

REVIEW ARTICLE

Numerical Bifurcation Methods and their Application to Fluid Dynamics: Analysis beyond Simulation

Henk A. Dijkstra^{1,*}, Fred W. Wubs², Andrew K. Cliffe³, Eusebius Doedel⁴,
Ioana F. Dragomirescu⁵, Bruno Eckhardt⁶, Alexander Yu. Gelfgat⁷,
Andrew L. Hazel⁸, Valerio Lucarini^{9,10}, Andy G. Salinger¹¹,
Erik T. Phipps¹¹, Juan Sanchez-Umbria¹², Henk Schuttelaars¹³,
Laurette S. Tuckerman¹⁴ and Uwe Thiele¹⁵

¹ *Institute for Marine and Atmospheric Research Utrecht, Utrecht University, The Netherlands.*

² *Department of Mathematics and Computer Science, University of Groningen, Groningen, The Netherlands.*

³ *School of Mathematical Sciences, University of Nottingham, Nottingham, UK.*

⁴ *Department of Computer Science, Concordia University, Montreal, Canada.*

⁵ *National Centre for Engineering Systems of Complex Fluids, University Politehnica of Timisoara, Romania.*

⁶ *Fachbereich Physik, Philipps-Universität Marburg, Marburg, Germany.*

⁷ *School of Mechanical Engineering, Faculty of Engineering, Tel-Aviv University, Tel-Aviv, Israel.*

⁸ *School of Mathematics, University of Manchester, Manchester, UK.*

⁹ *Meteorological Institute, Klimacampus, University of Hamburg, Hamburg, Germany.*

¹⁰ *Department of Mathematics and Statistics, University of Reading, Reading, UK.*

¹¹ *Sandia National Laboratories, Albuquerque, USA.*

¹² *Departament de Física Aplicada, Universitat Politècnica de Catalunya, Barcelona, Spain.*

¹³ *Department of Applied Mathematical Analysis, TU Delft, Delft, the Netherlands.*

¹⁴ *PMMH-ESPCI, Paris, France.*

¹⁵ *Department of Mathematical Sciences, Loughborough University, Loughborough, UK.*

Received 24 September 2012; Accepted (in revised version) 18 June 2013

Available online 26 July 2013

Abstract. We provide an overview of current techniques and typical applications of numerical bifurcation analysis in fluid dynamical problems. Many of these problems are characterized by high-dimensional dynamical systems which undergo transitions as parameters are changed. The computation of the critical conditions associated with

*Corresponding author. *Email address:* H.A.Dijkstra@uu.nl (H. A. Dijkstra)

these transitions, popularly referred to as ‘tipping points’, is important for understanding the transition mechanisms. We describe the two basic classes of methods of numerical bifurcation analysis, which differ in the explicit or implicit use of the Jacobian matrix of the dynamical system. The numerical challenges involved in both methods are mentioned and possible solutions to current bottlenecks are given. To demonstrate that numerical bifurcation techniques are not restricted to relatively low-dimensional dynamical systems, we provide several examples of the application of the modern techniques to a diverse set of fluid mechanical problems.

AMS subject classifications: 37H20, 35Q35, 76-02, 37M, 65P30

Key words: Numerical bifurcation analysis, transitions in fluid flows, high-dimensional dynamical systems.

Contents

1	Introduction	2
2	The methodology of continuation	4
3	Computation of bifurcation diagrams	10
4	Highlights of results	20
5	Summary and outlook	36

1 Introduction

Transition phenomena in flows of liquids and gases are of great fundamental interest. A prominent classical example is the sudden transition from laminar to turbulent flow which takes place in a circular pipe (Poiseuille flow) when the speed in the center of the pipe exceeds a critical value [51]. Other classical examples are the flow between rotating cylinders (Taylor-Couette flow), which undergoes successive transitions when the rotation rate of the inner cylinder is increased, and convection in a liquid layer heated from below (Rayleigh-Bénard-Marangoni flow), which shows a fascinating and rich set of flow patterns once a critical vertical temperature gradient is exceeded [62]. Central issues when studying these flows are the characterization of the range of conditions over which particular flows exist and the mechanisms of transition between the different flow patterns. In their classical treatise on Fluid Mechanics, [67] provide an elegant and general framework on the loss of stability of a general fluid flow due to changes in the parameters of the system.

Transitions in industrial and environmental flows are of great practical interest. Critical conditions in such flows, at which they may undergo a large qualitative change, are associated with what has been recently called a “tipping” point [48]. Examples are boundary flows, which may undergo qualitative changes in separation behavior and turbulence intensity, mean flows in turbulent buoyancy driven convection, which change

their overall pattern, and plasma flows in a Tokamak, which suddenly show strong oscillatory behavior. An example of a highly relevant environmental flow is the possible transition of the large-scale Atlantic ocean circulation once a critical freshwater perturbation is exceeded [30]. From a practical point of view, one would like to identify these critical conditions and understand how to avoid undesirable transitions.

To tackle flow transition problems, numerical models are essential and many finite element, finite difference (volume) and spectral codes for solving the Navier-Stokes equations and their extensions (such as Boussinesq equations) are available. Starting from a certain initial condition and given model parameters, the model is integrated forward in time and the long-time behavior of quantities of interest are studied; we will refer to these methods below as initial value problem (IVP) methods. To determine transition behavior and critical conditions, parameters are subsequently changed and the transient and asymptotic behavior of the model solutions is studied. In this way, transitions between different types of equilibrium behavior (steady or time-dependent) are found.

Since the primary interest is in changes in asymptotic behavior when parameters are changed, another class of numerical methods can be used, that focuses directly on the computation of the asymptotic flow states in the models. These may be steady states, periodic orbits, quasi-periodic orbits or more complicated states, usually referred to as attractors of the model. The issue of finding critical conditions for transitions is then rephrased in terms of dynamical systems theory to that of finding the parameter values at which bifurcations exist. The methods in the numerical bifurcation toolbox (NBT), in particular continuation techniques, consist of efficient numerical schemes to determine attractors as a function of parameters. This toolbox has originally been developed for tackling transition problems in models consisting of a small number of ordinary differential equations (ODEs), hence with a small number of degrees of freedom. Over the last decades the methods have been extended and applied to a number of flow problems governed by systems of discretized partial differential equations (PDEs) such as Navier-Stokes (and Boussinesq) equations with a large number of degrees of freedom.

The strong advantages of using NBT methods are (i) the systematic and efficient approach to the computation of attractors and bifurcation points and (ii) a detailed framework for understanding the physics of the flow transitions. In addition, the mathematical theory of dynamical systems provides a priori knowledge (e.g. based on the presence of symmetry in the problem) through a catalogue of possible behavior [50]. The NBT methods also provide information concerning the parameters and initial conditions which can lead to interesting transient behavior. Finally, NBT methods are also efficient in providing information to control possible dynamical behavior, e.g. how parameters can be adjusted to avoid certain transitions.

A disadvantage of the NBT methods is that they are more complicated, especially in their need for sophisticated numerical linear algebra and, as a result, their application in fluid flows is more limited than those of IVP methods. Although they can be quite relevant in detecting whether a bifurcation (or a series thereof) leads the system to chaos, NBT methods cannot be used to determine so-called strange attractors, characterized by

chaotic, aperiodic behavior. It is important to underline that a comprehensive understanding of the mechanisms leading a system to exhibit turbulent behavior is one of the grand challenges of physical and mathematical sciences [13].

The aim of the paper is to give an overview of the state-of-the-art of NBT methods and their application to flow transition problems. In Section 2, we will start by presenting the basic scheme followed in NBT methods together with a taxonomy of the different computational strategies. Section 3 is concerned with the state-of-the-art linear algebra methods applied in NBT methods. Subsequently, in Section 4, we provide an overview of current codes and their applications to fluid flow problems. Section 5 addresses current bottlenecks and the outlook for the use of NBT methods in the near future.

