

An efficient solver for the fully-coupled solution of large-displacement fluid-structure interaction problems

IN PRINT IN: COMPUTER METHODS IN APPLIED
MECHANICS AND ENGINEERING
(JANUARY 2003)

Matthias Heil

*Department of Mathematics, University of Manchester, Oxford Road, Manchester
M13 9PL, U.K.*

Abstract

This paper is concerned with the fully-coupled ('monolithic') solution of large-displacement fluid-structure interaction problems by Newton's method. We show that block-triangular approximations of the Jacobian matrix, obtained by neglecting selected fluid-structure interaction blocks, provide good preconditioners for the solution of the linear systems with GMRES. We present an efficient approximate implementation of the preconditioners, based on a Schur complement approximation for the Navier-Stokes block and the use of multigrid approximations for the solution of the computationally most expensive operations. The performance of the the preconditioners is examined in representative steady and unsteady simulations which show that the GMRES iteration counts only display a mild dependence on the Reynolds number and the mesh size. The final part of the paper demonstrates the importance of consistent stabilisation for the accurate simulation of fluid-structure interaction problems.

Key words:

Fluid-structure interaction, fully-coupled solution, finite elements, preconditioning, multigrid, stabilisation.

1 Introduction

Large-displacement fluid-structure interaction problems arise in many applications in which thin-walled shell structures interact with fluid flows. Examples from engineering include flutter analyses in aeroelasticity [1], the inflation of parachutes [2], the deformation of structures by explosions [3] and the design of sails [4]. Numerous applications exist in physiological fluid mechanics and include the flow of blood in the veins and arteries and the flow of air in the pulmonary airways [5], the dynamics of heart valves [6] and many others.

The computational modelling of large-displacement fluid-structure interaction problems typically requires the solution of the Navier-Stokes equations in a moving domain, coupled to the equations of large-displacement shell theory. The coupling between the two domains often leads to highly nonlinear behaviour and the development of robust and efficient solution techniques for such problems presents one of the great challenges in computational mechanics; see, e.g., the recent special volume of this journal [7]. Most computational methods can be classified as partitioned or fully-coupled ('monolithic') schemes. Partitioned methods utilise separate solvers for the fluid and solid domains and attempt to obtain a coupled solution via a fixed-point iteration. Starting with an initial guess for the wall shape, the fluid equations are solved and the traction that the fluid exerts on the wall is evaluated. This new traction is then used to update the wall shape and the iteration is continued until convergence. This approach is relatively easy to implement and is convenient in situations in which specialised solvers for the fluid and solid problems are readily available; indeed, if compatibility with existing commercial 'black box' fluid and solid solvers is required, this method presents the only feasible approach. However, partitioned methods also have serious drawbacks. Fixed-point iterations tend to converge slowly (though Aitken extrapolation can lead to a certain improvement; see, e.g., [8] and, more recently, [9]) and, in the presence of strong fluid-solid coupling, the iteration can diverge even if a good initial guess for the solution is available. When applied to time-dependent problems, partitioned solvers lead to staggered time-integration schemes, [1], [10]. The implementation of such methods as 'weakly coupled' schemes is straightforward but it is often difficult to ensure their overall stability and temporal accuracy. 'Strong coupling', which requires costly sub-iterations at every time-step, can improve the temporal accuracy, but, of course, only if the sub-iterations converge.

In a fully coupled ('monolithic') approach, the fluid and solid equations are

* Tel. +44 (0)161 275 5808

Fax. +44 (0)161 275 5819

Email address: M.Heil@maths.man.ac.uk (Matthias Heil).

URL: <http://www.maths.man.ac.uk/~mheil/> (Matthias Heil).

discretised and solved simultaneously. For highly transient fluid-structure interaction problems which involve compressible fluids (e.g. the modelling of explosions, [3]), the use of fully-coupled but explicit time-stepping schemes can be efficient. However, if the problem is steady and/or the fluid is incompressible then the solution of a large system of coupled nonlinear algebraic equations is required. Since good initial guesses for the solutions tend to be readily available (either from the solution at a previous time-step or from the solution at a previous value of the control parameter during steady-state parameter studies), the solution of the nonlinear system by Newton’s method is attractive since it provides a robust and rapidly converging scheme. Over the past years we have successfully used this approach in a wide variety of large-displacement fluid-structure interaction problems ([11], [12], [13], [14], [15]). The main computational cost of this method arises from the repeated assembly of the Jacobian matrix and the solutions of the associated linear systems for the Newton corrections.

In this paper, we develop an efficient preconditioning technique that allows the rapid iterative solution of the linear systems that arise in the Newton method. The preconditioners are based on block-triangular approximations of the Jacobian matrix and are obtained by neglecting selected fluid-structure interaction blocks. The structure of the paper is as follows. Section 2 introduces the model problem that we used to assess the efficiency of the numerical schemes. In section 3 we present the coupled finite-element discretisation of the governing equations and discuss the stabilisation of the fluid equations on coarse meshes. After presenting some representative simulations in section 4, we show in section 5 that the direct use of the block-triangular approximate Jacobians in the Newton method causes a severe deterioration of its convergence rate. We then demonstrate that the approximate Jacobians provide good preconditioners for the solution of the linear systems with GMRES. Next, we reduce the cost of the preconditioning operation by approximating the Navier-Stokes block with a pressure Schur-complement preconditioner and present an efficient implementation based on multigrid approximations for the computationally most expensive operations. Finally, we discuss the adaptive choice of the convergence tolerance for the GMRES iteration and demonstrate the importance of consistent stabilisation in the numerical simulation of fluid-structure interaction problems.

2 The model problem

We will assess the efficiency of the numerical schemes in the model problem sketched in Fig. 1. Viscous fluid is driven through a two-dimensional channel in which part of one wall is replaced by a pre-stressed elastic membrane. The membrane is loaded by the external pressure p_{ext}^* and by the fluid traction. This

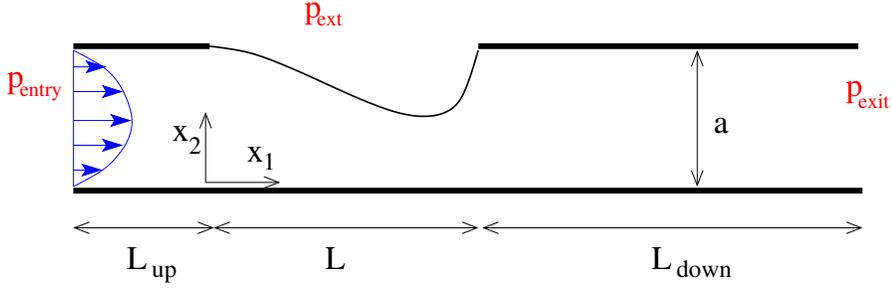


Fig. 1. Sketch of the model problem considered in this study. Viscous fluid is driven through a channel in which part of one wall is replaced by a pre-stressed elastic membrane.

system has been studied by many authors (see e.g., [5], for a recent review) and has features that are representative of many other large-displacement fluid-structure interaction problems. Specifically, the system has highly nonlinear flow characteristics and a propensity to develop large-displacement self-excited oscillations; see section 4.

We model the flow by the unsteady Navier-Stokes equations

$$Re \left(\frac{\partial u_i}{\partial t} + u_j \frac{\partial u_i}{\partial x_j} \right) = -\frac{\partial p}{\partial x_i} + \frac{\partial^2 u_i}{\partial x_j^2} \quad \text{and} \quad \frac{\partial u_j}{\partial x_j} = 0, \quad (1a,b)$$

where we have scaled all lengths on the channel width a , the velocities on the average inflow velocity U and the pressure on the viscous scale, $\mu U/a$. Throughout this paper, the summation convention is used and, unless stated otherwise, all subscripts range from 1 to 2. The Reynolds number is $Re = \rho a U/\mu$. We impose steady, fully-developed inflow at the upstream end of the channel, $\mathbf{u} = (6x_2(1-x_2), 0)$ at $x_1 = -L_{up}/a$ and assume parallel, axially traction-free outflow at the downstream end $x_2 = (L + L_{down})/a$. This is equivalent to setting $p_{exit} = 0$. No-slip conditions apply on all channel walls.

We model the elastic section of the wall as a pre-stressed, thin-walled elastic beam whose midplane we parametrise by the non-dimensional Lagrangian coordinate $\zeta \in [0, L/a]$. In its undeformed configuration, the wall is located at $\mathbf{r}_w = (\zeta, 1)^T$ and the displacement field $\mathbf{v}(\zeta, t)$ displaces its material particles to their new position $\mathbf{R}_w = \mathbf{r}_w + \mathbf{v}$. In the undeformed position, the elastic wall is subject to an initial stress $\sigma_0^* = T_0/h_0$ (where h_0 is the wall thickness and T_0 the initial longitudinal tension) which we assume to be much larger than the additional stress generated by the wall's deformation. This allows us to assume incrementally linear elastic behaviour which implies that $\sigma = \sigma_0 + \gamma$, where $\sigma = \sigma^*/E$ is the dimensionless second Piola-Kirchhoff stress and $\sigma_0 = \sigma_0^*/E$ the dimensionless pre-stress. E is the incremental Young's modulus and γ is the geometrically nonlinear extensional strain, $\gamma = v_{,\zeta}^1 + \frac{1}{2} \left((v_{,\zeta}^1)^2 + (v_{,\zeta}^2)^2 \right)$, where a comma denotes a derivative. The wall deformation is governed by the

principle of virtual displacements, [16],

$$\int_0^{L/a} \left[(\sigma_0 + \gamma) \delta\gamma + \frac{1}{12} \left(\frac{h_0}{a} \right)^2 \kappa \delta\kappa - \left(\frac{a}{h_0} \right) \mathbf{f} \cdot \delta \mathbf{R}_w \Delta \right] d\zeta = 0, \quad (2)$$

where $\mathbf{f} = \mathbf{f}^*/E$ is the non-dimensional traction acting on the wall, $\Delta = \sqrt{(1 + v_{,\zeta}^1)^2 + (v_{,\zeta}^2)^2}$ and

$$\kappa = (v_{,\zeta\zeta}^2 (1 + v_{,\zeta}^1) - v_{,\zeta\zeta}^1 v_{,\zeta}^2) / \Delta, \quad (3)$$

is the wall curvature. The first two terms in (2) represent the variation in the wall's strain energy due to its extension and bending, respectively. The last term represents the virtual work done by the load \mathbf{f} acting on the deformed wall. We assume that the elastic membrane is pinned at both ends and impose $\mathbf{v} = \mathbf{0}$ at $\zeta = 0, L/a$.

Fluid and solid interact via the no-slip condition

$$\mathbf{u} = \frac{\partial \mathbf{R}_w}{\partial t} \quad \text{on the wall}, \quad (4)$$

and via the (combined) traction on the wall

$$f_i = -p_{ext} N_i + Q \left(p \delta_{ij} - \frac{\partial u_i}{\partial x_j} - \frac{\partial u_j}{\partial x_i} \right) N_j \quad (5)$$

where $\mathbf{N} = (-v_{,\zeta}^2, 1 + v_{,\zeta}^1)^T / \Delta$ is the outer normal on the wall. The non-dimensional parameter

$$Q = \frac{\mu U}{aE} \quad (6)$$

represents the ratio of the viscous stresses to the elastic modulus of the wall and provides a measure of the strength of fluid-structure interaction; if $Q = 0$, the wall shape only depends on the transmural pressure difference.

