

***IFISS: A computational laboratory for
investigating incompressible flow problems***

Elman, Howard C. and Ramage, Alison and Silvester,
David J.

2012

MIMS EPrint: **2012.81**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

Reports available from: <http://eprints.maths.manchester.ac.uk/>

And by contacting: The MIMS Secretary
School of Mathematics
The University of Manchester
Manchester, M13 9PL, UK

ISSN 1749-9097

IFISS: A COMPUTATIONAL LABORATORY FOR INVESTIGATING INCOMPRESSIBLE FLOW PROBLEMS

HOWARD C. ELMAN*, ALISON RAMAGE†, AND DAVID J. SILVESTER‡

Abstract. The IFISS *Incompressible Flow & Iterative Solver Software* package contains software which can be run with MATLAB or Octave to create a computational laboratory for the interactive numerical study of incompressible flow problems. It includes algorithms for discretization by mixed finite element methods and *a posteriori* error estimation of the computed solutions, together with state-of-the-art preconditioned iterative solvers for the resulting discrete linear equation systems. In this paper we give a flavour of the code's main features and illustrate its applicability using several case studies. We aim to show that IFISS can be a valuable tool in both teaching and research.

Key words. Finite elements, incompressible flow, iterative solvers, mathematical software

AMS subject classifications. 97N80, 76D99, 65M60, 65N22

1. Introduction. Research in computational mathematics is often motivated by the results of numerical experiments. Frequently, a result is conjectured based on behaviour observed using a software environment such as MATLAB long before it is supported by any formal analytic results. In this sense, the “computational laboratory” plays just as important a role in modern mathematics as physical laboratories do in physics, chemistry, biology and engineering. A similar observation applies in the context of teaching computational mathematics: carrying out investigative numerical experiments on particular topics will help students learn how to formulate hypotheses, design simple experiments to test them and interpret the resulting data. As well as developing important deduction and interpretation skills, this hands-on approach is often more useful in helping students remember critical ideas over a significant period of time than a traditional textbook only method.

With this in mind, the *Incompressible Flow & Iterative Solver Software* (IFISS) toolbox [36] has been developed as a computational laboratory for the interactive numerical study of incompressible flow problems.¹ The software package includes algorithms for discretization by mixed finite element methods and *a posteriori* error estimation of the computed solutions, together with state-of-the-art preconditioned iterative solvers for the discrete linear systems that arise. The current release is IFISS 3.2. Version 2.0 was released in conjunction with the monograph *Finite Elements and Fast Iterative Solvers* [14] and used to perform the computational experiments described therein. An overview of this earlier version of IFISS can be found in [12]. The aim of this review article is therefore not to give a technical description of the software, but instead, to give a flavour of its usefulness in the sense of driving our research and in enhancing our teaching. In the sequel we focus on four specific examples of how we have used IFISS in this way: both as a teaching and research tool, and as a starting

*Department of Computer Science, University of Maryland, College Park, MD 20742, USA (elman@cs.umd.edu). Supported in part by the US National Science Foundation under grant DMS1115317 and the US Department of Energy under grant DE-SC0009301.

†Department of Mathematics and Statistics, University of Strathclyde, Glasgow G1 1XH, UK (A.Ramage@strath.ac.uk).

‡School of Mathematics, The University of Manchester, Manchester M13 9PL, UK (d.silvester@manchester.ac.uk).

¹The software package is freely available, see <http://www.manchester.ac.uk/ifiss> or <http://www.cs.umd.edu/~elman/ifiss>, and can be run using MATLAB or Octave. The source may be redistributed and/or modified under the terms of the GNU Lesser General Public License.

point for developing code for more specialized research applications.

2. Overview. The IFISS toolbox is structured as a stand-alone package for studying discretization algorithms for partial differential equations (PDEs) and for exploring and developing algorithms in numerical linear and nonlinear algebra for solving the associated discrete systems. There are currently twenty eight built-in test problems all based on common PDEs. Seventeen of these are elliptic PDE problems, and are associated with the Poisson equation or steady-state versions of the convection-diffusion equation, the Stokes equations or the Navier–Stokes equations. The remaining test problems are parabolic and involve the heat equation, advection-diffusion, or the Navier–Stokes equations modelling unsteady incompressible flow.

The first main feature of the package concerns problem specification and finite element discretization. For each PDE, IFISS offers a choice of two-dimensional domains on which the problem can be posed, along with boundary conditions and other aspects of the problem, and a choice of finite element discretizations on uniform or stretched quadrilateral element meshes. The package therefore makes it easy to study the accuracy of finite element solutions, different choices of elements, and *a posteriori* error estimates (see §3.1). In addition, special features associated with individual problems can be explored. These include the effects of boundary layers on solution quality for the convection-diffusion equation, and the effects of discrete inf-sup stability and adaptive time stepping on accuracy for the Stokes and Navier–Stokes equations.

Another important aspect of IFISS concerns iterative solution of the discrete algebraic systems, with emphasis on preconditioned Krylov subspace methods. The key for fast solution lies in the choice of effective preconditioning strategies. The IFISS package offers a range of options, including algebraic methods such as incomplete LU factorisations and algebraic multigrid, as well as more sophisticated state-of-the-art methods designed to take advantage of the specific structure of discrete flow equations. In addition, there is a choice of iterative strategies, Picard iteration or Newton’s method, for solving the nonlinear algebraic systems arising from Navier–Stokes problems. Again, IFISS offers scope for investigating issues like the choice of finite element approximation and the influence of the time step on the convergence of iterative solvers in a straightforward way (see §3.2).

