calculation of Inviscid Transonic Flow over a Complete Aircraft", *AIAA paper 86-0103,* January 1986.
2. Stoufflet, B., Periaux, J., Fezoui, F., and Dervieux, A., "Numerical Simulation of 3-D Hypersonic Euler Flows Around Space Vehicles Using Adapted Finite Elements", *AIAA paper 87-0560* January 1987.
3. Lohner, R., Morgan, K., Peraire, J., Vahdati, M., "Finite Element Flux-Corrected Transport for the Euler and Navier-Stokes Equations", *Int. J. Num. Meth. Fluids,* Vol 7, 1987, pp. 1093-1109
4. Mavriplis, D. J., "Multigrid Solution of the Two-Dimensional Euler Equations on Unstructured Triangular Meshes", *AIAA Journal,* Vol 26, No. 7, July 1988, pp.824-831
5. Mavriplis, D. J., and Jameson, A., "Multigrid Solution of the Euler Equations on Unstructured and Adaptive Meshes", *ICASE Report No. 87-53, Proceedings of the Third Copper Mountain Conference on Multigrid Methods, Lecture Notes in Pure and Applied Mathematics,* Ed S. F. McCormick, Marcel Dekker Inc., April 1987, pp. 413-430.
6. Berger, M. J., and Jameson, A., "An Adaptive Multigrid Method for the Euler Equations", *9th Int. Conf. on Num. Methods in Fluid Dynamics,* Lecture Notes in Physics, Vol 218, 1984.
7. Usab, W. J., "Embedded Mesh Solutions of the Euler Equations Using a Multiple-Grid Method", *PhD Thesis,* Massachusetts Institute of Technology, December 1983.
8. Dannenhoffer, J. F., "Grid Adaptation for Complex Two-Dimensional Transonic Flows", *Sc.D Thesis,* Massachusetts Institute of Technology, August 1987.
9. Ni, R. H., "A Multiple-Grid Scheme for Solving the Euler Equations", *AIAA Journal,* Vol 20, No. 11, November 1982, pp. 1565-1571
10. McCormick, S., and Thomas, J., "The Fast Adaptive Composite Grid (FAC) Method for Elliptic Equations", *Mathematics of Computations,* Vol 46, 1986, pp. 439-456.
11. Heroux, M., McCormick, S. F., McKay, S., and Thomas, J. W., "Applications of the Fast Adaptive Composite Grid Method", *Proceedings of the Third Copper Mountain Conference on Multigrid Methods, Lecture Notes in Pure and Applied Mathematics,* Ed S. F. McCormick, Marcel Dekker Inc., April 1987, pp. 251-266.
12. Mavriplis, D. J., Jameson, A., Martinelli, L., "Multigrid Solution of the Navier-Stokes Equations on Triangular Meshes" *AIAA paper 89-0120, ICASE Report No. 89-11* January 1989.

