INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS, VOL. 15, 1333–1355 (1992)

SEGMENTED MULTIGRID DOMAIN DECOMPOSITION PROCEDURE FOR INCOMPRESSIBLE VISCOUS FLOWS

KUMAR SRINIVASAN*

Department of Aerospace Engineering and Engineering Mechanics, ML No. 343, University of Cincinnati, Cincinnati, OH 45221, U.S.A.

AND

STANLEY G. RUBIN†

Department of Aerospace Engineering and Engineering Mechanics, ML No. 343, University of Cincinnati, Cincinnati, OH 45221, U.S.A.

SUMMARY

The application of grid stretching or grid adaptation is generally required in order to optimize the distribution of nodal points for fluid-dynamic simulation. This is necessitated by the presence of disjoint high gradient zones, that represent boundary or free shear layers, reversed flow or vortical flow regions, triple deck structures, etc. A domain decomposition method can be used in conjunction with an adaptive multigrid algorithm to provide an effective methodology for the development of optimal grids. In the present study, the Navier–Stokes (NS) equations are approximated with a reduced Navier–Stokes (RNS) system, that represents the lowest-order terms in an asymptotic *Re* expansion. This system allows for simplified boundary conditions, more generality in the location of the outflow boundary, and ensures mass conservation in all subdomain grid interfaces, as well as at the outflow boundary. The higher-order (NS) diffusion terms are included through a deferred corrector, in selected subdomains, when necessary. Adaptivity in the direction of refinement is achieved by grid splitting or domain decomposition in each level of the multigrid procedure. Normalized truncation error estimates of key derivatives are used to determine the boundaries of these subdomains. The refinement is optimized in two co-ordinate directions independently. Multidirectional adaptivity eliminates the need for grid stretching so that uniform grids are specified in each subdomain. The overall grid consists of multiple domains with different meshes and is, therefore, heavily graded. Results and computational efficiency are discussed for the laminar flow over a finite length plate and for the laminar internal flow in a backward-facing step channel.

KEY WORDS Multigrid Domain decomposition Flux splitting Deferred corrector

1. INTRODUCTION

Numerical solution procedures that are designed to compute accurately and efficiently steady, high Reynolds number (*Re*) flows necessitate the use of local grid refinement in regions of the flow field wherein the velocity or pressure gradients can be very large, e.g. boundary, shear or vortical layers, shocks waves, etc. These regions are generally confined to relatively small portions of the overall flow domain and are not known *a priori*. An adaptive multigrid domain decomposition

* Graduate Research Assistant.
† Professor.

0271–2091/92/231333–23$16.50

Received 14 June 1991
Revised 15 June 1992

procedure is shown herein to be effective for identifying and resolving all local, and also globally influential, phenomena. This procedure also ensures that an optimal level of grid refinement is provided in the differing domains of the flow field. Mass conservation is maintained throughout, on all grid levels, and boundary conditions, including those at the outflow, are accurately satisfied.

Local direction-selective refinement, that is driven by specified flow parameters and accuracy limits, is achieved by sequentially splitting the overall flow domain into a variety of subdomains. In the present analysis, the grid refinement process is performed, independently, in two co-ordinate directions by combining a domain decomposition strategy with an adaptive multigrid algorithm. In this approach, each grid in the multigrid hierarchy is of equal or lesser extent than all of its coarser predecessors. The subgrids are split into several multidimensional subdomains that are defined by specified directional and global resolution requirements. A similar approach has been presented in Reference 1 for steady-state problems; although, no attempt was made to meet the disparate need for refinement in two different co-ordinate directions. Thus, when solving a simple boundary layer flow, wherein the gradients are dominant in the 'normal' direction, the 'streamwise' grid may also be refined unnecessarily. In the present investigation, the 'subdomaining' process is based on requirements that allow for segmentally varying grid resolution in multiple (two) directions and throughout the flow field. Truncation errors in key flow parameters, such as the pressure gradient and the vorticity gradient, are utilized for this purpose.

The domain decomposition approach allows for the flexibility of employing different solvers in different subdomains of the flow field. In some regions, where interactions are very strong, a direct solver can significantly improve the efficiency of relaxation methods. In addition, unsteady phenomena are often confined to small regions of the flow-field. These can be identified through the domain decomposition procedure, and time-accurate methods can be used in these subdomains. Although only steady-state solutions are presented in the current work, the same can be extended to unsteady flows by treating each time-consistent step as a standard steady-state calculation. The segmented grid from the previous time level can be used initially for the present time step and the grid can be adjusted to accommodate the changes in the solution at the new time level.

In the present investigation several two-dimensional, steady, laminar, incompressible-flow examples have been considered. The reduced Navier–Stokes (RNS) system with an explicit deferred corrector (DC) for full Navier–Stokes analysis has been applied implicitly with a pressure-based form of flux vector discretization[2] to compute the triple-deck structure surrounding the trailing edge of a finite flat plate, and the reversed flow regions that develop in a backstep channel. The pressure-based flux-split formulation leads to a precise prescription of surface normal boundary conditions on all local subdomain boundaries as well as the outflow boundary. This ensures that global mass conservation requirements are satisfied automatically. This is generally not the case with most characteristic-based Navier–Stokes schemes, where special care has to be taken in order to ensure global mass conservation on each grid. The primitive variable RNS solver is also directly applicable and totally consistent on non-staggered grids. This is contrary to many other incompressible, primitive-variable, Navier–Stokes finite difference formulations, that require pressure Poisson solver or artificial compressibility concepts.

The RNS pressure-based flux-split discretization allows for a clear prescription of the outflow boundary condition. A condition is required only for the pressure at the outflow. The non-reflectivity of this boundary condition has been tested by truncating the domain length that is used to compute the flow in a backstep channel, i.e. bringing the outflow boundary much closer to the inflow. In addition, the use of the adaptive multigrid domain decomposition procedure allows the outflow boundary to be placed far downstream without increasing the computational cost

significantly. This removes any uncertainty concerning the influence of the outflow boundary location on the solution accuracy.

Another advantage of the present method is that grid independence of the solution is ensured in a very efficient manner. In non-adaptive calculations, fine grid computations become very expensive and establishing grid convergence can be quite costly. This is not the case in the present method, as the refinement is performed selectively and, hence, the use of very fine grids, in subdomains requiring such resolution, becomes feasible.

2. GOVERNING EQUATIONS AND DISCRETIZATION

The lowest-order RNS system of equations, in arbitrary sheared Cartesian co-ordinates and in non-conservation form, is written as follows:

continuity

$$u_\xi + v_\eta = 0 \tag{1a}$$

ξ-momentum

$$uu_\xi + uy'_b(V + y'_b u)_\xi + V[(1 + y'^2_b)u_\eta + y'_b V_\eta] + p_\xi = \frac{1}{Re} u_{\eta\eta} + \mathrm{DC} \tag{1b}$$

η-momentum

$$u(V + y'_b u)_\xi + V(V + y'_b u)_\eta + p_\eta = \mathrm{DC} \tag{1c}$$

where $\xi = x$, $\eta = y - y_b(x)$ and $V = v - y'_b u$ is the contravariant velocity component in the η or normal direction (for $y'_b(x) \ll 1$); $y_b(x)$ is the surface definition. (u, v) are the Cartesian velocities in the (x, y) directions, respectively. This system of equations is obtained by neglecting the ξ-diffusion terms in the surface-oriented ξ-momentum equation and all diffusion terms in the 'normal' or η-momentum equation, i.e. the DC deferred corrector terms. These terms are higher-order in the parameter expansion of the full Navier–Stokes (NS) equations and are retained in an explicit DC when full NS solutions are required. Generally speaking, the effect of the DC is negligible for the problems considered herein. The lower-order RNS system represents a composite of the Euler and second-order boundary layer equations.[3] This system has been used effectively for the computation of two-three-dimensional, incompressible/compressible, steady/unsteady, laminar/turbulent flows and for a variety of geometries.[4] The surface-oriented direction ξ is generally defined by an appropriate co-ordinate transformation and does not have to be in the direction of the mainstream, e.g. in the stagnation region the mainstream is in the surface normal (η) direction,[5] and for a driven cavity and Cartesian co-ordinates, the ξ direction changes from x along the horizontal surfaces to η along the vertical side-walls, where the RNS system is now defined to include the v_{xx} term in the y-momentum equation. The RNS methodology fails for very low values of Re or for high-frequency waves, where the ξ-diffusion terms are not higher-order and must be retained in full NS solvers.

System (1) is discretized with a pressure-based form of flux-vector splitting for the axial or ξ derivatives, which are entirely inviscid in character, and a 'box' (or trapezoidal) scheme for the continuity and η-momentum equations in the normal or η direction.[2] Therefore, the discrete continuity equation is centred at $(i, j-\frac{1}{2})$. The ξ-momentum equation is centred at (i, j) and the η-momentum equation at $(i, j+\frac{1}{2})$. Trapezoidal two-point $(i, j \pm \frac{1}{2})$, in (1a) and (1c), or three-point (i, j), in (1b), central differencing is used for the normal derivatives; all convective axial derivatives are upwind-differenced. The equations are centred at different grid locations, namely, $(i, j-\frac{1}{2})$ for

continuity, (i, j) for the ξ-momentum and $(i, j+\frac{1}{2})$ for the η-momentum equations; however, the flow variables are evaluated at the non-staggered grid point location (i, j). The p_ξ term in the ξ-momentum equation and the u_ξ term in the continuity equation are flux-vector-differenced with the pressure-based form of flux-vector splitting.[2] In this technique, the p_ξ term is represented by

$$p_\xi = \omega_{i-1/2}(p_i - p_{i-1})/\Delta\xi_i + (1-\omega_{i+1/2})(p_{i+1} - p_i)/\Delta\xi_i, \tag{2}$$

where

$$\omega_i = [\gamma M_\xi^2/(1+(\gamma-1)M_\xi^2), 1]_{\min},$$

$$\omega_{i\pm1/2} = (\omega_i + \omega_{i\pm1})/2.$$

Here M_ξ is the streamwise Mach number and γ is the ratio of specific heats. This reduces to a simple 'forward' difference for the incompressible flows considered herein. This introduces an elliptic acoustic interaction or upstream pressure influence, through the $p_\xi = (p_{i+1} - p_i)/(\Delta\xi)_i$ contribution in (2). The calculated pressure p_i for incompressible flow ($\omega_i \equiv 0$) is influenced only by the downstream pressure p_{i+1} and upstream velocities u_{i-1}, v_{i-1}. In fact, the inflow pressure is not prescribed but determined from the solver. This reflects the upstream acoustic influence that is included in the RNS procedure. For compressible flows, equation (2) shows both upstream and downstream pressure influences. For supersonic flow ($\omega_i \equiv 0$), the upstream influence vanishes as it should. This form of pressure-based flux discretization plays an important role in determining the required surface, upper, outflow and interface boundary conditions, and in satisfying interface mass conservation requirements. All diffusion terms, including those in the DC, are given by conventional central differences.