2 The methodology of continuation

Almost any computational fluid dynamical problem starts from the Navier-Stokes equations, possibly supplemented by conservation equations for scalar quantities. When these PDEs are discretized using a spectral, finite element or finite difference method, the resulting problem can be written in the general form

$$\mathcal{M} \frac{d\mathbf{u}}{dt} = \Phi(\mathbf{u}, \mathbf{p}) = \mathcal{L}\mathbf{u} + \mathcal{N}(\mathbf{u}), \quad (2.1)$$

where $\mathbf{u}(t) \in \mathbb{R}^n$ corresponds to a discretized solution of the original PDEs, \mathcal{M} is the mass matrix, \mathcal{L} the discretized linear operator, $\mathcal{N}(\mathbf{u})$ the discretized nonlinear operator. In general, the operators depend on a set of parameters, indicated in (2.1) by the vector \mathbf{p} . Note that \mathcal{M} is typically not invertible because of algebraic constraints such as the continuity equation in incompressible flows. We may also implicitly include the boundary conditions in the formulation (2.1). For instance, Dirichlet conditions correspond to \mathcal{M} and $\mathcal{N}(\mathbf{u})$ being zero on the boundary of the domain and \mathcal{L} being the identity operator.

2.1 Transition behavior and bifurcations

In many of the applications (as we will see in Section 4) one is interested in the changes in the long-time solutions of (2.1), when one (or more) parameter values are changed. In dynamical systems theory, these solutions are referred to as attractors of the high-dimensional dynamical system which is defined by (2.1). The simplest attractors of a model are fixed points, also called steady states. Complexity of the behavior of the solutions increases due to two factors: (i) the type of attractor and (ii) the number of attractors. After fixed points, the types of attractors are, in order of complexity, periodic orbits, orbits on a torus (quasi-periodic behavior) and a general class referred to as chaotic attractors.

Transitions between attractors when parameters are varied occur through bifurcations. The simplest bifurcations involve instabilities of steady states when one parameter p is varied. For each of these, there is a system of ODEs, called the normal form, which

contains the particular bifurcation in its most elementary form [53]. For easy reference in this and subsequent sections, we list these bifurcations below.

- (i) Saddle-node or fold bifurcation, also called a turning point: the normal form is given by

$$\frac{du}{dt} = p \pm u^2. \quad (2.2)$$

The sign characterizes supercriticality ($du/dt = p - u^2$) or subcriticality ($du/dt = p + u^2$). In the supercritical case, the branch of solutions $u = \sqrt{p}$ is linearly stable and the branch $u = -\sqrt{p}$ is unstable (see Fig. 1a). In the subcritical case, it is just the other way around (Fig. 1b).

- (ii) Transcritical bifurcation: in this case the normal form is given by

$$\frac{du}{dt} = pu \pm u^2. \quad (2.3)$$

In both subcritical and supercritical cases, there is an exchange of stability from stable to unstable fixed points and vice versa as the parameter p is varied through the bifurcation at $p = 0$ (Fig. 1c-d).

- (iii) Pitchfork bifurcation: the normal form is

$$\frac{du}{dt} = pu \pm u^3. \quad (2.4)$$

In the supercritical situation ($du/dt = pu - u^3$), stability is transferred from the symmetric solution $u = 0$ to the pair of conjugated solutions $u = \pm\sqrt{p}$ (Fig. 1e). In this case the system remains in a neighborhood of the equilibrium so that one observes a soft or non-catastrophic loss of stability. In the subcritical case ($du/dt = pu + u^3$), the situation is very different, as can be seen in Fig. 1f. The domain of attraction of $u = 0$ is bounded by the unstable fixed points and shrinks as p decreases, disappearing at $p = 0$. The system is thus pushed out from the neighborhood of the now-unstable fixed point leading to a sharp or catastrophic loss of stability. Decreasing the parameter again to negative values will not return the system to the previously stable fixed point since it will have already left its domain of attraction.

- (iv) Hopf bifurcation: It is also possible for a steady solution to transfer its stability to a limit cycle. The normal form can be written in polar coordinates (r, θ) as

$$\frac{dr}{dt} = pr \pm r^3, \quad (2.5a)$$

$$\dot{\theta} = \omega. \quad (2.5b)$$

The sign again determines whether the Hopf bifurcation is supercritical (Fig. 1g) or subcritical and the discussion is similar to the case of the pitchfork bifurcation.

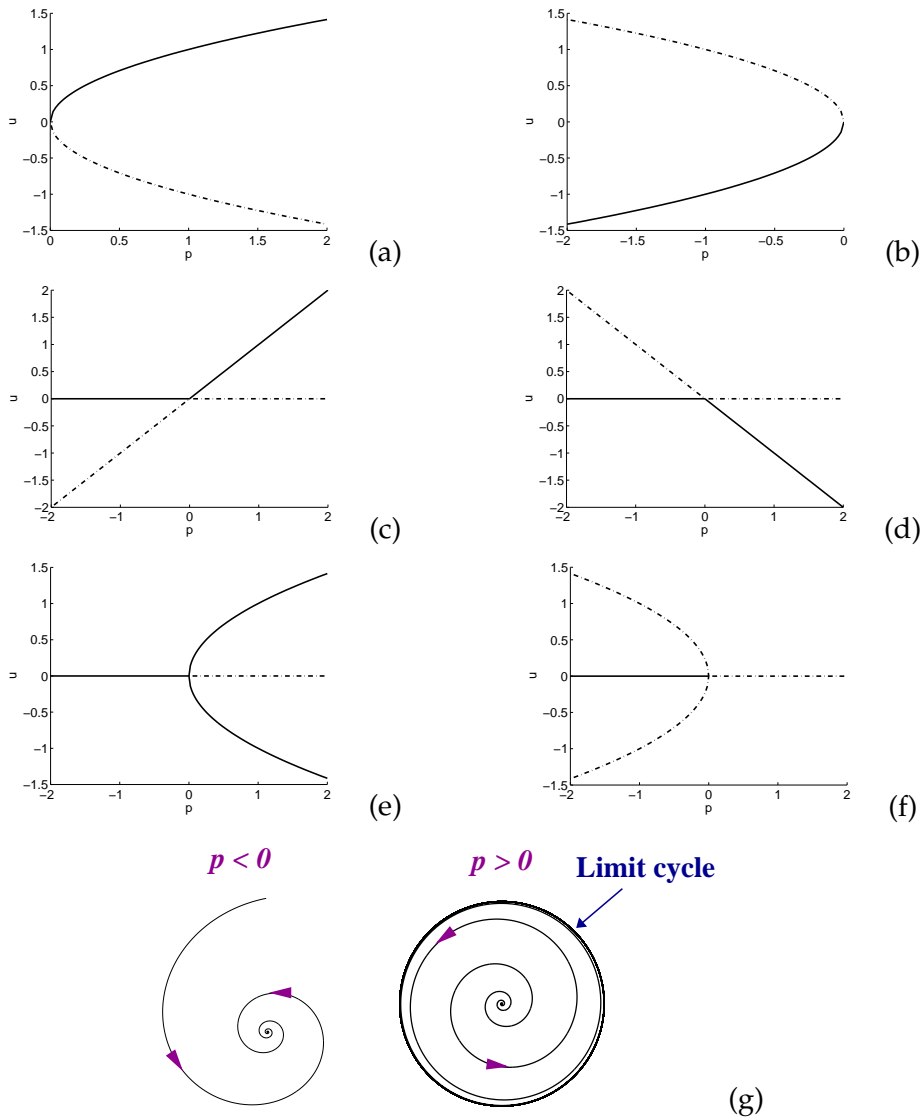


Figure 1: Supercritical (a) and subcritical (b) saddle-node bifurcation. Supercritical (c) and subcritical (d) transcritical bifurcation. Supercritical (e) and subcritical (f) pitchfork bifurcation. The solid (dash-dotted) branches indicate stable (unstable) solutions. (g) Phase space trajectories associated with a supercritical Hopf bifurcation at $p=0$. For $p < 0$, there is only one stable fixed point (left panel), whereas a stable limit cycle appears for $p > 0$ (right panel).

