

AN ALGEBRAIC MULTIGRID SOLVER FOR NAVIER–STOKES PROBLEMS IN THE DISCRETE SECOND-ORDER APPROXIMATION

R. WEBSTER

Roadside, Harpsdale, Halkirk, Caithness, KW12 6UL, U.K.

SUMMARY

An algebraic multigrid scheme is presented for solving the discrete Navier–Stokes equations to second-order accuracy using the defect correction method. Solutions for the driven cavity and asymmetric, sudden expansion test problems have been obtained for both structured and unstructured meshes, the resolution and resolution grading being controlled by global and local mesh refinements.

The solver is efficient and robust to the extent that, for problems attempted so far, no underrelaxation of variables has been required to ensure convergence. Provided that the computational mesh can resolve the flow field, convergence characteristics are almost mesh-independent. Rates of convergence actually improve with refinement, asymptotically approaching mesh-independent values. For extremely coarse meshes, where dispersive truncation errors would be expected to prevent convergence (or even induce divergence), solutions can still be obtained by using explicit underrelaxation in the iterative cycle.

KEY WORDS: algebraic multigrid; defect correction; unstructured meshes

1. INTRODUCTION

Solution of the equations of motion for viscous fluids in the discrete approximation demands powerful computing resources. This is because the flow fields of practical interest are invariably complex, requiring a high degree of spatial resolution. Resolution of length scales spanning many orders of magnitude may be necessary even for stable laminar flows. If Q is some measure of the linear resolving power of a discretization (such as an appropriately scaled inverse of the nodal separation), then the number of discrete equations to be solved, N , will scale as

$$N \sim Q^d, \quad (1)$$

where d is the number of spatial dimensions. Since, moreover, the computational work will scale as N^β , where β depends on the solution method ($\beta \geq 1.0$), then the required computing time T will scale as

$$T \sim Q^{\beta d}. \quad (2)$$

Clearly T can be a very strong function of the required resolution. For example, for equation sets that require direct solution methods (such as Gaussian elimination), the exponent can be as large as nine (i.e. $\beta = 3$, $d = 3$). Since in fluid dynamics we are looking for orders-of-magnitude improvements in resolution, it is essential to develop efficient solvers with optimum scaling ($\beta = 1.0$). It is also important

that this scaling hold good for non-uniform, unstructured meshes so that the nodal economy can be maximized by matching the density of nodes to the required resolution, which may be both anisotropic and inhomogeneous.

In a previous paper¹ a new iterative solver was presented for the discrete Navier–Stokes equations in the first-order approximation which addressed these requirements. The method was based on a fully implicit algebraic multigrid (AMG) scheme. This paper describes changes to the scheme which can virtually eliminate the need for underrelaxation in the iterative cycle. Performance data are provided for the driven cavity and asymmetric, sudden expansion test problems for both structured and unstructured computational meshes. The solutions for these tests have been obtained to second-order accuracy using the defect correction method.

2. THE DISCRETE APPROXIMATIONS

The discrete equation sets for the flow variables are derived from a finite volume discretization of a finite element mesh by enforcing the conservation of mass and momentum for an incompressible fluid. The simplest possible linear element is used, the triangle (in 2D), which is capable of giving second-order-accurate equations. Control volumes are constructed round each vertex node by joining the centroid of each element to the centre of each side (Figure 1). Within any given element, just one flux value is used for the control surfaces so formed and this is obtained by a special interpolation, the centroid providing the single interpolation point. A second discretization within the element is used to derive the interpolation equation. Figure 2 shows three examples of the subcontrol volumes that have been used; the smallest is the one chosen for this work. The scheme is similar to those proposed by Prakash,² Hookey,³ Schneider and Raw⁴ and Burns and Lonsdale.⁵

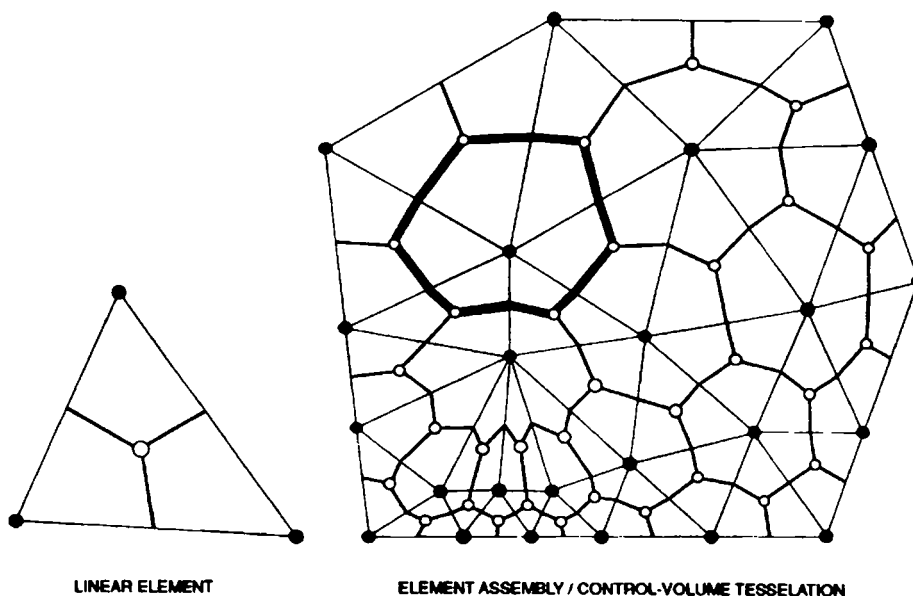


Figure 1. Linear triangular element, element assembly and construction of 'control volume tessellation' (one control volume is highlighted): ○, location of element interpolation point for element velocities v_e ; ●, location of nodal velocities v and pressures p




Figure 2. Interpolation for element velocities: three subcontrol volumes that have been used for local discrete solution of equation of motion

If \mathbf{v} represents the set of nodal velocities, \mathbf{v}_e the set of interpolated velocities within elements and \mathbf{p} the set of nodal pressures, then enforcing the conservation laws for both nodal control volumes and element subcontrol volumes delivers the following set of algebraic equations:

$$\mathbf{A}(\mathbf{v}_e)\mathbf{v} + \mathbf{G}\mathbf{p} = \mathbf{s}, \tag{3}$$

$$\mathbf{A}_e(\mathbf{v}_e)\mathbf{v}_e + \mathbf{F}(\mathbf{v}_e)\mathbf{v} + \mathbf{G}_e\mathbf{p} = \mathbf{s}_e, \tag{4}$$

$$\mathbf{D}\mathbf{v}_e = 0, \tag{5}$$

where \mathbf{A} and \mathbf{G} are the nodal advection–diffusion and gradient operators respectively, \mathbf{A}_e and \mathbf{F} are each part of the advection–diffusion operator for elements, \mathbf{G}_e is the element gradient operator, \mathbf{D} is the nodal divergence operator and \mathbf{s} and \mathbf{s}_e represent the momentum source/sink arrays for the nodal control volumes and element subcontrol volumes respectively.

The matrix \mathbf{A}_e is diagonal, so the solution of equation (4) is trivial, i.e.

$$\mathbf{v}_e = \mathbf{A}_e^{-1}(\mathbf{s}_e - \mathbf{F}\mathbf{v} - \mathbf{G}_e\mathbf{p}). \tag{6}$$

Direct substitution into equation (5) enables the following subset of coupled equations to be formed for the nodal variables:

$$\begin{bmatrix} \mathbf{A}(\mathbf{v}_e) & \mathbf{G} \\ \mathbf{D}\mathbf{A}_e^{-1}\mathbf{F} & \mathbf{D}\mathbf{A}_e^{-1}\mathbf{G}_e \end{bmatrix} \begin{bmatrix} \mathbf{v} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{s} \\ \mathbf{D}\mathbf{A}_e^{-1}\mathbf{s}_e \end{bmatrix}. \tag{7}$$

The solution of equations (6) and (7) is obtained by direct iteration using a predictor–corrector strategy for \mathbf{v}_e and $[\mathbf{v}\mathbf{p}]$, the AMG solver providing the coupled solution of equation (7) for $[\mathbf{v}\mathbf{p}]$.

