



NIS
IN-02-CR

QUINVED


067663

AIAA-98-0612

**Directional Agglomeration Multigrid
Techniques for High Reynolds
Number Viscous Flow Solvers**

D. J. Mavriplis
Institute for Computer Applications in Science
and Engineering (ICASE)
NASA Langley Research Center
Hampton, VA

**36th Aerospace Sciences
Meeting & Exhibit**
January 12-15, 1998 / Reno, NV



DIRECTIONAL AGGLOMERATION MULTIGRID TECHNIQUES FOR HIGH REYNOLDS NUMBER VISCOUS FLOW SOLVERS

D. J. Mavriplis

Institute for Computer Applications in Science and Engineering
MS 403, NASA Langley Research Center
Hampton, VA 23681-0001

1. Abstract

A preconditioned directional-implicit agglomeration algorithm is developed for solving two- and three-dimensional viscous flows on highly anisotropic unstructured meshes of mixed-element types. The multigrid smoother consists of a pre-conditioned point- or line-implicit solver which operates on lines constructed in the unstructured mesh using a weighted graph algorithm. Directional coarsening or agglomeration is achieved using a similar weighted graph algorithm. A tight coupling of the line construction and directional agglomeration algorithms enables the use of aggressive coarsening ratios in the multigrid algorithm, which in turn reduces the cost of a multigrid cycle. Convergence rates which are independent of the degree of grid stretching are demonstrated in both two and three dimensions. Further improvement of the three-dimensional convergence rates through a GMRES technique is also demonstrated.

2. Introduction

The goal of this work is the development of an efficient solver for compressible steady-state high Reynolds number Navier-Stokes flows on unstructured meshes. The overall strategy is based on a multigrid approach. Multigrid methods form the basis of some of the most efficient available solvers for such problems, both on structured and unstructured grids. For inviscid transonic flow problems, multigrid methods can deliver converged solutions in under 100 cycles [1]. However, for high-Reynolds number Navier-Stokes problems, and for flows involving large regions of low velocity fluid, multigrid convergence rates degrade seriously. This degradation is due partly to the stiffness induced by the highly stretched grids which are required to resolve

efficiently the thin boundary layers and wakes which occur at high Reynolds numbers. Additional stiffness is induced in regions of low Mach number flow, due to the disparity in eigenvalues corresponding to the acoustic and convective wave speeds, as the Mach number tends to zero.

The construction of an efficient solver requires simultaneous treatment of these effects. Semi-coarsening multigrid techniques as well as implicit line-solvers can be used effectively on structured grids to relieve the stiffness associated with highly stretched meshes [2, 3]. The basic semi-coarsening strategy consists of constructing coarser multigrid levels by coarsening the original grid in the coordinate direction normal to the grid stretching, rather than in all directions simultaneously. When conflicting stretching directions exist, multiple coarse grids must be constructed, each generated by a coarsening in a particular coordinate direction [4]. However, when a single stretching direction can be identified, only one family of directionally coarsened grids is required [5].

Semi-coarsening techniques can be generalized to unstructured meshes as directional coarsening methods [6, 7, 8, 9]. Graph algorithms can be constructed to remove mesh vertices based on the local degree and direction of anisotropy in either the grid or the discretized equations. This is achieved by basing point-removal decisions on the values of the discrete stencil coefficients. This is the basis for algebraic multigrid methods [9], which operate on sparse matrices directly, rather than on geometric meshes. These techniques are more general than those available for structured meshes, since they can deal with multiple regions of anisotropies in conflicting directions.

One of the drawbacks of semi- or directional-coarsening techniques is that they result in coarse grids of higher complexity. While a full-coarsening approach reduces grid complexity between successively coarser levels by a factor of 4 in 2D, and 8 in 3D, semi-coarsening techniques only achieve a grid

complexity reduction of 2, in both 2D and 3D. This increases the cost of a multigrid V-cycle, and makes the use of W-cycles impractical. Perhaps more importantly for unstructured mesh calculations, the amount of memory required to store the coarse levels is dramatically increased, particularly in 3D.

An alternative to semi-coarsening is to use an implicit line-solver in the direction normal to the grid stretching coupled with a regular full coarsening multigrid algorithm. Although predetermined grid lines do not exist in an unstructured mesh, such lines can be constructed by identifying and grouping together neighboring mesh edges using a graph algorithm [10, 11]. By using a weighted graph algorithm with edge weights which reflect the degree of coupling in the discretization between neighboring grid points, sets of lines which propagate in the direction of strong grid coupling can be constructed [12].

The solution strategy described in this paper addresses the anisotropy-induced stiffness problem through a combination of implicit line solvers coupled with directional coarsening multigrid. This coupled algorithm permits faster coarsening rates which result in more optimal coarse grid complexities. The low Mach number stiffness problem is addressed using preconditioning techniques [13, 14, 15, 16], which are integrated into the overall directional implicit multigrid algorithm. The combination of these three techniques into a single solver has previously been demonstrated in the context of geometric multigrid for two-dimensional problems [12]. The current work represents an extension of this strategy to the more practical agglomeration or algebraic multigrid approach for unstructured meshes, as well as the extension to three dimensions.

