INTERNATIONAL JOURNAL FOR NUMERICAL METHODS IN FLUIDS, VOL. 4, 147–158 (1984)

THE EFFECT OF THE CHOICE OF THE REFERENCE PRESSURE LOCATION IN NUMERICAL MODELLING OF INCOMPRESSIBLE FLOW

C. P. JACKSON

Theoretical Physics Division, AERE Harwell, Harwell, Oxon, OX11 ORA, U.K.

SUMMARY

The discontinuity of a finite-element pressure field that is sometimes present in the neighbourhood of the pressure-specification-point is shown to arise either from round-off, or from mistakes in modelling. The implications of this are considered. In particular it restricts grid refinement near the pressure-specification-point. The analysis can be extended to finite-difference calculations, and to other fields governed by equations similar to Poisson's equation.

KEY WORDS Finite Element Incompressible Flow Pressure-specification-point

1. INTRODUCTION

It is well known that it is difficult and costly to calculate numerically pressure fields in incompressible materials. The problems are caused by the infinite range and infinite speed of propagation of the pressure field in such materials. Thus, for example, iterative methods are slow to converge because they correspond to finite propagation speeds, and hence direct solvers such as fast Poisson solvers or Gaussian elimination are often used.

We report here on yet another consequence of the infinite range of the pressure, namely the sensitivity of the pressure field to the effects of round-off or mistakes in modelling. This manifests itself in the following manner. In incompressible materials the pressure field is only determined by the governing equations up to an additive constant, which may be fixed by specifying the value of the pressure at one point. If the location of this point is changed then, of course, the analytic pressure field is merely shifted by a constant. This is not true for the numerically determined pressure field. It is true that, away from the neighbourhood of the pressure-specification points, the change is simply a constant shift, but in the neighbourhood of the pressure-specification point the change is more complex. We shall show in this paper that nearly all of the round-off accumulates at the pressure specification point, perturbing the fields at that point.

A specific example may help to make this clearer. We first discovered the phenomenon in finite element modelling of the 'driven cavity' problem,[1] of incompressible laminar flow in a square cavity, across the top of which a lid slides at constant speed (see Figure 1). If the pressure-specification point was at the centre of the cavity then the pressure field was effectively smooth everywhere (except in the neighbourhood of the two corners under the lid, where there are of course logarithmic singularities in pressure), but if the pressure-specification point was at one of the lower two corners, say, then the pressure field was such

0271–2091/84/020147–12$01.20

Received 2 July 1982
Revised 21 November 1982

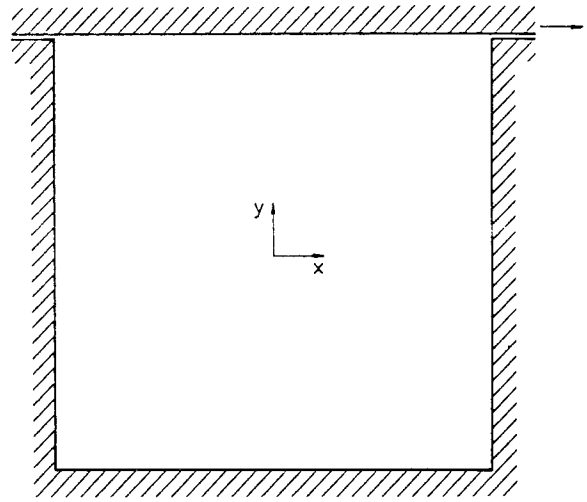

Figure 1. The 'driven-cavity' problem. The lid slides across the top of the cavity with constant velocity of 1·0

that there existed a constant which, when added to the pressure field, gave the same values as before everywhere except at the pressure-specification point where there was a sudden jump. Further, if the finite-element mesh in the neighbourhood of the pressure-specification point was refined then the magnitude of the jump increased, contrary to one's natural expectation that grid refinement improves the accuracy of numerical solutions.

In this paper we show that the phenomenon is a result of round-off, which effectively introduces into the finite-element model a flow imbalance that leaves the model through the pressure-specification point, thus perturbing there the pressure field and to a lesser extent the velocity field.

It is also possible to introduce such a flow imbalance by making a mistake in modelling a problem, and then the flow imbalance again effectively leaves the model through the pressure-specification point, perturbing the pressure and velocity fields there.