If upstream values are used in the enforcement of momentum conservation for the nodal control volumes, then equation (7) will be first-order-accurate. For this work a second-order approximation is also required. The simplest possible second-order approximation was adopted using equal proportions of upstream and downstream values for the advected momentum across the control surfaces, equivalent to the central differencing of finite difference methods.

3. THE ITERATIVE SOLUTION METHOD

3.1. Defect correction method: outer iterative cycle

Writing equations (6) and (7) in the more concise form

$$\mathbf{v}_e = \mathbf{A}_e^{-1}(\mathbf{s}_e - \mathbf{H}\boldsymbol{\varphi}), \quad (8)$$

$$\mathbf{L}(\mathbf{v}_e)\boldsymbol{\varphi} = \mathbf{f}, \quad (9)$$

where

$$\mathbf{L}(\mathbf{v}_e) = \begin{bmatrix} \mathbf{A}(\mathbf{v}_e) & \mathbf{G} \\ \mathbf{D}\mathbf{A}_e^{-1}\mathbf{F} & \mathbf{D}\mathbf{A}_e^{-1}\mathbf{G}_e \end{bmatrix}, \quad \boldsymbol{\varphi} = \begin{bmatrix} \mathbf{v} \\ \mathbf{p} \end{bmatrix}, \quad \mathbf{H} = [\mathbf{F} \ \mathbf{G}_e], \quad \mathbf{f} = \begin{bmatrix} \mathbf{s} \\ \mathbf{D}\mathbf{A}_e^{-1}\mathbf{s}_e \end{bmatrix},$$

and writing the first- and second-order approximations of $\mathbf{L}(\mathbf{v}_e)$ and \mathbf{f} as \mathbf{L}_1 , \mathbf{L}_2 and \mathbf{f}_1 , \mathbf{f}_2 respectively, the following iterative procedure can be constructed:⁶

$$\begin{aligned} \mathbf{v}_e^0 &= \mathbf{0}, & \boldsymbol{\varphi}^0 &= \mathbf{0}, \\ \mathbf{v}_e^n &= \mathbf{A}_e^{-1}(\mathbf{s}_e - \mathbf{H}\boldsymbol{\varphi}^n), & n > 0, \\ \mathbf{L}_1(\mathbf{v}_e^n)\boldsymbol{\varphi}^{n+1} &= \mathbf{f}_1^n, & n \leq m, \\ \mathbf{L}_1(\mathbf{v}_e^n)\boldsymbol{\varphi}^{n+1} &= \mathbf{f}_2^n + [\mathbf{L}_1(\mathbf{v}_e^n) - \mathbf{L}_2(\mathbf{v}_e^n)]\boldsymbol{\varphi}^n, & n > m, \end{aligned} \quad (10)$$

where m marks a suitable point in the iteration sequence for switching on the defect correction $\{[\mathbf{L}_1(\mathbf{v}_e^n) - \mathbf{f}_1^n] - [\mathbf{L}_2(\mathbf{v}_e^n) - \mathbf{f}_2^n]\}\boldsymbol{\varphi}^n$ ($m=2$ has been found satisfactory for the test problems described in Sections 4 and 5). At convergence, $\boldsymbol{\varphi}^{n+1} \approx \boldsymbol{\varphi}^n \approx \boldsymbol{\varphi}$ and the second-order equation

$$\mathbf{L}_2(\mathbf{v}_e)\boldsymbol{\varphi} = \mathbf{f}_2^n \quad (11)$$

will be satisfied within the permitted tolerance. The convergence should, moreover, proceed at a rate determined more by the properties of \mathbf{L}_1 than those of \mathbf{L}_2 (see Section 4).

3.2. Linear solver

The algebraic grid and point relaxation methods. The equation system

$$\mathbf{L}_1(\mathbf{v}_e^n)\boldsymbol{\varphi}^{n+1} = \mathbf{f}^n, \quad (12)$$

where \mathbf{f}^n is now understood to include the defect correction if $n > m$, may be represented graphically as a connected nodal network with a one-to-one correspondence between variables (equations) and nodes, the connections between nodes representing the coupling between equations. To distinguish this nodal network from the computational mesh, it will be referred to as the 'algebraic grid' or simply the grid. For a given variable there will be a one-to-one correspondence between grid nodes and mesh nodes and similarly between grid connections and mesh connections (i.e. the edges of elements). For the coupling between unlike variables, grid connections may be regarded as displacements in an abstract dimension.

In an iterative solution procedure based on point relaxation, each node of the grid is visited in turn and that variable is updated/corrected entirely on the basis of local information (i.e. from those neighbours to which the node has direct connections). Because of this, a single sweep through the grid system will

only see changes propagating short distances, i.e. of the order of one nodal spacing. Long-range propagation is a diffusion-like process that requires many iterative sweeps. If λ_i is a relevant propagation distance expressed in units of nodal spacing, then the number of iterations required, n_i , will scale as

$$n_i \sim \lambda_i^2 = (Q/Q_i)^2, \quad (13)$$

where Q is the maximum resolving power of the mesh and Q_i is the minimum resolving power required for the resolution of λ_i . Since, moreover, the computational cost of one iteration will scale as N , the total number of nodes to be visited, the required computing time will scale as

$$T \sim NQ^2 = Q^{d+2} \quad (14)$$

and thus, from the grid system equivalent of equation (2),

$$\beta = 1 + 2/d. \quad (15)$$

Clearly, solvers based on point/local relaxation can scale poorly, with $\beta = 2$ and $5/3$ for 2D and 3D problems respectively. To achieve optimum $\beta = 1$ scaling, it is clearly necessary to have an efficient propagation of corrections over all length scales simultaneously. This requires multigrid methods.

Algebraic multigrid. AMG methods^{7,8} exploit a hierarchy of reduced equation sets (coarse grids) derived from and including the base set (fine grid). Ideally, coarse grid generation proceeds recursively such that each successive grid is a consistent representation of the problem at a reduced scale of resolution, Q_i , associated with length scale λ_i . Just one sweep of a relaxation procedure at this level will be sufficient to propagate changes over λ_i , i.e. $Q = Q_i$ and hence, from equation (13), $n_i \approx 1$. With a sufficient number of grids spanning the complete range of length scales relevant to the problem, an efficient propagation over all length scales can take place simultaneously within one relaxation sweep. Thus, considering the first level of coarsening, if \mathbf{K} is a suitably chosen restriction operator, it may be applied to the base set (12) to form the reduced system

$$\mathbf{L}_1^c \boldsymbol{\varphi}^c = \mathbf{r}^c, \quad (16)$$

where

$$\mathbf{L}_1^c = \mathbf{K} \mathbf{L}_1 \mathbf{K}^T.$$

If \mathbf{r}^c is derived on the basis of the residual \mathbf{r} , i.e.

$$\mathbf{r}^c = \mathbf{K} \mathbf{r} = \mathbf{K}(\mathbf{f} - \mathbf{L}_1 \boldsymbol{\varphi}), \quad (17)$$

then a solution of equation (16) provides a correction $\boldsymbol{\varphi}^c$ that can be used to improve $\boldsymbol{\varphi}$, i.e.

$$\boldsymbol{\varphi} \rightarrow \boldsymbol{\varphi} + \mathbf{K}^T \boldsymbol{\varphi}^c \quad (18)$$

The procedure is thus (i) restrict residual errors to the coarse grid using equation (17), (ii) reduce the coarse grid (long-range) errors by applying local relaxation methods to equation set (16), (iii) prolongate the coarse grid correction and update the fine grid solution using (18) and finally (iv) reduce the fine grid (short-range) errors by applying local relaxation to equation set (12). Clearly equation (16) has the same form as equation (12), so the procedure can be applied recursively to generate smaller equation sets for successively coarser scale corrections. In this way a 'multiscale' correction $\mathbf{K}^T \boldsymbol{\varphi}^c$ can be assembled for updating $\boldsymbol{\varphi}$.