3. Discretization

The governing equations are discretized using a finite-volume approach. Flow variables are stored at the vertices of the mesh, and control volumes are formed by the median-dual graph of the original mesh, as shown in Figure 1. A control-volume flux balance is computed by summing fluxes evaluated along the control volume faces, using the average values of the flow variables on either side of the face in the flux computation. This construction of the convective terms corresponds to a central difference scheme which requires additional dissipation terms for stability. These may either be constructed explicitly as a blend of a Laplacian and biharmonic operator, or may be obtained by writing the resid-

ual of a standard upwind scheme as the sum of a convective term and dissipation term:

$$\sum_{k=1}^{neighbors} \frac{1}{2} (\mathbf{F}(w_i) + \mathbf{F}(w_k)) \cdot \mathbf{n}_{ik} - \frac{1}{2} |\mathbf{A}_{ik}| (w_L - w_R) \quad (1)$$

where the convective fluxes are denoted by $\mathbf{F}(w)$, \mathbf{n}_{ik} represents the normal vector of the control volume face separating the neighboring vertices i and k , and \mathbf{A}_{ik} is the flux Jacobian evaluated in the direction normal to this face. w_L and w_R represent extrapolated flow values at the left and right hand sides of the control volume face respectively. For a first order-scheme, these are taken as the values at the vertices to the left and right of the control volume interface, whereas for a second-order scheme, these are extrapolated from the corresponding vertex values using solution gradients pre-computed at these vertices.




FIG. 1. Median control-volumes for stretched quadrilateral and triangular elements

In this work, a matrix artificial dissipation is employed. The matrix-based artificial dissipation scheme is obtained by utilizing the same transformation matrix $|\mathbf{A}_{ik}|$ as the upwind scheme, but using this to multiply a difference of blended first and second differences (i.e. blended laplacian and biharmonic operator) rather than a difference of reconstructed states at control-volume boundaries. The traditional scalar artificial dissipation scheme [17, 18, 19] is obtained by replacing the four eigenvalues $u, u, u+c, u-c$ in the $|\mathbf{A}_{ik}|$ matrix of the matrix dissipation model by the maximum eigenvalue $|u|+c$, where u and c denote local fluid velocity and speed of sound, respectively. This matrix dissipation construction has been found to deliver accuracy comparable to an upwind scheme, while eliminating the need to compute and store flow gradients at mesh vertices.

The thin-layer form of the Navier-Stokes equations is employed in all cases, and the viscous terms are discretized to second-order accuracy by finite-difference approximation. For multigrid calculations, a first-order discretization is employed for the convective terms on the coarse grid levels.

The single equation turbulence model of Spalart and Allmaras [20] is utilized to account for turbulence effects. This equation is discretized and solved in a manner completely analogous to the flow equations, with the exception that the convective terms are only discretized to first-order accuracy.

This particular discretization is designed to enable the use of mixed element meshes in two dimensions (quadrilaterals and triangles) and three-dimensions (tetrahedra, prisms, pyramids, hexahedra). Meshes of differing element types are handled by employing a single edge-based data structure to assemble the fluxes across all element types [21]. In two dimensions, quadrilateral elements are employed in the regions of high mesh stretching, while triangular elements are employed in isotropic regions of the mesh. In three dimensions, hexahedra or prisms are employed in regions near the wall, while tetrahedra are generally employed elsewhere. The use of different element types in regions of high mesh stretching enables a more complete decoupling of the discretization in the stretching and normal directions, as discussed in section 5.

4. Preconditioned Smoothing

Once the governing equations are discretized, they must be integrated in time to obtain the steady-state solution. This is achieved using a preconditioned multi-stage time-stepping scheme. An explicit k-stage scheme can be written as:

$$\begin{aligned}
 \mathbf{w}_i^{(0)} &= \mathbf{w}_i^{(n)} \\
 &\vdots \\
 \mathbf{w}_i^{(q)} &= \mathbf{w}_i^{(0)} + \Delta t_i \times \mathbf{R}_i(\mathbf{w}^{(q-1)}) \\
 \mathbf{w}_i^{(q+1)} &= \mathbf{w}_i^{(0)} + \Delta t_i \times \mathbf{R}_i(\mathbf{w}^{(q)}) \\
 &\vdots \\
 \mathbf{w}_i^{(n+1)} &= \mathbf{w}_i^{(q=k)}
 \end{aligned} \tag{2}$$

where Δt represents the scalar time step estimate. While such a scheme is commonly used for scalar artificial dissipation discretizations, for upwind or matrix dissipation discretizations substantial increases

in efficiency can be obtained by using a Jacobi preconditioning approach in conjunction with the multi-stage scheme [22, 23, 24, 25]. The $(q + 1)$ th stage of a Jacobi preconditioned multi-stage scheme can be written as:

$$\mathbf{w}_i^{(q+1)} = \mathbf{w}_i^{(0)} + [\mathbf{D}_i]^{-1} \times \mathbf{R}_i(\mathbf{w}^{(q)}) \tag{3}$$

where the scalar time step Δt from equation (2) is replaced by the matrix time step given by the inverse of the matrix

$$[\mathbf{D}_i] = -\frac{\partial \mathbf{R}_i(\mathbf{w})}{\partial \mathbf{w}_i} = \left(\sum_{k=1}^{neighbors} \frac{1}{2} |\mathbf{A}_{ik}| \right) \tag{4}$$

which is a 5×5 matrix (4×4 in two dimensions) corresponding to the pointwise Jacobian of the residual. Note that for a scalar dissipation scheme, this matrix becomes diagonal, and the scalar time-step estimate is recovered, thus reducing the scheme to the standard explicit multi-stage scheme.

