

Adaptive Stokes Preconditioning for Steady Incompressible Flows

Cédric Beaume*

School of Mathematics, University of Leeds, Leeds LS2 9JT, United Kingdom.

Received 16 November 2016; Accepted (in revised version) 19 December 2016

Abstract. This paper describes an adaptive preconditioner for numerical continuation of incompressible Navier–Stokes flows based on Stokes preconditioning [42] which has been used successfully in studies of pattern formation in convection. The preconditioner takes the form of the Helmholtz operator $I - \Delta t L$ which maps the identity (no preconditioner) for $\Delta t \ll 1$ to Laplacian preconditioning for $\Delta t \gg 1$. It is built on a first order Euler time-discretization scheme and is part of the family of matrix-free methods. The preconditioner is tested on two fluid configurations: three-dimensional doubly diffusive convection and a two-dimensional projection of a shear flow. In the former case, it is found that Stokes preconditioning is more efficient for $\Delta t = \mathcal{O}(1)$, away from the values used in the literature. In the latter case, the simple use of the preconditioner is not sufficient and it is necessary to split the system of equations into two subsystems which are solved simultaneously using two different preconditioners, one of which is parameter dependent. Due to the nature of these applications and the flexibility of the approach described, this preconditioner is expected to help in a wide range of applications.

AMS subject classifications: 35Q30, 35Q35, 37M20, 37N10, 65F08, 65P30, 76D05, 76E05, 76E06

Key words: Fluid dynamics, hydrodynamic stability, numerical continuation, preconditioning, Stokes preconditioning, doubly diffusive convection, shear flows.

1 Introduction

The development of specialized numerical methods and the increase in available computing resources have helped to make substantial progress in understanding many non-linear problems as dynamical systems. The most basic tool available to that end is time integration which simulates the temporal evolution of an initial condition, thereby emulating an experimental or natural realization. Time integration provides access to the preferred transient and end state, however, it does not (necessarily) provide access to information regarding the origin of these end states. One way to understand how these

*Corresponding author. *Email address:* c.m.l.beaume@leeds.ac.uk (C. Beaume)

states are formed is to compute unstable solutions. Unstable states cannot be obtained (or in some rare cases with great difficulty) using time integration but help provide a complete picture of the dynamical system: they can gain stability or lead to the creation of new solutions or new transients under parametric changes. Numerical continuation has been developed to complement time-integration in that respect and has become an essential part of the toolkit of the nonlinear dynamicist.

Pioneered by Keller [26], these methods help compute steady solutions of a system of ordinary differential equations (ODE) and their evolution as parameter values are changed. They are designed to *continue* a fixed point in parameter space in order to draw its *branch* and unfold the *bifurcation diagram*, thereby revealing its formation mechanisms. Continuation methods typically consist in a two-step algorithm comprising a prediction phase based on previous iterates along the branch and a correction phase involving a fixed point method [39]. Due to their nature, these methods can compute exact solutions regardless of their stability and provide information on the effect of parametric changes on them. Numerical continuation has become a popular tool, broadly used in many different fields [27, 40] and a myriad of packages have been developed and released in the public domain [14, 17, 18, 28, 44].

Fluid dynamics has seen much progress with the help of continuation methods. Intricate pattern formation problems have been elucidated such as that of Rayleigh–Bénard convection rolls in Cartesian [41], cylindrical [11] and spherical shell geometries [20], doubly diffusive convection [10] and free surface binary fluid convection [9]. Spatially localized pattern formation, involving large aspect-ratio domains, has also been investigated with great success: a collection of spatially localized convective states has been found in two-dimensional large aspect-ratio binary fluid convection [33, 34], rotating convection [3, 7] and magnetoconvection [29]. Despite the successful and reliable use of continuation methods in two-dimensional and small three-dimensional domains, the extension to more complex geometries constitutes a major challenge. The most noticeable attempts concern doubly diffusive convection in a three-dimensional domain of square cross section and large transverse direction [5, 8] and porous medium convection in domains extended in two directions [30], each of these problems involving $\mathcal{O}(10^6)$ degrees of freedom. These studies involved long simulations and require a high level of experience in the use of numerical continuation.

Another area of fluid dynamics that has benefited from the developments of numerical continuation is that of transition to turbulence. Shear flows such as plane Couette flow or pipe flow are subcritical flows, i.e., the trivial laminar solution is stable and coexists with turbulence, a state in which the flow displays spatial and temporal complexity, above a threshold value of the parameter. The pioneering discovery of unstable nonlinear solutions in plane Couette flow [35] has drawn a great deal of attention and meticulous studies of this unstable state have provided crucial understanding of transition in small domains [25]. Numerical continuation has also led to the discovery of a number of new solutions [24, 32, 46] that taken together provide a comprehensive picture of transitional phenomena. Similar studies have taken place in other shear flows and hinted at a com-

mon mechanism for transition to turbulence in subcritical shear flows [19, 36, 48]. Lastly, a recent tour de force involving numerical continuation on large domains has revealed families of spatially localized states of different kinds [12, 23, 38].

There exists only a few numerical continuation methods in fluid systems [16, 22, 37, 45] and the method described below possesses several advantages, amongst which are ease of use and flexibility. The aim of this paper is thus to describe this preconditioner for the Navier–Stokes equation governing steady incompressible flows. The basic principles of numerical continuation and the description of the preconditioner are summarized in Section 2. In Section 3, the preconditioner is tested on two examples: three-dimensional coupled convection and a shear flow. A short conclusion terminates the paper.

2 Numerical method

2.1 Continuation methods

We consider the simple dynamical system:

$$\partial_t \mathbf{u} = \mathbf{F}(\mathbf{u}, \lambda), \quad (2.1)$$

where t represents time, $\mathbf{u} \in \mathcal{R}^n$ is the solution vector of dimension n , $\mathbf{F}: \mathcal{R}^{n+1} \rightarrow \mathcal{R}^n$ is a nonlinear operator and λ the continuation or free parameter. We seek solutions that satisfy $\partial_t \mathbf{u} = 0$ or equivalently $\mathbf{F}(\mathbf{u}, \lambda) = 0$. A continuation method consists in parameterizing and continuing the branch of solutions of Eq. (2.1): $\mathcal{B}(s) = (\mathbf{u}(s), \lambda(s))$, where s is the arclength along the branch. This is done in two steps: a prediction based on the extrapolation of previous results along the branch is made and then converged using a fixed point method based on Eq. (2.1).

Prediction can be achieved in different ways and be based on a polynomial extrapolation of previous solutions or include a condition on a norm of the prediction vector (originating from the last solution and pointing to the predicted one). To converge this prediction, an additional condition is required. One popular choice, *fixed parameter continuation*, is to fix the value of λ in Eq. (2.1). This fails in the presence of a saddle-node. To overcome this, the algorithm is complemented with a routine that detects saddle-nodes and switches to another continuation mode where a carefully chosen element u_k of \mathbf{u} is kept fixed, thus allowing λ to be corrected and allowed to continue saddle-nodes. Another popular method, *pseudo arc-length continuation*, consists of adding a condition ensuring that the correction is made along a vector orthogonal to the prediction vector. The basic differences between these two methods are highlighted in Fig. 1. More details on these methods and other types of continuation methods can be found in [26, 39].

The prediction/correction loop just described is usually optimized through the adjustment of the distance between the last solution and the prediction. When this distance is too small, the correction step is too simple and typically takes one or two Newton iterations. The continuation then consists in many small trivial steps that waste computation




Figure 1: Two different types of continuation methods for a scalar unknown u against parameter λ : fixed parameter continuation (a) and pseudo-arclength continuation (b). The former consists in correcting either at a fixed parameter value (see correction from u_{i+1}^p) or at a fixed value of the solution (or one of its components, see correction from u_{i+1}^p). The latter consists in correcting along a direction orthogonal to the prediction direction.

time. When the prediction distance is too large, the correction is tedious and can take a large (non-optimal) number of Newton iterations or simply fail due to the prediction being outside the basin of attraction of the solution for the Newton method. Continuation algorithms are designed to start with a small distance prediction which is then increased (resp. decreased) when the continuation is easy (resp. difficult).

To assess the performance of the preconditioner in Section 3, we use fixed parameter continuation away from saddle-nodes.

