MULTIGRID SOLUTION STRATEGIES FOR ADAPTIVE MESHING PROBLEMS

Dimitri J. Mavriplis[*]

Institute for Computer Applications in Science and Engineering

NASA Langley Research Center

Hampton, VA 23681-0001

ABSTRACT

This paper discusses the issues which arise when combining multigrid strategies with adaptive meshing techniques for solving steady-state problems on unstructured meshes. A basic strategy is described, and demonstrated by solving several inviscid and viscous flow cases. Potential inefficiencies in this basic strategy are exposed, and various alternate approaches are discussed, some of which are demonstrated with an example. Although each particular approach exhibits certain advantages, all methods have particular drawbacks, and the formulation of a completely optimal strategy is considered to be an open problem.

INTRODUCTION

Although most work on adaptive meshing methods has concentrated on the logistics of refining the mesh and the formulation of suitable refinement criteria, efficient solution techniques for the resulting discrete equations are also required in order to enable both fast and accurate solutions. The use of multigrid methods as fast solvers for computational fluid dynamics problems on both structured and unstructured meshes is now well established. Adaptive meshing in particular provides a natural setting for the use of multigrid solvers. The various refined meshes generated from the adaptive process can be used to form the set of coarse and fine meshes of the multigrid sequence. The multigrid algorithm can then be used to accelerate the convergence to steady-state of the discrete equations on the finest adaptive mesh. In fact, the synergy between the two techniques is greater than may be initially apparent, and has roots in the ideas of multi-resolution (see Figure 1). The role of the adaptation process is to identify regions of the domain where the resolution of smaller scales is required and to generate these required new mesh levels, while the role of the multigrid solver is to eliminate the various high and low frequency errors of the solution on the grid level which best represents them. This has led to the development of methods such as the FAC (Full Adaptive Composite) method, [1], and to the notion of the dealgebraization of multigrid, as described by Brandt [2], where the multigrid procedure is no longer viewed as simply a fast solver for discrete equation sets, but rather as part of a complete strategy for approximating the solution to the continuous partial differential equation. Spatial convergence is achieved by the adaptation process, while temporal or numerical convergence is achieved by the multigrid procedure. Additionally, the multigrid defect-correction (*i.e.* coarse grid source term in the multigrid formulation) can be used to devise a refinement criterion.

Although these ideas are appealing, their application to systems of non-linear equations such as those found in computational fluid dynamics is still a relatively unexplored research area. In the present work, various adaptive-meshing multigrid strategies are proposed, and evaluated both in practical terms (*i.e.* speed of convergence, complexity of V or W cycle), and in terms of how well they obey the principles of multi-resolution.

DESCRIPTION OF BASE STRATEGY

The first adaptive-meshing multigrid strategy employed is denoted as the "basic strategy". This method has been found to perform well in practice, and has been used to solve a number of inviscid and viscous steady-state cases. The approach relies exclusively on the use of unstructured meshes which greatly simplifies the task of adaptation.

Single Grid Solver

The Euler (inviscid) or Navier-Stokes (viscous) equations are discretized using a Galerkin finite element approach [3]. In the inviscid case, this reduces to a finite-volume scheme where the flow variables are stored at the vertices of the mesh, and the control volumes are formed by the union of all triangles which touch the considered vertex. This corresponds to a central difference scheme, and additional dissipative terms must be added in order to preserve stability. These are constructed as a blend of an undivided Laplacian and biharmonic operator, with the Laplacian terms used to suppress oscillations near shocks, and the biharmonic terms used to prevent odd-even decoupling in regions of smooth flow. These discrete equations are integrated in time using a five-stage time-stepping scheme devised specifically to damp high frequency error modes (as is required in a multigrid scheme). Integration to steady-state is accelerated by the use of local time-stepping and residual averaging [3,4,5].