There are of course many other finite element packages that could be used to investigate any of the above topics. As well as commercial codes (for example, ABAQUS, COMSOL and LS-DYNA) there are a number of excellent free software packages: in particular, ALBERTA [33], DUNE [6] and FEniCS [24]. A distinctive feature of IFISS is that it runs within a MATLAB environment. This means that the programming level is relatively high, so the source code is readable, portable and easily adaptable. In this sense, MATLAB is an ideal computational laboratory. Our finite element implementation is also very efficient. Throughout IFISS all local calculations (for example, quadrature in generating element matrices, application of essential boundary conditions, *a posteriori* error estimation) are vectorized over the elements—leading to efficient computation on the current generation of multicore processor architectures. We speculate that the majority of the IFISS code runs as fast as if it were coded in C or FORTRAN.

The design of IFISS not only facilitates the study of alternative discretizations and iterative solution algorithms, it also makes it easy to examine the interaction between the two and hence determine the overall solution cost. This feature is illustrated later using the example of the development of an optimal iterative solver for symmetric indefinite systems (see §3.3). As well as being a useful source of benchmark problems,

the package provides a convenient starting point for research projects that require the construction of new problem classes, discretizations or solution algorithms. A detailed description of how new problems can be incorporated into IFISS is given in the *Installation and Software Guide* [13]. An example of this type of “spin-off” development is described in §3.4.

3. Illustrative examples. In this section we describe four situations where IFISS has been used as a teaching or research tool. The descriptions given here are, by necessity, somewhat abbreviated. Readers who would like to learn more about the specific details should consult the cited references.

3.1. Investigating finite element convergence. When teaching finite element analysis, it is common to establish *a priori* and *a posteriori* error bounds which describe the convergence of the finite element approximation as the underlying grid is refined (see for example [3, 14, 21, 41] for an introduction to finite element approximation and [4] for a detailed derivation of classical error bounds). Having done this in class, it is easy to reinforce the practical importance of these ideas using numerical experiments within IFISS. (The following example uses diffusion test problems 3 and 4. Figures 3.1 and 3.3 can be generated using the IFISS function `sirev_ex3_1`).

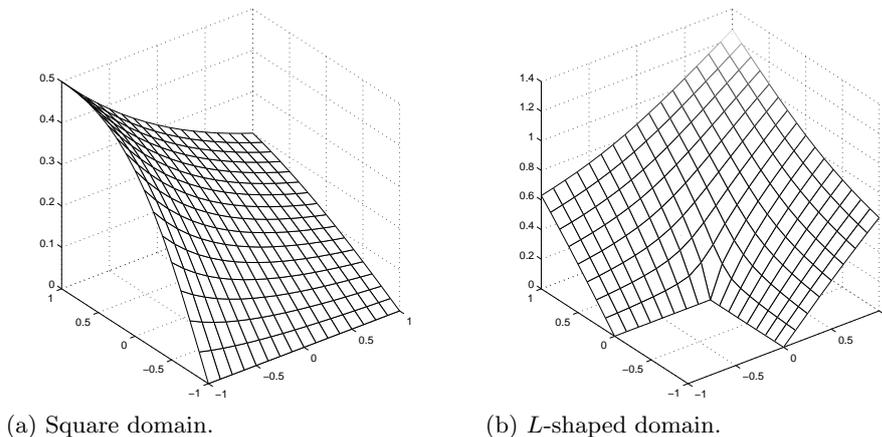


FIG. 3.1. Sample finite element approximations of analytic solutions.

Let us consider the problem of solving Laplace’s equation on the domain $[-1, 1] \times [-1, 1]$ for a smooth solution whose analytic form is known ([14], Example 1.1.3). A three-dimensional surface plot of a finite element solution computed on a 16×16 square grid is shown in Figure 3.1(a). The IFISS toolbox automatically provides *a posteriori energy estimates* of the approximation error for both bilinear Q_1 and biquadratic Q_2 approximation. These energy estimates are plotted in Figure 3.2(a) for a sequence of square grids (Q_1 : solid line, Q_2 : dashed line). It can be readily seen that the Q_1 error is halved with every successive grid refinement whereas the Q_2 error is reduced by a factor of four. This behaviour is consistent with classical energy estimates, see e.g. (1.80) and (1.81) in [14].²

Next, let us repeat the experiment for a more interesting test problem with a singular solution. This time we cut out the region $[-1, 0] \times [-1, 0]$ and solve Laplace’s

²The *a posteriori* estimates are remarkably accurate—the effectivity index (the ratio of the estimated and the exact energy errors) is 0.9992 for the 16×16 grid shown in Figure 3.1(a).

equation on the resulting L -shaped domain. The problem now has a singularity at the re-entrant corner (see Figure 3.1(b)). The errors for this problem are plotted in Figure 3.2(b).

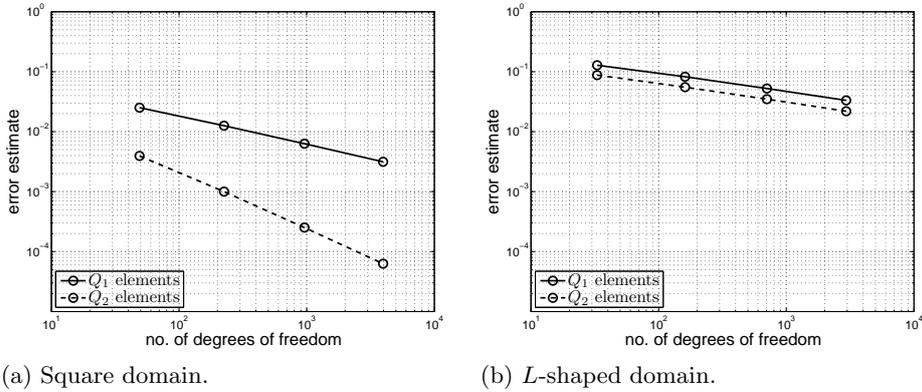


FIG. 3.2. Behaviour of the *a posteriori* error estimates.