3 Discretisation

3.1 Solid Equations

Carrying out the variations with respect to the displacements v^i and their derivatives transforms equation (2) into a variational equation of the form

$$\int_0^{L/a} \left(\phi_i^{(0)} \delta v^i + \phi_i^{(1)} \delta v_{,\zeta}^i + \phi_i^{(2)} \delta v_{,\zeta\zeta}^i \right) d\zeta = 0. \quad (7)$$

We discretise this equation by displacement-based finite elements. The ϕ -terms contain up to second derivatives of the displacements, therefore we need shape functions with continuous first derivatives across the element boundaries. Isoparametric Hermite elements with nodal displacements and slopes as independent degrees of freedom [17] were chosen such that the displacements v^i were interpolated as

$$v^i = \sum_j V^{ij} \psi_j^{(S)}, \quad (8)$$

where the $\psi_j^{(S)}$ are piecewise cubic Hermite polynomials; see [18] for further details of the implementation. This transforms (7) into

$$\left\{ \int_0^{L/a} \left(\phi_i^{(0)} \psi_j^{(S)} + \phi_i^{(1)} \psi_{j,\zeta}^{(S)} + \phi_i^{(2)} \psi_{j,\zeta\zeta}^{(S)} \right) d\zeta \right\} \delta V^{ij} = 0. \quad (9)$$

The variations of those V^{ij} that are not determined by the boundary conditions are arbitrary and the expressions multiplied by the corresponding δV^{ij} have to vanish. This provides a system of non-linear algebraic equations for the unknown V^{ij} . These equations still contain the load terms \mathbf{f} , which depend on the fluid pressure and the velocity gradients at the wall; see (5). The strong nonlinearity of the solid equations can result in complicated load-displacement paths. In particular, multiple equilibrium configurations can co-exist at a given external pressure, p_{ext} . Consequently, continuation techniques, such as Keller’s arclength method [19] are often required to compute strongly deformed equilibrium configurations. In the present problem, we implemented a simple displacement-control technique by adding the equation

$$f_{ctrl} = v^2(\zeta_{ctrl}) - V_{ctrl} = 0 \quad (10)$$

to the equations (9). Upon choosing a control point (characterised by the value of its Lagrangian coordinate ζ_{ctrl}) on the elastic membrane, this equation prescribes the vertical wall displacement at that point and allows us to treat the external pressure p_{ext} as an unknown in steady simulations.

3.2 Fluid Equations

3.2.1 Galerkin discretisation

We discretise the Arbitrary Lagrangian-Eulerian (ALE) form of the Navier-Stokes equations (1a,b) on a body-fitted moving mesh and perform the mesh update in response to changes in the wall position by a modified version of Kistler & Scriven’s ‘Method of Spines’ [20]. Briefly, we associate each nodal point j in the fluid mesh with a fixed material point on the wall which we identify by its Lagrangian coordinate $\zeta_j^{(ref)}$. As the wall deforms, each fluid node remains located at a fixed, pre-determined fraction $\omega_j \in [0, 1]$ along the

straight line from the spine origin $\mathbf{R}_0(\zeta) = (\zeta, 0)^T$ to $\mathbf{R}_w(\zeta)$ such that the position vector to fluid node j is given by

$$\mathbf{R}_j = \mathbf{R}_0(\zeta_j^{(ref)}) + \omega_j \left(\mathbf{R}_w(\zeta_j^{(ref)}) - \mathbf{R}_0(\zeta_j^{(ref)}) \right). \quad (11)$$

Fig. 2 illustrates the resulting meshes. Using standard isoparametric Taylor-Hood elements [21], the velocities, the global coordinates and the pressure are represented by

$$\mathbf{u}_i = \sum_j U^{ij} \psi_j^{(F)}, \quad x_i = \sum_j X^{ij} \psi_j^{(F)} \quad \text{and} \quad p = \sum_j P^j \psi_j^{(P)}, \quad (12)$$

where the $\psi_j^{(F)}$ and $\psi_j^{(P)}$ are bi-quadratic and bi-linear shape functions in the local element coordinates, respectively. X^{ij} are the nodal coordinates, specified by (11). We characterise the fluid meshes by h_{FEM} , the length of the element edges in the undeformed configuration.

The residuals of the momentum equations (1a) are weighted by the velocity shape functions $\psi_l^{(F)}$. We perform the usual integration by parts and obtain

$$f_{il}^{(F)} = \iint_{\Omega_f} \left[Re \left(\frac{\partial u_i}{\partial t} + (u_j - u_j^{(M)}) \frac{\partial u_i}{\partial x_j} \right) \psi_l^{(F)} - p \frac{\partial \psi_l^{(F)}}{\partial x_i} + \frac{\partial u_i}{\partial x_j} \frac{\partial \psi_l^{(F)}}{\partial x_j} \right] dv = 0, \quad (13)$$

where $\iint_{\Omega_f} dv$ is the integral over the fluid domain Ω_f and

$$u_j^{(M)} = \sum_j \frac{\partial X^{ij}}{\partial t} \psi_j^{(F)} \quad (14)$$

is the mesh velocity. All time-derivatives were evaluated by a second-order backward Euler scheme (BDF2) with constant time-step Δt . The fully implicit treatment of the momentum equations leads to a consistent mass representation. Similarly, the continuity equation (1b) is weighted by the bilinear pressure shape functions $\psi_l^{(P)}$ which yields

$$f_l^{(P)} = \iint_{\Omega_f} \frac{\partial u_j}{\partial x_j} \psi_l^{(P)} dv = 0. \quad (15)$$

3.2.2 Stabilisation by Petrov-Galerkin weighting

If the mesh Reynolds number Re_M satisfies

$$Re_M = \frac{1}{2} Re \bar{h} |\mathbf{u} - \mathbf{u}^{(M)}| < 1 \quad (16)$$

in all elements (where \bar{h} is a characteristic length of the element), the Galerkin discretisation is known to yield smooth solutions for the velocity and pressure

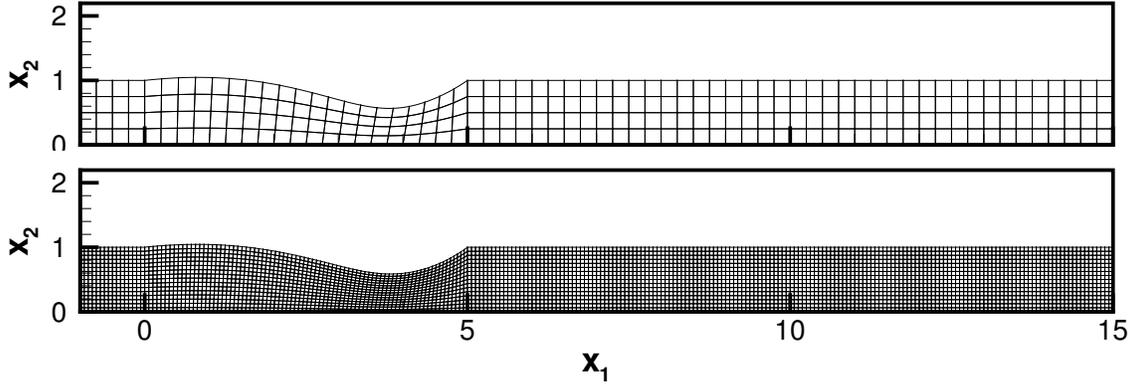


Fig. 2. Illustration of the coarsest ($h_{\text{FEM}} = 1/4$) and finest ($h_{\text{FEM}} = 1/16$) fluid meshes used in the computations.

fields; see, e.g., [22]. For the cases considered in this study, it is possible to use sufficiently fine meshes to avoid the occurrence of the spurious ‘wiggles’ that develop in the solution if the mesh Reynolds number becomes too large. However, the use of multigrid methods for preconditioning purposes (introduced in section 5.2.2 below) requires the solution of the fluid equations on coarse meshes on which the criterion (16) is severely violated. The development of ‘wiggles’ can be suppressed with a Petrov-Galerkin formulation in which the weighting functions $\psi_l^{(F)}$ for the momentum equations are replaced by

$$\psi_l^{(F,stab)} = \psi_l^{(F)} + \tau_{stab} (u_k - u_k^{(M)}) \frac{\partial \psi_l^{(F)}}{\partial x_k}. \quad (17)$$

τ_{stab} is a stabilisation parameter, which should be (just) large enough to suppress the ‘wiggles’. Following the ideas in [23], [24] and [25] we set

$$\tau_{stab} = \begin{cases} \bar{h} \left\{ 2 |\mathbf{u} - \mathbf{u}^{(M)}| \left(1 + \frac{1}{Re_M} + \frac{\bar{h}}{\Delta t |\mathbf{u} - \mathbf{u}^{(M)}|} \right) \right\}^{-1} & \text{if } Re_M > 1, \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

At large Re_M , this choice is consistent with the ‘optimal’ stabilisation for steady advection diffusion problems (see, e.g., [26] and [27]) and it leads to small values of τ_{stab} in problems which are dominated by the unsteady terms since $\tau_{stab} \sim \Delta t$ as $\Delta t \rightarrow 0$.

Consistent use of the weighting functions $\psi_i^{(F,stab)}$ adds the terms

$$f_{il}^{(F,stab)} = \iint_{\Omega_f} \tau_{stab} \left[\underbrace{Re \left(\frac{\partial u_i}{\partial t} \right)}_I + \underbrace{(u_j - u_j^{(M)}) \frac{\partial u_i}{\partial x_j}}_{II} + \underbrace{\frac{\partial p}{\partial x_i}}_{III} - \underbrace{\frac{\partial^2 u_i}{\partial x_j^2}}_{IV} \right] (u_k - u_k^{(M)}) \frac{\partial \psi_l^{(F)}}{\partial x_k} dv \quad (19)$$

to the weak form (13) of the momentum equations. The second-order term IV is typically evaluated element-by-element. For the purposes of stabilisation

(i.e. the suppression of ‘wiggles’) it is sufficient to include only the advective term II into $f_{il}^{(F,stab)}$. This is equivalent to the classical SUPG stabilisation of Hughes & Brooks [28]. However, if the terms I, III and IV are neglected in (19), the weak form of the equations becomes inconsistent in the sense that an exact solution of (1a,b) no longer satisfies the weak equations. The choice (18) for τ_{stab} ensures that the stabilisation terms disappear as \bar{h} and $\Delta t \rightarrow 0$. However, as long as $\tau_{stab} > 0$ anywhere in the domain, the inconsistency introduces additional errors into the solution. We will demonstrate in section 5.3 that, in fluid-structure interaction problems, these errors can be highly significant.

3.3 The coupled solution

We solve the coupled system of non-linear algebraic equations (9), (10), (13) and (15) by Newton’s method. This requires the repeated assembly and solution of linear systems $\mathcal{J}\mathbf{x} = -\mathbf{r}$, with the following block structure

$$\left[\begin{array}{c|cc} \mathbf{S} & \mathbf{C}_{su} & \mathbf{C}_{sp} \\ \hline \mathbf{C}_{us} & \mathbf{F} & \mathbf{G} \\ \mathbf{C}_{ps} & \mathbf{D} & \mathbf{0} \end{array} \right] \begin{pmatrix} \mathbf{x}_s \\ \mathbf{x}_u \\ \mathbf{x}_p \end{pmatrix} = - \begin{pmatrix} \mathbf{r}_s \\ \mathbf{r}_u \\ \mathbf{r}_p \end{pmatrix}. \quad (20)$$

Here \mathbf{r}_s , \mathbf{r}_u and \mathbf{r}_p represent the current residuals of the solid, momentum and continuity equations, respectively and \mathbf{x}_s , \mathbf{x}_u and \mathbf{x}_p are the associated Newton corrections. The diagonal block \mathbf{S} in the Jacobian matrix has two contributions

$$\mathbf{S} = \mathbf{S}^{(solid)} + \mathbf{S}^{(FSI)}, \quad (21)$$

where $\mathbf{S}^{(solid)}$ is the classical tangent stiffness matrix of the wall equations. The additional contribution, $\mathbf{S}^{(FSI)}$, arises from the fluid-structure interaction: A change in a solid degree-of-freedom affects the velocity gradients in (5) because the nodal positions in the fluid mesh are determined by the wall displacement field. The diagonal block

$$\mathbf{F} = Re \left(\frac{1}{\Delta t} \mathbf{M} + \mathbf{N} \right) + \mathbf{A} \quad (22)$$

represents the fluid momentum matrix. Its three components \mathbf{M} , \mathbf{N} and \mathbf{A} arise from the discretisation of the time-derivative, the ALE advection operator and the viscous term in the Navier-Stokes equations, respectively. The off-diagonal blocks \mathbf{G} and \mathbf{D} are the discrete gradient and divergence operators and without stabilisation we have $\mathbf{G} = \mathbf{D}^T$. The term $1/\Delta t$ in the above equation indicates the dependence on the time-step as $\Delta t \rightarrow 0$. The mass

matrix \mathbf{M} has two contributions

$$\mathbf{M} = \widehat{\mathbf{M}} + \mathbf{M}^{(stab)}, \quad (23)$$

where $\widehat{\mathbf{M}}$ represents the standard mass matrix corresponding to (13) and $\mathbf{M}^{(stab)}$ arises from the differentiation of the stabilisation term I in (19). Thus, consistent stabilisation makes \mathbf{M} asymmetric. However, $\mathbf{M}^{(stab)} \sim \Delta t$ as $\Delta t \rightarrow 0$, so for small Δt , the mass matrix is dominated by the symmetric positive definite contribution $\widehat{\mathbf{M}}$.