Adaptive Meshing Procedure

Adaptively refined meshes are generated by inserting new points into the existing mesh in regions of large gradients, and connecting them to existing mesh points by Delaunay triangulation. The refinement criterion is based on simple undivided differences of one or more flow variables. The difference of the flow variables across each mesh edge is compared to the average difference across all edges of the mesh. When the difference along a given edge is larger than some fraction of the average difference, a new mesh point is added midway along the edge. If one or more edges of a given triangle are flagged for refinement in this manner, then all three edges are refined. This ensures an isotropic refinement strategy, which is necessary to guarantee high quality meshes when using Delaunay triangulation. Once all the new mesh points have been determined, they are inserted into the existing mesh sequentially using Bowyer's algorithm for Delaunay triangulation [6]. Given an initial Delaunay triangulation, this method enables the insertion of a point anywhere in the mesh, and determines the reconnection of this point to the existing points, which is the Delaunay triangulation of this newly augmented point set. As illustrated in Figure 2, Bowyer's algorithm first identifies all triangles whose circumcircle is intersected by the new point. These triangles are then removed creating a polygonal cavity, and the new triangulation is formed by joining the new point to all vertices of the polygonal cavity. New boundary points are repositioned onto the spline curves which define the geometry of the boundaries. After all points have been inserted, the mesh is smoothed and edges are swapped in order to preserve the Delaunay property [4,5,7]. Several passes of smoothing and swapping are usually performed. The use of Bowyer's algorithm in this manner is ideally suited for adaptive meshing problems, since new meshes are constructed through local modifications of an existing mesh, which is much

more efficient than global mesh regeneration. Furthermore, the Delaunay construction of the adaptive meshes prevents the appearance of degenerate connectivities which can arise with simple refinement schemes such as triangle subdivision. Although a reverse Bowyer's algorithm is simple to formulate in two-dimensions, provisions for point removal have not been implemented, since the applications here concern exclusively steady-state problems. For transient problems, point removal capabilities are essential.

Multigrid Approach

There are various possible strategies for implementing a multigrid method with adaptive meshing techniques for unstructured meshes. One approach consists of using the adaptively refined meshes as the multigrid levels themselves [4,5]. If, for example, adaptively refined meshes are created by simply subdividing the appropriate mesh triangles into four finer nested triangles, multiple adaptive refinement passes result in a sequence of fully nested adaptive meshes to which multigrid can be applied in a straight-forward manner using simple restriction and prolongation (inter-grid) operators. This approach has been pursued by several authors in the literature [8,9]. One of the drawbacks of this approach is to restrict the type of adaptive refinement strategies which may be employed, and to tightly couple the multigrid process with the adaptive mesh generation procedure. Furthermore, if the initial unadapted grid is relatively fine (which is most often required to resolve initial flow features), multigrid efficiency will be limited by the ability to efficiently solve the discrete equations of this mesh.

The multigrid approach adopted in this work relies on a sequence of coarse and fine meshes which are essentially independent from one another [3,4,5,11]. The various meshes of the sequence are not required to be nested, or even to have common points. They simply must discretize the same physical domain. Linear interpolation is used to transfer flow variables, residuals and connections between the various meshes of the sequence. The intergrid transfer operators must be formed in a preprocessing operation, where for each vertex of a given grid, the enclosing triangle on the next coarser (or finer) grid must be determined. Once this information has been determined, grid transfer addresses and weights can be determined and stored for later use in the multigrid solution cycles. This multigrid strategy enables the adaptively refined meshes to be constructed by any means available, even global mesh regeneration. The Delaunay construction employed here, and described in the previous section, generally results in non-nested meshes, and meshes with no coincident points (due to the mesh smoothing operation which displaces the mesh points). Furthermore, additional coarser grids may be utilized to accelerate the solution of the initial grid itself. These are generated using the same global mesh generation procedure as the initial mesh, but with

3

lower resolution throughout the domain. The basic procedure consists of generating the initial mesh, and several coarser meshes. The flow solution on the initial mesh is then obtained using this sequence of meshes in the multigrid procedure. A new adaptively refined mesh is then constructed, based on the solution on the initial mesh, and this mesh is then added as a new finer mesh to the current stack of multigrid levels. The restriction and prolongation operators between the new and the initial mesh are then computed and stored. The flow solution is interpolated from the initial mesh to the new finer mesh using these operators, and multigrid cycling resumes, using the newly augmented sequence of meshes. This procedure can be repeated, each time adding a new finer mesh to the sequence, until the desired level of accuracy is obtained, as depicted in Figure 3.