Additional preconditioning of the type described in [13, 14, 15, 16] must be implemented in order to address the stiffness problems induced by regions of low Mach number flows. Traditionally, such preconditioners are described as a matrix multiplying an explicit updating scheme, and a similar matrix-based modification to the dissipation terms, which improves the accuracy at low Mach numbers. Thus, the $(q + 1)$ th stage of the standard multi-stage scheme (c.f. equation (2)) is rewritten as:

$$\begin{aligned}
 \mathbf{w}_i^{(q+1)} &= \mathbf{w}_i^{(0)} + \mathbf{P} \Delta t_i \times \\
 &\left(\sum_{k=1}^{neighbors} \frac{1}{2} (\mathbf{F}(w_i^q) + \mathbf{F}(w_k^q)) \cdot \mathbf{n}_{ik} \right. \\
 &\left. - \frac{1}{2} \mathbf{P}^{-1} |\mathbf{P} \mathbf{A}_{ik}| (w_L^q - w_R^q) \right)
 \end{aligned} \tag{5}$$

In the present work, we wish to implement this type of “preconditioner” in the context of a point-implicit (Jacobi-preconditioned) or line-implicit scheme. Since the low Mach number preconditioning matrix is a point-wise matrix, its implementation for point-implicit schemes is similar as for line-implicit, or any implicit scheme. The approach taken, which was originally described in [26, 21], is to modify the dissipation terms in the discretization, as per equation (5), and then simply take this modification into account in the point-wise linearization that is required for the point-implicit Jacobi scheme. Thus, the $(q + 1)$ th stage of the low Mach number

preconditioned Jacobi multi-stage scheme becomes:

$$\begin{aligned} \mathbf{w}_i^{(q+1)} = & \mathbf{w}_i^{(0)} + \\ & \left[\sum_{k=1}^{neighbors} \frac{1}{2} \mathbf{P}^{-1} |\mathbf{P}\mathbf{A}_{ik}| \right]^{-1} \times \\ & \left(\sum_{k=1}^{neighbors} \frac{1}{2} (\mathbf{F}(w_i^q) + \mathbf{F}(w_k^q)) \cdot \mathbf{n}_{ik} \right. \\ & \left. - \frac{1}{2} \mathbf{P}^{-1} |\mathbf{P}\mathbf{A}_{ik}| (w_L^q - w_R^q) \right) \end{aligned} \quad (6)$$

In regions where the Mach number is relatively large, the low Mach number preconditioning matrix \mathbf{P} becomes the identity matrix, and effect of the preconditioner vanishes. In this case, the above scheme reverts to the Jacobi preconditioned scheme of equations (3). Likewise, for scalar dissipation discretizations (i.e. when $|\mathbf{P}\mathbf{A}_{ik}|$ is approximated as a diagonal matrix), this scheme reverts to the low Mach number preconditioned schemes characterized by equation (5) and described in [13, 14, 16]. The particular form of the preconditioning matrix \mathbf{P} employed is that described in [27]. The implementation described therein is attractive because it can be achieved without any change of variables in the original discretization.

Equation (6) represents the scheme used in isotropic regions of the mesh. In regions of large mesh stretching, this pointwise scheme is replaced by a line implicit scheme, operating on grid lines which are pre-constructed in the grid. The implicit system generated by the set of lines can be viewed as a simplification of the general Jacobian obtained from a linearization of a backwards Euler time discretization, where the Jacobian is that obtained from a first-order discretization, assuming a constant Roe matrix in the linearization. For block-diagonal preconditioning, all off-diagonal block entries are deleted, while in the line-implicit method, the block entries corresponding to the edges which constitute the lines are preserved. The line-implicit solver is introduced into the current solution strategy as an extension of the Jacobi preconditioner. At each stage in the multi-stage scheme, the corrections previously obtained by multiplying the residual vector by the inverted block-diagonal matrix are replaced by corrections obtained by solving the implicit system of block-tridiagonal matrices generated from the set of lines. This implementation has the desirable feature that it reduces exactly to the block-diagonal preconditioned multi-stage scheme when

the line length becomes one (i.e. 1 vertex and zero edges), as is the case in isotropic regions of the mesh.

In summary, the final scheme, which is used as a smoother for multigrid on all levels, results in a point-implicit low-Mach number preconditioned multi-stage scheme in isotropic regions of the mesh, and a line-implicit low-Mach number preconditioned multi-stage scheme in regions of high mesh stretching. A three-stage multistage scheme with stage coefficients optimized for high frequency damping properties [28], and a CFL number of 1.8 is used in all computations.

5. Directional Agglomeration and Line Construction

The stiffness due to grid anisotropy is addressed by a directional agglomeration multigrid strategy coupled with a line-implicit smoother. The combination of these two strategies into a single algorithm has been found to result in a more robust and efficient solution method than the use of either strategy alone [29, 12].

In regions of high grid stretching, standard directional agglomeration (i.e. coarsening) results in the removal of one grid point for every retained coarse grid point. This produces a sequence of coarse grid levels for which the complexity between successive levels decreases by a factor of 2. Isotropic agglomeration, on the other hand, produces a coarse grid complexity reduction of 4:1 in 2D and 8:1 in 3D. The higher complexity of the directionally coarsened levels greatly increases memory overheads, particularly in three dimensions, and makes the use of the multigrid W-cycle impractical, since the operation count of the W-cycle becomes unbounded in such cases as the number of grid levels is increased.

The implicit-line solver achieves superior smoothing of error components along the direction of the implicit lines, as compared to a regular explicit scheme. This in-turn permits the use of an accelerated coarsening schedule by the agglomeration multigrid algorithm. However, since the implicit line-solver is only effective at smoothing error components along the implicit lines, multigrid coarsening must proceed precisely along the direction of these lines. This requires a close coupling between the directional agglomeration algorithm and the line construction algorithm. Both techniques are based on weighted graph algorithms, and must employ the same definition of the graph weights.