A similar phenomenon occurs for finite-element pressure fields in other problems. It occurs whenever the governing partial differential equations only determine the pressure up to an additive constant, as for example if the pressure is determined by Poisson's equation. The same effect can also occur in finite-difference calculations.[2]

It must be stressed that we are not saying that the finite-element pressure problem is ill-posed or badly conditioned, but rather that the finite-element pressure field may be inaccurate in the neighbourhood of the pressure-specification point.

2. THE FINITE ELEMENT EQUATIONS

The Navier–Stokes equations describing incompressible laminar fluid flow in two dimensions are in non-dimensional form

$$\begin{aligned} u\frac{\partial u}{\partial x}+v\frac{\partial u}{\partial y}+\frac{\partial p}{\partial x}-\frac{1}{Re}\nabla^2 u=0 \\ u\frac{\partial v}{\partial x}+v\frac{\partial v}{\partial y}+\frac{\partial p}{\partial y}-\frac{1}{Re}\nabla^2 v=0 \\ \frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}=0 \end{aligned} \tag{1}$$

where u, v are velocity components in the x, y directions, respectively, p is pressure and Re is Reynolds number. In any given problem appropriate boundary conditions have to be specified. For example, in the driven cavity problem the boundary conditions are

$$\begin{aligned} &u = 1, && \text{for} \quad x = 0 \\ &u = v = 0, && \text{for} \quad y = 0 \\ &u = v = 0, && \text{for} \quad x = 1 \\ &u = 1, \quad v = 0, && \text{for} \quad y = 1 \end{aligned} \tag{2}$$

where the cavity is

$$0 \leqslant x \leqslant 1, \qquad 0 \leqslant y \leqslant 1$$

in non-dimensional units. Note that since only gradients of pressure appear in the equations then the pressure is only determined up to an additive constant. This may be fixed by specifying the pressure value at one point.

The finite-element method for (1) is described in detail by many authors, and so we simply present here the standard form of the finite-element equations. It is of course necessary to use 'mixed interpolation' to ensure that the finite-element pressure field is uniquely determined.[3]

The finite-element equations are

$$\begin{aligned} &\int_{\hat{\Omega}} \psi_i \left(\hat{u} \frac{\partial \hat{u}}{\partial x} + \hat{v} \frac{\partial \hat{u}}{\partial y} + \frac{\partial \hat{p}}{\partial x} \right) + \int_{\hat{\Omega}} \frac{1}{Re} \nabla \psi_i \cdot \nabla \hat{u} = 0 \\ &\int_{\hat{\Omega}} \psi_j \left(\hat{u} \frac{\partial \hat{v}}{\partial x} + \hat{v} \frac{\partial \hat{u}}{\partial y} + \frac{\partial \hat{p}}{\partial y} \right) + \int_{\hat{\Omega}} \frac{1}{Re} \nabla \psi_j \cdot \nabla \hat{v} = 0 \\ &\int_{\hat{\Omega}} \phi_k \left(\frac{\partial \hat{u}}{\partial x} + \frac{\partial \hat{v}}{\partial y} \right) = 0 \end{aligned} \tag{3}$$

where $\{\psi_i\}$ is the set of piecewise biquadratic polynomials used to interpolate the finite-element velocity fields $\hat{u}, \hat{v}$, and $\{\phi_i\}$ is the set of piecewise bilinear polynomials used to interpolate the finite-element pressure field $\hat{p}$, and $\hat{\Omega}$ is the finite element approximation to Ω.

One way of implementing Dirichlet boundary conditions is simply to discard the finite-element equations corresponding to the boundary nodes, and use instead the Dirichlet boundary condition at the nodes. This ensures that there are still the same number of equations as unknowns.

Thus at the pressure-specification-point Q the finite-element continuity equation corresponding to node Q is replaced by an equation of the form

$$p_Q = \text{value}$$

which defines the reference level for the pressure.

3. THE PRESSURE-SPECIFICATION POINT

Although at first sight the pressure-specification point Q has been singled out for special treatment as regards the continuity equation, it can be shown that the finite-element continuity equation for node Q is implied by the other finite-element continuity equations.