A third multigrid approach for unstructured meshes consists of constructing the sequence of coarse level meshes automatically, given a fine grid. This approach is embodied in algebraic multigrid methods [12], agglomeration strategies [13,14,15], as well as automated coarsening methods used in conjunction with the independent-mesh multigrid approach described above. Thus, in the context of adaptive meshing, each time a new finer mesh is generated, the history of adaptive refinement which resulted in this mesh is ignored, and an automated algorithm is used to generate a complete set of coarse mesh levels based on the new mesh. The philosophy in this approach is to employ multigrid simply as a fast solver for discrete equation sets, in the same manner as an implicit method or direct solver may be used to solve the fine grid equations. Since the history of refinement is not utilized as part of the solution strategy, the multi-resolution concepts discussed previously are not exploited. Such methods have, however, proved to be advantageous, and will be discussed in more detail in the section on Adaptive Multigrid Issues.

RESULTS

The finite-volume method described above, combined with the non-nested multigrid strategy and the Delaunay point-insertion adaptive mesh-refinement technique has been used to solve various inviscid and viscous flow cases. These techniques have been implemented in a single FORTRAN code, which takes as input a sequence of coarse initial meshes, the desired number of adaptive levels, the number of cycles on each level and the refinement criteria for each level, and outputs the sequence of adaptive meshes generated and the solution obtained on the finest mesh.

Inviscid Flow Case 1

The first case consists of the inviscid transonic flow over a NACA 0012 airfoil at Mach number 0.8 and 1.25 degrees incidence. For this case, the outer boundary was approximately

4

circular and placed at a distance of 100 chords from the airfoil. The initial mesh contained 2,112 points, and 5 coarser mesh levels were generated to accelerate the solution on this mesh. The coarsest mesh of this sequence contains only 40 points. Three levels of adaptivity were employed for this calculation. The final mesh is shown in Figure 4, and the solution in terms of Mach contours is depicted in Figure 5. This mesh contains a total of 14,219 points. Mesh refinement is evident in the region of expansion near the leading-edge, and in the vicinity of both the upper and the weak lower shock. The slip line at the trailing edge of the airfoil is however poorly resolved. The undivided gradient of density was used as the refinement criterion. Figures 6 and 7 depict the computed surface pressures and entropy for this case. The shocks are well resolved, and the lift coefficient of 0.3587 is in agreement with previously reported values [16]. Entropy, computed as

$$s = \frac{\frac{P}{P_\infty}}{\frac{\rho}{\rho_\infty}} - 1$$

should be zero for inviscid flow ahead of the shock waves. As can be seen from Figure 7, the computed values near the leading edge are well below 1%, a good indication of the local accuracy of this solution. The convergence rate of the entire adaptive process is shown in Figure 8. At each state of adaptivity, 25 W-multigrid cycles were used to converge the solutions, and 100 W-cycles were used on the final level, in order to demonstrate the asymptotic convergence rate of this method. The residuals were reduced by 6 orders of magnitude in 100 cycles, which corresponds to an average reduction rate of 0.89. This case was performed in about 45 minutes of CPU time on an SGI Indigo R4000 workstation.