The idea of continuation methods is to compute families of steady states or periodic orbits as a parameter is varied. It enables one to obtain meaningful and generic information on the local dynamics of the PDE model (2.1) for a large range of parameter values. Although time integration of the model may ultimately be needed to detect more complicated bifurcations as well as statistical properties of the flow, continuation methods are

able to determine the first bifurcations from the branches of steady states and periodic orbits in an efficient way. When interested in steady states, continuation naturally avoids potentially long time integrations for many parameter values.

2.2 Pseudo-arclength continuation

We will first illustrate continuation methods through the computation of steady states and their stability. The basic technique involves two main steps. In the first step, one advances along a branch of steady states as a parameter is varied. In the second step, a linear stability analysis of the most recently computed steady state is performed [109].

Finding steady states of the system (2.1) amounts to solving

$$\Phi(\mathbf{u}, p) = 0. \quad (2.6)$$

The idea of pseudo-arclength continuation [59, 109] is to parametrize branches of solutions $\Gamma(s) \equiv (\mathbf{u}(s), p(s))$ with an arclength parameter s (or an approximation of it, thus the term ‘pseudo’) and choose s as the continuation parameter instead of p . An additional equation is obtained by approximating the normalization condition of the tangent $\dot{\Gamma}(s) = (\dot{\mathbf{u}}(s), \dot{p}(s))$ to the branch $\Gamma(s)$, where the dot refers to the derivative with respect to s , with $|\dot{\Gamma}|^2 = 1$. More precisely, for a given solution (\mathbf{u}_0, p_0) , the next solution (\mathbf{u}, p) is required to satisfy the constraint

$$\dot{\mathbf{u}}_0^T(\mathbf{u} - \mathbf{u}_0) + \dot{p}_0(p - p_0) - \Delta s = 0, \quad (2.7)$$

where $\dot{\Gamma}_0 = (\dot{\mathbf{u}}_0, \dot{p}_0)$ is the normalized direction vector of the solution family $\Gamma(s)$ at (\mathbf{u}_0, p_0) and Δs is an appropriately small step size. Eq. (2.7) stipulates that the projection of $(\mathbf{u}, p) - (\mathbf{u}_0, p_0)$ onto $(\dot{\mathbf{u}}_0, \dot{p}_0)$ has the value Δs .

In order to compute the tangent $\dot{\Gamma}_0(s)$, one differentiates $\Phi(\mathbf{u}(s), p(s)) = 0$ with respect to s at (\mathbf{u}_0, p_0) to find

$$\begin{bmatrix} \Phi_{\mathbf{u}}^0 & \Phi_p^0 \end{bmatrix} \dot{\Gamma}_0(s) = \begin{pmatrix} \frac{\partial \Phi_1^0}{\partial u_1} & \dots & \frac{\partial \Phi_1^0}{\partial u_n} & \frac{\partial \Phi_1^0}{\partial p} \\ \frac{\partial \Phi_n^0}{\partial u_1} & \dots & \frac{\partial \Phi_n^0}{\partial u_n} & \frac{\partial \Phi_n^0}{\partial p} \end{pmatrix} \dot{\Gamma}_0(s) = 0, \quad (2.8)$$

where the Jacobian matrix $\Phi_{\mathbf{u}}^0$ and the derivative Φ_p^0 are evaluated at (\mathbf{u}_0, p_0) . If (\mathbf{u}_0, p_0) is not a bifurcation point, then $\dim \ker [\Phi_{\mathbf{u}}^0, \Phi_p^0] = 1$ and therefore $[\Phi_{\mathbf{u}}^0, \Phi_p^0]$ has rank n . We may thus determine $\dot{\Gamma}_0(s)$ as the null vector of the $n \times (n+1)$ matrix $[\Phi_{\mathbf{u}}^0, \Phi_p^0]$. In practice, this can be done by upper triangulating the matrix $[\Phi_{\mathbf{u}}^0, \Phi_p^0]$ and solving a $(n+1) \times (n+1)$ extended system with a one in the right-lower corner and right-hand side. Then the normalization condition $|\dot{\Gamma}_0|^2 = 1$ is used [109].

Another possibility, which avoids solving linear systems once the continuation has been started, is to compute the tangent to the curve $\dot{\Gamma}_0$ by interpolation from a set of

previously computed solutions along the curve $\Gamma(s)$. This is important in the case of large-scale computations.

A predictor solution of (2.7) is given by

$$\mathbf{u}^0 = \mathbf{u}_0 + \Delta s \dot{\mathbf{u}}_0, \quad (2.9a)$$

$$p^0 = p_0 + \Delta s \dot{p}_0. \quad (2.9b)$$

The next step is then to project the predictor solution (\mathbf{u}^0, p^0) back onto the branch in a direction orthogonal to the tangent $\dot{\Gamma}_0$. This is called the corrector algorithm. It should rely on a robust nonlinear solver for the system (2.6). The one in common use is the Newton-Raphson method. This method converges quadratically, provided that the initial starting solution is close enough to the solution and that the Jacobian is non-singular [109]. The predictor step provides an adequate initial starting solution if Δs is sufficiently small. The Newton-Raphson iterations, with iteration index $k=0,1,\dots$, can then be written as

$$\begin{pmatrix} \Phi_{\mathbf{u}}(\mathbf{u}^k, p^k) & \Phi_p(\mathbf{u}^k, p^k) \\ \dot{\mathbf{u}}_0^T & \dot{p}_0 \end{pmatrix} \begin{pmatrix} \Delta \mathbf{u}^{k+1} \\ \Delta p^{k+1} \end{pmatrix} = \begin{pmatrix} -\Phi(\mathbf{u}^k, p^k) \\ r^k \end{pmatrix}, \quad (2.10)$$

where $r^k = \Delta s - \dot{\mathbf{u}}_0^T(\mathbf{u}^k - \mathbf{u}_0) - \dot{p}_0(p^k - p_0)$.

Once $(\Delta \mathbf{u}^{k+1}, \Delta p^{k+1})$ is found, one sets

$$\mathbf{u}^{k+1} = \mathbf{u}^k + \Delta \mathbf{u}^{k+1}, \quad (2.11a)$$

$$p^{k+1} = p^k + \Delta p^{k+1}. \quad (2.11b)$$

In practical situations, it is sometimes better to solve two $n \times n$ linear systems instead of directly solving (2.10), namely

$$\Phi_{\mathbf{u}}(\mathbf{u}^k, p^k) \mathbf{z}_1 = -\Phi(\mathbf{u}^k, p^k), \quad (2.12a)$$

$$\Phi_{\mathbf{u}}(\mathbf{u}^k, p^k) \mathbf{z}_2 = \Phi_p(\mathbf{u}^k, p^k). \quad (2.12b)$$

Then, the solution of (2.10) is

$$\Delta p^{k+1} = \frac{r^k - \dot{\mathbf{u}}_0^T \mathbf{z}_1}{\dot{p}_0 - \dot{\mathbf{u}}_0^T \mathbf{z}_2}, \quad \text{and} \quad \Delta \mathbf{u}^{k+1} = \mathbf{z}_1 - \Delta p^{k+1} \mathbf{z}_2. \quad (2.13)$$

The method is illustrated geometrically by Fig. 2, where the initial solution is indicated by $\mathbf{M}_0 = (\mathbf{u}_0, p_0)$, the predictor with \mathbf{M}_1 and the final converged solution by \mathbf{M}^* .