The fluid and solid variables are coupled through the interaction blocks \mathbf{C}_{**} . Their physical origin is as follows: \mathbf{C}_{sp} and \mathbf{C}_{su} represent the respective effects of the fluid pressure and the viscous fluid stresses on the wall equations and arise through the load terms (5). The interaction blocks

$$\mathbf{C}_{us} = \mathbf{C}_{us}^{(s)} + \frac{1}{\Delta t} \mathbf{C}_{us}^{(t)} \quad \text{and} \quad \mathbf{C}_{ps} \quad (24)$$

represent the effect of the wall displacement field on the discretised momentum and continuity equations, respectively. In steady problems, this interaction occurs via the algebraic mesh-update (11): A change in the wall displacement field changes the fluid mesh and thus the Jacobian of the isoparametric mapping in the fluid elements. The term $\frac{1}{\Delta t} \mathbf{C}_{us}^{(t)}$ arises from the fact that the time-derivative of the wall displacement field affects the momentum equations by setting the fluid velocities along the wall via (4) and by setting the mesh velocities $\mathbf{u}^{(M)}$ via (14). For small time-steps, the Jacobian matrix therefore has large off-diagonal entries.

3.4 Implementation

The repeated assembly of the Jacobian matrix involves a significant amount of computational work and needs to be optimised carefully. We assembled the blocks \mathbf{F} , \mathbf{D} and \mathbf{G} by the standard finite element assembly procedure, using the analytical expressions for the integrand and three point Gauss rules for the numerical integration. The remaining entries in the Jacobian matrix were generated by finite differencing equations (9), (10), (13) and (15). The number of elements used for the discretisation of the wall equations (9) was matched to the number of elements in the fluid domain such that their respective nodal points coincided. This resulted in a finer-than-necessary discretisation of the wall equations but ensured that every solid degree of freedom only affected the position of a small number of fluid nodes. The (modest) increase in the total number of degrees of freedom was more than compensated for by a dramatic increase in the efficiency of the sparse mesh-update-algorithm which is used frequently during the finite difference operations. The stabilisation parameter τ_{stab} was computed element-wise by evaluating (18) at the element centroids

and its value was kept constant during the Newton iteration. \bar{h} was determined by the procedure proposed in reference [27] to ensure that it represents the element length in the direction of the ‘wind’ $\mathbf{u} - \mathbf{u}^{(M)}$.

4 Typical results

Figs. 3-5 present results from representative steady and unsteady simulations which will be used as test cases for the numerical schemes presented below. All simulations were performed with the parameters $L_{up}/a = 1$, $L/a = 5$, $L_{down}/a = 10$, $Q = 10^{-2}$, $\sigma_0 = 10^3$, $p_{exit} = 0$, $h_0/a = 10^{-2}$ and $\zeta_{ctrl} = 3.5$.

4.1 Steady flows: Multiple solutions

Fig. 3 illustrates the changes in the flow-field and the wall deformation when the external pressure is varied while the flow rate and the fluid pressure at the downstream end of the rigid segment are held constant. The viscous pressure drop in the downstream rigid segment causes the fluid pressure underneath the membrane to exceed the fluid pressure $p_{exit} = 0$ at the outflow. At $Re = 500$, an external pressure of $p_{ext} = p_{ext}^{(flush)} = 1.524$ is required to maintain the membrane in an approximately flush state (see Fig. 3a). An increase in external pressure increases the compressive load on the wall and causes it to collapse more strongly near its downstream end. At finite Reynolds number, the Bernoulli effect leads to a local reduction in fluid pressure which augments the wall compression in the region of strongest collapse. The reduction in channel width also increases the overall viscous flow resistance and requires an increase in upstream pressure to maintain the same flow rate. The resulting increase in fluid pressure in the upstream half of the elastic segment causes it to bulge out; see Fig. 3b. The flow separates downstream of the point of strongest collapse and a long region of reversed flow develops.

Fig. 4 illustrates the pressure/deformation characteristics of the system at three different Reynolds number by plotting $V_{ctrl} = v^2(\zeta = 3.5)$, the vertical wall displacement at 70% of the membrane’s length, as a function of the external pressure. At zero Reynolds number, an increase in external pressure leads to an approximately proportional increase V_{ctrl} . At finite Reynolds number, the additional local pressure drop generated by the Bernoulli effect increases the collapse at a given value of p_{ext} . At $Re = 500$, this destabilising effect is so strong that limit points develop in the $V_{ctrl}(p_{ext})$ curve. Hence, small changes in p_{ext} can lead to large changes in the (steady) wall shapes. For instance, a small increase in the external pressure beyond $p_{ext}^{(flush)}$ forces the system to jump (dynamically) from its ‘flush’ state into a strongly collapsed equilibrium

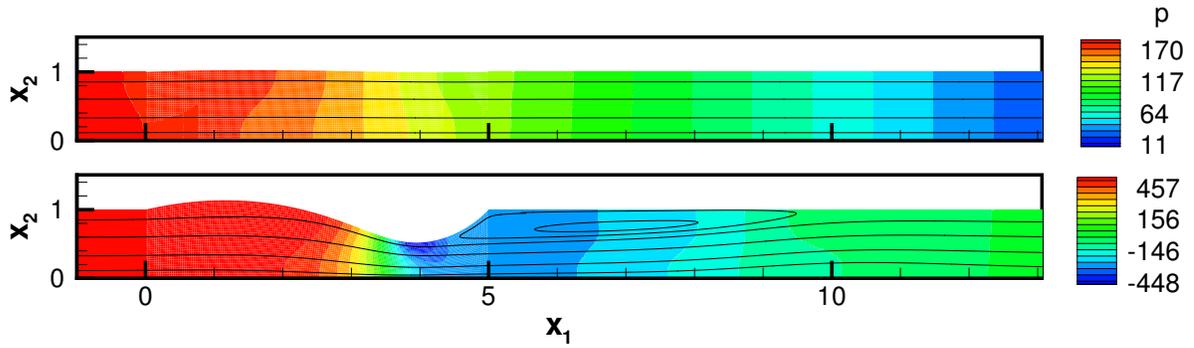


Fig. 3. Streamlines and pressure contours for steady flow at $Re = 500$. $p_{ext} = p_{ext}^*/E = 1.524$ (top) and $p_{ext} = 3.247$ (bottom). The fluid pressure is scaled on the viscous scale, $p = p^*a/(\mu U)$.

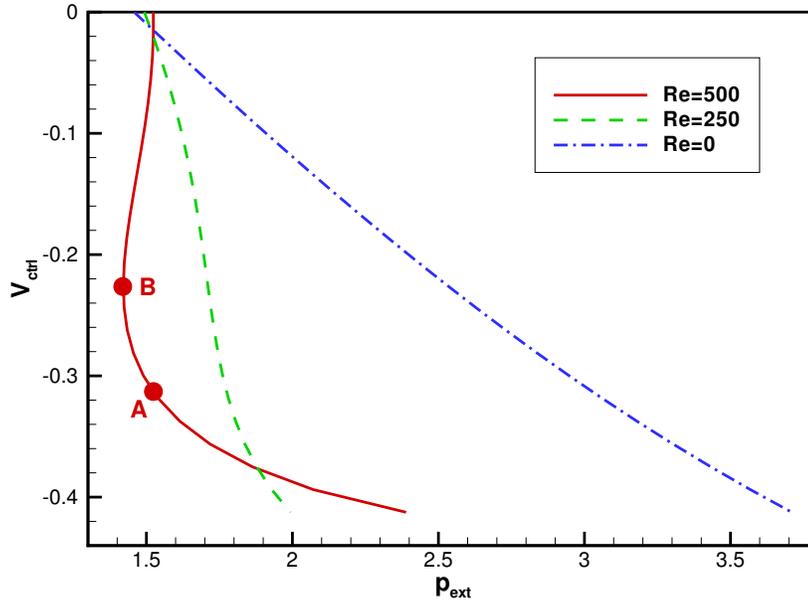


Fig. 4. Pressure-displacement characteristics for steady flow at various Reynolds numbers. The wall deformation is characterised by the vertical wall displacement $V_{ctrl} = v^2(\zeta = 3.5)$ at 70% of the membrane's length. At large Reynolds number, limit points develop in the pressure-displacement curve.

configuration in which $V_{ctrl} = -0.313$ (marked as point A in Fig. 4). Upon reducing the external pressure, the system evolves quasi-steadily along the curve from A to the limit point B where it has to undergo a further discontinuous transition to a strongly inflated state with $V_{ctrl} > 0$ (not shown in the diagram).

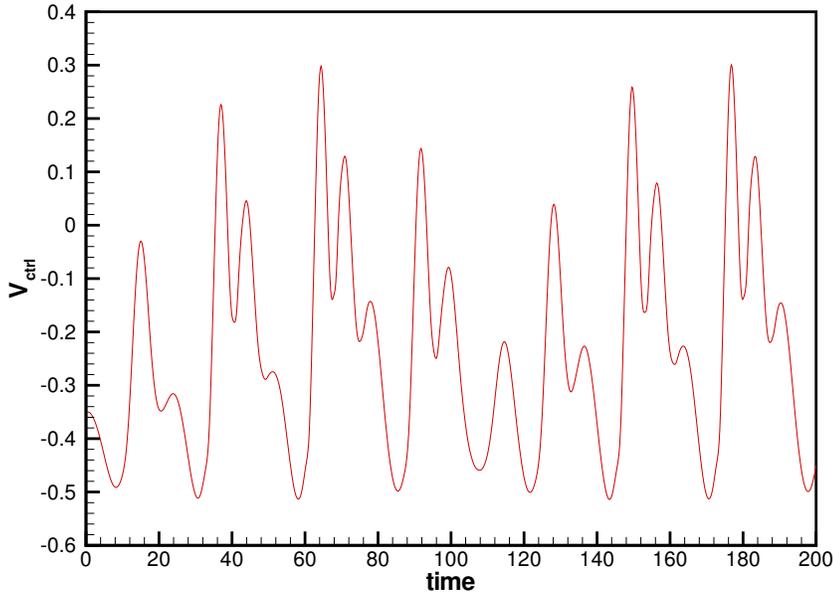


Fig. 5. Time history of the vertical wall displacement $V_{ctrl}(t) = v^2(\zeta = 3.5, t)$ for a self-excited oscillation at $Re = 500$. $p_{ext} = p_{ext}^*/E = 2.5$.

4.2 Unsteady flows – self-excited oscillations

As mentioned in the introduction, fluid-conveying elastic vessels are highly susceptible to the development of large-amplitude self-excited oscillations. An example of such oscillations is shown in Figs. 5 and 6. The simulation was started from the steady equilibrium configuration for $V_{ctrl} = 0.35$, $p_{ext} = 1.666$. For $t \geq 0$, the external pressure was set to $p_{ext} = 2.5$ and the flow rate was held constant at the upstream end of the channel. The time trace of the control displacement $V_{ctrl}(t) = v^2(\zeta = 3.5, t)$ shown in Fig. 5 illustrates that the system performs sustained large-amplitude oscillations with a fundamental period of $\mathcal{T} \approx 25$, with superimposed higher harmonics which create an oscillation of great complexity.

Fig. 6 illustrates the interaction between the wall displacement and the fluid flow during the oscillation. In Figs. 6(a-b), the wall collapses inwards and the size of the separation bubble at the downstream end of the elastic segment increases. During the second half of the oscillation, when the downstream half of the membrane moves outwards again, Figs. 6(c-d), the separation bubble detaches from the elastic wall, and a large-amplitude wave pattern (a ‘vorticity wave’ [29]) develops and travels downstream.

