INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS, VOL. 11, 1075–1095 (1990)

GLOBAL CONVERGENCE OF INEXACT NEWTON METHODS FOR TRANSONIC FLOW

DAVID P. YOUNG, ROBIN G. MELVIN AND MICHAEL B. BIETERMAN

Boeing Computer Services, M/S: 7L-21, PO Box 24346, Seattle, WA 98124-0346, U.S.A.

AND

FORRESTER T. JOHNSON AND SATISH S. SAMANT

Boeing Commercial Airplanes, M/S: 7C-36, PO Box 3707, Seattle, WA 98124-2207, U.S.A.

SUMMARY

In computational fluid dynamics, non-linear differential equations are essential to represent important effects such as shock waves in transonic flow. Discretized versions of these non-linear equations are solved using iterative methods. In this paper an inexact Newton method using the GMRES algorithm of Saad and Schultz is examined in the context of the full potential equation of aerodynamics. In this setting, reliable and efficient convergence of Newton methods is difficult to achieve. A poor initial solution guess often leads to divergence or very slow convergence. This paper examines several possible solutions to these problems, including a standard local damping strategy for Newton's method and two continuation methods, one of which utilizes interpolation from a coarse grid solution to obtain the initial guess on a finer grid. It is shown that the continuation methods can be used to augment the local damping strategy to achieve convergence for difficult transonic flow problems. These include simple wings with shock waves as well as problems involving engine power effects. These latter cases are modelled using the assumption that each exhaust plume is isentropic but has a different total pressure and/or temperature than the freestream.

KEY WORDS Inexact Newton methods Global convergence Finite elements Full potential equation Damping strategies

1. INTRODUCTION

Much recent computational work in aerodynamics has dealt with the steady state (time-independent) full potential, Euler or Navier–Stokes equations. The resulting boundary value problems are highly non-linear. The discretized versions of these problems and are often solved with time-marching-like algorithms such as ADI, Runge–Kutta methods or multigrid.[1–5] An alternative sometimes used because of its locally quadratic convergence is Newton's method. Shubin *et al.*[6] applied Newton's method to the Euler equations for one-dimensional duct flow. The supersonic blunt body problem is discussed in Reference 7. Continuation methods are suggested for difficult problems using the freestream Mach number, the ratio of specific heats of the gas, or the body shape as parameters. Early applications to viscous incompressible flow problems include those described in References 8 and 9. Childs and Pulliam[10] have applied Newton's method to the Euler equations using the factored implicit algorithm accelerated by multigrid. More recent applications to aerofoil calculations include the work of Drela, Giles and Thompkins.[11–14] They developed a new algorithm for solving the steady state two-dimensional

0271–2091/90/161075–21$10.50

Received 19 October 1989
Revised 24 January 1990

Euler equations that uses Newton's method with direct solution of the linearized systems. A good initial guess is supplied by using an elliptic grid generator, one set of whose grid lines correspond to streamlines for incompressible flow. The solution of each linearized system is accomplished by a sparse direct block tridiagonal matrix solver. This algorithm has proven very reliable and accurate for aerofoil calculations. Wigton[15] has extended this approach to multi-element aerofoils and demonstrated the use of the symbolic manipulation package MACSYMA to set up the Jacobian matrices required and a sparse matrix solver (a relative of the one used in TRANAIR described below) with nested dissection ordering of the unknowns. Continuation using the artificial viscosity parameter (viscosity damping) was also used to achieve more reliable global convergence. Venkatakrishnan[16] applied Newton's method to the Euler equations discretized on structured grids in two dimensions. Techniques such as grid sequencing (obtaining initial guesses for Newton's method by solution interpolation onto successively finer grids) were used to improve the reliability of convergence. Sparse matrix solvers were used to solve the linearized problems. Venkatakrishnan and Barth[17] applied Newton's method to the Navier–Stokes equations discretized on unstructured grids in two dimensions using similar methods.

Newton's method can be shown to be globally convergent only under very special assumptions. Many local damping strategies have been devised to increase the range of problems for which Newton's method converges.[18] These techniques have usually been developed in the context of numerical optimization and include Levenberg–Marquardt strategies, backtracking and dog-leg strategies. Several of these are examined in Reference 19.

In this paper we will compare various damping strategies used with an inexact Newton method[20] for transonic flow problems. In this situation, local damping strategies were found to be of limited value and strategies based on the nature of the boundary value problem had to be introduced to achieve reliable convergence. The strategies for improved global convergence have been implemented in a general three-dimensional geometry full potential method developed at Boeing called TRANAIR.[21] In this paper we describe results of applying these strategies with TRANAIR, discuss their relative merits and demonstrate in specific cases why they work.

In two space dimensions it is feasible to use Newton's method with exact generation and factorization of the sparse Jacobian matrices. In three space dimensions, however, typical computational stencils are larger and the total number of grid points (and hence unknowns) is much larger. This makes Newton's method practical only if suitable iterative methods are available for approximately solving the linearized systems. Krylov subspace methods have been used in this context of inexact Newton methods for numerical solution of ordinary differential equations[22–24] and for semiconductor device simulation.[25] Brown and Saad considered damping procedures such as Powell's dog-leg strategy and line search back tracking for general non-linear systems of equations.[26] A general non-linear local relaxation preconditioner was considered and tested by Chan and Jackson.[27] This preconditioner has the advantage that it can be implemented using only residual evaluations and the diagonal entries of the Jacobian matrix.

In the work described here an inexact Newton method is used, with the linear problems being solved approximately with a preconditioned GMRES algorithm.[28] The local damping strategy of Bank and Rose[19] is implemented in this inexact Newton method. General continuation in some parameter for the problem is included. The resulting FORTRAN subroutine is about 150 lines of executable code. In this context, problem-dependent damping strategies (such as those described below) are easy to implement. We have found these strategies to be indispensible in increasing the reliability of convergence for difficult transonic flow problems. Two such strategies will be examined in detail, namely grid sequencing and viscosity damping. Grid sequencing involves finding an initial guess by interpolation from a solution on a coarse grid. Viscosity damping is a continuation process using a parameter measuring artificial viscosity.

2. INEXACT NEWTON METHODS

A damped Newton method can be described as follows. Suppose we wish to solve the non-linear system of equations

$$F(x)=0. \tag{1}$$

Given an initial approximate solution x^0, for $n = 0, 1, 2, \ldots$ until the norm of the residual $F(x^n)$ is sufficiently small, set

$$x^{n+1} = x^n + \lambda(\delta x^{n+1}), \tag{2}$$

where λ is the step length and δx^{n+1} is the solution of the linear system

$$\bar{F}_{x^n}(\delta x^{n+1}) = -F(x^n). \tag{3}$$

Here $\bar{F}_{x^n}$ is the Jacobian for F linearized about x^n. This linear operator can be defined by giving its action on any vector y:

$$\bar{F}_x(y) = \lim_{\varepsilon \to 0} \frac{F(x+\varepsilon y)-F(x)}{\varepsilon}. \tag{4}$$

The step length λ can be selected in a variety of ways. One possibility is to choose λ so that in some appropriate norm, $\| F(x^{n+1}) \| < \| F(x^n) \|$. More sophisticated strategies will be discussed in Section 7. The GMRES algorithm[28] is used in the present method to solve equation (3). This algorithm requires only the ability to calculate the action of the linear operator $\bar{F}_x$ on any vector y. Equation (4) is used to approximate this action:

$$\bar{F}_x(y) \simeq \frac{F(x+\varepsilon y)-F(x)}{\varepsilon}, \tag{5}$$

where ε is small in some appropriate sense. Thus the linear problem, equation (3), is solved without ever explicitly generating the Jacobian for the full non-linear problem.

To control the cost of the present method, system (3) is solved only approximately with GMRES, i.e. δx^{n+1} satisfies

$$\| \bar{F}_{x^n}(\delta x^{n+1}) + F(x^n) \| / \| F(x^n) \| < \eta. \tag{6}$$

This makes the method an inexact Newton method.[20] In general, if η is constant, local linear convergence is guaranteed. If $\eta \to 0$ as convergence takes place, the convergence is superlinear locally.[20]

This algorithm can be viewed as an acceleration procedure for an existing iterative method for solving equation (1). Such methods can be thought of as taking some approximation x^n to the solution and producing a hopefully better approximation

$$x^{n+1} = M(x^n). \tag{7}$$

If the iterative method is consistent, the solution of equation (1) is then a fixed point of M so that we can replace F by $M-I$ in the above discussion. In practical situations, where complicated codes already exists to evaluate M, a simple implementation of the above inexact Newton method is possible to accelerate M without changing the already existing code. This formulation is equivalent to Newton's method for $F(x) = M(x)-x$ with each linearized problem solved using GMRES. Each linearized problem is

$$\bar{M}(\delta x)-I(\delta x) = -F(x^n). \tag{8}$$

$\bar{M}$ could be the iteration matrix for a stationary iterative method, e.g. SSOR, for solving some underlying problem (F already includes preconditioning in this case). In this instance, $I-\bar{M}=C^{-1}\bar{H}$, where C is a linear preconditioner and $\bar{H}$ is the Jacobian for the underlying problem at x^n. In special cases, M could be constructed by applying a stationary iterative method to the underlying problem $H(x)=0$. In simple cases (e.g. if H is symmetric, positive definite and linear[29,30] all the theory for the preconditioned conjugate residuals method is applicable. In this way the existing algorithm and code to evaluate M need not be changed. This approach has been quite successful in accelerating existing computational fluid dynamics codes.[31] Indeed, adequate preconditioning is quite important to the success of this approximate Newton method. This 'black box' approach is attractive because it automatically utilizes any working iterative method as a preconditioning. Also, the method easily takes advantage of large exterior storage devices such as the Cray SSD.

