JOURNAL OF COMPUTATIONAL PHYSICS **142,** 581–610 (1998)
ARTICLE NO. CP985915

Comparison of Multigrid Methods for Neutral and Stably Stratified Flows over Two-Dimensional Obstacles

M. F. Paisley and N. M. Bhatti

Department of Mathematics and Statistics, School of Computing, Staffordshire University, Stafford, United Kingdom ST18 0AD
E-mail: cmtmfp@soc.staffs.ac.uk

Received July 15, 1997; revised December 12, 1997

Coupled and decoupled methods for solving the Navier–Stokes equations are compared as underlying smoothers for a multigrid algorithm. Numerical results for the benchmark problem of the lid-driven cavity confirm that residual reduction factors per multigrid cycle for the coupled method are superior to those of the decoupled method. The line-wise implementation of the coupled method is shown to be more efficient than the cell-wise while retaining good convergence rates and is the fastest method for this problem. Both approaches are applied to the more challenging problem of homogeneous and inhomogeneous viscous flow past obstacles (a vertical barrier and a cosine-shaped bump), where the flow is largely unidirectional and good convergence rates for the coupled method can now only be achieved by solving the coupled equations in vertical lines. The convergence rates of both methods are shown to deteriorate for these flows, compared with those for the lid-driven cavity, but the deterioration is generally less for the decoupled method, however, and the relative efficiency of the decoupled method means that execution times are significantly less than those required for the coupled method. In the case of stratified flows convergence difficulties are found for the coupled approach when a high order discretisation is used for the density transport equation. Strategies developed to overcome this, based on the use of double discretisation techniques, are described.

Key Words: incompressible Navier–Stokes, multigrid, stratified flow.

1. INTRODUCTION

Since the pioneering work of Brandt [1] multigrid methods have become an accepted numerical procedure in many fields on application and certainly that of incompressible fluid mechanics. Along with several decisions to be made regarding the discretisation of the continuous problem is the decision concerning the nature of the smoother used in the

0021-9991/98 $25.00

multigrid algorithm. Of the many variants proposed, two basic approaches have emerged, namely decoupled (segregated) and coupled (desegregated). In the former approach, updates to each velocity are computed over the entire domain using the respective momentum equations in turn and the localised coupling between the velocities, and the pressure is then achieved using the continuity equation. Brandt's DGS (distributed Gauss Seidel) method [2] is an early example of a multigrid method of this type, although the most widely used now are based on pressure correction methods, with the SIMPLE [3] algorithm being the best known. In the latter approach, on the other hand, all flow variables are updated simultaneously in localised sets and global coupling is then achieved by sweeping over the computational domain in a preordered manner as in [4]. Although developed somewhat later than the pressure correction methods, coupled methods were immediately incorporated as smoothers in multigrid algorithms and several demonstrations of the effectiveness of the SCGS (symmetric coupled Gauss Siedel) method were given for the lid-driven cavity problem [5, 6] and for more general flows in two and three dimensions [7, 8]. Results of corresponding multigrid computations for the SIMPLE pressure correction method followed [9, 10] and numerous computations using both approaches have been reported since. See, for example, [11–14] for those using the coupled approach and [15–17], the decoupled method.

Recently multigrid computations were reported for the Navier–Stokes equations describing density-stratified flow past two-dimensional obstacles using the SIMPLE pressure correction method as smoother [18]. This work represents one of the first applications of a multigrid method to flows which are primarily of meteorological interest. Multigrid methods had been applied previously to flows described by the same equation set but in much simpler geometry [19], or to flows in more complex geometry but only requiring the solution to a simpler model, such as a Poisson equation [20], for example. Grid-independent convergence rates were demonstrated in [18] for computations of steady and unsteady flows past a vertical barrier and a cosine-shaped obstacle, under conditions of neutral and stably stratified flow, at low and high Reynolds number. In the absence of consensus on the issue of choice of the smoother, the SIMPLE pressure correction method was chosen arbitrarily, largely for reasons of convenience. The purpose of this paper is to provide a comparison between those results and new results obtained with a coupled method, with a view to providing some guidance on the matter, at least for these kinds of flows. The question is not merely academic, for flows of practical interest in many fields, including meteorology, are three-dimensional, time-dependent, and turbulent, and each of these generalisations tends to reduce the effectiveness of multigrid. It is clearly important to be using the most effective multigrid smoother when computing flows of all kinds, but especially those for which the benefit of multigrid is expected to be least.

The comparisons previously made are few and somewhat limited in scope. A practical and theoretical comparison for the lid-driven cavity [21] concluded that the coupled approach had consistently better smoothing rates than the pressure correction approach. This generally led to faster execution times for the coupled method for a given level of convergence for the standard driven cavity problem, although for the higher Reynolds numbers and the finest grid there was little to choose between the methods. Other comparisons for the same test problem [22] confirmed this. The flows of interest here, however, are quite different from those of this idealised test problem which is discretised on a square uniform grid and has simple boundary conditions. Flows past obstacles are largely undirectional and necessitate grids that are nonuniform and which may contain cells of high aspect ratio.

Early attempts with the cell-wise implementation of the coupled method (SCGS) for such strongly aligned flows in ducts [7] suffered from poor asymptotic convergence rates. For such flows, however, the coupled method should be implemented in a line-wise manner, so that variables belonging to entire lines of cells are updated simultaneously. Implemented in this way, it has been demonstrated that the coupled method can overcome the poor convergence properties associated with the cell-wise approach, and both the SCAL (symmetrically coupled alternating line) [12] and CLGS (collective line Gauss Seidel) [23] methods achieve this. Although there are other methods of implementing a coupled approach, such as those based on ILU factorisation, the results of the comparison in [23] suggests that the line-wise sweeping is probably the most efficient.

Given the apparent superiority of the coupled approach over the decoupled for the lid-driven cavity problem, the question naturally arises regarding which is best when flows are strongly aligned. The conclusion in [22] appeared to be that for flow over a backward facing step the line-wise coupled approach was consistently faster than the SIMPLE pressure correction method. The outcome of such comparisons may well be problem-dependent, but the relative performance of the pressure correction approach and the line-wise coupled approach implemented in the context of neutral and buoyant flow past obstacles is the main subject of interest in this paper.

The addition of an extra transport equation should not in itself affect the performance of a multigrid algorithm, and indeed, convergence rates in terms of cycles required for convergence in [18] for stratified flows were similar to those for the corresponding neutral flow. However, for one set of parameters a convergent multigrid iteration could not be obtained. Further work has shown that this was partly related to the small diffusion coefficient in the transport equation of the density equation, and the consequent lack of ellipticity in the discrete equations. This has been overcome subsequently using the ideas of double discretisation [24] and in fact this technique appears to be an essential element for the successful application of the coupled method in this context.

The organisation of the paper is as follows. The mathematical formulation of the flows under consideration is described first in Section 2 and the underlying numerical procedures and smoothers for a single grid computation are given in Section 3. Section 4 contains details of the multigrid algorithm and the results of the investigations are discussed in Section 5. The paper ends with some conclusions in Section 6.

2. MATHEMATICAL MODEL

The mathematical model under consideration has already been given in [18] and is only described here briefly. Confining our attention to laminar flow and using the Boussinesq approximation, in which density variations are neglected in the advection terms in the momentum equations, the equations of motion describing steady two-dimensional density stratified flow are

$$u\frac{\partial u}{\partial x} + w\frac{\partial u}{\partial z} = -\frac{\partial p}{\partial x} + \frac{1}{Re}\left\{\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial z^2}\right\} \tag{1}$$

$$u\frac{\partial w}{\partial x} + w\frac{\partial w}{\partial z} = -\frac{\partial p}{\partial z} - \frac{1}{F_h^2}\vartheta + \frac{1}{Re}\left\{\frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial z^2}\right\} \tag{2}$$

$$\frac{\partial u}{\partial x} + \frac{\partial w}{\partial z} = 0 \tag{3}$$

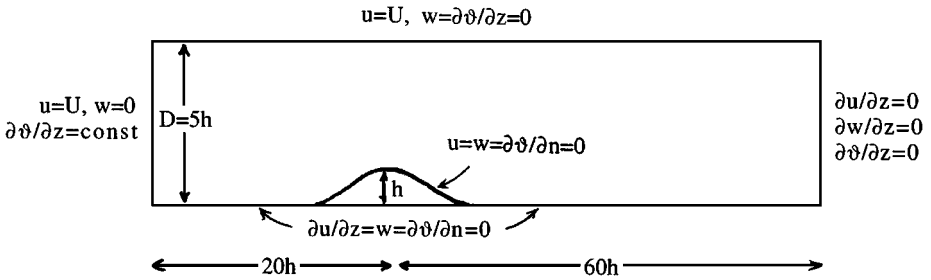

FIG. 1. Computational domain and boundary conditions for flow over smooth geometry.

$$u\frac{\partial \vartheta}{\partial x} + w\frac{\partial \vartheta}{\partial z} = \frac{1}{Re.Sc}\left\{\frac{\partial^2 \vartheta}{\partial x^2} + \frac{\partial^2 \vartheta}{\partial z^2}\right\}. \tag{4}$$

Equations (1)–(4) respectively express conservation of momentum in the two coordinate directions, the incompressibility constraint, and the transport of the scalar variable responsible for changes in density. F_h is the Froude number, given by $F_h = U/Nh$, where U is the freestream velocity, h is the obstacle height, and N is the bouyancy frequency (constant for a linear density gradient). ϑ is the density scalar given by $\vartheta = D(\rho - \rho_0)/h\Delta\rho$, where ρ is the physical density and $\Delta\rho$ the magnitude of the density change over the nondimensional domain height D/h. Sc is the Schmidt number (the ratio of viscous to diffusive effects) and is taken to be 1000. The equations have been nondimensionalised by U, h, and the reference density ρ_0.

For stratified flows of finite depth, there is an additional controlling parameter, namely K, defined by $K = D/\pi h F_h$, which is the ratio of the speeds of the fastest internal mode and the freestream. (Neutral flow corresponds to an infinite Froude number and $K = 0$.) For buoyant flow when $0 < K < 1$ all lee waves are swept downstream, while for $K > 1$ lee waves may exist in the wake of an obstacle and the flow may become unsteady, in addition [25]. The case $K = 1$ is the transition between these two regimes. It was shown in [18] that the SIMPLE algorithm was generally found to be an efficient smoother over the range $0 \le K \le 1$.

The geometry of interest here is flow in a channel where the ratio of channel depth to obstacle height is given by $D/h = 5$. The computational domain and boundary conditions are the same as in [18] and are illustrated in Fig. 1 for the case of smooth geometry. Uniform horizontal velocity and a linear density profile are specified at inflow, while simple zero gradient conditions are applied for all variables at outflow. A "moving wall" condition is applied at the top of the domain. Two obstacle shapes have been used: a vertical barrier for which the computational grid is Cartesian, and an obstacle with smooth profile given by $h(x) = 0.5(1 + \cos(\pi x/1.8))$ for which the grid is curvilinear. No-slip boundary conditions are applied to the obstacle, and zero stress conditions upstream and downstream. In the case of the cosine-shaped obstacle, no-slip conditions are applied for $|x| \le 8$, modelling the geometry of an obstacle mounted on a base plate often used in laboratory experiments [26]. The initial condition for all computations is that of uniform horizontal velocity, and for the stratified cases the initial density profile varies linearly with z.

Laboratory experiments such as those described in [26] have shown that stratification in the flow over a two-dimensional obstacle for $K \le 1$ acts to suppress vertical motion in the turbulent wake and that lee side separation is inhibited as a consequence. As K ranges from zero to unity the separation length reduces from its value in neutral flow,

leading to a fall in the value of the corresponding pressure drag. It was shown in [26] that these qualitative features of the real turbulent flow could be reproduced in the results of low Reynolds number computations with the laminar Navier–Stokes equations for an appropriate choice of Reynolds number. Although no formal comparison of quantities such as separation lengths and drag coefficients has been made, the low Reynolds number flows described here are, thus, not disimilar to the real flows.