Inviscid Flow Case 2

The second case consists of transonic flow over a NACA 0012 airfoil at a freestream Mach number of 0.95 and 0 degrees incidence. For this case, two oblique shock waves and a normal shock wave are set up downstream of the airfoil. The position of the normal shock is very sensitive to the accuracy of the solution. A similar strategy to that discussed for the previous case is employed; *i.e.*, five initial meshes, three levels of adaptivity, undivided difference of density as a refinement criterion. The final mesh and solution are depicted in Figures 9 and 10 respectively. This mesh contains approximately 16,000 points. The normal shock shown in Figure 10 is located 3.06 chord lengths downstream of the trailing edges, which is slightly ahead of that reported elsewhere [17]. In this case, the outer boundary was located 130 chords away from the airfoil leading edge. A previous run on a similar mesh with the outer boundary located at 42 chords yielded a normal shock position of 2.6 chords. This highlights the sensitivity of the solution to the position of the outer boundary. The use of a simple

undivided difference as refinement criterion may also be partly responsible for the inexact shock location in this case.

This case should be perfectly symmetric about the $y = 0$ axis, since the NACA 0012 profile is symmetric, and the flow incidence is zero. An appealing feature of the present adaptation strategy is that, in such cases, given an initial symmetric grid, the adaptively refined grids remain perfectly symmetric, as can be seen from Figure 9. In the final solution, the lift coefficient remained zero, to 6 significant figures.

Inviscid Flow Case 3

The third test case involves the inviscid subsonic flow over the Sudhoo-Hall four element airfoil. The freestream Mach number is 0.2, and the incidence is 0 degrees. Three levels of adaptivity were used for this case, beginning with an initial mesh of 6,466 points. Four coarser meshes were employed to accelerate the convergence on the initial mesh. Thus, a total of 8 mesh levels were used in the final phase of the calculations. The final mesh contained a total of 22,792 points, and is depicted in Figure 11. Figures 12 and 13 depict the computed surface pressures and surface entropy on the finest mesh. As can be seen, the entropy is less than 0.1% over the entire configuration, indicating a good level of local accuracy in the solution. The lift and drag coefficients for this case were 4.9245 and -0.0038 respectively. For inviscid isentropic flows, the overall drag should vanish. Thus the drag value of -38 counts is a good indication of the global accuracy of the solution. Figure 14 depicts the convergence rate for this case, where 100 multigrid W-cycles were performed at each level of adaptivity. The slopes of the various multigrid convergence histories are nearly identical on the four different mesh levels, demonstrating the mesh independent convergence property of the multigrid algorithm. Convergence on the final mesh is only slightly slower than that on the initial levels, resulting in an average reduction rate of 0.925. The convergence history of the computed lift coefficient is also plotted. On each mesh level, the lift coefficient comes very close to its final value in less than 50 cycles. The effect of grid convergence can also be seen by the diminishing differences between the final lift values on consecutively finer meshes. Figure 14 thus illustrates the concept of using adaptive-multigrid as a method of solving for the continuous set of partial differential equations, with the lift coefficient converging to the infinite resolution value, and the multigrid procedure driving the numerical solution on each level. This entire run, including all mesh adaptivity, was achieved in approximately 2 hours on an SGI Indigo R4000 workstation.

Viscous Flow Case

This case consists of viscous turbulent flow over a three-element high-lift airfoil section. The far-field boundary was placed at a distance of 50 chords away from the airfoil (wind-tunnel walls were not modeled in this case). The finite-element discretization of the Navier-Stokes equations described previously was employed, and the single equation turbulence model of Spalant-Allamaras [18] was implemented to account for turbulence effects. The same multigrid strategy described previously was employed to solve both the flow equations and the turbulence equation in a loosely coupled approach. The mesh refinement procedure required some modification for the highly-stretched meshes which are typically used for viscous flows. The Delaunay in-circle criterion described above is used in a mapped space, (resulting in a Delaunay in-ellipse criterion) for both the initial mesh construction, and subsequent adaptive refinement operations [19]. When new boundary points are generated by the refinement procedure, these must be displaced in order to coincide with the surface splines which define the body shape. Whereas in the inviscid case this was easily achieved, in the viscous case, this displacement can require the restructuring of many layers of grid cells near the boundary. This is due to the possibility of the boundary point displacement being much larger than the local normal grid spacing for highly stretched meshes. Thus, a system of pointers is managed, in order to enable local mesh reconstruction near the boundary [19].

