INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS, VOL. 9, 127–150 (1989)

PERFORMANCE OF ITERATIVE METHODS FOR NEWTONIAN AND GENERALIZED NEWTONIAN FLOWS

G. F. CAREY, K. C. WANG AND W. D. JOUBERT
The University of Texas at Austin, Austin, Texas 78712, U.S.A.

SUMMARY

We consider the use of accelerated gradient-type iterative methods for solution of Newtonian and certain non-Newtonian (power-law and Bingham models) viscous flow problems. The formulations are based on penalty and mixed finite element methods, and such factors as the effect of the penalty parameter, asymmetry, continuation and preconditioning are examined.

KEY WORDS Iterative solution Viscous flow Generalized conjugate gradient

INTRODUCTION

Both primitive variable (mixed) and penalty finite element methods have been extensively applied to Navier–Stokes problems. In this study we consider the application of these methods to generalized Newtonian fluids (such as power-law fluids) and the use of iterative techniques for the solution of the associated Jacobian systems for Newton iteration. Most finite element codes employ sparse elimination solvers for system solution, and there are several open issues related to the use of iterative methods; such as, how efficient are modern iterative methods for Navier–Stokes problems and for generalized Newtonian fluids? What complications arise for mixed or penalty methods, and how is the convergence influenced by choice of penalty parameter? For Navier–Stokes problems the inertial term $\mathbf{u}\cdot\nabla\mathbf{u}$ for velocity $\mathbf{u}$ introduces a non-symmetric contribution in the discrete problem. This in turn will influence the choice of iterative method and performance of the method. In addition, this non-linear term will become more significant as the flow velocity increases (higher-Reynolds-number flows). Convergence of the outer Newton non-linear iteration is sensitive to the choice of starting solution iterate—the initial iterate should lie in a 'ball of attraction' of the solution to the non-linear problem. Parameter and arc-length continuation techniques may determine a sequence of solutions on a continuation path, thereby providing improved starting iterates and better iterative performance. In the present study we conduct several exploratory investigations and numerical experiments concerning these issues.

PENALTY METHOD

We begin with the simpler case of linear Stokes flow and the penalty method. Stationary Stokes flow of an incompressible viscous fluid is governed by

$$-\nu\Delta\mathbf{u}+\nabla p=\mathbf{f} \quad \text{in } \Omega, \tag{1}$$

$$\nabla\cdot\mathbf{u}=0 \quad \text{in } \Omega, \tag{2}$$

0271–2091/89/020127–24$12.00
© 1989 by John Wiley & Sons, Ltd.
Received 20 August 1987
Revised 23 February 1988

where $\mathbf{u}$ is velocity, p is pressure, $\mathbf{f}$ is the body force, ν is the kinematic viscosity, and fluid density has been normalized as unity.

We consider essential boundary conditions

$$\mathbf{u} = \mathbf{g} \quad \text{on } \partial\Omega. \tag{3}$$

The penalized variational formulation of the Stokes problem may be obtained by introducing the perturbed Lagrangian

$$L(\mathbf{u}, p) = \int_\Omega \left(\frac{\nu}{2} \nabla\mathbf{u} : \nabla\mathbf{u} - p\nabla\cdot\mathbf{u} - \frac{\varepsilon}{2} p^2 - \mathbf{f}\cdot\mathbf{u} \right) dx, \tag{4}$$

where ε is a small parameter ($0 < \varepsilon \ll 1$) and $\nabla\mathbf{u} : \nabla\mathbf{u}$ denotes the dyadic product. At the stationary point $(\mathbf{u}, p)$ of L in (4) we have, on taking variations with respect to $\mathbf{u}$ and p and setting $\delta L = 0$ for arbitrary admissible $\mathbf{v} = \delta\mathbf{u}$, $q = \delta p$, respectively,

$$\int_\Omega (\nu \nabla\mathbf{u} : \nabla\mathbf{v} - p\nabla\cdot\mathbf{v} - \mathbf{f}\cdot\mathbf{v})\, dx = 0, \tag{5}$$

$$\int_\Omega (-\varepsilon p - \nabla\cdot\mathbf{u}) q\, dx = 0. \tag{6}$$

From (6), $p = -(1/\varepsilon)\nabla\cdot\mathbf{u}$ and substituting in (5) we obtain the penalty form: find $\mathbf{u}_\varepsilon \in V$ with $\mathbf{u}_\varepsilon = \mathbf{g}$ on $\partial\Omega$ such that

$$\int_\Omega \left(\nu \nabla\mathbf{u}_\varepsilon : \nabla\mathbf{v} + \frac{1}{\varepsilon} (\nabla\cdot\mathbf{u}_\varepsilon)(\nabla\cdot\mathbf{v}) \right) dx = \int_\Omega \mathbf{f}\cdot\mathbf{v}\, dx \tag{7}$$

for all $\mathbf{v} \in V$ with $\mathbf{v} = \mathbf{0}$ on $\partial\Omega$. For sufficiently regular data $\mathbf{f}$, the solution $\mathbf{u}_\varepsilon$ converges to $\mathbf{u}$ as $\varepsilon \to 0$.

The penalized statement (7) was the basis of several finite element studies from which it became apparent that the penalty term should be 'underintegrated' appropriately if the saddle-point properties required by existence/uniqueness theory are to go over to the discrete problem.[1-6] This discrete LBB or inf-sup condition then implies the approximate problem statement: find $\mathbf{u}_h \in V^h$ such that

$$\int_\Omega \nu \nabla\mathbf{u}_{h_\varepsilon} : \nabla\mathbf{v}_h\, dx + I\left[\frac{1}{\varepsilon} (\nabla\cdot\mathbf{u}_{h_\varepsilon})(\nabla\cdot\mathbf{v}_h) \right] = \int_\Omega \mathbf{f}\cdot\mathbf{v}_h\, dx \tag{8}$$

holds for all $\mathbf{v}_h \in V^h$, $\mathbf{v}_h = \mathbf{0}$ on $\partial\Omega$, where $I[\cdot]$ denotes reduced numerical integration. Here $\mathbf{u}_h$ denotes the usual finite element representation for the velocity components, and previous studies have identified appropriate choices of reduced integration rules for the penalty term. For a more detailed discussion, see, for instance, the treatment in Carey and Oden[7,8] and the references cited therein.

Introducing the finite element expansion for components of $\mathbf{u}_h$ and test basis functions for components of $\mathbf{v}_h$ in (8), we obtain a linear algebraic system of the form

$$\nu\mathbf{A}\mathbf{u}^* + \frac{1}{\varepsilon}\mathbf{B}\mathbf{u}^* = \mathbf{f}^*, \tag{9}$$

where $\mathbf{u}^*$ is the vector of $2N$ nodal velocity unknowns and $\mathbf{f}^*$ the right-side vector; $\mathbf{A}$ is $2N \times 2N$, symmetric positive definite and $\mathbf{B}$ is $2N \times 2N$, symmetric and positive but of reduced rank. Hence

the final system is

$$\mathbf{C}\mathbf{u}^* = \mathbf{f}^*, \tag{10}$$

where $\mathbf{C} = \nu\mathbf{A} + \varepsilon^{-1}\mathbf{B}$, with $\mathbf{C}$ symmetric positive definite (SPD).

PENALTY RESULTS

There are well established convergence results for SPD matrix systems.[9] However, numerical experiments performed in this study reveal that the conjugate gradient method and other iterative methods for solving (10) converge very slowly and in practice may fail. Closer examination of the intermediate calculations in the iterative algorithms reveals that the system (10) appears to be almost singular; that is, (10) is very ill conditioned. This result can be interpreted directly by noting that $\mathbf{B}$ is of reduced rank (singular) and ε^{-1} is large for $0 < \varepsilon \ll 1$. Thus $\varepsilon^{-1}\mathbf{B}$ dominates $\nu\mathbf{A}$ for computations in finite precision arithmetic. We found that diagonal (Jacobi) preconditioning and other standard preconditioning strategies, including even extensive use of incomplete factorization, all failed to resolve the problem. Some representative iterative performance studies are summarized in Table I for a generalized conjugate gradient method with various preconditionings. The results are discouraging and demonstrate that standard iterative methods and preconditioning accelerators do not work with the penalized Stokes problem (9).

The source of the difficulty is evidently the penalty parameter which implicitly enforces the incompressibility constraint ($\nabla\cdot\mathbf{u} \to 0$ as $\varepsilon \to 0$). We examined the sensitivity of iterative solution to choice of ε and found that the methods performed poorly at 'reasonable' ε-values. Dividing by ν in (9), we get $\mathbf{A} + (\nu\varepsilon)^{-1}\mathbf{B}$ so that the relative size of the viscosity will, of course, have some bearing on the conditioning and these results, but this does not change the trend for practical flow calculations. Continuation solution with incremental adjustment of ν and/or ε did little to alleviate the problem.