Coarsening procedure. A coarsening procedure based on that devised by Lonsdale⁹ for scalar field variables has been used to generate the reduced equation sets. This consists of seeking out the equations with the strongest coupling (the largest off-diagonals in the \mathbf{L} -matrices) and joining them

together by adding the corresponding matrix coefficients. Some care is required in implementing the procedure.⁹ The elementary matrix representation of Lonsdale's restriction operator \mathbf{K} (dimension $N_i \times N_j$, where $N_i < N_j \leq N$), if required, can be formed by simply adding the appropriate rows of the $N_j \times N_j$ unit matrix. The reduction factors N_i/N_j may be freely chosen, though values of about 0.5 are usually used.

Since here the equation system is for coupled vector and scalar fields, the procedure is implemented in a way which preserves the block structure of the \mathbf{L} -matrix operator. Combining equations for different field variable types is thus forbidden; coarsening is only permitted in 'real space', equivalent to choosing a block diagonal \mathbf{K} -matrix. Note that this does not prevent different coarsening for different field variables.

The process can be terminated when no further reduction in the number of equations is possible and the matrix dimension is then equal to the number of continuum flow variables. In Reference 1 and in this work, however, the process is actually terminated earlier, at between about 30 and 60 equations.

Restriction and prolongation. The elementary \mathbf{K} -matrix restriction combines equations in equal proportion. However, a restriction that permits equations to be combined with different weights, reflecting the different relative importance of each at the lower level of resolution, should produce a more representative coarse grid approximation. For AMG solvers this is particularly important for both uniform and non-uniform discretizations alike, because, even if an initial fine grid is a regular array of identical nodes, the algebraic coarsening process is unlikely to preserve such uniformity. Here, therefore, provision is made for a more general, weighted restriction. Thus, if \mathbf{R} and \mathbf{P} are the actual restriction and prolongation operators to be used, then fine grid and coarse grid weighting operators \mathbf{W} and \mathbf{W}^c are introduced such that

$$\mathbf{R} = [\mathbf{W}^c]^{-1} \mathbf{K} \mathbf{W} \quad (19)$$

subject to the scaling rule

$$\mathbf{R} \mathbf{I} \mathbf{P} = \mathbf{I}^c, \quad (20)$$

where the unit operator \mathbf{I} for the fine grid transforms under the action of \mathbf{R} and \mathbf{P} into the unit operator \mathbf{I}^c for the coarse grid. Combining these equations gives

$$\mathbf{W}^c = \mathbf{K} \mathbf{W} \mathbf{P}. \quad (21)$$

For computational expediency, \mathbf{P} has been chosen to be simply \mathbf{K}^T in this work so that the coarse grid weighting operator is simply the fine grid operator transformed using elementary restriction and prolongation.

For a finite volume discretization a natural choice for \mathbf{W} is the diagonal operator formed from the set of nodal control volumes. Equation (21) can then be simply interpreted as control volume agglomeration and the restriction procedure \mathbf{R} defined by equation (19) as (i) conversion of the fine grid equations into the more naturally additive net flux form (operator \mathbf{W}), (ii) formation of the coarse grid equations (operators \mathbf{K} and \mathbf{K}^T) and (iii) conversion of the coarse grid equations back to the normal form (operator $[\mathbf{W}^c]^{-1}$). The coarse grid approximation so produced results in a robust and efficient solution algorithm.

F-cycle schedule. Following the \mathbf{R} -restriction of residual errors down through the grid hierarchy, ν_1 relaxation sweeps at each level, the multiscale correction is assembled by the reverse procedure of the upward \mathbf{P} -prolongation of solutions (possibly scaled by σ , equation (23), this time applying ν_2 relaxation sweeps following each prolongation. This is the well known V-cycle schedule $V(\nu_1, \nu_2)$. In this work, however, the full multigrid cycle $F(\nu_1, \nu_2)$ has been adopted in which the upward leg of each




Figure 3. F-cycle strategy for transferring residuals and corrections

cycle itself contains nest V-cycles (Figure 3). Furthermore, because the coarsest grid only contains between 30 and 60 nodes, a direct solver is used to obtain an accurate solution.

Local relaxation. Two relaxation schemes have been adopted, both based on point Gauss–Seidel (PGS) relaxation. For the intermediate coarse grids, PGS with optimum damping is used. If $L_1^c = L + D + U$ is the standard splitting for Gauss–Seidel relaxation, L being the lower triangular block, U the upper triangular block and D the diagonal of L_1^c , then the algorithm for ν relaxation sweeps is

$$\begin{aligned}
 & \mathbf{r}^0 = \mathbf{r}^c - L_1^c \boldsymbol{\varphi}^c, & \mathbf{d}^{(0)} = 0, \\
 & i = 0 \\
 & \text{REPEAT} \\
 & \quad i = i + 1, \\
 & \quad \mathbf{d}^{(i)} = (\mathbf{L} + \mathbf{D})^{-1} (\mathbf{r}^{(i-1)} - \mathbf{U} \mathbf{d}^{(i-1)}), \\
 & \quad \mathbf{z}^{(i)} = L_1^c \mathbf{d}^{(i)}, \\
 & \quad \sigma^{(i)} = \langle (\mathbf{z}^{(i)})^T, \mathbf{r}^{(i-1)} \rangle / \langle (\mathbf{z}^{(i)})^T, \mathbf{z}^{(i)} \rangle, \\
 & \quad \boldsymbol{\varphi}^{\alpha(i)} = \boldsymbol{\varphi}^{\alpha(i-1)} + \sigma^{(i)} \mathbf{d}^{(i)}, \\
 & \quad \mathbf{r}^{(i)} = \mathbf{r}^{(i-1)} - \sigma^{(i)} \mathbf{z}^{(i)}, \\
 & \text{UNTIL } i = \nu.
 \end{aligned} \tag{22}$$

Before prolongation the coarse grid corrections are also scaled by the factor

$$\sigma = \langle (L_1^c \boldsymbol{\varphi}^c)^T, \mathbf{r}^c \rangle / \langle (L_1^c \boldsymbol{\varphi}^c)^T, L_1^c \boldsymbol{\varphi}^c \rangle. \tag{23}$$

For the fine grid an approximate four-direction, point Gauss–Seidel algorithm for unstructured meshes is used (4-PGS). This involves some preprocessing for the formation of four continuous line orderings of nodes such that each node is visited once only within each line and lines attempt wherever possible to pass through each node from different directions.

Residual reduction factors. The residual reduction factor or fractional error reduction each F-cycle, μ , depends on the efficiency of the local relaxation process (smoothing) and on the quality of the coarse grid approximation.^{7,8,10} In Section 4, empirical μ -factors are defined and results presented for several test problems.

The system matrix L_1 . Though L_1 does not have to be positive definite, it must have diagonal block matrices that are suitable for solution by scalar AMG methods;⁷ diagonal blocks must be at least positive semidefinite. The first-order discretization based on the advection of upstream momentum (Section 2) produces block diagonal matrices for the velocity component equations that should satisfy that requirement. The block diagonal matrix for the pressure equations is positive semidefinite in any case.

Boundary conditions are implicitly contained in L_1 . At least one pressure node is implicitly fixed in all calculations. No special measures are necessary for dealing with boundary conditions at the lower levels of the grid system. The necessary information is automatically transferred by the restriction operator.

Implicit underrelaxation. Implicit underrelaxation of both velocity and pressure is commonly used to ensure convergence of Navier–Stokes linear solvers. For this coupled AMG linear solver, underrelaxation has not been necessary. Provided that the above-described scaled, weighted restriction is employed, no underrelaxation has been required for any problems tackled so far. However, a small amount of underrelaxation can improve the rates of convergence for both inner and outer iterations. It can be implemented without prejudicing the long-range spatial coupling as follows. All entries in the off-diagonal blocks of L_1 are reduced by a factor ω and/or all entries in the diagonal blocks are increased by $1/\omega$, with appropriate compensations of the right-hand side of the equation sets, evaluated using previous iterates, φ^n . Optimum convergence rates occur for ω -values in the range $1.0 \geq \omega \geq 0.9$.

Note that it is also possible to relax the coupling between like variables by increasing just the diagonal entries of the relevant diagonal block and making the appropriate right-hand-side compensations. This is not recommended. It loosens the spatial coupling that AMG is supposed to be dealing with, resulting in a degradation of convergence performance (including the scaling).