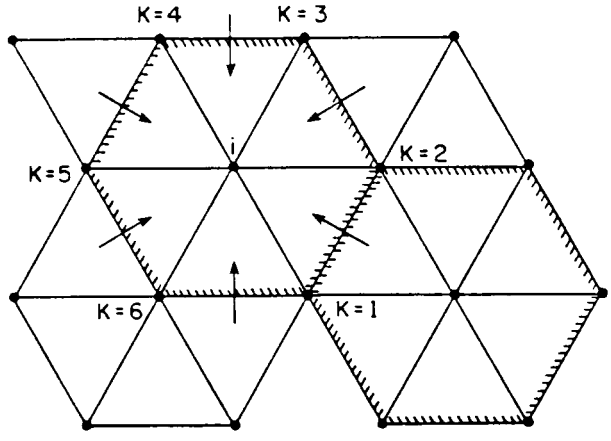

Figure 1
Control Volume for the Vertex Discretization Scheme

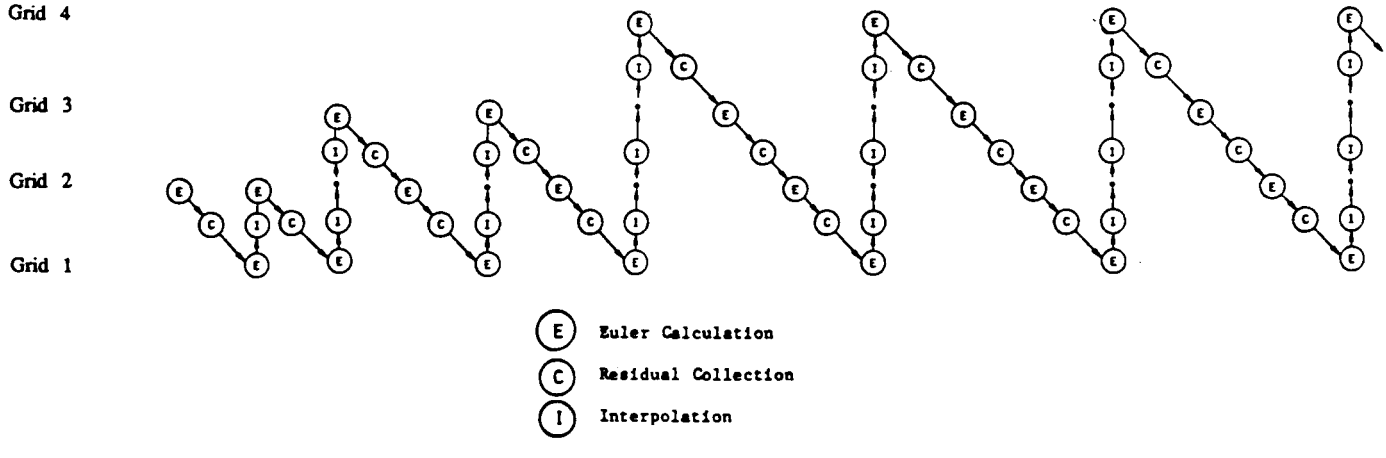

Figure 2
Full Multigrid Strategy using the Saw-Tooth Cycle

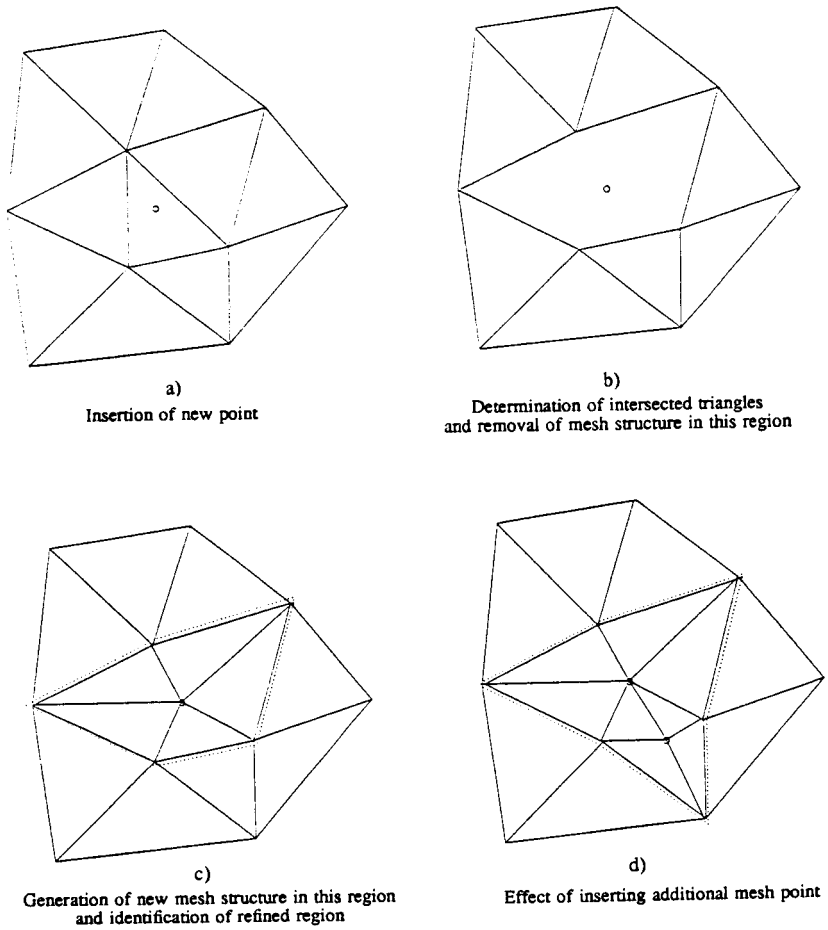

Figure 3
Various Steps Illustrating the Local Adaptive Refinement of the Unstructured Triangular Mesh
Circles Represent Newly Inserted Mesh Points
Dotted Line Delineates the Newly Refined Region

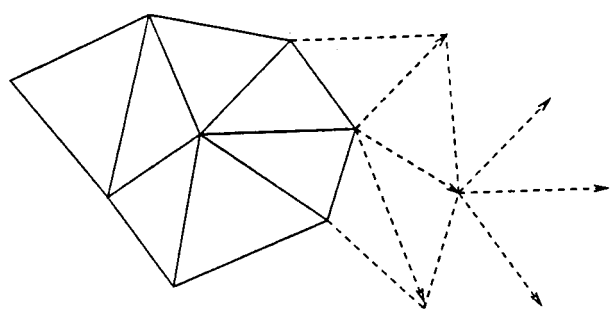

Figure 4
Tree-Search Outwards From the Boundary of a Zonal Mesh to Determine Extra Layers of Boundary Points
Solid Lines Represent Zonal Mesh
Dashed Lines Represent Regions of Extra Boundary Points

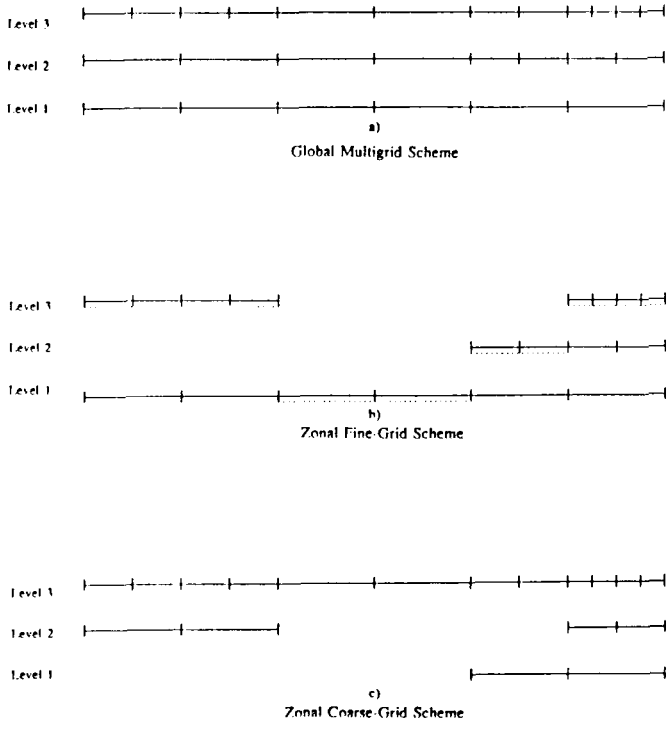

Figure 5
Illustration of the Three Different Multigrid Strategies for the Simplified Case of a Three Level System of Linear One-Dimensional Meshes
Dotted Underline in Part b) denotes Final Composite Mesh