The advantage of pseudo-arclength methods over traditional continuation methods is that the Jacobian of the extended system $[\Phi_{\mathbf{u}}, \Phi_p]$ has rank n , even at folds where $\Phi_{\mathbf{u}}$ becomes singular. Hence, one can easily continue around folds. Another advantage of pseudo-arclength continuation is that it can compute branches of unstable as well as stable solutions [109].




Figure 2: Sketch of the pseudo-arclength method. The initial solution is indicated by $\mathbf{M}_0=(\mathbf{u}_0, p_0)$, the predictor by \mathbf{M}_1 and the final converged solution by \mathbf{M}^* .

Once a steady state $\bar{\mathbf{u}}$ is computed, its linear stability must be assessed. Setting $\mathbf{u}(t) = \bar{\mathbf{u}} + \tilde{\mathbf{u}}(t)$, linearizing (2.1) around $\bar{\mathbf{u}}$ and separating $\tilde{\mathbf{u}}(t) = \hat{\mathbf{u}} e^{\lambda t}$ gives a generalized eigenvalue problem of the form

$$\Phi_{\mathbf{u}}(\bar{\mathbf{u}})\hat{\mathbf{u}} = \lambda \mathcal{M}\hat{\mathbf{u}}, \tag{2.14}$$

where $\Phi_{\mathbf{u}}(\bar{\mathbf{u}})$ is again the Jacobian matrix evaluated at $\bar{\mathbf{u}}$. When all eigenvalues of the generalized eigenvalue problem (2.14) have negative real part, then the steady solution is linearly stable. If at least one eigenvalue has a positive real part, then the steady state is unstable.

2.3 Large-scale applications lead to special problems

It is fair to say that in the context of PDEs, continuation method algorithms are basically the same as those for ODEs. The difficulty rather lies in the accuracy of the Jacobian, the solution of the nonlinear systems of equations, and the leading eigenvalues of the linear stability problem.

The large linear systems which arise through the Newton-Raphson iteration can in general no longer be solved by direct solvers. To date, there also exists no general iterative method that solves all sparse linear systems efficiently. Hence, often the linear algebra is tuned to the problem. The same holds for the eigenvalue problem solvers, where QZ methods [49] must be replaced by tailored techniques, as will be discussed in Section 3.

All continuation methods rely essentially on a correct estimation of the extended Jacobian matrix $[\Phi_{\mathbf{u}}, \Phi_p]$. For PDEs this can be done in at least four ways: 1) by explicitly determining and coding the linearization, 2) by automatic differentiation, 3) by symbolic

computation or 4) by using the approximation

$$(\Phi_{\mathbf{u}})_{ij}(\bar{\mathbf{u}}) \simeq \frac{\Phi_i(\bar{\mathbf{u}} + \epsilon \mathbf{e}_j) - \Phi_i(\bar{\mathbf{u}})}{\epsilon}, \quad (2.15)$$

for small $\epsilon > 0$, where \mathbf{e}_j is the unit vector in the j^{th} direction. Each of these methods has specific difficulties, as will be discussed in more detail in Section 3.2 below.

In applications, one often does not want to make changes in a supplied application code that is highly tested and verified. If a code provides a right-hand-side Φ and possibly a Jacobian matrix $\Phi_{\mathbf{u}}$, then it is preferable to not make changes in the code in order to carry out bifurcation analysis. Thus, it is advantageous for the continuation code to be independent of the application code. A special challenge is then to design the bifurcation analysis algorithm so that it does not control the linear algebra objects, but just operates on them through an interface.

For systems of PDEs, the calculations require much more computing power than for ODE systems, and hence the numerical methods should be efficient. One has to exploit parallel computing and fast operations for data in cache. This requires special data structures and leads to complex software, easily growing to millions of lines of code.

3 Computation of bifurcation diagrams

The methods for computing the branches of equilibrium solutions as a function of parameters can be divided roughly into two types: methods for which only Φ , defined in (2.6), is explicitly available and methods for which both Φ and its Jacobian matrix $\Phi_{\mathbf{u}}$ are available.

3.1 Matrix-free and time-integration based methods

Using \mathbf{u} to denote the vector of independent variables and p to denote a parameter, the system (2.1) is written as:

$$\mathcal{M} \frac{d\mathbf{u}}{dt} = \Phi(\mathbf{u}, p), \quad (\mathbf{u}, p) \in \mathcal{U} \subset \mathbf{R}^n \times \mathbf{R}, \quad (3.1)$$

with $n \gg 1$ again being the dimension of the system. In the following, we also need the associated linearized system given by

$$\mathcal{M} \frac{d\mathbf{y}}{dt} = \Phi_{\mathbf{u}}(\mathbf{u}, p)\mathbf{y} + \Phi_p(\mathbf{u}, p)\mu, \quad (\mathbf{y}, \mu) \in \mathbf{R}^n \times \mathbf{R}, \quad (3.2)$$

for initial conditions $\mathbf{y}(0) = \mathbf{y}_0$ and fixed μ . We assume that a time stepper is available to solve the initial value problems (3.1) and (3.2) for given initial conditions.

A time-discretization scheme can be seen as replacing (3.1) by a map

$$\mathbf{u} \rightarrow \mathcal{G}(\mathbf{u}, p, T), \quad (\mathbf{u}, p) \in \mathcal{U} \subset \mathbf{R}^n \times \mathbf{R}, \quad (3.3)$$

where T is the total integration time. We are interested in the stable or unstable steady states of (3.1), which are fixed points of (3.3), and this will be discussed in the Subsections 3.1.1 and 3.1.2 below. Other invariant manifolds of (3.1), in particular periodic orbits and invariant tori, can be defined as fixed points of other maps \mathcal{G} , which we will discuss in Subsections 3.1.3 and 3.1.4.

3.1.1 Computation of fixed points by Newton-Krylov methods

We first discuss the general problem of computing fixed points of maps $\mathcal{G}(\mathbf{u}, p)$ via a combination of Newton-Raphson and Krylov subspace methods, or Newton-Krylov methods. We will omit T from the argument list because in some cases described below, the map will not depend explicitly on T .

When a continuation method is used to obtain sets (\mathbf{u}, p) of fixed points, systems of the form

$$\begin{pmatrix} \mathbf{u} - \mathcal{G}(\mathbf{u}, p) \\ m(\mathbf{u}, p) \end{pmatrix} = 0 \in \mathbf{R}^n \times \mathbf{R} \quad (3.4)$$

must be solved. In the case of arclength continuation, we have

$$m(\mathbf{u}, p) \equiv \mathbf{v}_u^\top (\mathbf{u} - \mathbf{u}^0) + v_p (p - p^0), \quad (3.5)$$

where (\mathbf{u}^0, p^0) (similar as in (2.9)) and (\mathbf{v}_u, v_p) are predictions of a new point on the curve of solutions and the curve's tangent. Another possibility for $m(\mathbf{u}, p)$ is to fix the value of p , switching to fixing a component of \mathbf{u} near a fold.

The system (3.4) is solved by Newton's method. At each step, the linear system

$$\begin{pmatrix} I - \mathcal{G}_u(\mathbf{u}^i, p^i) & -\mathcal{G}_p(\mathbf{u}^i, p^i) \\ \mathbf{v}_u^\top & v_p \end{pmatrix} \begin{pmatrix} \Delta \mathbf{u}^i \\ \Delta p^i \end{pmatrix} = \begin{pmatrix} -\mathbf{u}^i + \mathcal{G}(\mathbf{u}^i, p^i) \\ -m(\mathbf{u}^i, p^i) \end{pmatrix} \quad (3.6)$$

is solved iteratively by matrix-free methods and the estimated values are updated via

$$(\mathbf{u}^{i+1}, p^{i+1}) = (\mathbf{u}^i, p^i) + (\Delta \mathbf{u}^i, \Delta p^i). \quad (3.7)$$

Iterative methods such as the generalized minimal residual method (GMRES) or the bi-conjugate gradient method (BiCGStab) [94] require only the computation of matrix products, i.e., the calculation of actions of the form

$$\mathcal{G}_u(\mathbf{u}^i, p^i) \Delta \mathbf{u} + \mathcal{G}_p(\mathbf{u}^i, p^i) \Delta p \quad (3.8)$$

for given $\Delta \mathbf{u}$ and Δp .