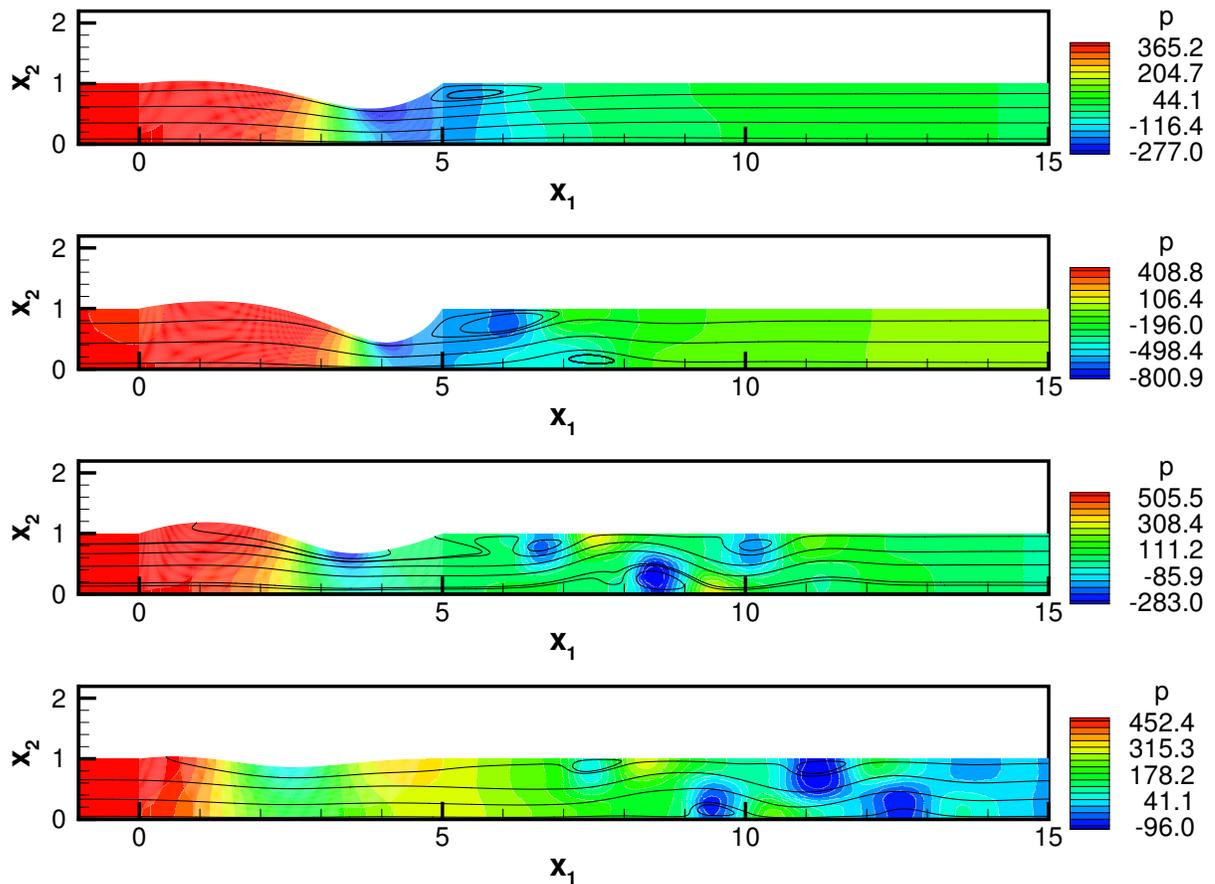


Fig. 6. Instantaneous streamlines and pressure contours in an unsteady flow undergoing a large-amplitude self-excited oscillation. $t = 2.0$ (top), 8.0 , 12.0 , 14.0 (bottom). $Re = 500$, $p_{ext} = p_{ext}^*/E = 2.5$. The fluid pressure is scaled on the viscous scale, $p = p^*a/(\mu U)$.

5 Solution strategies

For two-dimensional problems of moderate size, the direct solution of the linear system (20) is perfectly feasible, both in terms of memory requirements and CPU times. In the simulations presented above, the linear solves were performed with a frontal method (MA42 from the HSL2000 library [30]) and the Newton iteration typically converged in 5-7 steps, reducing the maximum residual from $\mathcal{O}(1)$ to below 10^{-8} . For larger problems, particularly in 3D, the memory requirements of direct solvers tend to become prohibitive – mainly because of the fill-in that occurs during the computation of the LU factors. In fluid-structure interaction problems, a large amount of the fill-in is created by the coupling between the fluid and solid variables. We will now explore alternative approaches which avoid the direct solution of the full system (20).

The additional fill-in associated with the fluid-structure interaction can be avoided by replacing the exact Jacobian matrices by suitable approximations. We consider the three block-triangular approximations

$$\mathcal{J}_{sup} = \left[\begin{array}{c|cc} \mathbf{S} & \mathbf{C}_{su} & \mathbf{C}_{sp} \\ \hline \mathbf{0} & \mathbf{F} & \mathbf{G} \\ \hline \mathbf{0} & \mathbf{D} & \mathbf{0} \end{array} \right] \quad \mathcal{J}_{sub} = \left[\begin{array}{c|cc} \mathbf{S} & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{C}_{us} & \mathbf{F} & \mathbf{G} \\ \hline \mathbf{C}_{ps} & \mathbf{D} & \mathbf{0} \end{array} \right] \quad \mathcal{J}_{diag} = \left[\begin{array}{c|cc} \mathbf{S} & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{F} & \mathbf{G} \\ \hline \mathbf{0} & \mathbf{D} & \mathbf{0} \end{array} \right], \quad (25)$$

the use of which reduces the solution of the linear system (20) to the independent solution of two smaller linear systems and an intermediate matrix-vector product. For instance, if we replace \mathcal{J} by \mathcal{J}_{sub} , the Newton correction \mathbf{x} can be computed via the sequence

$$\text{solve } \mathbf{S} \mathbf{x}_s = -\mathbf{r}_s \quad \text{for } \mathbf{x}_s, \quad (26)$$

$$\text{compute } \begin{pmatrix} \widehat{\mathbf{r}}_u \\ \widehat{\mathbf{r}}_p \end{pmatrix} = \begin{pmatrix} \mathbf{r}_u \\ \mathbf{r}_p \end{pmatrix} + \begin{bmatrix} \mathbf{C}_{us} \\ \mathbf{C}_{ps} \end{bmatrix} \mathbf{x}_s, \quad (27)$$

$$\text{solve } \begin{bmatrix} \mathbf{F} & \mathbf{G} \\ \mathbf{D} & \mathbf{0} \end{bmatrix} \begin{pmatrix} \mathbf{x}_u \\ \mathbf{x}_p \end{pmatrix} = - \begin{pmatrix} \widehat{\mathbf{r}}_u \\ \widehat{\mathbf{r}}_p \end{pmatrix} \quad \text{for } (\mathbf{x}_u, \mathbf{x}_p)^T. \quad (28)$$

A similar sequence of steps is required if the sub-diagonal blocks are neglected. If all fluid-structure interaction blocks are discarded, the intermediate matrix-vector product (27) can be omitted.

If we also approximate \mathbf{S} in (26) by the tangent stiffness matrix $\mathbf{S}^{(solid)}$, the two linear solves in (26) and (28) involve the same matrices that would arise from the discretisation of the uncoupled fluid and solid problems. The solution of the two linear systems can therefore utilise well-established solution techniques for these sub-problems. We will exploit this in section 5.2.2.

Fig. 7 compares the convergence history of the exact and inexact Newton methods for the steady problem illustrated in Fig. 3. In all cases the number of Newton iterations required to reduce the maximum residual below 10^{-8} tends to increase with the degree of the channel collapse. Fig. 7 shows the convergence histories of the fastest and slowest converging Newton iteration over the whole range of the simulation (V_{ctrl} increasing from 0 to -0.4125). If the exact Jacobian is used, the Newton method displays the expected rapid convergence and the iteration count only increases very slightly (from 5 to 7) with the degree of collapse. Neglecting the super- or sub-diagonal interaction blocks leads to a noticeable deterioration in the convergence rate. Furthermore,

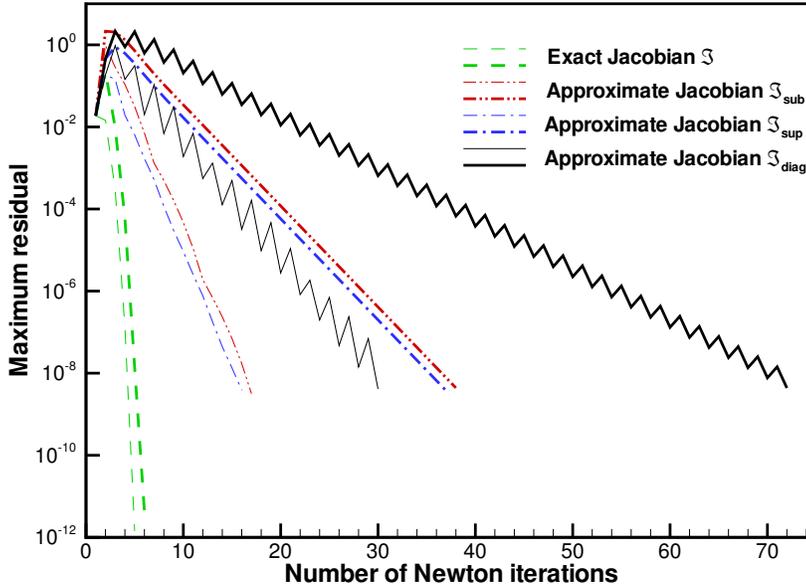


Fig. 7. Convergence histories of the Newton iteration with the exact and approximate Jacobian matrices for a steady computation at $Re = 500$. Thin lines: The solution for $V_{ctrl} = -0.0938$ is used as an initial guess for $V_{ctrl} = -0.1125$. Thick lines: The solution for $V_{ctrl} = -0.3938$ is used as an initial guess for $V_{ctrl} = -0.4125$.

the iteration count increases more rapidly with the degree of collapse (from ≈ 16 for $V_{ctrl} = -0.1125$ to ≈ 37 for $V_{ctrl} = -0.4125$). Neglecting all interaction blocks leads to a staggered scheme, in which the update of the fluid and solid variables is performed independently. This leads to a small saving in computational cost during each iteration because the matrix-vector product (27) can be avoided. However, the convergence rate of the Newton iteration deteriorates dramatically and also becomes extremely sensitive to the degree of collapse. Neglecting the block $\mathbf{S}^{(FSI)}$ did not affect the convergence rates in any of the four schemes. Overall, this behaviour is consistent with the observations in [31] where only the case of neglecting the sub-diagonal block (in our notation) was considered.

When the approximate Jacobian matrices (25) were used in the unsteady simulation shown in section 4.2, the Newton method failed to converge. Only the blocks \mathbf{C}_{su} and $\mathbf{S}^{(solid)}$ could be neglected without causing immediate divergence. However, in that case, no useful simplification of the block structure is achieved.

The numerical experiments presented in the previous section show that the use of approximate Jacobian matrices in the Newton iteration can lead to a severe deterioration of its convergence rate. An alternative approach is to use the approximate Jacobians as preconditioners in the iterative solution with a Krylov subspace solver, such as GMRES. The iterative solution of the (right-)preconditioned version of the linear system (20), $\mathcal{J}\mathcal{P}^{-1}\mathcal{P}\mathbf{x} = -\mathbf{r}$, involves matrix-vector products with the exact Jacobian matrix \mathcal{J} and solutions of linear systems with the preconditioning matrix \mathcal{P} . If \mathcal{P} is a good approximation to \mathcal{J} , we can expect rapid convergence of the ‘inner’ GMRES iteration to the exact solution of the original system (20), while maintaining the quadratic convergence rate of the ‘outer’ Newton iteration. Furthermore, the convergence tolerance of the ‘inner’ iteration can be coupled to the size of the residual in the Newton iteration to avoid ‘over-solving’; see [32]. We will explore this in section 5.2.3.

5.2.1 Preconditioning with the block-triangular approximate Jacobians

To illustrate the efficiency of the approximate Jacobians as preconditioners for the iterative solution of the linear system (20), Fig. 8 shows representative convergence histories of the GMRES iterations in a steady simulation at $Re = 500$. The data was obtained from a computation in which the solution for $V_{ctrl} = -0.21$ was used as the initial guess for $V_{ctrl} = -0.28$. In this particular case, the Newton iteration converges in five steps and the lines in Fig. 8 show the convergence histories of the GMRES iterations for the first and last Newton steps. During the first Newton steps, the GMRES convergence histories have an initial stagnant period (which lasts for approximately 50% of the total number of iterations) during which the scaled residual $r = |\mathcal{J}\mathbf{x} + \mathbf{r}|/|\mathbf{r}|$ only decays very little. This is followed by a period of rapid convergence during the second half of the iteration. For subsequent Newton steps, the initial stagnant period disappears. Overall, the preconditioner \mathcal{J}_{sup} produces slightly faster convergence rates than \mathcal{J}_{sub} , especially during the final Newton steps. The block diagonal preconditioner \mathcal{J}_{diag} requires approximately twice as many GMRES iterations as the two block-triangular versions.