NON-LINEAR NON-SYMMETRIC PROBLEM

The preceeding difficulty clearly persists for the non-linear Navier–Stokes problem

$$-\nu\Delta\mathbf{u} + \mathbf{u}\cdot\nabla\mathbf{u} + \nabla p = \mathbf{f} \quad \text{in } \Omega, \tag{11}$$

$$\nabla\cdot\mathbf{u} = 0 \quad \text{in } \Omega, \tag{12}$$

Table I. Performance of ORTHOMIN[10] with various preconditioners; Stokes flow, 5 × 5 mesh of biquadratic elements for driven cavity flow

Preconditioner	Number of iterations	CPU time
Line Jacobi	Fails to converge	—
Incomplete Cholesky	Fails to converge	—
Modified incomplete Cholesky	30	0·7062
Blocked incomplete Cholesky (v.1)	120	0·6552
Blocked incomplete Cholesky (v.2)	143	0·4587
Modified blocked incomplete Cholesky (v.1)	114	0·7218
Modified blocked incomplete Cholesky (v.2)	175	0·5850

which leads to the discrete penalty variational problem: find $\mathbf{u}_{h_\varepsilon} \in V^h$ such that

$$\int_\Omega (\nu \nabla \mathbf{u}_{h_\varepsilon} : \nabla \mathbf{v}_h + \mathbf{u}_{h_\varepsilon} \cdot \nabla \mathbf{u}_{h_\varepsilon} \cdot \mathbf{v}_h)\,dx + I\left[\frac{1}{\varepsilon}(\nabla \cdot \mathbf{u}_{h_\varepsilon})(\nabla \cdot \mathbf{v}_h)\right] = \int_\Omega \mathbf{f} \cdot \mathbf{v}_h\,dx \tag{13}$$

for all $\mathbf{v}_h \in V^h$, $\mathbf{v}_h = \mathbf{0}$ on $\partial\Omega$. This yields the non-linear system

$$\nu \mathbf{A}\mathbf{u}^* + \frac{1}{\varepsilon}\mathbf{B}\mathbf{u}^* + \mathbf{g}(\mathbf{u}^*) = \mathbf{f}^* \tag{14}$$

and Newton iteration: given $\mathbf{u}_0$, for $n = 0, 1, 2, \ldots$ solve

$$\mathbf{J}_n(\mathbf{u}_{n+1} - \mathbf{u}_n) = -\mathbf{F}_n, \tag{15}$$

where

$$\mathbf{J}_n = \nu \mathbf{A} + \varepsilon^{-1}\mathbf{B} + (\partial \mathbf{g}/\partial \mathbf{u})|_{\mathbf{u}_n} \tag{16}$$

and

$$\mathbf{F}_n = (\nu \mathbf{A} + \varepsilon^{-1}\mathbf{B})\mathbf{u}_n + \mathbf{g}(\mathbf{u}_n) - \mathbf{f}^*. \tag{17}$$

The matrix contribution $(\partial \mathbf{g}/\partial \mathbf{u})$ at $\mathbf{u}_n$ is non-symmetric. Moreover, as noted in the Introduction, the iteration (15) fails to converge if initial iterate $\mathbf{u}_0$ is not sufficiently close to the solution $\mathbf{u}^*$.

To circumvent the latter problem, continuation techniques may be introduced to generate a sequence of solutions on a continuation path to $\mathbf{u}^*$. Parameter ν enters naturally in the formulation, and ν^{-1} is proportional to the Reynolds number. The simplest continuation technique is then an incremental continuation in ν^{-1}. Beginning with Stokes flow (the linear problem considered earlier), we solve for $\mathbf{u}_0$ and then use this as a starting iterate to solve for $\mathbf{u}_1$ at a small increment in Reynolds number, and so on.

As noted previously, the penalized Stokes system is ill conditioned for ε small and clearly this applies equally for the Jacobian systems in (15). Moreover, attempts to use both simple preconditioners and also incomplete factorization did not significantly improve the situation. Since this is attributable to the penalty term, which is present for all Jacobian systems (15), we next examine the effect of 'full' factorization of the Stokes problem to precondition iterative solution of the subsequent Jacobian systems in the non-linear iteration. That is, we have the algorithm:

(1) Factor (10) to $\mathbf{C}\mathbf{u}^* = \mathbf{L}\mathbf{D}\mathbf{L}^{\mathrm{T}}\mathbf{u}^* = \mathbf{f}^*$.

(2) Use the factored Stokes system to precondition (15):

$$\mathbf{Q}^{-1}\mathbf{J}(\mathbf{u}_{n+1} - \mathbf{u}_n) = \mathbf{Q}^{-1}\mathbf{F}^*$$

(with $\mathbf{Q} = \mathbf{C}$) and apply a generalized conjugate gradient scheme.

This scheme should be very efficient for low-Reynolds-number flow since $\mathbf{C} \sim \mathbf{J}$ in this case. Our first objective then is to ascertain the influence of full factorization of $\mathbf{C}$ on alleviating the ill conditioning due to the penalty in the non-linear problem. The second aspect to be considered is the treatment of the asymmetry arising from the term $\partial \mathbf{g}/\partial \mathbf{u}$. As the Reynolds number increases, this asymmetry will become more pronounced.

Biconjugate gradient[11, 12] has been applied successfully to convection-dominated transport processes and similar non-symmetric problems. Several other accelerated iterative techniques have been developed more recently for treatment of non-symmetric problems. Some of the predominant methods are ORTHOMIN, ORTHORES, ORTHODIR and the generalized minimal residual methods.[13] In this study we mainly use ORTHOMIN as our accelerated

iterative scheme with the Stokes preconditioning strategy specified above. ORTHOMIN provides a monotone decreasing residual; furthermore, if the symmetric part of the preconditioned matrix is positive definite—a condition routinely satisfied for sufficiently small Reynolds number—the method is guaranteed to converge. The truncated algorithm, ORTHOMIN(s) is given by

$$\mathbf{p}^{(n)} = \mathbf{z}^{(n)} + \sum_{i=\max\{0,\,n-s\}}^{n-1} \alpha_{n,i}\mathbf{p}^{(i)}, \quad n \geqslant 0,$$

$$\alpha_{n,i} = -\frac{(\mathbf{Q}^{-1}\mathbf{A}\mathbf{p}^{(i)}, \mathbf{Q}^{-1}\mathbf{A}\mathbf{z}^{(n)})}{(\mathbf{Q}^{-1}\mathbf{A}\mathbf{p}^{(i)}, \mathbf{Q}^{-1}\mathbf{A}\mathbf{p}^{(i)})},$$

$$\mathbf{u}^{(n+1)} = \mathbf{u}^{(n)} + \lambda_n \mathbf{p}^{(n)}, \quad n \geqslant 1,$$

$$\lambda_n = \frac{(\mathbf{Q}^{-1}\mathbf{A}\mathbf{p}^{(n)}, \mathbf{z}^{(n)})}{(\mathbf{Q}^{-1}\mathbf{A}\mathbf{p}^{(n)}, \mathbf{Q}^{-1}\mathbf{A}\mathbf{p}^{(n)})} = \frac{(\mathbf{Q}^{-1}\mathbf{A}\mathbf{z}^{(n)}, \mathbf{z}^{(n)})}{(\mathbf{Q}^{-1}\mathbf{A}\mathbf{p}^{(n)}, \mathbf{Q}^{-1}\mathbf{A}\mathbf{p}^{(n)})}. \tag{18}$$

Here $\mathbf{u}^{(0)}$ is given and pseudoresidual $\mathbf{z}^{(n)} = \mathbf{Q}^{-1}(\mathbf{b} - \mathbf{A}\mathbf{u}^{(n)})$; the variable $s \geqslant 0$ is a fixed integer specifying the number of previous direction vectors to be used.

If the preconditioned system is sufficiently far from symmetric positive definiteness, the truncated ORTHOMIN scheme is not guaranteed to converge. The scheme may stagnate (i.e. no improvement in the solution as n increases). However, periodic 'restarting' as stagnation is detected usually enables the method to proceed onward to convergence. In this case the residual is monitored and when stagnation occurs the old direction vectors are rejected and the process reinitiated. We show some examples later.

The GMRES method[14] is another useful algorithm for non-symmetric problems. When run as a periodically restarted method, GMRES produces the same iterates as the restarted and non-truncated ORTHOMIN method (i.e. restarted ORTHOMIN (∞)); however, restarted GMRES is less susceptible to breakdown and stagnation. As a truncated method, the GMRES method applied to a symmetric system is equivalent to the MINRES method,[15] suitable for symmetric indefinite problems.

NAVIER–STOKES RESULTS

The unit-driven cavity problem was again employed for the numerical experiments. Plots of CPU time on the UT System CRAY X-MP/24 on meshes of 5×5, 7×7, 10×10 and 15×15 biquadratic (nine-node) elements are given in Figures 1–4 and Table II. Calculations were made with penalty parameter $\varepsilon = 10^{-6}$ and Reynolds numbers up to 500. A truncated five-term ORTHOMIN was used in this study (unless otherwise specified). The tolerance for the outer Newton iteration and inner gradient iteration was 10^{-4} on successive iterates (unless otherwise specified). The algorithm is compared here with frontal elimination solution for the same incremental continuation history.

From the figures we see that at small Reynolds numbers the iterative method is very efficient as expected. At the higher range of Re considered, asymmetry becomes more pronounced, the preconditioner becomes less effective and the iterative performance deteriorates. Furthermore, on the finer grids the 'cross-over point' between the methods moves towards the higher Reynolds numbers. The slight dip in the graph near $Re = 350$ is due to variation in the number of Newton iterations for convergence. Values are given in Table III for the CPU time as we reduced the number of ORTHOMIN iterations per Newton step for the 10×10 and 15×15 meshes at $Re = 400$. Even though the ORTHOMIN iteration is not convergent to the full tolerance at each Newton step, the final solution at the desired Reynolds number is fully converged to the specified