This time, using higher-order elements does not improve the asymptotic convergence rate, and both discretization methods have *exactly the same* rate of convergence. To emphasize the importance of the smoothness in the solution, we re-solve the second problem using a stretched grid, whose successive element edges are a factor of 1.5 times longer than the adjacent edge, see Figure 3.3(a). The surface plot of the estimated error which IFISS provides, shown in Figure 3.3(b), clearly shows the difficulty encountered at the interior corner. By comparing the estimated errors, it can be seen that for Q_2 approximation, the overall accuracy obtained with 161 degrees of freedom is comparable with that obtained using a uniform grid with 2945 degrees of freedom. Of course the challenge here is to determine the optimal stretching a priori!

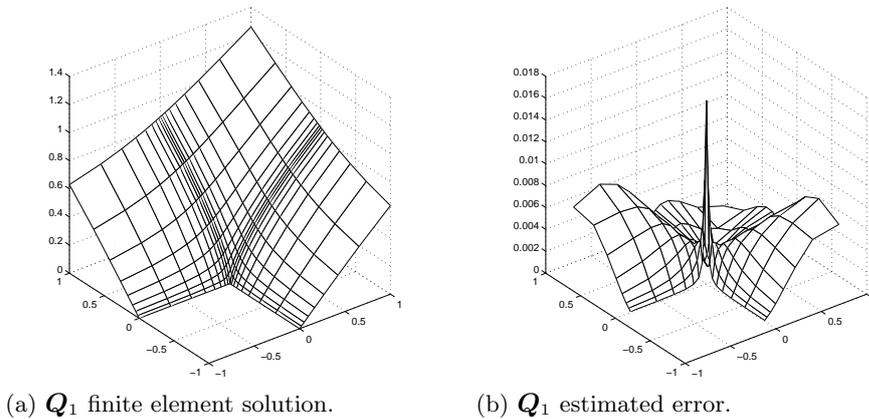


FIG. 3.3. L -shaped domain with an optimally-stretched rectangular grid.

3.2. Comparing iterative solvers for parabolic PDE problems. The IFISS toolbox is an excellent tool for studying unsteady flow problems. To show this, let us consider solving a Navier–Stokes flow problem with a view to exploring the effect of

the size of the time step on the convergence of the linear solvers used at each time level. Details of the mixed finite element methodology required for such situations can be found in [5, 14, 15, 19]. As a representative problem, we take the example of spinning up from a quiescent state to a steady lid-driven cavity flow in a square-shaped domain. The viscosity parameter ν is set to $1/200$ and the flow Reynolds number is 400. The spatial discretization is a mixed finite element method with a stable $\mathbf{Q}_2\text{-}\mathbf{Q}_1$ mixed approximation of the velocity and pressure, respectively. (See [14], Example 7.1.3 for details of the steady-state version of this problem). Snapshots of the computed flow (generated by the IFISS function `sirev_ex3.2`) are shown in Figure 3.4.

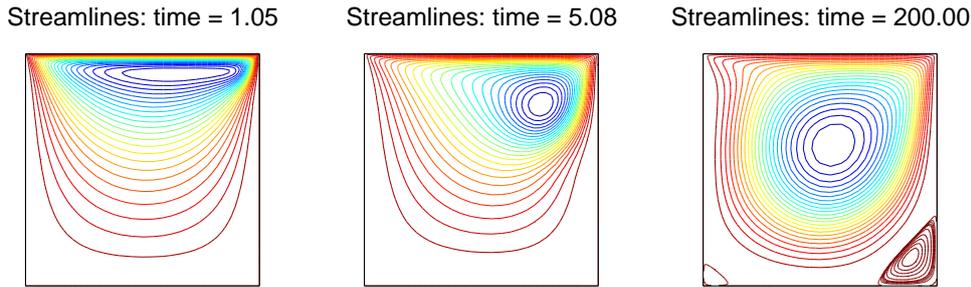


FIG. 3.4. Representative snapshots of the lid-driven cavity flow.

The self-adaptive time stepping algorithm used in IFISS involves an implicit time integrator using a stabilized trapezoidal rule with an explicit Adams–Bashforth method for error control (see [22] for details). The evolution of the time step with the accuracy error tolerance ε_t set to 10^{-4} is shown in Figure 3.5(a). At each individual time step, a linear algebra system with a distinctive saddle-point structure, see [14, section 7.3]),

$$\begin{pmatrix} F & B^T \\ B & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{p} \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{g} \end{pmatrix},$$

must be solved. We present results here using a preconditioned GMRES linear solver [32], comparing the performance of three alternative preconditioners (additional methods are available in IFISS). The choices are

- ILU Incomplete LU factorisation with zero fill-in (using the MATLAB function `ilu` with the `'nofill'` option), see e.g. [30, chapter 10].
- PCD Pressure Convection-Diffusion preconditioning described in [14, 23, 35].
- LSC Least-Squares Commutator preconditioning described in [9, 10, 14].