3. DISCRETISATION

Expressed in the form of steady two-dimensional convection–diffusion with source term, the transport equations in the set (1)–(4) are

$$\frac{\partial}{\partial x}(u\varphi) + \frac{\partial}{\partial z}(w\varphi) = \frac{\partial}{\partial x}\left(\Gamma\frac{\partial\varphi}{\partial x}\right) + \frac{\partial}{\partial z}\left(\Gamma\frac{\partial\varphi}{\partial z}\right) + S, \tag{5}$$

where φ may stand for u, w, or ϑ. Discretising on rectangular Cartesian grids with control volumes of area ΔV and using the finite volume approach gives

$$\sum\left[\left(u\varphi - \Gamma\frac{\partial\varphi}{\partial x}\right)\Delta z - \left(w\varphi - \Gamma\frac{\partial\varphi}{\partial z}\right)\Delta x\right] = S\Delta V, \tag{6}$$

where the summation is taken over the four sides of the control volume. Staggered grids are used with the usual arrangement of variables and associated control volumes, Fig. 2, and the problem reduces to that of providing estimates of flow quantities on cell sides. Γ is constant for laminar flows and local interpolation is used for u and w, while the diffusive terms are obtained by standard two-point differencing. The advective terms, on the other hand, are dealt with using a second-order flux-limited scheme designed to prevent spurious oscillation, originally developed for unsteady inviscid compressible flows [27] and adapted as follows [28].

The flow across cell boundaries is considered to be one-dimensional in directions normal to the cell faces. The value of the flow variable (u, w, or ϑ) at the centre of the cell under

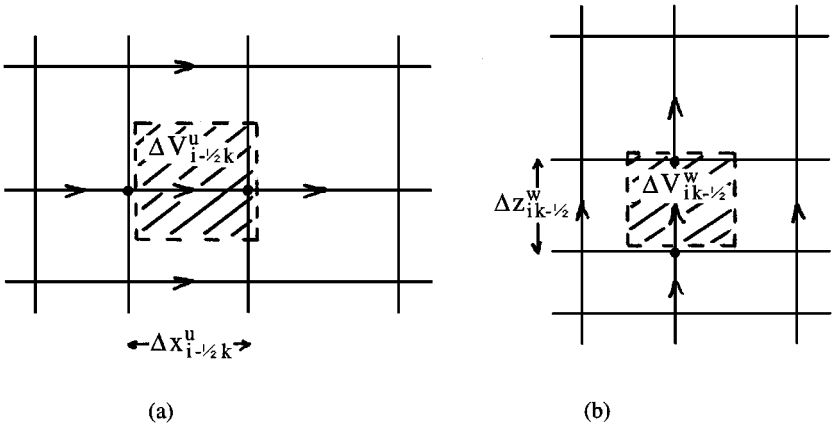

FIG. 2. Staggered grid arrangement and control volumes for (a) the horizontal and (b) the vertical momentum equations.

consideration is denoted ϕ_p, with ϕ_u and ϕ_d denoting the corresponding values at the centres of the cells immediately upwind and downwind, respectively. The estimate on the downwind cell face, ϕ_f, midway between ϕ_p and ϕ_d is based on the ratio of differences,

$$\hat{\phi}_p = \frac{\phi_p - \phi_u}{\phi_d - \phi_u},$$

and is taken to be the harmonic mean of the values given by averaging neighbouring values (central differencing) and second-order upwind interpolation, that is,

$$\phi_f = \frac{1}{2}(\phi_p + \phi_d)\hat{\phi}_p + \frac{1}{2}(3\phi_p - \phi_u)(1 - \hat{\phi}_p).$$

This can be expressed as

$$\phi_f = \phi_p + (\phi_d - \phi_p)\,\hat{\phi}_p,$$

from which the monotonicity property is easily deduced, since $\phi_p \le \phi_f \le \phi_d$ is clearly satisfied if $0 \le \hat{\phi}_p \le 1$. In this form the estimate now consists of the value immediately upstream ("first-order upwinding") plus a second-order flux-limited correction. Values of $\hat{\phi}_p$ outside the range $0 \le \hat{\phi}_p \le 1$ correspond to extreme points in the flow, and only the value immediately upstream is used.

By collecting the convective and diffusive contributions from the four cell faces in this manner Eq. (5) can be expressed in the form

$$a_{pq}\phi_{pq} = \sum a_{mn}\phi_{mn} + S^{\phi}_{pq}, \tag{7}$$

where the summation is taken over the centres of the four neighbouring cells and the multiplying coefficients contain the convective and diffusive flow rates. The source term includes the second-order corrections from the convective scheme (defect correction) and, in the case of both momentum equations, the pressure-gradient term is discretised using two-point differencing and, in the case of the vertical momentum equation, the bouyancy source term, too.

A sequence of grids is defined where the maximum number of lines on grid N_g ($N_g = 1$ corresponds to the finest grid) in the x- and z-directions are $IMAX_{N_g}$ and $KMAX_{N_g}$, respectively. To avoid wasteful use of storage space, single-dimensional arrays are used to store all variables in rows in the x-direction, so that consecutive elements contain horizontal neighbours in general. A pointer system is used to address the elements at position (I, K) on grid N_g of the form

$$M = I + (K - 1) \times IMAX_{N_g} + NG_{N_g},$$

where NG_{N_g} is a summation term over the grids defined by

$$NG_1 = 0, \quad NG_{N_g} = NG_{N_g - 1} + IMAX_{N_g - 1} \times KMAX_{N_g - 1}, \quad N_g > 1.$$

A Cartesian grid is used for the flows over the vertical barrier and care is taken to ensure that the location of the barrier coincides with the same vertical line of horizontal velocities on each of the grids in the sequence. For flows over smooth geometry, however, a simple

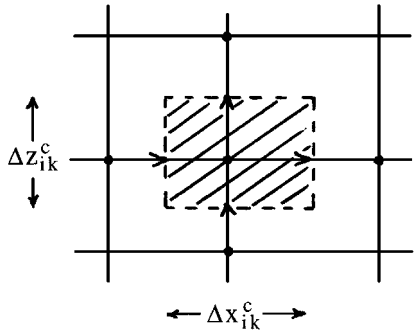

FIG. 3. Control volume for the continuity and density transport equations.

transformation was used to stretch the vertical coordinates. This introduces a volume scale factor and other metric terms into the momentum and scalar transport equations and additional source terms arising from the curvinlinear coordinate derivatives. Care must be taken in this case to ensure that the portion of the lower boundary, where no-slip conditions are applied, is the same on all grids.

Smoothing Algorithms

This paper focuses on two popular current methods for solving the Navier–Stokes equations discretised as above. Discussion of both the decoupled and the coupled methods proceeds best by reference to a typical continuity control volume, Fig. 3. We suppose that the exact discrete values u, w, and p satisfy the discrete steady momentum equations

$$a^u_{i-1/2k}u_{i-1/2k} = \sum a^u_{mn}u_{mn} + A^u_{i-1/2k}(p_{i-1k} - p_{ik}) + S^u_{i-1/2k} \tag{8}$$

$$a^u_{i+1/2k}u_{i+1/2k} = \sum a^u_{mn}u_{mn} + A^u_{i+1/2k}(p_{ik} - p_{i+1k}) + S^u_{i+1/2k} \tag{9}$$

$$a^w_{ik-1/2}w_{ik-1/2} = \sum a^w_{mn}w_{mn} + A^w_{ik-1/2}(p_{ik-1} - p_{ik}) + S^w_{ik-1/2} \tag{10}$$

$$a^w_{ik+1/2}w_{ik+1/2} = \sum a^w_{mn}w_{mn} + A^w_{ik+1/2}(p_{ik} - p_{ik+1}) + S^w_{ik+1/2}, \tag{11}$$

where $A^u_{i-1/2k} = \Delta V^u_{i-1/2k}/\Delta x^u_{i-1/2k}$ etc. and the source terms have been multiplied by the areas of the respective control volumes. The discrete continuity equation is

$$\left(u_{i+1/2k} - u_{i-1/2k}\right)\Delta z^c_{ik} + \left(w_{ik+1/2} - w_{ik-1/2}\right)\Delta x^c_{ik} = 0. \tag{12}$$

In the case of stratified flow the sources in the discrete vertical momentum equations (10), (11) are modified to include the bouyancy source term discretised as

$$\frac{\Delta V^w_{ik-1/2}}{F_h^2}\frac{1}{2}(\vartheta_{ik-1} + \vartheta_{ik}),$$

and the discrete scalar transport equation is solved, in addition,

$$a^\vartheta_{ik}\vartheta_{ik} = \sum a^\vartheta_{mn}\vartheta_{mn} + S^\vartheta_{ik}. \tag{13}$$

Decoupled Smoother (SIMPLE)

The SIMPLE (semi-implicit method for pressure linked equations) pressure correction algorithm [1] has been discussed many times. Briefly, expression (7) comprises a diagonally dominant set for each of the variables u, w, and ϑ. The coefficients in the discrete momentum equations are calculated for the entire domain and, using an estimated pressure field, these are solved in turn via an ADI sweep to yield updates to the current global velocity field. The rest of the iteration, which is the bulk of the computational work, provides the localised coupling. Relations between corrections to velocities and adjacent pressure values (see Fig. 3) are derived from the discrete momentum equations of the forms

$$u'_{i-1/2k} = \frac{A^u_{i-1/2k}}{a^u_{i-1/2K}}(p'_{i-1k} - p'_{ik}), \quad w'_{ik-1/2} = \frac{A^w_{ik-1/2}}{a^w_{ik-1/2}}(p'_{ik-1} - p'_{ik}), \tag{14}$$

and the discrete continuity equation is written terms of velocity corrections,

$$\left(u'_{i+1/2k} - u'_{i-1/2k}\right)\Delta z^c_{ik} + \left(w'_{ik+1/2} - w'_{ik-1/2}\right)\Delta x^c_{ik} = -R^c_{ik}, \tag{15}$$

where R^c_{ik} is the residual of the continuity equation for the current velocity field. Relations (14) are substituted into (15) to derive a Poisson-type equation for the pressure corrections:

$$A_{ik}p'_{ik} - \sum A_{mn}p'_{mn} = -R^c_{ik}. \tag{16}$$

Equation (16) is solved using four ADI sweeps to yield the pressure corrections, which are used to update the velocities through relationships (14). This completes one pressure correction iteration, and in the case of stratified flows the additional transport equation (13) would now be solved. For flows with an internal boundary, such as the vertical barrier, velocities on the barrier are forced to zero by appropriate modification of source terms, and neighbouring equation coefficients are ammended as necessary. Underrelaxation is required for a convergent iteration, with typical values for the momentum and density transport equations being 0.7 and 0.9, respectively. The pressure correction equation itself is not underrelaxed, but only a fraction of the resulting pressure corrections is added to the current pressure field, the value of which is typically 0.3.

The algorithm is summarised as the sequence of the following steps over the entire domain:

1. Calculate the coefficients for each u-momentum equation. Solve the equations globally and update the value of each u-velocity.
2. Calculate the coefficients for each w-momentum equation. Solve the equations globally and update the value of each w-velocity.
3. Calculate the current continuity residuals. Solve the pressure correction equation globally and calculate the corresponding velocity corrections.
4. Update the value of each variable.
5. If stratified calculate the coefficients for the density transport equation. Solve the equations globally and update the value of each density.