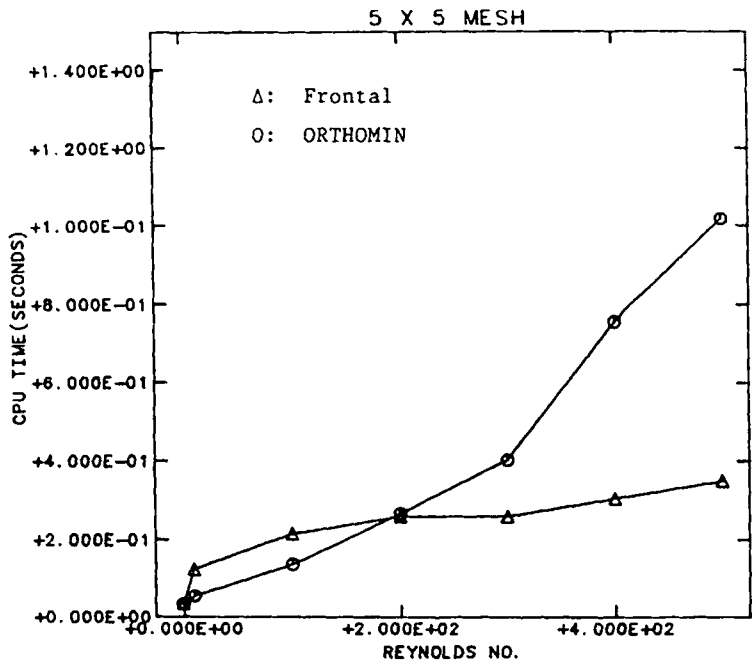

Figure 1. Comparison of CPU time for 5 × 5 mesh

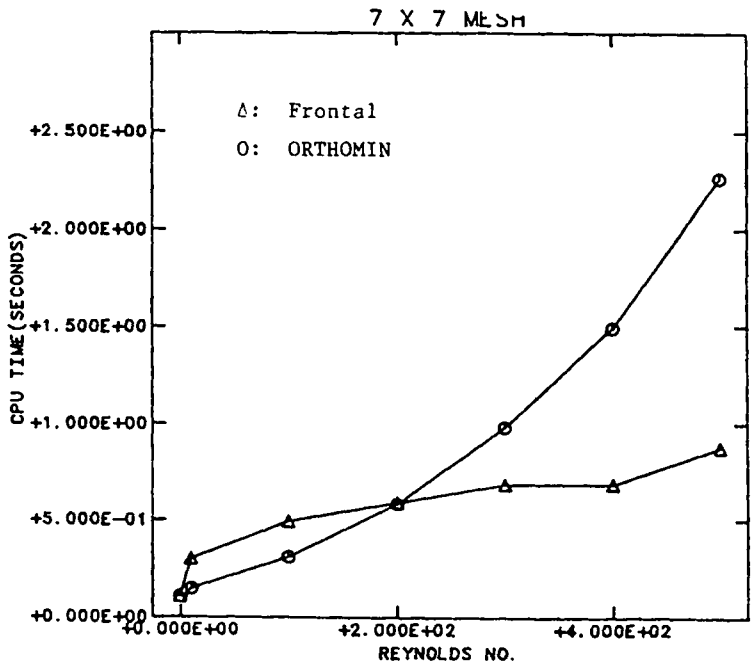

Figure 2. Comparison of CPU time for 7 × 7 mesh

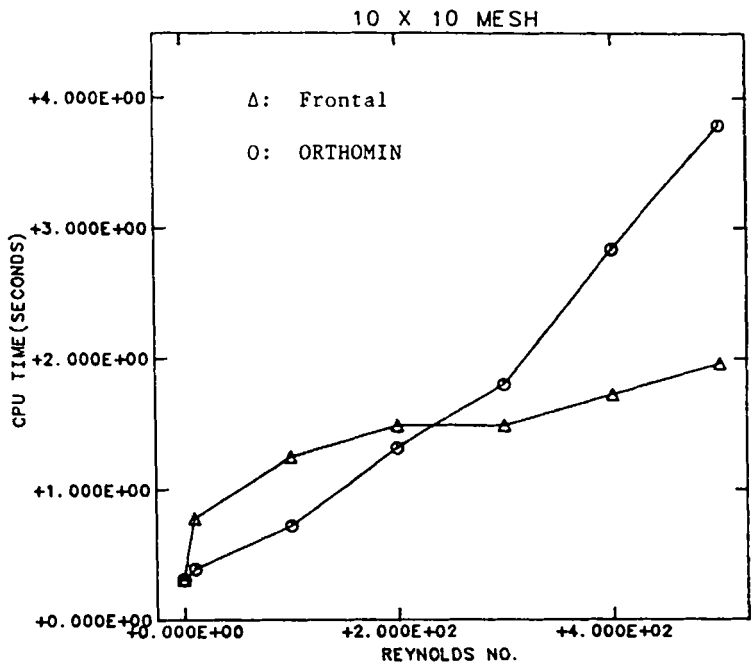

Figure 3. Comparison of CPU time for 10 × 10 mesh

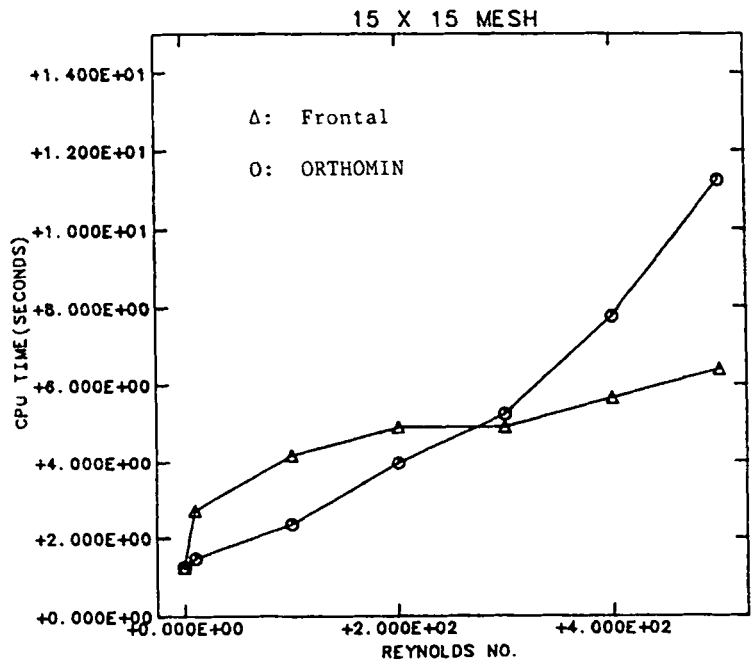

Figure 4. Comparison of CPU time for 15 × 15 mesh

Table II. CPU time for ORTHOMIN and frontal solver

Mesh size	*Re*	CPU time frontal solver	CPU time ORTHOMIN	Newton iterations/ORTHOMIN iterations†
5 × 5	0	0·03530	0·03530	1/1
	10	0·1250	0·05643	3/9 (1, 4, 4)
	100	0·2147	0·1366	5/48 (1, 12, 11, 12, 12)
	200	0·2596	0·2663	6/115 (1, 21, 21, 24, 23, 25)
	300	0·2596	0·4021	6/185 (1, 32, 37, 36, 39, 40)
	400	0·3044	0·7560	7/367 (1, 44, 96*, 47, 61, 57, 61)
	500	0·3493	1·0174	8/450 (1, 55, 97*, 64, 70, 70, 72, 71)
7 × 7	0	0·1068	0·1068	1/1
	10	0·3004	0·1486	3/9 (1, 4, 4)
	100	0·4941	0·3097	5/50 (1, 12, 12, 13)
	200	0·5909	0·5838	6/123 (1, 22, 21, 26, 26, 27)
	300	0·6877	0·9851	7/232 (1, 39, 37, 36, 37, 40, 42)
	400	0·6877	1·4926	7/367 (1, 75*, 63, 59, 56, 46, 57)
	500	0·8813	2·2597	9/571 (1, 63*, 120*, 56*, 67, 58, 74, 64, 68)
10 × 10	0	0·3077	0·3077	1/1
	10	0·7804	0·3903	3/9 (1, 4, 4)
	100	1·2531	0·7241	5/51 (1, 13, 12, 12, 13)
	200	1·4895	1·3204	6/132 (1, 25, 22, 27, 28, 29)
	300	1·4895	1·7993	6/199 (1, 43, 36, 38, 43, 38)
	400	1·7259	2·8420	7/337 (1, 43*, 61, 51, 62, 54, 55)
	500	1·9622	3·7920	8/467 (1, 53*, 43*, 67, 64, 91*, 76, 72)
15 × 15	0	1·2530	1·2530	1/1
	10	2·7200	1·4804	3/9 (1, 4, 4)
	100	4·1870	2·3630	5/52 (1, 13, 12, 12, 14)
	200	4·9206	3·9779	6/133 (1, 25, 22, 27, 28, 30)
	300	4·9206	5·2425	6/200 (1, 44, 37, 38, 41, 39)
	400	5·6541	7·7505	7/328 (1, 58*, 53, 49, 57, 56 54)
	500	6·3876	11·2390	8/506 (1, 53*, 48*, 74, 61*, 85, 104*, 80)

† Numbers in parentheses indicate number of ORTHOMIN iterations for each Newton step.
* Indicates ORTHOMIN iteration failed to converge to specified tolerance (residual stagnate without improvement) at indicated outer Newton step.

Table III. CPU time for ORTHOMIN iteration; $Re = 400$ and varying iteration limit ITMAX

Mesh size	ITMAX	CPU time	Newton iterations/ORTHOMIN iterations
10 × 10	100	2·8420	7/337 (1, 53, 61, 51, 62, 54, 55)
	50	2·5981	7/301 (1, 50, 50, 50, 50, 50, 50)
	25	1·6952	8/176 (1, 25, 25, 25, 25, 25, 25, 25)
	10	1·1026	10/91 (1, 10, 10, 10, 10, 10, 10, 10, 10, 10)
15 × 15	100	7·7505	7/337 (1, 53, 61, 51, 62, 54, 55)
	50	7·2684	7/301 (1, 50, 50, 50, 50, 50, 50)
	25	4·9248	8/176 (1, 25, 25, 25, 25, 25, 25, 25)
	10	3·3571	10/91 (1, 10, 10, 10, 10, 10, 10, 10, 10, 10)

tolerance. This approach of relaxing the ORTHOMIN iteration tolerance at intermediate Newton iterates reduces the CPU time by 50%. The midplane velocity profiles are given in Figure 5 and are indistinguishable from those obtained using full ORTHOMIN convergence as expected. For $Re = 1000$ continuation is necessary to obtain convergence of the Newton iteration. CPU times are given in Table IV for calculations with Reynolds number increment 200.

The deterioration of the preconditioner at higher Reynolds numbers is mainly due to the increasing asymmetry of the problem, since the penalty is accommodated through the Stokes preconditioner. To improve conditioning, a complete refactorization was made at each new Reynolds number with the (ORTHOMIN) iterative solution in the subsequent Newton solution steps. This approach improves the rate of convergence at high Re and compares favourably with the full elimination solution and no inner iteration. The midplane velocity profiles for $Re = 1000$ are compared in Figure 6. Figures 7–10 indicate CPU time versus Re for each of the meshes in the cases $Re = 300$ and 500 in the continuation scheme. Comparing the CPU times with those presented in earlier figures, it is clear that continuation enhances the iterative performance.

To assess the effect of penalty parameter at higher Re (since ν has changed), we made calculations with continuation in ε from $\varepsilon = 100$ (where penalty and viscous terms are of the same order) to $\varepsilon = 10^{-6}$. Results for this calculation are shown in Table V. Clearly this does not improve the behaviour. Augmented Lagrangian methods[16] also appeared to give little improvement.