3. BOUNDARY VALUE PROBLEM

The full potential equation of aerodynamics is

$$\mathscr{F}(\Phi)\equiv\vec{\nabla}\cdot\rho\vec{\nabla}\Phi=0, \tag{9}$$

where the density is given by

$$\rho=\rho_\infty\left[1+\frac{\gamma-1}{2}M_\infty^2\left(1-\frac{q^2}{q_\infty^2}\right)\right]^{1/(\gamma-1)}. \tag{10}$$

Here, with $\vec{V}_\infty$ taken to be a uniform onset flow, Φ is the total velocity potential to be determined, $q=\|\vec{\nabla}\Phi\|_2$ is the local speed, $q_\infty=\|\vec{V}_\infty\|_2$ is the freestream speed, ρ_∞ is the freestream density, M_∞ is the freestream Mach number and γ is the ratio of specific heats. Equation (9) describes irrotational compressible flow. In addition, boundary conditions are required to define a well posed problem. For external flow problems the far-field condition is

$$\phi=\mathcal{O}(1/R) \tag{11}$$

as $x\to-\infty$, i.e. upstream of the object. Here the perturbation potential is given by $\phi=\Phi-\Phi_\infty$, where $\vec{\nabla}\Phi_\infty=\vec{V}_\infty$. On impermeable surfaces the normal mass flux condition is $\rho(\partial\Phi/\partial n)=0$. On other surfaces, such as engine inlets, the normal mass flux is a specified function, $\rho(\partial\Phi/\partial n)=g_1$. On other surfaces we impose the Dirichlet condition $\Phi=g_3$. On engine exhaust surfaces, tangential flow can be prohibited by specifying g_3 to be constant. Wake surfaces must extend downstream from lifting components such as wings. These surfaces allow non-zero circulation in potential flow and can be thought of as thin sheets of concentrated vorticity.[32] The boundary conditions on a wake are

$$\hat{n}\cdot\Delta(\rho\vec{\nabla}\Phi)=0, \tag{12}$$

$$\Delta p=0, \tag{13}$$

where the pressure

$$p=p_\infty\left[1+\frac{\gamma-1}{2}M_\infty^2\left(1-\frac{q^2}{q_\infty^2}\right)\right]^{\gamma/(\gamma-1)} \tag{14}$$

and Δ represents the jump across the wake surface. Equation (12) is an expression of conservation of mass across the wake. Equation (13) is required for conservation of normal momentum. Equation (13) is often linearized about the freestream pressure $p=p_\infty$ assuming a small perturba-

tion velocity $\vec{\nabla}\phi$. This leads to the equivalent Dirichlet condition that $\Delta\Phi$ is constant along the wake in the direction of $\vec{V}_\infty$. The circulation μ at the trailing edge is determined by a Kutta condition imposed there.

The Bateman variational principle, namely that the integral of pressure over the flow field is stationary,[33] can be used to derive finite element formulae for the full potential equation. In subsonic flow this integral is maximized. A generalization of the Bateman variational principle which incorporates the boundary conditions described above is that the functional

$$J = \int_\Omega p\,\mathrm{d}V + \int_{\partial\Omega_1} g_1\Phi\,\mathrm{d}S - \int_{\partial\Omega_2} \alpha\left(\rho\frac{\partial\Phi}{\partial n}\right)(\Delta\Phi-\mu)\,\mathrm{d}S + \int_{\partial\Omega_3} \rho\frac{\partial\Phi}{\partial n}(\Phi-g_3)\,\mathrm{d}S \tag{15}$$

is stationary. Here g_1 is the given mass flux data on $\partial\Omega_1$, $\Delta\Phi$ is the jump in Φ across the wake surface $\partial\Omega_2$, μ is the unknown representing the jump in Φ on $\partial\Omega_2$ determined by equation (13), α denotes the average of the upper surface and lower surface values and g_3 is the given Dirichlet data on $\partial\Omega_3$. The function μ is itself unknown and is determined by equation (13). To achieve a stable numerical formulation, the treatment of Dirichlet boundary conditions and wake surface conditions must be modified. In addition, the natural Neumann condition must be modified to account for boundary curvature, since the solution is often sensitive to this quantity and the boundary is discretized using flat panels. These modifications are described in detail in Reference 21.

A modification of the above formulation allows the simulation of flows involving regions of differing total temperature and pressure. The flow in each separate region is still potential as long as total temperature and pressure are constant in the region, but pressure and density must be redefined in the following way:

$$p = p_\infty r_\mathrm{p}\left[1+\frac{\gamma-1}{2}M_\infty^2\left(1-\frac{q^2}{q_\infty^2 r_\mathrm{T}}\right)\right]^{\gamma/(\gamma-1)}, \tag{16}$$

$$\rho = \rho_\infty \frac{r_\mathrm{p}}{r_\mathrm{T}}\left[1+\frac{\gamma-1}{2}M_\infty^2\left(1-\frac{q^2}{q_\infty^2 r_\mathrm{T}}\right)\right]^{1/(\gamma-1)}. \tag{17}$$

Here r_p is the ratio of total pressure in the region to freestream total pressure and r_T is the ratio of total temperature in the region to freestream total temperature. The regions are assumed to be separated by fixed wake surfaces on which two jump boundary conditions are applied. The first is the standard static pressure continuity condition, equation (13). If the total pressure and/or temperature differences across the wake are large, the pressure formula cannot be linearized, i.e. μ is not constant in the downstream direction. The second condition is similar to equation (12) but requires a modification to make the answer less sensitive to wake position when total pressure and temperature differences are large. Equation (12) is replaced by

$$\hat{n}\cdot\Delta\vec{W}^* = 0, \tag{18}$$

where

$$\vec{W}^* = \frac{\rho_\infty q_\infty}{\rho_0 q_0}\rho\vec{\nabla}\Phi. \tag{19}$$

Here q_0 is the velocity magnitude which makes $p = p_\infty$ in the given region and ρ_0 is the density at this velocity. Equation (12) becomes a natural jump boundary condition for $\Phi^* = q_\infty\Phi/q_0$ if the

Bateman principle is modified so that

$$J = \int_{\Omega} p^* \, \mathrm{d}V, \tag{20}$$

where

$$p^* = p \frac{\rho_\infty q_\infty^2}{\rho_0 q_0^2}. \tag{21}$$

4. DISCRETIZATION

The boundary surfaces of objects to be modelled are described using networks of piecewise flat surface patches called panels. This input format allows relatively simple specification of complicated surfaces. A representation for a sphere that contains 1600 panels is shown in Figure 1.

The volume grid is generated automatically and is controlled using certain criteria, the most recent of which is a solution-adaptive procedure. We start with a coarse uniform rectangular grid, called the global grid, that contains all the boundary surfaces but is otherwise independent of them. This global grid is used to enforce the far-field condition as well as for the Prandtl–Glauert preconditioner (see Section 5). The global grid is refined locally in a hierarchical manner, i.e. any grid box can be refined into eight geometrically similar boxes of equal volume. This process is repeated to give a grid with any desired local resolution. De-refinement is also possible by locally reversing the process, i.e. removing the eight son boxes of some coarser box.

The non-linear boundary value problem is discretized on this grid using a finite element method. The potential ϕ is defined at the grid points and these are the fundamental unknowns. In addition, doublet parameters must be introduced at the trailing edges of lifting surfaces to simulate lift. The velocity is computed at centroids of the elements and used to compute a density that is constant on each element. If the density were constant everywhere (corresponding to incompressible flow) and there were no boundaries nearby, the result would be the standard seven-point finite difference formula for Poisson's equation. In the non-linear case the density is treated as constant in each region so that it can be factored out of the element stiffness matrix. This matrix is the same up to a constant depending on refinement level for any element not cut by a boundary surface. When assembled, the result is a 27-point finite difference formula. Elements cut by boundaries have special element stiffness matrices as well as special density formulae. The

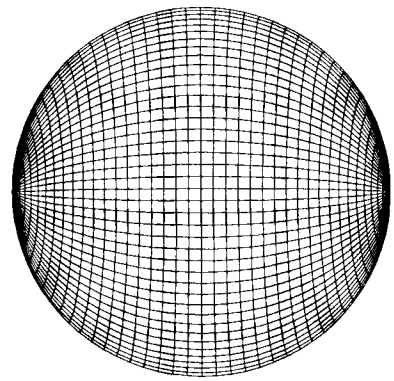

Figure 1. Panelling for sphere.

finite element integrals are computed on the portion of the box cut off by the boundary surface. The unknowns are located at the corners of the box element. In regions of supercritical flow, artificial dissipation must be introduced to rule out expansion shocks. This is done by upwinding the density using neighbouring elements. This discretization is second-order-accurate in the potential and is described in detail in Reference 21. This discrete non-linear operator will be denoted by F. In the present context these exact discretization details are not crucial. It is important to note, however, that in the presence of strong shocks F has large second derivatives, so that gradient information is only of value locally. This is also true of other discretizations of the full potential equation.