Convergence is facilitated if the spectrum of \mathcal{G}_u evaluated at the fixed points is clustered around the origin, more precisely if very few of its eigenvalues are located outside a disk centered at the origin and of a radius less than one. This is indeed the case for maps like (3.3) which involve the time integration of systems of elliptic-parabolic PDEs. In this case the map \mathcal{G} is a strong contraction to the fixed points, except along the unstable manifold if the fixed point is unstable. We assume that the dimension of this manifold is very

small compared to the dimension n of the system. In summary, the system (3.6) needs no preconditioning if the Jacobian $\mathcal{G}_{\mathbf{u}}(\mathbf{u}^i, p^i)$ has the kind of spectrum described above, and we can expect fast convergence in $\mathcal{O}(10)$ iterations [98].

To compute the stability of a fixed point, $\bar{\mathbf{u}}$, we only need to use any variant of the power method (subspace iteration, Arnoldi methods, etc.) to compute the leading eigenvalues (largest in modulus) of $\mathcal{G}_{\mathbf{u}}(\bar{\mathbf{u}})$ [2, 82].

We will now describe, for the three types of invariant objects under consideration, the corresponding map \mathcal{G} and the action of its Jacobian.

3.1.2 Continuation of fixed points based on time integration

In order to compute steady states of (3.1), we propose the map

$$\mathcal{G}(\mathbf{u}, p) = \boldsymbol{\varphi}(T, \mathbf{u}, p), \quad (3.9)$$

with $\boldsymbol{\varphi}(T, \mathbf{u}, p)$ the solution of (3.1) at time T with initial condition \mathbf{u} . It is clear that

$$\Phi(\mathbf{u}, p) = 0 \Rightarrow \mathbf{u} - \boldsymbol{\varphi}(T, \mathbf{u}, p) = 0. \quad (3.10)$$

The matrix products required to solve the linear systems (3.6) are obtained by integrating (3.2) with initial condition $\mathbf{y}_0 = \Delta \mathbf{u}$, and $\mu = \Delta p$. Each matrix product then requires the time integration of a system of $2n$ equations, (3.1) for \mathbf{u} and (3.2) for \mathbf{y} . This will also be the case for periodic orbits and invariant tori.

If the fixed points are stable, the method can be seen as an acceleration of the time evolution towards the steady state. If they are unstable, it is similar to a stabilization method like the Recursive Projection Method of [111]. This method was used in [81] to compute solution branches in a two-dimensional annular region subjected to radial gravity and differential heating.

The integration time T is chosen to optimize the computation time. Increasing T concentrates the spectrum of $\mathcal{G}_{\mathbf{u}}$ at the origin, reducing the number of matrix evaluations (3.9) required to solve (3.6). However, increasing T (for a fixed time step Δt) also increases the computation time to carry out each matrix evaluation (3.9). A few numerical experiments are usually required to determine the optimal value of T .

In another approach, a single large time step of size T is used, and so no tradeoff is necessary to select the value of T . In this case, \mathcal{G} no longer represents accurate time stepping of (3.2) but its fixed points can nonetheless be steady states of (3.2) as we will now demonstrate. A mixed time discretization is used to integrate (2.1), with \mathcal{L} and \mathcal{N} integrated implicitly and explicitly, respectively:

$$\mathcal{M}(\mathbf{u}(t + \Delta t) - \mathbf{u}(t)) = \Delta t (\mathcal{L}\mathbf{u}(t + \Delta t) + \mathcal{N}(\mathbf{u}(t))). \quad (3.11)$$

Eq. (3.11) is rearranged to yield the time-stepping scheme

$$\mathbf{u}(t + \Delta t) = (\mathcal{M} - \Delta t \mathcal{L})^{-1} (\mathcal{M} + \Delta t \mathcal{N}) \mathbf{u}(t), \quad (3.12)$$

where $\mathcal{M} - \Delta t \mathcal{L}$ can be inverted inexpensively. In any tensor-product geometry, such as Cartesian, cylindrical, or spherical, the inversion can be carried out directly, while in more complicated geometries, techniques such as multigrid methods are available for solving elliptic problems quickly. Thus (omitting the parameter \mathbf{p} in the notation),

$$\mathcal{G}(\mathbf{u}) = \boldsymbol{\varphi}(\mathbf{u}, T) \equiv (\mathcal{M} - T\mathcal{L})^{-1}(\mathcal{M} + T\mathcal{N})\mathbf{u}, \quad (3.13)$$

and we seek solutions of

$$\begin{aligned} 0 &= \mathcal{G}(\mathbf{u}) - \mathbf{u} \\ &= (\mathcal{M} - T\mathcal{L})^{-1}(\mathcal{M} + T\mathcal{N})\mathbf{u} - \mathbf{u} \\ &= (\mathcal{M} - T\mathcal{L})^{-1}(\mathcal{M} + T\mathcal{N} - (\mathcal{M} - T\mathcal{L}))\mathbf{u} \\ &= (\mathcal{M} - T\mathcal{L})^{-1}T(\mathcal{N} + \mathcal{L})\mathbf{u}. \end{aligned} \quad (3.14)$$

This demonstrates that steady states of (2.1) are indeed fixed points of \mathcal{G} for the time-stepping scheme (3.11), for any value of T . For large T , (3.14) shows that

$$\mathcal{G}(\mathbf{u}) - \mathbf{u} \approx \mathcal{L}^{-1}(\mathcal{N} + \mathcal{L})\mathbf{u}, \quad (3.15a)$$

$$\mathcal{G}_{\mathbf{u}} - \mathcal{I} \approx \mathcal{L}^{-1}(\mathcal{N}_{\mathbf{u}} + \mathcal{L}), \quad (3.15b)$$

where \mathcal{I} is the identity matrix. Consequently \mathcal{L}^{-1} plays the role of a preconditioner for $\mathcal{N}_{\mathbf{u}} + \mathcal{L}$, and $\mathcal{G}_{\mathbf{u}} - \mathcal{I}$ is essentially a preconditioned version of $\mathcal{N}_{\mathbf{u}} + \mathcal{L}$ [15, 19, 134]. This method has been used to calculate bifurcation diagrams in many physical systems, such as spherical Couette flow [73], cylindrical Rayleigh-Bénard convection [14, 133], Bose-Einstein condensation [57], and binary fluid convection [1, 3].

Far away from $\mathbf{u} = 0$, the linear solver may fail to converge, and other preconditioners of $\mathcal{N}_{\mathbf{u}} + \mathcal{L}$ may be sought. If finite differences, volumes or elements are employed, incomplete LU decomposition or other techniques described later can be used as preconditioners to accelerate the convergence (see [79, 95] for examples of this methodology). For spectral methods, it is not easy to find good preconditioners. Finite difference or finite element versions of the problem have been successfully used as preconditioners [17, 18] but the coding becomes much more complicated.