2.2 Stokes preconditioning

In general, incompressible fluid flows are modeled using the Navier–Stokes equation together with the continuity equation. Numerical continuation of the discretized version of these equations yields large dynamical systems, often exceeding 10^5 degrees of freedom. The Jacobian resulting from the Newton method used in the continuation method is thus difficult to handle and the use of iterative methods is necessary. In fact, the storage of the Jacobian would require memory space that is hardly affordable[†], so that methods that do not necessitate the evaluation of the Jacobian are preferred. *Matrix-free methods* revolve around the ability to express the product between the Jacobian and an arbitrary vector without ever evaluating the Jacobian itself. Another issue may arise when dealing with incompressible flows: the Jacobian is ill-conditioned when the Laplacian is predominant in its spectrum. Stokes preconditioning has been developed to overcome this issue [31, 42]. As we shall see, this preconditioner is naturally implementable within a matrix free method which makes it a technique of choice for continuation of steady incompressible flows.

[†]The storage of a $10^5 \times 10^5$ double precision matrix takes about 80GB.

Table 1: Preconditioner P and constant c for the two limits and intermediate values of Δt .

$\mathbf{u}^{t+\Delta t} - \mathbf{u}^t = cP^{-1} [N(\mathbf{u}^t) + L\mathbf{u}^t]$		
$\Delta t \ll 1$	$\Delta t = \mathcal{O}(1)$	$\Delta t \gg 1$
$c = \Delta t$	$c = \Delta t$	$c = 1$
$P \rightarrow I$	$P = I - \Delta t L$	$P \rightarrow -L$

We consider the following dynamical system:

$$\partial_t \mathbf{u} = N(\mathbf{u}) + L\mathbf{u}, \quad (2.2)$$

where t is the time, \mathbf{u} is the solution field, $N(\mathbf{u})$ represents the nonlinear term and $L\mathbf{u}$ the linear term. The method is based on the use of the first order Euler scheme, treating the linear part of the right-hand-side implicitly and the nonlinear part explicitly:

$$\frac{\mathbf{u}^{t+\Delta t} - \mathbf{u}^t}{\Delta t} = N(\mathbf{u}^t) + L\mathbf{u}^{t+\Delta t}, \quad (2.3)$$

where Δt is the timestep and \mathbf{u}^t is the evaluation of \mathbf{u} at time t . On expressing $\mathbf{u}^{t+\Delta t}$, we get:

$$\mathbf{u}^{t+\Delta t} = (I - \Delta t L)^{-1} [\mathbf{u}^t + \Delta t N(\mathbf{u}^t)]. \quad (2.4)$$

The preconditioner is obtained by subtracting \mathbf{u}^t from $\mathbf{u}^{t+\Delta t}$:

$$\begin{aligned} \mathbf{u}^{t+\Delta t} - \mathbf{u}^t &= (I - \Delta t L)^{-1} [\mathbf{u}^t + \Delta t N(\mathbf{u}^t)] - \mathbf{u}^t \\ &= \Delta t (I - \Delta t L)^{-1} [N(\mathbf{u}^t) + L\mathbf{u}^t], \end{aligned} \quad (2.5)$$

where the right-hand-side is the evaluation at time t of the right-hand-side of Eq. (2.2) premultiplied by $\Delta t (I - \Delta t L)^{-1}$.

The general form for the preconditioned equation is then:

$$\mathbf{u}^{t+\Delta t} - \mathbf{u}^t = cP^{-1} [N(\mathbf{u}^t) + L\mathbf{u}^t], \quad (2.6)$$

where c is a constant and P the preconditioner. The values taken by these quantities are summarized in Table 1 as a function of Δt . For $\Delta t \ll 1$, P approaches the identity operator and the equation is solved without effective preconditioning. On the other hand, $\Delta t \gg 1$ yields $(I - \Delta t L)^{-1} \approx -(\Delta t L)^{-1}$ and expression (2.5) becomes:

$$\mathbf{u}^{t+\Delta t} - \mathbf{u}^t \approx -L^{-1} [N(\mathbf{u}^t) + L\mathbf{u}^t]. \quad (2.7)$$

This limit of Stokes preconditioning was first used by Mamun & Tuckerman [31] to study symmetry breaking instabilities in spherical Couette flow. It proved efficient even at relatively large Reynolds numbers[‡] [31]. Stokes preconditioning has been nearly exclusively

[‡]The authors computed solutions up to $Re = 2200$.

used to study pattern formation in problems involving convection [2–5, 7–11, 20, 29, 33, 34, 41]. However, it has been used very little to study shear flows. In particular, Stokes preconditioning has been reported unsuccessful for the computation of steady states of the three-dimensional flow past a step [1]. Other configurations cannot be solved without a well-designed preconditioner [6]. To address this issue, Stokes preconditioning is extended to non-asymptotic values of Δt . The resulting preconditioner, $P = (I - \Delta t L)$, with $c = \Delta t$ and such that neither $\Delta t \ll 1$ nor $\Delta t \gg 1$ holds, proves efficient for a wide range of systems, as shown in Section 3.

To search for stationary flow solutions, we consider the following Newton method:

$$cP^{-1}[\delta N(\mathbf{u}) + L]\delta \mathbf{u} \approx cP^{-1}(N(\mathbf{u}) + L\mathbf{u}), \quad (2.8)$$

with correction $\mathbf{u} = \mathbf{u} - \delta \mathbf{u}$. The right-hand-side of Eq. (2.8) is evaluated using Eq. (2.5) and the left-hand-side of Eq. (2.8) is computed using a linearized version of the same equation. The inversion of the Jacobian is performed using iterative methods such as the biconjugate gradient stabilized method [15].

3 Examples

In this section, two examples of continuation of incompressible fluid flows are considered. Numerical continuation is performed using the preconditioner presented in Section 2.2. The first flow considered is doubly diffusive convection, in which the preconditioner is tested in a three-dimensional configuration. The second test flow is a two-dimensional model of shear flow in which a parameter dependent use of the preconditioner is prescribed and for which the system of equations is split into two sets preconditioned in different manners but solved simultaneously.

3.1 Doubly diffusive convection

We consider a Boussinesq fluid consisting of two components, the heavier of which is referred to as a salt. The fluid is placed within a three-dimensional enclosure of square horizontal cross-section and aspect ratio $L = 19.8536$. The flow is driven by buoyancy through the imposition of large scale horizontal gradients of temperature and concentration: one vertical wall is maintained at a higher temperature and salinity than the opposite one. The other walls are modeled using no flux boundary conditions for the temperature and concentration. No-slip boundary conditions are imposed on all walls.

The nondimensional equations governing the dynamics of this flow are:

$$Pr^{-1}[\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u}] = -\nabla p + Ra(T - C)\hat{\mathbf{x}} + \nabla^2 \mathbf{u}, \quad (3.1a)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (3.1b)$$

$$\partial_t T + (\mathbf{u} \cdot \nabla) T = \nabla^2 T, \quad (3.1c)$$

$$\partial_t C + (\mathbf{u} \cdot \nabla) C = \tau \nabla^2 C, \quad (3.1d)$$

where t is the time, $\mathbf{u} = (u, v, w)$ is the velocity field in the Cartesian frame $(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}})$, p is the pressure and T and C are linear rescalings of the fluid temperature and salt concentration in the Boussinesq approximation. Here, $\hat{\mathbf{x}}$ represents the vertical unit vector in the ascending direction. In addition to these quantities, three nondimensional parameters are introduced. The Prandtl number $Pr = 1$ is the ratio of the kinematic viscosity over the thermal diffusivity, the inverse Lewis number $\tau = 1/11$ is the ratio of the salt diffusivity over the thermal diffusivity and the Rayleigh number Ra quantifies the buoyancy strength and will thereafter be used as the continuation parameter. These equations are complemented with boundary conditions:

$$\text{at } x = \{0, L\} \quad \text{or } y = \{0, 1\}: \quad u = v = w = \partial_n T = \partial_n C = 0, \quad (3.2a)$$

$$\text{at } z = 0: \quad u = v = w = T = C = 0, \quad (3.2b)$$

$$\text{at } z = 1: \quad u = v = w = T - 1 = C - 1 = 0, \quad (3.2c)$$

where the operator ∂_n represents the spatial derivative in the direction normal to the wall. More details on the physical setup are available in [5].