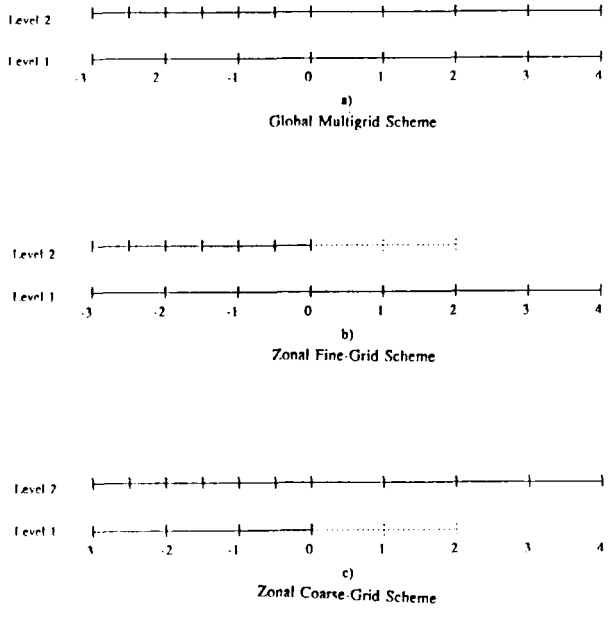

Figure 6
Illustration of the Three Different Multigrid Strategies for the Simplified Case of a Two Level System of Linear One-Dimensional Meshes Showing the Use of Extra Layers of Boundary Points for the Two Zonal Schemes

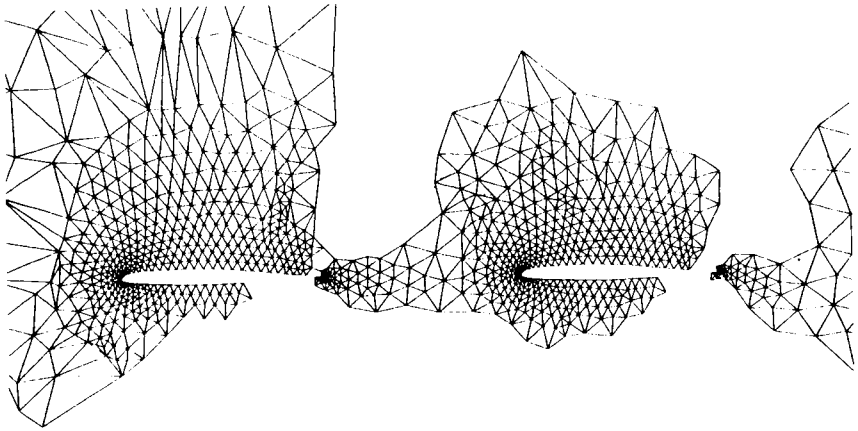

a)
Zonal Fine-Grid Level 4

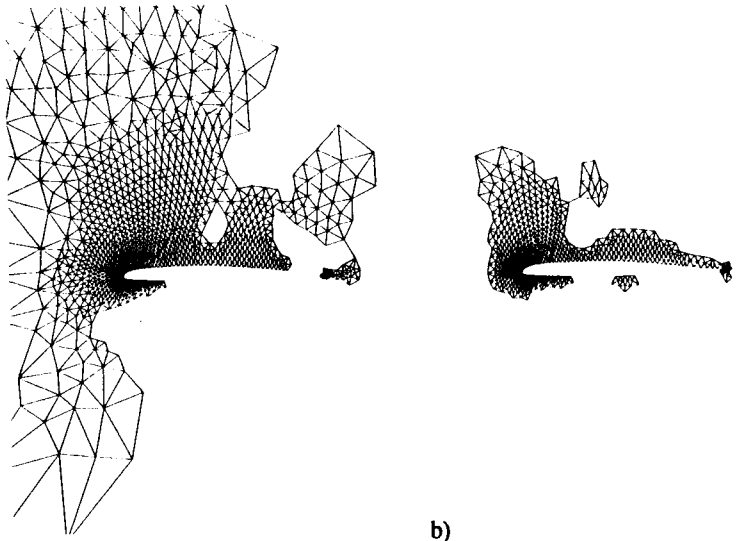

b)
Zonal Fine-Grid Level 5

c)
Zonal Fine-Grid Level 6

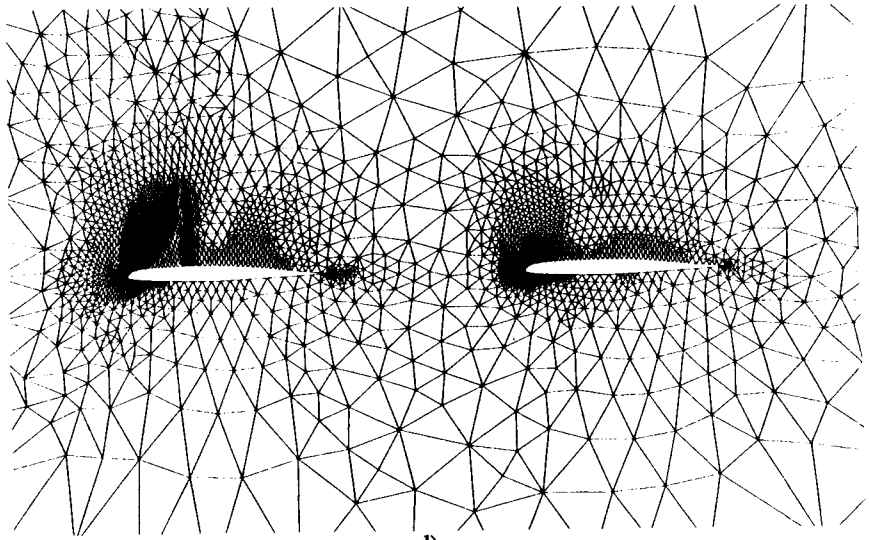

d)
Final Composite Mesh
(8th Level in Global Multigrid Scheme)

Figure 7
Illustration of Several of the Various Mesh Levels Employed in the Zonal Fine-Grid Scheme
Including the Final Composite Mesh which is equivalent to the finest (8th) Mesh Level
in the Global Multigrid Scheme and in the Zonal Coarse-Grid Scheme