Agglomeration multigrid may be viewed as a simplified algebraic multigrid strategy. Coarse level grids are constructed by fusing together or agglomerating neighboring control volumes to form a coarser set of larger but more complex control volumes. In the algebraic interpretation of agglomeration multigrid, the coarse levels are no longer geometric grids, but represent groupings of fine grid equations which are summed together to form the coarse grid equations sets [6, 30]. Therefore, it is important to base the directional agglomeration and line construction graph weights on algebraic quantities such as stencil coefficients, rather than geometric quantities such as edge lengths, which may be ill-defined on the coarse levels. However, a one-to-one correspondence between stencil coefficients and grid edges only exists for scalar equations and is not possible for systems of equations. For this reason, the edge weights for the line-construction algorithm and the directional agglomeration algorithm are taken as the stencil coefficients of a scalar convection equation discretized on the fine grid using the finite-volume approach. On the fine level, these correspond to the area-weighted normals of the control volume faces delimiting two neighboring vertices. On the coarser levels, these are constructed by summing the constituent fine level face normals.

For highly stretched quadrilateral cells, this results in large weights being associated with grid edges normal to the direction of stretching, and small edge weights in the direction parallel to the stretching, as can be inferred from the relative sizes of the control volume faces in Figure 1. However, for stretched triangular cells, the diagonal grid edges result in weights which may be comparable in the two directions. This weaker decoupling of the normal and stretching directions for triangular elements in two dimensions can produce undesirable results in the line and agglomeration algorithms. Therefore, we employ quadrilateral elements in two dimensions in regions of high mesh stretching, and prismatic (or hexahedral) elements in highly stretched regions for three dimensional meshes. An alternate approach would be to employ a different control volume definition, such as a containment-dual based control volume [31], and retain simplicial elements in these regions, although this has not been attempted to date.

The line construction algorithm begins by pre-computing the ratio of maximum to average adjacent edge weight for each vertex. The vertices are then sorted according to this ratio. The first vertex in

this ordered list is then picked as the starting point for a line. The line is built by adding to the original vertex the neighboring vertex which is most strongly connected to the current vertex, provided this vertex does not already belong to a line, and provided the ratio of maximum to minimum edge weights for the current vertex is greater than α , (using $\alpha = 4$ in all cases). The line terminates when no additional vertex can be found. If the originating vertex is not a boundary point, then the procedure must be repeated beginning at the original vertex, and proceeding with the second strongest connection to this point. When the entire line is completed, a new line is initiated by proceeding to the next available vertex in the ordered list. Ordering of the initial vertex list in this manner ensures that lines originate in regions of maximum anisotropy, and terminate in isotropic regions of the mesh. The algorithm results in a set of lines of variable length. In isotropic regions, lines containing only one point are obtained, and the point-implicit or Jacobi pre-conditioned scheme is recovered.

The agglomeration algorithm consists of choosing a seed point (i.e. a control-volume) which initiates a local agglomeration, and then agglomerating the neighboring control volumes to the seed point. The isotropic version of this algorithm [32, 33, 18] constitutes an unweighted graph algorithm. In this version of the algorithm, each time a seed point is chosen, all neighboring points are agglomerated to this point. The directional agglomeration algorithm is based on a weighted graph technique. The edge weights are defined in the same manner as for the line construction algorithm. Once a seed point is chosen, only those neighboring points that are connected to the seed point through an edge of weight greater than $\beta \times \max_{weight}$ are agglomerated, where \max_{weight} denotes the maximum edge weight incident to the seed point. Taking $\beta = 0.5$ reproduces the isotropic agglomeration algorithm in regions where all edge weights are close in size. However, in regions where one edge weight is much larger than the others, a directional coarsening is achieved. This results in a 2:1 coarsening ratio in such regions. In order to obtain a 4:1 coarsening ratio, the process must be repeated. This will result in the agglomeration of points or control volumes which were not originally neighbors of the initial seed point. This type of aggressive coarsening can only be tolerated in regions where the implicit line solver is used as a smoother. Therefore, the coarsening process is repeated only if the agglomerated control volume is joined to the

current seed point by an edge which is part of an implicit line. The process is repeated until four control volumes are agglomerated together, or until no line edges can be found.

From the above description, it is evident that the line construction and coarsening process are closely coupled and must be carried out simultaneously. The edge weights, once defined on the finest level, are computed on the fly for each coarser level as they are created. The whole process is performed in a preprocessing phase, and the output, consisting of sets of lines for each level and coarse grid groupings, is passed to the flow solver.




FIG. 2. *Unstructured Grid Used for Computation of Transonic Flow Over RAE 2822 Airfoil. Number of Points = 16167, Wall Resolution = 10^{-6} chords*

As an example, the directional implicit agglomeration multigrid algorithm has been applied to the grid of Figure 2. The lines created on the finest grid level are depicted in Figure 3. The first coarse agglomerated level is illustrated in Figure 4, depicting the agglomerated cells in the boundary-layer region near the leading edge, where a 4:1 directional coarsening is observed. Table 1 documents the complexity of the coarse grid levels using the isotropic agglomeration algorithm of [18], as well as the coarse grid complexity achieved using the current directional agglomeration multigrid algorithm. The resulting complexity for a multigrid W-cycle is just 15 % larger for the directionally agglomerated grids than for the isotropically agglomerated grids.