3.1.1 Numerics

The physical domain is meshed using 16 identical spectral elements of size $l_x \approx 1.24$, $l_y = l_z = 1$. Each element is meshed using Gauss–Lobato–Legendre points in all three directions: 21 in x , 19 in y and 19 in z . The discretization strategy is illustrated in Fig. 2 and yields $16 \times 21 \times 19 \times 19 \times 6 = 727,776$ degrees of freedom (counting in the pressure).

Time is discretized using the first order Euler scheme of Eq. (2.3):

$$T^{(n)} = (I - \Delta t \nabla^2)^{-1} \left(T^{(n-1)} - \Delta t [(\mathbf{u} \cdot \nabla) T]^{(n-1)} \right), \quad (3.3a)$$

$$C^{(n)} = (I - \Delta t \tau \nabla^2)^{-1} \left(C^{(n-1)} - \Delta t [(\mathbf{u} \cdot \nabla) C]^{(n-1)} \right), \quad (3.3b)$$

where $T^{(n)}$ stands for the evaluation of the temperature at the n -th timestep, I is the identity operator and Δt is the timestep.

The incompressible Navier–Stokes equation is discretized using a standard first order splitting method. An intermediate velocity is predicted which takes into account buoyancy and advection:

$$\hat{\mathbf{u}} = \mathbf{u}^{(n-1)} - \Delta t [(\mathbf{u} \cdot \nabla) \mathbf{u}]^{(n-1)} + \Delta t Ra (T - C)^{(n-1)} \hat{\mathbf{x}}. \quad (3.4)$$

It is corrected using the incompressibility condition by introducing the velocity $\hat{\hat{\mathbf{u}}}$:

$$\hat{\hat{\mathbf{u}}} = \hat{\mathbf{u}} - \Delta t \nabla p^{(n)}, \quad (3.5)$$




Figure 2: Example of discretization for the doubly diffusive convection problem. For the sake of the representation, only 6 out of the 16 elements are shown with only the first (bottom left) element displaying its surface mesh. The inner grid is hidden to avoid overloading the figure.

where the pressure is defined by the Poisson problem obtained by taking the divergence of Eq. (3.5):

$$\nabla^2 p^{(n)} = \frac{1}{\Delta t} \nabla \cdot \hat{\mathbf{u}} \tag{3.6}$$

complemented with the boundary condition:

$$\partial_n p^{(n)} = \left((T - C)^{(n-1)} \hat{\mathbf{x}} - [(\mathbf{u} \cdot \nabla) \mathbf{u}]^{(n-1)} - \nabla \times \nabla \times \mathbf{u}^{(n-1)} \right) \cdot \hat{\mathbf{n}}, \tag{3.7}$$

where $\hat{\mathbf{n}}$ represents the vector normal to the boundary and the last term is the reduction of the Laplacian term using the incompressibility condition. The time-step is completed via the following operation:

$$\mathbf{u}^{(n)} = (I - \Delta t \nabla^2)^{-1} (\Delta t \hat{\mathbf{u}}), \tag{3.8}$$

where the original boundary conditions (3.2)-(3.2b) are used.

Time-stepping then only requires the inversion of Helmholtz operators coming from the spatial discretization. As a consequence of the use of spectral elements, the Helmholtz operators yield sparse tensors and we use a Schur decomposition to invert them efficiently. Continuation is performed on the solution vector $(\mathbf{u}, T, C) = (u, v, w, T, C)$ using the temporal schemes above and the same Δt for all equations.

3.1.2 Results

This flow configuration exhibits localized pattern formation at onset through a subcritical bifurcation. The branches emerging from this bifurcation produce well-bounded back and forth oscillations in parameter space in a behavior known as *snaking*. More detailed information is available in [5].

We focus here on one branch of spatially localized states. The branch is shown in Fig. 3. To design a test for the preconditioner, we select two segments along the branch

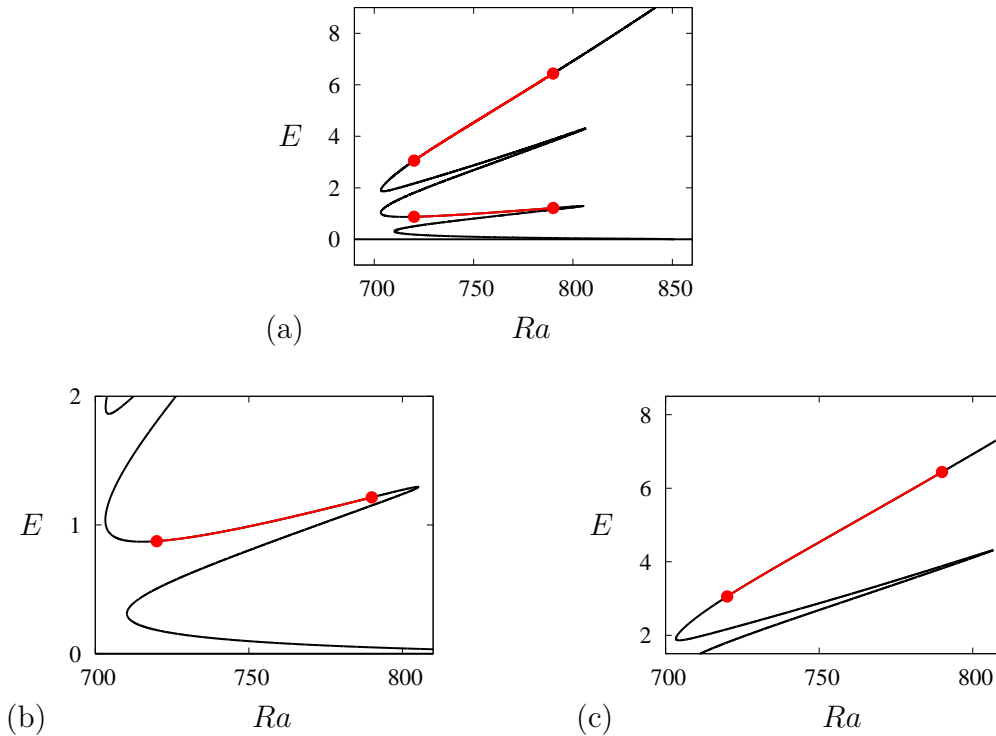


Figure 3: (a) Bifurcation diagram showing the kinetic energy E of the trivial motionless fluid branch (horizontal line) and of one of the localized state branches as a function of the Rayleigh number Ra . (b) Enlargement of the first region of interest: the test of the preconditioner is run through leftward continuation along the red portion of the branch delimited by the dots. The solutions at the dots are represented in Figs. 4(a), (b). (c) Enlargement of the second region of interest where the tests is run rightwards. The solutions at the dots are shown in Figs. 4(c), (d). The red segments are also reported in panel (a) for consistency.

and use fixed parameter continuation to avoid intricacies related to continuation with a variable parameter. One of these segments consists of solutions that are spatially localized in which most of the domain has $u \approx v \approx w \approx T \approx C \approx 0$ while the other segment consists of domain-filling solutions. The saddle-nodes occur at $Ra \approx 703$ and $Ra \approx 807$ so we restrict our simulations to $720 < Ra < 790$. The solutions at the extrema of both segments are shown in Fig. 4. These solutions display a number of similar convection rolls centered in the domain. In each of these rolls, the fluid goes up along the hot and saltier wall at $z=1$ and down along the opposite wall. The flow is not purely two-dimensional: due to the presence of walls at $y=0$ and $y=1$, a weak flow in the y direction is generated. This flow is typically an order of magnitude lower than in the two other directions and is therefore not shown here. As the branch is continued along the lower segment, the Rayleigh number is decreased and the solution changes from the right panels to the left panels of Figs. 4(a), (b), thereby adding one roll on either side of the central roll. Along the upper segment, the solution is continued in the direction of increasing Rayleigh numbers




Figure 4: (a), (b) Solutions at the extrema of the bottom test segment, shown in Figs. 3(b). (c), (d) Solutions at the extrema of the top test segment, shown in Fig. 3(c). The solutions are represented using isosurfaces of the x -velocity $u = \pm 0.5$ (a), (c) and of the z -velocity $w = \pm 0.2$ (b), (d). In each panel, the left solution is taken at $Ra \approx 720$ and the right solution at $Ra \approx 790$. The light (resp. dark) color indicates the positive (resp. negative) contour.