The first of these is a standard ‘black-box’ preconditioner, while the latter two methods are specifically designed for saddle-point systems coming from discretization of linearized Navier–Stokes equations. Both of these involve using an upper block-triangular preconditioner,

$$\mathcal{M} = \begin{pmatrix} F & B^T \\ 0 & -M_S \end{pmatrix},$$

where M_S is an approximation to the Schur complement operator $S = BF^{-1}B^T$. The two alternatives differ in the way in that this approximation is made.

Figure 3.5(b) shows the number of preconditioned GMRES iterations taken for the linear solve at each time step for the lid-driven cavity test problem with 8450

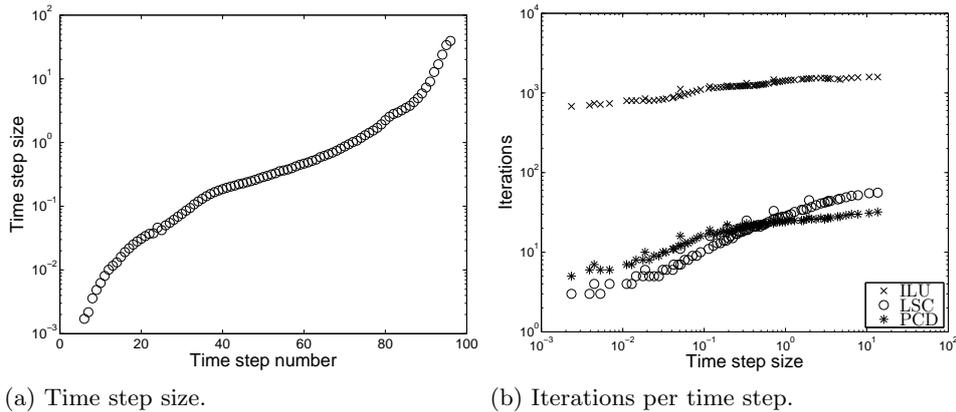


FIG. 3.5. Time stepping data for the representative lid-driven cavity flow problem.

degrees of freedom (these individual systems are generated using the IFISS function `snapshot_solveflow`). The horizontal axis shows the time step size and the vertical axis shows the number of GMRES iterations required for convergence with a tolerance of 10^{-6} at that time step. From this plot it is evident that the performance of all the preconditioned linear solvers tested deteriorates as the time step size increases—this is a known issue with unsteady simulations. The interesting thing, however, is the degree of deterioration, which is clearly different for each preconditioner. Such considerations can be important in the design of an efficient adaptive solver, suggesting the possibility of automatically switching the type of preconditioner used as the time step size changes.

3.3. Linking discretizations and solvers. One of the more unusual features which IFISS offers is easy access to the interplay between finite element discretizations and iterative solvers. Details of how this can be done are given in the *Installation and Software Guide* [13]. As an example of how the package has been used effectively in this way, we now briefly summarize the work of Silvester and Simoncini [38] on the development of an optimal solver for symmetric linear systems associated with mixed approximation of PDEs. Specifically, we consider this work in the context of solving a discrete Stokes problem, namely, the slow-flow version of the regularized lid-driven cavity problem described in the previous section (see for example [5, 14, 15, 19] for background material on Stokes problems).

The new optimal solver described in [38] has three key ingredients. Firstly, it uses a block preconditioner which gives convergence at a rate independent of the problem parameters. Secondly, it uses an adaptation of the usual MINRES algorithm of Paige and Saunders [27] which takes advantage of certain aspects of the finite element discretization in the design of the iterative solution method. The resulting EST_MINRES algorithm ([38], Fig. 10), instead of minimizing in the standard Euclidean norm of the residual vector, minimizes in a computable monotonically decreasing norm that is equivalent to the natural norm for error estimation of the discrete solution. The final new feature is the incorporation of *a posteriori* error estimation functionality so that a precise stopping criterion can be identified which balances the algebraic error with the PDE approximation error. The combination of these ideas clearly shows why considering the PDE origin of a linear system is essential in the design of efficient algebraic solvers.

To fix notation, we write the linear system arising from mixed finite element approximation of our test Stokes cavity flow problem as

$$\begin{bmatrix} A & B^T \\ B & 0 \end{bmatrix} \begin{bmatrix} \mathbf{u} \\ \mathbf{p} \end{bmatrix} = \begin{bmatrix} \mathbf{f} \\ \mathbf{g} \end{bmatrix} \quad \text{or} \quad K\mathbf{x} = \mathbf{b}, \quad (3.1)$$

with residual and algebraic error at the k th step given by $\mathbf{r}^{(k)} = \mathbf{b} - K\mathbf{x}^{(k)}$ and $\mathbf{e}^{(k)} = \mathbf{x} - \mathbf{x}^{(k)}$, respectively (where \mathbf{x} represents the exact solution to (3.1)). Now suppose system (3.1) is preconditioned with a symmetric and positive definite preconditioner, M say. That is, conceptually we think of solving $MK\mathbf{x} = M\mathbf{b}$ instead, choosing M so that MINRES convergence is faster than for (3.1). One natural choice for a Stokes problem is the *canonical* preconditioner

$$M = \begin{bmatrix} A^{-1} & 0 \\ 0 & Q^{-1} \end{bmatrix} \quad (3.2)$$

(where Q represents the pressure mass matrix) that maps the dual finite element space into the original (see [25, 29] for details). Note that the quantity which is minimized at each step by standard preconditioned MINRES iteration, $\|\mathbf{r}^{(k)}\|_M$, is measured in a *preconditioner-dependent* norm. In terms of the underlying PDE, as discretization error for the Stokes problem is measured in the L_2 norm for the gradient of the velocity and in the L_2 norm for pressure (see [14, §5.4]), the natural matrix norm here is $\|\mathbf{e}^{(k)}\|_E$ where

$$E = \begin{bmatrix} A & 0 \\ 0 & Q \end{bmatrix}.$$

Ideally, we would like to relate the quantity computed by MINRES to this natural norm which arises naturally from the PDE. For the Stokes problem, we assume that the finite element discretization used is *inf-sup stable* (see e.g. [5, (2.29)]). With this assumption, we can use the heuristic

$$\|\mathbf{e}^{(k)}\|_E \leq \frac{\sqrt{2}}{\gamma^2} \|\mathbf{r}^{(k)}\|_M \quad (3.3)$$

where γ is the so-called *inf-sup constant* (see [38]).