FIG. 3. *Directional Implicit Lines Constructed on Grid of Figure 2 by Weighted Graph Algorithm*




FIG. 4. *First Agglomerated Multigrid Level Constructed on Grid of Figure 2 Illustrating 4:1 Directional Coarsening in Boundary Layer Region*

Mesh Level	Regular AMG		Directional AMG	
	Nnode	Ratio	Nnode	Ratio
1	16167	1	16167	1
2	4074	3.99	4476	3.61
3	1038	3.92	1383	3.24
4	268	3.87	585	2.36
W-Cycle Complexity	1.89		2.18	

Table 1: Comparison of Coarse Grid Complexity and Resulting W-cycle Complexity for Regular Isotropic Agglomeration and Directional Agglomeration Multigrid

6. Two Dimensional Results

The combined directional-implicit agglomeration multigrid algorithm produces convergence rates independent of the degree of grid anisotropy. This is demonstrated in two dimensions by solving the transonic flow over an RAE 2822 airfoil on three different grids. All three grids contain the same distribution of boundary points, but different resolutions in the direction normal to the boundary and wake regions. The first grid contains a normal wall spacing of 10^{-5} chords, and a total of 12,568 points, while the second grid contains a normal wall spacing of 10^{-6} chords, and 16,167 points, and the third grid a normal wall spacing of 10^{-7} chords, and 19,784 points. The cells in the boundary layer and wake regions are generated using a geometric progression of 1.2 for all three grids. The second grid, depicted in Figure 2, contains what is generally regarded as suitable normal and streamwise resolution for accurate computation of this type of problem, while the first and third grids are most likely under-resolved and over-resolved in the direction normal to the boundary layer, respectively.




FIG. 5. Computed Mach Contours on Grid of Figure 2. Mach = 0.73, Incidence = 2.31 degrees, Re = 6.5 million




FIG. 6. Comparison of Explicit Isotropic and Directional-Implicit Agglomeration Multigrid Algorithm Convergence Rates on 3 Grids of Varying Normal Resolution




FIG. 7. Comparison of Low-Mach Number Preconditioned and Unpreconditioned Directional-Implicit Agglomeration Multigrid Algorithm Convergence Rates for Various freestream Mach Numbers

The Mach number for this case is 0.73, the incidence is 2.31 degrees, and the Reynolds number is 6.5 million. The computed solution on the grid with normal wall spacing of 10^{-6} chords is depicted in Figure 5. The flow is transonic and the low Mach number preconditioning matrix reverts to the identity matrix for this case. With the effect of this preconditioning removed, a more direct comparison between the directional implicit multigrid and the previously developed unpreconditioned full coarsening multigrid method [18, 19] is possible. The convergence rates of both methods on all three grids are shown in Figure 6. The explicit full coarsening multigrid solver produces convergence rates which decay substantially as the grid stretching is increased. In fact, the asymptotic rate of this scheme for the most highly stretched grid is almost two orders of magnitude slower than that achieved on the least stretched grid. On the other hand, the directional-implicit agglomeration scheme produces convergence to machine zero in under 600 cycles and is essentially unaffected by the degree of grid anisotropy. This comparison represents the best possible performance for each scheme. The explicit full-coarsening multigrid algorithm employs a five stage time-stepping scheme which is augmented with implicit residual smoothing and is used to solve a scalar dissipation dis-

cretization. The use of more accurate matrix dissipation with the explicit full-coarsening multigrid algorithm produces slower and less robust convergence rates. The directional implicit agglomeration algorithm operates on the matrix dissipation discretization and uses a three stage time-stepping scheme with no residual smoothing but with point- or line-preconditioning where the jacobians are evaluated and inverted only at the first stage of the scheme and then frozen for the remaining stages. Although Figure 6 compares the two schemes in terms of multigrid cycles, the cost per cycle of both schemes is relatively close, the directional implicit agglomeration scheme being about 15% more expensive per cycle, which is mainly due to the added work for the evaluation of the matrix dissipation.

The benefits of low Mach-number preconditioning are demonstrated in Figure 7, where the flow over an RAE 2822 airfoil at varying Mach numbers has been computed on the grid of Figure 2 using the directional implicit agglomeration algorithm with and without the low Mach number preconditioner. For the transonic case, the preconditioner is not active, and both cases give identical convergence. However, as the Mach number is lowered, the convergence rate degrades substantially for the cases run with no preconditioning, while the preconditioned cases all converge to machine zero in approximately 300 cycles. This example demonstrates the importance of employing both techniques simultaneously (low-Mach number preconditioning and directional implicit agglomeration) in order to obtain rapid convergence rates for subsonic Navier-Stokes flows.

The computation of high-lift flows simultaneously involves regions of low velocity fluid and high grid anisotropy, therefore providing a good demonstration of the current algorithm. Figure 8 depicts an unstructured grid about a three-element airfoil high-lift configuration. The grid contains a total of 61,104 points and a normal spacing of 10^{-6} chords at the surface of each airfoil element. The implicit lines generated by the graph algorithm for this case are depicted in Figure 9, and a qualitative view of the solution as a set of Mach contours is given in Figure 10. The freestream Mach number for this case is 0.2, the incidence is 16 degrees, and the Reynolds number is 9 million. The convergence rates of the directional-implicit agglomeration scheme and the explicit full-coarsening agglomeration scheme are compared on the basis of CPU time in Figure 11. The explicit full-coarsening scheme employs a five

stage time-stepping scheme and residual smoothing and solves the scalar dissipation discretization, while the directional-implicit agglomeration scheme employs the preconditioned three-stage time-stepping scheme with Jacobians frozen at the second and third stages, and solves the matrix dissipation discretization. As can be inferred from Figure 11, the directional-implicit agglomeration scheme achieves a 6 order of magnitude residual reduction in approximately one quarter of the time required by the explicit full-coarsening approach, which permits engineering solutions to be obtained in approximately 1.5 hours on an inexpensive workstation.