3. BOUNDARY CONDITIONS

The boundary conditions are prescribed as follows:

(a) At the inflow, u and v are specified; p is determined from the solution.
(b) No slip is imposed at all solid walls; p is determined from the solution.
(c) For external flow, u and p (generally, $p=0$) are imposed at the outer (in η) boundary; v is determined from the solution through mass conservation or the continuity equation (1a).
(d) If at the outflow (in ξ) $u>0$, p_ξ is prescribed for external flows, while for internal flows; only p is prescribed. This is a result of upwinding of the convective terms. If $u<0$ at the outflow (in ξ), e.g. subdomain within a separated flow region, then the negative convective fluxes are assumed to be higher-order and are neglected. This is similar to the FLARE approximation used in separated boundary layer theory. With this approximation, only the pressure or pressure gradient is specified at the outflow boundary.

These are the natural boundary conditions on the global flow domain. In the adaptive multigrid domain decomposition process, the flow field is split into a number of subdomains. These have local boundaries and, therefore, it becomes necessary to impose modified boundary conditions at the subdomain interfaces. The individual subdomains communicate with each other through boundary condition interchange. This procedure will be explained further in the section on interdomain transfers and conservation at grid interfaces.

The pressure flux-split discretization and the RNS approximation play a key role in determining the boundary conditions. Significantly, at the outflow boundary, for internal flows only a simple Dirichlet condition on the pressure is required. No condition is required on the velocities, which are satisfied exactly from mass and momentum conservation. This differs from

the boundary conditions that are usually specified with most other full NS solvers. For this reason, the outflow boundary can now be moved much farther upstream. This is possible when the pressure can be prescribed accurately, without distorting the solution near the outflow boundary.

A downstream boundary condition is required for the velocities only if reversed flow occurs at the outflow. In this case, the requirement for a velocity boundary condition can be removed by neglecting the negative convective fluxes at the outflow station. For weak reversed flows this is an excellent approximation, as the small negative velocities have only a minor influence on the solution. Similarly, for external flows, with $u>0$, only a Neumann condition on the pressure is required. The terms in the full NS equations, that are added explicitly through the deferred corrector, do not alter these boundary conditions in any way. Therefore, the advantages of the RNS discretization, i.e. simplicity of boundary conditions and efficiency of the solution procedure, are retained even for full NS solutions with a small DC correction.

4. GRID STRUCTURE

Standard multigrid methods use a heirarchy of grids ranging from coarse to fine. The present adaptive method aims at limiting the extent of the fine grids to regions where the flow gradients are large. In this manner, excessive grid points can be eliminated. In general, the Nth multigrid level consists of several subdomains. Each multigrid level has an equal or lesser extent than the next coarser grid in the multigrid hierarchy. This extent is determined by truncation error estimates of key derivatives.

The first two (or possibly three) grid levels cover the entire computational domain. The mesh size is initially quite coarse in all directions in which adaptivity is to be prescribed. If unidirectional, e.g. ξ, adaptivity is to be implemented, then a preset stretched grid is used along the co-ordinate direction which is not adaptively refined (typically, η), and a coarse grid is employed in the ξ direction. Each multigrid level comprises several subdomains, which derive part of their topology from the subdomaining pattern of the coarser predecessor. Within each subdomain, of a given multigrid level, the refinement is specified independently, although the refinement criterion is the same for each co-ordinate direction. Thus, each subdomain of a multigrid level can act as a parent for a subdomain or for subdomains of the next finer multigrid level. If, at a given multigrid level, a particular subdomain is refined in only one direction, e.g. η, then on subsequent multigrid levels, further refinement within this subdomain is performed only in the η direction. A similar strategy is adopted for the ξ direction. Only subdomains that result from refinement of a parent subdomain in both the ξ and η directions require further decomposition according to the direction-selective refinement specifications.

5. REFINEMENT STRATEGY

In most adaptive gridding methods, on any grid level, an estimate of the truncation error of the discretized system of equations is used to identify those regions that require finer grid resolution.[1,6] The overall truncation error estimates, however, do not provide information on the specific direction(s) that require refinement. Therefore, for regions requiring improved resolution, the grid is refined in both directions, even though only one co-ordinate gradient may be significant. In order to achieve directional refinement adaptivity, it is necessary to monitor the truncation error of selected gradients or derivatives. For the two problems considered herein, the truncation error in the pressure and vorticity gradients, e.g. p_ξ and $u_{\eta\eta}$, are monitored in order to define the regions that require refinement in ξ and η directions, respectively.

The truncation error estimate is obtained from the solution on two successive grids of the multigrid hierarchy. In order to determine the truncation error in a ξ (and/or η) derivative, a finer grid must be used in the ξ (and/or η) direction. Thus, truncation error estimates would be available only for those points that are on the fine grid. Although the p_ξ and $u_{\eta\eta}$ terms are the key derivatives for the present analysis, the truncation error of these terms alone will not suffice to ensure that uniform accuracy is achieved throughout the flow domain. The global truncation error for the full discrete system of equations is monitored for this purpose. The truncation error in the global discrete system provides a good 'stopping' criterion for the refinement process. A second approach which has also been tested is grid convergence of the solution itself. The maximum change in the solution is monitored and a tolerance is set for that quantity, e.g. 10^{-2}.

Two types of adaptive calculations are performed for the geometries considered herein.

(a) One-dimensional adaptive calculation or semicoarsening multigrid, with adaptivity in the ξ direction only and with a preset stretched (adapted) η grid.
(b) Two-dimensional adaptive calculation, in which the refinement is automated in both directions and uniform grids are used in each subdomain. Grid stretching is not applied since the grids change discretely from subdomain to subdomain. Thus, the overall grid is heavily graded.

The underlying procedure for both the methods is identical. However, they are described separately for clarity.

5.1. *One-dimensional (semicoarsening) adaptivity*

The solution is first obtained on a grid that is coarse in the ξ direction, but stretched in the η direction. This means one co-ordinate direction is well-resolved whereas the other is not. The grid is then refined over the entire domain (in ξ), and an improved solution is obtained. From these two solutions, the truncation error of the key derivative, e.g. p_ξ, and of the global discrete system is estimated using a Richardson extrapolation procedure. For example, the truncation error in the p_ξ term is calculated as

$$\tau = (p_{i+1} - p_i)/\Delta\xi - (p_{m+1} - p_m)/(2\Delta\xi).$$

Here i represents the count on the fine grid and m represents the count on the coarse grid for the same physical location. Hence, the distance between points i and $i+1$ is half the distance between m and $m+1$. Typically, only one subdomain results for a given multigrid level. This decreases in extent as the grid level increases. The extent not only decreases in the ξ direction, but also in the η direction. This clearly identifies the region influenced by the pressure interaction. For the problems considered herein, the significant flow gradients in ξ are centred around a region $|\xi| < \xi_0$. For more complicated flows, it is possible that even with a one-dimensional adaptive calculation, several disjoint subdomains will be necessary to significantly improve the efficiency relative to non-adaptive calculations.

5.2. *Two-dimensional adaptivity*

For the two-dimensional adaptive calculation, different regions will have different refinement requirements in the two co-ordinate directions; therefore, it is necessary to define regions that have disparate grid requirements. Subdomains requiring refinement in the η direction, in the ξ direction, or in both the (ξ, η) directions, are identified using the truncation error estimates for the key derivatives. As a result, a given multigrid level will typically have more than one subdomain. Once again, p_ξ is used to identify regions for ξ refinement and $u_{\eta\eta}$ is used as the

parameter to identify regions where η refinement is required. The regions that require refinement in the ξ and η directions are identified independently. The intersection of the two regions is refined in both directions and the remaining portions are refined in one of the two co-ordinate directions. Different grid sizes are used in different regions; however, within each subdomain uniform grids are specified. This procedure is applied on the third and higher levels of the multigrid hierarchy. The calculation proceeds with intergrid multigrid transfers. On convergence, the truncation error estimation process is repeated using the Nth and the stored $(N-1)$th multigrid level grid solutions.

When two-dimensional adaptivity is implemented, the grid that is generated at each multigrid level reflects very clearly the flow physics. In addition, the procedure identifies those regions that need further refinement in the two co-ordinate directions. For example, for the finite flat-plate geometry, the gradients in the normal direction are larger near the trailing edge than in other regions. This indicates that the scaling of the flow field in the region near the trailing edge, i.e. the triple deck, is different from simple boundary layer behaviour. For a separated flow calculation, e.g. the flow past a trough geometry, it was found in Reference 7 that the region where η refinement was required was much larger in the vicinity of the separation bubble than elsewhere along the boundary layer. This reflects the increase in boundary layer thickness at separation. In this manner, the adaptive grid generation procedure automatically provides an insight into the nature of the flow field, and the result is an appropriate mesh that resolves such behaviour.

The two-dimensional adaptive calculation, by virtue of using only uniform grids, recovers the advantages of the standard full multigrid procedure. The procedure starts in a full multigrid mode, in which the grid is coarsened in both directions when restricting to a coarser level. As different types of refinement (ξ, η, or both) are specified in different subdomains, the choice of the multigrid mode is fixed. In subdomains in which the grid is only refined along one co-ordinate direction, the semicoarsening mode is employed for the appropriate direction. Thus, the standard full multigrid mode is employed wherever possible. This is described in greater detail in the section on multigrid implementation.