Coupled Smoother—Cell-wise (SCGS)

The starting point for Vanka's SCGS (symmetric coupled Gauss Seidel) method [5] is the equation set (8)–(13). The current velocity and pressure fields, denoted here with asterisks,

are deemed to satisfy these equations up to a residual, so that (8), for example, is written

$$a^u_{i-1/2k}u^*_{i-1/2k} = \sum a^u_{mn}u^*_{mn} + A^u_{i-1/2k}(p^*_{i-1k} - p^*_{ik}) + S^u_{i-1/2k} + R^u_{i-1/2k}. \tag{17}$$

Corrections to $u^*_{i-1/2k}$ and p^*_{ik} are sought to produce zero residual, so that

$$a^u_{i-1/2k}u_{i-1/2k} = \sum a^u_{mn}u^*_{mn} + A^u_{i-1/2k}(p^*_{i-1k} - p_{ik}) + S^u_{i-1/2k}, \tag{18}$$

and subtraction of (17) from (18) yields an equation linking local velocity and pressure corrections of the form

$$a^u_{i-1/2k}u'_{i-1/2k} + A^u_{i-1/2k}p'_{ik} = -R^u_{i-1/2k}. \tag{19}$$

The equations for the three other velocities associated with the continuity control volume are derived similarly, and, along with the continuity expression (15), the set of five equations for neutral flow is written in the following matrix form

$$\begin{bmatrix} a^u_{i-1/2k} & 0 & 0 & 0 & A^u_{i-1/2k} \\ 0 & a^u_{i+1/2k} & 0 & 0 & -A^u_{i+1/2k} \\ 0 & 0 & a^w_{ik-1/2} & 0 & A^w_{ik-1/2} \\ 0 & 0 & 0 & a^w_{ik+1/2} & -A^w_{ik+1/2} \\ -\Delta z^c_{ik} & \Delta z^c_{ik} & -\Delta x^c_{ik} & \Delta x^c_{ik} & 0 \end{bmatrix} \begin{bmatrix} u'_{i-1/2k} \\ u'_{i+1/2k} \\ w'_{ik-1/2} \\ w'_{ik+1/2} \\ p'_{ik} \end{bmatrix} = \begin{bmatrix} -R^u_{i-1/2k} \\ -R^u_{i+1/2k} \\ -R^w_{ik-1/2} \\ -R^w_{ik+1/2} \\ -R^c_{ik} \end{bmatrix}. \tag{20}$$

The matrix is inverted analytically by treating it as bordered to yield the required updates which are immediately added to the values of the current solution. Underrelaxation is implemented by adding a fraction of the changes calculated to the respective field variables—typical values for the fraction being in the range 0.7–0.9 for the velocities and 0.8–1.0 for the pressure.

The case of stratified flow would be similar, with the additional variable increasing the dimension of the matrix as

$$\begin{bmatrix} a^u_{i-1/2k} & 0 & 0 & 0 & A^u_{i-1/2k} & 0 \\ 0 & a^u_{i+1/2k} & 0 & 0 & -A^u_{i+1/2k} & 0 \\ 0 & 0 & a^w_{ik-1/2} & 0 & A^w_{ik-1/2} & B^\vartheta_{ik-1/2} \\ 0 & 0 & 0 & a^w_{ik+1/2} & -A^w_{ik+1/2} & B^\vartheta_{ik+1/2} \\ -\Delta z^c_{ik} & \Delta z^c_{ik} & -\Delta x^c_{ik} & \Delta x^c_{ik} & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & a^\vartheta_{ik} \end{bmatrix} \begin{bmatrix} u'_{i-1/2k} \\ u'_{i+1/2k} \\ w'_{ik-1/2} \\ w'_{ik+1/2} \\ p'_{ik} \\ \vartheta'_{ik} \end{bmatrix} = \begin{bmatrix} -R^u_{i-1/2k} \\ -R^u_{i+1/2k} \\ -R^w_{ik-1/2} \\ -R^W_{ik+1/2} \\ -R^c_{ik} \\ -R^\vartheta_{ik} \end{bmatrix}, \tag{21}$$

where $B^\vartheta_{ik-1/2} = \Delta V^w_{ik-1/2}/2F^2_h$ and $B^\vartheta_{ik+1/2} = \Delta V^w_{ik+1/2}/2F^2_h$. This matrix too is easily inverted by first calculating the density changes $\vartheta'_{ik} = -R^\vartheta_{ik}/a^\vartheta_{ik}$ and then rewriting the system as

$$\begin{bmatrix} a^u_{i-1/2k} & 0 & 0 & 0 & A^u_{i-1/2k} \\ 0 & a^u_{i+1/2k} & 0 & 0 & -A^u_{i+1/2k} \\ 0 & 0 & a^w_{ik-1/2} & 0 & A^w_{ik-1/2} \\ 0 & 0 & 0 & a^w_{ik+1/2} & -A^w_{ik+1/2} \\ -\Delta z^c_{ik} & \Delta z^c_{ik} & -\Delta x^c_{ik} & \Delta x^c_{ik} & 0 \end{bmatrix} \begin{bmatrix} u'_{i-1/2k} \\ u'_{i+1/2k} \\ w'_{ik-1/2} \\ w'_{ik+1/2} \\ p'_{ik} \end{bmatrix} = \begin{bmatrix} -R^u_{i-1/2k} \\ -R^u_{i+1/2k} \\ -R^w_{ik-1/2} - B^\vartheta_{ik-1/2}\vartheta'_{ik} \\ -R^w_{ik+1/2} - B^\vartheta_{ik+1/2}\vartheta'_{ik} \\ -R^c_{ik} \end{bmatrix}, \tag{22}$$

so that the only difference from the case of neutral flow is a change to the residual terms on the right-hand side of the vertical momentum equations. In fact the cell-wise SCGS method was never used for stratified flows over obstacles, but the principle has been described to illustrate the procedure for the linewise approach which was used and is discussed later.

The algorithm for neutral flow can be summarised as the sequence of the following steps for each continuity control volumes in the sweep:

1. Calculate the coefficients and the residuals of the u-momentum equation for both u-velocities.
2. Calculate the coefficients and the residuals of the w-momentum equation for both w-velocities.
3. Calculate the current continuity residual.
4. Invert the matrix and update the values of all variables.

Although the philosophy of the SCGS method is entirely opposite to that of the SIMPLE method, there are similarities between the two. For example, both methods share the feature that the velocities are updated twice while the pressures are updated once, and that local pressure/velocity coupling is achieved by a similar relationship in each case which is based on the discrete momentum equations (compare (19) with (14)). In the case of SCGS, however, immediate local linkage is provided between variables and the global coupling is achieved by sweeping through the domain in a prescribed manner, such as forward lexicographic ordering followed by backward. Continuity control volumes are visited in turn and all the coefficients of the transport equations are recalculated to take account of the updated values of flow variables. The bulk of the work is in assembling the equation coefficients needed to calculate the current value of the residuals, and this outweighs the overhead of inverting the matrix. An indication of the work required is obtained by considering the number of flux evaluations needed to calculate a set of updates for a single continuity control volume. For a two-dimensional rectangular grid the horizontal momentum equation at two neighbouring points requires flux evaluations at each of the four faces of the two respective control volumes. One of these faces is shared, however, so the total number of evaluations is seven. Correspondingly, the vertical momentum equation also requires seven, so the total number of flux evaluations for one set of updates is 14. For a square grid containing n^2 control volumes, therefore, $14n^2$ flux evaluations are required for one sweep across the grid, while for the usual forward and backward sweep, this doubles to $28n^2$. This is considerably greater than the number of flux evaluations required for one sweep of the decoupled method, where, since the values of all flow variables are considered to be at the same stage of the computation, all cell faces between neighbouring control volumes can be considered shared. On the same grid, therefore, each momentum equation requires only $2n^2$ flux evaluations, combining to give a total of $4n^2$, a factor of seven fewer than for the coupled method.

Coupled Smoother—Line-wise (CLGS)

Solving the coupled system in a line-wise fashion can be implemented in several ways. The SCGS/LS algorithm of Shah *et al.* [29], for example, solves only for the pressures along gridlines and this appeared to be more efficient than the cell-wise approach for the driven cavity problem. Although this algorithm was applied successfully to flow over a backward facing step [30], the more conventional view [31] is that for unidirectional flows all variables should be updated simultaneously in lines, as achieved in the SCAL (symmetrically coupled

alternating line) [12], CLGS (collective line Gauss Siedel) [23] algorithms, as well as that described (but only briefly) in [22]. SCAL differs from CLGS only in the sweeping pattern used to produce updates—alternating zebra in the former and line-Gauss-Seidel in the latter. Although algorithms of both these kinds have been tried in the present work, the best has been found to be the CLGS method implemented by sweeping in vertical lines, and the description that follows is in these terms.

With reference to Fig. 4, the variables along the i^{th} vertical line are grouped in sets of four and are written as the vector

$$\left[\ldots, \left(w_{ik-3/2}, u_{i-1/2k-1}, u_{i+1/2k-1}, p_{ik-1}\right), \left(w_{ik-1/2}, u_{i-1/2k}, u_{i+1/2k}, p_{ik}\right), \left(w_{ik+1/2}, u_{i-1/2k+1}, u_{i+1/2k+1}, p_{ik+1}\right), \ldots\right]^T.$$

The form of the discrete update equations for the vertical momentum equation changes slightly with vertically adjacent pressures now linked, while the form for the horizontal momentum equation remains unchanged. In the case of neutral flow the equations along a vertical line to be solved simultaneously have the following form, with the right-hand side given by the vector of corresponding residuals:

$$\left(\begin{array}{cccc|cccc|cccc}
a^w_{ik-3/2} & 0 & 0 & +A^w_{ik-3/2} & 0 & 0 & 0 & 0 & & & & \\
0 & a^u_{i-1/2k-1} & 0 & A^u_{i-1/2k-1} & 0 & 0 & 0 & 0 & & & & \\
0 & 0 & a^u_{i+1/2k-1} & -A^u_{i+1/2k-1} & 0 & 0 & 0 & 0 & & & & \\
-\Delta x^c_{ik-1} & -\Delta z^c_{ik-1} & \Delta z^c_{ik-1} & 0 & \Delta x^c_{ik-1} & 0 & 0 & 0 & & & & \\
\hline
0 & 0 & 0 & -A^w_{ik-1/2} & a^w_{ik-1/2} & 0 & 0 & +A^w_{ik-1/2} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & a^u_{i-1/2k} & 0 & A^u_{i-1/2k} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & 0 & a^u_{i+1/2k} & -A^u_{i+1/2k} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & -\Delta x^c_{ik} & -\Delta z^c_{ik} & \Delta z^c_{ik} & 0 & \Delta x^c_{ik} & 0 & 0 & 0 \\
\hline
 & & & & 0 & 0 & 0 & -A^w_{ik+1/2} & a^w_{ik+1/2} & 0 & 0 & +A^W_{ik+1/2} \\
 & & & & 0 & 0 & 0 & 0 & 0 & a^u_{i-1/2k+1} & 0 & A^u_{i-1/2k+1} \\
 & & & & 0 & 0 & 0 & 0 & 0 & 0 & a^u_{i+1/2k+1} & A^u_{i+1/2k+1} \\
 & & & & 0 & 0 & 0 & 0 & -\Delta x^c_{ik+1} & -\Delta z^c_{ik+1} & \Delta z^c_{ik+1} & 0
\end{array}\right)$$

$$\times \begin{pmatrix} \vdots \\ w'_{ik-3/2} \\ u'_{i-1/2k-1} \\ u'_{i+1/2k-1} \\ p'_{ik-1} \\ w'_{ik-1/2} \\ u'_{i-1/2k} \\ u'_{i+1/2k} \\ p'_{ik} \\ w'_{ik+1/2} \\ u'_{i-1/2k+1} \\ u'_{i+1/2k+1} \\ p'_{ik+1} \\ \vdots \end{pmatrix} = \begin{pmatrix} \vdots \\ -R^w_{ik-3/2} \\ -R^u_{i-1/2k-1} \\ -R^u_{i+1/2k-1} \\ -R^c_{ik-1} \\ -R^w_{ik-1/2} \\ -R^u_{i-1/2k} \\ -R^u_{i+1/2k} \\ -R^c_{ik} \\ -R^w_{ik+1/2} \\ -R^u_{i-1/2k+1} \\ -R^u_{i+1/2k+1} \\ -R^c_{ik+1} \\ \vdots \end{pmatrix}^T.$$

The structure of this matrix is block-triadiagonal and is solved using a specially written routine based on the Thomas algorithm. In terms of flux evaluations, more cell faces can be considered shared than before and the work count reduced, compared with the cell-wise version. In the vertical implementation described here, for example, the horizontal

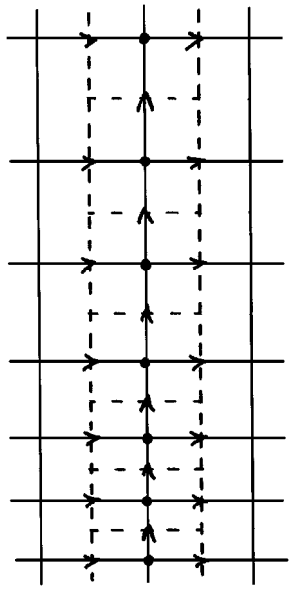

FIG. 4. Variables to be updated when the equations are coupled in vertical lines.

momentum equation is required in two neighbouring lines of n cells. Of the 8n cell faces involved, 3n are shared, so only 5n flux evaluations are required. The vertical momentum equation, on the other hand, is required at points in a single line of n cells, requiring 3n flux evaluations. Combining these gives 8n evaluations for a set of changes along a line of cells and, hence, $8n^2$ for the complete grid, or $16n^2$ for a forward sweep followed by a backward sweep. The work required to solve the matrix system of equations is about the same as that of the cell-wise approach (46 floating point operations per continuity control volume in each case), but the reduced number of flux evaluations represents a saving over the coupled method. In practise, however, the benefit of this is offset by the fact that array variables are required to store the coefficients in the line-wise implementation, with additional associated computational cost, while they are not for the cell-wise.