We have

$$\int_{\hat{\Omega}} \phi_i \left(\frac{\partial \hat{u}}{\partial x} + \frac{\partial \hat{v}}{\partial y} \right) = 0, \qquad i \neq Q \tag{4}$$

Now the space of finite-element test functions contains 1, usually in the form

$$\sum_i \phi_i = 1 \tag{5}$$

(The modifications to the argument below if 1 is some different linear combination of the ϕ_i are trivial). Thus summing (4) over $i \neq Q$ and using (5) we have

$$\int_{\hat{\Omega}} (1 - \phi_Q) \left(\frac{\partial \hat{u}}{\partial x} + \frac{\partial \hat{v}}{\partial y} \right) = 0$$

and so

$$\begin{aligned} \int_{\hat{\Omega}} \phi_Q \left(\frac{\partial \hat{u}}{\partial x} + \frac{\partial \hat{v}}{\partial y} \right) &= \int_{\hat{\Omega}} \left(\frac{\partial \hat{u}}{\partial x} + \frac{\partial \hat{v}}{\partial y} \right) \\ &= \oint_{\partial \hat{\Omega}} \hat{\mathbf{u}} \cdot \mathbf{n} \\ &= 0 \end{aligned} \tag{6}$$

provided that

$$\int_{\hat{\Omega}} \hat{\mathbf{u}} \cdot \mathbf{n} = 0$$

This last condition is just the requirement that there is no net influx or outflux of material into the region $\hat{\Omega}$. Thus, in a problem with Dirichlet boundary conditions on velocity on all the boundaries, the finite-element continuity equation at node Q is implied by the other finite-element continuity equations, provided that the boundary conditions are chosen carefully so that material is conserved in the sense that $\oint \hat{\mathbf{u}} \cdot \mathbf{n} = 0$. This requires a little more care than at first sight appears necessary, since it is a requirement upon the finite-element field $\hat{\mathbf{u}}$ and not upon $\mathbf{u}$.

For example, consider the driven cavity problem. Although $\oint \mathbf{u} \cdot \mathbf{n}$ is zero because $\mathbf{u} \cdot \mathbf{n}$ is zero on all the walls of the cavity, $\hat{\mathbf{u}} \cdot \mathbf{n}$ is not necessarily zero on all the walls and so $\oint \hat{\mathbf{u}} \cdot \mathbf{n}$ is only zero if care is taken. In one possible formulation of the boundary conditions $\hat{u}$ on the left-hand and right-hand edges is 0 at all the nodes except the ones on the lid, where it is 1. This then implies that $\hat{u}$ is non-zero on the left-hand and right-hand edges, on the elements immediately below the lid, since it is a low order polynomial fit to the nodal values. Thus $\hat{\mathbf{u}} \cdot \mathbf{n}$ is non-zero on these two elements and zero everywhere else on the boundary, and so $\oint \hat{\mathbf{u}} \cdot \mathbf{n}$ is only zero if the sizes of the two elements in question are the same.

Now the above proof that the continuity equation for node Q is implied by the remaining continuity equations assumed exact arithmetic, but of course computer calculations use finite-precision arithmetic. Perhaps the easiest way to understand the effect of this is as follows. Finite precision arithmetic leads to nodal values of the fields $\mathbf{u}'_i$, p' differing slightly from $\hat{\mathbf{u}}_i$, $\hat{p}$ the solutions in exact arithmetic of the finite element equations. Thus writing

$$\mathbf{u}' = \sum \phi_i \mathbf{u}'_i \tag{7}$$

we obtain

$$\int_{\Omega} \phi_i \left(\frac{\partial u'}{\partial x} + \frac{\partial v'}{\partial y} \right) = \eta_i \tag{8}$$

say, where we have used exact arithmetic in the evaluation of the integral in (8), (and inexact values $\mathbf{u}'$ of course). The magnitude of η_i depends upon the program in use. For example, in the program that we were using when we first discovered the effect, all calculations were performed in IBM single precision (with approximately 7 significant figures) except the solution of the linear equations which was done in IBM double precision (with approximately 16 significant figures) by the frontal method of Gaussian elimination. In this case η_i is dominated by the round-off in the calculation of the matrix and residual for the linear equations, and so is of order of machine precision times a typical value for $\partial u/\partial x$ or $\partial v/\partial y$, times a small number (possibly as large as 5·0) determined by the number of arithmetic operations required in the calculation of the matrix and residual. On the other hand, if all calculations had been carried out in single or double precision, then η_i would be dominated by the round-off growth in the solution of the linear equations, and could be considerably larger, particularly if the problem was poorly conditioned. Note that η_i must not be confused with the value calculated in the computer for the integral in (8), which will differ from η_i, since additional round-off will occur in its calculation, although it will be of the same order of magnitude as η_i. Hence, following the same line of argument as before

$$\int_{\hat{\Omega}} (1 - \phi_Q) \left(\frac{\partial u'}{\partial x} + \frac{\partial v'}{\partial y} \right) = \sum_{i \neq Q} \eta_i \tag{9}$$

and so

$$\begin{aligned} \int_{\hat{\Omega}} \phi_Q \left(\frac{\partial u'}{\partial x} + \frac{\partial v'}{\partial y} \right) &= \int_{\hat{\Omega}} \left(\frac{\partial u'}{\partial x} + \frac{\partial v'}{\partial y} \right) - \sum_{i \neq Q} \eta_i \\ &= \oint_{\partial \hat{\Omega}} \mathbf{u}' \cdot \mathbf{n} - \sum_{i \neq Q} \eta_i \end{aligned}$$