(from the left to the right panel in Figs. 4(c), (d)) and the rolls grow in size and amplitude without any nucleation or annihilation.

The algorithmic parameters are kept at the values used during the original study [5] and which were determined using a combination of intuition and trial and error: the tolerance of the BCGStab is fixed at 10^{-2} , the continuation step is initialized at $\Delta Ra = 10^{-3}$ and allowed to increase to 1. The convergence of the solution is assessed by calculating the following quantity:

$$\mathcal{L}(\mathbf{u}) = \left| (1 + \Delta t) (I - \Delta t L)^{-1} [N(\mathbf{u}) + L\mathbf{u}] \right|_{L_2}, \quad (3.9)$$

computed by stepping forward once in time (see Eq. (2.6)), multiplying by $(1 + \Delta t) / \Delta t$ and then taking the L_2 -norm of the resulting vector. The aforementioned scalar multiple is taken so that the relative convergence of $N(\mathbf{u}) + L\mathbf{u}$ does not depend strongly on the value of Δt . If more than 4 Newton iterations are necessary to obtain $\mathcal{L}(\mathbf{u}) < 10^{-7}$, the continuation step remains untouched, otherwise ΔRa is multiplied by 1.2. If the Newton iteration fails to converge in 9 iterations or if the number of conjugate gradient iterations needed to invert the Jacobian is greater than 5000, the current step is cancelled and ΔRa decreased by 10%.

A number of attempts were made for both continuation segments, with Δt ranging from 10^{-4} to 10^8 . The smallest values of Δt did not allow the algorithm to converge, implying that this problem does indeed need preconditioning. I report here the successful simulations, for $\Delta t \geq 10^{-3}$. The basic algorithmic behavior is illustrated in Fig. 5 on sim-




Figure 5: (a) Evolution of the parameter step ΔRa as a function of the solution index i for $\Delta t=0.06$. Most of the other simulations yielded the same results. The limiting value $\Delta Ra=1$ is imposed. (b) Number of conjugate gradient iterations η needed to converge as a function of the solution index i for $i \geq 50$ and for $\Delta t=0.003$ (upper curve displaying large variations), $\Delta t=0.06$ (lowest curve) and $\Delta t=10^5$ (upper curve with small variations). These results have been obtained during the continuation of the lower segment from Fig. 3(b).

ulations carried out on the lower segment in Fig. 3(b). Panel (a) shows the continuation step as a function of the solution index i and reports a steady acceleration until $\Delta Ra=1$ is reached, indicating that ΔRa could potentially be increased further. Nearly all the simulations run provided the same results here. The few exceptions were generally obtained for $\Delta t \leq 0.01$ for which the preconditioner is nearly the identity. These cases display a few solutions at which more than 4 Newton iterations are made and where ΔRa is not increased. Fig. 5(b) shows, for three different values of Δt , the typical values taken by η , the total number of conjugate gradient iterations needed to converge a solution (summed up on all the Newton steps required). These results are reported for $i \geq 50$ where continuation is done with $\Delta Ra=1$. For $\Delta t=0.003$, the convergence speed displays large variations, with η varying between 265 and 1062 and averaging $\bar{\eta} \approx 471$ with a standard deviation of $\sigma \approx 179$ over the last 50 solutions. Results at $\Delta t=10^5$ do not suffer from such large oscillations, $\sigma \approx 32$ and their average is significantly lower: $\bar{\eta} \approx 349$ but the best results were obtained for $\Delta t=0.06$ (lowest curve in Fig. 5(b)): $\bar{\eta} \approx 140$ with $\sigma \approx 17$.

The results of all the simulations are compiled in Fig. 6 and presented through the average number of conjugate gradient iterations per solution $\bar{\eta}$ and the normalized standard deviation from this result $\sigma/\bar{\eta}$ for the last 50 solutions. Fig. 6(a) clearly indicates the presence of three distinct regions where the preconditioner behaves differently. The first region is defined for $\Delta t < 0.01$: for such low values of Δt , there is effectively little preconditioning in place and the continuation struggles as shown by the large increase of $\bar{\eta}$ as Δt decreases. For $\Delta t \leq 0.001$, continuation is impossible, confirming the need for a preconditioner. A second region of interest is obtained at large Δt . In fact, little difference is observed for simulations run with $\Delta t \geq 10^2$, indicating that the asymptotic regime is reached. There, converging a solution costs about 350 conjugate gradient iterations for the lower segment and 455 for the upper one. Between these two regions lies an interval located around $\Delta t=0.06$ where the preconditioner is at its peak performance. The




Figure 6: Compilation of the results showing (a) the average number of conjugate gradient iterations $\bar{\eta}$ needed to converge the Newton method on a sample consisting of the last 50 solutions as a function of Δt . The right panel (b) shows the standard deviation σ associated with these simulations normalized by $\bar{\eta}$ as a function of Δt . In both panels, the red dashed curve corresponds to the upper continuation segment results while the blue dashed curve corresponds to the lower continuation segment results. The thick black curve indicates the average between both segments.

most efficient continuation of the lower segment was obtained for $\Delta t = 0.06$ with an average of $\bar{\eta} \approx 140$ conjugate gradient iterations per solution. Similarly, the most efficient continuation of the upper segment is found for $\Delta t = 0.03$ and $\bar{\eta} \approx 253$.

Fig. 6(b) shows the standard deviation associated with the number of conjugate gradient iterations normalized by $\bar{\eta}$ to quantify the variation significance. The results indicate that the asymptotic regime of large Δt yields a steady-paced continuation with a relative standard deviation of about 10% of $\bar{\eta}$ for $\Delta t \geq 10^2$. As Δt decreases, $\sigma/\bar{\eta}$ increases to reach values above 1 for the smallest successful Δt . This highlights one particular characteristic: although the preconditioner is at its best in sheer performance for values of Δt around 0.06, continuation does not progress at such a steady pace as for large Δt . This accounts for the fact that the results presented in Fig. 6(a) are smooth at large Δt but display some anomalies at lower Δt . In particular, the anomaly reported at $\Delta t = 0.1$ along the upper segment is due to a solution necessitating more than 5000 conjugate gradient iterations.

One may wish to bound the continuation step, as we have done above with $\Delta Ra = 1$, in order to obtain a satisfactory discretization of the solution branch and to avoid jumping onto another branch in a case of an imperfect bifurcation. However, a most economical calculation of the branch can be obtained when the continuation step is unbounded. Fig. 7 shows the total number of conjugate gradient iterations undergone from the startpoint until the endpoint of the continuation segments with Δt unbounded. These results are very similar to those shown in Fig. 6(a) and confirm the previous observations.

3.2 Shear flow

We now consider a three-dimensional fluid confined between two parallel plates of infinite extent. The flow is driven by the imposition of a sinusoidal volume force creating a shear across the fluid layer in a configuration known as *plane Waleffe flow*. The Navier–




Figure 7: Total number of conjugate gradient iterations η_{tot} needed to complete the lower (upper) continuation segment in dashed blue (red) lines versus the algorithmic parameter Δt . The solid black line represents the average between the two data sets.