5. PRECONDITIONING

Because problems of practical interest are large and not well conditioned, a preconditioned GMRES algorithm is used to solve equation (3).

For the non-linear discrete full potential operator it is now practical in some cases to use a sparse solver as a preconditioner. In two dimensions the linearized problem could be solved directly using a nested dissection ordering based on the computational grid.[15] However, this is not practical for large three-dimensional problems. In TRANAIR a combination of preconditioners is employed, one of which is the Jacobian for F restricted to a part of the computational grid. This restricted Jacobian matrix often remains a good preconditioner for several Newton steps.

To conveniently handle such a preconditioner, a general package has been constructed for input and decomposition of sparse matrices. It has a sorting and merging feature to accumulate contributions to a single matrix element. This process corresponds to constructing the global stiffness matrix from the element stiffness matrices in the finite element context and is completely automated in this sparse matrix package. The package requires an externally generated ordering if fill is to be kept minimal. In the case considered here it is relatively easy to generate a nested dissection ordering using the grid structure and the unknown locations. Incomplete factorization using a dynamic drop tolerance is often helpful for large three-dimensional problems and is a feature of the sparse matrix package. More details about this preconditioner can be found in References 21 and 34.

The other preconditioner is a fast Poisson solver applied to the Prandtl–Glauert equation

$$\mathscr{T}\Phi = (1 - M_\infty^2)\Phi_{xx} + \Phi_{yy} + \Phi_{zz} \tag{22}$$

discretized on a global uniform grid. This equation is a linearization of the full potential equation (9) about $\vec{V}_\infty$. The discrete operator T is the standard seven-point finite difference discretization of $\mathscr{T}$. More details can be found in References 21 and 35.

For this combination of preconditioners the linear GMRES iteration converges to five digits in 10–40 iterations, depending on the number of unknowns in the problem, the size of the supercritical flow region, and the drop tolerance used in the sparse solver preconditioner.

6. CONVERGENCE FOR SUBSONIC PROBLEMS

Newton's method is rarely globally convergent. Also, its convergence rate is generally quadratic only sufficiently close to the solution. In the full potential case the initial iterate is taken to be $\phi = 0$, corresponding to freestream flow. This initialization is frequently used in aerodynamics but is not a good approximation to the solution in many cases. Since the full potential operator linearized about $\vec{V}_\infty$ (this corresponds to linearizing about $\phi = 0$) is the Prandtl–Glauert equation

(9), the solution after one step of Newton's method is the solution to the Prandtl–Glauert equation on the given grid. For problems in which the flow is subsonic this is often a good starting point.

If the linearized equation convergence tolerance η of equation (6) is very small, the convergence of the method should be close to that of an exact Newton method with $\lambda = 1{\cdot}0$ (see equation (2)). This provides a valuable consistency check for the whole code. Any inconsistency in the evaluation of the function F would destroy quadratic convergence. The convergence of Newton's method was tested for the case of a sphere of radius 0·8 with an onset flow condition of $M_\infty = 0{\cdot}2$. In this case the peak Mach number in the solution is 0·3. The grid used in this case contained about 150000 finite elements. The value of η was chosen to be 10^{-7} and the Newton method was converged to 10 digits. Table I gives the norm of the residual as a function of the Newton iteration number. Quadratic convergence since to take place after step 2 since $\log(R_3) \simeq 2\log(R_2)$, where R_n is the L_2-norm of the residual for the non-linear equation (1) at iteration n.

For this case we also examined the total CPU time required on a single-processor CRAY X-MP to reduce the non-linear residual by 10 orders of magnitude as a function of η. This CPU time excludes all overhead costs, including the generation and factorization of the partial Jacobian matrix. The relevant data are shown in Table II for several values of η. Note that minimizing the number of Newton steps does not necessarily minimize the computer time. For minimum computer time, η should start at a fairly large value and decrease as convergence takes place. In practice we have found it convenient to fix $\eta = 0{\cdot}00001$.

7. DAMPING STRATEGIES FOR THE FULL POTENTIAL EQUATION

The situation as described in the previous section is quite satisfactory for subsonic problems. However, when shocks or large stagnation regions are present, convergence of Newton's method is not reliable with the $\phi = 0$ initialization or any other problem-independent initialization.

Table I. Newton method residuals for sphere in subsonic flow

n	R_n	$\log(R_n)$
0	1.0	0·0
1	0·06	−1·2
2	0·0000181	−4·7
3	0·000000000290	−9·5

Table II. Solution CPU cost as a function of η

η	CPU seconds	Newton steps
0·01	419	5
0·0001	416	4
0·000001	429	3
0·0000001	443	3

Damping of Newton's method must be introduced for large transonic problems to prevent divergence or very slow convergence.

Various damping strategies have been tested in the present method. One due to Bank and Rose[19] for determining the step length λ is based on the residual for equation (1). This strategy was chosen because of its simplicity and ease of implementation. It provides adequate local damping in many cases. We tried another trust region method based on a Levenburg–Marquardt method but found it no more effective.

Another strategy that has proved useful in some cases is to further limit λ so that the solution iterate x^{n+1} does not have local Mach numbers greater than some prescribed cut-off value. This prevents spurious large velocities from causing stagnation of convergence. In the ONERA M6 wing results reported below, this strategy was used with a local Mach number cut-off of $\sqrt{5}$.

However, local damping procedures of this kind are only adequate by themselves in cases that almost converge anyway. In difficult transonic cases, convergence of Newton's method can stagnate owing to the formation of a steep shock in the wrong location early in the iterative process. Once this occurs, a local method can rarely move the shock more than one grid point per iteration, resulting in very slow convergence. This situation seems to be due to the fact that the residual is much larger near the shock than elsewhere.

To improve convergence in the presence of shock waves, a problem-dependent dissipation parameter is introduced. This parameter is used in a continuation process. To describe this process, we must first describe the nature of the artificial viscosity used to rule out non-physical expansion shocks in the full potential formulation.

Standard first-order upwinding of the density is used to produce the artificial viscosity required when supersonic flow is present.[4,36] Such an upwinding is given by replacing ρ in the full potential equation with

$$\tilde{\rho} = \rho - \nu \hat{V} \cdot \vec{\nabla}_{-} \rho, \tag{23}$$

where $\hat{V}$ is the normalized local velocity and $\vec{\nabla}_{-}\rho$ is an upwind undivided difference. In equation (23) ν is the switching function given by

$$\nu = \max(0, 1 - M_c^2/M^2), \tag{24}$$

where M is the local Mach number and M_c is the cut-off Mach number assigned the value $M_c^2 = 0{\cdot}95$ chosen to introduce dissipation just below Mach 1·0. More details can be found in Reference 21.

The continuation process can now be described. Initially the discrete problem is modified by multiplying the switching function of equation (24) by a moderate constant (1·5–3·0) and by reducing the cut-off Mach number. This has the effect of increasing the amount of artificial viscosity and applying it to a larger part of the flow field. When the non-linear residual is reduced by two orders of magnitude, the problem is modified by reducing the multiplying factor and raising the cut-off Mach number. This process is repeated until the desired level of dissipation is reached. This continuation process works very well in many cases. It has the effect of locating the supersonic zone and the shock position fairly early in the process, even though the shock is quite smeared. The effect of viscosity damping can be seen in the case of the ONERA M6 wing at $M_\infty = 0{\cdot}84$ and angle of attack $\alpha = 3{\cdot}06°$ on a grid having about 311000 elements. Results are presented for this case in the next section and show a strong shock outboard as well as an oblique supersonic-to-supersonic shock. If Newton's method is used with an initial iterate $\phi = 0$ and the Bank–Rose strategy for limiting λ, the convergence stagnates at the iterate shown in Figure 2. The final converged solution is shown for reference. If λ is further limited to control local Mach