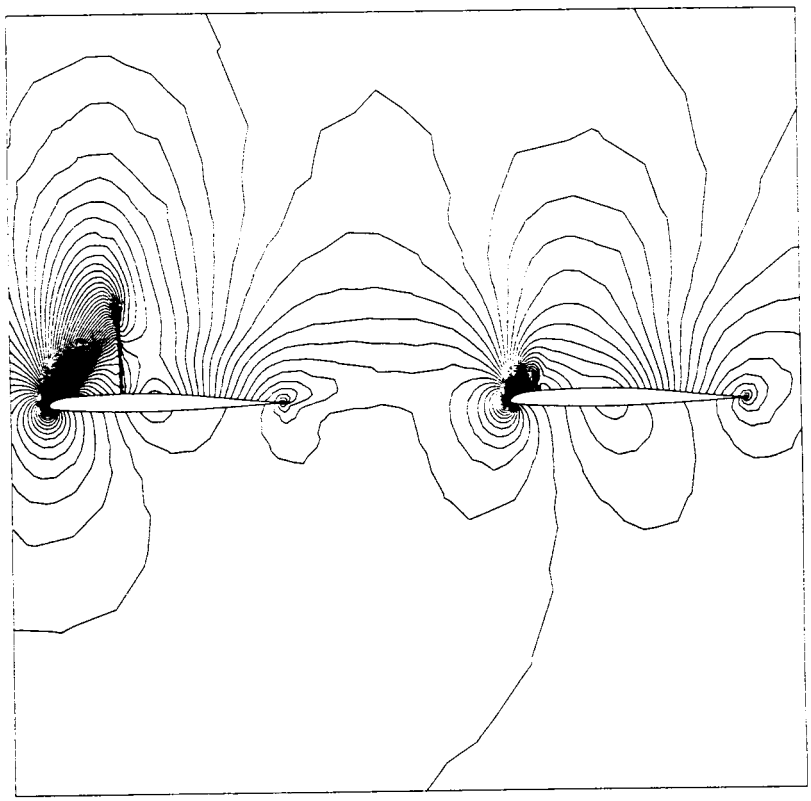

Figure 8
Flow-Field Mach Contours of the Solution of Transonic Flow over Two Airfoils in Tandem
Computed on the Mesh of Figure 7 d) Mach = 0.7, Incidence= 3°

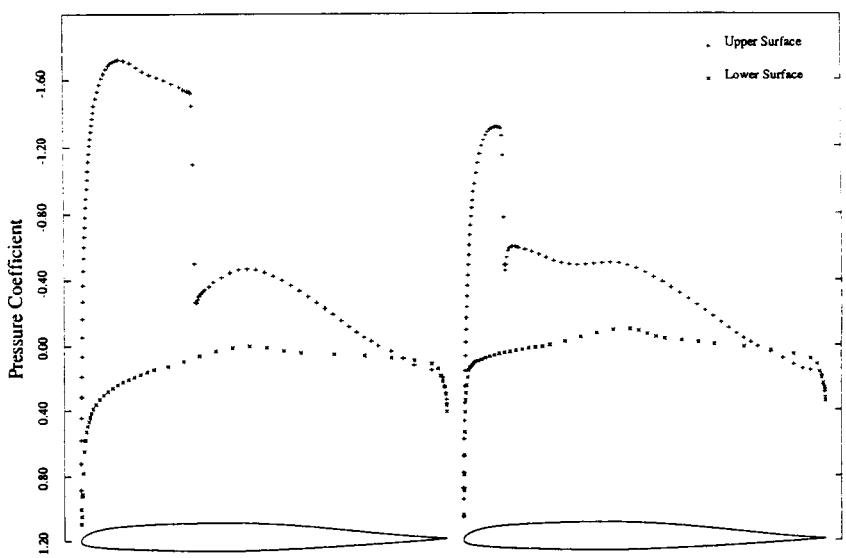

Figure 9
Surface Pressure Distribution of the Transonic Flow Solution over Two Airfoils in Tandem
Computed on the Mesh of Figure 7 d) Mach = 0.7, Incidence= 3°

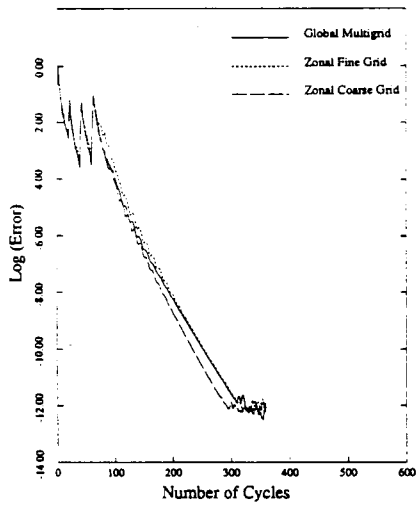

Figure 10
Convergence Rate of the Three Multigrid Methods for the Solution of Transonic flow over Two Airfoils in Tandem as Measured by the RMS Average of the Density Residuals Throughout the Flow-Field versus the Number of Multigrid Cycles

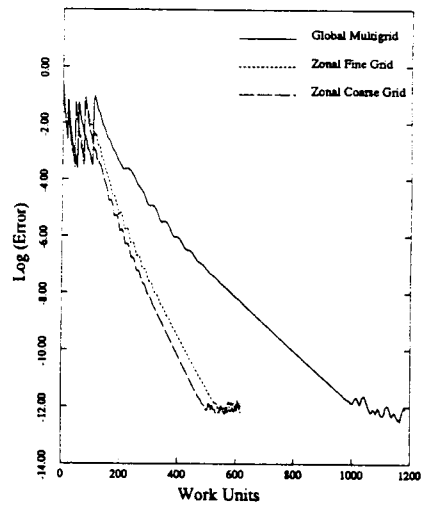

Figure 11
Convergence Rate of the Three Multigrid Methods for the Solution of Transonic flow over Two Airfoils in Tandem as Measured by the RMS Average of the Density Residuals Throughout the Flow-Field versus the Number of Work Units

	Number of Mesh Points		
Mesh Level	Global Multigrid Scheme	Zonal Fine Grid Scheme	Zonal Coarse Grid Scheme
8	5632	1239	5632
7	5038	1604	645
6	4323	2496	889
5	2872	2605	1045
4	1413	1329	1146
3	702	702	618
2	190	190	190
1	55	55	55
Total	20225	10220	10220
Fraction	100%	50.5%	50.5%