The choice of a refinement criterion is flexible. In Reference 8 a strategy by which the tolerance is changed so that a certain fixed percentage of points are refined during a refinement process is adopted. Two types of refinement criteria are specified herein. In one procedure, a tolerance is set directly for the raw truncation error and, in the other, a tolerance is set for the local truncation error normalized by the maximum value in the flow domain. The results obtained through the two methods were quite similar. The regions that require refinement in the respective direction(s) are identified through these tolerances. The choice of tolerance is constrained by the gradient of the truncation error parameter. Typically, for the problems considered herein, the pressure field is fairly smooth and as such the truncation error in p_ξ also varies smoothly. On the other hand, the vorticity gradients are quite sharp. A much smaller tolerance has to be chosen if the standard deviation of this truncation error field is large. For a normalized p_ξ truncation error, typical tolerances in the range 0·01–0·05 are specified. But for the vorticity, values in the range 0·0001–0·0005 are required to achieve the required level of accuracy.

6. SOLUTION PROCEDURE AND MULTIGRID IMPLEMENTATION

A global pressure relaxation procedure[9] is generally used to solve the discrete system of equations. This is essentially a line SOR procedure. The RNS discretization necessitates that the pressure field, in unseparated regions, or both the pressure and velocity, in separated regions, is relaxed. The solution is marched from inflow to outflow, with an initial guess required for the pressure. The multigrid method inherently allows the provision of a good initial guess as the

calculation progresses from the coarser to the finer grids. This further enhances the convergence process.[10]

For the RNS system of equations, a semicoarsening multigrid procedure has been presented in Reference 11 to accelerate the convergence of the global pressure relaxation (in ξ) procedure.[9] A von Neumann analysis of the linearized discrete form of the RNS system shows that the rate of convergence of the global relaxation procedure is dictated by the maximum eigenvalue λ, as given by

$$\lambda \sim 1 - c_1 \pi^2 (\Delta\xi)^4 N_\xi^2 / \eta_M^4,$$

where c_1 is a constant of $O(1)$, N_ξ is the number of stations in the ξ direction, η_M is the normal boundary location, and $\Delta\xi$ is the axial step size. The convergence rate is significantly improved if the extent of the domain in the two directions is reduced or the number of stations N_ξ is increased. The current multigrid domain decomposition procedure, in effect, decrease λ by reducing η_M whenever a fine $\Delta\xi$ is specified. Therefore, on very fine grids, where $\Delta\xi \ll 1$, N_ξ or η_M is reduced and λ is comparable to the values of the coarser grids.

In the present method, each multigrid level is comprised of one or more subdomains. Within each subdomain, the global pressure line relaxation procedure is applied. The use of a direct solver in certain strong interaction regions has been considered in Reference 12. The calculation proceeds from inflow to outflow, and, through the various subdomains, passes through a number of interdomain boundaries. Boundary conditions are exchanged at subdomain interfaces during the marching procedure. A minimal amount of overlap is allowed for this purpose. This is explained in greater detail in Section 7.

In the present application, the multigrid method is implemented in a full approximation storage (FAS) mode. The global pressure relaxation technique reduces to a block SOR procedure in ξ for the pressure in attached flow regions, and for the pressure and upwinded velocities in reversed flow regions. At each station, an implicit, fully coupled tridiagonal system is inverted.

When highly stretched grids are used in η to resolve the boundary layer, the standard non-adaptive semi-coarsening mode of the multigrid method has been shown to be more effective than the standard full coarsening mode. In this mode, the streamwise grid alone is coarsened when the calculation shifts to coarser grids. The same η grid is retained. Significant gains in the overall effort have been achieved with this approach.[9] The full coarsening mode of the multigrid method is preferable, however, as coarsening in both directions significantly reduces the computational effort required to obtain coarse-grid solutions. The advantage of the full coarsening multigrid mode is restored with the two-dimensional adaptive calculation, as uniform grids are now considered in each of the individual subdomains, i.e. there are no stretched grids in any subdomain. Grid changes are reflected by the appearance of a new subdomain. The calculation is initiated with the full coarsening mode of the multigrid method. As finer multigrid levels develop through the refinement strategy, the semicoarsening multigrid mode is applied in those subdomains, for which refinement has been reduced to one direction relative to the parent subdomain.

The smoothing properties of the relaxation operator, as in all multigrid methods, play an important role in determining the overall performance of the procedure. In Reference 11 it was found that with semicoarsening, the standard discrete global pressure relaxation method does not provide very good smoothing properties. This is due to flux-split discretization procedure, which leads to significant over-relaxation of the pressure field. This is especially true for high *Re*, viscous/inviscid interacting flows. For inviscid flows the smoothing properties of the relaxation operator are found to provide smooth residual fields, even in the over-relaxation mode so that satisfactory performance of the multigrid method results. Some form of under-relaxation is

required in order to improve the smoothing properties of the pressure relaxation procedure. This is achieved by incorporating a source term, first introduced by Israeli and Rosenfeld,[13] in the ξ momentum equation. This leads to much smoother residual fields and improved representation on the coarser grids. With the additional source term, the expression for the pressure gradient $p_{\xi\xi}$, as given by (2), becomes

$$\Delta\xi^2 p_{\xi\xi} \approx p_{i+1}^{n-1} - (1+\varepsilon)p_i^n + p_{i-1}^n - (1-\varepsilon)p_i^{n-1}.$$

Therefore, $\varepsilon < 1$ corresponds to over-relaxation, $\varepsilon > 1$ corresponds to under-relaxation and $\varepsilon = 1$ is Gauss–Seidel relaxation. For $\varepsilon = 0$, which corresponds to the case of gross over-relaxation, the standard form of the pressure relaxation procedure is recovered and the residual field is insufficiently smoothed.

Although under-relaxation of the pressure field leads to slower asymptotic convergence on a single grid, the domain decomposition procedure alleviates this problem to a great extent. Since the truncation error in the p_ξ term is used to determine those regions requiring refinement in the ξ direction, subdomains in which the grid is refined only in the η direction will be influenced by a nearly converged pressure field. Therefore, it is possible to proceed with the calculation in these subdomains without any form of under-relaxation. Non-smooth residual fields are generated when the pressure field changes significantly from one multigrid level to the next. Regions that do not require refinement in ξ are those for which the pressure field is already grid-independent; these regions will have smooth residual fields and the source term or under-relaxation is not required.

In the present calculation, the one-dimensional adaptive calculation adds an element of subdomaining to the semi-coarsening analysis presented in Reference 11. Only portions of the global domain now require fine grid resolution in the ξ direction. For the two-dimensional adaptive calculation, the form of the multigrid algorithm can be different for different subdomains. The standard full coarsening mode (semi-coarsening mode) is employed in subdomains requiring bidirectional (unidirectional) refinement. One fine-grid work-unit is comprised of one sweep in each subdomain belonging to a given multigrid level. This includes the interdomain transfer processes. The decision to move the calculation back to a coarser grid is based on the rate of convergence on each subdomain. If the ratio of the residual norm between two successive global iterations, in any subdomain belonging to that multigrid level, falls below a certain value, then the calculation is restricted to the coarser level. Values ranging from 0·85–0·95 were used for the ratio of the residual norms. In flows with strong separation, it was necessary to use a higher value (closer to unity). This ensures that more iterations are performed on the fine grid before restriction to the coarse grid. In this way the separated flow regions can become reasonably established on the fine grid. If the calculation is switched to the coarser grid prematurely, it is found that the coarse-grid calculations can become unstable.

A criterion for shifting back to the fine grid is also required. The fine-grid solution is not corrected until the residuals of the coarse-grid subdomains are all reduced to a value one order of magnitude lower than the maximum residual for all subdomains of the finer level.

The multigrid components are summarized as follows:

(a) Relaxation is given by $u_n^k = S^k u_{n-1}^k$, where S^k is the global pressure relaxation operator and u^k on convergence satisfies $L^k u^k = f^k$. Here k represents the present or finest multigrid level and n represents the iterate. $L^k u^k = f^k$ is the discrete approximation of the continuous problem $Lu = f$.

(b) Restriction to the coarse grid is represented by the following equations:

$$L^{k-1} u^{k-1} = I_k^{k-1} r_n^k + L^{k-1} \hat{I}_k^{k-1} u_n^k$$

for points on the coarse grid which lie within the fine grid and

$$L^{k-1}u^{k-1}=f^{k-1}$$

for points on the coarse grid that lie outside the extent of the fine grid. Here $r_n^k=f^k-L^k u_n^k$. I_k^{k-1} and $\hat{I}_k^{k-1}$ are fine-to-coarse grid transfer operators. The full-weighting operator recommended in Reference 14 is used to transfer the residuals, and the solution is restricted by using a simple injection operator.

(c) Prolongation or correction, where the fine-grid solution is corrected with the solution from the coarse-grid modified problem is represented by

$$u_{n+1}^k=u_n^k+I_{k-1}^k(u^{k-1}-I_k^{k-1}u_n^k),$$

where I_{k-1}^k is a coarse-to-fine interpolation operator.

In both the adaptive semi-coarsening multigrid calculation and the two-dimensional adaptive calculation, the fine multigrid levels are of limited extent and these lead to interior boundaries. Boundary conditions for these subdomains are prescribed from the coarser-grid values. Therefore, the multigrid transfer operations act in a manner similar to that of interacting boundary layer procedures, i.e. the transfer processes involve communication between the viscous and inviscid layers. The multigrid restriction process yields information about the viscous layers (where most of the η refinement is performed) to the inviscid regions. Similarly, the prolongation or correction process yields information about the inviscid region, e.g. the location of the outer boundary, etc., to the viscous layers. This exchange of information provides accurate solutions in all subdomains on all grid levels.

It should be noted that in the present calculation, the multigrid transfer operations play a dual role. In addition to accelerating the convergence of the relaxation procedure, they also provide information from finer to coarser grids and, thus, improve the accuracy of the solution in regions of the coarser grids where refinement was not required. The second term in the multigrid restriction process acts as a truncation error injection term and improves the discrete approximation on the coarse grid. Thus, on the coarser grids, instead of solving $L^{k-1}u^{k-1}=f^{k-1}$ everywhere, we solve $L^{k-1}\hat{u}=\tau$ in the portion of the domain where $\tau=L^{k-1}\hat{I}_k^{k-1}u^k$. This is closer to the continuous problem $Lu=f$. Here L is the continuous counterpart of the discrete operator L^{k-1} and u is the exact solution to the continuous problem; u^{k-1} is the exact solution to the discrete problem and $\hat{u}$ is the improved solution due to the modified right-hand side of the discrete aproximation.