Table 1 provides a more comprehensive overview of the iteration counts for the three preconditioners and presents the average number of GMRES iterations, N_{GMRES} , for steady computations at different Reynolds numbers. The average was taken over all Newton steps in a computation in which V_{ctrl} was increased from 0 to -0.35 in five steps. In all cases, the GMRES iteration was continued until $r < \text{TOL}_{GMRES} = 10^{-6}$. Table 1 shows that an increase in Reynolds number leads to a slight increase in N_{GMRES} and confirms \mathcal{J}_{sup} as the most efficient pre-

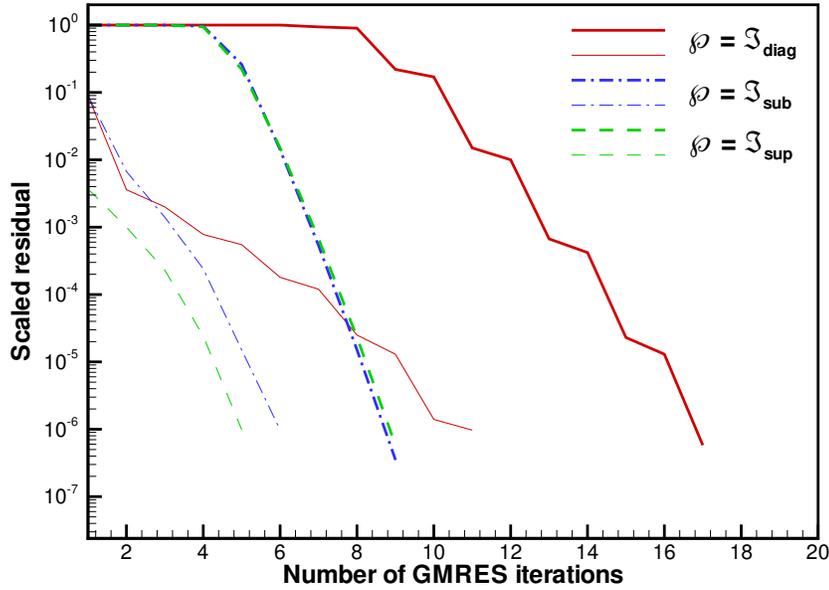


Fig. 8. Convergence histories of the preconditioned GMRES iteration for a steady computation at $Re = 500$. The solution for $V_{ctrl} = -0.21$ is used as the initial guess for $V_{ctrl} = -0.28$. Thick lines: first Newton step; thin lines: last (=fifth) Newton step.

conditioner. The data shown was obtained with $h_{FEM} = 1/8$ but computations at different spatial resolutions displayed no evidence of any mesh-dependence of the convergence rates.

	$Re = 0$	$Re = 250$	$Re = 500$
$\mathcal{P} = \mathcal{J}_{sup}$	3.2	5.4	6.2
$\mathcal{P} = \mathcal{J}_{sub}$	3.8	6.2	7.0
$\mathcal{P} = \mathcal{J}_{diag}$	6.4	11.0	12.4

Table 1

Average GMRES iteration counts for the steady simulations shown in section 4.1. The approximate Jacobians \mathcal{J}_{sup} , \mathcal{J}_{sub} and \mathcal{J}_{diag} are used as right preconditioners for the iterative solution of the linear system (20). $TOL_{GMRES} = 10^{-6}$.

Table 2 presents the average iterations counts for the three preconditioners in the unsteady simulation at $Re = 500$, shown section 4.2. Here N_{GMRES} represents the average over all Newton iterations at 20 different time-steps during the system's evolution. The data was obtained by restarting the simulation from the four representative configurations shown in Fig. 6 and continuing the time-integration for five time-steps. Selected computations were performed over several periods of the oscillation to confirm that the averages shown in Table

2 are representative.

Table 2 shows that, as in the steady case, the performance of \mathcal{J}_{sup} and \mathcal{J}_{sub} is very similar while the the block-diagonal preconditioner \mathcal{J}_{diag} requires approximately twice as many iterations. The iteration count increases slightly with a reduction in the time-step, $N_{\text{GMRES}} \sim \Delta t^{-0.3}$ so that reducing the time-step by a factor of 10 approximately doubles the number of iterations. This is documented in more detail in Fig. 9 which compares representative convergence histories of the GMRES iterations (right-preconditioned with $\mathcal{P} = \mathcal{J}_{sup}$) for different values of Δt . A reduction in Δt only leads to a slight reduction in the asymptotic convergence rate of the GMRES iteration but significantly extends the length of the initial stagnant period. Furthermore, at small Δt the initial stagnant period is longer for the final rather than the initial Newton steps.

The increase in N_{GMRES} with a decrease in Δt presents a significant difference to the behaviour typically observed in the iterative solution of pure fluid or solid problems. In uncoupled problems, a reduction in the time-step tends to increase the diagonal dominance of the Jacobian matrix and this typically accelerates the convergence of iterative solvers. In the present problem, the entries in the sub-diagonal block \mathbf{C}_{us} grow like $1/\Delta t$ as $\Delta t \rightarrow 0$ and, in fact, tend to become the dominant entries in the matrix. Interestingly, the fact that the block \mathbf{C}_{us} is retained in \mathcal{J}_{sub} does not significantly increase the efficiency of this preconditioner at small Δt .

	$\Delta t = 2 \times 10^{-1}$	$\Delta t = 2 \times 10^{-2}$	$\Delta t = 2 \times 10^{-3}$
$\mathcal{P} = \mathcal{J}_{sup}$	11.5	23.4	45.2
$\mathcal{P} = \mathcal{J}_{sub}$	12.4	24.2	41.2
$\mathcal{P} = \mathcal{J}_{diag}$	22.7	46.2	87.2

Table 2

Average GMRES iteration counts for the unsteady simulation shown in section 4.2. The approximate Jacobians \mathcal{J}_{sup} , \mathcal{J}_{sub} and \mathcal{J}_{diag} are used as right preconditioners for the iterative solution of the linear system (20). $Re = 500$. $\text{TOL}_{\text{GMRES}} = 10^{-6}$.

5.2.2 A pressure Schur complement preconditioner

The results shown in the previous section demonstrate that the approximate Jacobians present efficient preconditioners for the iterative solution of the linear system (20). However, in their exact form, the application of the preconditioners is very expensive and the computational cost is dominated by the solution of the Navier-Stokes system (28). For the purposes of preconditioning, its solution does not have to be computed exactly but can be replaced by a computationally cheaper approximation. For this purpose, we replace the Navier-Stokes block in (25) by a global pressure Schur complement preconditioner

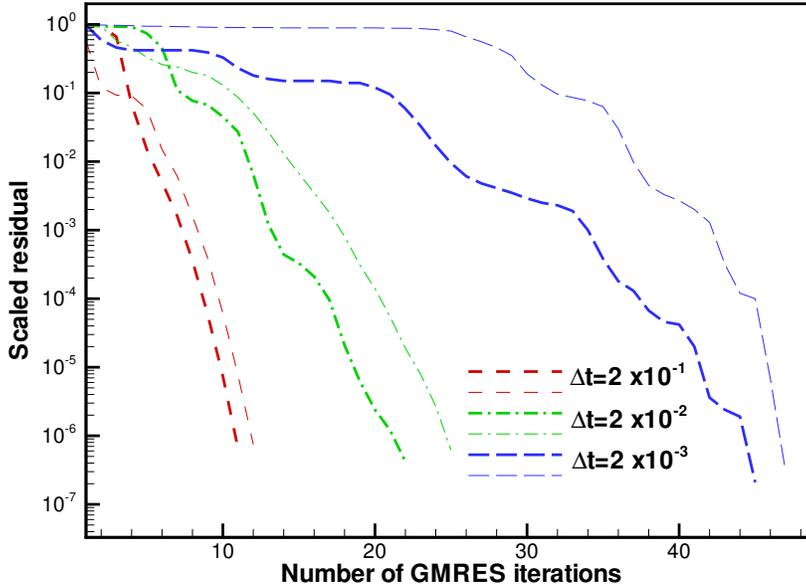


Fig. 9. Convergence histories of the GMRES iteration for different time-steps Δt . \mathcal{J}_{sup} is used as the right preconditioner and the results are for the unsteady simulation of section 4.2 at $t = 14$. $Re = 500$. Thick lines: first Newton step; thin lines: last (=fifth) Newton step.

tioner [33]. This results in preconditioners of the form

$$\mathcal{P}_{sup} = \left[\begin{array}{c|cc} \mathbf{S} & \mathbf{C}_{su} & \mathbf{C}_{sp} \\ \hline \mathbf{0} & \mathbf{F} & \mathbf{G} \\ \mathbf{0} & \mathbf{0} & -\mathbf{X} \end{array} \right] \quad \mathcal{P}_{sub} = \left[\begin{array}{c|cc} \mathbf{S} & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{C}_{us} & \mathbf{F} & \mathbf{G} \\ \mathbf{C}_{ps} & \mathbf{0} & -\mathbf{X} \end{array} \right] \quad \mathcal{P}_{diag} = \left[\begin{array}{c|cc} \mathbf{S} & \mathbf{0} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{F} & \mathbf{G} \\ \mathbf{0} & \mathbf{0} & -\mathbf{X} \end{array} \right], \quad (29)$$

where \mathbf{X} is (an approximation of) the pressure Schur complement $\mathbf{DF}^{-1}\mathbf{G}$. In the following we will use Elman's BFBt approximation [34]

$$\mathbf{X} = (\mathbf{DG}) (\mathbf{DFG})^{-1} (\mathbf{DG}), \quad (30)$$

though other approximations could be used in its place. The action of the preconditioners defined by (29) and (30) can then be obtained in a number of sub-steps. For instance, the solution of the system $\mathcal{P}_{sub}\mathbf{x} = \mathbf{r}$ can be computed via

$$\text{solve } \mathbf{S} \mathbf{x}_s = -\mathbf{r}_s \text{ for } \mathbf{x}_s, \quad (31)$$

$$\text{compute } \begin{pmatrix} \widehat{\mathbf{r}}_u \\ \widehat{\mathbf{r}}_p \end{pmatrix} = \begin{pmatrix} \mathbf{r}_u \\ \mathbf{r}_p \end{pmatrix} + \begin{bmatrix} \mathbf{C}_{us} \\ \mathbf{C}_{ps} \end{bmatrix} \mathbf{x}_s, \quad (32)$$

$$\text{solve } (\mathbf{DG}) \mathbf{y} = \widehat{\mathbf{r}}_p \text{ for } \mathbf{y}, \quad (33)$$

$$\text{compute } \mathbf{z} = (\mathbf{DFG}) \mathbf{y}, \quad (34)$$

$$\text{solve } (\mathbf{DG}) \mathbf{x}_p = \mathbf{z} \text{ for } \mathbf{x}_p, \quad (35)$$

$$\text{compute } \mathbf{b}_u = -\widehat{\mathbf{r}}_u - \mathbf{G} \mathbf{x}_p, \quad (36)$$

$$\text{solve } \mathbf{F} \mathbf{x}_u = \mathbf{b}_u \text{ for } \mathbf{x}_u. \quad (37)$$

Similar sequences arise for \mathcal{P}_{sup} and \mathcal{P}_{diag} . The procedure can be viewed as an extension of the block preconditioning strategy for ‘pure’ fluids problems proposed in reference [35]. In this form, the application of the preconditioner involves the solution of four linear systems and (up to) three matrix-vector products.

The matrix \mathbf{S} in (31) arises from the discretisation of a spatially one-dimensional problem and is therefore much smaller than the matrices in (33), (35) and (37). The linear system (31) can therefore be solved efficiently by a sparse direct solver; we used Demmel *et al.*’s SuperLU solver [36]. The computationally most expensive part of the preconditioning operation is the solution of the three linear systems (33), (35) and (37). We will now replace the exact solves for these systems by suitable approximations to improve the overall efficiency of the algorithm.