4. PERFORMANCE

The solver has been applied to two well-established test problems: flow in a driven square cavity and flow in a channel with a sudden asymmetric expansion. These problems incorporate several features of complex fluid behaviour that can present difficulties for solvers, particularly at high Reynolds number, e.g. singularities, recirculation, boundary layers, entering flows, outlet flows. Some of these features have been isolated for special investigation by those involved in the development of multigrid methods.^{11,12} Here the problems are presented direct to the solver; apart from considerations of mesh refinement and the need for spatial resolution, no special treatment is afforded to any particular flow feature.

Of interest are the quality of the second-order solutions, the rates of convergence and, in particular, the mesh dependence of both these aspects of performance. To assist in the presentation and analysis of results, it will be useful to introduce mesh resolution and grading factors and to define the convergence factors.

Mesh resolution and grading factors

The inverse nodal separation (linear resolution) and its variation with direction and position (grading) are used to characterize the meshes. The global extremes of the resolution and grading will be sufficient for most purposes. Thus reference is made to the maximum linear resolving power Q , the maximum global grading factor Γ and the maximum local grading factor γ . Q is defined as the ratio of the largest characteristic length scale divided by the closest nodal spacing. Γ is defined as the ratio of the maximum to the minimum nodal separation for elements in the mesh regardless of their position. The local grading factor for any node in the mesh is the ratio of the largest to the smallest separation of the node from its immediate neighbours (i.e. for elements common to the node). Directional aspects are thus largely ignored, except where reference is made to longitudinal and transverse resolution and grading factors Q_x , Γ_x , γ_{xx} and Q_y , Γ_y , and γ_{yy} respectively. The aspect ratio γ_{xy} will also be referred to. In this case the nodal separations in any chosen element are both selected and weighted according to their degree of alignment with the relevant direction.

Convergence factors

Convergence characteristics will be quantified in terms of the convergence factor $\rho^{(n)}$, where $\rho^{(n)}$ is defined as

$$\rho^{(n)} = \|\delta\varphi^{(n)}\|_{\infty} / \|\delta\varphi^{(n-1)}\|_{\infty}, \quad (24)$$

with $\delta\varphi^{(n)}$ the multiscale correction for the iteration index n . Thus the larger the rate of convergence, the smaller is the convergence factor. The average convergence factor ρ for a sequence of N Navier–Stokes (i.e. outer) iterations is

$$\rho = (\|\delta\varphi^{(N)}\|_{\infty} / \|\delta\varphi^{(0)}\|_{\infty})^{1/N} = \left(\prod_0^N \rho^{(n)} \right)^{1/N}. \quad (25)$$

The residual reduction factors μ and $\mu^{(i)}$ for inner iterations are defined similarly but in terms of the Euclidean norm of the residual errors, i.e.

$$\mu^{(i)} = \|\mathbf{r}^{(i)}\|_2 / \|\mathbf{r}^{(i-1)}\|_2, \quad (26)$$

where in this case $\mathbf{r}^{(i)}$ is the residual following the F-cycle of index (i) .

Various F-cycle schedules have been tried from F(1,0) to F(8,2). On the fine grid, $\nu_2 = 1$ actually corresponds to one application of the 4-PGS smoother.

In practice, the important convergence parameter is the fractional reduction of error per unit of computing time, which may not be quite the same as the reduction of error per iteration as defined in equation (25). However, with a fixed number ν of F-cycles per iteration, the computing time per iteration will be more or less constant; then, as long as $\mu^{\nu} \ll \rho$, ρ will be equivalent to the convergence rate in time for all practical purposes. The number of F-cycles does not have to be large to satisfy this requirement. Also, there is little if anything to be gained by insisting that μ^{ν} be extremely small, since much of the work done will be immediately undone when the non-linear terms are updated in the outer iteration.

4.1. Driven cavity test problem

4.1.1. *Problem definition and computational meshes.* A two-dimensional fluid-filled square cavity has three solid, no-slip sides. The remaining side has a uniform tangential velocity which drives a circulation of the fluid in the cavity. This primary circulation may be accompanied by secondary and tertiary circulations at sufficiently high Reynolds numbers. Results have been published by many authors for Reynolds numbers up to 10^4 , on which there would appear to be general agreement.^{13,14}

During the course of this investigation, calculations have been made for a large number of computational meshes and a wide range of values for Q , Γ and γ , each factor having been varied by more than one order of magnitude. The results for just two basic mesh types have been selected for presentation. The prototype triangulations for uniform and non-uniform nodal distributions are shown in Figure 4 as mesh types A and B respectively. Actual meshes used are derived from these two prototypes by a regular nested splitting of each proto-element into q^2 congruent triangles, where q is the level of refinement. Thus for mesh type A, $Q = 2q$, while for mesh type B, $Q = \Gamma q$ ($Q_{\min} = q$) if no smoothing of the mesh is employed. Smoother gradings of mesh B (γ reduction/minimization) are obtained by simple point relaxation, whereby the position of each node is moved to the centre of mass of its nearest neighbours, nodes being weighted uniformly or non-uniformly as required. Boundary nodes are not of course permitted to move out of the boundary surfaces.

Both mesh types have been arbitrarily chosen. No attempt has been made to optimize them for the problem, other than to place the smaller elements for mesh B in the regions of strongest turning flow and steep pressure gradient.

Dirichlet boundary conditions for velocity were enforced on all surfaces, with the non-corner nodes on the driving boundary having the same non-zero tangential component. For pressure, all but one boundary node were free.

4.1.2. *Linear solver: reduction factors.* Reduction factors were typical of those to be expected for the PGS smoothers used.¹⁵ Values averaged over F(1,0)-cycles ranged between 0.08 and 0.3, depending on Re and the mesh. Thus less than about five F-cycles would usually be sufficient to reduce residuals by at least three orders of magnitude, which is more than sufficient (a 10^4 -fold reduction in the residual has actually been observed to slow up the overall rate of convergence). Typical overall averages for mesh type B ($Re = 10$) are given in Table I. The low-Reynolds-number case is given




Figure 4. Prototype triangulations used for driven cavity test problem: for mesh type A, $Q_{\min} = 2q$, $Q = 2q$, $\Gamma = 1$; for mesh type B, $Q_{\min} = q$, $Q = \Gamma q$, $\Gamma = 10$

Table I. Reduction factors for driven cavity test problem: mesh type B ($Re = 10$)

| N | 273 | 573 | 1488 | 2838 | 5673 | 8103 |
|---------------|-------|-------|-------|-------|-------|-------|
| Q | 4 | 6 | 10 | 14 | 20 | 24 |
| $\mu[F(1,0)]$ | 0.088 | 0.093 | 0.104 | 0.125 | 0.145 | 0.178 |

because this is considered to be the most demanding from the point of view of multigrid efficiency for the simple intergrid transfer operators that are being exploited. A robust efficiency is evident, the reduction factors are small and the mesh dependence is weak.

4.1.3. Non-linear solver: convergence performance

Uniform meshes. Calculations have been performed for Reynolds numbers in the range 10–1000 for discretizations with resolving powers Q ranging from 10 to 50. For the boundary layers to be adequately resolved, Q should satisfy

$$Q > Re^{1/2}. \quad (27)$$

In Figure 5 the convergence characteristics for three cases ($Re = 100, 400$ and 1000) are shown for meshes with Q -values in the range 10–50. It is clear that the better the resolution, the better is the rate of convergence, as if $\beta < 1.0$, i.e. an apparently superior scaling to the $\beta = 1.0$ target. However, this is somewhat misleading. Comparing the ρ -values at $Re = 100$ with those from Reference 1 (Table II) reveals a performance penalty for the second-order solution that would appear to be associated with truncation errors (since it reduces with mesh refinement). As the resolution improves, the convergence factors approach those for the first-order method; for $Q \gg Re^{1/2}$, $\beta \rightarrow 1.0$. Note that the ρ^1 results are between first- and second-order, becoming second-order for $Q \gg Re^{1/2}$.