In the case of density-stratified flow, the method is modified in a manner similar to that described for the cell-wise version. The buoyancy source term in each update equation for the vertical momentum includes vertically adjacent densities while the update equations for the densities themselves form a tridiagonal set. The structure of the middle row of the block matrix above becomes

$$\begin{pmatrix}
0 & 0 & 0 & -A^{w}_{ik-1/2} & B^{\vartheta}_{ik-1/2} & | & a^{w}_{ik-1/2} & 0 & 0 & +A^{w}_{ik-1/2} & B^{\vartheta}_{ik-1/2} & | & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & | & 0 & a^{u}_{i-1/2k} & 0 & A^{u}_{i-1/2k} & 0 & | & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & | & 0 & 0 & a^{u}_{i+1/2k} & -A^{u}_{i+1/2k} & 0 & | & 0 & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & 0 & | & -\Delta x^{c}_{ik} & -\Delta z^{c}_{ik} & \Delta z^{c}_{ik} & 0 & 0 & | & \Delta x^{c}_{ik} & 0 & 0 & 0 & 0 \\
0 & 0 & 0 & 0 & a^{\vartheta}_{ik-1} & | & 0 & 0 & 0 & 0 & a^{\vartheta}_{ik} & | & 0 & 0 & 0 & 0 & a^{\vartheta}_{ik+1}
\end{pmatrix}.$$

The density updates are easily calculated and used to modify the residuals in the vertical momentum equation as before, and the solution proceeds as for homogenous flow. Under-relaxation is again achieved by adding a fraction of the computed changes as described for the cell-wise version.

The algorithm can be summarised as the sequence of the following steps for each line of continuity control volumes in the sweep:

1. Calculate the coefficients and the residuals for the u-momentum equations for each u-velocity.
2. Calculate the coefficients and the residuals for the w-momentum equations for each w-velocity.
3. Calculate each current continuity residuals.
4. If stratified calculate the coefficients and the residuals for the density transport equation for each densities.
5. Invert the block matrix and update the values of all variables.

4. MULTIGRID ALGORITHM

The procedures used were given in [18] and only a brief outline is given here. Denoting the linear and nonlinear discrete operators on the fine grid to be L_h and N_h, respectively, the discrete horizontal momentum equation (8), for example, can be written

$$N_h u_h + L_h p_h = 0. \tag{23}$$

The current approximations $\tilde{u}_h$ and $\tilde{p}_h$ satisfy Eq. (23) to the extent of a residual $\tilde{R}_h$,

$$N_h \tilde{u}_h + L_h \tilde{p}_h = \tilde{R}_h, \tag{24}$$

and subtraction of (24) from (23) yields

$$N_h u_h = N_h \tilde{u}_h - L_h p'_h - \tilde{R}_h, \tag{25}$$

where $p'_h = p_h - \tilde{p}_h$. Equation (25) is the basis of the coarse grid equations. Coarse grids are generated so that a coarse grid continuity cell is the sum of four fine grid cells—see Fig. 5. Flow variables and residuals are transferred to the coarse grid via a restriction operator I_h^{2h} (weighted means for flow variables and area-scaled sums for residuals) so that the coarse

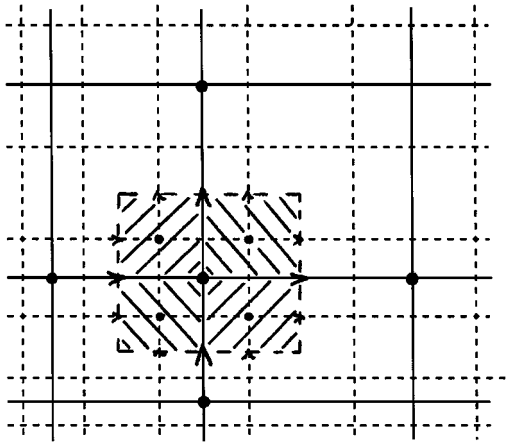

FIG. 5. Relationship of fine and coarse grids—a continuity control volume on the coarse grid is composed of four on the fine grid.

grid momentum equations take the form

$$N_{2h}u_{2h} = N_{2h}I_h^{2h}\tilde{u}_h - L_{2h}p'_{2h} - I_h^{2h}\tilde{R}_h. \tag{26}$$

Starting from the initial conditions $u_{2h} = I_h^{2h}\tilde{u}_h$ and $p'_{2h} = 0$, coarse grid iteration using any of the smoothers yields new coarse grid velocities u_{2h} and coarse grid pressure corrections p'_{2h}. The changes in the coarse grid velocities and the pressure corrections are then tranferred back to the fine grid using prolongation operator I_{2h}^h (bilinear interpolation) and the fine grid solution is updated as

$$u_h = \tilde{u}_h + I_{2h}^h\left(u_{2h} - I_h^{2h}\tilde{u}_h\right), \quad p_h = \tilde{p}_h + I_{2h}^h p'_{2h}. \tag{27}$$

Convergence on the fine grid implies that the residual forcing term in (26) is zero, and the equation is satisfied by $u_{2h} = I_h^{2h}\tilde{u}_h$ and $p'_{2h} = 0$, in which case the calculated changes are zero and interpolation leaves the fine grid solution unchanged.

Both the decoupled and coupled smoothers proceed on the coarse grid in a manner similar to the implementation on the fine grid described in the previous section. Each procedure needs the modification to allow the use of curvilinear grids described in [18]. The continuity expression corresponding to (15) for coarse grids should be

$$C_{2h}u'_{2h} = -R^c_{2h}, \tag{28}$$

where the right-hand side is given by

$$R^c_{2h} = C_{2h}u^*_{2h} - C_{2h}I_h^{2h}\tilde{u}_h + I_h^{2h}R^c_h, \tag{29}$$

where the first term on the right-hand side represents the continuity residual of the current coarse grid velocity field. As explained in [18], when Cartesian grids are used, the restriction operator preserves mass fluxes, so that the last two terms here cancel. When grids are curvilinear, however, mass fluxes are not preserved, and all three terms must be evaluated. When convergence is achieved on the fine grid (the last term on the right-hand side is zero), no changes result from the coarse grid momentum equations, leaving the continuity residuals unchanged from their values calculated after restriction (the second term on the right-hand side), and hence, a zero right-hand side as required.

As mentioned in the Introduction, convergence for stratified flows was only obtained in some cases by utilising a technique known as double discretisation [24], where the discrete operator used to approximate the differential equation is kept distinct from that used to smooth the solution to the discrete equations so generated. That this is feasible arises from the observation that the solution to the discrete fine grid equations is governed solely by the operator used to generate the fine grid residuals. In a multigrid scheme these residuals are restricted to the coarse grid, but the discrete operator of the coarse grid equation—the N_{2h} on each side of (26)—may be chosen in principle to be anything that will smooth the error. This idea can be applied on the fine grid too, as part of a multigrid cycle but not necessarily so. All that is required is that the fine grid equation is written as in (25), so that the residuals on the right-hand side are generated using the scheme desired for the fine grid solution, while the operator N_h may be chosen to be anything else. What is typically chosen here is to generate the fine grid residuals for all transport equations using the Van Leer scheme described earlier, and then to use a much more dissipative scheme such

as first-order upwinding in the coarse (or fine) grid density equation. Clearly there is no advantage if a convergent iteration can be obtained without doing this, but in other cases this technique can be used to generate a multigrid iteration which converges to a solution of desirable accuracy.

5. RESULTS

Lid-driven Cavity

The relative merits of the smoothing algorithms were first tested using the standard problem of the lid-driven cavity. Although only a model problem with many idealisms not found in practical flows, this exercise enables comparison to be made with previous data and the present implementations to be validated.

Prior to the discussion of the multigrid convergence rates, however, we give an indication of the accuracy of the discretisation scheme. Although the solution to the lid-driven cavity problem has been given in detail elsewhere, a simple measure sometimes used to quantify the overall accuracy of a solution is the maximum negative velocity on the centre-line of the cavity. Figure 6 shows this quantity obtained in the computations here plotted against the square of the grid spacing for the two cases $\mathrm{Re} = 100$ and $\mathrm{Re} = 1000$, and four grids, with 32×32, 64×64, 128×128, and 256×256 nodes. Shown for comparison are Vanka's [5] data for the hybrid discretisation on staggered grids as here, and Lien and Leschziner's [16] data for the MUSCL and QUICK schemes on a collocated grid. The higher order schemes demonstrate second-order behaviour, while the hybrid scheme clearly does not. Of the higher order schemes, Van Leer's harmonic scheme would appear to be less dissipative than the others in that the peak values attained with this scheme are consistently higher than those of the other schemes for both Reynolds numbers.

The convergence rate of the SIMPLE method has been compared with four implementations of the coupled approach. These are SCGS, CLGS implemented in vertical lines (denoted CLGS-V), CLGS implemented in horizontal lines (denoted CLGS-H), and SCAL. Forward and backward lexicographic sweeping is used for SCGS, while symmetric sweeping is used for CLGS (that is, a sweep in one direction followed by another in the reverse direction). The order of the sweeping in SCAL is even horizontal lines then odd horizontal lines followed by even vertical lines then odd vertical lines. To facilitate the comparison of the rates of convergence with the work of other authors we employ the commonly used convergence criterion that the L_2 residual,

$$\mathrm{R} = \left[\frac{\sum (\mathrm{R}^{\mathrm{u}})^2 + \sum (\mathrm{R}^{\mathrm{w}})^2 + \sum (\mathrm{R}^{\mathrm{c}})^2}{3 \times \mathrm{IMAX}_{\mathrm{N_g}} \times \mathrm{KMAX}_{\mathrm{N_g}}} \right]^{1/2},$$

reach 10^{-4}. (Here the individual equation residuals are normalised by control volume areas so that the dimensions are those of the original continuous equations.) Convergence data for the methods is summarised in Table 1, where the number of cycles required for convergence is given, along with the computing time taken on a Silicon Graphics Indy workstation in appropriate units. Identical cycling patterns were used in each case, which was not the case in either of the comparisons [21, 22], namely fixed W-cycles, with the equations on the coarsest grid solved using 5 cycles with $\mathrm{Re} = 100$, 10 with $\mathrm{Re} = 400$, and 15 with $\mathrm{Re} = 1000$. One postsmoothing and one presmoothing iteration were used, with residuals reevaluated prior to restriction.