In the absence of stabilisation, the matrix \mathbf{DG} is symmetric and positive definite and plays a role similar to a pressure Poisson operator; see [34]. This suggests that the systems (33) and (35) can be solved efficiently by multigrid, [37], and that suitable approximations can be obtained by performing a fixed number of multigrid cycles. We performed three V-cycles, using point Gauss-Seidel (with two pre- and two post-smoothing steps on each level) as the smoother. The interpolation was based on the bilinear pressure interpolation, full weighting was used for the injection and the linear system on the coarsest mesh was solved with SuperLU.

At finite Reynolds number, the matrix \mathbf{F} in (37) is non-symmetric and has a complex spectrum. Consequently, at large Reynolds number, multigrid with simple point Gauss-Seidel smoothing does not act as an efficient solver. However, replacing the exact solution of (37) by a fixed number of multigrid V-cycles can still provide an efficient preconditioner, provided a sufficient amount of stabilisation is added on each multigrid level. The adaptive choice of the stabilisation parameter τ_{stab} via (18) ensures that τ_{stab} can be relatively large on the coarser multigrid levels (on which the solutions only provide corrections for the finer levels) but remains small enough not to affect the solution at the finest level, provided consistent stabilisation is used; see section 5.3. The implementation of the multigrid solver was based on interpolation consistent with the bi-quadratic variation of the velocity degrees of freedom and full weighting was used for the injection. SuperLU was used as the direct solver on the coarsest level. ‘With-the-wind’ numbering of the equations is known to be beneficial for the efficiency of Gauss-Seidel smoothing but would be difficult

and costly to implement for the highly complex and time-varying flow fields encountered in the present problem; see, e.g., Fig. 6. In each smoothing step, we performed four consecutive point Gauss-Seidel sweeps over all variables, sweeping along the four principal directions of the structured mesh [38].

In this form, the computational work associated with a single application of the preconditioner scales linearly with the number of unknowns in the problem. Hence, the ‘optimality’ of the solver is directly related to the number of GMRES iterations required to reduce the scaled residual r below the convergence tolerance.

Table 3 shows the average values for N_{GMRES} for steady computations at various Reynolds numbers and spatial resolutions. We only present data for \mathcal{P}_{sup} since it tended to be the most efficient preconditioner throughout. The two numbers shown for each combination of Reynolds number and mesh size correspond to the cases where the linear systems (33), (35) and (37) were solved exactly (by SuperLU) and approximately (by the fixed-cost multigrid approximations). The use of the multigrid approximations typically increased the number of GMRES iterations by about 30% (interestingly, in one case it actually reduced N_{GMRES}) but this was more than compensated for by a significant decrease in the overall CPU times. For instance, for the case $Re = 500$ and $h_{\text{FEM}} = 1/16$, the 30% increase in N_{GMRES} still allowed a reduction in the total CPU time by a factor of 2.6.

The relative increase in N_{GMRES} with an increase in Reynolds number from $Re = 250$ to $Re = 500$ is comparable to that observed for the ‘exact’ preconditioner \mathcal{J}_{sup} in Table 1, whereas at smaller Reynolds numbers, an increase Re has a more detrimental effect. The iteration count scales approximately like $N_{\text{GMRES}} \sim h_{\text{FEM}}^{-0.7}$ which is comparable to the mesh dependence reported in reference [34], where the performance of the BFBt preconditioner was examined in an Oseen problem.

	$Re = 0$	$Re = 250$	$Re = 500$
$h_{\text{FEM}} = 1/4$	14.4 (15.8)	41.2 (57.0)	46.8 (71.2)
$h_{\text{FEM}} = 1/8$	25.0 (25.8)	71.8 (86.8)	81.2 (104.0)
$h_{\text{FEM}} = 1/12$	35.6 (31.2)	91.0 (116.6)	111.4 (141.4)
$h_{\text{FEM}} = 1/16$	44.8 (46.4)	104.6 (130.2)	130.8 (173.0)

Table 3

Average GMRES iteration counts for the steady simulations shown in section 4.1. \mathcal{P}_{sup} is used as the right preconditioners for the iterative solution of the linear system (20). $\text{TOL}_{\text{GMRES}} = 10^{-6}$. The numbers in brackets are for the case when the fixed-cost multigrid approximations for the systems (33), (35) and (37) are used.

Fig. 10 shows that the (significant) differences in the convergence histories at

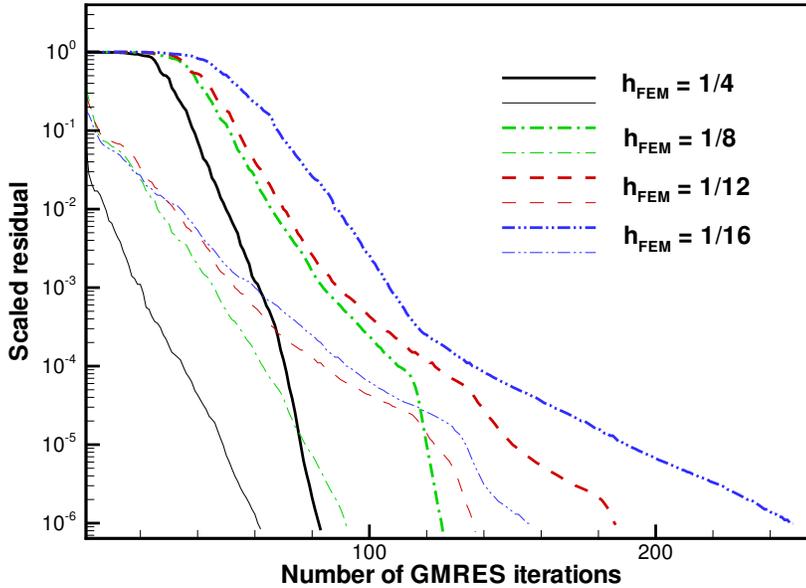


Fig. 10. Convergence histories of the preconditioned GMRES iteration for a steady computation at $Re = 500$. \mathcal{P}_{sup} with fixed-cost multigrid approximations for (33), (35) and (37) is used as the right preconditioner. The solution for $V_{ctrl} = -0.21$ is used as the initial guess for $V_{ctrl} = -0.28$. Thick lines: first Newton step; thin lines: last (=fifth) Newton step.

different stages of the Newton iteration are similar to those observed for \mathcal{J}_{sup} in Fig. 8. However, the deterioration in the asymptotic convergence rate with a reduction in h_{FEM} makes the presence of the initial stagnant period relatively less important for simulations on fine meshes.

Table 4 presents the corresponding data for the unsteady simulations of section 4.2 at $Re = 500$. The iteration counts are very similar to those presented in Table 2 for \mathcal{J}_{sup} , and display very little mesh-dependence. In fact, the iteration counts for $\Delta t = 2 \times 10^{-2}$ are practically independent of h_{FEM} . Even the detailed convergence histories shown in Fig. 11 are very similar to those presented for \mathcal{J}_{sup} in Fig. 9, indicating that the approximate implementation of the preconditioner has little effect on its efficiency.

5.2.3 The choice of the convergence tolerance for the GMRES iteration

In order to allow a meaningful comparison between the various preconditioners, we used a fixed convergence tolerance of $TOL_{GMRES} = 10^{-6}$ for all studies in the previous sections. Dembo *et al.* [39] showed that the adaptive choice

$$TOL_{GMRES} = \mathcal{O}(|\mathbf{r}|) \quad (38)$$

	$\Delta t = 2 \times 10^{-1}$	$\Delta t = 2 \times 10^{-2}$	$\Delta t = 2 \times 10^{-3}$
$h_{\text{FEM}} = 1/4$	17.0 (17.6)	23.0 (24.2)	25.0 (28.0)
$h_{\text{FEM}} = 1/8$	17.9 (23.1)	24.2 (29.4)	46.0 (50.5)
$h_{\text{FEM}} = 1/12$	22.6 (28.3)	24.2 (30.4)	53.9 (64.6)
$h_{\text{FEM}} = 1/16$	28.2 (44.9)	23.9 (35.6)	55.8 (76.4)

Table 4

Average GMRES iteration counts for the unsteady simulation of section 4.2. \mathcal{P}_{sup} is used as the right preconditioners for the iterative solution of the linear system (20). $\text{TOL}_{\text{GMRES}} = 10^{-6}$. The numbers in brackets are for the case when the fixed-cost multigrid approximations for the systems (33), (35) and (37) are used.

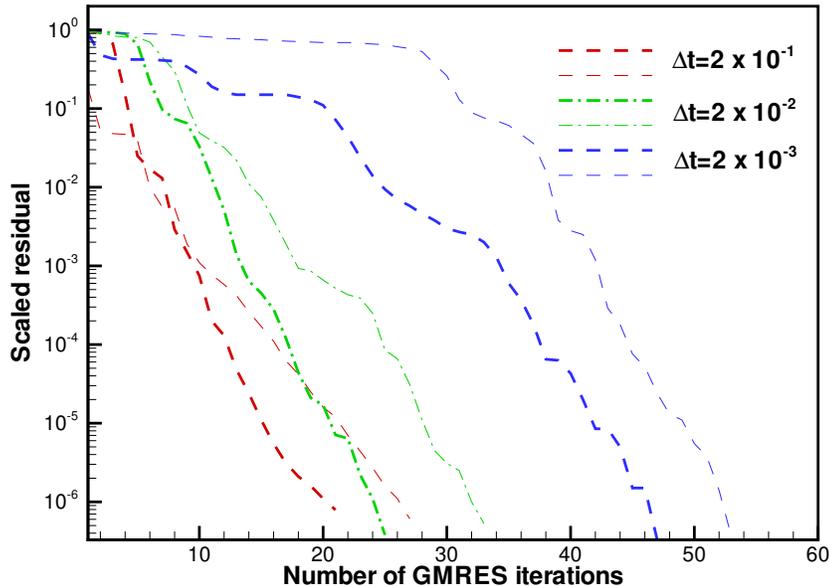


Fig. 11. Convergence histories of the GMRES iteration for different time-steps Δt . \mathcal{P}_{sup} with fixed-cost multigrid approximations for (33), (35) and (37) is used as the right preconditioner and the results are for the unsteady simulation of section 4.2 at $t = 14$. $h_{\text{FEM}} = 1/12$, $Re = 500$. Thick lines: first Newton step; thin lines: last (=fifth) Newton step.

is sufficient to maintain the (asymptotically) quadratic convergence rate of the Newton iteration. (38) shows that the linear system (20) needs to be solved more accurately as the Newton iteration approaches convergence. This also implies that a natural lower limit for $\text{TOL}_{\text{GMRES}}$ is given by the convergence tolerance of the Newton method itself.

In practice, the adaptive choice of $\text{TOL}_{\text{GMRES}}$ can be rather delicate and problem-dependent since the analysis underlying the derivation of (38) only applies

close to the converged solution. If the initial guess for the Newton iteration has a large residual then the large value of $\text{TOL}_{\text{GMRES}}$ suggested by (38) can cause the Newton method to diverge. Furthermore, any saving in computational work resulting from the reduction in N_{GMRES} must be compared to the potential additional work created by a possible small increase in the total number of Newton steps. This is particularly important if, as in the present case, good (but costly to construct) preconditioners already lead to small numbers of GMRES iterations. In such cases, even the cost of constructing a single additional Jacobian matrix and the associated preconditioning operators can outweigh the savings due to the reduction in N_{GMRES} . Finally, the potential benefits resulting from the adaptive adjustment of $\text{TOL}_{\text{GMRES}}$ depend strongly on the convergence characteristics of the GMRES iteration. If the GMRES iteration has a long initial stagnant period then an increase in $\text{TOL}_{\text{GMRES}}$ can only lead to a small reduction in N_{GMRES} .

We performed extensive numerical experiments to determine an ‘optimal’ strategy for the choice of $\text{TOL}_{\text{GMRES}}$ and found that, for steady problems, the choice

$$\text{TOL}_{\text{GMRES}} = \begin{cases} 10^{-4} & \text{if } |\mathbf{r}| > 10^{-4} \\ |\mathbf{r}| & \text{if } 10^{-6} \leq |\mathbf{r}| \leq 10^{-4} \\ 10^{-6} & \text{if } |\mathbf{r}| < 10^{-6} \end{cases} \quad (39)$$

resulted in modest savings in CPU times (typically about 20-30% compared to a fixed tolerance of $\text{TOL}_{\text{GMRES}} = 10^{-6}$) while ensuring robustness in the sense that it only rarely resulted in costly additional Newton steps. The benefits of the adaptive adjustment of $\text{TOL}_{\text{GMRES}}$ are limited by the fact that, in steady problems, the GMRES iteration displays the initial stagnant period during the first few Newton steps. In unsteady problems, the initial stagnant period is more pronounced during the final Newton steps but it tends to persist for longer, particularly at small time-steps. Furthermore, the number of GMRES iterations tends to be much smaller than in steady problems, and the adaptive adjustment of $\text{TOL}_{\text{GMRES}}$ in unsteady problems only led to small savings in CPU time.