Rather than employ ORTHOMIN, BCG or similar non-symmetric iterative solvers, we considered the possibility of symmetrizing the non-linear iteration scheme in a modified variant of successive approximation. Recall that in successive approximation the non-linear term is typically linearized as $\mathbf{u}_n \cdot \nabla \mathbf{u}_{n+1}$, where $\mathbf{u}_n$ is the known solution at the previous iterate. This convection term then yields a non-symmetric contribution. Alternatively, one can write for each component of $\mathbf{u} \cdot \nabla \mathbf{u}$ with $\mathbf{u} = (u, v)$

$$\begin{aligned} &u_{n+1}\left(\frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}\right)_n + v_{n+1}\left(\frac{\partial u}{\partial y}+\frac{\partial v}{\partial x}\right)_n - \left(u\frac{\partial v}{\partial y}+v\frac{\partial v}{\partial x}\right)_n, \\ &u_{n+1}\left(\frac{\partial u}{\partial y}+\frac{\partial v}{\partial x}\right)_n + v_{n+1}\left(\frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}\right)_n - \left(u\frac{\partial u}{\partial y}+v\frac{\partial u}{\partial x}\right)_n \end{aligned} \tag{19}$$

respectively. However, this symmetric formulation yields solutions in a reasonable number of iterations only for Reynolds numbers up to 60. Even with continuation, the rate of convergence was too slow for this formulation to be seriously considered.

MIXED METHOD

Next we considered the use of iterative methods in conjunction with a mixed formulation. The variational statement of the Navier–Stokes problem is: find $\mathbf{u} \in V$, $p \in P$ such that

$$\int_\Omega (\nu \nabla \mathbf{u} : \nabla \mathbf{v} + \mathbf{u} \cdot \nabla \mathbf{u} \cdot \mathbf{v} + p \nabla \cdot \mathbf{v})\, dx = \int_\Omega \mathbf{f} \cdot \mathbf{v}\, dx, \tag{20}$$

$$\int_\Omega q \nabla \cdot \mathbf{u}\, dx = 0 \tag{21}$$

for all admissible test functions $(\mathbf{v}, q)$. Introducing the finite element model (C^0 biquadratic velocities and C^0 bilinear pressures), we obtain the sparse non-linear system

$$\nu \mathbf{A}\mathbf{u} + \mathbf{g}(\mathbf{u}) - \mathbf{B}^{\mathrm{T}}\mathbf{p} = \mathbf{f}^*, \tag{22}$$

$$\mathbf{B}\mathbf{u} = \mathbf{0}. \tag{23}$$

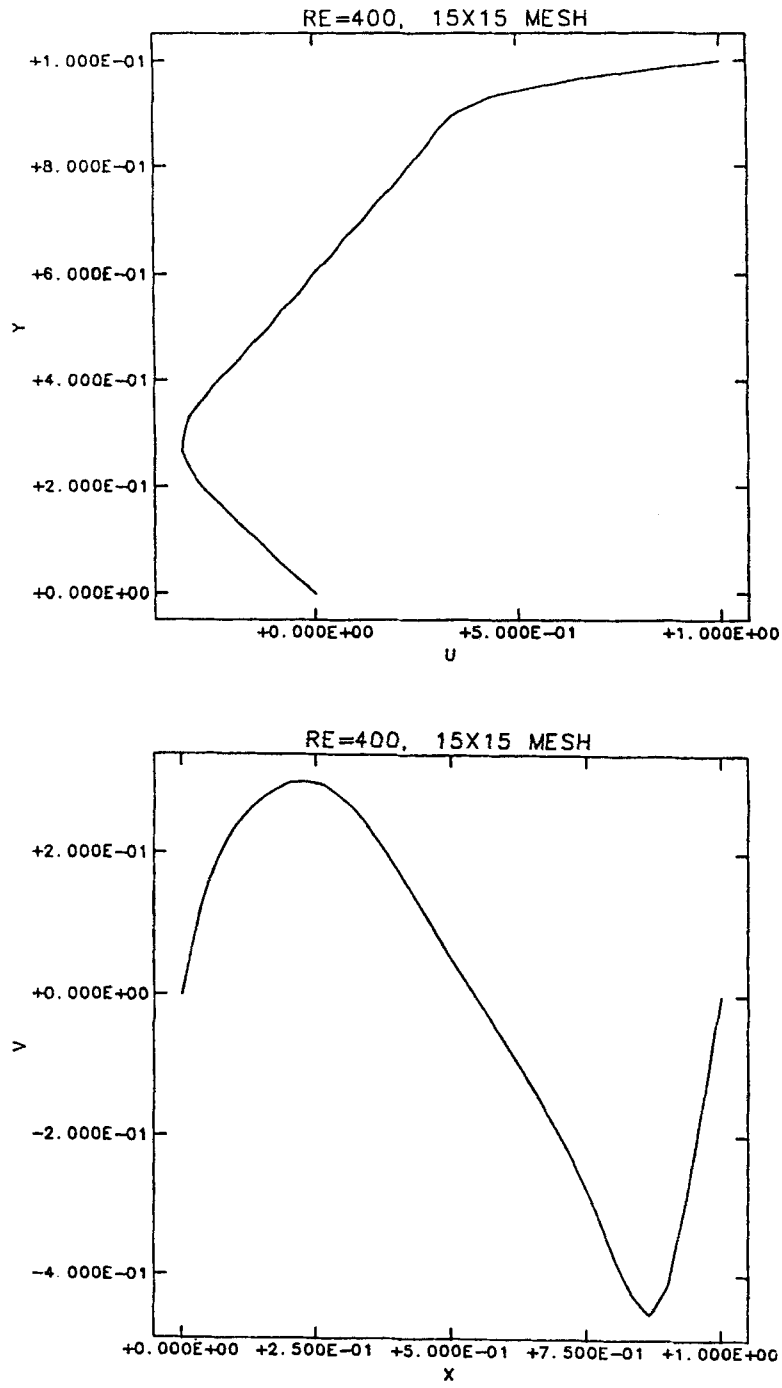

Figure 5. Midplane velocity profiles; $Re = 400$

Table IV. CPU time for $Re = 1000$ with incremental continuation in Re

Mesh size	CPU time/iteration count ORTHOMIN	CPU time/iteration count frontal solver
5 × 5	0·6548/115, 50, 23, 18, 16	1·0226/6, 5, 4, 4, 4
7 × 7	1·4851/123, 55, 23, 17, 16	2·2370/6, 5, 4, 4, 4
10 × 10	3·5622/132, 53, 24, 18, 16	5·5078/6, 5, 4, 4, 4
15 × 15	11·4806/133, 47, 24, 18, 16	17·3904/6, 5, 4, 4, 4

For stokes flow $\mathbf{g}(\mathbf{u}) = \mathbf{0}$ and, if we write (23) as $-\mathbf{Bu} = \mathbf{0}$, we obtain a sparse symmetric linear block system for $(\mathbf{u}, \mathbf{p})$. However, the system is not positive definite. In particular there is a zero diagonal block in (23). Hence the previous problem associated with the penalty formulation is not present, but the zero diagonal entries in (23) prevent conjugate gradient (CG) solution with Jacobi (diagonal) preconditioning. If we perturb the zero block by a small non-zero diagonal of order 10^{-10}, CG still fails. For systems that are not SPD, the generalized conjugate gradient methods such as ORTHOMIN are appropriate. A driven cavity Stokes flow was computed using truncated ORTHOMIN, ORTHORES and GMRES accelerators with Jacobi preconditioning. Performance results are presented in Table VI for the various methods with tolerances 10^{-6} and 10^{-8} on the inner iteration. There is no apparent difference between the velocity profile obtained by using either tolerance on the inner iteration, and results agree closely with the elimination solution.

In Figures 11–13 we plot the residual norm (an error measure for the iteration) against the number of iterations. The residual norm decreases monotonically for ORTHOMIN and GMRES but oscillates for ORTHORES. We also considered the non-symmetric form (20) directly. Again we used the non-symmetric accelerator with Jacobi preconditioning to solve the driven cavity Stokes flow. Performance results are presented in Table VII. ORTHOMIN yielded convergent solutions in some cases. In most cases the error norm was observed to 'level off' and failed to converge. Plots of residual norm against number of iterations for the five-term truncated formula are shown in Figures 14–16. Restarting the algorithm when this occurs appears to improve the behaviour as seen in Figure 17. In other research on non-symmetric problems, restarting has also been shown to be effective. Performance results for the driven cavity problem are presented in Table VIII. The iterative methods perform less well than the elimination solution in all cases, even for low-Reynolds-number flow. This was due mainly to the use of a non-symmetric diagonal storage method which is less suitable when the very sparse matrix $\mathbf{B}$ for the pressure unknowns must be accommodated.