In standard non-adaptive multigrid methods, since the fine-grid extent is the same as that of the coarse grid, once the fine-grid residual norm reduces below a certain tolerance, e.g. 10^{-4}, the calculation can be terminated. But in adaptive multigrid calculations, a second tolerance has to be set for the corrections from the coarse to the fine grid. Small values for the residuals on the fine grid do not guarantee an accurate solution; the boundary condition imposed on the fine grid could still be in error. Since some of these boundary conditions are updated during the multigrid prolongation process, it becomes necessary to set a tolerance for the correction $I_{k-1}^k(u^{k-1}-I_k^{k-1}u_n^k)$. Typically, values ranging from 10^{-3} to 10^{-4} are used for these computations.

For the backstep geometry, the explicit deferred corrector (DC) is included, for interior points, as part of the discrete system. Since the DC terms are added explicitly, they are treated as known functional fields on the right-hand side of the governing equations. The effect of the DC becomes significant only on finer grids. The DC is included when the solution attains a specified level of convergence. The introduction of the DC at interior points changes the character of the governing

system. The marching procedure now requires some under-relaxation. Therefore, the solution is allowed to converge to a certain tolerance level, e.g. 10^{-3} without the DC terms and then the DC, which is evaluated using this base solution, is treated as a known explicit term on the right-hand side of the equations. The decision as to when the DC is to be included is determined by monitoring the DC truncation error. If this changes more than a prescribed tolerance at any given grid level, then the DC is retained. If the primary DC terms are included implicity, as for a direct solver, this becomes unnecessary.

7. INTERDOMAIN TRANSFER OF BOUNDARY CONDITIONS AND CONSERVATION AT GRID INTERFACES

For a given subdomain, the following boundary conditions are prescribed:

$$u=v=0 \text{ at } \eta=0;\ u=1,\quad p=0 \text{ at } \eta=\eta_{\max};\ p_\xi=0 \text{ at } \xi=\xi_{\max};$$

u and v are free-stream values at $\xi=0$. Here $\eta_{\max}$ and $\xi_{\max}$ represent the boundaries of the global domain. If a subdomain has its outflow at some $\xi<\xi_{\max}$, then the boundary condition, on pressure, changes from Neumann to Dirichlet type at this location. If the lower boundary of a subdomain is at some $\eta>0$, then non-zero velocities from the coarser grid are prescribed. In time-dependent, characteristic-based, Navier–Stokes computations, that use locally embedded grids, boundary conditions are required for all variables, i.e. u, v and p at all boundaries. Special care has to be taken to ensure that mass conservation is not violated locally or globally at the interfaces and, in particular, in a multigrid procedure. In the present pressure-based flux-split RNS formulation, this difficulty does not occur as the normal velocity (v in η, or u in ξ) is not prescribed at the upper or lower or local subdomain outflow boundaries. Only the tangential component u is prescribed at the upper interface or interdomain boundary. The pressure-based flux-split box differencing allows the calculation of the normal velocity at the outer boundaries and the pressure at the body surface. The normal velocity is computed from the continuity equation (1a) at every grid point and, therefore, mass conservation is automatically satisfied on all levels and for all subdomains. This eliminates the need for special interpolation formulae to ensure conservation of mass, when the boundary conditions are prescribed from the coarse-grid solution. Thus, weak instabilities, that arise when such methods are applied without satisfying mass conservation, do not appear in the present method. Direct evaluation of the pressure at inflow or lower boundaries also eliminates the need for special pressure boundary conditions.

The calculation is performed sequentially rather than in parallel in the various subdomains. As such, the boundary conditions at the inflow and outflow stations for each subdomain are updated with the latest available values. The overlap allowed in the subdomaining process follows the following rules:

(a) The last station of any subdomain, which is at some $\xi<\xi_{\max}$, coincides with the first station of the subdomain to its right (if one exists), where the pressure is computed.
(b) Similarly, the inflow station of any subdomain, which is at some $\xi>0$, coincides with the last station on the subdomain to its left (if one exists), where the velocities are computed.
(c) If the inflow station or the outflow station of a given subdomain coincides with the physical boundaries of the global flow field, then the boundary conditions of Section 3 are applied.
(d) If there are no subdomains to the right for the cases in (a), or if there are no subdomains to the left for the cases in (b), then these boundaries are updated using coarse-grid values during the multigrid prolongation process.

In the vertical direction, no overlap is necessary. If a subdomain has only one of its horizontal boundaries in common with another subdomain, then updating the boundary conditions along this edge, after one sweep in all subdomains, leads to iterative divergence on this subdomain. This influence gradually filters through to other subdomains. If these boundaries are updated through the multigrid transfer process, then the calculation is convergent. This reflects the fact that updating just one boundary after each sweep, with the other three updated only during the multigrid transfer process, leads to an inconsistency. This constraints the variables from adjusting to changes that occur dynamically, i.e. as the solution evolves in the various subdomains. It is possible to perform the calculation in parallel on a given multigrid level that comprises more than one subdomain. This requires that the updating process be lagged by one sweep. This approach has not been implemented in the present computations.

8. RESULTS AND DISCUSSIONS

All the calculations presented herein are initiated with uniform-flow initial conditions on the coarsest grid. On finer grids, the interpolated coarse-grid solution fields are used as the starting values. The better the initial guess, the better is the convergence to the final solution.[10] The multigrid framework provides this element in a natural fashion. For both of the examples considered herein, Reynolds number continuation is not necessary in order to initiate the calculation at any prescribed value of Re.

Example 1: Flow past a finite flat plate

Figure 1 depicts the grid obtained with the one-dimensional adaptive calculation. Note that the finer grids zoom in around the trailing edge, which is located at $x = 1{\cdot}0$. The finer grids also reduce in extent in the η direction, even though adaptivity is required only in the ξ direction. Each

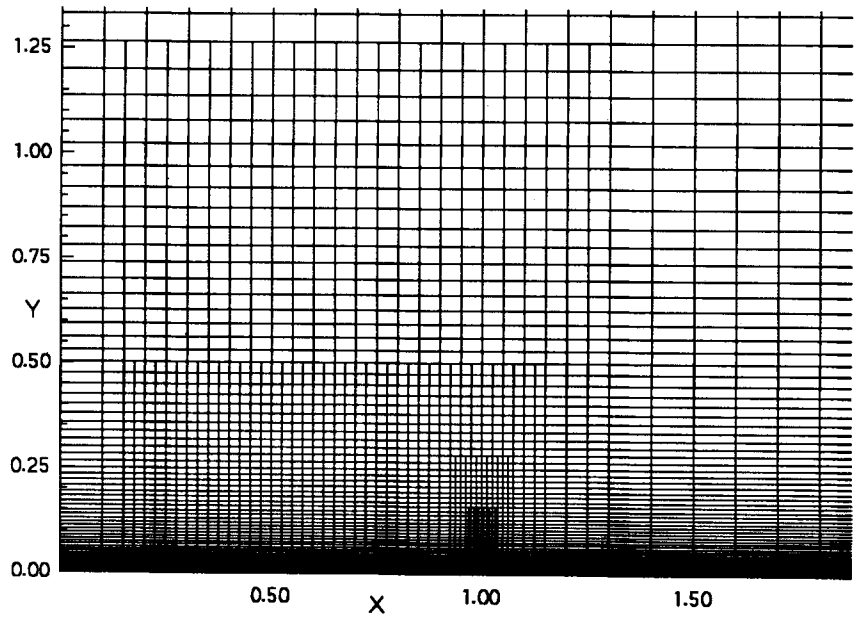

Figure 1. Multigrid levels (one-dimensional semicoarsening adaptivity); trailing-edge flow ($Re = 10^5$)

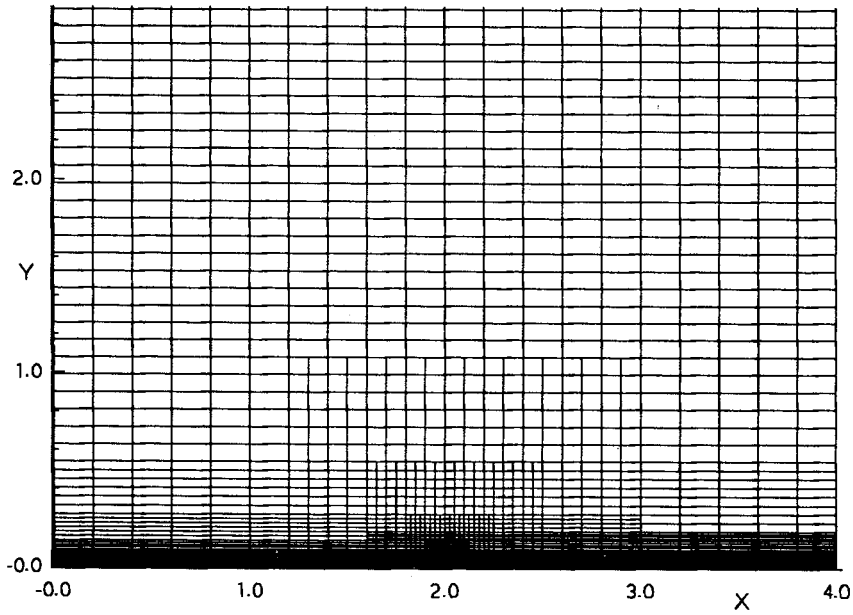

Figure 2. Multigrid levels (two-dimensional adaptivity); trailing-edge flow ($Re = 10^5$)

multigrid level contains only one subdomain that requires further refinement on subsequent levels.

Figure 2 shows the composite grid obtained with two-dimensional adaptivity. (The x-co-ordinate is scaled by a factor of two.) Note that, within each subdomain, uniform grids are prescribed. Figure 2 is an overlay of seven multigrid levels, each comprising several subdomains. In each level it is found that the subdomain, in which refinement in both directions is performed, is centred around the trailing edge. This validates the refinement strategy. Both of these adaptive computations, i.e. semi-coarsening and two-dimensional adaptivity, are compared with non-adaptive semi-coarsening multigrid calculations. A uniform fine grid in ξ and a stretched η grid is prescribed. The grid stretch factor for the latter is chosen by specifying the minimum and maximum $\Delta\eta$ values and the location of η_{max} as applied in the two-dimensional adaptive study. The same η grid is employed for adaptive semi-coarsening. Figure 3 shows a comparison of the pressure coefficient C_p for the three calculations. Note that there is good agreement in the pressure variation and, in particular, the predicted peak pressures. Table I summarizes the computer memory and CPU requirements. These are given as percentages of the non-adaptive semi-coarsening calculation. Note that the memory requirement or the one- and two-dimensional adaptive calculations are quite similar. This signifies that the specified η stretching for the one-dimensional calculation is quite good.