5.3 *The importance of consistent stabilisation*

The use of multigrid approximations for the solution of the linear systems (33), (35) and (37) presents a key step in the efficient implementation of the preconditioning strategy. The stabilisation of the fluid momentum equations, introduced in section 3.2.2, is necessary to allow the solution of the linear system (37) on the coarse meshes. The adaptive choice (18) for the stabilisation parameter τ_{stab} ensures that only the minimum amount of stabilisation is applied on each multigrid level and that the stabilisation (whether applied

consistently or inconsistently) becomes insignificant as the element size h_{FEM} and/or the time-step Δt tend to zero. However, (18) shows that $\tau_{stab} \sim h_{\text{FEM}}$ as $h_{\text{FEM}} \rightarrow 0$ whereas, in the absence of stabilisation, the finite element expansions (12) ensure the quadratic convergence of the solution under mesh refinement. Thus, as long as $Re_M > 1$ anywhere in the domain, any errors that are introduced by inconsistent stabilisation reduce the rate of convergence under mesh refinement. We will now demonstrate that in fluid-structure interaction problems, the use of inconsistent stabilisation can introduce significant errors.

Figs. 12(a) and (c) present results of representative steady and unsteady simulations at different spatial resolutions. (On the coarsest mesh, the flow is severely under-resolved and the Newton method diverges at $t = 8$.) The computations were performed with a direct solver and no stabilisation was applied. The figures demonstrate that the finest mesh ($h_{\text{FEM}} = 1/16$) is adequate to resolve the results to within plotting accuracy. Figs. 12(b) and (d) present the corresponding results (computed only on the finest mesh) with various forms of stabilisation. Inconsistent stabilisation can be seen to introduce significant errors even though, in the absence of stabilisation, the solution is fully converged on this mesh.

The explanation for the strong effect of inconsistent stabilisation in the present problem is that classical SUPG stabilisation (implemented by including only term II into (19)) introduces artificial dissipation into the problem. This implies that, for a given velocity field, a larger pressure drop is required to drive the flow. In fluid-structure interaction problems, this can introduce large errors into the coupled solution since the wall deformation depends strongly on the fluid pressure. To assess the significance of this effect in the present problem, Fig. 13 shows the results of a steady computation in which the wall displacement field was prescribed (a sinusoidal indentation with a maximum inwards displacement of 0.5; the resulting degree of collapse is comparable to that in the most strongly collapsed configuration in Fig. 12). The lines show the two velocity components u_1 and u_2 and the pressure p , plotted along the line $x_2 = 0.4$ (marked as the thick dashed line in the inset). The figure shows that the different versions of stabilisation have little effect on the velocity fields. However, the additional dissipation introduced by inconsistent stabilisation leads to a noticeable increase in the pressure drop along the ‘membrane’. With SUPG stabilisation, the fluid pressure at the upstream end of the ‘membrane’ is $\Delta p = 14.6$ higher than with consistent stabilisation (or without any stabilisation at all). If the wall was elastic, this increase in fluid pressure would reduce the degree of collapse. Conversely, to maintain the same degree of collapse, the external pressure would have to be increased by an amount comparable to Δp . The fluid pressures (non-dimensionalised on the viscous scale) and the external pressure (non-dimensionalised on the elastic scale) are related by the parameter Q , via $\Delta p_{ext} = Q \Delta p$; see (5). In the computations presented in Fig. 12, $Q = 10^{-2}$, so the increase in external pressure required to

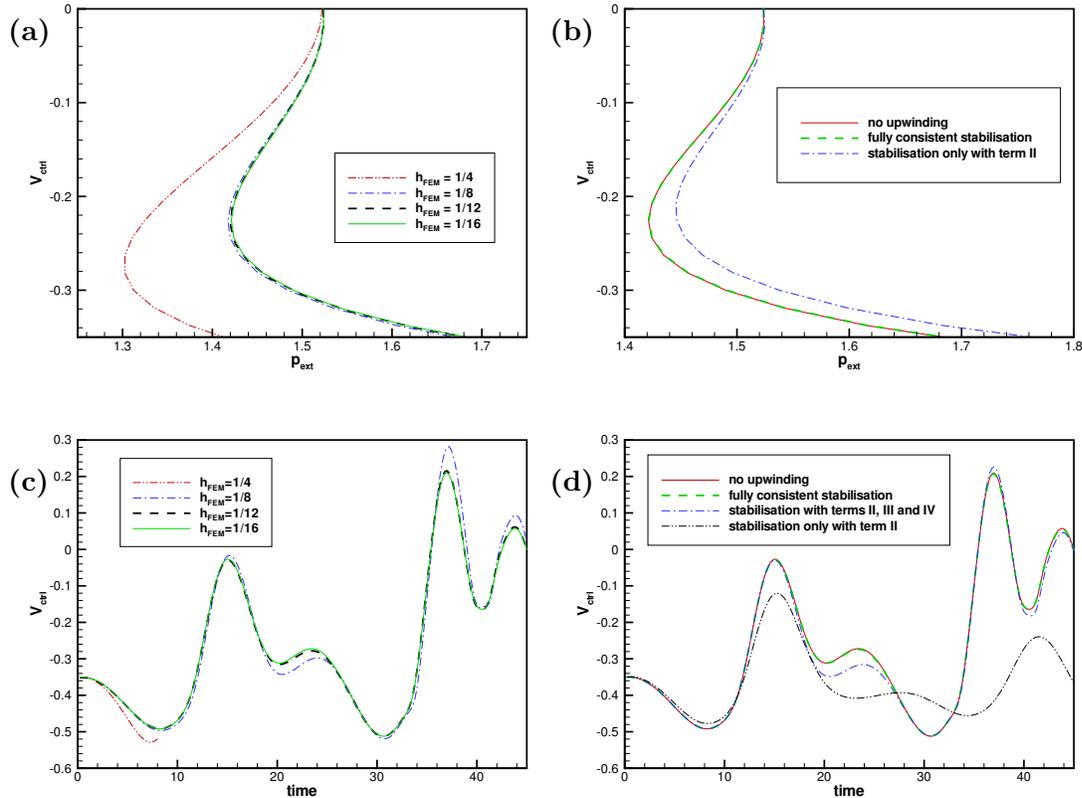


Fig. 12. Demonstration of the errors introduced by inconsistent stabilisation in steady (a and b) and unsteady (c and d) simulations at $Re = 500$. (a) and (c) document the mesh convergence of the solutions in the absence of stabilisation; Figs. (b) and (d) demonstrate the effect of inconsistent stabilisation on the finest mesh. The parameter values are the same as in section 4.

approximately compensate for the rise in fluid pressure is $\Delta p_{ext} = 0.146$ which is comparable to (in fact, nearly twice as large as) the error in p_{ext} observed in Fig. 12(b) where $\Delta p_{ext} = 0.08$ for $V_{ctrl} = 0.35$. This confirms that the effect of artificial dissipation introduced by inconsistent stabilisation is large enough to explain the differences observed in Fig. 12(b). In time-dependent simulations in which the external pressure is kept constant, the increase in dissipation constantly extracts energy from the system. Over sufficiently long time-scales, this can lead to a significant change in the system's behaviour. For instance, in the simulation shown in Fig. 12(d), the excessive dissipation causes the rapid decay of the oscillation.

Given that in the present problem, stabilisation is only required for the multi-grid preconditioning operations, it seems desirable to set $\tau_{stab} = 0$ on the finest mesh and use stabilisation (consistent or inconsistent) only for the auxiliary computations on the coarser meshes – following Gresho & Lee's maxim not to

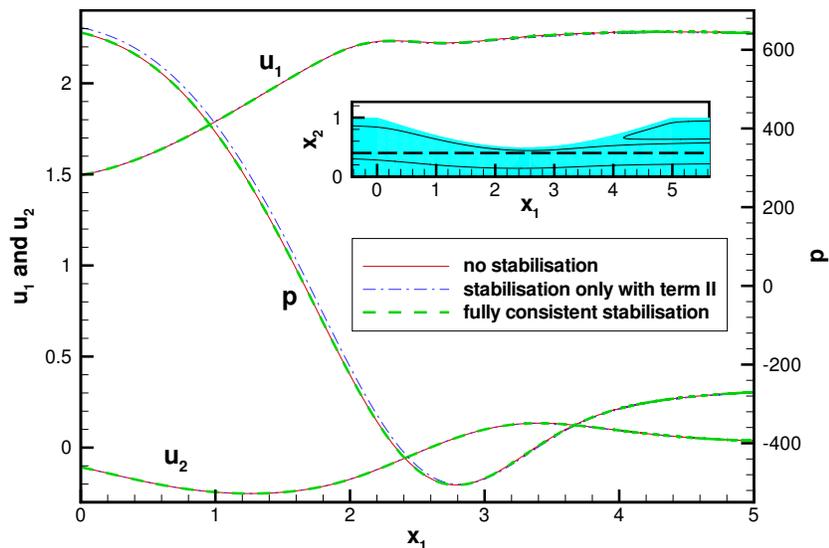


Fig. 13. The effect of inconsistent stabilisation on the velocities and pressures along the line $x_2 = 0.4$ (indicated by the dashed line in the inset) in a steady simulation at $Re = 500$. The wall shape is prescribed and the computation is performed on the finest mesh, $h_{\text{FEM}} = 1/16$.

“suppress the wiggles” [40]. This approach can be implemented in two ways: **(i)** Set $\tau_{stab} = 0$ for all computations on the finest mesh and use (18) to determine the stabilisation parameter for the coarser multigrid levels. **(ii)** Use (18) to determine the stabilisation parameter on all multigrid levels but set $\tau_{stab} = 0$ when computing the nonlinear residuals in the Newton method. Both methods ensure that the Newton method can only converge to the solution for $\tau_{stab} = 0$. However, numerical experiments showed that in case **(i)** the multigrid-preconditioned GMRES iteration fails to converge, indicating that the fine-grid problem with $\tau_{stab} = 0$ is poorly approximated by the coarse grid problems with $\tau_{stab} \neq 0$. With strategy **(ii)**, the GMRES iterations converged but the Newton iteration stagnated, indicating again that the stabilisation terms still play an important role on the finest meshes. Finally, we mention that the above results are independent of the pressure non-dimensionalisation – we observed exactly the same behaviour when the pressure was scaled on the inertial scale. We conclude that the preconditioning strategy proposed in this paper requires the use of consistent stabilisation on all multigrid levels.

6 Summary

We have developed an efficient preconditioning technique that allows the rapid iterative solution of the linear systems that arise in the fully coupled solution

of steady and unsteady large-displacement fluid-structure interaction problems with Newton’s method. The preconditioners were derived from block-triangular approximations of the Jacobian matrix. While the direct use of these matrices in the Newton method leads to an unacceptable deterioration of its convergence rate, they provide efficient preconditioners for the iterative solution of the linear system by GMRES. The block-triangular structure of the approximate Jacobians makes it possible to perform the preconditioning operation in a sequence of sub-steps. To improve the overall efficiency of the method, we developed approximate versions of the preconditioners, by replacing the Navier-Stokes blocks by a global pressure Schur complement preconditioner, using Elman’s BFBt approximation [34] for the pressure Schur complement. Finally, we replaced the most computationally expensive operations by multigrid approximations. In this form, the work involved in a single application of the preconditioner scales linearly with number of degrees of freedom in the problem. For steady problems, the mesh dependence of the iteration counts was found to be comparable to that observed in applications of the BFBt preconditioner in ‘pure’ fluids problems and the iteration counts only displayed a weak dependence on the Reynolds number. In time-dependent problems, the convergence rates were generally much faster and displayed very little mesh dependence, though the iterations counts increased slightly with a reduction in the time-step.