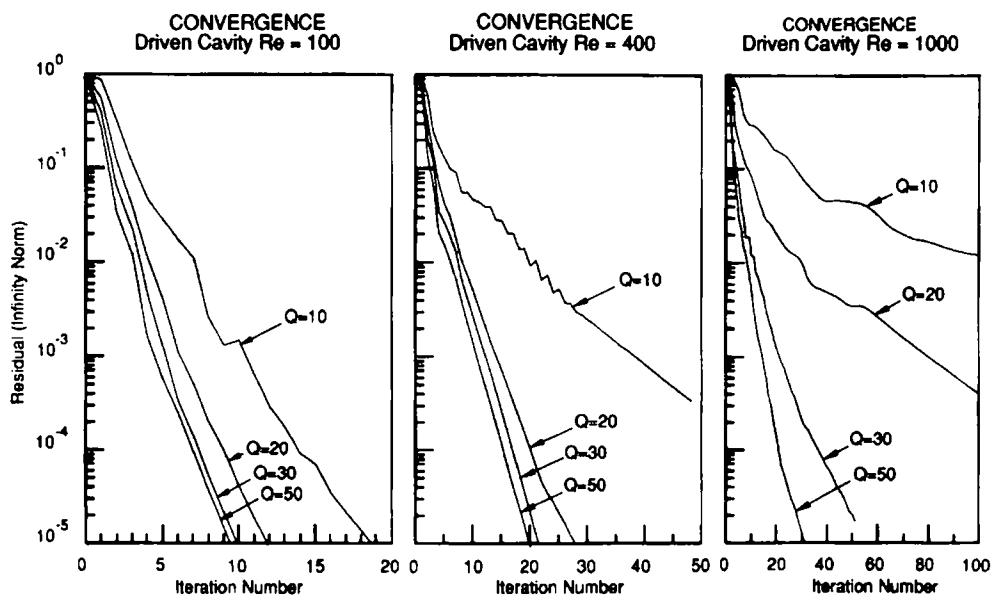


Figure 5. Convergence characteristics for driven cavity test problem at $Re = 100, 400$ and 1000 for uniform meshes with resolving powers in range $Q = 10-50$

Table II. Convergence factors for driven cavity test problem: uniform mesh type A ($Re = 100$; ρ^1 , mean convergence factor for first/second-order discretization; ρ , mean convergence factor for second-order discretization)

| N | 363 | 1323 | 2883 | 5043 | 7803 |
|----------|------|------|------|------|------|
| Q | 10 | 20 | 30 | 40 | 50 |
| ρ^1 | 0.30 | 0.28 | 0.29 | 0.28 | 0.26 |
| ρ | 0.54 | 0.37 | 0.30 | 0.29 | 0.26 |

Similar behaviour is evident at $Re = 400$ and 1000 . Table III compares the convergence factors for all three cases at different Q -values. Values in parentheses are for meshes with insufficient resolution ($Q < Re^{1/2}$), to be discussed further. Note that ρ -values are superior to those that could be expected from segregated solution methods. Even parenthesized results are better.

Non-uniform meshes. The results for B-type meshes are presented in Table IV, which may be compared with those for uniform meshes in Table III. In all cases the mesh grading has been smoothed to reduce γ -values below 1.4 (and Γ from 10 to 4). As for the uniform cases, convergence factors improve with improved resolution, becoming less mesh-dependent when the resolution is sufficient to resolve the flow field. Note also that the values at higher resolution compare well with those for uniform meshes. At lower resolutions, however, it will be evident that mesh-type-B convergence factors are superior for a given number of nodes, suggesting that the non-uniform nodal distribution is more suited to the problem and that with an appropriate adaptive meshing scheme the convergence performance could be optimized.

Further evidence of this is presented in Table V, where results are presented for mesh type B graded more steeply ($\gamma = 3$, $\Gamma = 10$) to give 2.5 times the resolution near the top corners of the cavity. For $Q_{\min} > 10$ there is a 10%–20% improvement in the convergence factors over those for smoother graded meshes (note that for $Q_{\min} \geq 20$ the $\rho(\Gamma = 10)$ results are better than the ρ^1 results in Table II).

Low-resolution meshes. Solutions obtained using computational meshes with insufficient resolution are of doubtful value. However, since the required resolution may not always be correctly anticipated, it is important to know what to expect should the resolving power prove to be inadequate.

Table III. Convergence factors for driven cavity test problem: uniform computational meshes type A

| N | 363 | 1323 | 2883 | 5043 | 7803 |
|-------------------|--------|--------|------|------|------|
| Q | 10 | 20 | 30 | 40 | 50 |
| $\rho(Re = 100)$ | 0.54 | 0.37 | 0.30 | 0.29 | 0.26 |
| $\rho(Re = 400)$ | (0.85) | 0.63 | 0.58 | 0.56 | 0.55 |
| $\rho(Re = 1000)$ | (0.97) | (0.91) | 0.75 | 0.69 | 0.64 |

Table IV. Convergence factors for driven cavity test problem: non-uniform computational meshes type B (global grading $\Gamma = 4$, local gradings $\gamma \leq 1.4$)

| N | 1488 | 2838 | 5673 | 8103 |
|-------------------|-------|-------|-------|-------|
| Q_{\min} | 10 | 14 | 20 | 24 |
| $\rho(Re = 100)$ | 0.33 | 0.32 | 0.29 | 0.28 |
| $\rho(Re = 400)$ | 0.604 | 0.557 | 0.549 | 0.533 |
| $\rho(Re = 1000)$ | 0.789 | 0.725 | 0.665 | 0.66 |

Table V. Convergence factors for driven cavity test problem: non-uniform computational meshes type B (global grading $\Gamma = 10$, local gradings $\gamma \leq 3.0$; global grading $\Gamma = 4$, local gradings $\gamma \leq 1.4$)

| N | 570 | 1488 | 2838 | 5673 | 8103 |
|---------------------|------|------|------|------|------|
| Q_{\min} | 6 | 10 | 14 | 20 | 24 |
| $\rho(\Gamma = 10)$ | 0.39 | 0.33 | 0.29 | 0.24 | 0.24 |
| $\rho(\Gamma = 4)$ | 0.42 | 0.33 | 0.32 | 0.29 | 0.28 |

In this finite volume formulation the L_2 -operator was formed using equal proportions of upstream and downstream values for the momentum advected across control surfaces, so the discretization error will be dispersive. For coarse meshes the dispersion will be large, resulting in solutions that may contain spatial instabilities (wiggles). These may interfere with convergence. Ghia *et al.*¹⁶ and Thompson and Ferziger¹³ have both noticed reduced rates of convergence for centred flux differencing on meshes with reduced resolution. The latter authors tentatively attribute this to the effects of wiggles in regions where the mesh spacing is insufficient to resolve the fine scale structure in the flow field, improvements in convergence rate with refinement being most evident at high Reynolds number.

In this work the improvement in convergence rate with refinement is also evident (Tables III and IV) and, like Thompson and Ferziger, no obvious evidence of spatial instability has been found in the solutions for which the resolution criterion $Q > Re^{1/2}$ is satisfied. On the other hand, when the criterion is not satisfied, a marked degradation in convergence performance has been observed. This is illustrated in Figure 6 for $Re = 1000$ and $Q = 22$ (curve 1), where it will be evident that the reduction of residuals can stall completely ($\rho_i \rightarrow 1.0$) at a residual level well above round-off error, sometimes preceded by an increase ($\rho_i > 1.0$). Large differences between the Euclidean and infinity norms of the nodal corrections can develop, with $\|\delta\varphi\|_{\infty}$ becoming as much as one order larger than $\|\delta\varphi\|_2$. This is just the type of response to be expected from a localized spatial instability and a close examination of the region of maximum residuals confirmed the development of wiggles.




Figure 6. Convergence characteristics for driven cavity test problem at $Re = 1000$ on uniform computational mesh with insufficient resolution ($Q < Re^{1/2}$): curve 1, with no explicit underrelaxation; curve 2, with explicit relaxation factor of 0.8

Even in such extreme cases, convergence can nevertheless be achieved if an explicit underrelaxation of velocity is used in the iterative cycle (Figure 6, curve 2). Thompson and Ferziger used an explicit underrelaxation of velocity in all their calculations, with relaxation factors as low as 0.2 for $Re = 5000$ (0.4 for $Re = 1000$). The parenthesized convergence factors in Table III were obtained in this way, a relaxation factor as low as 0.5 being required for the $Q = 10$, $Re = 1000$ case. Note, however, that a convergence so produced is not necessarily wiggle-free. The very fact that an explicit relaxation is required probably indicates the presence of wiggles.