In theory, inequality (3.3) could be used as the basis for a stopping criterion for MINRES iteration (by stopping when the quantity on the right-hand side of (3.3) is less than some fixed tolerance). However, there are still two issues to be addressed before this can be considered useful in practice. The first of these is that, for realistic problems, an “ideal” preconditioner of the form (3.2) would usually have to be replaced by a (cheaper) spectrally equivalent preconditioner, M_* say: the example given in [38] replaces the action of applying A^{-1} by a fixed number of Algebraic Multigrid V-cycles and the action of applying Q^{-1} by a fixed number of Jacobi iterations with Chebyshev acceleration. The authors show there that if the approximation of M by M_* is sufficiently accurate, then the impact on convergence of replacing the M -norm in (3.3) by the M_* -norm will be minimal.

The second practical issue with (3.3) is the need to explicitly compute the inf-sup constant γ . It turns out, however, that an accurate estimate of γ^2 can be automatically generated at every step of EST_MINRES by exploiting the connection between the iteration and the Lanczos estimates of the eigenvalues of the preconditioned matrix.

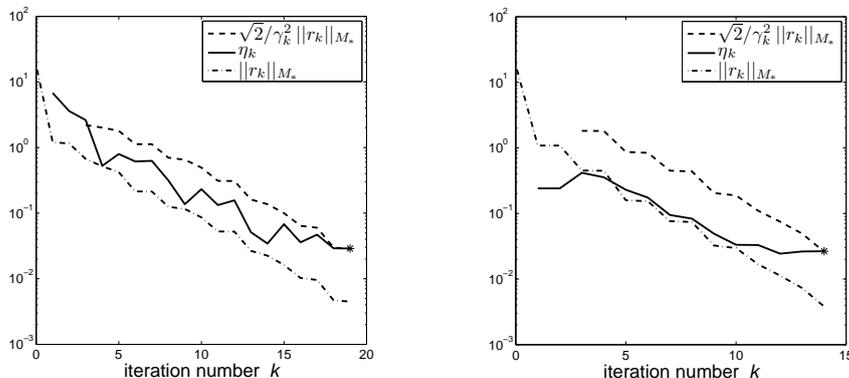
In addition, the algorithm gives a bonus in the form of a free estimate of the accuracy of M_* as an approximation to M : see [38], §4.1 for details.

The final step in producing an optimal parameter-free iterative method is to connect the actual stopping tolerance with the PDE discretization error. With established *a posteriori* error estimation technology, an estimate for the PDE discretization error at each step can be computed, and the iteration terminated when the appropriate measure of the absolute residual is less than this value. Denoting such an error estimate at EST_MINRES step k by $\eta^{(k)}$, the final practical stopping criterion suggested is

$$\frac{\sqrt{2}}{\gamma^2} \|\mathbf{r}^{(k)}\|_{M_*} \leq \eta^{(k)}, \quad (3.4)$$

which, in light of (3.3), forces the algebraic error to be comparable with the estimation of the discretization error.

Figure 3.6 shows some convergence plots for EST_MINRES applied to the regularized lid-driven cavity test problem with $\mathbf{Q}_2\text{-}\mathbf{P}_1$ approximation on a 64×64 stretched grid of square elements. In each case, the evolution of three quantities is shown: the



(a) One AMG V-cycle.

(b) Four AMG V-cycles.

FIG. 3.6. Convergence curves.

upper bound estimate for $\|\mathbf{e}^{(k)}\|_E$ from (3.3) (dashed line), the approximation error estimate $\eta^{(k)}$ (solid line) and the preconditioned residual error (dash-dot line). The (automatically computed) termination point associated with the stopping test (3.4) is marked with $*$. It can be seen from Figure 3.6(a) that with one AMG V-cycle in the preconditioner, the iteration converges in 19 iterations, whereas with 4 AMG V-cycles (Figure 3.6(b)), only 14 iterations are required. Other aspects of the performance of this optimal solver can be seen using the IFISS routine `stokesflowdemo`.

3.4. Applications and extensions. One immediate application of IFISS is as a convenient source of sample matrices for testing new ideas. This is not just limited to the development of algorithms specifically for problems in incompressible fluid flow [1, 2]. IFISS has also been used as a source of test problems in a much more general way, for example, in assessing the performance of new Krylov subspace solvers and preconditioners for a wider class of saddle-point problems [42, 45, 46], and for specific applications in optimization [7], signal processing [16] and eigenvalue computations [20, 28, 43]. Details of how the code can be used in this way are given in the *Installation and Software Guide* [13]. Since its creation in 1997, see [8], the IFISS package has

evolved to include many additional features such as new equation models, discretization methods, preconditioning methods and solver algorithms. Because IFISS has been deliberately designed in a modular way, it is relatively easy for the user to incorporate new developments and additional code, or to extract parts of IFISS to be used in combination with other software. Examples where IFISS has been used successfully in this way include the enhanced package PIFISS [37] for solving potential flow problems, stochastic finite element software used to study uncertainty quantification [34, 44] and a coupled flow-orientation solver in liquid crystal modelling [39].