NON-NEWTONIAN FLOWS

Part of the present research is directed towards the use of iterative solvers for 'generalized' Newtonian fluids, e.g. non-Newtonian flows of power-law type or Bingham plastic flow. The viscous stress for this type of fluid is given by $\tau_{ij} = 2\eta(I_2)D_{ij}$, where D_{ij} is the rate of deformation tensor, $I_2 = \frac{1}{2}\mathrm{tr}(D^2)$ is the second invariant of D_{ij} and η is the apparent viscosity. A variety of models have been proposed and correlated with experimental data.[17] In this study we consider power-law fluids so that η has the form $\eta = KI_2^{(n-1)/2}$, where K is the consistency factor and n is the power-law index. One of the significant features of power-law fluids is the shear-thinning effect when $n < 1$. This effect can best be illustrated by an example. The fully developed velocity profile

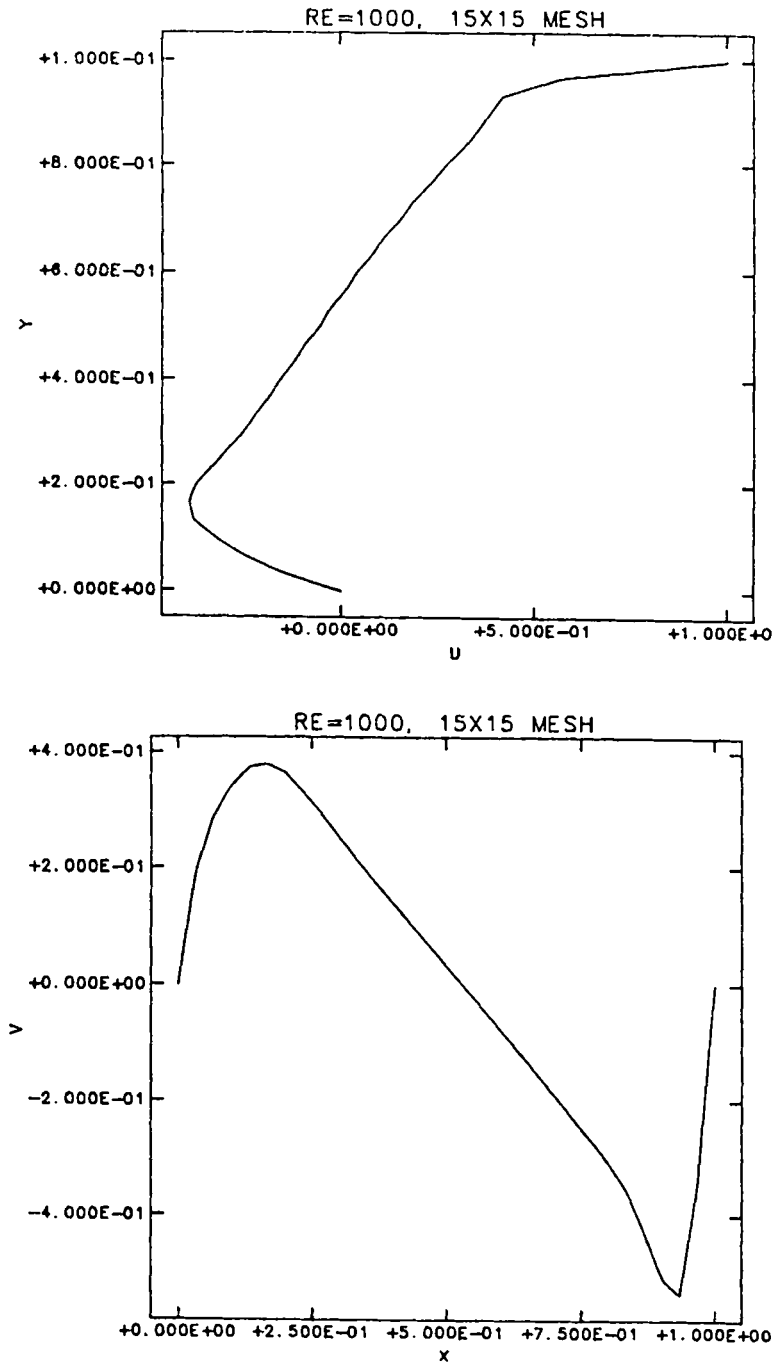

Figure 6. Midplane velocity profiles; $Re = 1000$

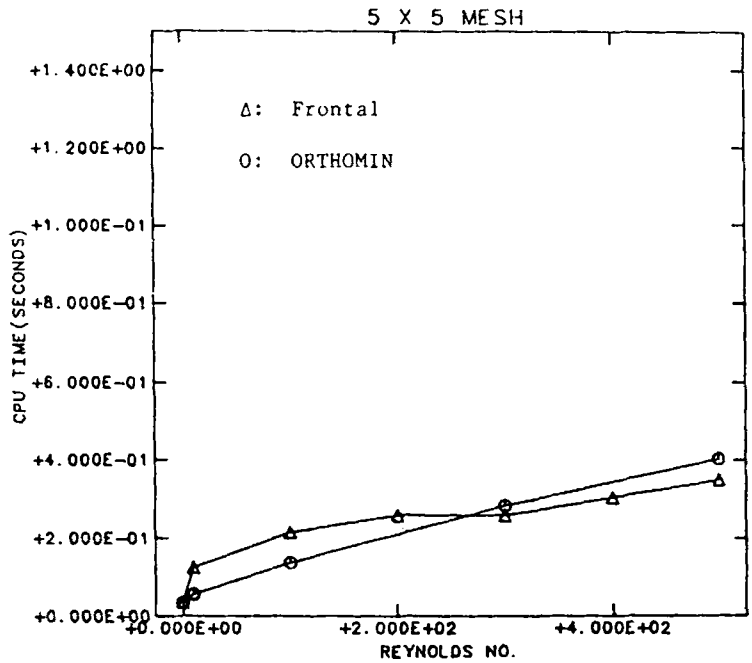

Figure 7. CPU time versus Reynolds number (incremental continuation in *Re*)

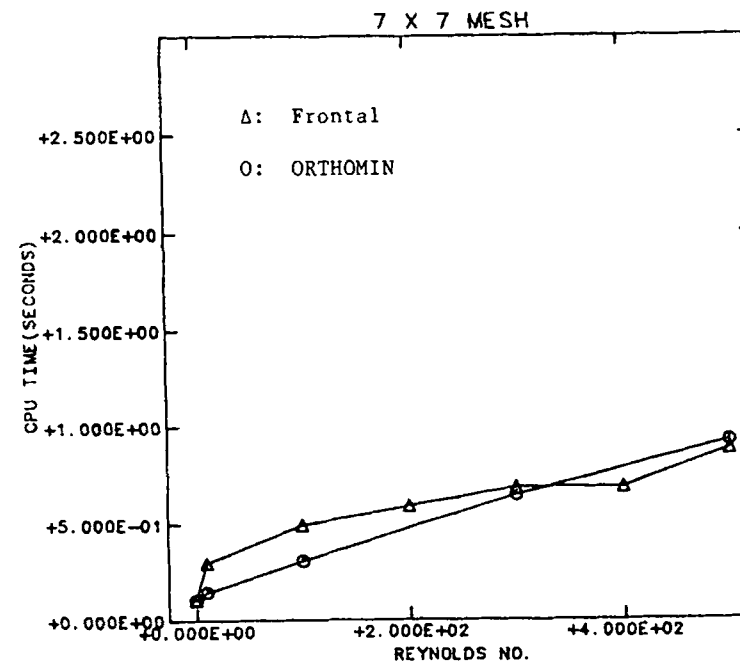

Figure 8. CPU time versus Reynolds number (incremental continuation in *Re*)

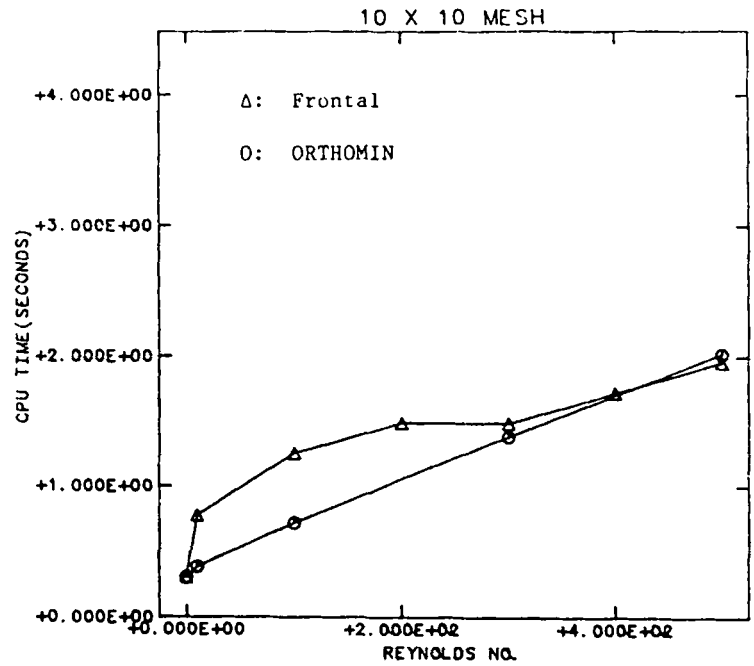

Figure 9. CPU time versus Reynolds number (incremental continuation in *Re*)

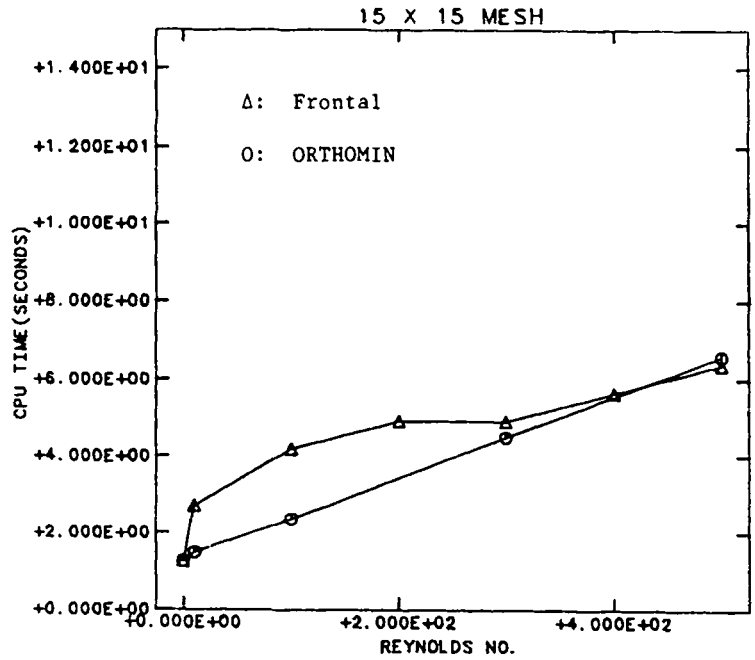

Figure 10. CPU time versus Reynolds number (incremental continuation in *Re*)

Table V. CPU time for $Re = 400$; continuation on penalty parameter from initial ε to $\varepsilon = 10^{-6}$

Mesh size	Initial ε	CPU time/iteration count
5 × 5	100	1·4957/218, 417, 18, 3, 1
	1	0·7158/264, 18, 3, 1
7 × 7	100	3·0785/289, 323, 12, 3, 1
	1	1·6713/306, 12, 3, 1
10 × 10	100	7·0542/321/378, 18, 3, 1
	1	3·6452/289, 18, 3, 1
15 × 15	100	21·0770/318, 401, 18, 3, 1
	1	14·8543/455, 18, 3, 1