FIG. 8. Unstructured Grid Used for Computation of Subsonic Flow Over Three-Element Airfoil Geometry. Number of Points = 61,104, Wall Resolution = 10^{-6} chords




FIG. 9. Implicit Lines Generated by Weighted Graph Algorithm on Grid of Figure 8




FIG. 10. Computed Mach Contours for Flow over Three-Element Airfoil. Mach = 0.2, Incidence = 16 degrees, Reynolds number = 9 million




FIG. 11. Comparison of Convergence Rates obtained for Flow Over Three-Element Airfoil in terms of CPU time on Workstation

7. Three Dimensional Results

In three dimensions, a directional coarsening ratio of 8:1 is required in order to match the coarse grid complexities achieved by an isotropic full coarsening algorithm. Unfortunately, robustness problems associated with 8:1 coarsening ratios have been encountered. Therefore, at present a 4:1 coarsening ratio is employed in three-dimensions, although faster coarsening ratios are still under investigation. This

results in a 50% increase in storage and cpu time per multigrid W-cycle as compared to that achieved by an 8:1 coarsening algorithm, but nevertheless results in an efficient solution procedure for three dimensional problems.




FIG. 12. *Illustration of Mixed Element Grid and Implicit Lines Used for Computation of two-dimensional Flow over three-dimensional Wing Geometry. Number of Grid Points: 177,837*




FIG. 13. *Comparison of Convergence Rate obtained by Three-Dimensional Directional-Implicit Multigrid Algorithm with Rate Obtained on Equivalent Two-Dimensional Problem*

The first three-dimensional test case involves the computation of a two-dimensional flow using a three-dimensional grid, in order to compare the performance of the three-dimensional code with that of the two-dimensional code. The grid of Figure 2 has been extruded in the spanwise direction, resulting in a three dimensional grid of 177,837 vertices. While the original two dimensional grid contained quadrilateral elements in the boundary and wake regions and triangular elements elsewhere, the three dimensional grid contains hexahedral elements in the viscous regions, and prismatic elements in the inviscid regions. The surface grid and the implicit lines generated by the three-dimensional graph algorithm are depicted in Figure 12. The prescribed coarsening ratio of 4:1 results in coarse levels which are very similar to those produced by the two-dimensional algorithm, at least near the wing surface. The convergence rates of the two- and three-dimensional codes are compared in Figure 13. The Mach number is 0.1, the incidence 2.31 degrees, and the Reynolds number is 6.5 million. The three-dimensional convergence curve is much faster than the convergence of the isotropic algorithm, reaching machine zero in just under 600 multigrid cycles. However, it is somewhat slower than the equivalent two-dimensional algorithm which reaches machine zero in just 300 iterations.

The next example demonstrates the insensitivity of the current three-dimensional algorithm to grid stretching. Three unstructured tetrahedral grids have been constructed about an ONERA M6 wing using the VGRID grid generation package [34]. These grids all contain the same surface resolution, but different normal resolutions near the wing surface. The first grid contains a normal wall spacing of 10^{-5} chords, and a total of 1.2 million points, while the second grid contains a normal wall spacing of 10^{-6} chords, and 1.6 million points, and the third grid a normal wall spacing of 10^{-7} chords, and 2 million points. The cells in the boundary layer and wake regions are generated using a geometric progression of 1.2 for all three grids. The second grid (i.e. 10^{-6} spacing) is depicted in Figure 14. As explained previously, prismatic elements are required in the boundary layer regions for the directional-implicit agglomeration algorithm. Since the VGRID grid generation package produces fully tetrahedral meshes, a mesh merging technique is employed to merge tetrahedra triplets into prisms in regions near the wall [21]. At this initial stage of development, a uniform height of prism layers is em-

ployed over the entire wing surface. This avoids the use of “hanging edges” or transitional elements such as pyramids when a variable number of prismatic layers are allowed in the grid. On the other hand, this restriction may result in the incomplete merging of some stretched tetrahedral elements. The convergence rates of the directional-implicit agglomeration algorithm on all three grids are depicted in Figure 15. A four level W-cycle was used for these computations. The coarsening ratios achieved between the first and second, second and third, third and fourth levels were 3.69:1, 3.16:1 and 2.2:1 respectively for the 10^{-6} spacing grid, with similar coarsening ratios obtained on the other grids. The Mach number is 0.1, the incidence is 2.0 degrees and the Reynolds number is 3 million.

Very similar convergence rates are achieved on all three grids. All cases exhibit a slowdown in convergence after approximately 6 or 7 orders of magnitude, but achieve close to an 8 order of magnitude reduction in 600 cycles. Considering that these three grids represent a two order of magnitude variation in the degree of grid stretching, the convergence rates can be qualified as independent of the grid stretching. As an example of the computational overheads incurred, the grid containing 1.6 million vertices required a total of 350 Mwords of memory and 53 cpu hours to converge 600 cycles. This case was run on 8 CPUS of the CRAY C90 and achieved a cpu to wall-clock time ratio of 7 on a lightly loaded machine.