Now, in a problem with Dirichlet boundary conditions on the velocity on all the boundary we have that $\mathbf{u}'$ is prescribed on the boundary and is such that

$$\oint_{\partial \hat{\Omega}} \mathbf{u}' \cdot \mathbf{n} = 0$$

provided of course that there are no mistakes in modelling. Thus

$$\int_{\hat{\Omega}} \phi_Q \left(\frac{\partial u'}{\partial x} + \frac{\partial v'}{\partial y} \right) = - \sum_{i \neq 0} \eta_i = E \tag{10}$$

say. One way of interpreting this result is to say that effectively all the round-off in the continuity equation is accumulated at the pressure-specification point.

The magnitude of E can be estimated in the case discussed above when all calculations except the solutions of the linear equations are carried out in IBM single precision. It is not unreasonable to postulate that each of the η_is can be treated as independent random variables, and so using the central limit theorem the magnitude of E is expected to be

$$\varepsilon \sqrt{N} \left(\text{typical value of } \frac{\partial u}{\partial x}, \frac{\partial v}{\partial y} \right) \times (\text{a small number}) \tag{11}$$

where N is the number of nodes in the grids and ε is machine precision. If the grid is uniform with spacing h then this can be rewritten as

$$\varepsilon \frac{1}{h}\left(\text{typical value of } \frac{\partial u}{\partial x}\right) \tag{12}$$

In other cases the η_i may be correlated and an estimate is harder to obtain.

It is reasonable to suggest that the effect of E is to perturb the velocity field in the neighbourhood of Q by a source/sink of fluid. Thus, using (10) and (4), and approximating the integrals, the velocity fields at Q will be perturbed by Δu $(=u'-\hat{u})$ and Δv $(=v'-\hat{v})$ such that

$$A_Q\left(\frac{\Delta u}{h_Q}+\frac{\Delta v}{h_Q}\right)\simeq E \tag{13}$$

where A_Q is the area of elements including node Q, and h_Q is the mesh spacing near Q. (The round-off in the computer calculation of u_Q', v_0' will be of order η_i and so considerably less than E and can be safely ignored in this estimate). Then the pressure at Q will be perturbed by Δp where $\Delta p = p' - \hat{p}$, such that

$$A_q \frac{\Delta p}{h_Q} \simeq A_Q \frac{1}{h_Q} \frac{(\Delta u \text{ or } \Delta v)}{h_Q} \frac{1}{Re} \tag{14}$$

using the velocity equations and neglecting the non-linear term for the moment. Thus

$$\Delta p \simeq \frac{1}{Re}\frac{E}{A_Q} \quad \text{or} \quad \frac{1}{Re}\frac{\Delta u}{h_Q} \tag{15}$$

That is the magnitude of the pressure jump is directly proportional to the flow imbalance, and inversely proportional to the area of the elements containing the pressure-specification point. It is easy to see that including the effect of the non-linear terms from the velocity equation merely changes the details of the argument above and not the magnitude of Δp.

If the flow imbalance E is due to a modelling error, rather than to the accumulation of round-off, then a similar argument can be stated exactly as before from equation (13) onwards.

There are several implications of the above. In the case when E is due to round-off, then if the grid is locally refined near Q this should only change the η_i corresponding to nodes near Q, and so E should remain more or less the same. Equally if E is due to a modelling error then local refinement near Q will in general leave E unchanged. Then in both cases Δp will increase in direct proportion as A_Q increases. Thus local refinement near the pressure specification point is not a good idea.

In either case if the grid is uniform and the pressure-specification point is in the middle of the flow then there will be more elements adjacent to it and A_Q will be larger than if the pressure-specification point was on an edge or a corner, and so Δp will be smaller. Also the velocity perturbation will be smaller relative to the true flow field.