3.1.3 Continuation of periodic orbits

For computing periodic orbits, two possibilities are available for the map \mathcal{G} . The first consists of a Poincaré map on a certain manifold. A hyperplane Σ_1 defined by

$$g_{\Sigma_1}(\mathbf{u}) = \mathbf{v}_1^\top (\mathbf{u} - \mathbf{u}_1^\Sigma) = 0 \quad (3.16)$$

was used in [98]. In this case $\mathcal{G}(\mathbf{u}, p) = \boldsymbol{\varphi}(T, \mathbf{u}, p)$, with $g_{\Sigma_1}(\boldsymbol{\varphi}(T, \mathbf{u}, p)) = 0$, is the first intersection close to \mathbf{u} of the trajectory starting at $\mathbf{u} \in \Sigma_1$, and T the arrival time. When $\mathbf{u} = \boldsymbol{\varphi}(T, \mathbf{u}, p)$ we have a periodic orbit of period T . The details on how to parameterize Σ_1 ,

which has dimension $n-1$, are given in [98]. The action of the Jacobian can be computed by the formula (see [112])

$$\mathcal{G}_{\mathbf{u}}(\mathbf{u}, p) \Delta \mathbf{u} + \mathcal{G}_p(\mathbf{u}, p) \Delta p = \mathbf{y} - \frac{\mathbf{v}_1^\top \mathbf{y}}{\mathbf{v}_1^\top \mathbf{z}} \mathbf{z}, \quad (3.17)$$

where $\mathbf{v}_1^\top \Delta \mathbf{u} = 0$, $\mathbf{z} = \Phi(\mathcal{G}(\mathbf{u}, p), p)$, and \mathbf{y} is the solution of the first variational equation (3.2) with initial condition $\mathbf{y}_0 = \Delta \mathbf{u}$, and $\mu = \Delta p$. Again, each matrix product requires the time integration of a system of $2n$ equations.

This formulation, together with Newton-Krylov methods, was used in thermal convection problems in [98] and [89], and in [100] to obtain a normal form of a periodic orbit with symmetries. The latter study also required the eigenfunctions of an adjoint operator at the specific periodic orbit. A multiple shooting variant using parallelism has been considered in [96]. The computations of segments of two-dimensional invariant manifolds in large-scale systems has been described in [135], generalizing the ideas of [64] and of [63].

The second formulation, without Poincaré maps, is that used, for instance, in AUTO [33–35], CONTENT [66], or MatCont [29] and consists of solving the bordered system

$$\mathbf{u} - \boldsymbol{\varphi}(T, \mathbf{u}, p) = 0, \quad (3.18a)$$

$$g_{\text{phase}}(\mathbf{u}) = 0, \quad (3.18b)$$

$$m(\mathbf{u}, p) = 0 \quad (3.18c)$$

for (\mathbf{u}, p, T) using Newton-Krylov methods. In this case the map depends explicitly on T , i.e., $\mathcal{G}(\mathbf{u}, p, T) = \boldsymbol{\varphi}(T, \mathbf{u}, p)$. In (3.18), $g_{\text{phase}}(\mathbf{u}) = 0$ is a phase condition which selects a single point on the periodic orbit, and $m(\mathbf{u}, p) = 0$ is a pseudo-arclength-like condition as in (3.4). The phase condition can be again by Eq. (3.16) or an integral constraint, as is done in AUTO. It is possible, but not strictly necessary, to scale the time as $t = T\tau$, with T the unknown period of the periodic orbit (so that now the period is $\tau = 1$), and to then rewrite the original system (3.1) as

$$\mathcal{M} \frac{d\mathbf{u}}{d\tau} = T\Phi(\mathbf{u}, p). \quad (3.19)$$

This is done when the periodic orbits are computed by collocation in time, in order to fix the endpoints of the time interval.

The action of the Jacobian can be computed as described in the case of fixed points by integrating the variational equation (3.2). The only difference is that the derivative with respect to T must be included. These algorithms have been used to find periodic orbits for plane channel and pipe flows in [37] and in [47]. In some cases the authors did not have a good initial condition to start the continuation and therefore used globalized Newton's methods to increase the size of the basin of attraction [87].

Newton-Krylov methods are most widely used, but other techniques have been also employed. In particular, the use of Newton-Picard iterations to compute periodic orbits was described in [72]. A library, PDECONT, is freely available and was used in [128] for the flow in a lid-driven cavity.

3.1.4 Computation of invariant two-dimensional tori

It is also possible to compute invariant tori using continuation methods [99]. The intersection of an invariant two-dimensional torus with a hyperplane Σ_1 which cuts it transversally is, near one of the points of the intersection, an arc of a curve (Fig. 3). Let $\mathcal{P}: \mathcal{V} \subset \Sigma_1 \rightarrow \Sigma_1$ be the Poincaré map defined on the hyperplane Σ_1 . The intersection arc will be invariant under this map. Let Σ_2 be another hyperplane, given by $\mathbf{v}_2^\top (\mathbf{u} - \mathbf{u}_2^\Sigma) = 0$, transverse both to Σ_1 and to the invariant two-torus. Then we define the map

$$\mathcal{G}(\mathbf{u}, p): \mathcal{U} \subset (\Sigma_1 \cap \Sigma_2) \times \mathbf{R} \rightarrow \mathbf{R}^{n-2} \tag{3.20}$$

as follows (see Fig. 3). If \mathbf{u} is a point on the intersection $\Sigma_1 \cap \Sigma_2$, and $B(\mathbf{u}, \varepsilon)$ is the ball of radius ε centered at \mathbf{u} , a time integration with initial condition \mathbf{u} is started to find the first $q+1$ powers of the Poincaré map

$$\mathbf{z}_j = \mathcal{P}^{k_j}(\mathbf{u}, \lambda), \quad \text{with } j=1, \dots, q+1 \quad \text{and } k_1 < k_2 < \dots < k_{q+1}, \tag{3.21}$$

which fall inside $B(\mathbf{u}, \varepsilon)$ ($q=3$ in Fig. 3). Then the intersection of Σ_2 with the polynomial which interpolates these points defines the map \mathcal{G} . Its fixed points are approximations of the points we are looking for. If

$$\mu_j = \mathbf{v}_2^\top (\mathcal{P}^{k_j}(\mathbf{u}, p) - \mathbf{u}_2^\Sigma), \quad j=1, \dots, q+1 \tag{3.22}$$

are the projections of the points \mathbf{z}_j onto the line $\mathbf{u} = \mathbf{u}_2^\Sigma + \mu \mathbf{v}_2$, then

$$\mathcal{G}(\mathbf{u}, p) = \sum_{j=1}^{q+1} l_j(0) \mathbf{z}_j, \tag{3.23}$$

where the $l_j(0)$ are the Lagrange interpolation polynomials of degree q , based on the nodes μ_j , evaluated at $\mu=0$.




Figure 3: Scheme of the map \mathcal{G} defining a point on the torus. The solid line represents an arc of an invariant curve in Σ_1 , and $\Sigma_1 \cap \Sigma_2$ is the vertical dashed line. The intersection of the invariant curve with Σ_2 is the point we want to approximate.

The computation of the action of the Jacobian of \mathcal{G} on a vector reduces to the case of the differential of the Poincaré map. As in the case of periodic orbits, $\Sigma_1 \cap \Sigma_2$ must be parameterized, and the radius ε defining \mathcal{G} must be varied adaptively during the continuation process. The method is only valid, in principle, for stable quasiperiodic tori, but practice has shown that weakly unstable and resonant tori can be also computed. More details are given in [99], where the method was applied to compute invariant tori in a binary-mixture convection problem. Very recently an efficient parallel implementation has been developed [97].

3.2 Matrix-based techniques

In this section we discuss the bifurcation analysis that can be carried out when the application code supplies a Jacobian matrix. A significant number of simulation codes for complex flow simulations compute a Jacobian matrix, defined as $\Phi_{\mathbf{u}}$ in the previous section. This code development is typically motivated by investigations of behaviors that occur on long time scales and at steady state. In these cases, Newton-Raphson based algorithms for implicit time integration and direct-to-steady solvers are the methods of choice for efficiency.