Table VI. Driven cavity Stokes flow (mixed formulation); 5 × 5 mesh

Method	Tolerance	Number of iterations	CPU time
ORTHOMIN	10^{-6}	115	0·3237
	10^{-8}	190	0·5305
ORTHORES	10^{-6}	102	0·2876
	10^{-8}	154	0·4352
GMRES	10^{-6}	99	0·2922
	10^{-8}	154	0·4536

for channel flow of a power-law fluid at $n = 0{\cdot}8$, $0{\cdot}5$, $0{\cdot}3$ and $0{\cdot}2$ is shown in Figure 18. The fully developed profile becomes flat in the central zone as n decreases and the shearing near the wall becomes more pronounced. The apparent viscosity across a section of the channel is shown in Figure 19 for the same fluid cases. At $n = 0{\cdot}2$ the viscosity is small near the wall where the shear stress is large, and the viscosity is large near the centre of the channel where the shear stress is small. Hence there is a large viscosity gradient across the channel, and this behaviour for small n influences the behaviour of the non-linear solution iteration. As the power-law index n decreases from 1 towards 0, the non-Newtonian flow becomes increasingly difficult to model. For non-Newtonian flows such as Bingham fluids and power-law fluids with small power-law index, Newton's method fails; for example, Gartling[18] indicates the failure of Newton's method for cavity flow problems with $n \leq 0{\cdot}25$. Numerical experiments indicate that even incremental continuation in n using the solution at $n = 0{\cdot}4$ as a starting iterate for flow at $n = 0{\cdot}3$, and so on, appears to be of little help, since the Newton correction step is too large owing to the extreme change associated with shear thinning. By introducing a line-search procedure for under-relaxation of the Newton correction, the method can be improved and a solution obtained. Here we introduce a one-dimensional line search based on a residual minimization process to under-relax the Newton iteration.

Let $\mathbf{u}_n$ be the current solution iterate and $\delta\mathbf{u}$ the correction vector obtained from solution of the Newton Jacobian system. Then the relaxation iterate is

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \omega\delta\mathbf{u}_n,$$

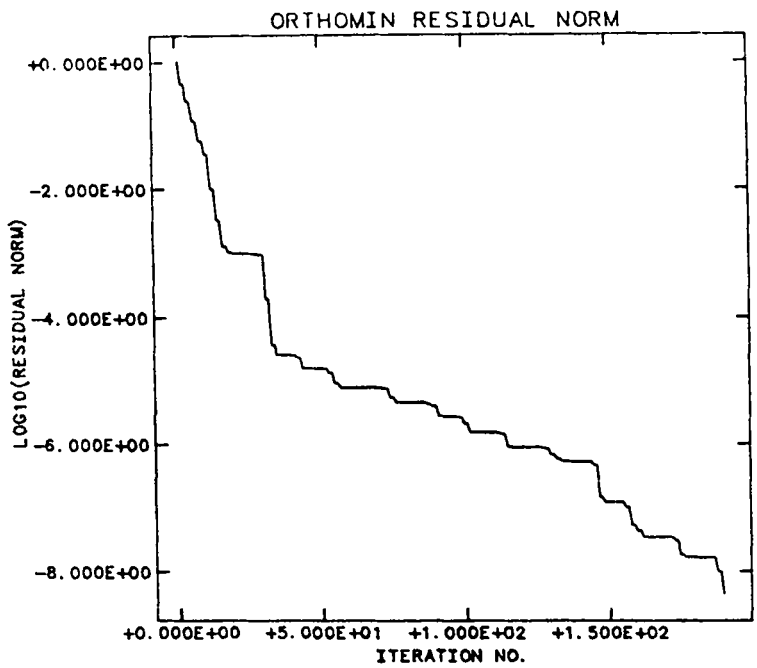

Figure 11. History of residual norm

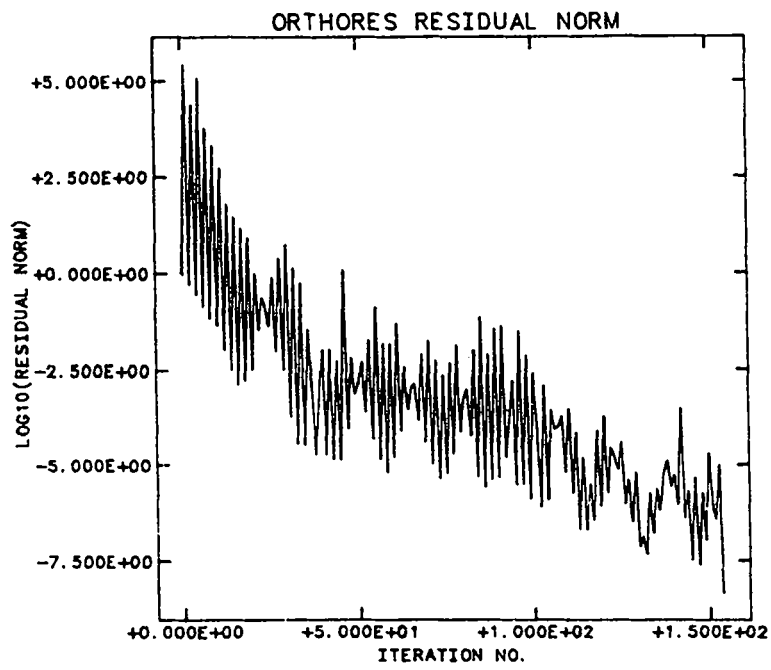

Figure 12. History of residual norm

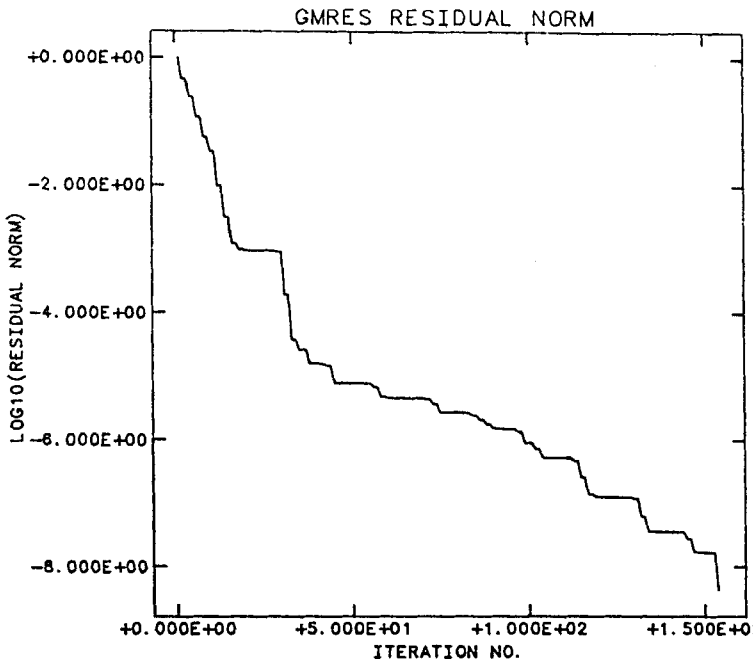

Figure 13. History of residual norm

Table VII. Driven cavity stokes flow; nonsymmetric mixed form, Jacobi preconditioning, 5 × 5 mesh

Accelerator	Number of terms	Number of iterations	CPU time
ORTHOMIN	2	> 400	NA
	3	320	0·8962
	4	> 400	NA
	5	364	1·0471
	7	294	0·8844
	8	> 400	NA
	9	> 400	NA
	10	> 400	NA
ORTHORES	2	> 400	NA
	5	> 400	NA
	6	> 400	NA
	7	> 400	NA
GMRES	2	> 400	NA
	4	> 400	NA
	5	> 400	NA
	6	> 400	NA
	7	> 400	NA
CG		Diverges	NA

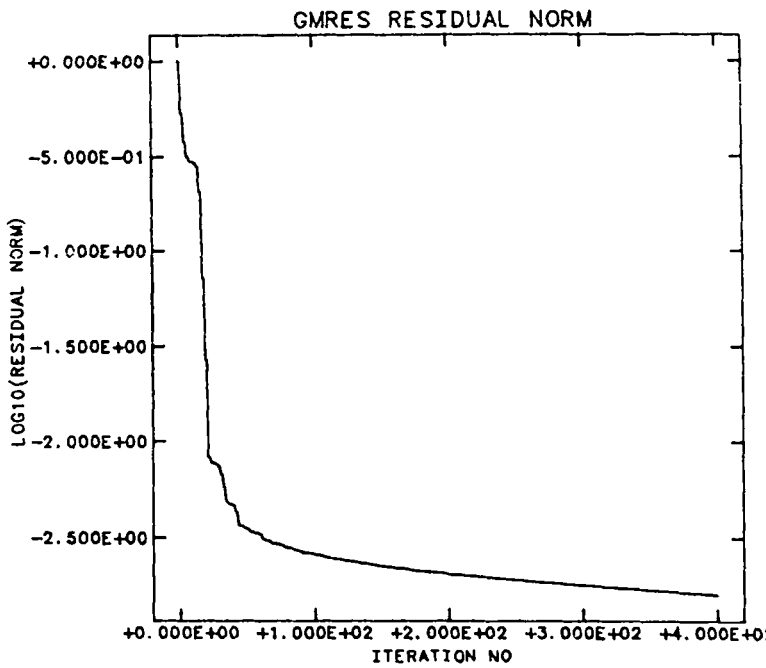

Figure 14. History of residual norm (non-symmetric case)