In the case when E is due to round-off then if Δp is compared on a sequence of uniform grids of different mesh spacing h then

$$\Delta p \simeq \frac{1}{Re}\frac{E}{h^2} \simeq \varepsilon h^{-3}\left(\text{typical value of } \frac{\partial u}{\partial x}\right) \times (\text{small number})$$

from (12) and (15), where the last equality holds only in the case when all calculations except matrix solutions are carried out in IBM single precision. Note that there will of course not be

exact proportionality between Δp and h^{-3} because of the random nature of round-off. The magnitude of the effect is proportional to machine precision, and so in IBM single precision (with approximately 7 significant figures) then the effect can become noticeable at quite low refinements, but in double precision, the effect should usually be quite negligible unless the pressure-specification point is in an extremely refined region, or there are strongly varying material properties.

4. NUMERICAL EXAMPLES

In order to test the theory presented above a large number of numerical calculations were made. Three problems were used, namely the driven cavity problem, Poiseuille flow and Couette flow, all at various Reynolds numbers. In all cases the results were in excellent agreement with the theory. We present here only a few of the results to save paper.

The first problem considered was the driven cavity. The finite-element equations were solved in double precision, but all other calculations were performed in IBM single precision, which means that numbers have about 7 significant digits. A variety of finite-element grids were used, both uniform grids of the form shown in Figure 2 with various grid spacings, and also such grids with the lower left corner refined as shown in Figure 3. The elements used were the six-node quadratic triangle for velocities, and the three-node linear triangle for pressure.

In Tables I and II we present the calculated pressure fields at Reynolds number 1·0 on a 17×17 uniform grid, with the pressure-specification point at the centre of the cavity and the lower left corner, respectively. The reference pressure value was taken to be 0·0 at the pressure specification point. (We use the notation that an $N \times N$ grid is one with N nodes in each direction). In Table III we present the difference between the two solutions. It can be seen that, everywhere except at the pressure-specification points, there is a constant difference of 0·32, approximately, and that the magnitude of the pressure discontinuity Δp is approximately $1{\cdot}8 \times 10^{-3}$ when the pressure specification point is at the lower left corner. In a similar manner the pressure discontinuity with the pressure-specification point at the lower left corner was obtained for 9×9, 13×13 and 21×21 grids. The results are plotted in Figure 4, against the inverse of the mesh spacing, on a log–log plot. A line of slope 3·0 is also drawn

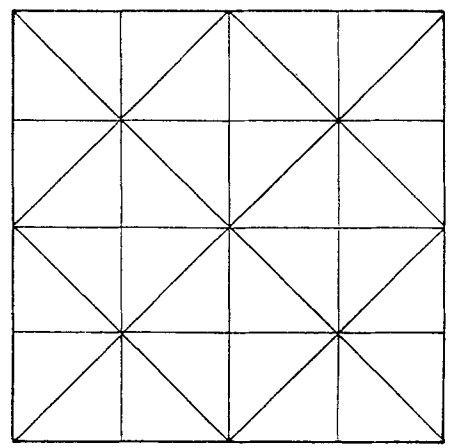

Figure 2. A 9×9 uniform finite-element grid of the sort used for the numerical calculations of this paper

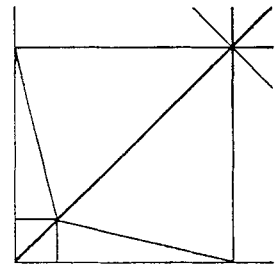

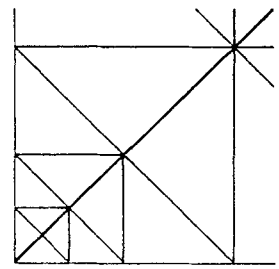

Figure 3. Two types of corner refinement used in the finite-element grids

Table I. Pressure values for the driven cavity problem on a 17×17 grid with the pressure-specification point at the centre of the grid

−38·070	−2·7006	−4·3529	−0·15627	−0·031841	0·14362	4·4146	2·8829	38·551
−15·011	−9·0836	−3·6213	−1·7055	−0·045934	1·6281	3·6398	9·2313	15·229
−5·2009	−5·1325	−3·1837	−1·5703	−0·055403	1·4900	3·1684	5·1932	5·2739
−1·4368	−2·5381	−2·0872	−1·0386	−0·027490	1·0006	2·0959	2·5786	1·4565
−0·79758	−1·1102	−1·1080	−0·65878	0·0	0·67005	1·1385	1·1480	0·83418
−0·00091011	−0·57863	−0·57902	−0·36171	0·017288	0·39911	0·62368	0·62627	0·042002
−0·067773	−0·28993	−0·38735	−0·22327	0·023553	0·27255	0·43894	0·33991	0·11755
−0·18001	−0·33594	−0·34708	−0·23091	0·025780	0·28272	0·39935	0·38817	0·23211
−0·32179	−0·44578	−0·53550	−0·36882	0·025579	0·42078	0·58827	0·49791	0·37340