Having a Jacobian matrix, and not merely an algorithm for applying the Jacobian operator, can lead to more robust and efficient linear algebra. For modest sized problems, a direct solver can be used. For large-scale applications where direct solvers are not practical, the entries of the matrix are used to generate a preconditioner to be used with a Krylov iteration method such as GMRES. Having a fully formed matrix also leads to more efficient application of the Jacobian operator for this computationally-intensive step of Krylov iterations. Thus, for stiff linear algebra problems for which effective preconditioners and many Krylov iterations are needed to solve the linear system, the cost of computing a Jacobian matrix is worthwhile for efficiency and even necessary for robustness.

3.2.1 Linear systems

In pseudo-arclength continuation, as well as in the eigensolution methods, *bordered* linear systems arise. These are of the form

$$\begin{bmatrix} J & V \\ W^T & C \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \mathbf{s} \end{bmatrix} = \begin{bmatrix} \mathbf{f}_x \\ \mathbf{f}_c \end{bmatrix}, \quad (3.24)$$

where J is a sparse Jacobian matrix, and V and W contain an equal number of vectors. The system (2.10) serves as an example with $J = \Phi_{\mathbf{u}}$, $V = \Phi_p$, $W = \dot{\mathbf{u}}_0$ and $C = \dot{p}_0$.

A standard approach is to carry out a block LU-factorization

$$\begin{bmatrix} J & V \\ W^T & C \end{bmatrix} = \begin{bmatrix} J & 0 \\ W^T & I \end{bmatrix} \begin{bmatrix} I & J^{-1}V \\ 0 & C - W^T J^{-1}V \end{bmatrix}. \quad (3.25)$$

This is fine as long as J is non singular. However, at bifurcation (e.g. fold) points the matrix J becomes singular, and applying its inverse leads to an unstable algorithm. Our experience is that this is not a problem when using pseudo-arclength continuation to pass around a fold point since the stepping algorithm will not land on the bifurcation point precisely. This does become an issue when attempting to converge to the bifurcation point. For the case of folds, there are essentially two approaches to deal with this situation. The first is to integrate the border rows and columns into the system and use multilevel preconditioners. The second one is to stabilize the block LU-factorization near the bifurcation point; we will only discuss the first one.

Block multilevel preconditioners treat the border as part of the matrix. For example, in a two-level case we get

$$\begin{aligned} & \begin{bmatrix} A_{11} & A_{12} & V_1 \\ A_{21} & A_{22} & V_2 \\ W_1^T & W_2^T & C \end{bmatrix} \\ &= \begin{bmatrix} A_{11} & 0 & 0 \\ A_{21} & I & 0 \\ W_1^T & 0 & I \end{bmatrix} \begin{bmatrix} I & A_{11}^{-1}A_{12} & A_{11}^{-1}V_1 \\ 0 & A_{22}-A_{21}A_{11}^{-1}A_{12} & V_2-A_{21}A_{11}^{-1}V_1 \\ 0 & W_2^T-W_1^T A_{11}^{-1}A_{12} & C-W_1^T A_{11}^{-1}V_1 \end{bmatrix}. \end{aligned} \quad (3.26)$$

As described in [117], on the last level a direct method with pivoting can be applied, which precludes an unstable factorization. Note that on the first level (and in general on all but the last level) the application of the preconditioner to the border is just the same as for the right hand side; it only differs for the last block. The implementation can take advantage of this [117].

3.2.2 The eigenvalue problem

For the linear stability problem (2.14) one is usually only interested in those in a special region, e.g., near the origin or near the imaginary axis. Fortunately the corresponding eigenmodes are in general quite smooth, since, due to diffusion, eigenvalues corresponding to high frequency eigenmodes have a large real negative part. The interesting eigenmodes are in the low frequency range and these are for a large part determined by the geometry. Moreover, in general the number of these is limited.

For the eigenvalue problem basically two methods exist: the Arnoldi method, [92] which is a Krylov subspace method, and the Jacobi-Davidson method [116], which is an accelerated Newton method. Both methods can be adjusted to target a desired set of eigenvalues through a spectral transformation. However, targeting means that a shifted system (such as the shift-invert transformation) needs to be solved. For Arnoldi methods such systems need to be solved to high accuracy in order to get an accurate Krylov subspace, which is immediately related to the accuracy to which one can find the eigenvalues. Newton methods, and thereby the Jacobi-Davidson method, do not require this. Inaccurate solves only influence the rate of convergence.

The shift-invert spectral transformation illuminates eigenvalues near the shift point for accelerated identification and convergence. Since linear solvers are typically only




Figure 4: Contour plot of the magnitude of an eigenvalue after undergoing a Cayley spectral transformation (3.27) as a function of the complex eigenvalue in the untransformed system for $\sigma = -\mu = 100$. For an eigenvalue λ located at any spot in this plane, the color contour indicates the magnitude of the corresponding eigenvalue γ after undergoing the Cayley transformation. This transformation selects those eigenvalues near or to the right of the imaginary axis, which is desirable for stability analysis.

available for real-valued matrices, one usually restricts the shift to a real value. In order to be able to compute eigenvalues that have significant imaginary components, a so-called Generalized Cayley Transformation [76], given by

$$C = (J - \sigma \mathcal{M})^{-1} (J - \mu \mathcal{M}), \quad (3.27)$$

is often applied. This transforms eigenvalues λ of the original generalized eigenvalue problem (2.14) ($J\mathbf{u} = \lambda \mathcal{M}\mathbf{u}$) to eigenvalues $\gamma = (\lambda - \mu) / (\lambda - \sigma)$ of the transformed problem $C\mathbf{u} = \gamma \mathbf{u}$. In Fig. 4, a contour plot over the complex plane is shown for $\sigma = -\mu = 100$. For an eigenvalue λ located at any spot in this plane, the color contour indicates the magnitude of the corresponding eigenvalue γ after undergoing the Cayley transformation.

The Arnoldi method will converge fastest to the eigenvalues with the largest norms after this transformation. It is evident that the Cayley transformation is desirable for problems whose eigenvalues have large imaginary components, since the region of relatively large magnitude extends up the imaginary axis (note the difference in scales for the real and imaginary axes). The target need not be very precise in order to locate the interesting eigenvalues, but it may influence the speed of convergence.

Analogies to time integration can be used to determine good choices for the parameters in (3.27). Since eigenvalues have units of inverse time, the appropriate value of the shift σ in the spectral transformations depends on the scaling of the particular problem. As a rule, if one is investigating an instability through transient simulation and would pick a time step of Δt (based on knowledge or intuition of the characteristic time scales of the problem), then an appropriate shift would be $\sigma = 2.0 / \Delta t$.

3.2.3 Implementation issues

Once a problem has been abstracted out to the linear algebra level, it is possible to use general-purpose algorithms for all phases of the analysis: preconditioning, iterative solvers, Newton solvers, eigensolvers, continuation algorithms, and bifurcation tracking algorithms. A large set of such algorithms have been developed in a modular, object-oriented environment in the Trilinos framework (trilinos.sandia.gov). Inside Trilinos there are separate software packages for the different phases of the calculations: LOCA (continuation and bifurcation analysis), NOX (nonlinear solvers), Anasazi (eigensolver), Epetra (distributed memory linear algebra objects), and many preconditioners and linear solvers. All of these packages are designed to interoperate, so the continuation algorithm can call the nonlinear solver, which in turn can call any linear solver. However, the connection is through abstract interfaces so that users can use a tailored algorithm in place of the Trilinos version, such as a user-defined physics-based preconditioner.

The bordered matrix systems that arise from the bifurcation algorithms can be presented using abstract linear algebra objects. Some of the algorithms, such as the block matrix LU approach, can be implemented at this abstract level. Others, such as when integrating the bordered system into the matrix, require direct manipulation of data and layouts (such as a vector length or matrix sparsity) that can only be implemented for concrete linear algebra objects. In the LOCA implementation in Trilinos, bordered algorithms are available at the generic level, while others (e.g. the Householder approach of [139]) are only available for codes that use the Trilinos Epetra linear algebra data structures.