FIG. 14. *Illustration of One of Three Unstructured Grids Employed For Computation of Flow Over ONERA M6 Wing: Number of vertices= 1.6 million, Wall Spacing = 10^{-6}*




FIG. 15. *Comparison of Convergence Rates Achieved by Directional Implicit Agglomeration Multigrid Algorithm on 3 Grids of Varying Normal Resolution for ONERA M6 Wing Geometry at Mach number = 0.1*

The final case consists of a three-dimensional high-lift application. The geometry involves a partial span flap unswept wing in a wind-tunnel. The unstructured grid employed for this case is depicted in Figure 16. This mesh contains a total of 549,176 points, and a normal spacing at the wing surface of 10^{-5} chords. As in the previous case, a uniform height of prismatic layers was created in the boundary layer regions using the mesh merging algorithm of [21]. The Mach number for this case is 0.2, the incidence is 10 degrees, and the Reynolds number is 3.7 million. The computed density contours for this case are depicted in Figure 17. A four level W-cycle was used for this computation. The coarsening ratios achieved between the first and second, second and third, third and fourth levels were 3.84:1, 3.43:1 and 2.23:1 respectively. The convergence obtained by the directional implicit agglomeration multigrid algorithm is compared with that achieved by the explicit full-coarsening agglomeration multigrid algorithm in Figure 18. As in the two-dimensional comparisons, this represents the best possible performance for each algorithm: the directional algorithm employs a three-stage time-stepping scheme and operates on the matrix dissipation discretization, while the isotropic algorithm employs a five-stage scheme with residual smoothing and operates

on the scalar dissipation discretization. The directional algorithm produces substantially faster convergence than the isotropic algorithm, although a slowdown is observed after 4 to 5 orders of magnitude. While the asymptotic rate of the directional algorithm is still substantially faster than that of the isotropic algorithm, the rate is much slower than that achieved in two dimensions.




FIG. 16. *Illustration of Unstructured Grid employed for Computation of Flow Over Partial-Span Flap Geometry. Number of Vertices = 549,176*




FIG. 17. *Illustration of Computed Density Contours for Flow Over Partial-Span Flap Geometry. Mach Number = 0.2, Incidence = 10 degrees, Reynolds Number = 3.2 million*

Further increases in the convergence rate can be achieved by resorting to a Krylov acceleration technique such as GMRES [35]. The preconditioned directional implicit agglomeration algorithm can be

employed as a preconditioner to GMRES [24, 12]. The current implementation uses a nonlinear GMRES solver [36] which computes Jacobian-vector products by finite differencing the residual. The addition of GMRES incurs little extra cpu time, measured on a multigrid cycle basis, but requires considerable additional storage, since a solution vector must be stored for each of the Krylov search directions. In the current implementation, 20 search directions are employed, resulting in a memory increase of 100 words per vertex (about 50% increase). The convergence rate using GMRES is depicted in Figure 18. The solver was run initially 150 multigrid cycles using the directional agglomeration multigrid algorithm alone, and then another 462 multigrid cycles using the preconditioned GMRES approach (i.e. 22 GMRES(20) cycles). The addition of GMRES is largely successful in overcoming the slowdown in the asymptotic convergence rate observed by the simpler directional implicit multigrid algorithm alone, achieving an overall residual reduction of 9 orders of magnitude over 600 cycles. This case required a total of 230 Mwords of memory and 20 cpu hours on the CRAY C90, and ran at a cpu-time to wall-clock ratio of 7 on 8 processors.




FIG. 18. *Comparison of Convergence Rates Achieved by Various Multigrid Schemes for Flow Over Partial Span Flap Geometry*

8. Future Work

The preconditioned directional agglomeration multigrid algorithm has been shown to produce grid-aspect-ratio independent convergence rates in both two and three dimensions. In particular, the two-dimensional implementation of the current algorithm has resulted in a very efficient solver for viscous flows. However the convergence rates obtained in three dimensions, while substantially faster than those achieved by the isotropic algorithm, are still slower than those observed in two dimensions. This may be due in part to the possibility of the existence of multidimensional stretching in three dimensions. In such cases, a modified line-construction and coarsening strategy may be required. Furthermore, the coarsening ratios achieved in three dimensions, which are often even lower than the prescribed 4:1 rate, result in additional cpu time per multigrid cycle and increase the overall memory requirements of the solver. Future work will focus on augmenting the three-dimensional coarsening ratios to approximate the 8:1 ratio observed in isotropic cases as closely as possible, and on accelerating the three-dimensional convergence rates through improved line construction and preconditioning. In order to isolate the effects of the turbulence model, the turbulence values have been frozen after 150 to 200 cycles in the in the computations of the preceding section. The efficient convergence of the combined system of flow and turbulence equations is also under investigation.

9. Acknowledgments

The author wishes to thank S. Pirzadeh for his time spent in generating the ONERA M6 wing and partial span flap unstructured meshes. This work was made possible in large part due to the computational resources provided by the NAS facility.

REFERENCES

- [1] A. Jameson. Artificial diffusion, upwind biasing, limiters and their effect on accuracy and multigrid convergence in transonic and hypersonic flow. AIAA paper 93-3359, Proceedings of the AIAA 11th Computational Fluid Dynamics Conference, Orlando, Florida, July 1993.
- [2] A. Brandt. Multigrid techniques with applications to fluid dynamics:1984 guide. In *VKI Lecture Series*, pages 1-176, March 1984.
- [3] S. R. Allmaras. Analysis of semi-implicit preconditioners for multigrid solution of the 2d compressible navier-stokes equations. In *Proceedings of the 12th AIAA CFD Conference, San Diego CA*, June 1995. AIAA Paper 95-1651-CP.
- [4] W. A. Mulder. A high resolution Euler solver based on multigrid semi-coarsening and defect correction. *Journal of Computational Physics*, 100:91-104, 1992.
- [5] N. Pierce, M. Giles, A. Jameson, and L. Martinelli. Accelerating three-dimensional navier-stokes calculations. In *Proceedings of the 13th AIAA CFD Conference, Snoemass Village, CO*, June 1997. AIAA paper 97-1953.
- [6] D. J. Mavriplis. Multigrid techniques for unstructured meshes. In *VKI Lecture Series VKI-LS 1995-02*, March 1995.
- [7] M. Raw. Robustness of coupled algebraic multigrid for the Navier-Stokes equations. AIAA paper 96-0297, January 1996.
- [8] J. Francescatto. Résolution de l'équation de Poisson sur des maillages étirés par une methode multigrille. INRIA Report No. 2712, November 1995.
- [9] J. W. Ruge and K. Stüben. Algebraic multigrid. In S. F. McCormick, editor, *Multigrid Methods*, SIAM Frontiers in Applied Mathematics, pages 73-131, Philadelphia, 1987. SIAM.
- [10] O. Hassan, K. Morgan, and J. Peraire. An adaptive implicit/explicit finite element scheme for compressible high speed flows. AIAA Paper 89-0363, January 1989.
- [11] D. Martin and Löhner. An implicit linelet-based solver for incompressible flows. AIAA Paper 92-0668, January 1992.
- [12] D. J. Mavriplis. Multigrid strategies for viscous flow solvers on anisotropic unstructured meshes. In *Proceedings of the 13th AIAA CFD Conference, Snowmass, CO*, pages 659-675, June 1997. AIAA Paper 97-1952-CP.
- [13] E. Turkel. Preconditioning methods for solving the incompressible and low speed compressible equations. *J. Comp. Phys.*, 72:277-298, 1987.
- [14] B. van Leer W. T. Lee and P. Roe. Characteristic time-stepping or local preconditioning of the Euler equations. In *Proceedings of the 10th AIAA CFD Conference, Hawaii*, pages 260-282, June 1991. AIAA Paper 91-1552-CP.