In this section we highlight one particular application [11] where IFISS has been used to explore a fundamental question in mathematical modelling concerning the stability of dynamical systems. When they arise from discretization of parameter-dependent PDEs, these systems have the form

$$Mu_t = f(u, \nu), \quad (3.5)$$

where $f : \mathbb{R}^n \times \mathbb{R} \rightarrow \mathbb{R}^n$ is a nonlinear mapping, $u = u^{(\nu)}(t) \in \mathbb{R}^n$ is a discrete state variable that depends on a scalar parameter ν , and (for finite elements) M is a mass matrix. The question of interest is: given a steady solution $u^{(\nu)}$, i.e., $u_t^{(\nu)} = 0$, and a small perturbation $\delta(0)$ at time $t = 0$, does $\delta(t)$ decay with time, in which case the steady solution is *stable*, or does it grow? A necessary condition for stability, in the study of *linear* stability analysis, is that all eigenvalues of the generalized eigenvalue problem

$$J\mathbf{x} = \mu M\mathbf{x} \quad (3.6)$$

have negative real parts, where $J = \frac{\partial f}{\partial u}(u, \nu)$ is the Jacobian matrix of f evaluated at the specific parameter value ν .

In cases of *Hopf bifurcation* [18] the eigenvalues of interest consist of a complex conjugate pair that, as ν varies from a stable regime to a critical value ν^* at which stability is lost, trace out paths moving from left to right in the complex plane until they cross the imaginary axis. The identification of ν^* is a difficult computational task because of complexities associated with the eigenvalue problem (3.6). For large-scale discretizations, J will be too large to use direct eigenvalue methods such as the QR algorithm [17, 40]. The alternative choice, iterative methods such as Arnoldi-based methods or subspace iteration [31, 40], can compute only a few selected eigenvalues. Without any information about the location of the important eigenvalues, iterative methods may have difficulty finding them.

IFISS provides an easy way to generate interesting benchmark problems and to incorporate them into the study of new algorithms. In particular, it is known that two classic problems, driven cavity flow and flow over an obstacle, exhibit Hopf bifurcations. Examples of the distributions of three hundred eigenvalues closest to the origin for these two problems are shown in Figure 3.7. The rightmost eigenvalues are highlighted with “×”. Note that for the cavity problem (shown on the left), there are many poorly separated eigenvalues, and an iterative method applied to $M^{-1}J$ will have great difficulty finding them. The alternative is to find the largest inverse eigenvalue of (3.6) using $J^{-1}M$. If there are many eigenvalues smaller in modulus than the critical ones, as in the case of the cavity problem, this also causes problems for (inverse) iterative methods. In [11], Elman et al. have used IFISS to explore and develop new algorithms to address these issues.

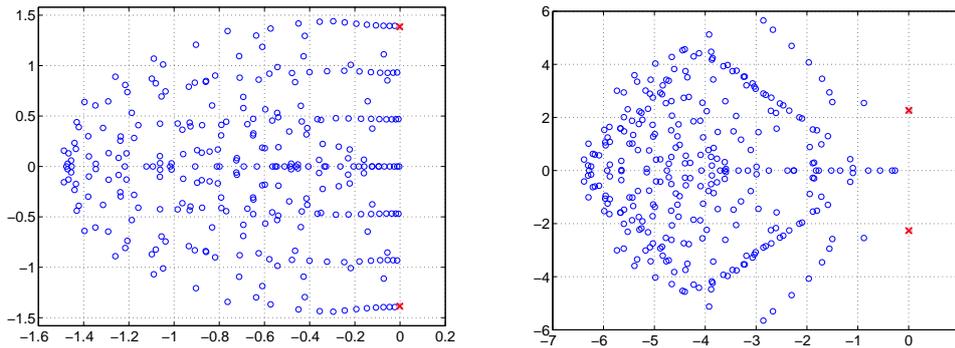


FIG. 3.7. Eigenvalues of the Jacobian system for driven cavity flow (left) and flow over an obstacle (right).

For models of incompressible flow, the Jacobian matrix has the structure

$$J(\nu) = \begin{pmatrix} \nu L + N(\nu) & B^T \\ B & 0 \end{pmatrix},$$

where ν is the kinematic viscosity. If we have an estimate ν_0 for the critical value ν^* , then

$$J(\nu^*) \approx \begin{pmatrix} \nu_0 L + N(\nu_0) & B^T \\ B & 0 \end{pmatrix} + (\nu^* - \nu_0) \begin{pmatrix} L + N'(\nu_0) & 0 \\ 0 & 0 \end{pmatrix} \equiv A + \lambda C,$$

where L is a discrete Laplacian and $N(\nu)$ is a parameter-dependent matrix that represents the effect of nonlinearity in the Navier-Stokes equations, including convection.³ When the Jacobian is of this form, (3.6) is a parameterized eigenvalue problem $(A + \lambda C)\mathbf{x} = \mu M\mathbf{x}$. It is shown in [26] that the solution to this problem can be obtained by solving a related eigenvalue problem having a special *Lyapunov structure*

$$AZM^T + MZA^T + \lambda(CZM^T + MZC^T) = 0. \quad (3.7)$$

The quantity sought is the smallest real parameter λ for which there is a solution Z in (3.7). (The matrix solution Z , which is also referred to as an eigenvector, is of rank 2 and has the form $Z = VDV^T$; see [11, 26] for details.) This has the desirable feature of requiring a small real eigenvalue, and a variant of inverse iteration can be used to find it. Once λ is obtained, the critical parameter is estimated as $\nu^* = \nu_0 + \lambda$. It is also shown in [26] that the rightmost eigenvalues of the parameterized eigenvalue problem can also be found easily by solving the 2×2 problem

$$V^T(A + \lambda C)\mathbf{y} = \mu V^T M V \mathbf{y}.$$

In [11], this idea was explored in detail for the benchmark problems above.