The adaptive grid of Figure 2 shows that the interaction near the trailing edge is of finite extent. From asymptotic triple-deck theory, three layers with different length scales can be identified,[15] i.e. a lower viscous rotational deck of $O(Re^{-5/8})$, a middle inviscid rotational deck of $O(Re^{-4/8})$ and an upper inviscid irrotational deck of $O(Re^{-3/8})$. Since the vorticity is zero in the upper deck, and since vorticity is the monitored parameter for refinement in the η direction, η refinement should not be necessary in this region. The grid obtained from the two-dimensional adaptive calculation displays this result clearly. In each multigrid level, there is a region away from the

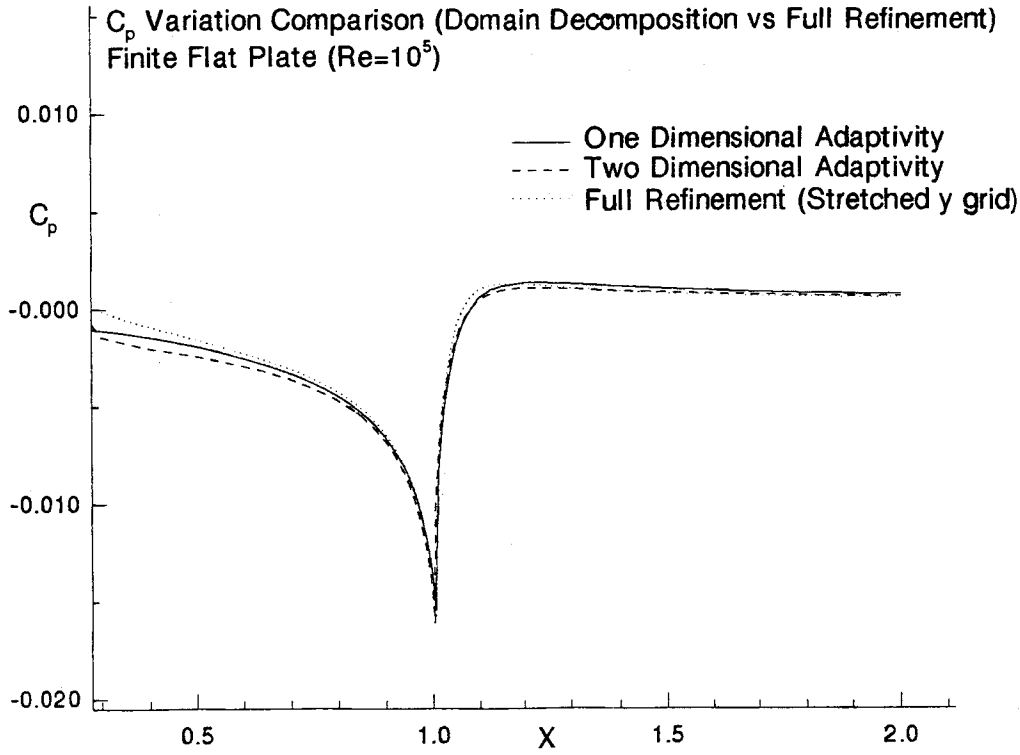

Figure 3. Comparison of C_p variation; trailing-edge flow ($Re = 10^5$)

Table I. Comparison of computer resource requirements for trailing-edge flow

Aspect	Two-dimensional adaptive (%)	One-dimensional adaptive (%)	Full refinment with stretched η grid (%)
CPU	18·03	15·10	100·0
Memory	12·90	13·22	100·0

body that is, in fact, only refined in the ξ direction. This region, in the finest multigrid level, represents the extent of the upper inviscid deck. Estimates for the extent of the other two decks can also be obtained from the grid structure. In more complicated flows, e.g. turbulent flow past the same geometry, for which analytical methods cannot be easily developed, the necessary resolution in appropriate regions will result automatically. A form of numerical asymptotic analysis results.

Example 2. Internal flow in a backstep channel

In this type of flow, which is dominated by rather large recirculation regions, it is still possible to initiate the calculation with uniform flow conditions for all Re (based on step height) values considered herein. Even for the relatively difficult case of $Re = 800$, which leads to two separation

bubbles, one on each wall, some form of Reynolds number continuation, as applied in several other reported NS solvers (e.g. Reference 16), is not required for the present calculations.

The reattachment length (X_r/h) for the primary recirculation zone is compared in Table II for a laminar Reynolds numbers. Comparisons are given for the present method, full refinement computations using the standard non-adaptive multigrid method, and earlier calculations by other researchers.

For adaptive refinement in the η direction the truncation error is scanned from the wall towards the outer boundary. For external flow, the vorticity gradient decreases exponentially and a thin layer near the wall, where refinement is maximum, can be identified. This region is specified by fixing the upper boundary at the furthest point, or the largest η value taken over all ξ locations, that satisfies the previously given truncation error tolerance. For internal flows, boundary layers where refinement in η is required, exist at both boundaries of the normal or η domain. Since the number of grid points that are necessary for resolving the gradients in the normal direction is moderate for the flow in a backstep channel, no attempt was made herein to adaptively refine in the η direction. However, a full two-directional multigrid procedure is applied wherever ξ refinement is necessary.

Table III displays the computer resource requirements for the backstep channel calculation. For each Reynolds number, the CPU and memory requirements are shown as percentages of the corresponding non-adaptive calculations. Note that as the Reynolds number increases from $Re=133$ to $Re=600$, the number of grid points required to resolve the flow field increases. This is expected, as the size of the separation bubble increases with Reynolds number. The number of required grid levels, as well as the finest mesh size for all Reynolds numbers up to $Re=600$, is identical in each direction. A total of five multigrid levels are defined for all Reynolds numbers up to $Re=600$; however, the subdomain extent for each multigrid level did vary for different

Table II. Comparison of reattachment length (X_r/h)

	Present calculations		Ferziger	Caruso	Sotiropoulos
Re	Adaptive	Non-adaptive	(Reference 1)	(Reference 17)	(Reference 18)
133	3·88	3·88	4·00	3·90	3·68
267	6·50	6·50	6·50	6·50	6·34
400	8·64	8·64	8·70	8·80	8·80
600	11·0	11·0	10·7	10·8	11·26

Table III. Comparison of computer resource requirements for backward-facing step

Aspect	*Re*	Adaptive/Non-adaptive calculation
CPU	133	35·49%
	267	36·15%
	400	46·23%
	600	50·40%
Memory	133	30·80%
	267	37·44%
	400	41·49%
	600	47·49%

Reynolds numbers. The extent of the finer grids is governed by the size of the recirculation zone, which increases as the Reynolds number is increased. For the $Re=800$ case, six multigrid levels are required, as the change in the solution from level 4 to level 5 is still significant. The CPU and memory gains for the $Re=800$ case are comparable to the $Re=400$ case. An increase in computational time, similar to that in memory requirements, is observed for the examined range of Reynolds numbers. This is due to the increase in the number of grid points. The time required for the full refinement calculation shows only a marginal change, as the Reynolds number is increased from 133 to 600. This is entirely due to the changing nature of the flowfield, since the same grid pattern was developed. When velocity relaxation is increased, i.e. as increasing reversed flow regions exist, the convergence rate slows down and additional iterations are required.

The location of the outflow boundary and the non-reflectivity of the outflow boundary condition are major aspects of this study. Since the adaptive multigrid domain decomposition procedure is initiated on very coarse grids, it is possible to place the outflow boundary very far downstream and still perform the calculations efficiently.

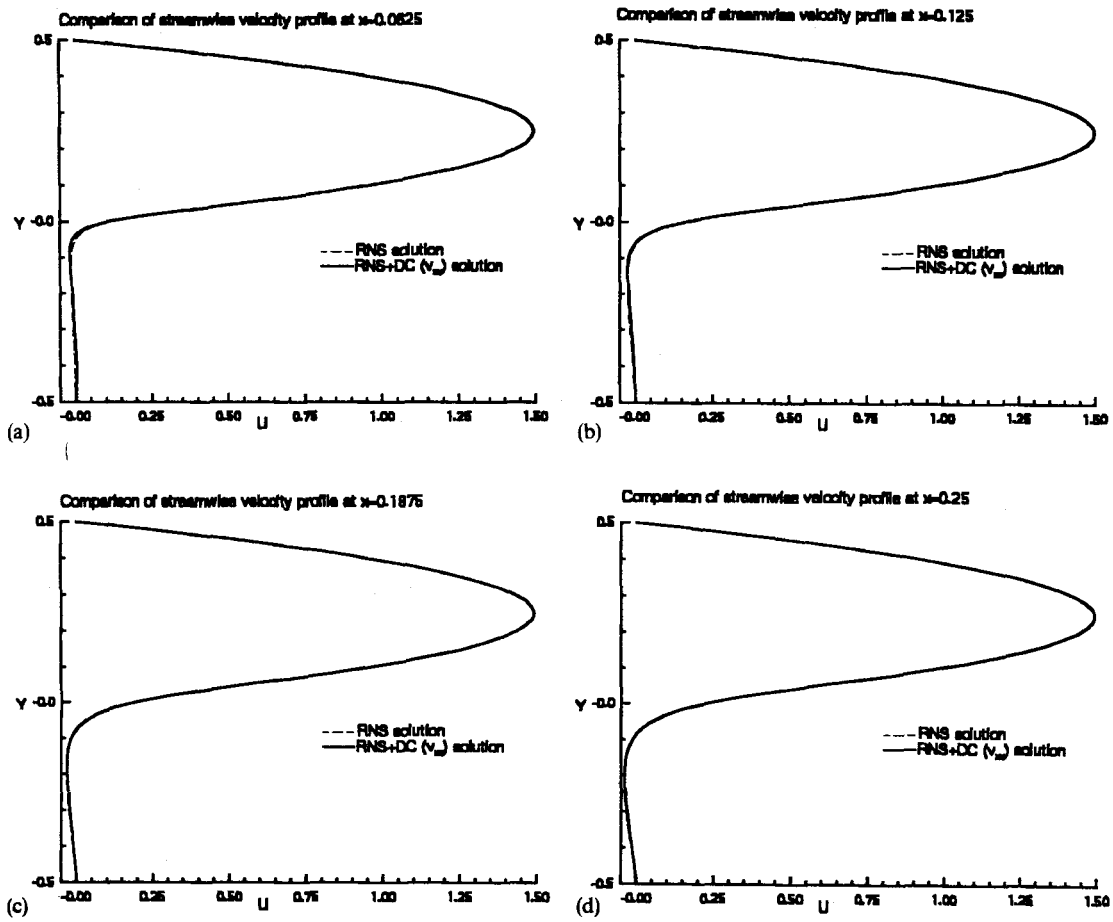

Figure 4. Comparison of streamwise velocity profiles at (a) $x=0{\cdot}0625$, (b) $x=0{\cdot}125$, (c) $x=0{\cdot}1875$ and (d) $x=0{\cdot}25$; backward facing step ($Re=600$)