Table II. Pressure values for the driven cavity problem on a 17×17 grid with the pressure-specification point at the lower left corner

−37·750	−2·3806	−4·0329	0·16373	0·28816	0·46362	4·7346	3·2029	38·871
−14·691	−8·7636	−3·3013	−1·3855	0·27407	1·9481	3·9598	9·5513	15·549
−4·8809	−4·8124	−2·8637	−1·2503	0·26460	1·8101	3·4884	5·5132	5·5939
−1·1168	−2·2181	−1·7672	−0·71855	0·29254	1·3206	2·4159	2·8986	1·7765
−0·47757	−0·79024	−0·78797	−0·33874	0·31991	0·99008	1·4585	1·4680	1·1542
0·31904	−0·25863	−0·25902	−0·041693	0·33732	0·71912	0·94368	0·94628	0·36201
0·25236	0·030061	−0·067329	0·096740	0·34355	0·59255	0·75895	0·65991	0·43755
0·13951	−0·015953	−0·027083	0·089093	0·34578	0·60273	0·71935	0·70817	0·55211
0·0	−0·12626	−0·21536	−0·048876	0·34559	0·74078	0·90827	0·81791	0·69341

Table III. Difference between Table I and Table II

0·320	0·3200	0·3200	0·32000	0·32000	0·32000	0·3200	0·3200	0·320
0·320	0·3200	0·3200	0·3200	0·32000	0·3200	0·3200	0·3200	0·320
0·3200	0·3201	0·3200	0·3200	0·32000	0·3201	0·3200	0·3200	0·3200
0·3200	0·3200	0·3200	0·3201	0·32003	0·3200	0·3200	0·3200	0·3200
0·32001	0·3200	0·3200	0·32004	0·31991	0·32003	0·3200	0·3200	0·3200
0·31995	0·32000	0·32000	0·32002	0·32003	0·32001	0·32000	0·32001	0·32001
0·32013	0·31999	0·32002	0·32001	0·31997	0·32000	0·32001	0·32000	0·32000
0·31952	0·32001	0·32000	0·32000	0·32000	0·32001	0·32000	0·32000	0·32000
0·32179	0·31952	0·32014	0·31994	0·32001	0·32000	0·32000	0·32000	0·32001

to indicate the expected behaviour. The agreement is excellent bearing in mind the (pseudo-) random nature of round-off.

Similarly, the pressure discontinuity with the pressure specification point in the lower left corner was obtained for a number of 21×21 grids with varying refinements in the corner. The results are plotted in Figure 5 against the inverse of the local mesh spacing, using a log–log plot. The slope of the line joining the points is 2·26 in good agreement with the expected value of 2·0. Note that the magnitude of the pressure discontinuity is 2·654 when the local mesh spacing is reduced by a factor of 16 compared to the uniform grid. This is of the same order of magnitude as the pressure field values over most of the cavity, so that the perturbation is not insignificant.

We also obtained good agreement with the predicted behaviour of the pressure discontinuity in the cases of Couette flow and Poiseuille flow. We report only a few of the results