In general, it can be difficult to identify the desired (rightmost) eigenvalue using a method such as implicitly restarted Arnoldi, especially when, as in Figure 3.7, there are many poorly separated eigenvalues. Using the alternative approach of trying to find the largest eigenvalue of the inverse eigenvalue problem is also difficult when the imaginary part of the eigenvalue sought is large. The method discussed above is potentially more robust than these conventional methods for identifying Hopf bifurcations. Having access to IFISS is what made this study possible.

³In [11], the matrix $C = \frac{dJ}{d\nu}(\nu_0)$ was approximated using a forward difference approximation.

4. Summary and future developments. The IFISS package is designed to be accessible—anyone with a basic knowledge of MATLAB can use it to gain understanding of the key concepts underlying the numerical modelling of PDEs. The case studies described in this review article highlight the potential of using the IFISS package to explore the numerical analysis of the PDEs underlying incompressible flow. The package is currently used in academic institutions around the world to support research and teaching. Examples of extensions to IFISS that are currently being developed include the incorporation of unstructured triangular meshes and the extension to PDE problems with random data. In the future, we also envisage extending the capabilities of IFISS to three dimensions, making use of parallel MATLAB constructs in order to take advantage of multicore architectures.

REFERENCES

- [1] M. BENZI AND X. P. GUO, *A dimensional split preconditioner for Stokes and linearized Navier-Stokes equations*, Appl. Numer. Math., 61 (2011), pp. 66–76.
- [2] M. BENZI, M. A. OLSHANSKII, AND Z. WANG, *Modified augmented Lagrangian preconditioners for the incompressible Navier-Stokes equations*, Int. J. Numer. Meth. Fl., 66 (2011), pp. 486–508.
- [3] D. BRAESS, *Finite Elements*, Cambridge University Press, London, 2007. Third Edition.
- [4] S. C. BRENNER AND L. R. SCOTT, *The Mathematical Theory of Finite Element Methods*, Springer-Verlag, New York, 2010. Third Edition.
- [5] F. BREZZI AND M. FORTIN, *Mixed and Hybrid Finite Element Methods*, Springer-Verlag, New York, 1991.
- [6] A. DEDNER, R. KLÖFKORN, M. NOLTE, AND M. OHLBERGER, *DUNE-FEM, the fem module of the distributed and unified numerics environment*, accessed 17–4–2013. <http://dune.mathematik.uni-freiburg.de>.
- [7] H. S. DOLLAR, N. I. M. GOULD, M. STOLL, AND A. J. WATHEN, *Preconditioning saddle-point systems with applications in optimization*, SIAM J. Sci. Stat. Comput., 32 (2010), pp. 249–270.
- [8] H. ELMAN, D. SILVESTER, AND A. WATHEN, *Iterative methods for problems in computational fluid dynamics*, in *Iterative Methods in Scientific Computing*, R. Chan, T. Chan, and G. Golub, eds., Springer, Singapore, 1997, pp. 271–327.
- [9] H. C. ELMAN, V. E. HOWLE, J. SHADID, R. SHUTTLEWORTH, AND R. TUMINARO, *Block preconditioners based on approximate commutators*, SIAM J. Sci. Comput., 27 (2006), pp. 1651–1668.
- [10] H. C. ELMAN, V. E. HOWLE, J. SHADID, D. J. SILVESTER, AND R. TUMINARO, *Least squares preconditioners for stabilized discretizations of the Navier-Stokes equations*, SIAM J. Sci. Comput., 30 (2007), pp. 290–311.
- [11] H. C. ELMAN, K. MEERBERGEN, A. SPENCE, AND M. WU, *Lyapunov inverse iteration for identifying Hopf bifurcations in models of incompressible flow*, SIAM J. Sci. Stat. Comput., 34 (2012), pp. A1584–A1606.
- [12] H. C. ELMAN, A. RAMAGE, AND D. J. SILVESTER, *Algorithm 886: IFISS, a Matlab toolbox for modelling incompressible flow*, ACM Trans. Math. Soft., 33 (2007).
- [13] ———, *Incompressible Flow Iterative Solution Software (IFISS) installation and software guide version 3.2*, 2012. http://www.manchester.ac.uk/ifiss/ifiss_guide.3.2.pdf.
- [14] H. C. ELMAN, D. J. SILVESTER, AND A. J. WATHEN, *Finite Elements and Fast Iterative Solvers with applications in incompressible fluid dynamics*, Oxford University Press, Oxford, UK, 2005.
- [15] V. GIRAULT AND P.-A. RAVIART, *Finite Element Methods for Navier-Stokes Equations*, Springer-Verlag, Berlin, 1986.
- [16] G. H. GOLUB, M. STOLL, AND A. J. WATHEN, *Approximation of the scattering amplitude and linear systems*, Electron. T. Numer. Ana., 31 (2008), pp. 178–203.
- [17] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, The Johns Hopkins University Press, Baltimore, third ed., 1996.
- [18] W. J. F. GOVAERTS, *Numerical Methods for Bifurcations of Dynamical Equilibria*, SIAM, Philadelphia, 2000.
- [19] M. D. GUNZBURGER, *Finite Element Methods for Viscous Incompressible Flows*, Academic Press, San Diego, 1989.