Stokes equation with the incompressibility constraint for this configuration read:

$$\partial_t \mathbf{u} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} + \frac{\sqrt{2}\pi^2}{4Re} \sin\left(\frac{\pi y}{2}\right) \hat{\mathbf{x}}, \quad (3.10a)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (3.10b)$$

where t is the time, $\mathbf{u} = (u, v, w)$ the velocity field in the (x, y, z) coordinate frame where x is the streamwise direction, y the wall-normal direction and z the spanwise direction, p is the pressure and Re is the Reynolds number which quantifies the imposed shear across the fluid. These equations are accompanied with periodic boundary conditions in x and z and free-slip boundary conditions in y :

$$\partial_y u = v = \partial_y w = 0 \quad \text{at } y = \pm 1. \quad (3.11)$$

This flow configuration is a close cousin of plane Couette flow and is studied to investigate transition to turbulence. A number of studies have revealed the influence of exact coherent states in the transition process [13, 24, 25, 35, 46, 47]. These states are exact solutions of the associated system of equations and some of them follow an asymptotic behavior as the Reynolds number is increased [47]. To approximate these solutions, we apply the following asymptotics for $\epsilon = Re^{-1} \ll 1$:

$$u(x, y, z, t) \sim u_0(y, z, T) + \epsilon \left(u_1(y, z, T) + u'_1(y, z, t, T) e^{i\alpha x} + c.c. \right), \quad (3.12a)$$

$$v(x, y, z, t) \sim \epsilon \left(v_1(y, z, T) + v'_1(y, z, t, T) e^{i\alpha x} + c.c. \right), \quad (3.12b)$$

$$w(x, y, z, t) \sim \epsilon \left(w_1(y, z, T) + w'_1(y, z, t, T) e^{i\alpha x} + c.c. \right), \quad (3.12c)$$

where u_0 , u_1 , v_1 and w_1 are real, u'_1 , v'_1 and w'_1 are complex, $T = \epsilon t$ is the slow time, α is the chosen wavelength in the streamwise direction and $c.c.$ represents the complex

conjugate. The pressure is expanded accordingly and a streamfunction ϕ_1 and a vorticity ω_1 are introduced such that: $v_1 = -\partial_z \phi_1$, $w_1 = \partial_y \phi_1$ and $\omega_1 = \nabla_{\perp}^2 \phi_1$.

The three-dimensional system (3.10a), (3.10b) reduces down to the following two-dimensional system:

$$\partial_T u_0 + J(\phi_1, u_0) = \nabla_{\perp}^2 u_0 + \frac{\sqrt{2}\pi^2}{4} \sin\left(\frac{\pi y}{2}\right), \quad (3.13a)$$

$$\partial_T \omega_1 + J(\phi_1, \omega_1) + 2(\partial_y^2 - \partial_z^2)(\mathcal{R}(v_1' w_1'^*)) + 2\partial_y \partial_z (w_1' w_1'^* - v_1' v_1'^*) = \nabla_{\perp}^2 \omega_1, \quad (3.13b)$$

$$(\alpha^2 - \nabla_{\perp}^2) p_1' = 2i\alpha (v_1' \partial_y u_0 + w_1' \partial_z u_0), \quad (3.13c)$$

$$\partial_t \mathbf{v}'_{1\perp} + i\alpha u_0 \mathbf{v}'_{1\perp} = -\nabla_{\perp} p_1' + \epsilon \nabla_{\perp}^2 \mathbf{v}'_{1\perp}, \quad (3.13d)$$

where $J(\phi_1, \cdot) = \partial_y \phi_1 \partial_z \cdot - \partial_z \phi_1 \partial_y \cdot$, $\nabla_{\perp} = (\partial_y, \partial_z)$, $\nabla_{\perp}^2 = \partial_y^2 + \partial_z^2$, $\mathcal{R}(\cdot)$ indicates the real part of \cdot , \cdot^* is the complex conjugate of \cdot and i is the unit imaginary number. The fields u_0 , ϕ_1 and ω_1 are real while $\mathbf{v}'_{1\perp} = (v_1', w_1')$ and p_1' are complex. The boundary conditions read:

$$\partial_y u_0 = \omega_1 = \phi_1 = v_1' = \partial_y w_1' = 0 \quad \text{at } y = \pm 1, \quad (3.14)$$

together with periodic boundary conditions in z . For more details on the derivation, see Beaume et al. [6].

3.2.1 Numerics

The physical domain is two-dimensional and has size $L_y = 2$ and $L_z = \pi$. It is meshed with 32 equidistributed points in each direction and the linear operators treated using the Fast Fourier Transform in z and either the Fast Cosine Transform I or the Fast Sine Transform I in y depending on the boundary condition [21]. The usual 2/3 dealiasing rule is applied to prevent frequency folding. The resulting number of unfiltered degrees of freedom is then: $32 \times 32 \times 8 = 8,192$.

The fluctuating pressure p_1' is solved for in a preliminary step:

$$p_1'^{(n-1)} = 2i\alpha (\alpha^2 - \nabla_{\perp}^2)^{-1} \left(v_1'^{(n-1)} \partial_y u_0^{(n-1)} + w_1'^{(n-1)} \partial_z u_0^{(n-1)} \right), \quad (3.15)$$

where the nonlinear right-hand-side is evaluated in physical space and the linear operator inverted in frequency space.

The same first order Euler scheme as for the doubly diffusive convection problem is used to treat time dependence in the remaining equations:

$$u_0^{(n)} = (I - \epsilon \Delta t_1 \nabla_{\perp}^2)^{-1} \left[u_0^{(n-1)} + \epsilon \Delta t_1 \left(-J(\phi_1^{(n-1)}, u_0^{(n-1)}) + \frac{\sqrt{2}\pi^2}{4} \sin\left(\frac{\pi y}{2}\right) \right) \right], \quad (3.16a)$$

$$\omega_1^{(n)} = (I - \epsilon \Delta t_1 \nabla_{\perp}^2)^{-1} \left[\omega_1^{(n-1)} + \epsilon \Delta t_1 \left(-J(\phi_1^{(n-1)}, \omega_1^{(n-1)}) + 2(\partial_z^2 - \partial_y^2)(\mathcal{R}(v_1'^{(n-1)} w_1'^{(n-1)*})) + 2\partial_y \partial_z (v_1'^{(n-1)} v_1'^{(n-1)*} - w_1'^{(n-1)} w_1'^{(n-1)*}) \right) \right], \quad (3.16b)$$

$$\mathbf{v}'_{1\perp}{}^{(n)} = (I - \epsilon \Delta t_2 \nabla_{\perp}^2)^{-1} \left(\mathbf{v}'_{1\perp}{}^{(n-1)} + \Delta t_2 \left(-i\alpha u_0^{(n-1)} \mathbf{v}'_{1\perp}{}^{(n-1)} - \nabla_{\perp} p_1'^{(n-1)} \right) \right), \quad (3.16c)$$

where Δt_1 and Δt_2 have been introduced as two non necessarily equal parameters for continuation. The preconditioning method is then written as follows:

$$u_0^{(n)} - u_0^{(n-1)} = \epsilon \Delta t_1 (I - \epsilon \Delta t_1 \nabla_{\perp}^2)^{-1} \left[\nabla_{\perp}^2 u_0^{(n-1)} + N_{11} \left(u_0^{(n-1)}, \phi_1^{(n-1)} \right) \right], \quad (3.17a)$$

$$\begin{aligned} \omega_1^{(n)} - \omega_1^{(n-1)} = \epsilon \Delta t_1 (I - \epsilon \Delta t_1 \nabla_{\perp}^2)^{-1} \\ \times \left[\nabla_{\perp}^2 \omega_1^{(n-1)} + N_{12} \left(\omega_1^{(n-1)}, \phi_1^{(n-1)}, \mathbf{v}'_{1\perp}{}^{(n-1)} \right) \right], \end{aligned} \quad (3.17b)$$

$$\begin{aligned} \mathbf{v}'_{1\perp}{}^{(n)} - \mathbf{v}'_{1\perp}{}^{(n-1)} = \Delta t_2 (I - \epsilon \Delta t_2 \nabla_{\perp}^2)^{-1} \\ \times \left[\epsilon \nabla_{\perp}^2 \mathbf{v}'_{1\perp}{}^{(n-1)} + N_2 \left(\mathbf{v}'_{1\perp}{}^{(n-1)}, u_0^{(n-1)}, p_1'^{(n-1)} \right) \right], \end{aligned} \quad (3.17c)$$

for which

$$N_{11} \left(u_0^{(n-1)}, \phi_1^{(n-1)} \right) = -J \left(\phi_1^{(n-1)}, u_0^{(n-1)} \right) + \frac{\sqrt{2}\pi^2}{4} \sin \left(\frac{\pi y}{2} \right), \quad (3.18a)$$