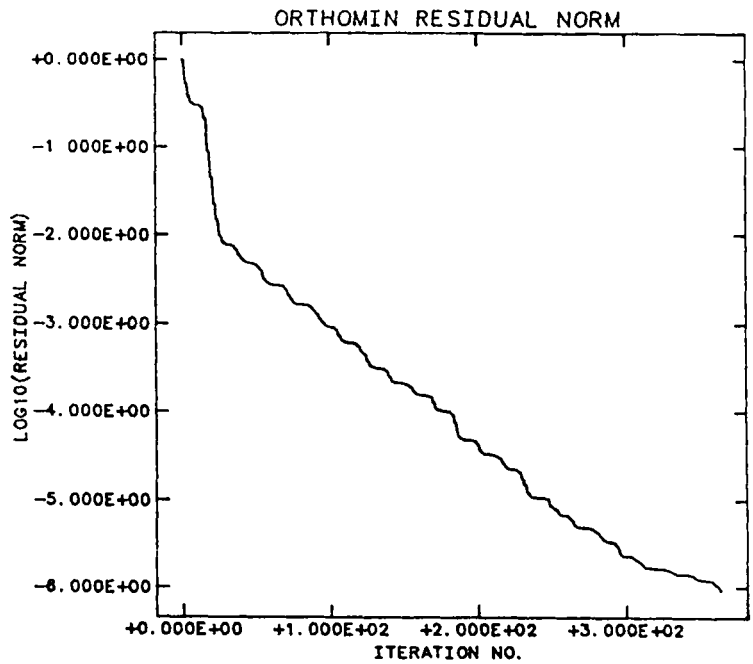

Figure 15. History of residual norm (non-symmetric case)

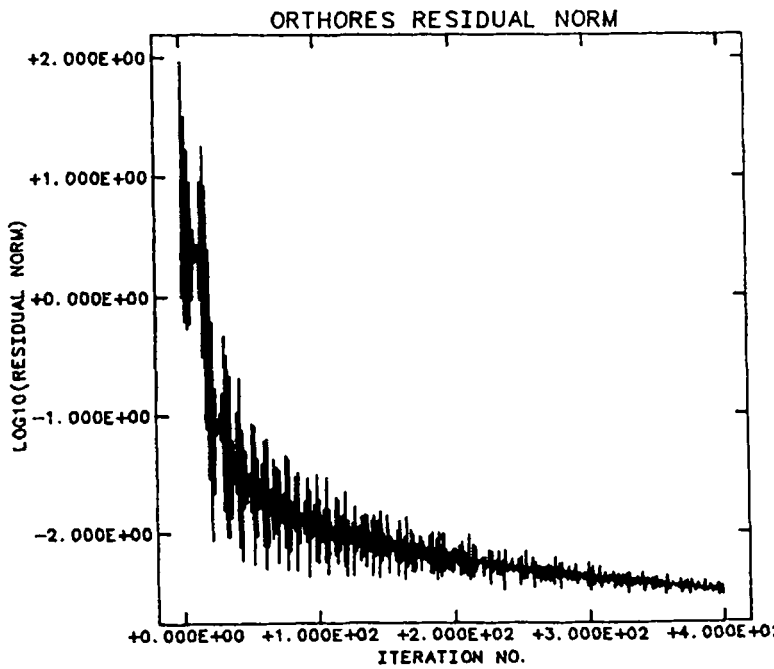

Figure 16. History of residual norm (non-symmetric case)

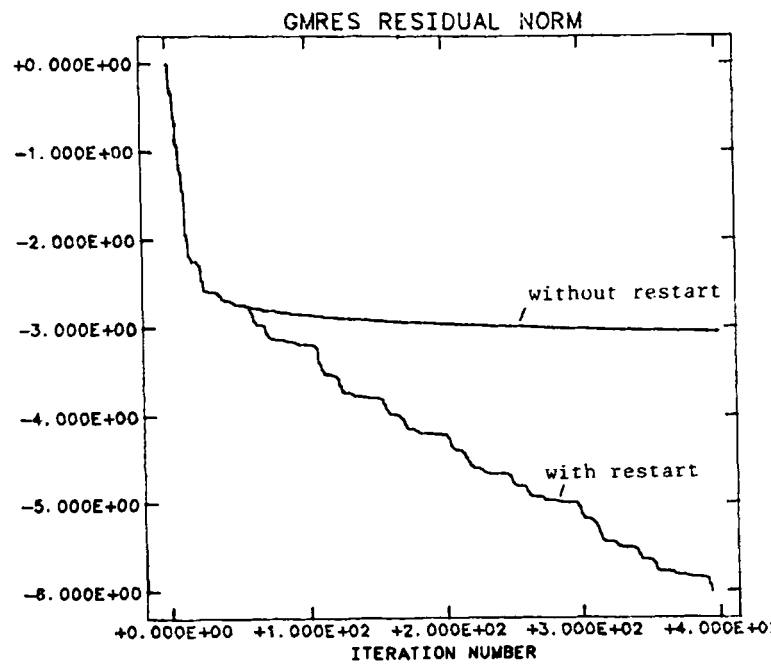

Figure 17. History of residual norm, with and without restart

Table VIII. Driven cavity flow (mixed formulation); 5 × 5 mesh

Re	CPU time frontal solver	CPU time ORTHOMIN	Newton iterations/ORTHOMIN iterations
10	0·1890	0·2500	3/9 (1, 4, 4)
100	0·3151	0·5147	5/60 (1, 15, 15, 15, 14)
200	0·3781	0·9705	6/150 (1, 28, 30, 31, 31, 29)
300	0·4411	1·6457	7/286 (1, 46, 49, 50, 48, 51, 41)
400	0·5042	3·1107	8/582 (1, 63, 80, 112, 96, 68, 83, 79)

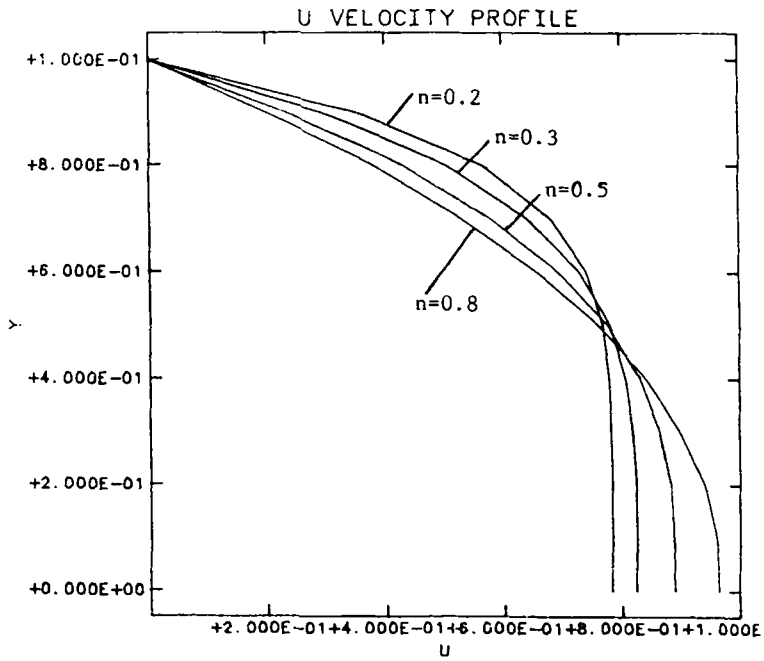

Figure 18. Velocity profile across the channel

where ω is the relaxation factor. The value of ω is computed by line search to minimize a parabolic fit to the L^2 norm of the residual $r(\omega)$ evaluated at $\omega = 0{\cdot}0$, $0{\cdot}5$ and $1{\cdot}0$. Let R denote $\|r\|$; minimizing R, we find $\omega = -\frac{1}{4}[4R(0{\cdot}5) - 3R(0) - R(1)]/D$ if $D = R(1) + R(0) - 2R(0{\cdot}5) > \tau > 0$ ($\tau \sim 10^{-3}$). If $R(0) > R(0{\cdot}5) > R(1{\cdot}0)$ (monotone decreasing), we take $\omega = 1$. If $R(0) < R(0{\cdot}5) < R(1)$ (monotone increasing), we take $\omega = 0{\cdot}25$ and check if $R(\omega)$ is less than $R(0)$; if $R(\omega)$ is still greater than $R(0)$, we bisect ω and repeat this process until the residual norm is decreased or ω is reduced to 0·03125. This limit is specified to avoid stagnation of the iteration with $\omega = 0{\cdot}0$.

To analyse the effect of shear thinning on the Newton iteration, let us examine in more detail the form of the Jacobian matrix. Considering the viscous term alone, the contribution to the finite

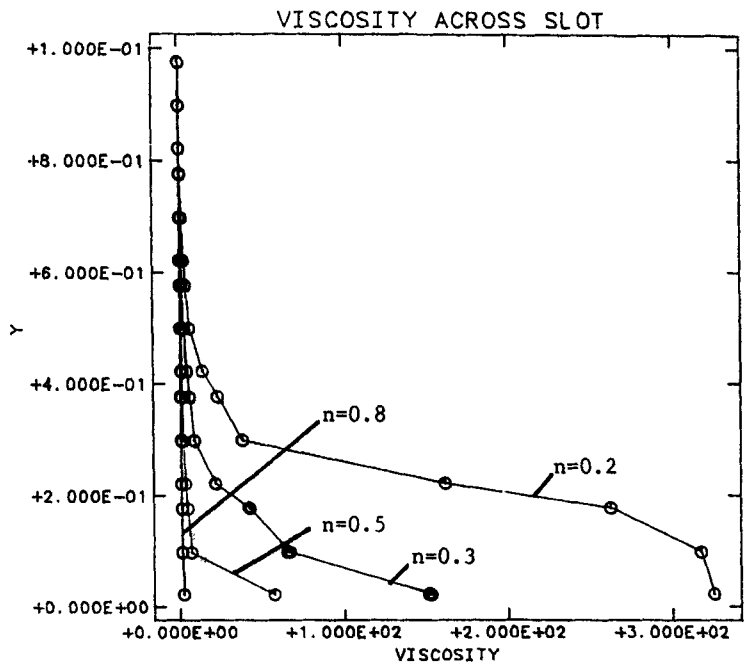

Figure 19. Viscosity profile across the channel

element system is

$$\mathbf{F}(\mathbf{u}) = \int_{\Omega} \eta \boldsymbol{\tau} : \nabla \mathbf{w} \, dx$$

and in the Jacobian system we have accordingly

$$J_{ij} = \frac{\partial F_i}{\partial u_j} = \int_{\Omega} \eta \frac{\partial}{\partial u_j} (\boldsymbol{\tau} : \nabla \mathbf{w}) \, dx + \int_{\Omega} \frac{\partial \eta}{\partial u_j} \boldsymbol{\tau} : \nabla \mathbf{w} \, dx.$$

Clearly the structure of the Jacobian will be strongly influenced by the behaviour of η and $\partial \eta / \partial u_j$ in the flow domain. Recalling Figure 19, we see that η is small in the shear-thinning layer and varies by two orders of magnitude across the channel. For this case it is straightforward to verify that $\partial \eta / \partial u_j$ behaves similarly. It follows that the entries in τ_{ij} will be small in the layer relative to the centre of the flow zone. If the nodes are ordered naturally (left to right and bottom to top), then entries of matrix $\mathbf{J}$ will range from large values in the upper left corner to small values in the lower right corner. This implies, in turn, that the conditioning of $\mathbf{J}$ will deteriorate as n decreases. The following mathematical analysis demonstrates that qualitatively the behaviour of the Newton iteration deteriorates as the conditioning worsens.