Table 1
Comparison of the Number of Mesh Points Employed on Each Mesh Level for the Three Multigrid Schemes

APPENDIX A

ON THE USE OF W-CYCLES FOR ADAPTIVE UNSTRUCTURED MULTIGRID METHODS

The results shown in the main body of this paper were obtained using a multigrid saw-tooth cycle, where one time-step is performed on each mesh level when proceeding from coarse to fine mesh levels, and prolongation of the corrections is performed, without any time-stepping when proceeding from the coarse mesh levels back to the fine mesh level. For Navier-Stokes calculations [12], the use of a multigrid W-cycle has often been found to provide more robust convergence properties. While a simple saw-tooth cycle (or V-cycle) results in the same number of time-steps or relaxation sweeps being performed on each mesh level within a single cycle, the W-cycle is a recursive strategy which performs a larger number of time steps on the coarser mesh levels. Thus, a W-cycle results in a single time step being performed on the finest mesh level, two time steps on the next coarser level, and in general, 2^{n-1} time steps on the nth coarser level, as shown inFigure A.1.

In section 5, it was pointed out that the main difference between the zonal fine-grid scheme and the zonal coarse-grid scheme was that, in the former case, the unrefined regions of the domain are assigned to the coarse mesh levels, whereas in the latter case, these are assigned to the fine mesh levels, as illustrated in Figure 6. When a saw-tooth or V-cycle is employed, this merely corresponds to changing the relative order in which time-stepping is performed in these different regions of the domain. However, with the use of a W-cycle, assigning a region or a portion of a zonal mesh to a different mesh level, also has the effect of altering the number of time steps performed on this portion of the mesh, within the multigrid cycle, since a W-cycle results in more frequent visits to the coarser mesh levels. It is therefore not surprising that the zonal coarse-grid scheme could not be made to converge with a W-cycle. The zonal fine-grid scheme, on the other hand, would converge using a W-cycle, provided all successive refinements were nested. The presence of non-nested refinements, as illustrated in Figure 5, requires communication between non-adjacent mesh levels, which is not compatible with a W-cycle strategy. This situation can be remedied by shifting the regions of non-nested refinements to lower mesh levels, such that the resulting sequence of meshes contains only nested refinements. However, for unstructured meshes, this requires the restructuring of various mesh levels, as well as the redefinition of the boundaries delimiting the refined and unrefined regions on each of these levels, as described in section 4. Thus it would seem that the zonal multigrid strategies described in this paper are not well suited for use with a W-cycle.

However, the straight forward application of a W-cycle to any adaptive multigrid strategy will in general yield low solution efficiencies. In such strategies, a large number of mesh levels are often employed. In the present work, eight levels were used, but problems with even more levels are not uncommon. This is possible when adaptive meshes are employed, since the number of mesh points does not necessarily increase rapidly with additional mesh levels. Since

the number of time steps or relaxation sweeps on the coarser mesh levels increases as 2^{n-1} for n mesh levels, the cost of a W-cycle can be seen to increase rapidly with an increasing number of mesh levels. When global (as opposed to locally adaptive) mesh refinement is employed in two dimensions, the number of mesh points on each successively coarser grid is four times smaller than that on the previous finer grid. Hence, normalizing by the amount of work required for a time step on the finest mesh level, a complete W-cycle involves an amount of work which is given by

$$1 + \frac{2}{4} + \frac{4}{16} + \cdots + \frac{2^n}{4^n} < 2$$

This amount of work is thus bounded and is always less than two fine-grid work units, regardless of the number of mesh levels. However, when the global adaptive multigrid strategy of [5] is employed, the number of mesh points on successively coarser grids will be much greater than 1/4 the number of points on the previous finer grid, and thus the above bound on the total amount of work does not hold any longer. To be sure, the amount of work in a single W-cycle on such meshes is still bound (< 2) with respect to a work unit performed on a globally refined mesh with the same resolution as the highest resolution of the finest adaptive mesh of the sequence. However, since such a mesh is never employed in the calculations, and may contain several orders of magnitude more points than the finest adaptive mesh of the sequence, this bound is of little significance. The same general arguments may be made in the case of the zonal multigrid strategies.

Hence, the most economical way to implement a W-cycle with a sequence of adaptively generated meshes is to pursue a coarsening strategy, which attempts to minimize the number of points on the coarser grid levels. This type of strategy may be implemented as follows. We begin with an initial coarse mesh, which is then adaptively refined to form a new finer mesh. This new fine mesh contains newly introduced grid points from the refinement procedure, as well as old mesh points which are common to the initial coarse mesh. The mesh may also be partitioned into regions which have been locally restructured upon refinement of the previous mesh, and regions in which the structure has not been altered. It is important to realize that the restructured regions contain both coarse grid points as well as adaptively introduced points, whereas the unaltered regions contain only points from the original coarse mesh. Similarly, a fine mesh which results from n-1 adaptive refinements of the initial coarse mesh, may be decomposed into n zones, with the existing structure in each zone k resulting from the adaptive refinement of the k-1st mesh. For example, the k=1 region is such that it contains the same resolution and structure as the original coarsest mesh of the sequence, and the k=2 zone contains only regions which have been adaptively refined once throughout the sequence, i.e upon the generation of the second coarse mesh. Hence, each region k may contain points which are common to any of the other k,k-1,..,2,1 grid levels. A new uniformly coarser mesh may be obtained by examining each region of the current fine mesh level n, and deleting only those mesh points in each region k which are also common to the kth mesh level. Since these points were added in the k-1st adaptive refinement step, they represent the most recently (finest) introduced mesh points in that region, and must therefore be therefore omitted. Once this operation has been performed in all such zones covering the entire domain, the remaining mesh points are fed back into the Delaunay triangulation algorithm, and a new mesh is triangulated. Each time a new n+1st mesh level is adaptively generated, prior to solving the equations on this mesh, a new sequence of n coarse meshes must be regenerated. For each coarser mesh, a lower resolution is obtained throughout the domain, except in regions where the original coar-

sest mesh resolution is reached, which represents the lowest permitted resolution.