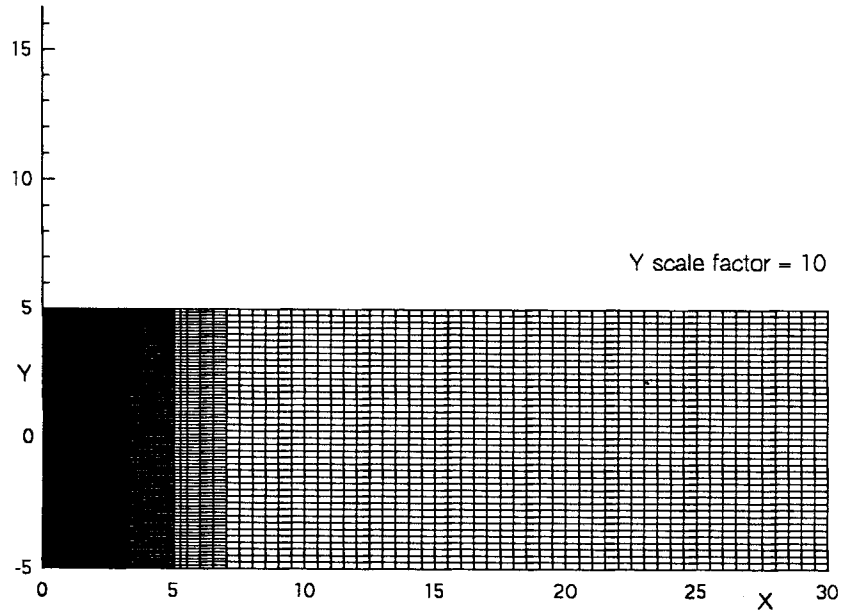

Figure 5. Multigrid levels; backward-facing step ($Re = 100$)

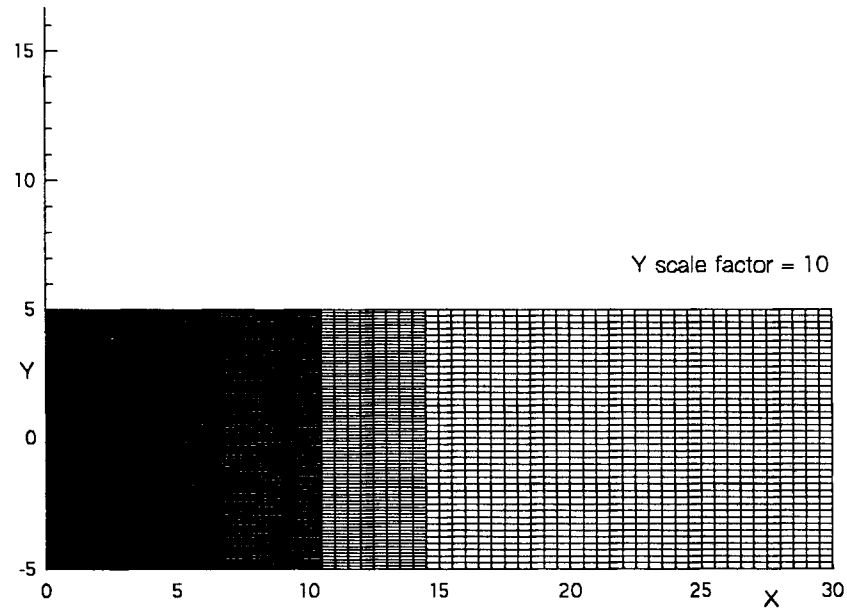

Figure 6. Multigrid levels; backward-facing step ($Re = 267$)

The solutions at the outflow are in almost perfect agreement with the analytical fully developed flow values. Although finer grids only appear in and near the reverse-flow regions the influence of the outflow boundary location is propagated through the multigrid transfers to and from the coarser grids.

For the laminar backstep, the DC terms neglected in the RNS approximation are included after obtaining the base solution for the reduced system. For this geometry, the vertical wall near the step corner is a region where the full Navier–Stokes terms might be important. Along the vertical wall, the v_{xx} term represents the vortical or diffusive boundary layer influence. It is found that the inclusion of the DC term does not have significant quantitative influence on the solution. But there is a significant qualitative change in the solution very near the step corner. Figures 4(a)–4(d) show the comparison of the streamwise velocity profile at four successive stations from the step corner for $Re = 600$. Note that at the first station, the streamwise velocity component is actually positive when the DC is included. This represents a counter-rotating secondary vortex within the large primary recirculation zone. The difference in the profiles diminishes further away from the step corner. The reattachment length remains unchanged. As expected, the influence of the DC is confined to a small region near the corner. The lowest-order RNS approximation is quite accurate throughout the flow field.

Figures 5 and 6 depict the grid generated by the adaptive strategy for the $Re = 100$ and $Re = 267$ cases, respectively. Note that the fine-grid region for the $Re = 267$ case is larger in extent as compared to the $Re = 100$ case. This reflects the fact that the recirculation region for the $Re = 267$ case is larger than the $Re = 100$ case. Note that near the outflow a grid as coarse as $\Delta x = 0{\cdot}5$ can be used and still the solution accuracy is not lost. Thus, the adaptive strategy makes full use of the fact that far downstream the flow is fully developed and the streamwise derivatives are very small. Thus, the truncation error even on a coarse grid is acceptably small.

The non-reflectivity of the Dirichlet pressure outflow boundary condition was tested for the $Re = 800$ case by performing the calculation on computational domains of three different lengths.

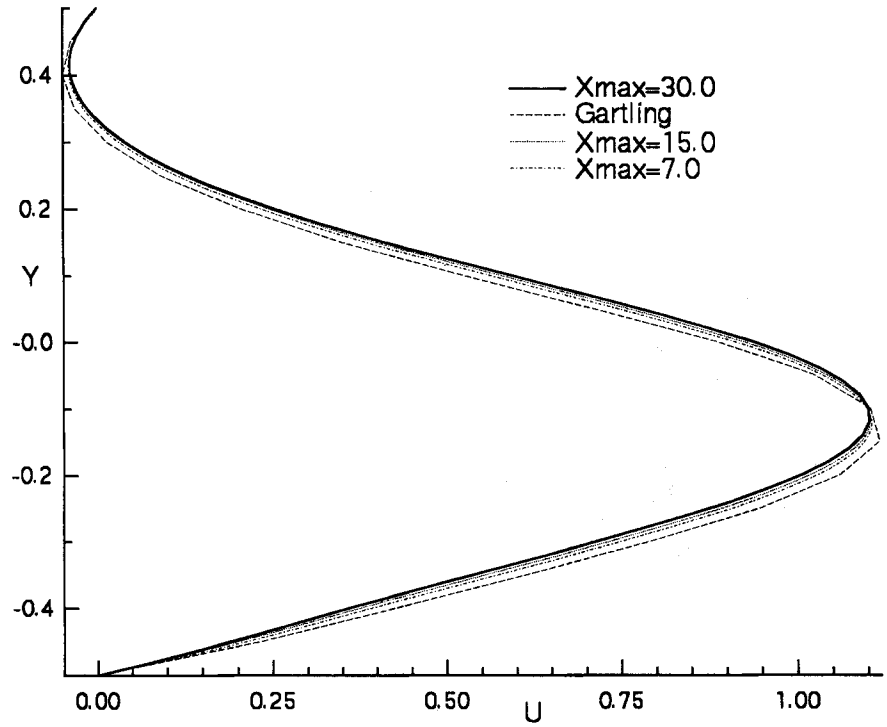

Figure 7. Comparison of streamwise velocity profile at $x = 7{\cdot}0$; backward-facing step ($Re = 800$)