$$\begin{aligned} N_{12} \left(\omega_1^{(n-1)}, \phi_1^{(n-1)}, \mathbf{v}'_{1\perp}{}^{(n-1)} \right) \\ = -J \left(\phi_1^{(n-1)}, \omega_1^{(n-1)} \right) - 2(\partial_y^2 - \partial_z^2) \left(\mathcal{R} \left(v_1'^{(n-1)} w_1'^{* (n-1)} \right) \right) \\ - 2\partial_y \partial_z \left(w_1'^{(n-1)} w_1'^{* (n-1)} - v_1'^{(n-1)} v_1'^{* (n-1)} \right), \end{aligned} \quad (3.18b)$$

$$N_2 \left(\mathbf{v}'_{1\perp}{}^{(n-1)}, u_0^{(n-1)}, p_1'^{(n-1)} \right) = -i\alpha u_0^{(n-1)} \mathbf{v}'_{1\perp}{}^{(n-1)} - \nabla_{\perp} p_1'^{(n-1)}, \quad (3.18c)$$

and for which $p_1'^{(n-1)}$ has already been evaluated in Eq. (3.15) and $\phi_1^{(n-1)} = (\nabla_{\perp}^2)^{-1} \omega_1^{(n-1)}$. The convergence criterion \mathcal{L} is obtained on the L_2 -norm of the right hand side of Eqs. (3.17a)-(3.17c) multiplied by $(1 + \epsilon \Delta t_i) / \Delta t_i$, where $i = 1$ for Eqs. (3.17a), (3.17b) and $i = 2$ for Eq. (3.17c).

Continuation is carried out on the vector $(u_0, \omega_1, \mathbf{v}'_{1\perp}) = (u_0, \omega_1, v_1', w_1')$. The presence of a small parameter in front of the linear operator in Eq. (3.13d) and the absence of scalar multiplier in front of that in Eqs. (3.13a), (3.13b) while all the other terms are $\mathcal{O}(1)$ leads to a different balance in these equations. In fact, as we shall see, they require different preconditioners that we can obtain by setting $\Delta t_1 \neq \Delta t_2$.

3.2.2 Results

The trivial solution of plane Waleffe flow $(u_0, \omega_1, v_1', w_1') = (\frac{\sqrt{2}\pi^2}{4} \sin(\frac{\pi y}{2}), 0, 0, 0)$ is linearly stable for all values of the Reynolds number which does not prevent the existence of other nonlinear solutions. These solutions are formed at saddle-node bifurcations at finite Re and take the form of upper and lower branches, the former being energetically farther from the trivial solution than the latter. We focus here on the most basic of these solutions obtained for a domain of size $L_x = 4\pi$ (equivalently $\alpha = 0.5$), $L_y = 2$ and $L_z = \pi$. The bifurcation diagram is shown in Fig. 8. The solution is formed at a saddle-node at $Re \approx 136$ and splits into a lower branch state, shown in Fig. 9, and an upper branch state, shown in




Figure 8: Bifurcation diagram representing twice the kinetic energy per unit volume associated with velocity u_0 : N_u versus the Reynolds number Re . The quantity N_u is defined as follows: $N_u = D^{-1} \int_{\mathcal{D}} u_0^2 dy dz$ with $D = \int_{\mathcal{D}} dy dz$ and $\mathcal{D} = [-1;1] \times [0;L_z]$ represents the domain of integration. The trivial solution of plane Waleffe flow, $N_u=1$, is not shown. The lower (resp. upper) branch solution is labeled L (resp. U).

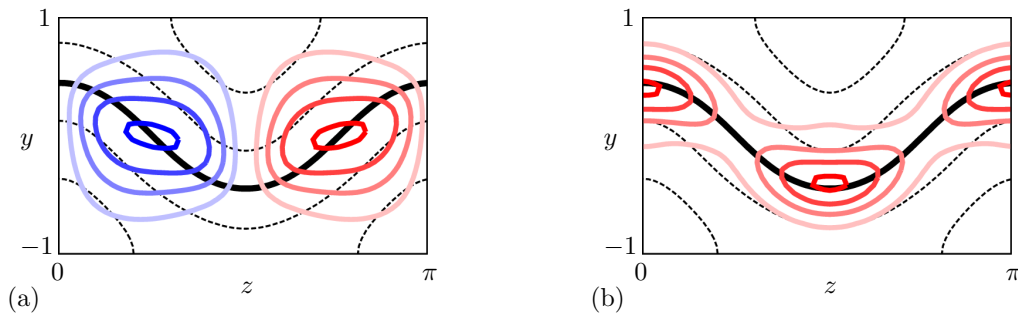


Figure 9: Solution from the lower branch L in Fig. 8 taken at $Re \approx 1000$. It is represented through equidistributed streamfunction ϕ_1 contours (with increments of 0.4) (a) and fluctuation amplitude $|(v'_1, w'_1)|$ contours (with increments of 1.75) (b) in the (y, z) -plane. Positive (negative) quantities are represented in red (blue) and are plotted on top of the equidistributed contours of streamwise-invariant streamwise velocity u_0 (with increments of 0.5) in black, with the thick solid black line representing the critical layer where $u_0 = 0$.

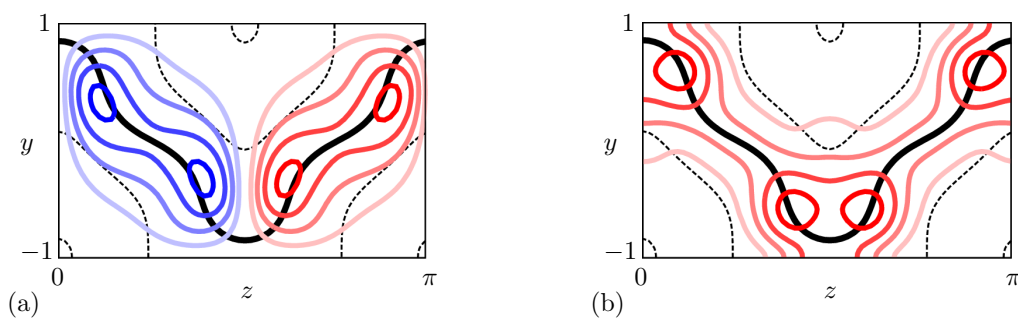


Figure 10: Same representation as in Fig. 9 but for a solution taken along the upper branch U at $Re \approx 1000$. The contours are equidistributed with increments of 0.7 for the streamfunction (a) and 2.5 for the fluctuation amplitude (b).

Fig. 10, both for $Re \approx 1000$. As the Reynolds number is increased along these branches, the pattern remains similar but the fluctuations become sharper.

The continuation code from [6] is modified to output relevant data for algorithmic comparison but no changes to the computational part of the code is made. In particular, the algorithmic constants are kept the same: the tolerance for the BCGStab iterations is 10^{-2} with a maximum number of iterations of 1000 before failure is declared and the Newton method is considered to have converged when $\mathcal{L} < 10^{-8}$.

Continuation was first tried with $\Delta t_1 = \Delta t_2$ but failed. A critical observation for a successful continuation method here is that the linear operator in Eq. (3.13d) is multiplied by a small quantity. Its influence on the Jacobian compared to that of the nonlinear term is weaker and we can anticipate that Δt_2 needs to be small and, thus, different from Δt_1 .

Much like for the doubly diffusive convection equations, the linear operator of the mean equations (3.13a), (3.13b) is not premultiplied. Given the results of Section 3.1, we assume that an intermediate preconditioner will be efficient here, so we set $\Delta t_1 = Re$ such that the term $I - \epsilon \Delta t_1 \nabla_{\perp}^2$ in Eqs. (3.17a), (3.17b) becomes $I - \nabla_{\perp}^2$. Other values of Δt_1 , including asymptotically large ones, have been tested unsuccessfully.

To investigate the optimal preconditioner, we set up a number of simulations consisting in computing one continuation step with $\Delta Re = 1$. These simulations are carried out for a range of Δt_2 with all other parameters fixed. The number of conjugate gradient iterations needed to converge is then recorded and reported in Fig. 11 for some representative simulations. We observe that the conjugate gradient fails to converge for small values of Δt_2 . Despite the small coefficient in front of the Laplacian, Eq. (3.13d) needs preconditioning. For large Δt_2 , the method also fails. More precisely, no simulation converged for $\Delta t_2 > 10^5$ and only a few successful events were recorded for $\Delta t_2 > 500$ and $Re = 500$. The asymptotic regime (Δt_2 arbitrarily large) of Stokes preconditioning is hence not applicable here.