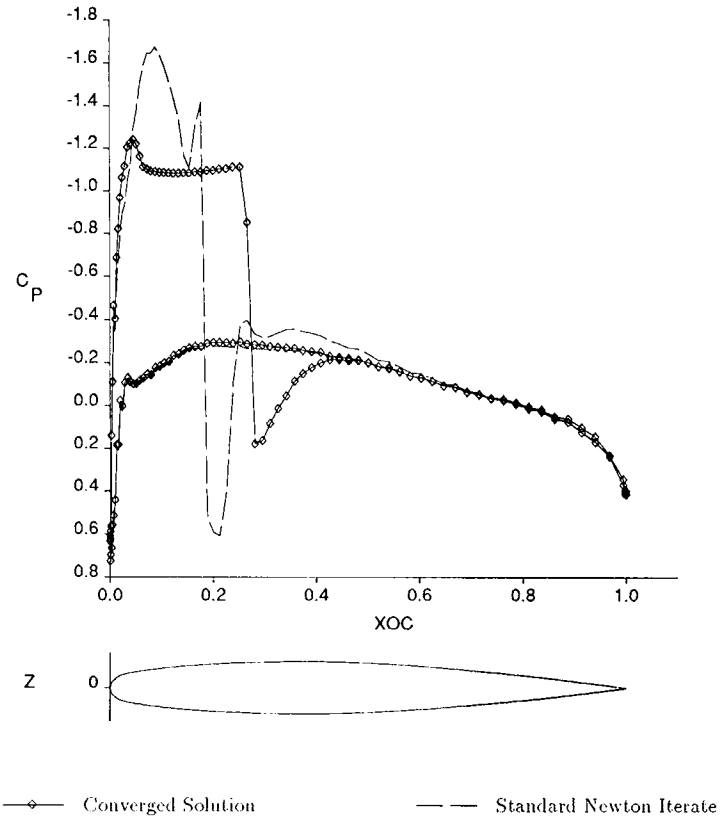

Figure 2. Iterate for Newton's method with residual damping for ONERA M6 wing case, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$, 91% span station.

numbers as described above, the convergence is still very slow. Figure 3 shows the Newton iterate after six and 12 Newton steps. The final converged solution is shown for reference. Newton's method is moving the shock towards the correct location very slowly. When, in addition, viscosity damping is used, convergence is rapid after the initial viscous problems are partially solved. A partially converged solution at the second continuation step (Newton step 7) is shown in Figure 4. Figure 5 shows the convergence histories for these runs. The residual jumps in this figure correspond to discrete changes in the continuation parameter. The drawback of this continuation approach is the high cost of even partially solving the viscous problems that are introduced.

Several continuation strategies were tried that were not very effective. The first was continuation in freestream Mach number (M_∞) and the second was continuation in the total pressure of the freestream. In both cases, shock location was sensitive to the continuation parameter.

8. GRID SEQUENCING

A strategy that has proven to be very reliable for ensuring convergence for difficult transonic problems is grid sequencing. In TRANAIR the final fine grid is currently specified by the user with tolerances based on local panel size and regions of interest as described in Reference 21. This final grid is constructed and then is de-refined one level wherever possible to create the next

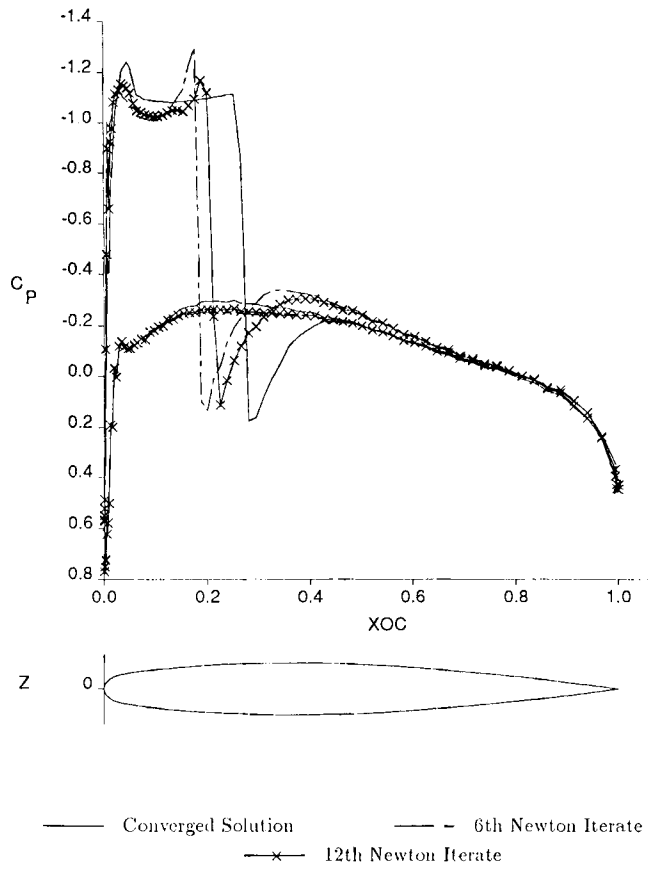

Figure 3. Iterates for Newton's method with residual and local Mach number damping for ONERA M6 wing case, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$, 91% span station.

coarser grid. This strategy results in rapid coarsening of the grids and contributes to computational efficiency. The global grid is also coarsened if possible in this process. This process is continued until a suitably coarse grid is reached. The problem is first solved on this coarse grid. This solution is then interpolated using trilinear interpolation to the next grid in the sequence and this interpolated solution used as an initial guess for Newton's method on this grid. This process is continued until the finest grid is reached. Thus on the fine grid a good initial iterate is obtained at low cost.

The benefits of this approach are more reliable convergence and lower computer cost. Below we present results for several cases that proved very difficult for Newton's method when the initial guess $\phi = 0$ was used. With grid sequencing, all converged rapidly. These cases include a sphere of radius 0·8 with a freestream Mach number of 0·7 ($M_\infty = 0{\cdot}7$), the ONERA M6 wing with $M_\infty = 0{\cdot}84$ and $\alpha = 3{\cdot}06°$, and a powered nacelle where the exhaust contains shock diamonds.

The first case is a sphere with radius 0·8 at $M_\infty = 0{\cdot}7$. At this condition the flow is transonic and contains a strong shock. This case was used to test the effectiveness of the upwinding used in

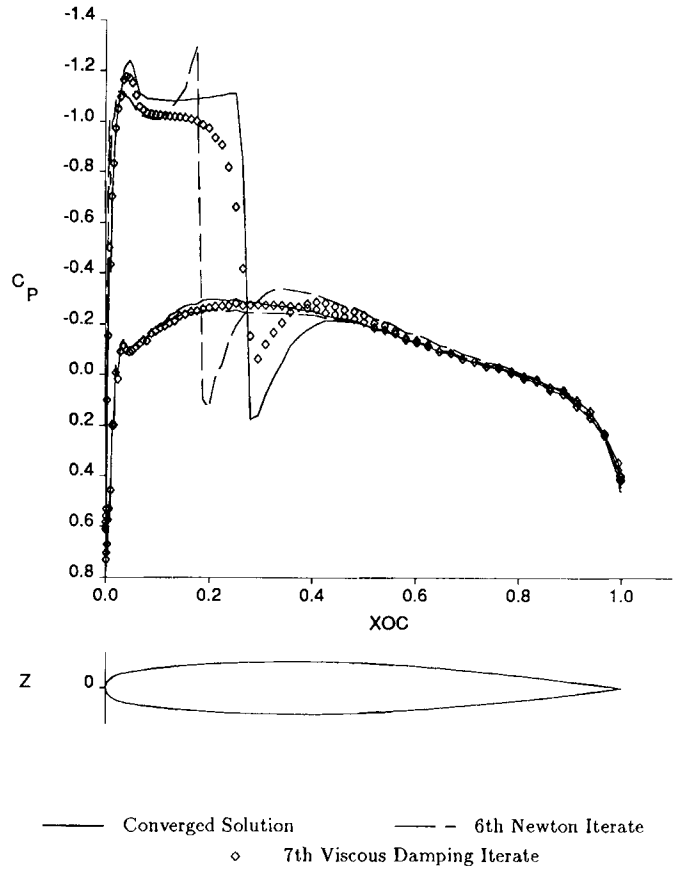

Figure 4. Partially converged iterate for the second continuation step using viscosity damping for ONERA M6 wing case, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$, 91% span station.

TRANAIR. The solution should be axially symmetric. A fine grid was used to test the accuracy of the TRANAIR discretization. The grid contained about 170000 elements and two planes of symmetry were used to reduce the size of the problem. A cut through this grid is shown in Figure 6. Figure 7 shows surface Mach numbers as a function of x. Values at all circumferential stations are plotted. The solution is quite symmetric and also captures the well known re-expansion phenomenon at the foot of the shock. Figure 8 shows the convergence history for this case with grid sequencing and with viscosity damping described in the previous section. Five continuation steps were needed to achieve convergence with viscosity damping in this case. Significant step size damping was required for the first viscous problem. No step size damping was needed with grid sequencing.