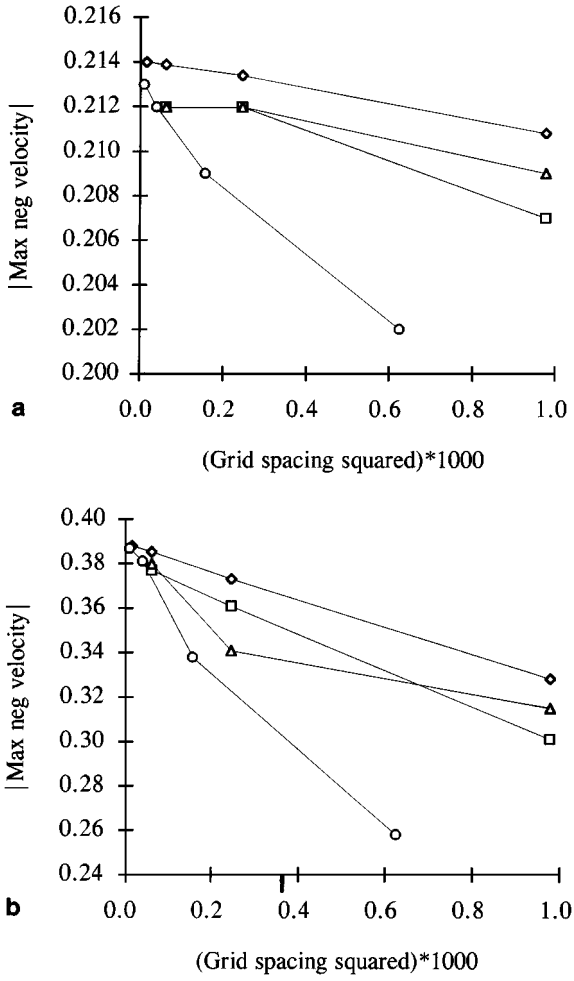

FIG. 6. Maximum negative velocity on the centre line of flow in the driven cavity. (a) Re = 100, (b) Re = 1000. ◇, Van Leer (present scheme); □, MUSCL (Lien & Leschziner, 1994); △, QUICK (Lien & Leschziner 1994); ○, Hybrid (Vanka 1986).

It is immediately seen from Table 1 that all the coupled methods require fewer cycles to meet the convergence criterion than SIMPLE. On the other hand, the coupled methods are more expensive per cycle, with the respective times per cycle for SIMPLE : SCGS : CLGS-V : CLGS-H : SCAL being in the approximate ratios 1.0 : 1.6 : 1.6 : 1.2 : 1.4. The computing time per cycle for CLGS-V is almost identical to that of SCGS. Note, however, that the computing time per cycle for the implementation with horizontal line sweeping is approximately 25% less than for vertical sweeping, despite the equivalence in the number of arithmetic operations. This is presumably a consequence of the method of data storage, where single-dimensional arrays are used to store all variables in horizontal rows (consecutive elements contain horizontal neighbours) as discussed in Section 3. Since SCAL uses both vertical and horizontal sweeps, the cycle time for SCAL lies between that for CLGS-V and CLGS-H, meaning that both SCAL and CLGS-H are cheaper than SCGS. This is in contrast to the findings in [12], where SCAL is reported to require 50% more computing time than SCGS.

TABLE 1
Numbers of Multigrid Cycles and Computing Times Taken to Reach Convergence for Different Smoothing Algorithms Applied to the Driven Cavity Problem

Re = 100	SIMPLE	SCGS	CLGS-V	CLGS-H	SCAL
32 × 32	9 (3.8s)	5 (3.7s)	5(3.7s)	5 (3.1s)	5 (3.4s)
64 × 64	8 (14.1s)	4 (11.8s)	4 (12.0s)	5 (11.8s)	5 (13.2s)
128 × 128	9(1m 9s)	4 (48.2s)	4 (50.2s)	5 (47.7s)	5 (54.7s)
Re = 400	SIMPLE	SCGS	CLGS-V	CLGS-H	SCAL
32 × 32	12 (5.0 s)	7 (6.6s)	8 (5.7s)	8 (4.8s)	7 (5.4s)
64 × 64	10 (17.6s)	5 (15.8s)	6 (17.8s)	7 (17.7s)	6 (17.2s)
128 × 128	9 (1m 9s)	5 (1m 2s)	5 (1m 3s)	6 (57.2s)	5 (56.7s)
Re = 1000	SIMPLE	SCGS	CLGS-V	CLGS-H	SCAL
32 × 32	17 (7.0s)	10 (7.1s)	14 (9.7s)	12 (8.1s)	15 (11.0s)
64 × 64	13 (22.8s)	8 (26.8s)	10 (29.5s)	12 (30.1s)	11 (31.4s)
128 × 128	11(1m 24s)	6 (1m 17s)	7 (1m 27s)	10(1m 39s)	7 (1m 19s)

The advantage of the coupled methods over SIMPLE is greatest at the lowest Reynolds number, which concurs with the results of [21, 22], although the advantage found here is rather less than in either of these other comparisons. The reasons for this may include the fact that the cycling patterns used here are identical, but there are other more definite reasons. First, the differencing scheme used here is more complicated and takes more arithmetic operations to implement than the hybrid scheme used in these other two comparisons. From the discussion in Section 3, it is clear that the coupled schemes require many more flux evaluations than does SIMPLE, and the relative costs of the additional work are bound to favour the scheme which requires fewest flux evaluations. Second, although the grids used are uniform and square for this problem, the code retains the capacity to treat grids which are not uniform. Array variables are therefore required for the dimensions of the different control volumes, instead of the one real scalar variable which would otherwise have sufficed and which would have surely been used in the other comparisons. The overheads associated with this would also favour the scheme with fewest flux evaluations.

The convergence histories for the five smoothers are compared in Fig. 7 for the computations on the 128 × 128 grid and the asymptotic convergence rates for these calculations are given in Table 2 (determined in each case by the average residual reduction over the final four cycles). As expected, convergence rates fall in each case with increasing Reynolds number, although the deterioration between Re = 100 and Re = 400 is slight, except for

TABLE 2
Asymptotic Convergence Rates (Residual Reductions per Cycle) for the Different Smoothing Algorithms Applied to the Driven Cavity Problem

Re	SIMPLE	SCGS	CLGS-V	CLGS-H	SCAL
100	0.391	0.084	0.117	0.144	0.136
400	0.397	0.112	0.202	0.360	0.119
1000	0.495	0.309	0.356	0.430	0.371

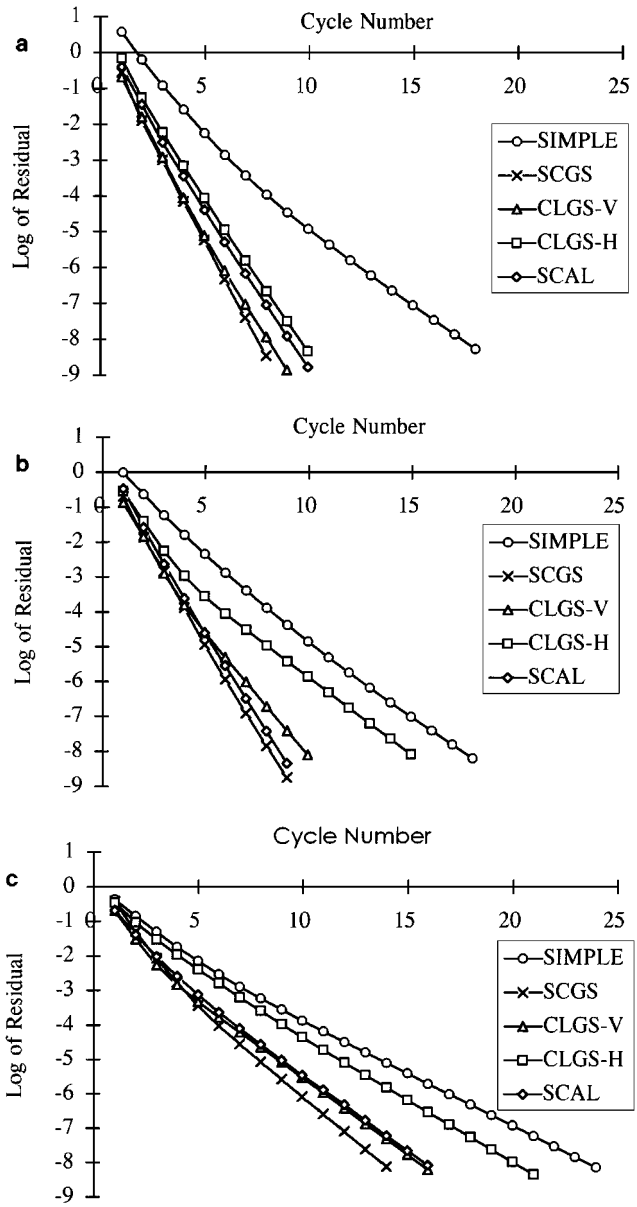

FIG. 7. Convergence histories of the computations of flow in the driven cavity. (a) Re = 100, (b) Re = 400, (c) Re = 1000.

CLGS-H. All the coupled methods have better convergence rates than the decoupled method SIMPLE and, of the coupled methods, SCGS consistently has the best convergence rate over CLGS-V, which in turn is better than SCAL. The worst of the coupled smoothers is CLGS-H.

Convergence rates are often quoted in terms of the practical smoothing rate λ, determined from the residual reduction per cycle κ by $\lambda = \kappa^{1/(\nu_1+\nu_2)}$, where ν_1 and ν_2 are the numbers of presmoothing and postsmoothing iterations, respectively, which are both unity in these computations. The asymptotic practical smoothing rates corresponding to these convergence

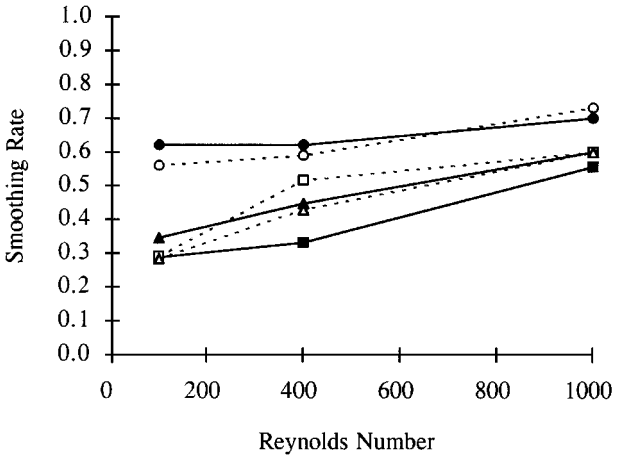

FIG. 8. Comparison of practical smoothing rates for the computation of flow in the driven cavity. ●, SIMPLE (Present); ○, MGPC (Sivaloganathan *et al.*, 1988); ■, SCGS (Present); □, BLIMM (Sivaloganthan *et al.*, 1988); ▲, CLGS (Present); △, SCGS/LS (Shah *et al.*, 1990).

rates are compared with those presented by other authors for this problem in Fig. 8. The rates obtained here with the SIMPLE and SCGS methods are compared with those in [21] for computations, which, ostensibly, are identical implementations of the methods used here. The rates for the vertical line solver CLGS-V are compared with those of the SCGS/LS algorithm [29]. These are not equivalent implementations for CLGS solves for all variables in lines, while SCGS/LS solves only for the pressures. It is clear from Fig. 8 that very similar convergence rates are being obtained in the present implementations to these previous results. This is significant, bearing in mind that the present results are asymptotic rates (it is not clear that this is so for the other data) and, more importantly, that the present results are for a convective discretisation scheme which is very much less disipative than the hybrid scheme used in all these other implementations.

The conclusion from this section is that the coupled methods certainly have superior convergence rates to those of the decoupled methods and, despite the greater computational cost per cycle, the computing times are superior to those of the decoupled method, especially at low Reynolds numbers. The line solver CLGS performs best when implemented in vertical lines. If the data structure were changed so that variables were stored in vertical rows, then the times for vertical sweeping in Table 1 would be around 25% less than the values given. This would then mean that the line solver with vertical sweeping would be the fastest method for this problem at all Reynolds numbers tried. Whether the effort of doing this would be repaid depends upon the relative performance obtained when applied to the computation of flows over obstacles, to which we now turn.