For relatively low values of Re , the continuation is rather forgiving and a wide range of Δt_2 can be used successfully. Increasing the Reynolds number has several effects. First, even for the optimal Δt_2 , the number of conjugate gradient iterations increases. The condition on Δt_2 for a successful continuation also becomes more stringent. The upper bound of the successful Δt_2 interval is about 10^5 at $Re = 500$ and decreases down to $\Delta t_2 \approx 100$ at $Re = 2000$. The lower bound displays more irregularities, with values around $\Delta t_2 \approx 0.02$ for $Re = 500$ and $Re = 1500$. Nonetheless, a trend might be observed on Figs. 11(a)-(d): the lower bound seems to increase slowly until $\Delta t_2 \approx 0.4$ for $Re = 3500$. As a result, the Δt_2 interval for which the continuation method works becomes narrower as Re increases. These tendencies are confirmed in Fig. 12 where some further characterisations of the success of the continuation as a function of Δt_2 are reported. In addition to the parameter interval in which the continuation method works, the figure indicates for which parameter values it is efficient. This interval corresponds to the set of Δt_2 for which η is at most twice as large as the optimal number of conjugate gradient iteration at that Reynolds number, η_{opt} . Knowledge of this interval is important as a factor of 5 in computation time can be observed between results at the optimum Δt_2 and those on




Figure 11: (a) Number η of conjugate gradient iterations needed to converge to a solution along the lower branch state (see Fig. 9) for $Re = 500, 1000, 1500, 2000, 2500, 3000, 3500$ as a function of Δt_2 with $\Delta Re = 1$. (b) Subset of (a) for $Re = 500$ (thick line), $Re = 1000$ (thin line) and $Re = 1500$ (dashed line). (c) Subset of (a) for $Re = 1500$ (thick line), $Re = 2000$ (thin line) and $Re = 2500$ (dashed line). (d) Subset of (a) for $Re = 2500$ (thick line), $Re = 3000$ (thin line) and $Re = 3500$ (dashed line).




Figure 12: Efficiency results for the continuation of the lower branch in the $(\Delta t_2, Re)$ plane. The outer lines indicate the working interval within which continuation works. The shaded region indicates the efficiency interval where continuation takes less than 200% of the number of conjugate gradient iterations needed at the optimal Δt_2 for the same value of Re .

the edge of the working interval (see Fig. 11). Despite the fact that the working interval shrinks as Re increases, the efficiency interval appears steady and spans approximately $0.5 \leq \Delta t_2 \leq 10$. The steadiness of this interval is a crucial result in the design of a robust preconditioner for this system.




Figure 13: Comparison of the continuation method efficiency between the lower branch solution (dashed lines, see also Fig. 11) and the upper branch solution (thick line). Results are shown for $Re=1000$ (a), $Re=1500$ (b) and $Re=2000$ (c).

The above observations are compared to results along the upper branch at representative values of the Reynolds number in Fig. 13. The continuation along the upper branch is more difficult than along the lower branch. For $Re = 1000$ and $Re = 1500$, one can see that for an optimal preconditioner at Δt_{2opt} , the upper branch is slightly more computationally demanding than the lower branch, however, away from the optimum, the upper branch solution takes much longer to compute. In addition to the number of conjugate gradient iterations η that increases more abruptly for the upper branch away from η_{opt} , Fig. 13 also indicates that the working interval of the preconditioner is narrower for the upper branch than for the lower branch. These observations are confirmed and enhanced as the Reynolds number increases, as shown for $Re = 2000$ in Fig. 13(c).

4 Discussion

In this paper, I have presented and studied the performance of a preconditioner for numerical continuation of viscous incompressible stationary flows based on Stokes preconditioning [42]. This preconditioner is based on a time-stepper and can be tuned via one

parameter, the time-step Δt . In the limit $\Delta t \rightarrow 0$, the preconditioner acts like the identity while in the case of large Δt , it approximates the linear operator. The intermediate case, which has not been explored in the literature, leads to the preconditioner: $I - \Delta t L$, where I is the identity, Δt the parameter and L the linear operator.

In the original Stokes preconditioning article [42], the author investigated the influence of the parameter Δt on the convergence to an exact solution of Rayleigh–Bénard convection in a cylindrical container. The results are summarized in the single figure of the paper and show that the method fails at low Δt but that it improves as Δt increases to reach a performance plateau for $\Delta t > 0.1$. Stokes preconditioning has also been applied to spherical Couette flow where once again, it is shown that convergence is improved as Δt increases until $\Delta t \approx 10$ when the convergence speed reaches a plateau [43]. Both these observations seem to imply that taking Δt large maximizes the performance of the preconditioner. This idea has fed a number of studies in convection where Stokes preconditioning is used with large Δt , see [4, 5] among other papers.

The preconditioner is applied here to two cases that provide complementary observations: the optimum preconditioner is not necessarily obtained for large Δt . In the first problem, doubly diffusive convection, Stokes preconditioning provides a good solution but it is possible to improve convergence speed by using non asymptotic values of Δt . The number of conjugate gradient iterations can be reduced by up to 60%: in the lower segment case, convergence using the optimal preconditioner required 140 iterations against 350 for the asymptotic preconditioner. The other problem, the reduced model of shear flow, is more complex to deal with and increasing Δt deteriorates the preconditioner until it fails. The set of equations comprising the reduced model of shear flow is treated (simultaneously) using two different preconditioners. A first subset of the model is solved using the preconditioner $I - L$ and reasonably robust preconditioners for the second subset take the form: $I - Re^{-1}L$, thus varying with Re . It is important to emphasize the differences between the values used for Δt when doing Stokes preconditioned continuation and those used when doing time integration. For the same values of the other parameters, time integration is typically done using $\Delta t = 10^{-3}$ for doubly diffusive convection, so with values typically 100 times smaller than for optimal Stokes preconditioning. In the shear flow problem, time integration requires $\Delta t \leq Re^{-1}$, so values that are at smaller by at least a factor of Re compared to Stokes preconditioning.

Lastly, the method explained here has a number of advantages. It is easy to implement as it is based on a first order Euler time-scheme and it is adaptive through fine-tuning of the parameter Δt . When searching for stationary solutions, the system of equations can be split into several subsystems, each of which can then be treated with a different Δt and parameter-dependent preconditioning is also easily implementable and proved efficient in one test case studied here. This method is believed to be widely applicable and will provide an easy and reliable tool to investigate pattern formation in a variety of nonlinear systems.

Acknowledgments

This author is grateful to L. S. Tuckerman for discussions and encouragement and thanks A. Bergeon for providing some of the routines used.