A standard aerodynamic test case is flown about the ONERA M6 wing at $M_\infty = 0{\cdot}84$ and $\alpha = 3{\cdot}06°$. This case has an oblique supersonic-to-supersonic shock as well as a normal supersonic-to-subsonic shock. Figure 14 (see later) shows surface pressures at four stations obtained from running TRANAIR on three grids. The fine grid had about 311000 elements and the TRANAIR solution on this grid agrees well with other inviscid codes.[21] As discussed in Section 7,

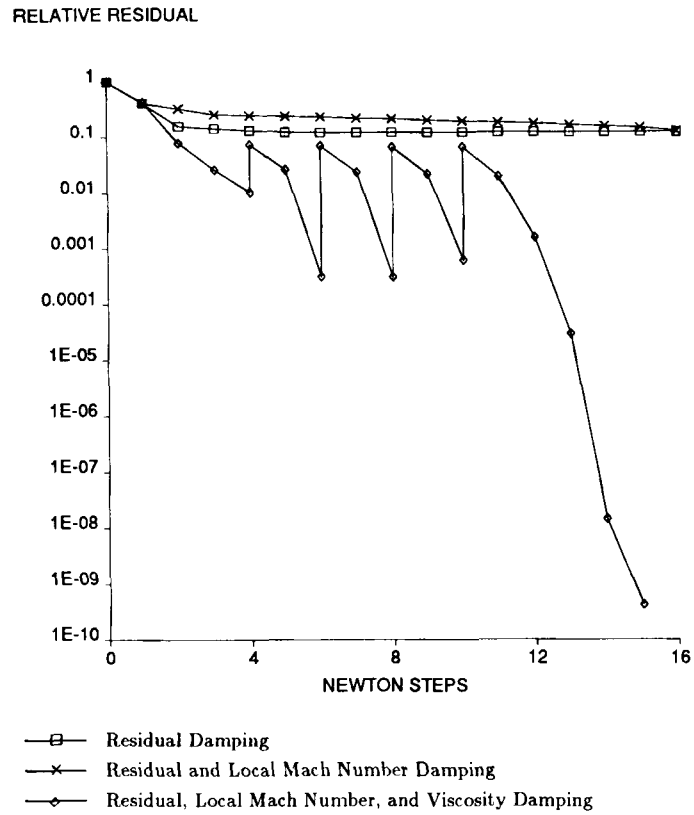

Figure 5. Convergence histories for Newton's method with various damping strategies for ONERA M6 wing case, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06$.

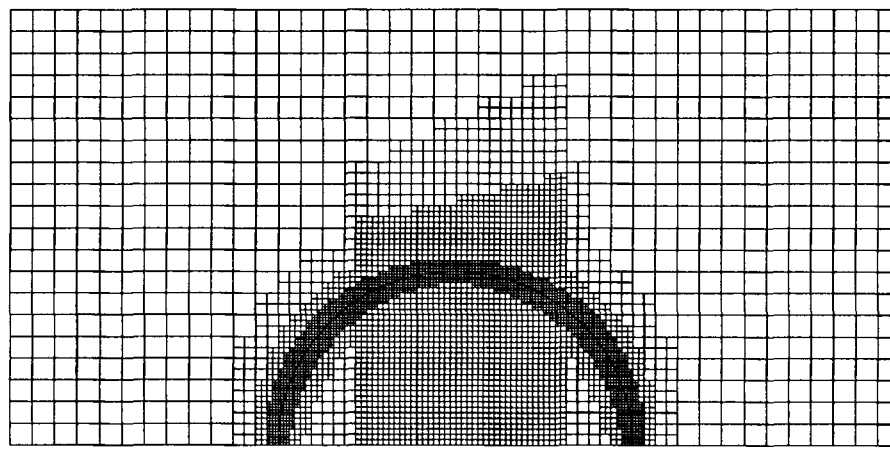

Figure 6. Cut through the grid for a sphere in transonic flow, $M_\infty = 0{\cdot}7$.

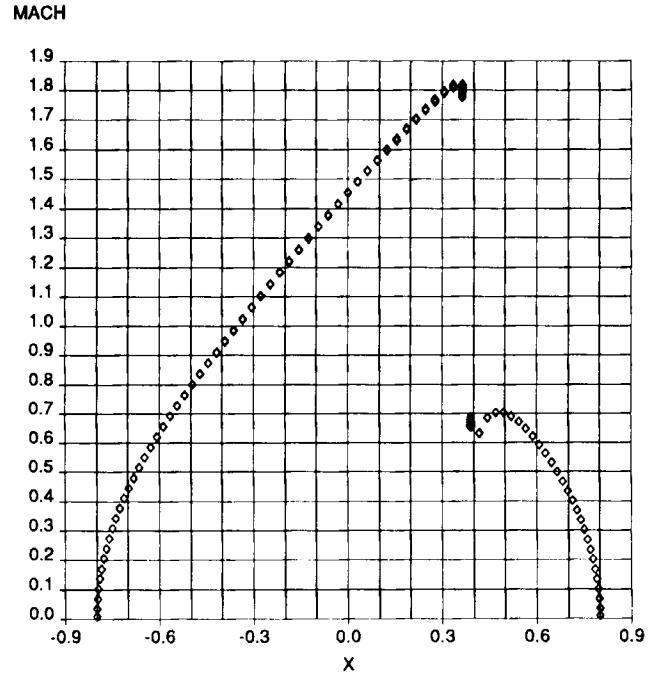

Figure 7. Surface Mach numbers for sphere, $M_\infty = 0{\cdot}7$.

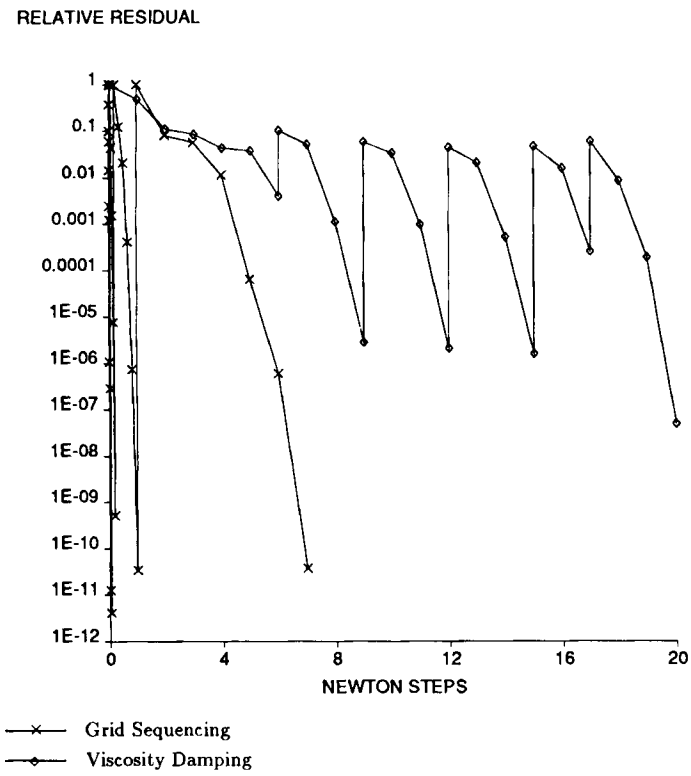

Figure 8. Convergence histories for viscosity damping method and grid sequencing method for sphere case, $M_\infty = 0{\cdot}7$.

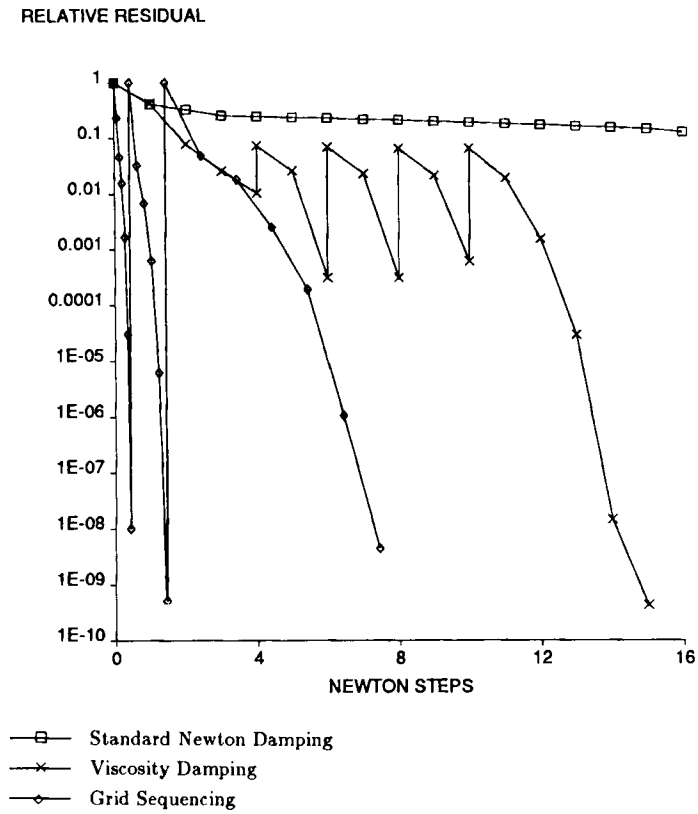

Figure 9. Convergence histories for Newton's method, Newton's method with viscosity damping, and grid sequencing for ONERA M6 wing case, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$.