Figure A.2 depicts a coarse mesh generated in this fashion for the test problem described in the main body of this paper. This represents the 6th mesh level in a sequence of eight meshes. The corresponding mesh employed in the global multigrid strategy (from [5]), is shown in Figure A.3. The two meshes exhibit similar resolution in certain regions of the domain (where refinement has taken place at every mesh level), but the mesh generated by means of the coarsening strategy exhibits lower resolution in the remaining regions, thus reducing the amount of work performed on the coarse levels. Table A.1 illustrates the reduction in the size of the coarse meshes obtained with the coarsening strategy for this case, as well as the amount of work required on each level within a single multigrid cycle. Whereas a single W-cycle performed on the original sequence of global meshes requires over 20 work units per cycle, this may be reduced to roughly 12 work units using the meshes of the sequence generated with the coarsening strategy. While this represents an improvement over the previous case, roughly 3.5 times the work per cycle is still required, as compared to the saw-tooth cycle. This is partly due to the fact that the coarsest permitted resolution in any region of the domain in the coarsening strategy was determined by the resolution of the third mesh level of the entire sequence, since this represents the coarsest mesh of the adaptively generated meshes (i.e k=1 level for coarsening), the other levels having been generated by global mesh regeneration. Although this mesh contains only 12% of the points of the finest mesh of the sequence, a single W-cycle requires 32 time-steps on this mesh, thus resulting in an expense of 4 work units for this level alone. Hence, to achieve better efficiency for a W-cycle, it appears an even coarser base grid resolution for the adaptive sequence of meshes (i.e. the coarsest permitted resolution of the coarsening strategy) must be employed.

Although this approach provides a framework for employing W-cycles in conjunction with adaptively generated meshes, it is rather cumbersome, since at each adaptive refinement of the current mesh, an entire sequence of coarse accelerator meshes must be regenerated. Furthermore, it has been found that the coarse meshes produced by this method may contain large variations in the sizes of neighboring elements, which may require a certain amount of postprocessing (i.e. smoothing of the mesh) to avoid deterioration of the overall multigrid convergence rate. Hence, in general, saw-tooth and V-cycles remain the preferred choice for adaptive multigrid strategies.