For this case, the freestream Mach number is 0.2, the incidence is 16 degrees, and the Reynolds number is 9 million. Three levels of adaptivity were employed. The initial mesh contained approximately 25,000 points, while the final adaptive mesh which is depicted in Figure 15, contains 120,307 points. This mesh exhibits very high resolution in the regions of rapid expansions and in boundary layer and wake regions. A combination of (undivided) pressure and Mach number gradients were employed to identify inviscid and viscous phenomena for refinement. The solution in terms of computed surface pressure, is depicted in Figure 16. This case involved a total of 7 multigrid levels (three adaptive levels, four initial levels). The solution was obtained by running 100 multigrid W-cycles on each mesh, and 300 cycles on the final mesh. The residuals were reduced by 2.5 orders of magnitude on the finest mesh in 300 cycles. This rate is substantially slower than for the inviscid cases, and is primarily due to the stiffness associated with high grid stretching. For the viscous flow cases, the mesh adaptivity operations are run as a separate job with a stand-alone code.

This case has been computed previously on non-adapted meshes of high resolution (up to 240,000 points) and compared extensively with experimental data [20]. Although the solution in Figure 16 appears well resolved, there are certain features, (such as the wake of the slat element for example), which are lost prematurely when compared with the results

of [20], due to inadequate grid resolution. This illustrates the difficulty in applying adaptive meshing to viscous flows, where features such as wakes are both spatially hyperbolic and nonisotropic, and highlights the need for better refinement criteria.

ADAPTIVE MULTIGRID ISSUES

Although the previous examples demonstrate the effectiveness of multigrid as an efficient solution strategy for adaptive meshing problems, certain characteristics of adaptive problems can degrade the overall efficiency of the above multigrid approach. These manifest themselves, not as degradations of the observed convergence rates, but rather as unwanted increases in complexity (number of operations) of the multigrid cycle. For example, in the non-adaptive two dimensional case, the complexity of a V-cycle is bounded by 4/3 work units, and that of a W-cycle by 2 work units, where a work unit is defined as the equivalent work of one fine grid iteration (see Figure 17 for the definition of these cycles). Here, the meshes are not generated adaptively, and the above bounds are computed assuming each coarser mesh level contains 1/4 the number of points of the previous level. In the case of adaptively generated meshes, where such relations between the complexities of the various mesh levels no longer hold, the V-cycle complexity becomes equal to the sum of the complexities of all meshes in the sequence, while the W-cycle complexity can become so high as to make it impractical.

Even the V-cycle complexity is much higher than it need be. For adaptively refined meshes, refinement only occurs in localized regions of the mesh, and there are large regions of the domain where the mesh resolution is essentially unaltered between mesh levels. Repeatedly time-stepping in these regions of the mesh on various levels represents a waste of computational effort. In this section, two strategies which overcome this increase in complexity for V-cycles are described. A third approach which results in optimum complexity, thus enabling the use of V or W cycles, is finally discussed.

The Zonal Fine Grid Scheme

The basic idea behind this scheme [21] is to omit time-stepping in regions of the mesh which have not been refined with regards to the previous level. A crude implementation consists of making use of the same multigrid strategy as described previously, but blanking out the appropriate vertices on each mesh level. In actual fact, the fine mesh consists only of the regions which have been refined, with possibly some extra buffer layers. The method can be implemented by only storing these regions at each level in order to save memory (although this has not been done in this work).