this case did not converge when residual and local Mach number damping was used with Newton's method with an initial guess of $\phi = 0$. Convergence was obtained in two ways. Initially, viscosity damping was used and it was found that four continuation steps were required. A detailed description of the result is contained in the preceding section. Partially converging the viscous problems accounted for about two-thirds of the TRANAIR iteration steps. With grid sequencing this case converged more rapidly and CPU times were proportionally reduced. Figure 9 compares convergence histories for these three methods. Iterations in the grid-sequencing run are scaled by the approximate size of the problem for the early small grids (this scaling corresponds approximately to CPU cost). Grid sequencing offers a substantial advantage in both rate of convergence and storage requirements. For the grid-sequencing run the CPU time was about half of that needed for the viscosity-damping run. About 40% of this time was overhead used to generate the grids and operators on those grids. SSD requirements were 174000 blocks (a block is 512 words), which is about 20% less than the requirement for the viscosity-damping run. Cuts through the three grids used are shown in Figures 10–13 . The final fine grid is the last of these three grids. The grids had about 19000, 56000 and 311000 elements respectively. Figure 14 shows surface pressures obtained on the three grids used in this case. On the coarser grids the shock is in the right location but smeared.

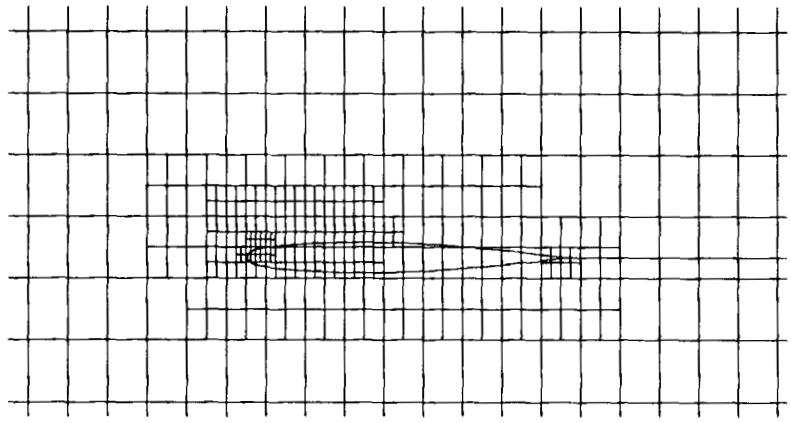

Coarse Grid

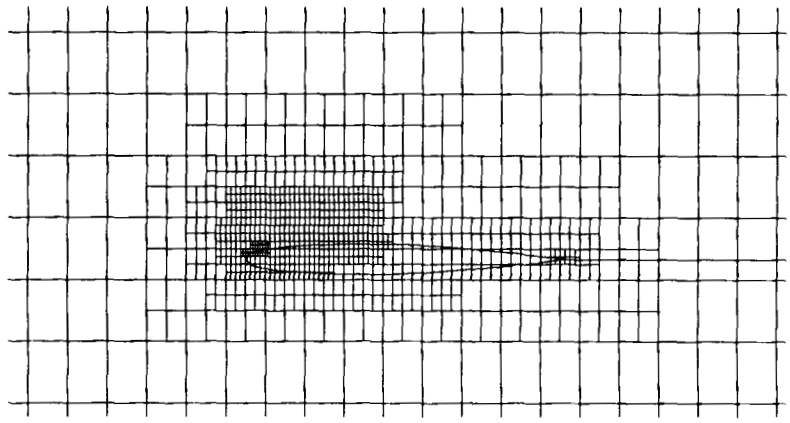

Medium Grid

Figure 10. Cuts through the coarse and medium grids generated by grid sequencing for ONERA M6 wing at 91% span, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$.

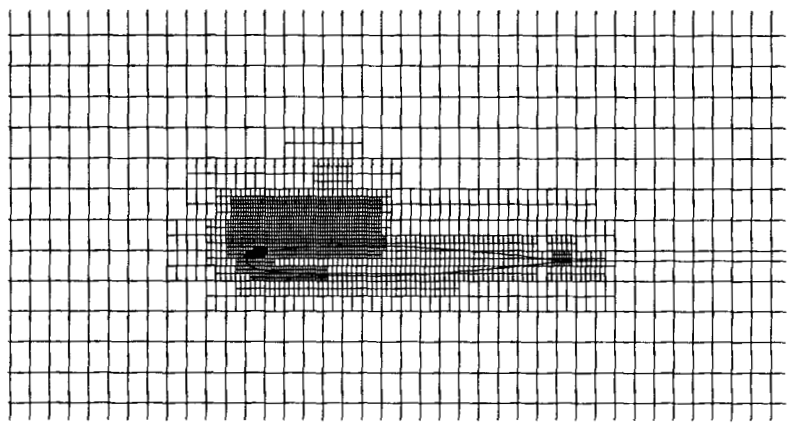

Fine Grid

Figure 11. Cut through the find grid generated by grid sequencing for ONERA M6 wing at 91% span, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$.

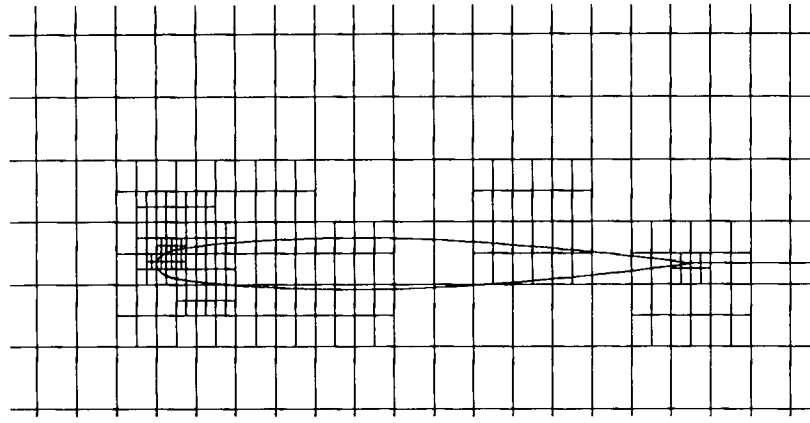

Coarse Grid

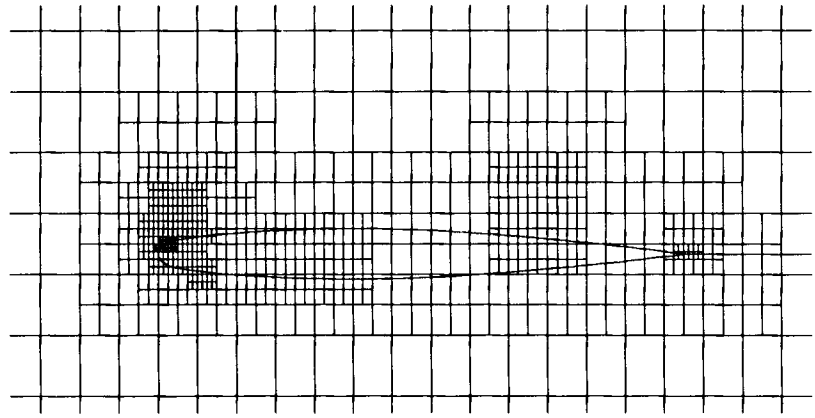

Medium Grid

Figure 12. Cuts through the coarse and medium grids generated by grid sequencing for ONERA M6 wing at the plane of symmetry, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$.

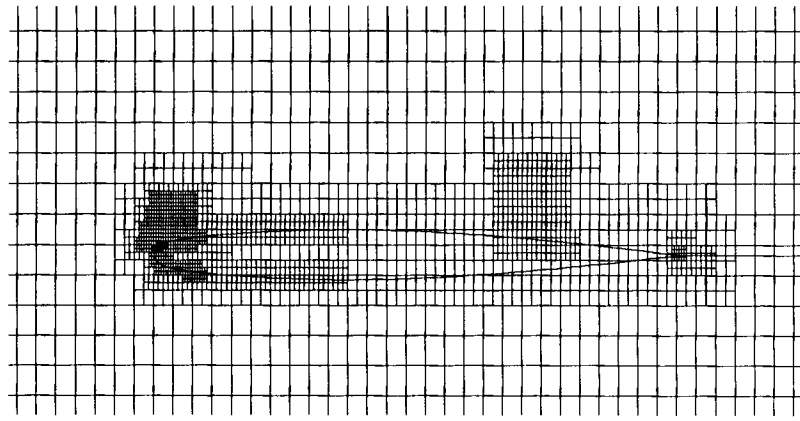

Fine Grid

Figure 13. Cut through the find grid generated by grid sequencing for ONERA M6 wing at the plane of symmetry, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$.