Improvement in resolution would of course be the preferred way to achieve both a better convergence and a better solution. Sometimes it may be possible to achieve this without actually increasing the total number of nodes. Merely redistributing the existing nodes may provide the required improvement in resolution in the critical regions. An example of this is shown in Figure 7, where just smoothing the grading of mesh type B restores an otherwise stalled convergence. Note also that the non-uniform meshes used for the calculations of Figure 7 contain fewer nodes than the uniform mesh used for the calculations of Figure 6.

4.1.4. Navier-Stokes performance. The solutions for the driven cavity at low ($Q < Re^{1/2}$) and intermediate ($Q \geq Re^{1/2}$) resolution are here compared with published results of others. The comparisons are with calculations made on computational meshes of either similar or greater resolution, the main objective being to show that the method gives accurate solutions when the discretization has a sufficient resolving power.

The results of Ghia *et al.*¹⁶ are used as a reference data set. Comparisons will be limited to the velocity profiles across the mid-axes; an example for $Re = 1000$ is given in Figure 8. The data points are from the tabulated results of Ghia *et al.* for $Q = 128$; the full curves are the profiles calculated here for $Q = 50$. To simplify other comparisons, just the extrema of these profiles, u_{min} , v_{min} and v_{max} , and their spatial location will be used. Ghia *et al.*'s reference values are reproduced in Table VI. For comparison, Tables VII and VIII give the results of this work for type A ($Q = 50$) and type B ($Q_{min} = 24$, $\Gamma = 4$) meshes respectively. For $Re < 3200$ both sets of results agree with those of Ghia *et al.* to within 1.2% of the




Figure 7. Convergence characteristics for driven cavity test problem at $Re = 1000$ on non-uniform computational meshes with insufficient spatial resolution, $Q_{min} = 10$. The curves are for the same mesh type B but with three different gradings γ

DRIVEN CAVITY VELOCITY PROFILES
 $Re = 1000$




Figure 8. Velocity profiles along horizontal and vertical lines through centre of driven cavity at $Re = 1000$. The full curves represent calculations made on a uniform mesh of resolving power $Q = 50$. The data points are those of Ghia *et al.*¹⁶ obtained using a uniform mesh of $Q = 128$

maximum reference velocity. At $Re = 3200$, where the resolution does not satisfy $Q > Re^{1/2}$, the magnitudes of v_{min} still agree to within 5.8% and 3.7% for the uniform and non-uniform meshes respectively. The positions of the extrema agree to within 1.5% and 1.2% of the cavity dimension for the two respective cases. Considering that the resolution of Ghia *et al.* is 2.25 times better than the type A mesh and up to 5.3 times better than for some regions of the type B mesh, this is a remarkable agreement that warrants closer examination.

Table VI. Extrema in driven cavity velocity profiles: reference data of Ghia *et al.*¹⁶ ($Q = 128$)

| Re | v_{min} ($y = 0.5$) | x | u_{min} ($x = 0.5$) | y | v_{max} ($y = 0.5$) | x |
|------|----------------------------|--------|----------------------------|---------|----------------------------|---------|
| 100 | -0.24533 | 0.3047 | -0.21090 | -0.0469 | 0.17527 | -0.2656 |
| 400 | -0.44993 | 0.3594 | -0.32726 | -0.2187 | 0.30203 | -0.2734 |
| 1000 | -0.51550 | 0.4063 | -0.38289 | -0.3281 | 0.37095 | -0.3437 |
| 3200 | -0.54053 | 0.4453 | -0.41933 | -0.3984 | 0.42768 | -0.4062 |

Table VII. Extrema in driven cavity velocity profiles: uniform mesh type A ($Q = 50$)

| Re | v_{min} | x | u_{min} | y | v_{max} | x |
|------|-----------|--------|-----------|---------|-----------|---------|
| 100 | -0.25244 | 0.3200 | -0.21275 | -0.0400 | 0.17873 | -0.2600 |
| 400 | -0.45029 | 0.3600 | -0.32408 | -0.2200 | 0.30065 | -0.2800 |
| 1000 | -0.52487 | 0.4200 | -0.38388 | -0.3400 | 0.37524 | -0.3400 |
| 3200 | -0.59815 | 0.4600 | -0.46181 | -0.4200 | 0.44384 | -0.4000 |

Table VIII. Extrema in driven cavity velocity profiles: non-uniform mesh type B ($Q_{\min} = 24$, $\Gamma = 4$)

| Re | v_{\min} ($y = 0.5$) | x | u_{\min} ($x = 0.5$) | y | v_{\max} ($y = 0.5$) | x |
|------|-----------------------------|--------|-----------------------------|---------|-----------------------------|---------|
| 100 | -0.25292 | 0.3170 | -0.21296 | -0.0510 | 0.17873 | -0.2715 |
| 400 | -0.45311 | 0.3590 | -0.32746 | -0.2200 | 0.30390 | -0.2740 |
| 1000 | -0.52763 | 0.4160 | -0.38563 | -0.3370 | 0.38017 | -0.3380 |
| 3200 | -0.57811 | 0.4510 | -0.42586 | -0.4070 | 0.43908 | -0.3970 |

Thus in Table IX the data for $Re = 1000$ ($Q = 50$ and 128) have been singled out for comparison with the higher-resolution data of Bruneau and Jouron¹⁴ ($Q = 256$). These authors also use a second-order-accurate discretization, but it is based on a second-order upwind differencing scheme for advection. The comparison confirms the accuracy of the present scheme. The extrema $v_{\max}(Q = 50)$ and $u_{\min}(Q = 50)$ are both within 2% of the corresponding values at $Q = 256$, while $v_{\min}(Q = 50)$ is within 0.7% of $v_{\min}(Q = 256)$, i.e. actually closer than it is to the result of Ghia *et al.*, $v_{\min}(Q = 128)$. The spatial positions of the extrema for $Q = 50$ agree with those for $Q = 256$ to within 1% of the cavity dimension, which is as good as could be expected for this resolution (i.e. 2% nodal spacing). Here again, some results are actually closer to the higher-resolution data than they are to the data of Ghia *et al.*

Velocity magnitudes for the present scheme are overpredicted, and as Q is reduced, the overprediction increases roughly as Q^{-2} (Table X), consistent with a dispersive second-order truncation error. Since the scheme of Bruneau and Jouron underpredicts velocity magnitudes,¹⁴ the above differences will represent conservative estimates of the accuracy of the present scheme.

4.2. Asymmetric, sudden expansion test problem

To test the solver on a problem with inflow and outflow boundary conditions, it has been applied to the asymmetric, sudden expansion problem. This is a high-aspect-ratio problem, so it offers a convenient test for the performance of the solver on meshes with highly elongated elements.

Table IX. Mid-axis velocity extrema in driven cavity: (a) this work; (b) data of Ghia *et al.*,¹⁶ (c) data of Bruneau and Jouron¹⁴ ($Re = 1000$)

| Q | v_{\min} ($y = 0.0$) | x | u_{\min} ($x = 0.0$) | y | v_{\max} ($y = 0.0$) | x |
|---------|-----------------------------|--------|-----------------------------|---------|-----------------------------|---------|
| 50 (a) | -0.5247 | 0.420 | -0.3839 | -0.34 | 0.3752 | -0.34 |
| 128 (b) | -0.5155 | 0.4063 | -0.3829 | -0.3281 | 0.3710 | -0.3437 |
| 256 (c) | -0.5208 | 0.4102 | -0.3764 | -0.3398 | 0.3665 | -0.3477 |

Table X. Mid-axis velocity extrema in driven cavity: dependence on mesh resolving power ($Re = 1000$)

| Q | v_{\min} ($y = 0.0$) | x | u_{\min} ($x = 0.0$) | y | v_{\max} ($y = 0.0$) | x |
|-----|-----------------------------|-------|-----------------------------|--------|-----------------------------|--------|
| 10 | -0.6191 | 0.4 | -0.4818 | -0.3 | 0.5361 | -0.3 |
| 20 | -0.5580 | 0.45 | -0.4247 | -0.35 | 0.4089 | -0.35 |
| 30 | -0.5316 | 0.433 | -0.3944 | -0.333 | 0.3843 | -0.333 |
| 40 | -0.5276 | 0.425 | -0.3862 | -0.325 | 0.3763 | -0.325 |
| 50 | -0.5247 | 0.420 | -0.3839 | -0.34 | 0.3752 | -0.34 |