8

As an example, consider the adaptive mesh used to compute the inviscid flow over a tandem airfoil configuration, shown in Figure 18. This mesh is the result of 6 levels of adaptivity. For the zonal fine grid scheme, the 3rd and 4th adaptive levels are depicted in Figure 19. Figure 20 compares the convergence rates of the zonal-fine grid scheme with that of the global multigrid scheme described previously for this case. There are in fact 8 mesh levels in both multigrid cases, 2 initial global levels, and 6 adaptively generated levels. (The global levels are identical for both schemes). The freestream Mach number is 0.7, and the incidence is 3 degrees. The resulting transonic flow solution is qualitatively depicted in Figure 21. Both multigrid schemes converge at nearly identical rates, in terms of residual reduction per cycle. This result verifies the fact that multigrid time-stepping in regions where no change in resolution occurs is unnecessary. The advantage of the zonal fine grid scheme is the result of the reduction in complexity of the multigrid cycle, as shown in Figure 20. For this case, the zonal fine grid scheme is seen to be roughly twice as efficient as the global multigrid approach.

This so-called zonal fine grid scheme developed in [21] is the unstructured mesh equivalent of the fast-adaptive-composite scheme (FAC) [1], and as such embodies the multi-resolution principles outlined in the introduction. Each mesh level is responsible for resolving a particular range of scales, and highly disparate length scales are not found on any common mesh, as is the case in a global mesh with localized regions of adaptive refinement.

One of the drawbacks of this method is that the final solution lies on a composite mesh which is spread over various multigrid levels. Aside from practical difficulties involved in postprocessing the solution, this complicates other issues, such as the requirement of constructing a conservative discretization in the final solution, as well as the use of different schemes on fine and coarse mesh levels.

<center>Zonal Coarse Grid Schemes</center>

The idea of the zonal coarse grid scheme is to overcome the difficulties encountered in the zonal fine grid scheme due to the composite nature of the final solution, by maintaining a global fine grid upon which the final solution is based. In order to maintain favorable complexity, time-stepping is omitted on the coarser meshes in regions of the domain where no mesh refinement takes place between two consecutive levels. This strategy is illustrated in Figure 22, using one-dimensional linear meshes, and compared to the zonal fine-grid and global multigrid strategies. The overall complexity of the zonal fine grid and coarse grid schemes are necessarily equivalent. As can be inferred from the figure, the zonal fine grid and coarse grid schemes are equivalent, except that in the former case the non refined mesh regions are represented on the coarse level meshes, whereas in the latter, these are assigned

<center>9</center>

to the finest possible mesh level. Hence, the zonal coarse grid scheme simply corresponds to a reordering of the local unrefined and refined mesh levels.

The convergence rate of the zonal coarse grid scheme is compared with that of the zonal find grid scheme and the global multigrid scheme for the transonic tandem-airfoil case on the mesh of Figure 18. As expected, all three methods yield similar convergence rates on a per cycle basis, while the zonal fine and coarse grid schemes achieve a factor two gain in efficiency over the global multigrid scheme in this case due to the reduction in complexity, as shown in Figure 20. Thus the zonal coarse grid scheme is equivalent to the zonal fine grid scheme in terms of efficiency, but enables the final solution to be computed on a global fine grid. The disadvantage of this approach is that each time a new adaptively refined mesh is generated, the zonal coarse meshes must be reassigned to the appropriate levels.

<p align="center">Aggressive Coarsening Strategies</p>

While the zonal fine and coarse grid schemes achieve substantial reduction in the complexity of a multigrid cycle for adaptively generated meshes, the use of a W-cycle with such schemes is still unpractical, due to the relative complexities of the various mesh levels. Since the W-cycle performs frequent visits to the coarse level meshes within a single cycle, the mesh complexity must be reduced by at least a factor of four when going to the next coarser level in order to guarantee a bound on the overall W-cycle complexity, as the number of mesh levels increases. Another characteristic of the zonal multigrid schemes described above is that they rely on the adaptive refinement history in order to identify the coarse and fine mesh levels. Such methods cannot be used effectively in the cases where this information is not available, or in the case of a mesh of arbitrary construction.