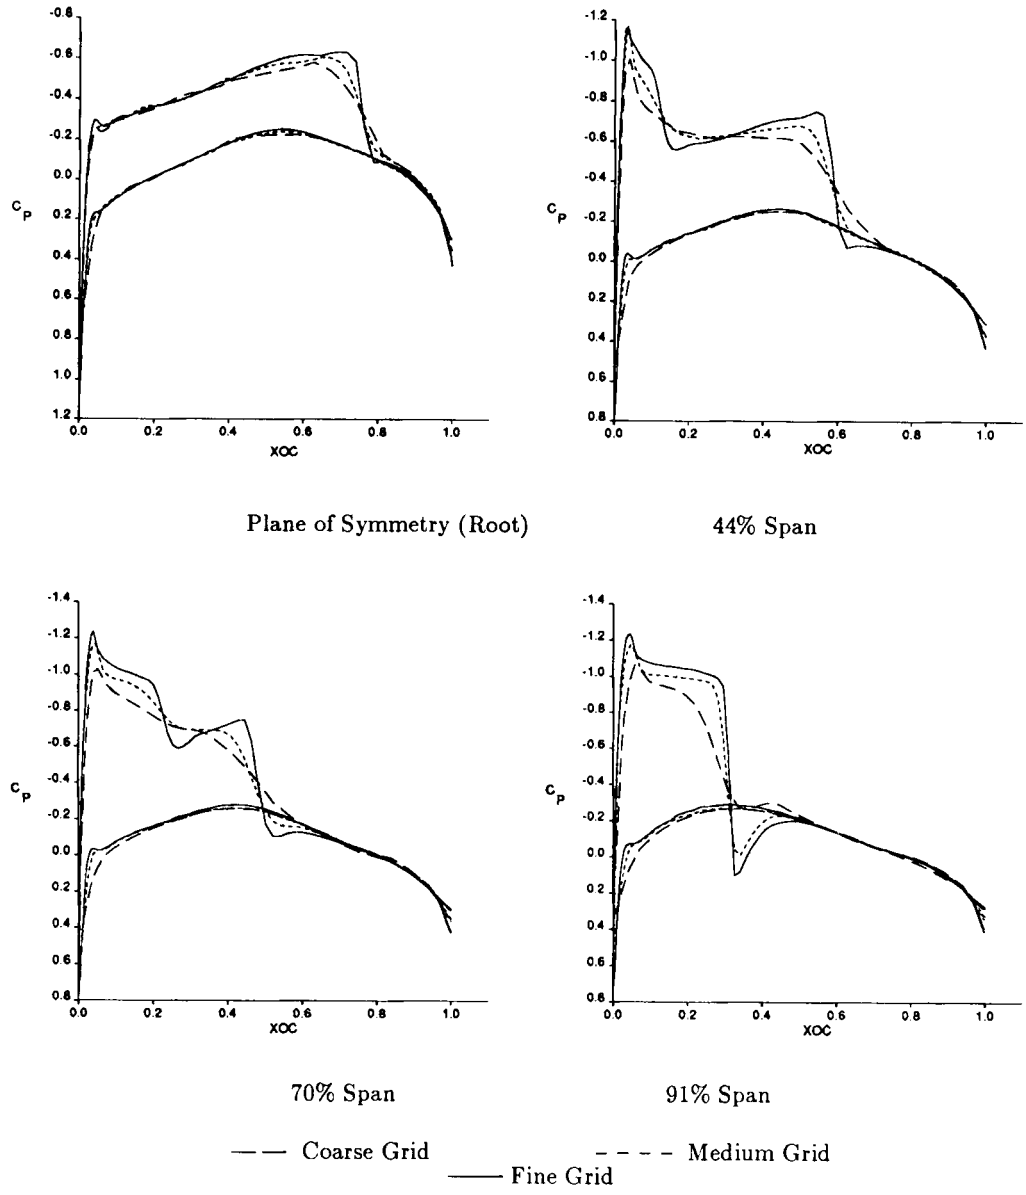

Figure 14. Surface pressure for the three grids generated by grid sequencing for ONERA M6 wing, $M_\infty = 0{\cdot}84$, $\alpha = 3{\cdot}06°$.

A third difficult case that was solved with grid sequencing illustrates the capability to model different total pressure in an axisymmetric nacelle for which static test data are available.[37] This configuration is shown in Figure 15 and had about 5000 panels. The wakes were panelled so that the powered streams maintain equal area downstream from the exits. Two planes of symmetry were used for this case. Results are shown for the final grid with about 86000 elements. Grid sequencing was necessary to obtain reliable convergence for this case, even for fairly coarse grids. Five grids were used, with 539, 880, 2828, 16030 and 85983 elements. The total pressure in the

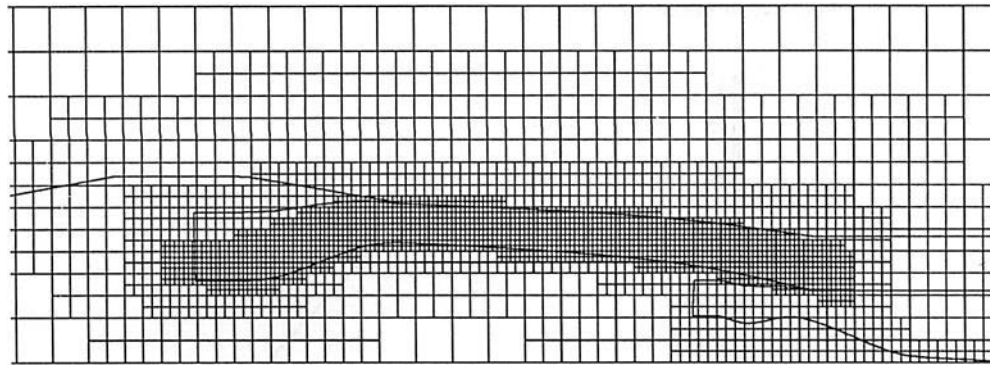

Figure 15. Cut through the grid for an axisymmetric powered nacelle, $M_\infty = 0{\cdot}1$.

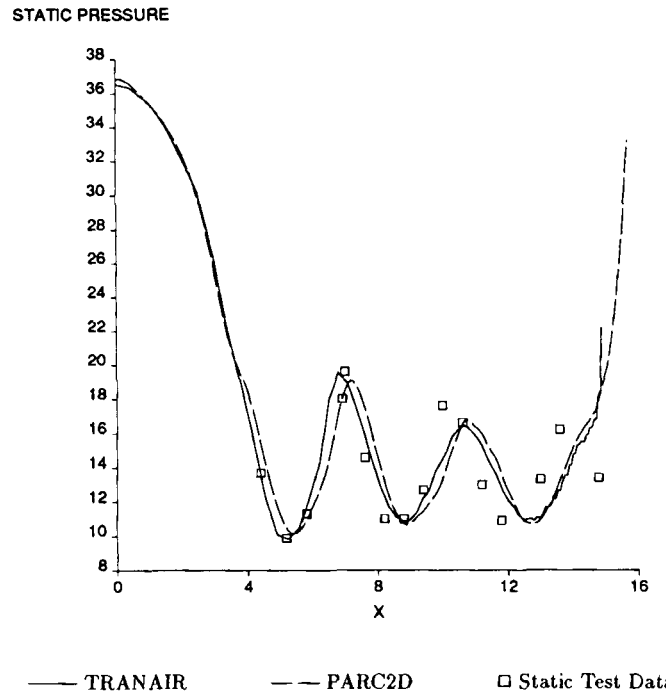

Figure 16. Static pressure for axisymmetric powered nacelle compared to experiment and Navier–Stokes code results, $M_\infty = 0{\cdot}1$.

powered stream was 2·807 and that in the primary stream was 2·3425. The static pressure on the core cowl is compared with experimental data and with the results of running a Navier–Stokes code PARC2D[37,38] in Figure 16. In this case PARC2D predicted no total pressure loss in the fan stream. Thus inviscid modelling can capture the major features of this flow. Finally, Figure 17 gives the convergence history for this case. Once again the steps on the coarser grids are scaled by the number of elements in the grid.