- [15] Y. Choi and C. Merkle. The application of preconditioning to viscous flows. *J. Comp. Phys.*, 105, 1993.
- [16] J. M. Weiss and W. A. Smith. Preconditioning applied to variable and constant density time-accurate flows on unstructured meshes. AIAA Paper 94-2209, June 1994.
- [17] A. Jameson, W. Schmidt, and E. Turkel. Numerical solution of the Euler equations by finite volume methods using Runge-Kutta time stepping schemes. AIAA Paper 81-1259, 1981.
- [18] D. J. Mavriplis and V. Venkatakrishnan. Agglomeration multigrid for two dimensional viscous flows. *Computers and Fluids*, 24(5):553-570, May 1995.
- [19] D. J. Mavriplis and V. Venkatakrishnan. A 3D agglomeration multigrid solver for the Reynolds-averaged Navier-Stokes equations on unstructured meshes. *International Journal for Numerical Methods in Fluids*, 23(6):527-544, 1996.
- [20] P. R. Spalart and S. R. Allmaras. A one-equation turbulence model for aerodynamic flows. *La Recherche Aéropatiale*, 1:5-21, 1994.
- [21] D. J. Mavriplis and V. Venkatakrishnan. A unified multigrid solver for the Navier-Stokes equations on mixed element meshes. AIAA Paper 95-1666, June 1995.
- [22] K. Riemsdagh and E. Dick. A multigrid method for steady Euler equations on unstructured adaptive grids. In *6th Copper Mountain Conf. on Multigrid Methods, NASA conference publication 3224*, pages 527-542, 1993.
- [23] E. Morano and A. Dervieux. Looking for $O(N)$ Navier-Stokes solutions on non-structured meshes. In *6th Copper Mountain Conf. on Multigrid Methods*, pages 449-464, 1993. NASA Conference Publication 3224.
- [24] C. Ollivier-Gooch. Towards problem-independent multigrid convergence rates for unstructured mesh methods I: Inviscid and laminar flows. In *Proceedings of the 6th International Symposium on CFD, Lake Tahoe, NV*, September 1995.
- [25] N. Pierce and M. Giles. Preconditioning on stretched meshes. AIAA paper 96-0889, January 1996.
- [26] E. Turkel. Preconditioning-squared methods for multidimensional aerodynamics. In *Proceedings of the 13th AIAA CFD Conference, Snowmass, CO*, pages 856-866, June 1997. AIAA Paper 97-2025-CP.
- [27] D. Jespersen T. Pulliam and P. Buning. Recent enhancements to OVERFLOW. AIAA Paper 97-0644, January 1997.
- [28] B. van Leer, C. H. Tai, and K. G. Powell. Design of optimally-smoothing multi-stage schemes for the Euler equations. AIAA Paper 89-1933, June 1989.
- [29] D. J. Mavriplis. Directional coarsening and smoothing multigrid strategies for anisotropic Navier-Stokes problems. In *8th Copper Mountain Conf. on Multigrid Methods*, 1997.
- [30] B. R. Hutchinson and G. D. Raithby. A multigrid method based on the additive correction strategy. *Numerical Heat Transfer*, 9:511-537, 1986.
- [31] T. J. Barth and S. W. Linton. An unstructured mesh newton solver for compressible fluid flow and its parallel implementation. AIAA paper 95-0221, January 1995.
- [32] M. Lallemand, H. Steve, and A. Dervieux. Unstructured multigridding by volume agglomeration: Current status. *Computers and Fluids*, 21(3):397-433, 1992.
- [33] W. A. Smith. Multigrid solution of transonic flow on unstructured grids. In *Recent Advances and Applications in Computational Fluid Dynamics*, November 1990. Proceedings of the ASME Winter Annual Meeting, Ed. O. Baysal.
- [34] S. Pirzadeh. Viscous unstructured three-dimensional grids by the Advancing-Layers Method. AIAA Paper 94-0417, January 1994.
- [35] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Sci. Stat. Comput.*, 7(3):856-869, 1986.
- [36] L. B. Wigton N. J. Yu and D. P. Young. GMRES acceleration of computational fluid dynamic codes. In *Proceedings of the 7th AIAA CFD Conference*, pages 67-74, July 1985. AIAA Paper 85-1494-CP.