- [20] C.-H. GUO, N. J. HIGHAM, AND F. TISSEUR, *An improved arc algorithm for detecting definite Hermitian pairs*, SIAM J. Matrix Anal. Appl., 31 (2009), pp. 1131–1151.
- [21] C. JOHNSON, *Numerical Solution of Partial Differential Equations by the Finite Element Method*, Dover Books on Mathematics, 2009.
- [22] D. KAY, P. M. GRESHO, D. F. GRIFFITHS, AND D. J. SILVESTER, *Adaptive time-stepping for incompressible flow; part II: Navier–Stokes equations*, SIAM J. Sci. Stat. Comput., 32 (2010), pp. 111–128.
- [23] D. KAY, D. LOGHIN, AND A. J. WATHEN, *A preconditioner for the steady-state Navier–Stokes equations*, SIAM J. Sci. Comput., 24 (2002), pp. 237–256.
- [24] A. LOGG, K.-A. MARDAL, AND G. WELLS, eds., *Automated Solution of Differential Equations by the Finite Element Method*, vol. 84 of Lecture Notes in Computational Science and Engineering, Springer, Berlin, 2012. <http://fenicsproject.org>.
- [25] K.-A. MARDAL AND R. WINTHER, *Preconditioning discretizations of systems of partial differential equations*, Numer. Linear Algebra Appl., 18 (2010), pp. 1–40.
- [26] K. MEERBERGEN AND A. SPENCE, *Inverse iteration for purely imaginary eigenvalues with application to the detection of Hopf bifurcation in large scale problems*, SIAM J. Matrix Anal. Appl., 31 (2010), pp. 1982–1999.
- [27] C. PAIGE AND M. SAUNDERS, *Solution of sparse indefinite systems of linear equations*, SIAM J. Numer. Anal., 12 (1975), pp. 617–629.
- [28] J. ROMMES, *Arnoldi and Jacobi–Davidson methods for generalized eigenvalue problems $Ax = \lambda Bx$ with singular B* , Math. Comp., 77 (2008), pp. 995–1015.
- [29] T. RUSTEN AND R. WINTHER, *A preconditioned iterative method for saddlepoint problems*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 887–904.
- [30] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, SIAM, Philadelphia, 2003. Second Edition.
- [31] Y. SAAD, *Numerical Methods for Large Eigenvalue Problems*, SIAM, Philadelphia, revised ed., 2011.
- [32] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
- [33] A. SCHMIDT AND K. G. SIEBERT, *ALBERTA, an adaptive hierarchical finite element toolbox*, accessed 17–4–2013. <http://www.alberta-fem.de>.
- [34] D. J. SILVESTER, A. BESPALOV, AND C. E. POWELL, *A framework for the development of implicit solvers for incompressible flow problems*, Discrete and Continuous Dynamical Systems — Series S (DCDS–S), 5 (2012), pp. 1195–1221.
- [35] D. J. SILVESTER, H. C. ELMAN, D. KAY, AND A. J. WATHEN, *Efficient preconditioning of the linearized Navier–Stokes equations for incompressible flow*, J. Comp. Appl. Math., 128 (2001), pp. 261–279.
- [36] D. J. SILVESTER, H. C. ELMAN, AND A. RAMAGE, *Incompressible Flow and Iterative Solver Software (IFISS) version 3.2*, May 2012. <http://www.manchester.ac.uk/ifiss/>.
- [37] D. J. SILVESTER AND C. E. POWELL, *Potential (Incompressible) Flow and Iterative Solver Software (PIFISS)*, February 2007. <http://www.manchester.ac.uk/ifiss/pifiss.html>.
- [38] D. J. SILVESTER AND V. SIMONCINI, *An optimal iterative solver for symmetric indefinite systems stemming from mixed approximation*, ACM Trans. Math. Soft., 37 (2010).
- [39] A. M. SONNET AND A. RAMAGE, *Computational fluid dynamics for nematic liquid crystals*, tech. rep., University of Strathclyde Department of Mathematics and Statistics Research Report 28-07, 2007.
- [40] G. W. STEWART, *Matrix Algorithms. Volume II: Eigensystems*, SIAM, Philadelphia, 2001.
- [41] G. STRANG AND G. J. FIX, *An Analysis of the Finite Element Method*, Wellesley–Cambridge Press, <http://www.siam.org/books/>, 2008. Second Edition.
- [42] L.-Y. SUN AND J. LIU, *Constraint preconditioning for nonsymmetric indefinite linear systems*, Numer. Linear Algebra., 17 (2010), pp. 677–689.
- [43] D. B. SZYLD AND F. XUE, *Efficient preconditioned inner solves for inexact Rayleigh quotient iteration and their connections to the single-vector Jacobi–Davidson method*, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 993–1018.
- [44] E. ULLMANN, H. C. ELMAN, AND O. G. ERNST, *Efficient iterative solvers for stochastic Galerkin discretizations of log-transformed random diffusion problems*, SIAM J. Sci. Stat. Comput., 34 (2012), pp. 659–682.
- [45] M. B. VAN GIJZEN AND P. SONNEVELD, *Algorithm 913: An elegant IDR(s) variant that efficiently exploits biorthogonality properties*, ACM Trans. Math. Soft., 38 (2011).
- [46] J. T. ZHOU AND Q. A. NIU, *Substructure preconditioners for a class of structured linear systems of equations*, Math. Comput. Model., 52 (2010), pp. 1547–1553.