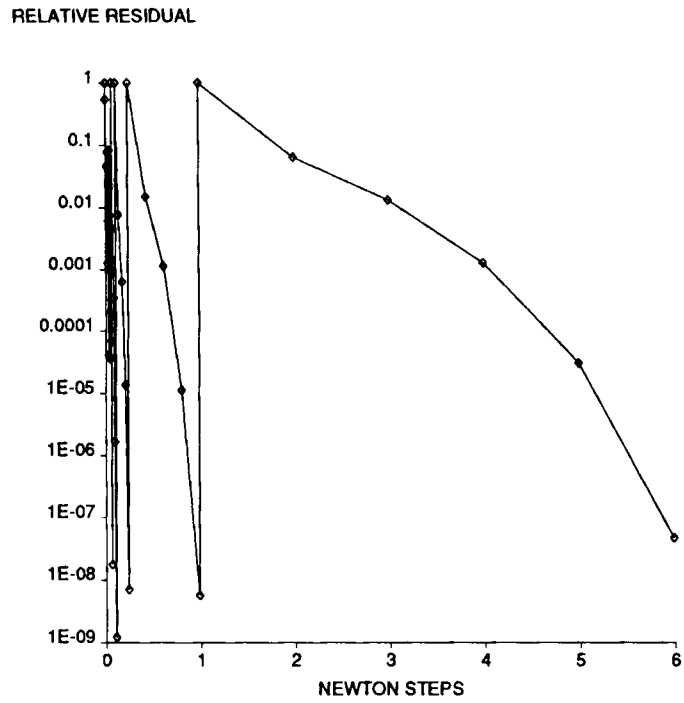

Figure 17. Convergence history for grid sequencing method for axisymmetric powered nacelle case, $M_\infty = 0{\cdot}1$.

9. CONCLUSIONS

The application of an inexact Newton method to the full potential equation of aerodynamics has been examined. Several strategies have been examined to extend the range of problems for which Newton's method converges and to accelerate convergence. These strategies are particularly helpful when strong shocks are present.

ACKNOWLEDGEMENTS

Many people have aided us in this work on TRANAIR, including Larry Wigton of Boeing Commercial Airplanes, John Bussoletti of Boeing Advanced Systems and Richard Burkhart of Boeing Computer Services. Dave Foutch of Boeing Commercial Airplanes supplied us with the axisymmetric nacelle configuration, test data and PARC2D results. This work was supported in part by National Science Foundation grant ASC-8519353 and NASA contract NAS2-12513.

REFERENCES

1. A. Jameson, 'Computational transonics', *Commun. Pure Appl. Math.*, **41**, 507–549 (1988).
2. A. Jameson, 'Transonic flow calculations', *Princeton University Department of Mechanical and Aerospace Engineering Report 1651*, 1983.
3. A. Jameson, W. Schmidt and E. Turkel, 'Numerical solutions of the Euler equations by finite volume methods using Runge–Kutta time-stepping schemes', *AIAA Paper 81-1259*, 1981.
4. T. L. Holst and W. F. Ballhaus, 'Fast conservative schemes for the full potential equation applied to transonic flows', *AIAA J.*, **17**, (1979).

5. T. A. Reyhner, 'Three-dimensional transonic potential flow about complex three-dimensional configurations', *NASA CR-3814*, 1984.
6. G. R. Shubin, A. B. Stephens and H. M. Glaz, 'Steady shock tracking and Newton's method applied to one-dimensional duct flow', *J. Comput. Phys.*, **39**, 364–374 (1981).
7. G. R. Shubin, A. B. Stephens, H. M. Glaz, A. B. Wardlaw and L. B. Hackerman, 'Steady shock tracking, Newton's method, and the supersonic blunt body problem', *SIAM J. Sci. Stat. Comput.*, **3**, 127–144 (1982).
8. B. Fornberg, 'A numerical study of viscous flow past a cylinder', *J. Fluid Mech.*, **98**, 819–855 (1980).
9. M. S. Engelman, G. Strang and K.-J. Bathe, 'The application of quasi-Newton methods in fluid mechanics', *Int. j. numer. methods eng.*, **17**, 707–718 (1981).
10. R. E. Childs and T. H. Pulliam, 'A Newton multigrid method for the Euler equations', *AIAA Paper 84-0430*, 1984.
11. M. Giles, M. Drela and W. T. Thompkins, 'Newton solution of direct and inverse transonic Euler equations', *AIAA Paper 85-1530*, 1985.
12. M. Giles and M. Drela, 'A two-dimensional transonic aerodynamic design method', *AIAA Paper 86-1793*, 1986.
13. M. Drela, 'Two-dimensional transonic aerodynamic design and analysis using the Euler equations', *Ph.D. Thesis*, MIT Department of Aeronautics and Astronautics, 1985.
14. M. B. Giles, 'Newton solution of steady two-dimensional transonic flow', *Ph.D. Thesis*, MIT Department of Aeronautics and Astronautics, 1985.
15. L. B. Wigton, 'Application of MACSYMA and sparse matrix technology to multielement airfoil calculations', *AIAA Paper 87-1142-CP*, 1987.
16. V. Venkatakrishnan, 'Newton solution of inviscid and viscous problems', *AIAA Paper 88-0413*, 1988.
17. V. Venkatakrishnan and T. J. Barth, 'Application of direct solvers to unstructured meshes for the Euler and Navier–Stokes equations using upwind schemes', *AIAA Paper 89-0364*, 1989.
18. R. Fletcher, *Practical Methods of Optimization, Vol. 1*, Wiley, New York, 1980.
19. R. E. Bank and D. J. Rose, 'Global approximate Newton methods', *Numer. Math.*, **37**, 279–295 (1981).
20. R. S. Dembo, S. C. Eisenstat and T. Steihaug, 'Inexact Newton methods', *SIAM J. Numer. Anal.*, **19**, 400–408 (1982).
21. D. P. Young, R. G. Melvin, M. B. Bieterman, F. T. Johnson, S. S. Samant and J. E. Bussoletti, 'A locally refined rectangular grid finite element method: application to computational fluid dynamics and computational physics', *Boeing Computer Services Technical Report SCA-TR-108*, 1989; to appear in *J. Comput. Phys.*
22. P. N. Brown and A. C. Hindmarsh, 'Matrix-free methods for stiff systems of ODE's', *SIAM J. Numer. Anal.*, **23**, 610–638 (1986).
23. P. N. Brown and A. C. Hindmarsh, 'Reduced storage matrix methods in stiff ODE systems', *Appl. Math. Comput.*, **31**, 40–91 (1989).
24. T. F. Chan and K. R. Jackson, 'The use of iterative linear equation solvers in codes for large systems of stiff IVP's for ODE's', *Technical Report 170/84*, Department of Computer Science, University of Toronto, 1984.
25. T. Kerkhoven and Y. Saad, 'Acceleration techniques for decoupling algorithms in semiconductor simulation', *Center for Supercomputer Research and Development Report CSRD 686*, University of Illinois at Urbana, 1987.
26. P. N. Brown and Y. Saad, 'Hybrid Krylov methods for nonlinear systems of equtions', *Center for Supercomputer Research and Development Report CSRD 699*, University of Illinois at Urbana, 1987.
27. T. F. Chan and K. R. Jackson, 'Nonlinearly preconditioned Krylov subspace methods for discrete Newton algorithms', *SIAM J. Sci. Stat. Comput.*, **5**, 533–542 (1986).
28. Y. Saad and M. H. Schultz, 'GMRES: a generalized minimal residual method for solving nonsymmetric linear systems', *Yale University Department of Computer Science Research Report YALEU/DCS/RR-254*, 1983; *SIAM J. Sci. Stat. Comput.*, **7**, 856 (1986).
29. O. Axelson and V. A. Barker, *Finite Element Solutions of Boundary Value Problems: Theory and Computation*, Academic Press, New York, 1984.
30. D. M. Young, *Iterative Solution of Large Linear Systems*, Academic Press, New York, 1971.
31. L. B. Wigton, N. J. Yu and D. P. Young, 'GMRES acceleration of computational fluid dynamics codes', *AIAA Paper 85-1494*, 1985.
32. P. E. Rubbert and G. R. Saaris, 'Review and evaluation of a three-dimensional lifting potential flow analysis method for arbitrary configurations', *AIAA Paper 72-188*, 1972.
33. H. Bateman, 'Irrotational motion of a compressible inviscid fluid', *Proc. Natl. Acad. Sci.*, **16**, 816–825 (1930).
34. D. P. Young, R. G. Melvin, F. T. Johnson, J. E. Bussoletti, L. B. Wigton and S. S. Samant, 'Application of sparse matrix solvers as effective preconditioners', *SIAM J. Sci. Stat. Comput.*, **10**, 1186–1199 (1989).
35. P. E. Rubbert, J. E. Bussoletti, F. T. Johnson, K. W. Sidwell, W. S. Rowe, S. S. Samant, G. SenGupta, W. H. Weatherill, R. H. Burkhart, B. L. Everson, D. P. Young and A. C. Woo, 'A new approach to the solution of boundary value problems involving complex configurations', in A. K. Noor (ed.), *Computational Mechanics — Advances and Trends*, The American Society of Mechanical Engineers, New York, 1986, pp. 49–84.
36. M. Hafez, E. M. Murman and J. C. South, 'Artificial compressibility methods for numerical solutions of transonic full potential equations', *AIAA Paper 78-1148*, 1978.
37. C. C. Clark and D. W. Foutch, 'PARC analysis of an axisymmetric turbofan nozzle', *Boeing Commercial Airplanes Technical Report PROP-BN31U-C89-020*, 1989.
38. G. K. Cooper, 'The PARC code: theory and usage', *Arnold Engineering Development Center Technical Report AEDC-TR-87-24*, 1987.