Neutral Flow Past Obstacles

In this section we compare the convergence rates of the multigrid methods with the different smoothers for laminar flow past two obstacles. High Reynolds number turbulent flows clearly have more application in practise and, indeed, a simple turbulence model was included in some of the previous computations [18]. However, it was felt that the scope of the current investigation was already wide enough without including turbulence here,

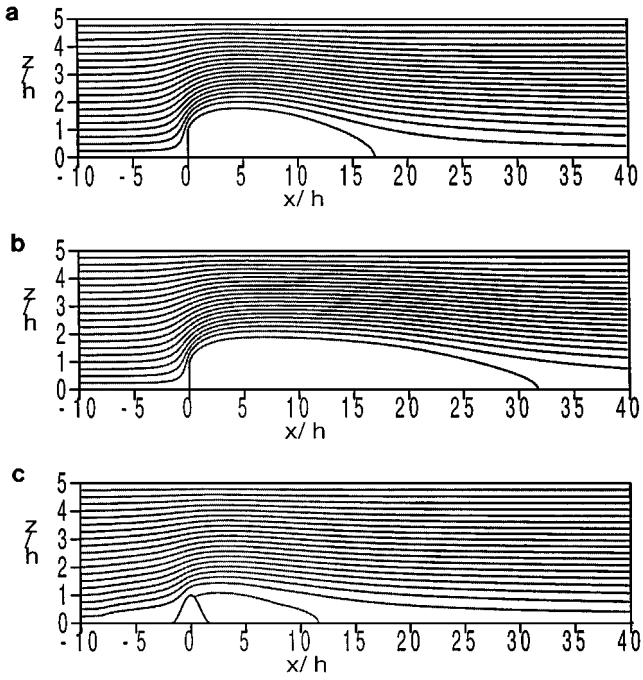

FIG. 9. Streamlines for neutral flow over obstacles: (a) the vertical barrier, Re = 50; (b) the vertical barrier, Re = 100; (c) the cosine obstacle, Re = 100.

and in fact, we restrict ourselves to relatively low Reynolds numbers. Although the performance for laminar flows at higher Reynolds numbers can clearly be tested, the results for the driven cavity problem above indicate that any advantage of the coupled method is likely to be greatest at low Reynolds number. There is a practical difficulty, too, in that for laminar flow the downstream extent of the recirculation region behind the obstacle grows linearly with increasing Reynolds number. We have, therefore, performed computations for flow over the vertical barrier at Re = 50 and Re = 100, and flow over the cosine hill at Re = 100. A set of five grids was used to cover the domain $[-20, 60] \times [0, 5]$, the finest of which contains 320×80 cells and the coarsest 20×5. The grids are nonuniform, but with expansion ratios nowhere exceeding 1.05. Streamlines for neutral flow over the vertical barrier and the cosine-shaped hill corresponding to the flows computed on the finest grid are shown in Fig. 9.

The multigrid computations proceed in a manner similar to those for the driven cavity, with the addition of more complicated boundary conditions. Fixed W-cycles are used as before, with one presmoothing, one postsmoothing iteration, 10 iterations on the finest grid and up to five multigrid levels. Initial computations were performed for flow over the vertical barrier, Re = 50, with the SIMPLE method and the three line-wise coupled smoothers CLGS-V, CLGS-H, and SCAL. The convergence histories using the single L_2 residual for these computations are shown in Fig. 10. CLGS-H performed the least well of the coupled methods for the driven cavity, and it is clearly performing very poorly for this problem. It would appear that sweeping in horizontal lines when the flow is aligned in that direction does not achieve effective smoothing. The incorporation of horizontal sweeping in a method utilising alternating directions presumably contributes to the relatively poor

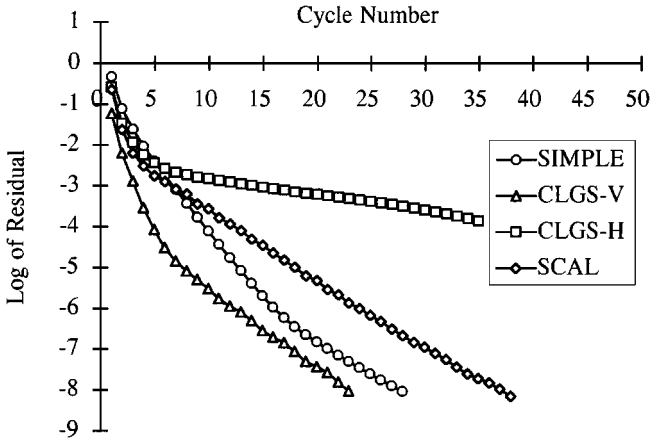

FIG. 10. Comparison of the convergence rates of SIMPLE with three coupled line solvers. Neutral flow over the vertical barrier, Re = 50.

performance of SCAL too. Sweeping across the flow direction, on the other hand, appears to be the most effective strategy for the coupled methods used here, and for this problem CLGS-V has the best asymptotic convergence rate, although it is only marginally better than that of the SIMPLE method. The cycle times of the coupled methods for these flows are somewhat longer relative to that of SIMPLE when compared with the ratios given earlier and for the four methods SIMPLE : CLGS-V : CLGS-H : SCAL are in the ratios 1.0 : 2.2 : 1.6 : 1.8. Although horizontal sweeping is again somewhat cheaper than vertical sweeping, the relatively poor convergence properties of CLGS-H and SCAL meant that no further use was made of them. The relative merits of SIMPLE and CLGS-V are discussed next in more detail, however.

Table 3 gives the performance of the SIMPLE method and CLGS implemented in vertical lines for these computations in terms of the number of multigrid cycles required to reach a prescribed level of convergence and the computing time taken. Convergence was monitored using sums of absolute values of residuals over the whole domain as well as the single L_2 residual described earlier. (The former residuals represent the net flux imbalance for each control volume and are not normalised by the control volume areas—because of this the value of the L_2 residual is usually at least an order of magnitude smaller than any of the sums of absolute values.) The convergence criterion used in Table 3 is that the maximum sum of absolute values over the three equations falls below 10^{-4}, the same criterion as [18]. It is seen that a reasonable level of grid independence is achieved for both smoothers, with

TABLE 3
Numbers of Multigrid Cycles and Computing Times Taken to Reach Convergence for Neutral Flow Past Obstacles

	Barrier, Re = 50		Barrier, Re = 100		Cosine, Re = 100	
	SIMPLE	CLGS	SIMPLE	CLGS	SIMPLE	CLGS
80 × 20	12 (9.3s)	11 (20.4s)	21 (20.0s)	20 (49.0s)	15 (17.2s)	13 (37.4s)
160 × 40	10 (30.1s)	9 (1m 3s)	15 (51.7s)	11 (1m 30s)	16 (1m 16s)	10 (1m 52s)
320 × 80	10 (2m 12s)	11 (5m 13s)	12 (2m 48s)	14 (7m 12s)	15 (5m 6s)	10 (7m 44s)

computing times increasing by approximately a factor of four as the number of grid cells is doubled. For the vertical barrier the number of cycles tends to rise with the Reynolds number, as might be expected. It is also noticable that for the vertical barrier the number of cycles required is roughly the same for SIMPLE as for CLGS and neither method establishes an advantage in terms of convergence rate. The longer time per cycle for CLGS therefore means that the SIMPLE method generally requires less than half the computing time. This is in contrast to the results for the driven cavity, where CLGS had a very much better convergence rate, leading to shorter computing times than SIMPLE, despite the greater cost per cycle.

For the flow over the cosine hill, on the other hand, CLGS is consistently better than SIMPLE in terms of convergence rate, requiring only two-thirds as many cycles for convergence on the finer grids as SIMPLE. The comparative cycle times for the two methods are still in the approximate ratio SIMPLE : CLGS of 1.0 : 2.2, however, and the shorter time per cycle for SIMPLE means that CLGS takes at least 50% longer. The full convergence histories are shown in Fig. 11 for both flows with Re = 100, where the residual used now is the L_2 norm used before. For the flow over the vertical barrier it can be seen

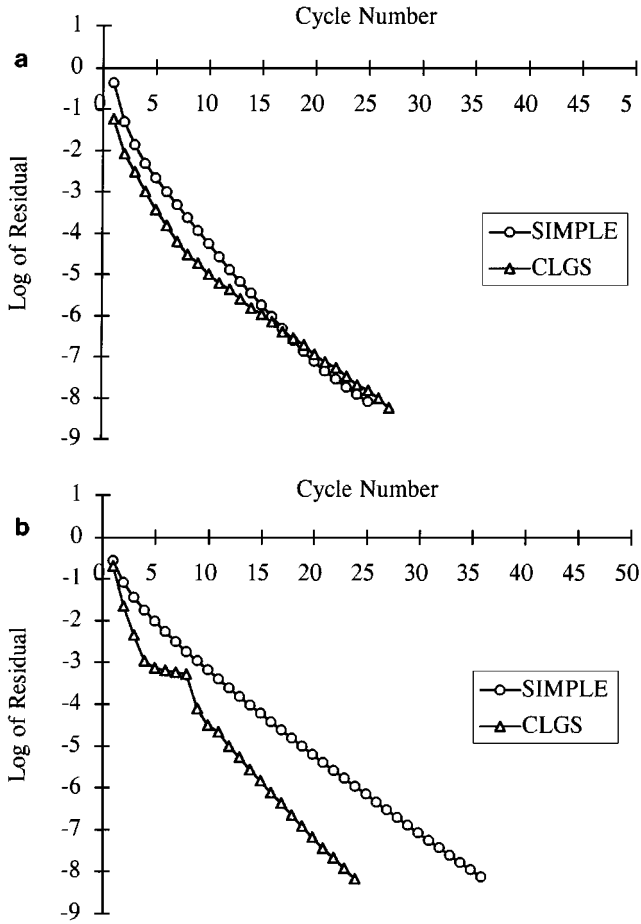

FIG. 11. Convergence histories of the computations of neutral flow over obstacles with Re = 100: (a) The vertical barrier; (b) the cosine obstacle.

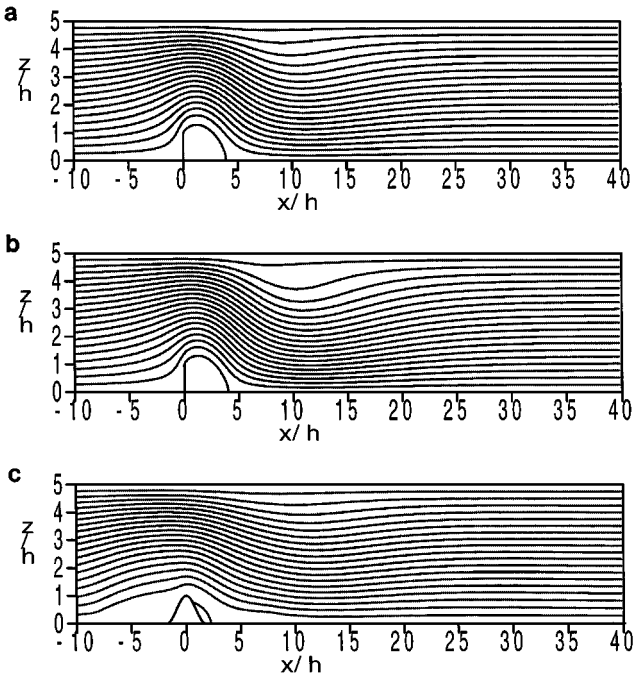

FIG. 12. Streamlines for stratified flow over obstacles with $F_h = 1.592$ ($K = 1.0$) and the hybrid scheme used for the density equation: (a) the vertical barrier, $Re = 50$; (b) the vertical barrier, $Re = 100$; (c) the cosine obstacle, $Re = 100$.

that the SIMPLE method actually has the better asymptotic rate of convergence, while for the flow over the cosine hill the asymptotic rate of CLGS is significantly better than that of SIMPLE.