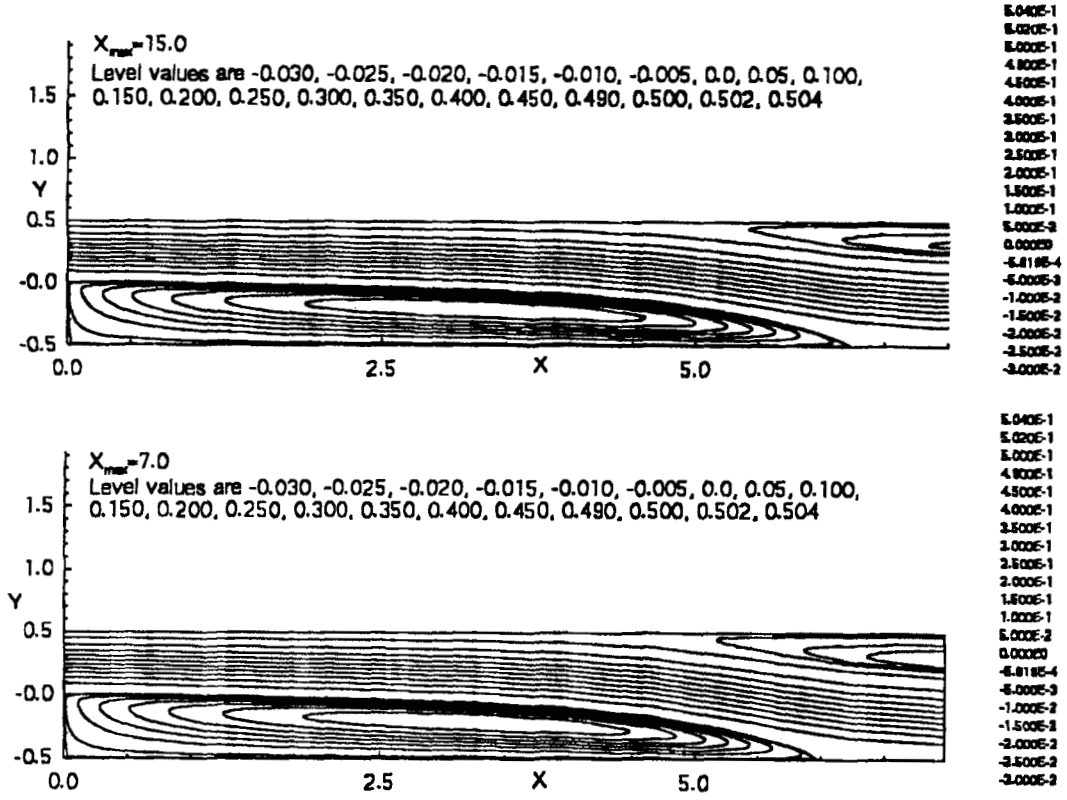

Figure 8. Comparison of streamfunction contours; effect of outflow boundary location; backward-facing step ($Re = 800$)

It was found that moving the outflow boundary inward did not have a significant effect on these solutions. Comparisons of the streamwise velocity profile at $x = 7$ are given in Figure 7. The outflow boundary is located at $x = 7$, $x = 15$ and $x = 30$ and the results clearly indicate that the effect of the outflow location on the solutions is minimal. Similarly, Figure 8 shows a comparison of the streamfunction contours obtained by placing the outflow at $x = 7$ and $x = 15$. The two contour patterns are identical. Further results can be found in Reference 19.

A benchmark solution for $Re = 800$ is available in Reference 14. The present solution, which is obtained using the adaptive multigrid domain decomposition procedure, is compared with these results. In Figure 7, note that reverse flow also occurs on the upper wall. The appearance of this upper separation bubble is thought to introduce three-dimensionality into the flow, and, for this reason, there is some disagreement between the experimental results and all the numerical solutions. The present results, which are totally grid-independent, agree quite well with several of the other computations. Due to the very fine meshes that have been prescribed with the multigrid domain decomposition procedure, the residuals and truncation errors are quite small and, therefore, these solutions are considered to be highly accurate. Figure 9 shows the nearly fully developed streamwise velocity profile at $x = 15{\cdot}0$. The grid spacing in the streamwise direction at this station is 0·5 but the agreement in the solutions is still very good. Figure 10 depicts a comparison of vorticity profiles at $x = 7{\cdot}0$. Figure 11 compares the normal velocity profile at $x = 15$, and Figure 12 depicts the vorticity contours obtained by placing the outflow boundary at

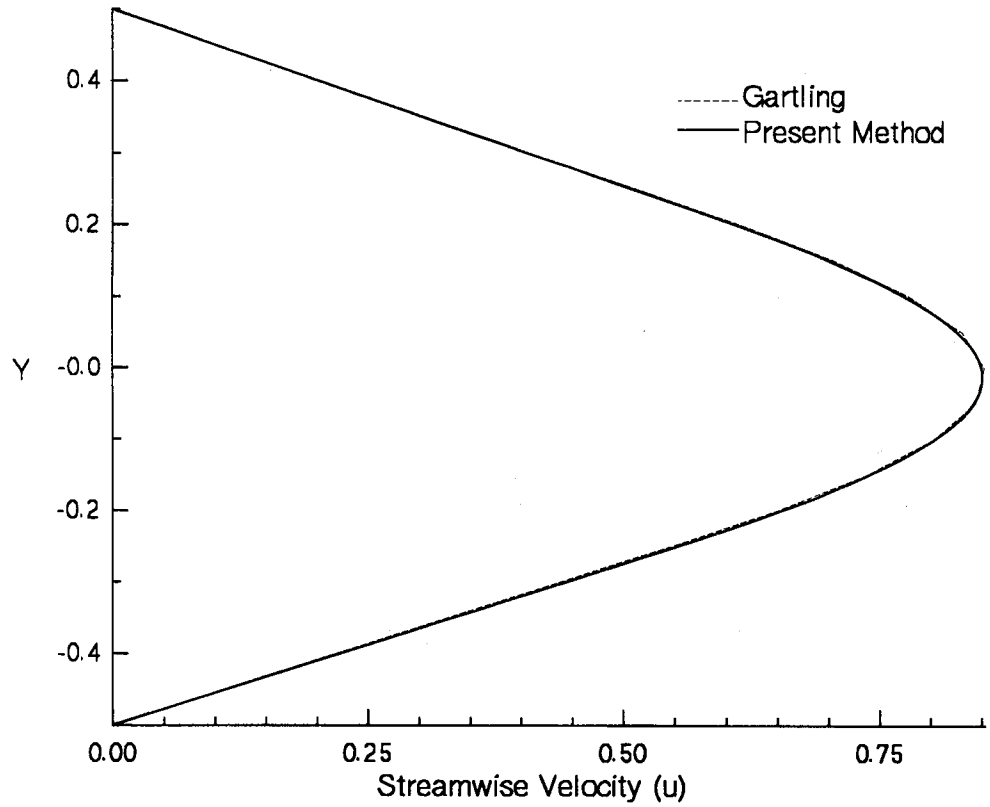

Figure 9. Comparison of streamwise velocity profile at $x = 15{\cdot}0$; backward-facing step ($Re = 800$)

$x = 7{\cdot}0$. Once again, the same contour levels as in Reference 16 are used. The agreement throughout is excellent. Figure 13 shows the convergence history on the finest grid level for the $Re = 800$ case. The number of iterations on coarser grids is scaled according to the number of grid points in each level.

8. CONCLUSIONS

An adaptive multigrid domain decomposition method has been used to compute incompressible laminar flows efficiently. A low-order RNS system of equations, a pressure–velocity non-staggered grid solver, a strong mass conservation technique at interfaces and the outflow boundary, and a pressure-based flux-split discretization are the key features of the procedure. Significant gains in computer resources have been achieved, compared to standard non-adaptive methods. Good agreement is obtained between the present solutions, standard non-adaptive full refinement computations and other published results. The computational cost is several times smaller than that required by most other NS methods.[16] For example, the CPU required for a backstep channel calculation, with Reynolds numbers in the range 100–600, varies between 5 and 10 min on an IBM 320 RISC/6000 workstation. For $Re = 800$, one additional multigrid level is added to ensure grid independence, and the CPU required is increased to approximately 30 min. All solutions are initiated with uniform flow approximations and Reynolds number

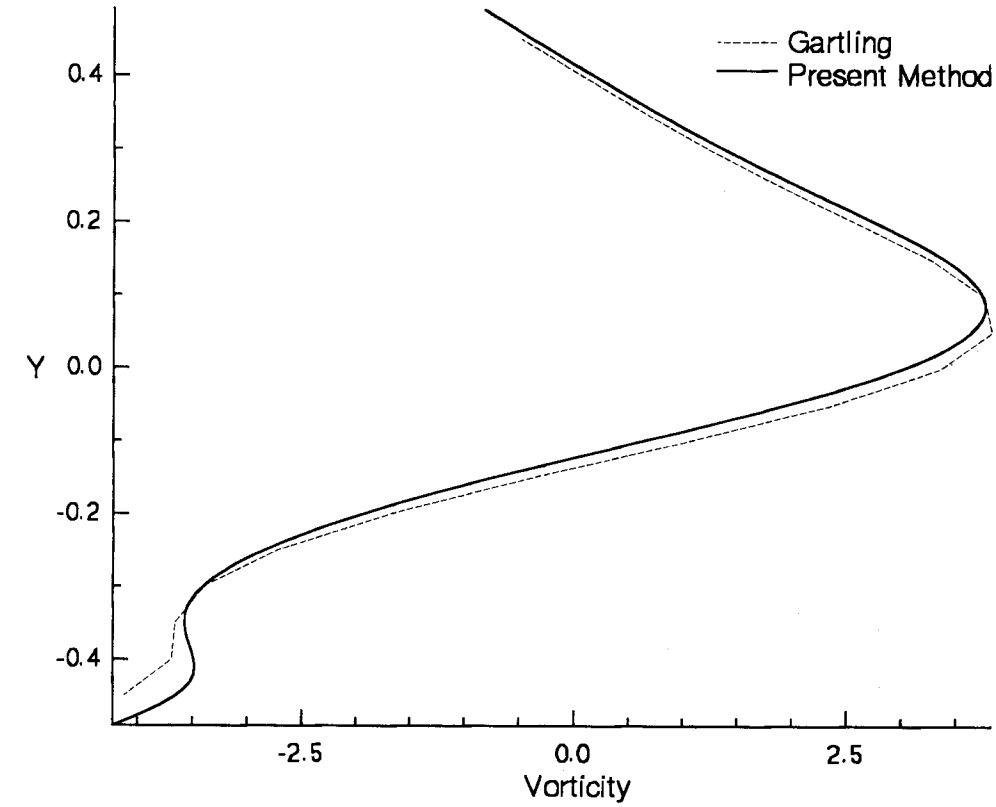

Figure 10. Comparison of vorticity profiles at $x = 7{\cdot}0$; backward-facing step ($Re = 800$)