Following Ortega and Rheinboldt,[19] the fixed-point condition for convergence of Newton iteration implies that the spectral radius satisfies

$$\rho[\mathbf{I} - \mathbf{J}^{-1}(\mathbf{u}_n)\mathbf{J}(\mathbf{u})] < 1,$$

where $\mathbf{u}$ is the solution to the system and $\mathbf{u}_n$ is the current iterate. Let $\mathbf{E}_n = \mathbf{J}(\mathbf{u}_n) - \mathbf{J}(\mathbf{u})$ so that

$$\mathbf{J}(\mathbf{u}_n) = \mathbf{J}(\mathbf{u}) + \mathbf{E}_n = \mathbf{J}(\mathbf{u})(\mathbf{I} + \mathbf{J}^{-1}(\mathbf{u})\mathbf{E}_n).$$

For $\mathbf{E}_n$ small we then have

$$\mathbf{J}^{-1}(\mathbf{u}_n) \approx (\mathbf{I} - \mathbf{J}^{-1}(\mathbf{u})\mathbf{E}_n)\mathbf{J}^{-1}(\mathbf{u})$$

so that the spectral condition reduces to

$$\rho(\mathbf{J}^{-1}(\mathbf{u})\mathbf{E}_n) < 1.$$

Next let $\|\mathbf{E}_n\| = e\,\|\mathbf{J}(\mathbf{u})\|$ with e of bounded size (e represents the 'relative error') so that

$$\rho(\mathbf{J}^{-1}(\mathbf{u})\mathbf{E}_n) \approx e \text{ cond } (\mathbf{J}(u)),$$

where cond (**J**) is the condition number of **J**. Using this approximation in the spectral inequality, we have the qualitative result for convergence

$$\text{cond } (\mathbf{J}(\mathbf{u})) \leqslant 1/e.$$

From numerical calculations of channel flow, the condition number of the Jacobian increases by two orders of magnitude as n changes from 0·8 to 0·2. Hence, as the condition number increases with shear thinning, the convergence deteriorates.

Solutions for cavity flow with a 5 × 5 mesh and the penalty method were first computed, and the performances of a frontal elimination solver and of ORTHOMIN for this problem are given in Table IX. Again the factored Stokes flow problem was used for preconditioning. The iterative methods are competitive over the range of n considered even for this coarse mesh size. As $n \to 0$, the non-linearity is strong so the preconditioning deteriorates and the iterative method performs less well.

The problem was also computed using the mixed formulation and ORTHOMIN to yield the iteration results in Table X. In this instance the iterative performance was inferior and failed for $n < 0{\cdot}3$.

CONCLUSIONS

Accelerated gradient-type iterative methods such as the Lanczos method, ORTHOMIN, ORTHORES and the generalized minimum residual methods have been the subject of considerable research and practical interest to numerical analysts. Here we apply these techniques to the Jacobian systems for non-linear Navier–Stokes problems and power-law generalized Newtonian flows. Both mixed and penalty formulations are considered and the influence of continuation in Reynolds number *Re* and power-law index are explored.

Table IX. CPU time of power-law cavity flow; 5 × 5 mesh

Power-law index	CPU time frontal solver	CPU time ORTHOMIN	Newton iterations/ORTHOMIN iterations
0·8	0·1750	0·06648	4/16 (1, 5, 5, 5)
0·5	0·2625	0·1472	6/50 (1, 8, 11, 10, 10, 10)
0·3	0·3531	0·2955	8/126 (1, 11, 19, 15, 16, 19, 22, 23)
0·2	0·6703	0·9116	13/428 (1, 15, 17, 19, 23, 25, 27, 47, 71, 53, 48, 47)
0·1	0·9078	2·2461	20/802 (1, 19, 27, 24, 27, 27, 52, 56, 46, 45, 50, 56, 42, 42, 41, 48, 60, 48, 48, 42)

Table X. Power-law cavity flow (mixed formulation); 5 × 5 mesh

Power-law index	CPU time frontal solver	CPU time ORTHOMIN	Newton iterations/ORTHOMIN iterations
0·8	0·2511	0·2933	4/17 (1, 5, 6, 5)
0·5	0·3767	0·6137	6/79 (1, 9, 20, 22, 13, 14)
0·3	0·4395	2·2476	10/409 (1, 17, 66, 42, 48, 43, 39, 44, 51, 58)
0·2	0·6906		6/171 (1, 20, 37*, 39*, 37*, 37*)
0·1	1·7578		Diverges

* See footnote to Table II.

Numerical experiments for the penalized Navier–Stokes solution of the driven cavity problem compare the frontal solution and iterative performance with ORTHOMIN. Preconditioning is seen to be an important and sensitive issue and iterative performance for the linear problem is poor. To accommodate this in the penalized non-linear problem, we precondition using the factored Stokes (linear) operator. ORTHOMIN is then seen to be superior at low *Re* and, provided incremental continuation in *Re* is employed, for higher values of *Re*. To accelerate the scheme as *Re* increases, a new preconditioner is constructed by complete refactorization at each continuation step. Performances of ORTHOMIN, ORTHORES and GMRES are compared for a test problem and the effect of restarting is seen to be important.

The iterative methods were also applied to power-law fluids exhibiting shear thinning. A line-search strategy allows computation of solutions at low power-law index. The effect of shear thinning on the Jacobian is analysed and related to the conditioning. This explains the deterioration in iterative performance as the power-law index is decreased.

ACKNOWLEDGEMENTS

This research has been supported in part by the Department of Energy and the Texas Advanced Technology Program (ATP). We also wish to thank Tom Oppe for helping us in the use of NSPCG and David Young and David Kincaid for their comments.

REFERENCES

1. I. Babuska, 'The finite element method with penalty', *Math Comput.*, **27**, 221–228 (1973).
2. T. J. R. Hughes, W. K. Liu and A. Brooks, 'Review of finite element analysis of incompressible viscous flows by the penalty function formulation', *J. Comput. Phys.*, **30**, 1 (1979).
3. G. F. Carey and R. Krishnan, 'Penalty finite element method for the Navier–Stokes equations', *Comput. Methods Appl. Mech. Eng.*, **42**, 183–224 (1982).
4. P. M. Gresho, R. L. Sani, R. L. Lee and D. F. Griffiths, 'The cause and cure (?) of the spurious pressures generated by certain FEM solutions of the incompressible Navier–Stokes equations: Part 1', *Int. j. numer. methods fluids*, **1**, 17 (1981).
5. P. M. Gresho, R. L. Sani, R. L. Lee, D. F. Griffiths and M. S. Engleman, 'The cause and cure (?) of the spurious pressures generated by certain FEM solutions of the incompressible Navier–Stokes equations: Part 2', *Int. j. numer. methods fluids*, **1**, 171 (1981).
6. J. T. Oden, N. Kikuchi and Y. J. Song, 'Penalty finite element methods for the analysis of Stokesian flows', *Comput. Methods Appl. Mech. Eng.*, **31**(3), 297–330 (1982).
7. G. F. Carey and J. T. Oden, *Finite Elements: A Second Course*, Prentice-Hall, Englewood Cliffs, NJ, 1983.
8. G. F. Carey and J. T. Oden, *Finite Elements: Fluid Mechanics*, Prentice-Hall, Englewood Cliffs, NJ, 1986.
9. L. Hageman and D. Young, *Applied Iterative Methods*, Academic Press, New York, 1981.
10. P. K. W. Vinsome, 'Orthomin, an iterative method for solving sparse sets of simultaneous linear equations', *4th Symp. Numerical Solutions of Reservoir Performance*, Soc. Petroleum Eng., ASME, Los Angeles, CA, paper SPE 5739 (1976).
11. C. Lanczos, 'Solution of systems of linear equations by minimized iterations', *J. Res. NBS*, **49**, 33–53 (1952).

12. R. Fletcher, 'Conjugate gradient methods for indefinite systems', *Lecture Notes in Mathematics, Vol. 506*, Springer-Verlag, Berlin and New York, 1976.
13. D. M. Young and K. C. Jea, 'Generalized conjugate-gradient acceleration of non-symmetrizable iterative methods', *Lin. Alg. and its Applications*, **34**, 159–194 (1980).
14. Y. Saad and M. H. Schultz, 'GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems', *SIAM J. Sci. Stat. Comput.*, **17**(3), 856–869 (1986).
15. C. C. Paige and M. A. Saunders, 'Solution of sparse indefinite systems of linear equations', *SIAM J. Numer. Anal.*, **12**(4), 617–629 (1975).
16. M. Fortin and R. Glowinski, *Augmented Lagrangian Methods: Applications to the Numerical Solution of Boundary-Value Problems*, North-Holland, Amsterdam, 1983.
17. R. B. Bird, R. C. Armstrong and O. Hassoger, *Dynamics of Polymetric Liquids, Vol. I, Fluid Mechanics*, Wiley, New York, 1977.
18. D. K. Gartling, 'Finite element methods for non-Newtonian flows', *Sandia Report SAD85-1704*, 1986.
19. J. M. Ortega and W. C. Rheinboldt, *Iterative Solution of Nonlinear Equations in Several Variables*, Academic Press, New York, 1970.