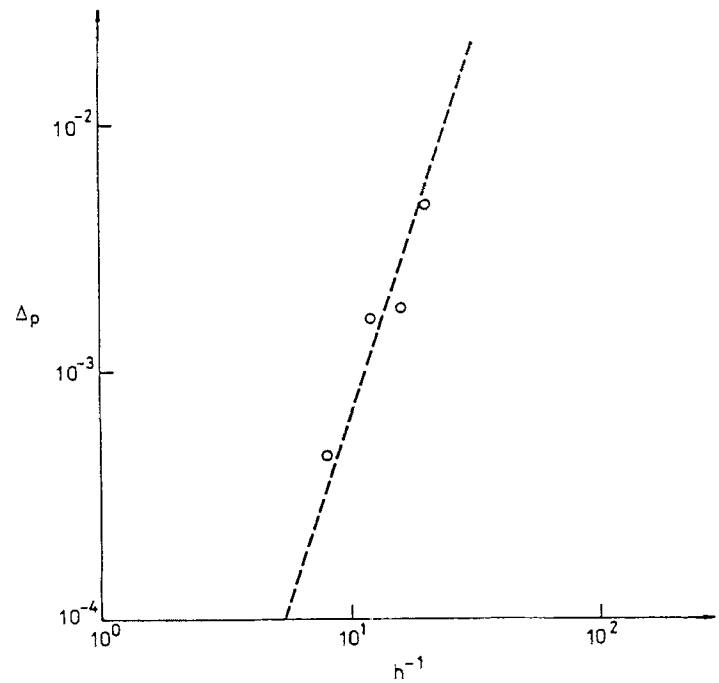

Figure 4. A log–log of the pressure discontinuity for the 'driven cavity' problem with the pressure-specification-point in the lower left corner, versus the inverse of the mesh spacing for various uniform grids. The line of slope 3·0 indicates the expected behaviour of the plot

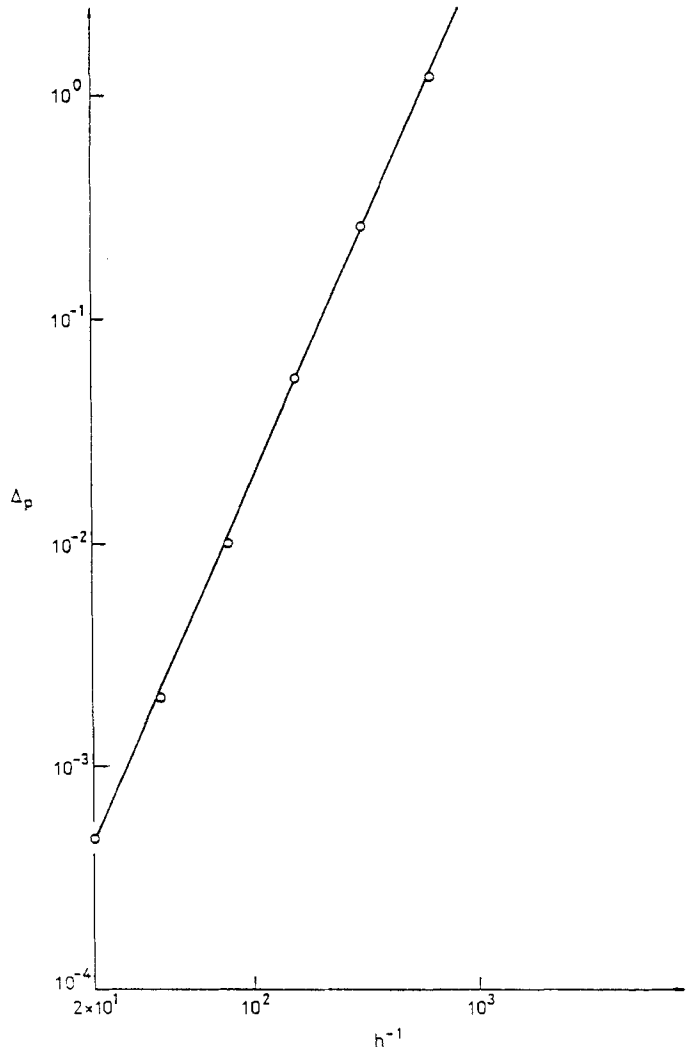

Figure 5. A log–log plot of the pressure discontinuity for the 'driven cavity' problem, with the pressure-specification point in the lower left corner, versus the inverse of the local mesh spacing for various refined 21×21 grids

here. The problem of Couette flow is that of flow in a straight channel with one wall fixed and one wall sliding at a constant velocity (see Figure 6). The boundary conditions are

$$\begin{aligned} u &= y, \quad v = 0, \quad \text{for} \quad x = 0 \\ u &= y, \quad v = 0, \quad \text{for} \quad x = 1 \\ u &= 0, \quad v = 0, \quad \text{for} \quad y = 0 \\ u &= 1, \quad v = 0, \quad \text{for} \quad y = 1 \end{aligned}$$

in non-dimensional units with the orientation of axes shown in Figure 6. This problem has the analytic solution