Automated coarsening strategies can be employed to overcome these difficulties. Given a fine mesh, these methods automatically generate coarser level meshes for use in the multigrid algorithm. Algebraic multigrid [12], and agglomeration multigrid [13,14,15] are examples of automated coarsening strategies. Automated coarsening algorithms have also been devised for use with the fully nested multigrid approaches [10] and the non-nested approach [22]. These methods are attractive because they are fully automated and can be applied to any given grid, regardless of its construction. These methods represent a philosophy in which multigrid is decoupled form the adaptive process, and employed simply as a fast solver for a discrete fine grid problem, much in the same manner as an implicit or direct solver would be employed.

Aggressive coarsening relates to the attempt in an automated coarsening process to optimize the complexity of the generated coarse mesh levels. For a multigrid smoother which

<p align="center">10</p>

is designed to damp high-frequency errors (as is usually the case), the optimal reduction in coarse grid complexity between two successive levels is 4:1 in two dimensions, and 8:1 in three dimensions. Aggressive coarsening strategies can be devised which result in such reductions of mesh complexity, thus resulting in an overall multigrid cycle of near optimal complexity, and enabling the use of V or W-cycles. Although the complexity of the multigrid cycle may be optimal, the overall solution efficiency can only be competitive provided the multigrid convergence rate does not degrade substantially. Figure 23 provides a comparison between the coarse mesh level obtained by two passes of aggressive coarsening on the fine mesh of Figure 18, and the equivalent mesh from the global multigrid sequence (6th level out of 8). Because each cell of the original grid is forced to "grow" at the same rate, the large outer boundary cells are seen to grow much more rapidly throughout the coarsening process than the small refined cells in the shock region of the fine mesh. This results in large discontinuities in cell size which become even more pronounced on the coarser levels. This in turn may degrade the observed convergence rate of a multigrid scheme based on these mesh levels. A similar behavior is observed for agglomeration multigrid methods [15]. Aggressive coarsening strategies are evidently in complete violation of the multi-resolution principle associated with adaptive multigrid methods, where each mesh level is responsible for a given range of scales. Not only does each mesh level contain a wide range of scales in the present approach, but the bandwidth of this range increases on the coarser mesh levels.

Nevertheless, for many problems, aggressive coarsening strategies are highly desirable, both due to their fully automatic nature, and their low complexity. Such methods could obviously be improved by trading off complexity for more regularity in the coarse mesh levels, and thus better multigrid efficiency. However, this task generally requires global information about the current fine mesh construction (*i.e.* in the adaptive mesh case the history of refinement). This has important implications for the future design of automated coarsening techniques, since at present, most of these methods (including algebraic multigrid methods) rely exclusively on local information for constructing coarser levels.

CONCLUSION

Multigrid methods and adaptive meshing techniques have been shown to be complimentary strategies which, when combined in the appropriate manner, can lead to a powerful method which enables rapid convergence, both numerically and spatially, to the continuous partial differential equation. Such methods naturally embody the principle of multi-resolution where each mesh level is responsible for the spatial and numerical resolution of given length scales. In practice, strict adherence to these principles is not always possible

or desirable. Successful methods must achieve a balance between complexity, convergence efficiency, practicality, and ease of implementation.

A non-nested multigrid approach which utilizes each new adaptively refined mesh as an additional multigrid level has been shown to work well in practice for a range of fluid dynamics problems. The simple refinement criterion based on gradients in the flow solution is not sufficiently reliable for application to all types of flows, particularly in the viscous case. Improved refinement criteria and/or better error estimates are sorely needed before adaptive meshing can be routinely used with confidence for complex viscous flows.

References

[1] McCormick, S., and Thomas, J., "The Fast Adaptive Composite Grid (FAC) Method for Elliptic Equations," Mathematics of Computations, Vol. 46, 1986, pp. 439-456.

[2] Brandt, A., "Multigrid Techniques with Applications to Fluid Dynamics: 1984 Guide" Lecture Notes for the Computational Fluid Dynamics Lecture Series, von-Karmen Institute for Fluid Dynamics, Belgium, March 1984.