Stratified Flow Past Obstacles

The effects on the flow field of reducing the Froude number so that $K\ (= D/\pi h F_h)$ ranges from zero to unity have been described before [18, 26] and, as mentioned earlier, essentially consist of a reduction in the separation length from its value in neutral flow and a corresponding reduction in the value of the pressure drag. The former effect is clearly seen in streamlines for the case $K = 1.0$ (Fig. 12) when compared with their counterparts in neutral flow (Fig. 9). It was shown in [18] that the SIMPLE method was an effective smoother over the range $0 \leq K \leq 1$ ($1.592 \leq F_h \leq \infty$) for laminar and turbulent flow over the two obstacles, with computations performed at Froude numbers $F_h = 3.180$ ($K = 0.5$), $F_h = 1.989$ ($K = 0.8$), and $F_h = 1.592$ ($K = 1.0$). Convergence was achieved so that sums of absolute values of residuals were reduced to at least 10^{-4} in all cases tried with the exception of laminar flow ($Re = 100$) over the cosine hill when $K = 1.0$. This was explained by the suggestion that this value of K marks the transition to unsteadiness and perhaps convergence difficulties therefore might be expected. That they were not encountered for the barrier, when the Reynolds number was lower ($Re = 50$), also suggests the possibility that discrete ellipticity may be lost when Sc.Re combination in the density transport equation becomes large.

The CLGS method with the addition of the modification for the density equation was used with vertical sweeping in an attempt to provide a comparison with the performance of

TABLE 4
Number of Multigrid Cycles and Computing Times Taken for Convergence for Stratified Flow Past Obstacles (F_h = 1.592, K = 1.0) with the Hybrid Scheme Used for the Density Equation

	Barrier, Re = 50		Barrier, Re = 100		Cosine, Re = 100	
	SIMPLE	CLGS	SIMPLE	CLGS	SIMPLE	CLGS
80 × 20	14 (16.2s)	13 (37.6s)	21 (23.2s)	16 (54.4s)	15 (25.7s)	13 (57.5s)
160 × 40	14 (1m 5s)	13 (2m 24s)	15 (1m 9s)	13 (2m 38s)	15 (1m 47s)	14 (4m 1s)
320 × 80	12 (4m 7s)	9 (8m 1s)	12 (4m 7s)	13 (12m 3s)	15 (7m 41s)	16 (18m 34s)

SIMPLE for stratified flows with the same Froude number values. However, convergence difficulties were immediately found with the CLGS method for both obstacles at all values of F_h tried, however—even $F_h = 3.182$, for which the value of K is well below the transition value. This is in direct contrast to the results obtained in [18] with the SIMPLE method, and strongly suggests that greater robustness is derived from being able to smooth the velocity field independently from the density field, as is achieved with the decoupled method.

To restore smoothing to the CLGS method and obtain a convergent multigrid algorithm it therefore appeared necessary to increase the effective diffusivity coefficient in the density transport equation. One way of achieving this is to relax the order of the discretisation used for we have already seen that in the case of the driven cavity the results given by the Van Leer scheme are comparatively very nondissipative. An obvious choice would be to use the hybrid scheme instead, which, in view of the small diffusive coefficient in the continuous equation, would effectively mean first-order upwinding for the density throughout the flow. This had a dramatic effect, and convergent multigrid solutions could be now be achieved for all three Froude numbers mentioned earlier. In order to produce a valid comparison the same modification was made to the code implementing the decoupled method and the multigrid performance of both methods is summarised in Table 4 for $F_h = 1.592$ ($K = 1.0$), which were typical of the results obtained. The number of cycles required is roughly the same for both methods and for the vertical barrier again tends to rise with the Reynolds number. For all cases the SIMPLE method is taking approximately half the time of CLGS or less. The asymptotic convergence rates for the two obstacles at this Froude number and $Re = 100$ are compared in Fig. 13, which confirm that for the barrier SIMPLE still has the better asymptotic rate, as was the case in neutral flow. For the flow over the cosine hill, the advantage of CLGS in terms of convergence rate is now much less than for the corresponding neutral flow.

Although multigrid convergence can be achieved in this way, it is clearly at the expense of spatial accuracy, for the use of a low order scheme for one equation is bound to degrade the accuracy of the solution as a whole. The double discretisation technique described in Section 4 has been used successfully to generate a convergent multigrid solution technique for all the cases described above with the accuracy of the Van Leer discretisation restored to the density equation. This includes the case of the cosine hill, $Re = 100$, $F_h = 1.592$, for which no converged multigrid solution was obtained in [18] and all the cases with the coupled solution method for which no converged solutions were obtained at all. The details of the cycles are as before with the exception that at least two postsmoothing iterations on the finest grid were found necessary for convergence to be achieved. Table 5 gives the

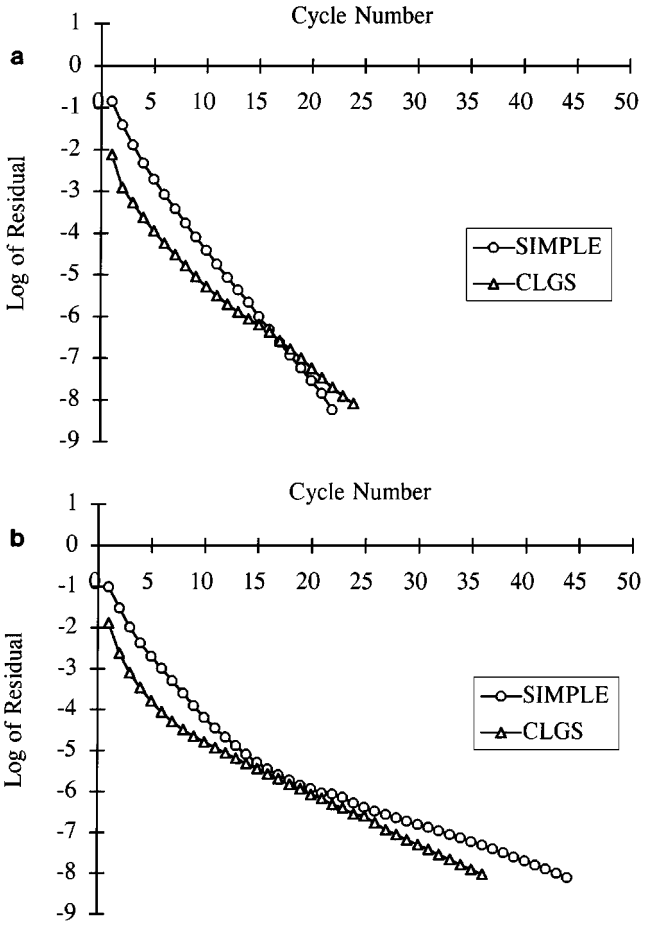

FIG. 13. Convergence histories for the computation of stratified flow over obstacles with $Re = 100$, $F_h = 1.592$ ($K = 1.0$) and the hybrid scheme used for the density equation: (a) the vertical barrier; (b) the cosine obstacle.

corresponding performance characteristics for computations with $K = 1.0$ and Fig. 14 gives the convergence histories with and without double discretisation for both methods on the finest grid. As expected convergence is not as rapid as with first order upwinding (Table 4) and at least half as many additional cycles are required to achieve the same convergence

TABLE 5

Number of Multigrid Cycles and Computing Times Taken for Convergence for Stratified Flow Past Obstacles ($F_h = 1.592$, $K = 1.0$) with the Van Leer Scheme and Double Discretisation Used for the Density Equation

	Barrier, Re = 50		Barrier, Re = 100		Cosine, Re = 100	
	SIMPLE	CLGS	SIMPLE	CLGS	SIMPLE	CLGS
80×20	25 (47.4s)	25 (1m 39s)	35 (1m 11s)	33 (2m 6s)	29 (1m 6s)	33 (3m 22s)
160×40	22 (2m 55s)	22 (5m 50s)	32 (4m 25s)	30 (8m 25s)	25 (4m 6s)	30 (12m 11s)
320×80	21 (12m 11s)	15 (18m 38s)	23 (13m 54s)	20 (25m 31s)	20 (14m 56s)	29 (52m 49s)

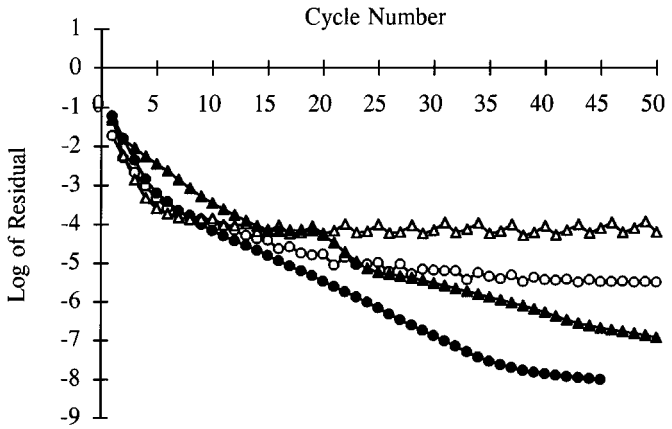

FIG. 14. Convergence histories for the computation of stratified flow over the cosine obstacle, Re = 100, $F_h = 1.592$ (K = 1.0) and the Van Leer scheme used for the density transport equation with and without double discretisation: ○, SIMPLE (without); ●, SIMPLE (with); △, CLGS (without); ▲, CLGS (with).

level. The additional postsmoothing iterations mean that computing times are between two and three times longer than those in Table 4. This procedure is somewhat ad hoc, however, and efficiencies could doubtless be made.

The differences in the flow field computed with and without the Van Leer correction for the density equation can be seen by comparing Fig. 12 with Fig. 15. Although there is little apparent difference for the flow over the vertical barrier at Re = 50, there is a significant

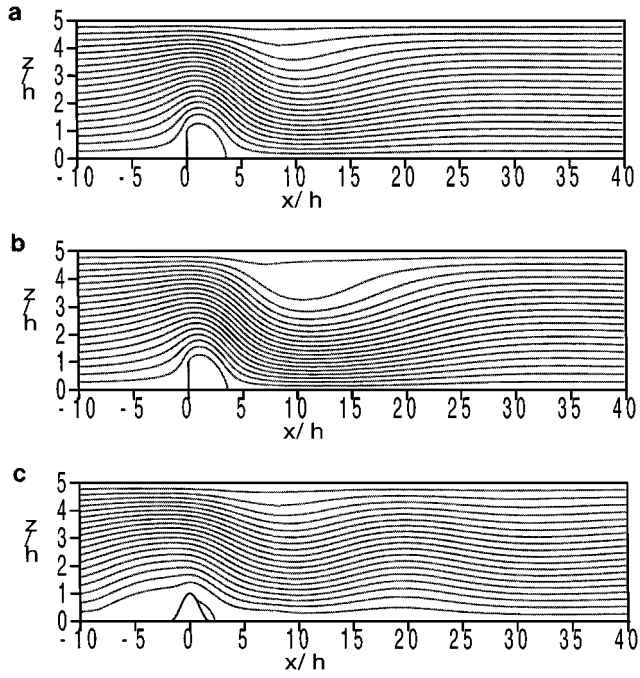

FIG. 15. Streamlines for stratified flow over obstacles with $F_h = 1.592$ (K = 1.0) and the Van Leer scheme and double discretisation used for the density equation: (a) the vertical barrier, Re = 50; (b) the vertical barrier, Re = 100; (c) the cosine obstacle, Re = 100.