Figure 9. Prototype triangulation for asymmetric, sudden expansion test problem consisting of 81 proto-elements:
 $Q_x = 5\Gamma_x q$; $Q_y = 3\Gamma_y q$; q , level of nested refinement; $\Gamma_x = 32$; $\Gamma_y = 4$; $\gamma_{xx} = 4$; $\gamma_{yy} = 4$; $\gamma_{xy} = 5.3$

4.2.1. Problem definition and mesh. Flow enters a two-dimensional channel with a parabolic inlet velocity profile. Some distance from the inlet there is a one-sided step increase in channel width to $3/2$ times the original. Flow separates at the re-entrant corner and a recirculation zone is established after the step. The axial extent of the circulation is marked by the point of reattachment or the point at which unidirectional flow is re-established across the entire width of the channel. This depends on the Reynolds number. Results have been published for Reynolds numbers up to and in some cases exceeding $Re = 250$, Re being based on step height and mean inlet velocity (note that this definition gives values six times smaller than those based on hydraulic diameter and maximum inlet velocity).

A significant length of the expanded channel (exceeding say three hydraulic diameters) needs to be modelled to ensure that the imposed outlet boundary condition does not unduly influence the behaviour upstream. Thus the problem is bound to be one of large aspect ratio (~ 10), and in view of the need for fine resolution near the point of separation, the discretization could prove to be nodally expensive if uniform meshes are used. Thus only non-uniform meshes have been adopted for this investigation and results for just one unstructured mesh type have been selected for presentation.

The prototype triangulation is illustrated in Figure 9. It consists of 81 proto-elements which have been assembled to give the highest resolution at the point of separation and so that the lateral resolving power Q_y is maintained moderately high up to the point of reattachment. The actual meshes used were obtained by a q -fold nested refinement of each proto-element into as many as $q^2 = 64$ congruent triangles, giving a finest mesh of 5184 elements (2717 nodes). The mesh is anisotropic and inhomogeneous, with grading factors $\gamma_{xx} = 4$, $\gamma_{yy} = 4$, $\gamma_{xy} = 5.3$, $\Gamma_x = 32$ and $\Gamma_y = 4$. Dirichlet boundary conditions for velocity and free pressure boundary conditions apply on all surfaces except the outlet. The latter (continuity and constant pressure) was placed 38 step lengths from the expansion.

4.2.2. Linear solver: reduction factors. The reduction factors for this test problem were again within the expected range for point Gauss-Seidel relaxation. Table XI gives the average values for a low Reynolds number. Both definitions of Reynolds number are used, i.e. the first, Re , based on step height and average inlet velocity, the second, Re^h , based on hydraulic diameter and maximum inlet velocity. The latter has been added to give a more appropriate pointer for comparisons with the driven cavity data (Table I).

Table XI. Reduction factors for asymmetric, sudden expansion test problem ($Re = 16.67$, $Re^h = 100$)

| n | 1236 | 2133 | 3273 | 4656 | 8151 |
|---------------|-------|-------|-------|-------|-------|
| q | 3 | 4 | 5 | 6 | 8 |
| $\mu[F(1,0)]$ | 0.109 | 0.159 | 0.184 | 0.215 | 0.306 |
| $\mu[F(2,1)]$ | 0.042 | 0.059 | 0.091 | 0.114 | 0.143 |

Table XII. Convergence factors for asymmetric, sudden expansion test problem (level of refinement, $q = 8$; number of unknowns, $N = 8151$)

| | | | | | |
|--------|-------|-------|-------|-------|-------|
| Re | 16.67 | 50 | 100 | 150 | 200 |
| Re^h | 100 | 300 | 600 | 900 | 1200 |
| ρ | 0.426 | 0.587 | 0.684 | 0.754 | 0.816 |

4.2.3. *Non-linear solver: convergence performance.* Convergence factors for the finest mesh for the same range of Reynolds numbers are presented in Table XII. Both definitions of Reynolds number are again used to facilitate comparisons with the driven cavity data (Table III). This reveals slower rates of convergence but, nevertheless, still better than those for segregated solution methods. The mesh dependence of the convergence (Table XIII) is similar to that for the driven cavity. In Table XIII, typical values for ρ are given at four different levels of refinement at just three selected Reynolds numbers.

Dick and Linden¹⁷ obtained second-order-accurate coupled solutions to the same test problem discretized using a flux-difference-splitting approach. They also used a defect correction scheme, but their solver was based on a geometric (FAS) multigrid method. Their published result for the case corresponding here to $Re = 100$ was $\rho = 0.81$, which compares with $\rho = 0.68$ in Table XII. The reason for the difference is not clear (however, see Section 5). Dick and Linden also reported a deterioration in convergence performance with mesh refinement which has not been observed in this work. The evidence is for constant or improving convergence rates with mesh refinement (Table XIII), as for the driven cavity test problem.

4.2.4. *Navier-Stokes performance.* The axial extent of the recirculation eddy following the step expansion will be used as the gauge for assessing the quality of the solutions. Experimental data are available, but not for a truly parabolic inlet velocity profile. Predictions of the experiment would therefore have to be based on the measured profile, which is known to result in a short eddy. Since overdiffusive calculational methods would tend to underpredict the eddy length anyway, there could well be fortuitously good first-order calculations of this experiment wherever a parabolic inlet velocity profile has been mistakenly used. Here such complications are avoided by assessing the performance against other calculations of the idealized problem only. Thus the results are compared with the higher-order-accurate calculations of Hutton and Smith¹⁸ and with the first- and second-order-accurate calculations of Shaw.¹⁹

For Reynolds numbers up to $Re = 200$ the resolution requirement should be satisfied for the mesh specified in Section 4.2.1. Results for the range $Re = 16.7-200$ are given in Figure 10 as the five full-circle data points. For comparison, two sets of data from Hutton and Smith are plotted, one as a full curve, obtained using a coarse mesh of 69 biquadratic rectangular elements (246 nodes), and the other as

Table XIII. Convergence factors for asymmetric, sudden expansion test problem (N , number of unknowns; q , level of refinement)

| N | 2133 | 3273 | 4656 | 8151 |
|-------------------|-------|-------|-------|-------|
| q | 4 | 5 | 6 | 8 |
| $\rho(Re = 16.7)$ | 0.464 | 0.432 | 0.426 | 0.426 |
| $\rho(Re = 50)$ | 0.602 | 0.608 | 0.587 | 0.587 |
| $\rho(Re = 150)$ | 0.911 | 0.807 | 0.771 | 0.754 |




Figure 10. Length of recirculation eddy versus Reynolds number: comparison with published results of Hutton and Smith¹⁸ and Shaw¹⁹

four open-circle data points, obtained using a finer mesh of 256 quadratic triangular elements (565 nodes). The agreement is within 2% in all cases. This is consistent with that achieved for the driven cavity.

The open-square data points are from the calculations of Shaw using 600 rectangular linear elements. The two lower points at $Re = 12.5$ and 25 are second-order-accurate and consistent with Hutton and Smith's data. The remaining three points were obtained using a first-order scheme for advection. They underpredict the length of the recirculation, by as much as 27% at $Re = 100$. Shaw attributed this to the coarseness of the mesh and the false numerical diffusion associated with the first-order upwind scheme.

5. DISCUSSION AND GENERAL COMMENTS

The above results are a representative sample of the tests to which the solver has been applied. On the basis of all tests the following general comments are made and the subsequent conclusions drawn.