Acknowledgements

Financial support from the EPSRC is gratefully acknowledged. The author wishes to thank David Silvester, Milan Mihajlovic and Andrew Hazel for many helpful and enjoyable discussions. The HSL library routine MA42 (a frontal solver for sparse, unsymmetric systems [30]) and the the GMRES solver from the CERFACS repository [41] were used in this work.

References

- [1] C. Farhat, M. Lesoinne, Two efficient staggered algorithms for the serial and parallel solution of three-dimensional nonlinear transient aeroelastic problems, *Computer Methods in Applied Mechanics and Engineering* 182 (2000) 499–515.
- [2] K. Stein, R. Benney, T. Tezduyar, J. Potvin, Fluid-structure interactions of a cross parachute: numerical simulation, *Computer Methods in Applied Mechanics and Engineering* 191 (2001) 673–687.
- [3] F. Casadei, J. P. Halleux, A. Sala, F. Chille, Transient fluid-structure interaction algorithms for large industrial applications, *Computer Methods in Applied*

- [4] T. Charvet, F. Hauville, S. Huberson, Numerical simulation of the flow over sails in real sailing conditions, *Journal of Wind Engineering and Industrial Aerodynamics* 63 (1996) 111–129.
- [5] M. Heil, O. E. Jensen, Flows in deformable tubes and channels – theoretical models and biological applications, in: T. J. Pedley, P. W. Carpenter (Eds.), *Flow in Collapsible Tubes and Past Other Highly Compliant Boundaries*, Kluwer, Dordrecht, Netherlands, 2003, p. (In print).
- [6] J. D. Hart, G. W. M. Peters, P. J. G. Schreurs, F. P. T. Baaijens, A three-dimensional computational analysis of fluid-structure interaction in the aortic valve, *Journal of Biomechanics* 36 (2003) 103–112.
- [7] R. Ohayon, C. Felippa, Special issue: Advances in computational methods for fluid-structure interaction and coupled problems, *Computer Methods in Applied Mechanics and Engineering* 190.
- [8] M. Heil, Stokes flow in an elastic tube – a large-displacement fluid-structure interaction problem, *International Journal for Numerical Methods in Fluids* 28 (1998) 243–265.
- [9] D. P. Mok, W. A. Wall, Partitioned analysis schemes for the transient interaction of incompressible flows and nonlinear flexible structures, in: W. A. Wall, K.-U. Bletzinger, K. Schweizerhof (Eds.), *Trends in Computational Structural Mechanics*, CIMNE, Barcelona, Barcelona, Spain, 2001.
- [10] C. A. Felippa, K. C. Park, C. Farhat, Partitioned analysis of coupled mechanical systems, *Computer Methods in Applied Mechanics and Engineering* 190 (2001) 3247–3270.
- [11] M. Heil, Finite Reynolds number effects in the propagation of an air finger into a liquid-filled flexible-walled channel, *Journal of Fluid Mechanics* 424 (2000) 21–44.
- [12] M. Heil, J. P. White, Airway closure: Surface-tension-driven non-axisymmetric instabilities of liquid-lined elastic rings, *Journal of Fluid Mechanics* 462 (2002) 79–109.
- [13] A. L. Hazel, M. Heil, Three-dimensional airway reopening: The steady propagation of a semi-infinite bubble into a buckled elastic tube, *Journal of Fluid Mechanics* 478 (2003) 47–70.
- [14] O. E. Jensen, M. Heil, High-frequency self-excited oscillations in a collapsible-channel flow, *Journal of Fluid Mechanics* 481 (2003) 235–268.
- [15] A. L. Hazel, M. Heil, Steady finite Reynolds number flow in three-dimensional collapsible tubes (2003).
- [16] G. A. Wempner, *Mechanics of solids with applications to thin bodies*, Sijthoff & Noordhoff, Alphen aan den Rijn, 1981.

- [17] F. K. Bogner, R. L. Fox, L. A. Schmit, A cylindrical shell discrete element, *AIAA Journal* 5 (1967) 645–750.
- [18] M. Heil, T. J. Pedley, Large axisymmetric deformations of cylindrical shells conveying viscous flow, *Journal of Fluids and Structures* 9 (1995) 237–256.
- [19] H. E. Keller, Numerical solution of bifurcation and non-linear eigenvalue problems, in: P. Rabinowitz (Ed.), *Applications of Bifurcation Theory*, Academic Press, New York, 1977, pp. 359–383.
- [20] S. F. Kistler, L. E. Scriven, Coating flows, in: J. Pearson, S. Richardson (Eds.), *Computational Analysis of Polymer Processing*, Applied Science Publishers, London, 1983.
- [21] C. Taylor, P. Hood, A numerical solution of the Navier-Stokes equations using the finite element technique, *Computers and Fluids* 1 (1973) 73–100.
- [22] T. J. R. Hughes (Ed.), *Finite element methods for convection dominated flows*, Vol. 34 of AMD, ASME, New York, 1979.
- [23] T. E. Tezduyar, D. K. Ganjoo, Petrov-Galerkin formulations with weighting functions dependent on spatial and temporal discretizations: Applications to transient convection-diffusion problems, *Computer Methods in Applied Mechanics and Engineering* 59 (1986) 49–71.
- [24] T. E. Tezduyar, Stabilized finite element formulations for incompressible flow computations, *Advances in Applied Mechanics* 28 (1992) 1–44.
- [25] T. E. Tezduyar, Y. Osawa, Finite element stabilisation parameters computed from element matrices and vectors, *Computer Methods in Applied Mechanics and Engineering* 190 (2000) 411–430.
- [26] B. Fischer, A. Ramage, D. J. Silvester, A. J. Wathen, On parameter choice and iterative convergence for stabilised discretisations of advection-diffusion problems, *Computer Methods in Applied Mechanics and Engineering* 179 (1999) 185–202.
- [27] A. Ramage, A multigrid preconditioner for stabilised discretisations of advection-diffusion problems, *Journal of Computational and Applied Mathematics* 101 (1999) 187–203.
- [28] T. J. R. Hughes, A. Brooks, A multi-dimensional upwind scheme with no crosswind diffusion, in: T. J. R. Hughes (Ed.), *Finite element methods for convection dominated flows*, Vol. 34 of AMD, ASME, New York, 1979, pp. 19–35.
- [29] T. J. Pedley, K. D. Stephanoff, Flow along a channel with a time-dependent indentation in one wall: the generation of vorticity waves, *Journal of Fluid Mechanics* 160 (1985) 337–367.
- [30] HSL 2000: A collection of Fortran codes for large scale scientific computation, <http://www.numerical.rl.ac.uk>.

- [31] O. Ghattas, X. Li, A variational finite-element method for stationary nonlinear fluid-solid interaction, *Journal of Computational Physics* 121 (1995) 347–356.
- [32] C. T. Kelley, *Iterative Methods for Linear and Nonlinear Equations*, SIAM, Philadelphia, 1995.
- [33] S. Turek, *Efficient solvers for incompressible flow problems – an algorithmic and computational approach*, Springer, Berlin, Heidelberg, New York, 1999.
- [34] H. C. Elman, Preconditioning the steady-state Navier-Stokes equations with low viscosity, *SIAM Journal of Scientific Computing* 29 (1999) 1299–1316.
- [35] H. C. Elman, D. J. Silvester, A. J. Wathen, *Finite Elements and Fast Iterative Solvers*, Oxford University Press, Oxford, 2003, to appear.
- [36] J. W. Demmel, S. C. Eisenstat, J. R. Gilbert, X. S. Li, J. W. H. R. Liu, A supernodal approach to sparse partial pivoting, *SIAM Journal on Matrix Analysis and Applications* 20 (1999) 720–755.
- [37] W. L. Briggs, V. E. Henson, S. F. McCormick, *A multigrid tutorial*, 2nd Edition, SIAM, Philadelphia, 2000.
- [38] D. J. Silvester, Personal communication (2002).
- [39] R. S. Dembo, S. C. Eisenstat, T. Steihaug, Inexact Newton methods, *SIAM Journal on Numerical Analysis* 19 (1982) 400–408.
- [40] P. M. Gresho, R. L. Lee, Don't suppress the wiggles – they're telling you something, *Computers and Fluids* 9 (1981) 223–253.
- [41] V. Frayss, L. Giraud, S. G. J. Langou, A set of GMRES routines for real and complex arithmetics on high performance computers, CERFACS Technical Report TR/PA/03/3, Public domain software available on www.cerfacs/algor/Softs (2002).

List of Figures

- 1 Sketch of the model problem considered in this study. Viscous fluid is driven through a channel in which part of one wall is replaced by a pre-stressed elastic membrane. 4
- 2 Illustration of the coarsest ($h_{\text{FEM}} = 1/4$) and finest ($h_{\text{FEM}} = 1/16$) fluid meshes used in the computations. 8
- 3 Streamlines and pressure contours for steady flow at $Re = 500$. $p_{\text{ext}} = p_{\text{ext}}^*/E = 1.524$ (top) and $p_{\text{ext}} = 3.247$ (bottom). The fluid pressure is scaled on the viscous scale, $p = p^*a/(\mu U)$. 12
- 4 Pressure-displacement characteristics for steady flow at various Reynolds numbers. The wall deformation is characterised by the vertical wall displacement $V_{\text{ctrl}} = v^2(\zeta = 3.5)$ at 70% of the membrane's length. At large Reynolds number, limit points develop in the pressure-displacement curve. 12
- 5 Time history of the vertical wall displacement $V_{\text{ctrl}}(t) = v^2(\zeta = 3.5, t)$ for a self-excited oscillation at $Re = 500$. $p_{\text{ext}} = p_{\text{ext}}^*/E = 2.5$. 13
- 6 Instantaneous streamlines and pressure contours in an unsteady flow undergoing a large-amplitude self-excited oscillation. $t = 2.0$ (top), 8.0, 12.0, 14.0 (bottom). $Re = 500$, $p_{\text{ext}} = p_{\text{ext}}^*/E = 2.5$. The fluid pressure is scaled on the viscous scale, $p = p^*a/(\mu U)$. 14
- 7 Convergence histories of the Newton iteration with the exact and approximate Jacobian matrices for a steady computation at $Re = 500$. Thin lines: The solution for $V_{\text{ctrl}} = -0.0938$ is used as an initial guess for $V_{\text{ctrl}} = -0.1125$. Thick lines: The solution for $V_{\text{ctrl}} = -0.3938$ is used as an initial guess for $V_{\text{ctrl}} = -0.4125$. 16
- 8 Convergence histories of the preconditioned GMRES iteration for a steady computation at $Re = 500$. The solution for $V_{\text{ctrl}} = -0.21$ is used as the initial guess for $V_{\text{ctrl}} = -0.28$. Thick lines: first Newton step; thin lines: last (=fifth) Newton step. 18

- 9 Convergence histories of the GMRES iteration for different time-steps Δt . \mathcal{J}_{sup} is used as the right preconditioner and the results are for the unsteady simulation of section 4.2 at $t = 14$. $Re = 500$. Thick lines: first Newton step; thin lines: last (=fifth) Newton step. 20
- 10 Convergence histories of the preconditioned GMRES iteration for a steady computation at $Re = 500$. \mathcal{P}_{sup} with fixed-cost multigrid approximations for (33), (35) and (37) is used as the right preconditioner. The solution for $V_{ctrl} = -0.21$ is used as the initial guess for $V_{ctrl} = -0.28$. Thick lines: first Newton step; thin lines: last (=fifth) Newton step. 23
- 11 Convergence histories of the GMRES iteration for different time-steps Δt . \mathcal{P}_{sup} with fixed-cost multigrid approximations for (33), (35) and (37) is used as the right preconditioner and the results are for the unsteady simulation of section 4.2 at $t = 14$. $h_{FEM} = 1/12$, $Re = 500$. Thick lines: first Newton step; thin lines: last (=fifth) Newton step. 24
- 12 Demonstration of the errors introduced by inconsistent stabilisation in steady (a and b) and unsteady (c and d) simulations at $Re = 500$. (a) and (c) document the mesh convergence of the solutions in the absence of stabilisation; Figs. (b) and (d) demonstrate the effect of inconsistent stabilisation on the finest mesh. The parameter values are the same as in section 4. 27
- 13 The effect of inconsistent stabilisation on the velocities and pressures along the line $x_2 = 0.4$ (indicated by the dashed line in the inset) in a steady simulation at $Re = 500$. The wall shape is prescribed and the computation is performed on the finest mesh, $h_{FEM} = 1/16$. 28