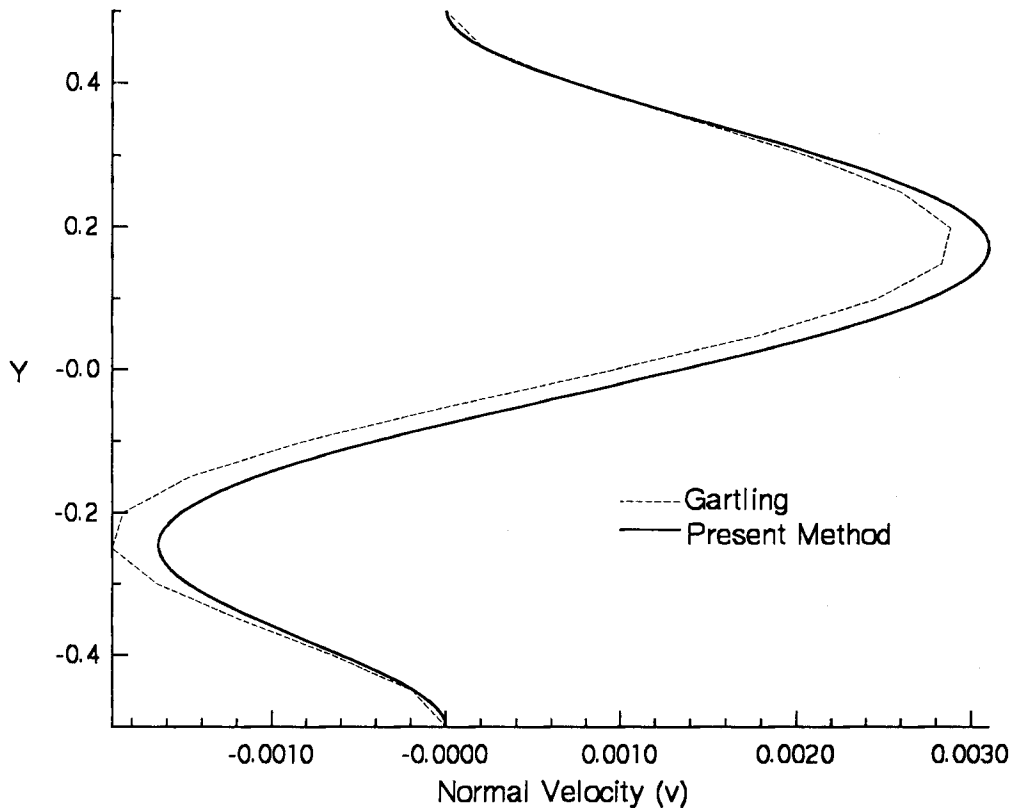

Figure 11. Comparison of normal velocity profiles at $x = 15{\cdot}0$; backward-facing step ($Re = 800$)

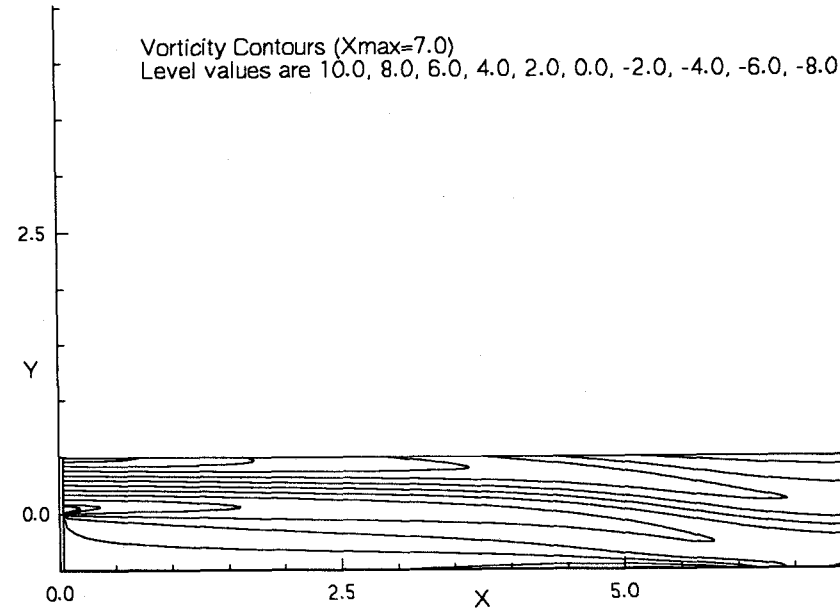

Figure 12. Vorticity contours (outflow at $x = 7{\cdot}0$); backward-facing step; ($Re = 800$)

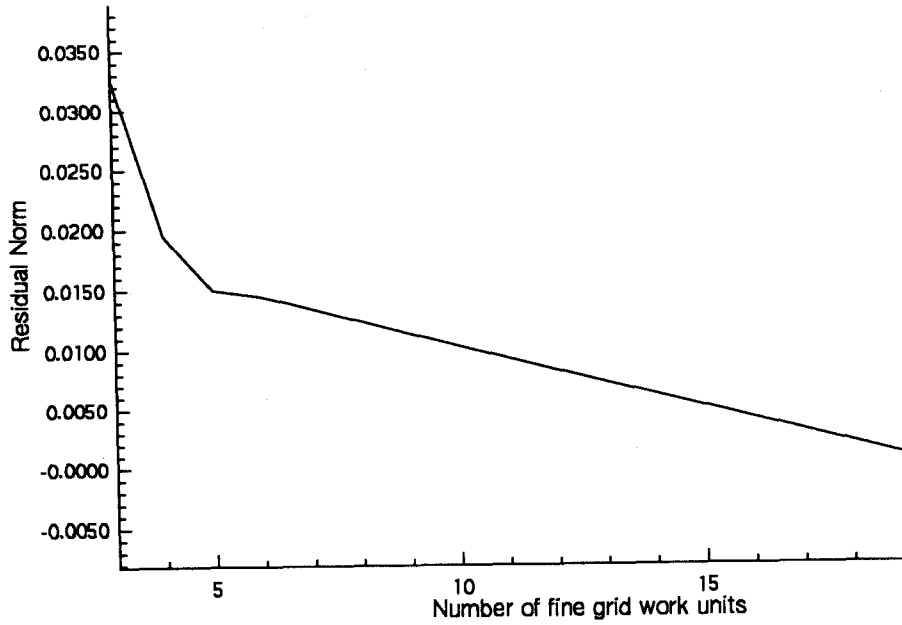

Figure 13. Convergence history on sixth segmented multigrid level; $Re = 800$.

continuation is not required, even for the complex $Re = 800$ case. Grid convergence has been established efficiently through the adaptive multigrid procedure. The outflow boundary condition has been shown to be non-reflective. In addition, it has been shown that the procedure is not very sensitive to the location of the outflow, i.e. far downstream or somewhat closer to the inflow. The flux-split discretization allows a direct computation of the normal velocity and, therefore, mass conservation at grid interfaces and subdomain boundaries is achieved in a simple fashion on non-staggered grids. Extension of the same philosophy to three-dimensional flows is currently in progress.

ACKNOWLEDGEMENTS

This research was supported in part by NASA, Grant No. NAG-397, and by AFOSR, Grant No. 90-0096. The Cray Y-MP at the Ohio Supercomputer Center was used for these computations.

REFERENCES

1. M. C. Thompson and J. H. Ferziger, 'An adaptive multigrid technique for the incompressible Navier–Stokes equations', *J. Comput. Phys.*, **82**, 94–121 (1989).
2. S. G. Rubin, 'RNS/Euler pressure relaxation and flux vector splitting', *Comput. Fluids*, **16**, 485–490 (1988).
3. R. T. Davis and S. G. Rubin, 'Non-Navier–Stokes viscous flow computations', *Comput. Fluids*, **8**, 101–122 (1980).
4. S. G. Rubin and J. C. Tannehill, 'Parabolized/reduced Navier Stokes computational techniques', *Ann. Rev. Fluid Mech.*, **24**, 117–144 (1992).
5. P. K. Khosla and S. G. Rubin, 'Pressure based flux vector splitting for blunt geometries', *Proc. 10th AIAA Computational Fluid Dynamics Conference*, CP914, AIAA, New York, NY, July 1991, pp. 975–976.
6. M. J. Berger and P. Colella, 'Local adaptive mesh refinement for shock hydrodynamics', *J. Comput. Phys.*, **82**, 64–84 (1989).
7. K. Srinivasan and S. G. Rubin, 'Adaptive multigrid domain decomposition solutions of the reduced Navier–Stokes equations', Presented at the *5th SIAM Conf. on Domain Decomposition Methods for Partial Differential Equations*, Norfolk VA, May 1991.
8. D. De Zeeuw and K. Powell, 'An adaptively refined Cartesian mesh solver for the Euler equations', AIAA paper (1542), Presented at the *10th Computational Fluid Dynamics Conference*, Hawaii, June 1991.
9. S. G. Rubin and D. R. Reddy, 'Analysis of global pressure relaxation for flows with strong interaction and separation', *Comput. Fluids*, **11**, 281–306 (1983).
10. S. G. Rubin and A. Himansu, 'Convergence properties of high Reynolds number separated flow calculations', *Int. j. numer. methods fluids*, **9**, 1395–1411 (1989).
11. A. Himansu and S. G. Rubin, 'Multigrid acceleration of a relaxation procedure for the reduced Navier–Stokes equations', *AIAA J.*, **26**, 1044–1051 (1988).
12. H. S. Pordal, P. K. Khosla and S. G. Rubin, 'Transient behaviour of supersonic flow through inlets', *AIAA J.*, **30**(3), 711–717 (1992).
13. M. Israeli and M. Rosenfeld, 'Numerical solutions of incompressible flows by a marching multigrid nonlinear method', *AIAA J.*, **25**, 641–647 (1987).
14. A. Brandt, 'Multi-level adaptive solutions to boundary value problems', *Math. Comput.*, **31**, 333–390 (1977).
15. K. Stewartson, 'Multistructured boundary layers on flat plates and related bodies', in *Advances in Applied Mechanics*, Vol. 14, Academic, New York, 1974, pp. 145–239.
16. D. K. Gartling, 'A test problem for outflow boundary condition—flow over backward facing step', *Int. j. numer. methods fluids*, **11**, 953–967 (1990).
17. S. Caruso, 'Adaptive grid techniques for fluid flow problems', *Ph.D. Thesis*, Thermosciences Division, Department of Mechanical Engineering, Stanford University, CA, 1985.
18. F. Sotiropoulos, 'A primitive variable method for the solution of external and internal incompressible flow fields', *Ph.D. Dissertation*, Department of Aerospace Engineering and Engineering Mechanics, 1991.
19. K. Srinivasan and S. G. Rubin, 'Flow over a backward facing step using the reduced Navier–Stokes equations', *Proc. Minisymposium on Outflow Boundary Conditions*, Stanford University, July 1991.