Inner iterations: linear solver performance

It has not been found necessary to use any underrelaxation of variables to ensure convergence of the linear solver. The rates of reduction of the residual errors within inner iterations are typical of those to be expected for the PGS-based relaxation methods used and the simple intergrid transfer operators being exploited. Note that from the point of view of the coarse grid approximation the values quoted are for the worst Navier–Stokes cases, low Reynolds numbers. They are nevertheless more than adequate for the problems attempted. The dependence of μ on mesh size, albeit weak, is an inevitable consequence of the primitive intergrid transfer operators used. It is sufficiently weak, however, to have little if any impact on the scaling of ρ . A higher-order interpolation would be required for a better coarse grid approximation and this is unlikely to be cost-effective.

Outer iterations: non-linear solver performance

Provided that the computational mesh has a sufficient resolving power for the problem, rapid convergence superior to that possible with segregated solution methods is achieved. When, however, the mesh has insufficient resolution (e.g. $Q < Re^{1/2}$ in the driven cavity case), the convergence can stall ($\rho \rightarrow 1$) unless an explicit underrelaxation of velocity is exploited. This is thought to be due to the influence of the dispersive truncation error on the convergence process. For finer meshes, explicit relaxation is not required and rates of convergence improve with refinement, asymptotically approaching mesh-independent values as the resolution is increased, i.e. $\beta \rightarrow 1$ as $Q \rightarrow \infty$. No evidence has been found for $\beta > 1$ in any applications so far. If this proves to be a better performance than that achieved with other defect correction, multigrid algorithms, the accuracy of the present discretization may be responsible.

Accuracy of solutions

In Section 4 the comparison of solutions for this discretization at $Q = 50$ with those for the high-resolution meshes of Bruneau and Jouron at $Q = 256$ suggested exceptional accuracy. To give some perspective to this, compare the driven cavity solutions at $Re = 1000$ with those for other discretizations at similar resolving powers.

Perhaps the most commonly exploited discretization is that based on the first-order upwind treatment of advection, popular because it produces diagonally dominant matrices that permit robust iterative solutions. Thus, for computational meshes with resolving power $Q = 40$, Mukhopadhyay *et al.*²⁰ underpredict v_{\min} by at least 30%, whereas with this discretization the overprediction is less than 1.5%. For a coarser mesh at $Q = 20$, Wang *et al.*²¹ underpredict v_{\min} by 74%, whereas here the overprediction is less than 7%. Note that these differences are consistent with the Q^{-1} and Q^{-2} dependences that would be expected for first-order diffusive and second-order dispersive truncation errors respectively. The large underpredictions are of course due to the false diffusion which is also responsible for the robustness of the iterative solution.

For second-order upwind differencing methods the error is smaller. Thus Bruneau and Jouron, for uniform meshes at $Q = 64$, underpredict v_{\min} , u_{\min} and v_{\max} by at least 5%, 6% and 6.4% respectively, whereas the present discretization (at a slightly lower resolving power of $Q = 50$) overpredicts these extrema by less than 0.7%, 2.0% and 2.4% respectively. Using Vanka's²² scheme on the $Q = 64$ mesh, Bruneau and Jouron find underpredictions of 15%, 18% and 17% for the respective extrema.

The exceptional accuracy of the present scheme is due to the special interpolation (equation (6)). One of the originators of this type of discretization, Schneider,²³ has also commented on the accuracy of the method.

6. CONCLUSIONS

An efficient and robust iterative numerical method is presented for solving the coupled equations of motion for viscous fluids in the discrete second-order approximation.

Provided that the discretization has sufficient spatial resolution for the flow field, a rapid convergence to machine accuracy is achieved which is almost mesh-independent in so far as the convergence rates either improve or are maintained for increased nodal concentration.

With sufficient resolution the method is also robust to the extent that no underrelaxation of flow variables has been required to ensure convergence. However, small amounts of underrelaxation can improve convergence rates. Converged solutions can also be obtained when the mesh resolution is insufficient to resolve the flow field, but in the more extreme cases of low resolution some explicit underrelaxation is necessary to prevent the non-linear solver from stalling.

The discretization provides accurate solutions on relatively coarse meshes. This is probably due to the interpolation scheme used for the momentum flux within elements, which is based on a local discrete solution of the equations of motion within the element.

ACKNOWLEDGEMENTS

I am grateful to Guy Robinson and Ben MacGregor for their continued interest, help and encouragement.

REFERENCES

1. R. Webster, 'An algebraic multigrid solver for Navier Stokes problems', *Int. j. numer. methods fluids*, **18**, 761–780 (1994).
2. C. Prakash, 'An improved control volume fine element method for heat and mass transfer and for fluid flow using equal-order, velocity–pressure interpolation', *Numer. Heat Transfer*, **9**, 253–276 (1986).
3. N. A. Hookey, 'Evaluation and enhancements of control volume finite element methods for two-dimensional fluid flow and heat transfer', *M.Eng. Thesis*, Department of Mechanical Engineering, McGill University, Montreal, 1986.
4. G. E. Schneider and M. J. Raw, 'Control volume finite element method for heat transfer and fluid flow using collocated variables', *Numer. Heat Transfer*, **11**, 363–390 (1987).
5. A. Burns and R. D. Lonsdale, personal communication, 1991.
6. P. W. Hemker, in *Lecture Notes in Mathematics*, Vol. 1228, Springer, New York, 1986, p. 150.
7. J. Ruge and K. Stuben, 'Algebraic multigrid', in S. Cormick (ed.), *Frontiers in Applied Mathematics*, Vol. 5, *Multigrid Methods*, SIAM, Philadelphia, 1987.
8. A. Brandt, 'Algebraic multigrid theory: the symmetric case', *Proc. 2nd Int. Multigrid Conf.*, SIAM, Copper Mountain, CO, April 1983.
9. R. D. Lonsdale, 'An algebraic multigrid solver for the Navier–Stokes equations on unstructured meshes', *Int. j. numer. methods heat fluid flow*, **3**, 3–14 (1993).
10. A. Brandt, 'Rigorous local mode analysis of multigrid', in *Prelim. Proc. Fourth Copper Mountain Conf. on Multigrid Methods*, SIAM, Copper Mountain, Colorado, 1989, Chap. 2.
11. A. Brandt and I. Yavneh, 'Inadequacy of first-order upwind difference schemes for some recirculating flows', *J. Comput. Phys.*, **93**, 128–143 (1991).
12. A. Brandt and I. Yavneh, 'On multigrid solution of high-Reynolds incompressible entering flows', *J. Comput. Phys.*, **101**, 151–164 (1992).
13. M. C. Thompson and J. H. Ferziger, 'An adaptive multigrid technique for incompressible Navier Stokes equations', *J. Comput. Phys.*, **82**, 94–121 (1989).
14. C. Bruneau and C. Jouron, 'An efficient scheme for solving steady incompressible Navier–Stokes equations', *J. Comput. Phys.*, **89**, 389–413 (1990).
15. P. Wesseling, *An Introduction to Multigrid Methods*, Wiley, Chichester, 1992.
16. U. Ghia, K. N. Ghia and C. T. Shin, 'High-resolutions for incompressible fluid flow using the Navier–Stokes equations and a multigrid method', *J. Comput. Phys.*, **48**, 387–411 (1982).
17. E. Dick and J. Linden, 'A multigrid method for steady incompressible Navier–Stokes equations based on flux difference splitting', *Int. j. numer. methods fluids*, **14**, 1311–1323 (1992).
18. A. Hutton and R. Smith, 'The prediction of laminar flow over a downstream facing step by the finite element method', *CEGB Rep. RD/B/N3660*, Berkeley, 1979.
19. C. T. Shaw, 'Using a segregated finite-element scheme to solve the incompressible Navier–Stokes equations', *Int. j. numer. methods fluids*, **12**, 81–92 (1991).
20. A. Mukhopadhyay, T. Sundararajan and G. Biswas, 'An explicit transient algorithm for predicting incompressible viscous flows in arbitrary geometry', *Int. j. numer. methods fluids*, **17**, 975–993 (1993).
21. Y. L. Wang, J. He and B. Q. Zhang, 'A calculation procedure for two-dimensional elliptic flows', *Int. j. numer. methods fluids*, **9**, 609–617 (1989).
22. S. P. Vanka, *J. Comput. Phys.*, **65**, 138 (1986).
23. G. E. Schneider, in W. J. Minkowycz *et al.* (eds), *Handbook of Numerical Heat Transfer*, Wiley, New York, 1988, p. 417.