$$u = y, \qquad v = 0, \qquad p = \text{constant}.$$

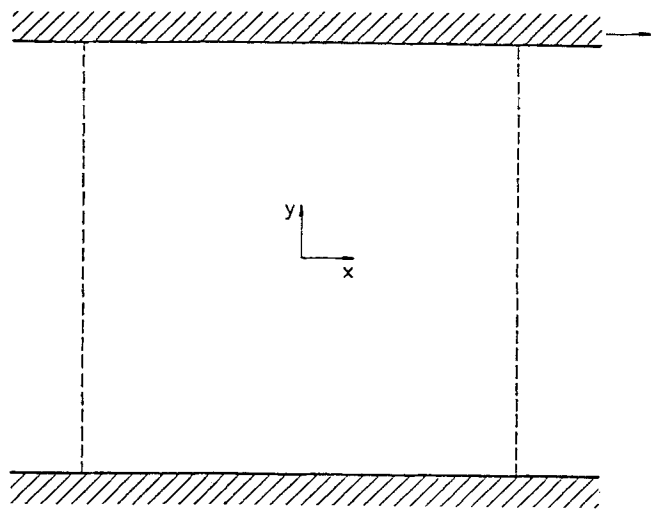

Figure 6. Couette flow. The upper channel wall moves with constant velocity 1·0

If a pressure discontinuity occurs, it is very easy to see in this problem because of the form of this solution. The pressure discontinuity behaved exactly as predicted. For example, on a 21×21 grid with $Re = 1$, if the pressure-specification point was at the centre of the right hand edge, then the pressure discontinuity was $1{\cdot}08 \times 10^{-3}$, whereas, if the pressure-specification point was at the bottom left-hand corner then the pressure discontinuity was $2{\cdot}71 \times 10^{-4}$, in agreement with the remarks made at the end of Section 3. One slightly amusing test that we carried out was to perturb the co-ordinates of the nodes by a (pseudo-) random perturbation on a scale of 10^{-4} times the grid scale. This changed the pressure discontinuity from $1{\cdot}08 \times 10^{-3}$ to $1{\cdot}68 \times 10^{-3}$ with the pressure specification point at the centre of the right hand edge. This is strong confirmation that round-off is the source of the trouble.

One last very sensitive test of the theory is simply to compare the same calculations in double and single precision. The magnitude of the pressure discontinuity should be smaller by a factor of about 10^9, when the calculation is performed in double precision on an IBM computer. This is exactly what was observed. The magnitude of the pressure discontinuity for the problem of Couette flow at $Re = 1$, on a 21×21 grid with the pressure specification point at the centre of the right hand edge is $1{\cdot}08 \times 10^{-3}$ when calculated in single precision, and $1{\cdot}26 \times 10^{-13}$ when calculated in double precision.

5. CONCLUSIONS AND REMARKS

The theory presented here provides a reasonable explanation of the pressure discontinuity sometimes seen near the pressure-specification point. The theory has been shown to be in excellent agreement with numerical experiments.

The basis of the argument is of course just a numerical analogue of the consistency condition for Neumann problems.

There are several implications of the paper. Firstly, it is perhaps desirable to perform all calculations in double precision. Secondly one should in any case avoid placing the pressure-specification point in regions of high refinement. Thirdly and perhaps most importantly, if

there is a significant discontinuity at the pressure-specification point after taking the above two remarks into consideration, then there is almost certainly a modelling error. This can be a very useful debugging aid.

It is easy to see that very similar results hold if the pressure is determined by the simpler Poisson-like equation

$$\frac{\partial}{\partial x}\left(a\frac{\partial P}{\partial x}\right)+\frac{\partial}{\partial y}\left(b\frac{\partial P}{\partial y}\right)=0$$

where a and b are given functions of position. The modifications to the argument of Section 3 in order to deal with this type of pressure equation are simple. A similar argument can also be made for any field, whose governing partial differential equations admit arbitrary additive shifts in the solution that need to be fixed by a boundary condition.

The argument is not restricted to finite-element discretizations. The same sort of argument can be made for finite-difference discretizations which are in conservative form.

ACKNOWLEDGEMENTS

I should like to thank my colleagues Dr. K. H. Winters and Mr. K. A. Cliffe for bringing this problem to my attention and for performing the numerical calculations, and I should also like to thank them and Dr. I. P. Jones for many helpful discussions.

REFERENCES

1. K. H. Winters and K. A. Cliffe, 'A finite-element study of driven laminar flow in a square cavity', *AERE Report R.9444* (1979).
2. I. P. Jones, Private Communication.
3. C. P. Jackson and K. A. Cliffe, 'Mixed interpolation in primitive variable finite-element formulations for incompressible flow', *Int. J. Num. Meth. Eng.*, **17** 1659–1688 (1981).