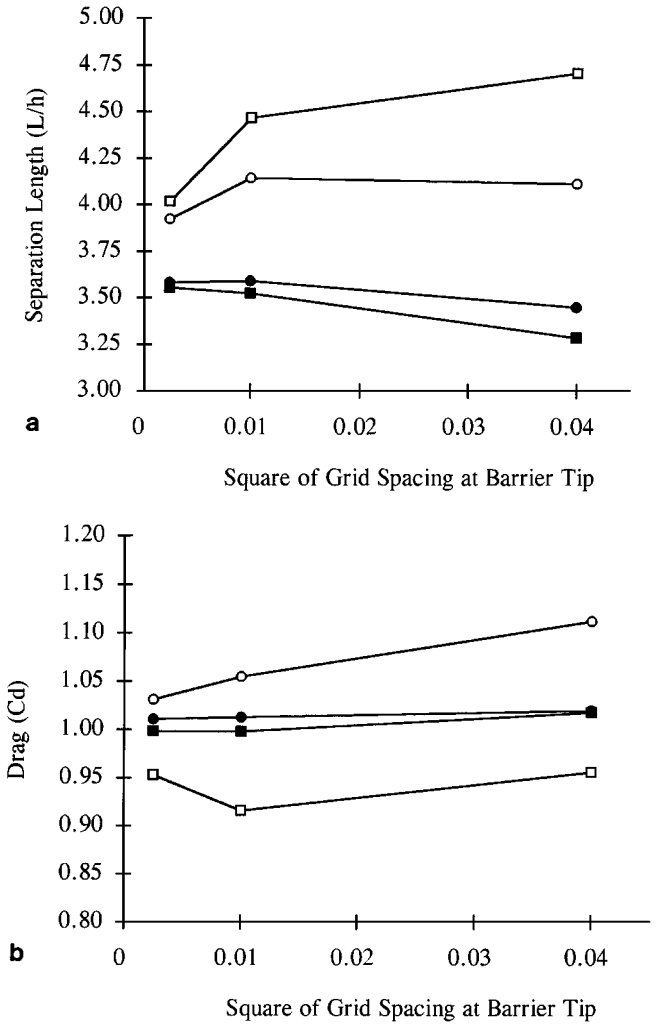

FIG. 16. Effects of grid refinement on the results of computations of flow over the vertical barrier with the two schemes for the density transport equation: (a) the length of the region of separated flow; (b) the drag. ○, Hybrid, Re = 50; ●, Van Leer, Re = 50; □, Hybrid, Re = 100; ■, Van Leer, Re = 100.

difference at Re = 100, with a rather more extensive region of slowed fluid given by the computation with the Van Leer scheme. This region is less extensive in both computations of the flow over the cosine obstacle, due to the influence of the surface boundary layer, but, whereas no lee wave is perceptible in the computation with the hybrid scheme, it is much more in evidence in the computation with the Van Leer scheme. That the global characteristics of the flow are also affected by the down-grading of the scheme for the density equation is illustrated for the flow over the vertical barrier in Fig. 16. Shown are the separation length, Fig. 16a, and the pressure drag, Fig. 16b, plotted against the square of the grid spacing at the barrier tip for the computations with and without the Van Leer correction at both Reynolds numbers. The results of the computations without the correction show typical first-order behaviour and undue grid dependence, while those with the correction show the second-order behaviour expected and much less dependence, especially between the two finest grids.

6. CONCLUSIONS

Multigrid computations with decoupled and coupled smoothing algorithms have been presented and compared for flow in a driven cavity and then neutral and stably stratified flow over two-dimensional obstacles. Convergence rates for the driven cavity problem were found to concur in general with those of other authors despite the less diffusive nature of the convective differencing scheme used, and for this problem the fastest convergence rates are those of the coupled methods. Implementation in lines increases the efficiency of the coupled method, however, and although convergence rates deteriorate somewhat, computing times per cycle and overall can be reduced. Vertical line sweeping was found to produce better convergence rates than horizontal line sweeping, although because of the manner of data storage, horizontal line sweeping requires less computing time per multigrid cycle. Changes to the data structure would mean that for the cavity problem vertical line sweeping would be the fastest method on the finest grid used at all Reynolds numbers tried.

For the flows over obstacles sweeping in horizontal lines was found to have very poor convergence characteristics and, although the scheme incorporating sweeps in alternating directions is a great improvement, the best by far of the coupled methods tried utilises sweeping across the flow direction in vertical lines. The reasons underlying the superiority of vertical line sweeping for both these kinds of flows are not clear at present. Theoretical analysis for the single convection-diffusion equation [31] indicates that symmetric (forward followed by backward) line sweeping should be robust when implemented in horizontal lines as well as vertical, but the very poor performance of horizontal line smoothing for the obstacle flows suggests that flow direction and/or grid nonuniformity are decisive factors. The relatively poor performance for the recirculating flows of the driven cavity is somewhat puzzling, however, but at least it confirms other numerical evidence for the same problem [23].

The rapid convergence rates of some of the coupled methods for the driven cavity problem were not reproduced for the flows over obstacles, however, even at the low Reynolds numbers used, and similar convergence rates are achieved for both coupled and decoupled methods. For flow over a vertical barrier, the asymptotic convergence rate of SIMPLE is actually better than that of the coupled method in both neutral and stably stratified flow. The shorter computing time per cycle of the SIMPLE algorithm means that convergence to a prescribed level of accuracy can be achieved in approximately half the computing time required for the coupled method. Even with a change of data structure the decoupled method would be faster.

For the stratified flows the coupled method was further found to be deficient in that convergence could not be achieved for any value of the Froude number when a high order differencing scheme was used for the density transport equation. This contrasts with the results of computations with the decoupled algorithm reported previously [18], where converged solutions were obtained in all but one case (although asymptotic convergence rates were not investigated). The reason for this is thought to relate to the small value of the diffusion coefficient in the density transport equation, where consequent lack of discrete ellipticity may lead to marginal smoothing. Decoupling the density field from the velocity field ensures that each can be smoothed independently of the other and to a sufficient degree. When the velocity and density fields are treated together as a locally coupled set, however, the results here suggest that difficulties in smoothing one variable will prevent effective smoothing for all. Although further work is necessary to fully confirm this, two

remedies to overcome the problem have been tried. Reducing the order of the discretisation to first-order upwinding allowed the computations to converge but at the expense of spatial accuracy. Accuracy and convergence can be achieved together by the use of double discretisation techniques, although computing times are now somewhat longer.

The multigrid scheme described here was used in [18] for stably stratified flows over two-dimensional obstacles which may be steady or unsteady, laminar or turbulent, and the extensions to deal with flows around three-dimensional obstacles are anticipated. The evidence given here suggests that, on grounds of efficiency and robustness, decoupled methods such as the SIMPLE algorithm are likely to continue to be the most appropriate choice of smoother for these kinds of flows. There may well be formulations of a coupled method which can outperform decoupled methods, but those implemented here do not appear to be among them.

ACKNOWLEDGMENTS

The second author wishes to acknowledge the financial support of the Natural Environment Research Council of the United Kingdom (Studentship GT4/95/291/MAS) while this work was undertaken.

REFERENCES

1. A. Brandt, Multi-level adaptive solutions to boundary-value problems, *Math. Comput.* **31**, 333 (1977).
2. A. Brandt and N. Dinar, Multigrid methods for flow problems, in *Numerical Methods for Partial Differential Equations* (Academic Press, New York, 1979).
3. S. V. Patankar, A calculation procedure for two-dimensional elliptic situations, *Numer. Heat Transfer* **4**, 409 (1981).
4. S. P. Vanka, Block implicit calculation of steady turbulent recirculating flows, *Int. J. Heat Mass Transfer* **28**, 2093 (1985).
5. S. P. Vanka, Block implicit multigrid solution of Navier–Stokes equations in primitive variables, *J. Comput. Phys.* **65**, 138 (1986).
6. P. Gaskell and N. Wright, Multigrids applied to an efficient fully coupled solution technique for recirculating fluid flow problems, in *Simulation and Optimisation of Large Systems* (IMA, Southend-on-Sea, 1988).
7. S. P. Vanka, Block implicit multigrid calculation of three-dimensional recirculating flows, *Comput. Meths. Appl. Mech. Engrg.* **58**, 29 (1986).
8. S. P. Vanka, A calculation procedure for three-dimensional steady recirculating flows using multigrid methods, *Comput. Meths. Appl. Mech. Engrg.* **55**, 321 (1986).
9. G. J. Shaw and S. Sivaloganathan, On the smoothing properties of the SIMPLE pressure correction algorithm, *Inter. J. Num. Methods for Fluids* **8**, 441 (1988).
10. S. Sivaloganathan and G. J. Shaw, A multigrid method for recirculating flows, *Inter. J. Num. Methods for Fluids* **8**, 417 (1988).
11. R. Webster, An algebraic multigrid solver for Navier–Stokes problems, *Inter. J. Num. Methods for Fluids* **18**, 761 (1989).
12. M. C. Thompson and J. H. Ferziger, An adaptive multigrid technique for the incompressible Navier–Stokes equations, *J. Comput. Phys.* **82**, 94 (1989).
13. S. Zeng and P. Wesseling, Multigrid solution of the incompressible Navier–Stokes in generalised coordinates, *SIAM J. Num. Anal.* **31**, 1764 (1994).
14. A. T. Degani and G. C. Fox, Parallel multigrid computation of the unsteady incompressible Navier–Stokes equations, *J. Comput. Phys.* **128**, 223 (1996).
15. F. S. Lien and M. A. Leschziner, Multigrid convergence acceleration for complex flow including turbulence, *Multigrid methods III* (Birkhauser, Boston, 1991), p. 277.

16. F. S. Lien and M. A. Leschziner, Multigrid acceleration for turbulent flow with a non-orthogonal collocated scheme, *Comp. Meth. Appl. Mech Engrg.* **118**, 351 (1994).
17. R. Jyotsna and S. P. Vanka, Pressure-based multigrid procedure for the Navier–Stokes equations on unstructured grids, in *Seventh Copper Mountain Conference on Multigrid Methods*, edited by N. D. Melson, T. A. Monteuffel, S. F. McCormick (NASA Conference Publication 3224, 1995), p. 691.
18. M. F. Paisley, Multigrid computation of stratified flow over two-dimensional obstacles, *J. Comput. Phys.* **136**, 411 (1997).
19. C. P. Thompson, G. K. Leaf, and S. P. Vanka, Application of a multigrid method to a bouyancy-induced flow problem, *Multigrid Methods*, edited by S. F. McCormick (Marcel Dekker, New York, 1988), p. 605. (*Lecture Notes in Pure and Applied Mathematics* **110**)
20. G. R. Stuhne and W. R. Peltier, Vortex erosion and amalgamation in a new model of large scale flow on the sphere, *J. Comput. Phys.* **128**, 58 (1996).
21. S. Sivaloganathan, G. J. Shaw, T. M. Shah, and D. F. Mayers, A comparison of multigrid methods for the incompressible Navier–Stokes equations, *Numerical Methods for Fluid Dynamics*, edited by K. W. Morton and M. J. Baines (Oxford University Press, 1988), p. 401.
22. C. H. Arakawa, A. O. Demuren, W. Rodi, and B. Schonung, Application of multigrid methods for the coupled and decoupled solution of the incompressible Navier Stokes equations, in *Notes on Numerical Fluid Dynamics 20* (Vieweg, 1988), p. 1.
23. S. Zeng and P. Wesseling, Numerical study of a multigrid method with four smoothing methods for the incompressible Navier–Stokes equations in general coordinates, in *Sixth Copper Mountain Conference on Multigrid Methods*, edited by N. D. Melson, T. A. Monteuffel, S. F. McCormick, C. C. Douglas (NASA Conference Publication 3339, 1996), p. 409.
24. A. Brandt, *Guide to Multigrid Development*, *Lecture Notes in Mathematics*, Vol. 960 (Springer–Verlag, New York/Berlin, 1981).
25. J. W. Rottman, D. Broutman, and R. Grimshaw, Numerical simulations of uniformly-stratified fluid flow over topography, *J. Fluid Mech.* **306**, 1 (1996).
26. I. P. Castro, W. H. Snyder, and P. G. Baines, Obstacle drag in stratified flow, *Proc. R. Soc. Lond. A* **429**, 119 (1990).
27. B. Van Leer, Towards the ultimate conservative difference scheme II. Monotonicity and conservation combined in a second order scheme, *J. Comput. Phys.* **14**, 361 (1974).
28. B. P. Leonard and S. Mokhtari, Beyond first-order upinding: The ultra-sharp alternative for non-oscillatory steady state simulation of convection, *Int. J. Numer. Methods Eng.* **30**, 729 (1990).
29. T. M. Shah, D. F. Mayers, and J. S. Rollett, Analysis and application of a line solver for recirculating flows using multigrid methods, in *Numerical Treatment of the Navier–Stokes Equations*, edited by W. Hackbusch and R. Rannacher (Vieweg, Braunschweig, 1990), p. 34. [*Notes on Numerical Fluid Mechanics* **30**]
30. T. M. Shah, D. Phil. thesis, Oxford University, 1990.
31. P. Wesseling, *An Introduction to Multigrid Methods* (Wiley, New York, 1991).