[3] Mavriplis, D. J., "Turbulent Flow Calculations Using Unstructured and Adaptive Meshes," Int. J. Numer. Methods Fluids, Vol. 13, No. 9, November 1991, pp. 1131-1152.

[4] 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.

[5] Mavriplis, D. J., "Accurate Multigrid Solution of the Euler Equations on Unstructured and Adaptive Meshes, AIAA Journal, Vol. 28, No. 2, February 1990, pp. 213-221.

[6] Bowyer, A., "Computing Dirichlet Tessalations," The Computer Journal, Vol. 24, No. 2, 1981, pp. 162-166.

[7] Lawson, C. L., "Transforming Triangulations," Discrete Mathematics, Vol. 3, pp. 365-372, 1972.

[8] Perez, E., "Finite Element and Multigrid Solution of the Two-Dimensional Euler Equations on a Non-Structured Mesh," INRIA Report No. 442, September 1985.

[9] Connell, S. D., and Holmes, D. G., "Three-Dimensional Unstructured Adaptive Multi-grid Scheme for the Euler Equations," AIAA Journal, Vol. 32, No. 8, pp. 1626-1632, 1994.

[10] Parthasarathy, V., and Kallinderis, Y., "New Multigrid Approach for Three-Dimensional Unstructured Adaptive Grids," AIAA Journal, Vol. 32, No. 5, May 1994.

[11] Mavriplis, D. J., "Multigrid Solution of the Two-Dimensional Euler Equations on Un-structured Triangular Meshes," AIAA Journal, Vol. 26, No. 7, July 1988, pp. 824-831.

[12] Ruge, J. W., and Stuben, K., Algebraic Multigrid, in Multigrid Methods, McCormick, S. F., Editor, SIAM Frontiers in Applied Mathematics, SIAM, Philadelphia, 1987, pp. 73-131.

[13] Lallemand, M. H., Steve, H., and Dervieux, A., "Unstructured Multigridding by Volume Agglomeration: Current Status," Computers and Fluids, Vol. 21, pp. 397-433, 1992.

[14] Smith, W. A., "Multigrid Solution of Transonic Flow on Unstructured Grids," Recent Advances and Applications in Computational Fluid Dynamics, Proceedings of the ASME Winter Annual Meeting, Ed. O. Baysal, November 1990.

[15] Venkatakrishnan, V. and Mavriplis, D., "Agglomeration Multigrid for the 3D Euler Equations," AIAA Paper 94-0069, January 1994.

[16] Pulliam, T. H., and Barton, J. T., "Euler Computations of AGARD Working Group 07 Airfoil Test Cases," AIAA Paper 85-0018, January 1985.

[17] Warren, G., Anderson, W. K., Thomas, J., and Krist, S., "Grid Convergence for Adaptive Methods," AIAA Paper 91-1592-CP, June 1991.

[18] Spalart, P. R., and Allmaras S. R., "A One-Equation Turbulence Model for Aerodynamic Flows," AIAA Paper 92-0439, January 1992.

[19] Mavriplis, D. J., "Unstructured and Adaptive Mesh Generation for High-Reynolds Number Viscous Flows," Proc. of the Third International Conference on Numerical Grid Generation in Computational Fluid Dynamics and Related Fields, Eds. A. S. Arcilla, J. Hauser, P. R. Eiseman, J. F. Thompson, North Holland, June 1991.

[20] Valarezo, W. O., and Mavriplis, D. J., "Navier-Stokes Applications to High-Lift Airfoil Analysis," AIAA Paper 93-3534, AIAA 11th Applied Aerodynamics Conference, Monterey CA, August 1993.

13

[21] Mavriplis, D. J., "Zonal Multigrid Solution of Compressible Flow Problems on Unstructured and Adaptive Meshes," ICASE Report No. 89-35, NASA CR-181848, May 1989.

[22] Guillard, H., "Node Nested Multigrid with Delaunay Coarsening," INRIA Report No. 1898, 1993.