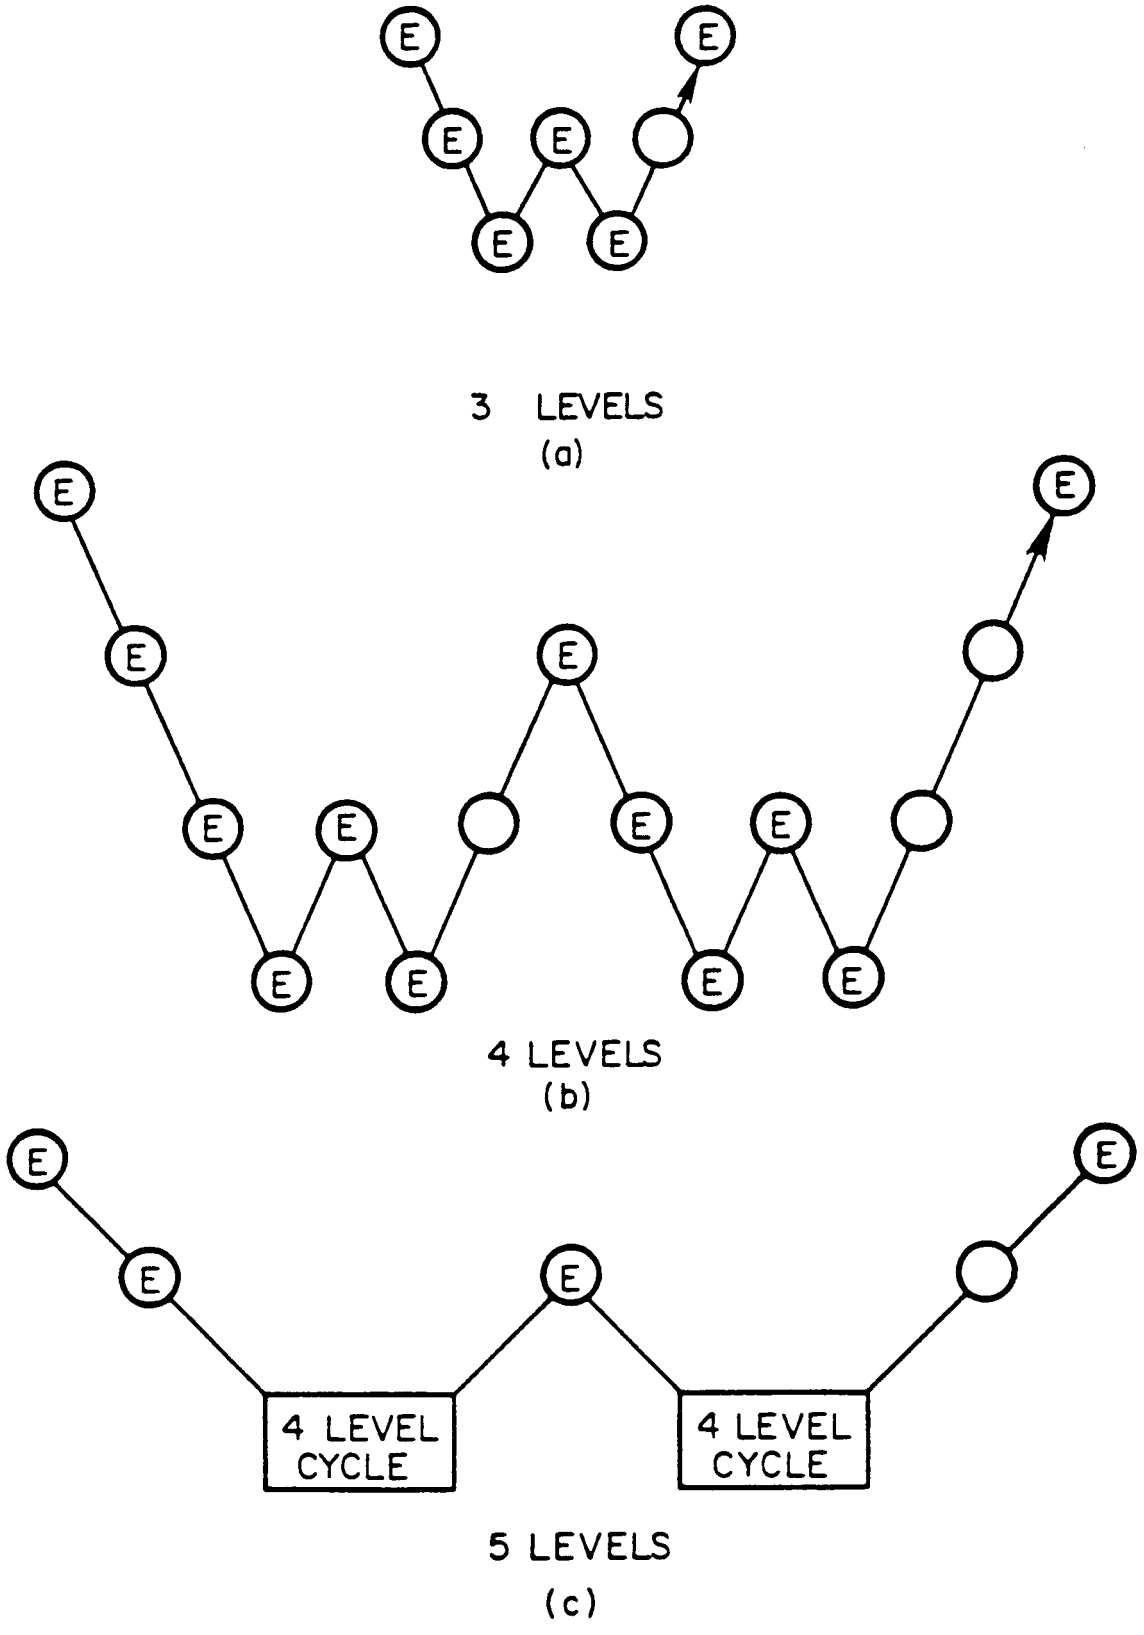

Figure A.1
The Multigrid W-Cycle

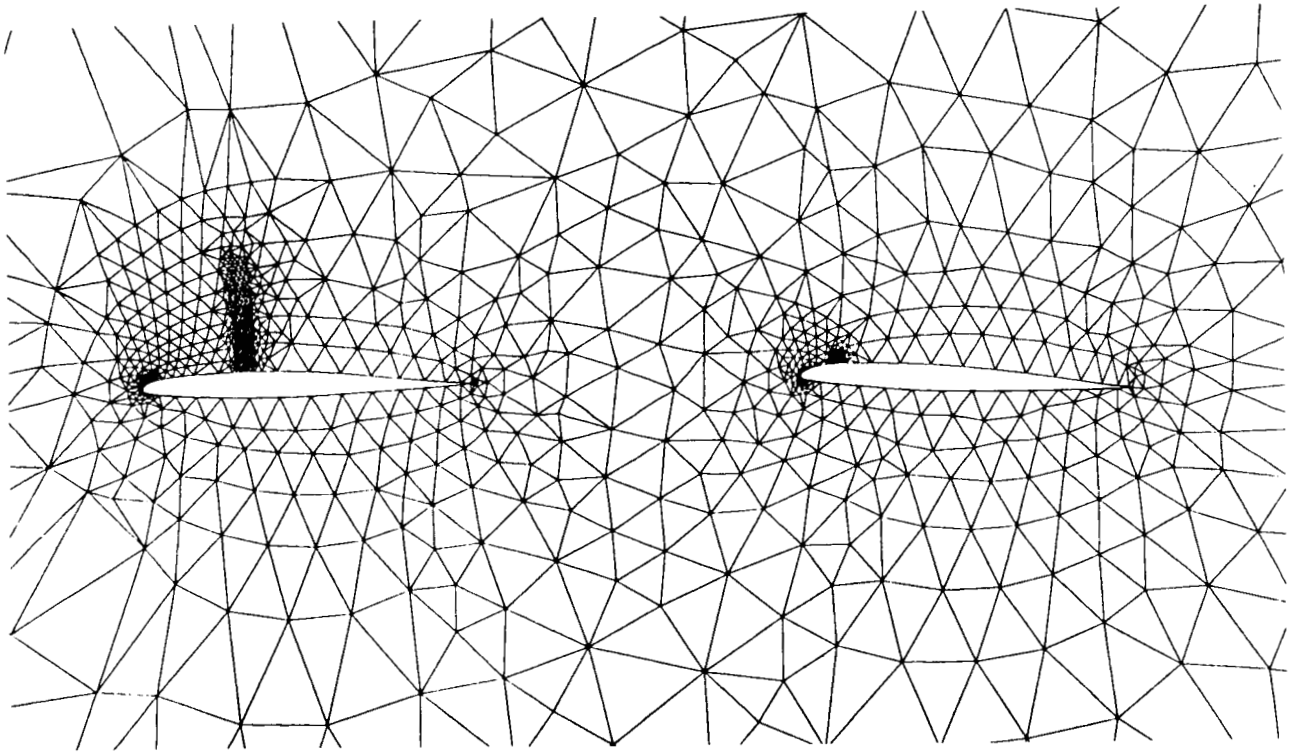

Figure A.2
Sixth Mesh Level Employed in the Multigrid Coarsening Strategy

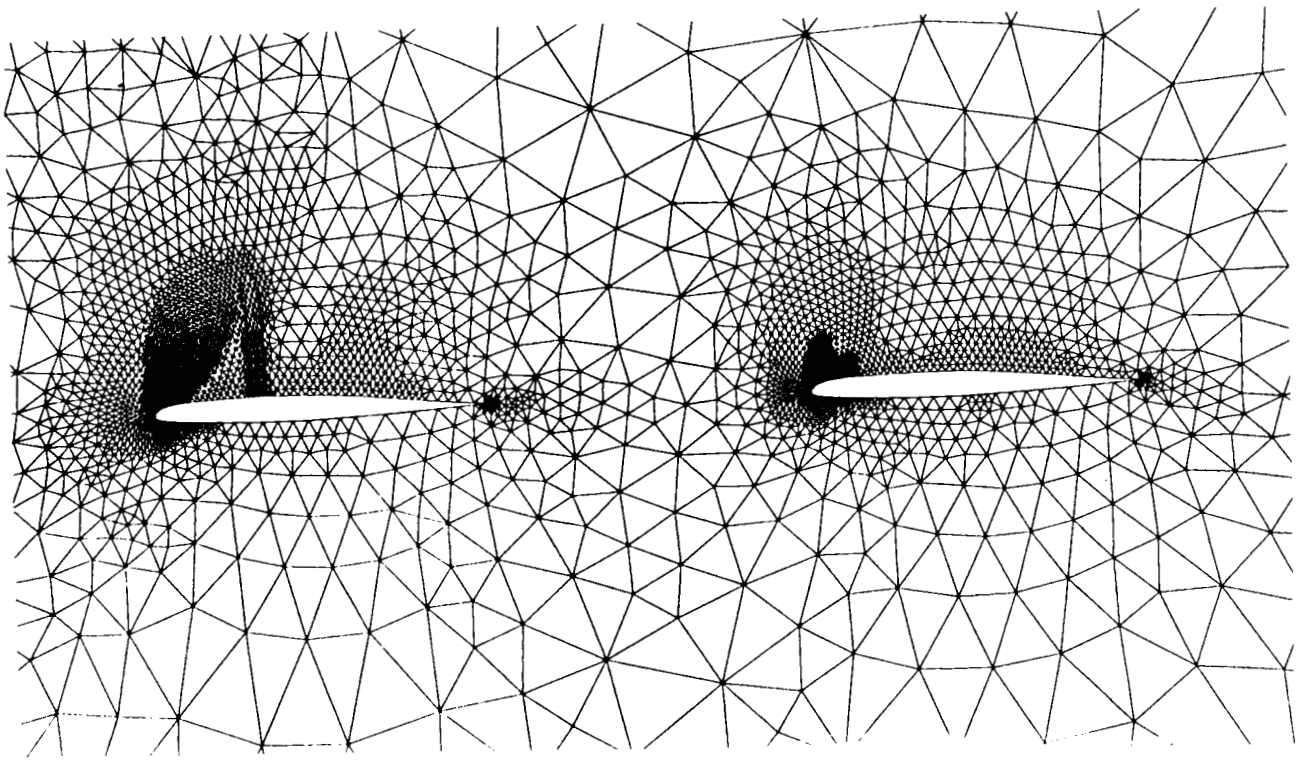

Figure A.3
Sixth Mesh Level Employed in the Regular Adaptive Multigrid Strategy

Mesh Level	Global Multigrid Scheme			Coarsening Scheme	
	Number of Mesh Points	Work units (Saw-Tooth Cycle)	Work Units (W-Cycle)	Number of Mesh Points	Work Units (W-Cycle)
8	5632	1	1	5632	1
7	5038	0.89	1.79	2868	1.02
6	4323	0.77	3.07	1205	0.86
5	2872	0.51	4.08	789	1.12
4	1413	0.25	4.01	712	2.02
3	702	0.12	3.99	702	3.99
2	190	0.03	2.16	190	2.16
1	55	0.01	1.25	55	1.25
Total	20225	3.58	21.35	12153	12.42

Table A.1
Comparison of the Number of Points and Required Amount of Work
on Each Mesh Level for the Global Multigrid Strategy and
the Coarsening Multigrid Strategy Employing a W-Cycle

NASA
National Aeronautics and Space Administration

Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA CR-181848 ICASE Report No. 89-35		

4. Title and Subtitle	5. Report Date
ZONAL MULTIGRID SOLUTION OF COMPRESSIBLE FLOW PROBLEMS ON UNSTRUCTURED AND ADAPTIVE MESHES	May 1989
	6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Dimitri J. Mavriplis	89-35
	10. Work Unit No.
	505-90-21-01

9. Performing Organization Name and Address	11. Contract or Grant No.
Institute for Computer Applications in Science and Engineering Mail Stop 132C, NASA Langley Research Center Hampton, VA 23665-5225	NAS1-18107 NAS1-18605
	13. Type of Report and Period Covered
12. Sponsoring Agency Name and Address	Contractor Report
National Aeronautics and Space Administration Langley Research Center Hampton, VA 23665-5225	14. Sponsoring Agency Code