References

- [1] D. Barkley, M. G. M. Gomes and R. D. Henderson, Three-dimensional instability in flow over a backward-facing step, *J. Fluid Mech.*, 473 (2002), 167–190.
- [2] O. Batiste, E. Knobloch, A. Alonso and I. Mercader, Spatially localized binary-fluid convection, *J. Fluid Mech.*, 560 (2006), 149–158.
- [3] C. Beaume, A. Bergeon, H.-C. Kao and E. Knobloch, Convectons in a rotating fluid layer, *J. Fluid Mech.*, 717 (2013), 417–448.
- [4] C. Beaume, A. Bergeon and E. Knobloch, Homoclinic snaking of localized states in doubly diffusive convection, *Phys. Fluids*, 23 (2011), 094102.
- [5] C. Beaume, A. Bergeon and E. Knobloch, Convectons and secondary snaking in three-dimensional natural doubly diffusive convection, *Phys. Fluids*, 25 (2013), 024105.
- [6] C. Beaume, G. P. Chini, K. Julien and E. Knobloch, Reduced description of exact coherent states in parallel shear flows, *Phys. Rev. E*, 91 (2015), 043010.
- [7] C. Beaume, H.-C. Kao, E. Knobloch and A. Bergeon, Localized rotating convection with no-slip boundary conditions, *Phys. Fluids*, 25 (2013), 124105.
- [8] C. Beaume, E. Knobloch and A. Bergeon, Nonsnaking doubly diffusive convectons and the twist instability, *Phys. Fluids*, 25 (2013), 114102.
- [9] A. Bergeon, D. Henry, H. Ben Hadid and L. S. Tuckerman, Marangoni convection in binary mixtures with Soret effect, *J. Fluid Mech.*, 375 (1998), 143–177.
- [10] A. Bergeon and E. Knobloch, Natural doubly diffusive convection in three-dimensional enclosures, *Phys. Fluids*, 14 (2002), 3233–3250.
- [11] K. Borońska and L. S. Tuckerman, Extreme multiplicity in cylindrical Rayleigh–Bénard convection II. Bifurcation diagram and symmetry classification, *Phys. Rev. E*, 81 (2010), 036321.
- [12] M. Chantry, A. P. Willis and R. R. Kerswell, Genesis of streamwise-localized solutions from globally periodic traveling waves in pipe flow, *Phys. Rev. Lett.*, 112 (2014), 164501.
- [13] R. M. Clever and F. H. Busse, Tertiary and quaternary solutions for plane Couette flow, *J. Fluid Mech.*, 344 (1997), 137–153.
- [14] R. H. Clewley, W. E. Sherwood, M. D. LaMar and J. M. Guckenheimer, PyDSTool, A Software Environment for Dynamical Systems Modeling, (2007).
- [15] H. A. V. der Vorst, Bi-CGSTAB: A fast and smoothly converging variant of bi-CG for the solution of nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.*, 13 (1992), 631–644.
- [16] H. A. Dijkstra, F. W. Wubs, A. K. Cliffe, E. Doedel, I. F. Dragomirescu, B. Eckhardt, A. Y. Gelfgat, A. L. Hazel, V. Lucarini, A. G. Salinger, E. T. Phipps, J. Sanchez-Umbria, H. Schutteleaars, L. S. Tuckerman and U. Thiele, Numerical bifurcation methods and their applications to fluid dynamics: analysis beyond simulation, *Commun. Comput. Phys.*, 15 (2014), 1–45.
- [17] E. J. Doedel, A. R. Champneys, F. Dercole, T. Fairgrieve, Y. Kuznetsov, B. Oldeman, R. Paffenroth, B. Sandstede, X. Wang and C. Zhang, AUTO-07P: Continuation and Bifurcation Software for Ordinary Differential Equations, 2008.
- [18] K. Engelborghs, T. Luzyanina and G. Samaey, DDE-BIFTOOL v. 2.00: a Matlab package for bifurcation analysis of delay differential equations, Technical Report TW-330, Dep. Comp.

- Sci., KU Leuven, (2001).
- [19] H. Faisst and B. Eckhardt, Traveling waves in pipe flow, *Phys. Rev. Lett.*, 91 (2003), 224502.
 - [20] F. Feudel, K. Bergemann, L. S. Tuckerman, C. Egbers, B. Fütterer, M. Gellert and R. Hollerbach, Convection patterns in a spherical fluid shell, *Phys. Rev. E*, 83 (2011), 046304.
 - [21] M. Frigo and S. G. Johnson, The design and implementation of FFTW3, *Proceedings of the IEEE*, 93 (2005), 216–231.
 - [22] J. F. Gibson, *Channelflow: A spectral Navier–Stokes simulator in C++*, tech. rep., U. New Hampshire, 2012. <http://channelflow.org/>.
 - [23] J. F. Gibson and E. Brand, Spanwise-localized solutions of planar shear flows, *J. Fluid Mech.*, 745 (2014), 25–61.
 - [24] J. F. Gibson, J. Halcrow and P. Cvitanović, Equilibrium and travelling-wave solutions of plane Couette flow, *J. Fluid Mech.*, 638 (2009), 243–266.
 - [25] G. Kawahara, M. Uhlmann and L. van Veen, The significance of simple invariant solutions in turbulent flows, *Annu. Rev. Fluid Mech.*, 44 (2012), 203–225.
 - [26] H. B. Keller, *Numerical Solutions of Bifurcation and Non-Linear Eigenvalues Problem: Application of Bifurcation Theory*, Academic Press New York, (1977).
 - [27] B. Krauskopf, H. M. Osinga and J. Galán-Viosque, *Numerical Continuation Methods for Dynamical Systems: Path Following and Boundary Value Problems*, Springer-Verlag, 2007.
 - [28] Y. A. Kuznetsov, A. Dhooge and W. Govaerts, Matcont: a MATLAB package for numerical bifurcation analysis of ODEs, *ACM Trans. Math. Softw.*, 29 (2003), 141–164.
 - [29] D. Lo Jacono, A. Bergeon and E. Knobloch, Magneto-hydrodynamic Convection, *J. Fluid Mech.*, 687 (2011), 595–605.
 - [30] D. Lo Jacono, A. Bergeon and E. Knobloch, Three-dimensional spatially localized binary-fluid convection in a porous medium, *J. Fluid Mech.*, 730 (2013), R2.
 - [31] C. K. Mamun and L. S. Tuckerman, Asymmetry and Hopf bifurcation in spherical Couette flow, *Phys. Fluids*, 7 (1995), 80–91.
 - [32] K. Melnikov, T. Kreilos and B. Eckhardt, Long-wavelength instability of coherent structures in plane Couette flow, *Phys. Rev. E*, 89 (2014), 043008.
 - [33] I. Mercader, O. Batiste, A. Alonso and E. Knobloch, Localized pinning states in closed containers: Homoclinic snaking without bistability, *Phys. Rev. E*, 80 (2009), 025201(R).
 - [34] I. Mercader, O. Batiste, A. Alonso and E. Knobloch, Convection, anticonvection and multi-convection in binary fluid convection, *J. Fluid Mech.*, 667 (2011), 586–606.
 - [35] M. Nagata, Three-dimensional finite-amplitude solutions in plane Couette flow: bifurcation from infinity, *J. Fluid Mech.*, 217 (1990), 519–527.
 - [36] C. C. T. Pringle and R. R. Kerswell, Asymmetric, helical and mirror-symmetric traveling waves in pipe flow, *Phys. Rev. Lett.*, 99 (2007), 074502.
 - [37] J. Sanchez, F. Marques and J. M. Lopez, A continuation and bifurcation technique for Navier–Stokes flows, *J. Comput. Phys.*, 180 (2002), 78–98.
 - [38] T. M. Schneider, J. F. Gibson and J. Burke, Snakes and ladders: Localized solutions of plane Couette flow, *Phys. Rev. Lett.*, 104 (2010), 104501.
 - [39] R. Seydel, *Practical Bifurcation and Stability Analysis, Interdisciplinary Applied Mathematics*, Springer, 2009.
 - [40] R. Seydel and V. Hlavacek, Role of continuation in engineering analysis, *Chem. Eng. Sci.*, 42 (1987), 1281–1295.
 - [41] J. F. Torres, D. Henry, A. Komiya and S. Maruyama, Bifurcation analysis of steady natural convection in a tilted cubical cavity with adiabatic sidewalls, *J. Fluid Mech.*, 756 (2014), 650–688.

- [42] L. S. Tuckerman, Steady-state solving via stokes preconditioning; recursion relations for elliptic operators, in 11th International Conference on Numerical Methods in Fluid Dynamics, D. Dwoyer, M. Hussaini and R. Voigt, eds., vol. 323 of Lecture Notes in Physics, Springer Berlin Heidelberg, 1989, 573–577.
- [43] L. S. Tuckerman, Laplacian preconditioning for the inverse Arnoldi method, Commun. Comput. Phys., 18 (2015), 1336–1351.
- [44] H. Uecker, D. Wetzel and J. D. M. Rademacher, Pde2path—a Matlab package for continuation and bifurcation in 2D elliptic systems, Numer. Math. Theor. Methods Appl., 7 (2014), 58–106.
- [45] D. Viswanath, The critical layer in pipe flow at high Reynolds numbers, Phil. Trans. R. Soc. A, 367 (2009), 561–576.
- [46] F. Waleffe, Homotopy of exact coherent structures in plane shear flows, Phys. Fluids, 15 (2003), 1517–1534.
- [47] J. Wang, J. Gibson and F. Waleffe, Lower branch coherent states in shear flows: Transition and control, Phys. Rev. Lett., 98 (2007), 204501.
- [48] H. Wedin and R. R. Kerswell, Exact coherent structures in pipe flow: travelling wave solutions, J. Fluid Mech., 508 (2004), 333–371.