15. Supplementary Notes

Langley Technical Monitor:
Richard W. Barnwell

SIAM Journal on Scientific and Statistical Computing

Final Report

16. Abstract

This work is concerned with the simultaneous use of adaptive meshing techniques with a multigrid strategy for solving the two-dimensional Euler equations in the context of unstructured meshes. To obtain optimal efficiency, methods capable of computing locally improved solutions without recourse to global recalculations are pursued. A method for locally refining an existing unstructured mesh, without regenerating a new global mesh is employed, and the domain is automatically partitioned into refined and unrefined regions. Two multigrid strategies are developed. In the first, time-stepping is performed on a global fine mesh covering the entire domain, and convergence acceleration is achieved through the use of zonal coarse grid accelerator meshes, which lie under the adaptively refined regions of the global fine mesh. Both schemes are shown to produce similar convergence rates to each other, and also with respect to a previously developed global multigrid algorithm, which performs time-stepping throughout the entire domain, on each mesh level. However, the present schemes exhibit higher computational efficiency due to the smaller number of operations on each level.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
multigrid, unstructured, adaptive, Euler	02 - Aerodynamics 64 - Numerical Analysis Unclassified - Unlimited

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of pages	22. Price
Unclassified	Unclassified	25	A03

NASA FORM 1626 OCT 86