At the beginning of this section we stipulated that we are dealing with application codes with a Jacobian matrix. Here we will briefly discuss the approaches that are used to generate this matrix. We can divide the approaches into *analytic* Jacobians which are exact derivatives up to floating point precision, and *numerical* Jacobians which are generated with finite difference approximations to the derivatives.

Analytic Jacobian matrices can be programmed by hand. For example, the linear part of the matrix is computed and used for both the construction of the right-hand-side and the Jacobian. In each update of the Jacobian matrix only the new nonlinear part has to be added. So this approach is strongly matrix oriented. Analytic Jacobian matrices can also be computed with Automatic Differentiation, a symbolic approach for generating derivatives. With Automatic Differentiation, the code for the analytic derivatives is created, either by source transformation using a pre-processor, for example, ADIFOR, <http://www.mcs.anl.gov/research/projects/adifor> and OpenAD <http://www.mcs.anl.gov/OpenAD> projects or through the operator overloading approach for C++ codes (<http://trilinos.sandia.gov/packages/sacado>). Automatic Differentiation can lead to a considerable decrease in code development time over hand-coding the analytic Jacobian. If the PDEs are fixed, this might be a modest gain, but it becomes a tremendous benefit when part of the research effort involves changes to the form of the equations (such as investigations of source terms, solution-dependent properties, and equations of state). The numerical efficiency depends on many factors, but an

Automatic Differentiation-generated Jacobian could take 50% longer to compute than a hand-coded Jacobian matrix. For large-scale problems where the linear solver time dominates, this difference decreases in significance.

Numerical Jacobian matrices can be calculated by finite difference approximation by repeated calculation of the right-hand-side. Using the approximation of a directional derivative (2.15), the numerical Jacobian matrix can be constructed. In order to use numerical differentiation for large problems, the sparsity pattern of the Jacobian matrix must be exploited. In general it is not possible to evaluate a certain entry of the right-hand-side function. Usually there is a vector input and a vector output. So it is crucial to limit the number of calls to the right-hand-side function; one should try to get as many derivatives as possible from one call to the right-hand-side function. For instance, if the Jacobian matrix is just a diagonal matrix then due to the independence only two calls are needed to get all the values. In general, unknowns can be subdivided by a coloring algorithm in such a way that a small number of right-hand-side evaluations yields all of the coefficients [27]. A more efficient approach, particularly suited to finite element calculations, is to perform the finite difference loop at a local (element) level. If an element has N degrees of freedom, then the dense element stiffness matrix can be computed by $N+1$ residual evaluations for this element. This is slightly more invasive in the code, but tends to be much easier to implement and requires far fewer residual evaluations than the coloring approach.

To our knowledge, the LOCA package of the Trilinos framework is the only general-purpose bifurcation analysis package that targets large-scale applications. The main obstacle involved in devising such a library is that the performance and robustness often depend largely on the inner linear solver method, and its application-specific data structures and preconditioners. So, a general purpose package must use abstract layers around the linear algebra so the bifurcation analysis algorithms can be written independently of the linear algebra. This makes it harder to write, and to interface to, than might be expected.

4 Highlights of results

To demonstrate the use of numerical bifurcation analysis in solving fluid dynamical problems, we provide an overview of several cases which have been studied in detail over the years.

4.1 Exact coherent states in shear flows

The transition to turbulence in plane Couette flow and pipe flow shares some of the phenomenology of subcritical transitions: turbulence can be found while the laminar flow is stable against infinitesimal perturbations, so that a finite amplitude disturbance is required to initiate turbulence. This looks like a subcritical transition but it differs from the

standard cases discussed in Section 2 in that there is no laminar state from which a sub-critical bifurcation could be tracked [41, 51]. Nevertheless, numerical and experimental studies of pipe flow show that there are persistent three-dimensional coherent structures that provide a scaffold for the turbulence in the sense that they appear transiently in the turbulent velocity fields [40, 41, 43, 56, 105, 140]. These flow states are exact, persistent solutions to the Navier-Stokes equations. They typically are fully three-dimensional, i.e. all velocity components are active and vary in all three space dimensions. They are known as ‘exact coherent states’ or ECS [138].

The exact coherent states can be fixed points of the Navier-Stokes equations in the original frame or in a co-moving frame of reference (in which case they become travelling waves in the laboratory frame) or real periodic states. For fixed points, an interesting variant on continuation methods is the Newton hook-step method introduced in [136, 137]. Let $\Phi(\mathbf{u}, \mathbf{p})=0$ again be the equation to be solved for the fixed point \mathbf{u} as a function of the external parameters \mathbf{p} (which in addition to Reynolds number could also contain streamwise and/or spanwise wave numbers in periodic domains). Given an approximation \mathbf{u}^k to the fixed point, an improved one, $\mathbf{u}^{k+1} = \mathbf{u}^k + \Delta\mathbf{u}^{k+1}$, is obtained again by solving the linear equation

$$\Phi_{\mathbf{u}}(\mathbf{u}^k, \mathbf{p})\Delta\mathbf{u}^{k+1} = -\Phi(\mathbf{u}^k, \mathbf{p}). \quad (4.1)$$

In the hook-step method, solving this linear system is done by considering the usually much better behaved optimization problem

$$\min_{\Delta\mathbf{u}^{k+1}} \|\Phi_{\mathbf{u}}(\mathbf{u}^k, \mathbf{p})\Delta\mathbf{u}^{k+1} + \Phi(\mathbf{u}^k, \mathbf{p})\|. \quad (4.2)$$

The GMRES method is then again used to solve this problem [130]. This by itself is not enough, as the correction vector $\Delta\mathbf{u}^{k+1}$ could still have a large norm. Therefore, the search for an optimum in the Krylov subspace is restricted to solutions where $\|\Delta\mathbf{u}^{k+1}\|$ is small.

For the above problem, the gain in convergence is enormous: while a direct method typically requires that the deviation of the initial estimate from the fixed point be as small as 10^{-6} , the Newton hook-step method can converge if the initial state is within 10^{-3} of the fixed point. This gain in efficiency suffices to find fixed points by starting from points during the time-evolution where $d\mathbf{u}/dt$ is small. Other methods use embeddings into a family of flows within which the desired states can be tracked via homotopy [26, 43, 80, 138].

In systems with a subcritical transition, the perturbations must exceed a critical level before they can induce turbulence. Moving away from the laminar profile, one finds that the time it takes to return increases, and it diverges when the boundary is crossed. In the shear flows in which this method for determining the boundary was introduced, the turbulent state was not persistent but transient [56], and the lifetimes showed a sensitive dependence on initial conditions. Since the boundary separates a region with regular variations in lifetimes from one with chaotic variations it was called the ‘edge of chaos’ [115]. States on the boundary do not become turbulent and they do not return to the laminar state. They are permanently trapped in this manifold, but within it converge to

relative attractors (relative, because they are attracting only within this manifold). The states thus reached are called 'edge states'. The boundary between laminar and turbulent is then formed by the stable manifold of these edge states. The motion within the edge of chaos can be quite involved [39].

Tracking edge states is relatively straightforward, and requires only a program for time integration [115, 129]. It has been used in various flow systems, including pipe flow [104, 106, 137], plane Couette flow [107], in boundary layers [25, 38] and also in plasma physics [20]. In the cases of the usual saddle-node bifurcation with a real node and a saddle of co-dimension one, the edge state is the saddle and the boundary is the stable manifold. In more complicated cases, where many states appear in saddle-node bifurcations, there usually is only one saddle with a stable manifold of co-dimension one. Edge tracking will then select this state among all the lower-branch states [107]; an example is shown in Fig. 5.

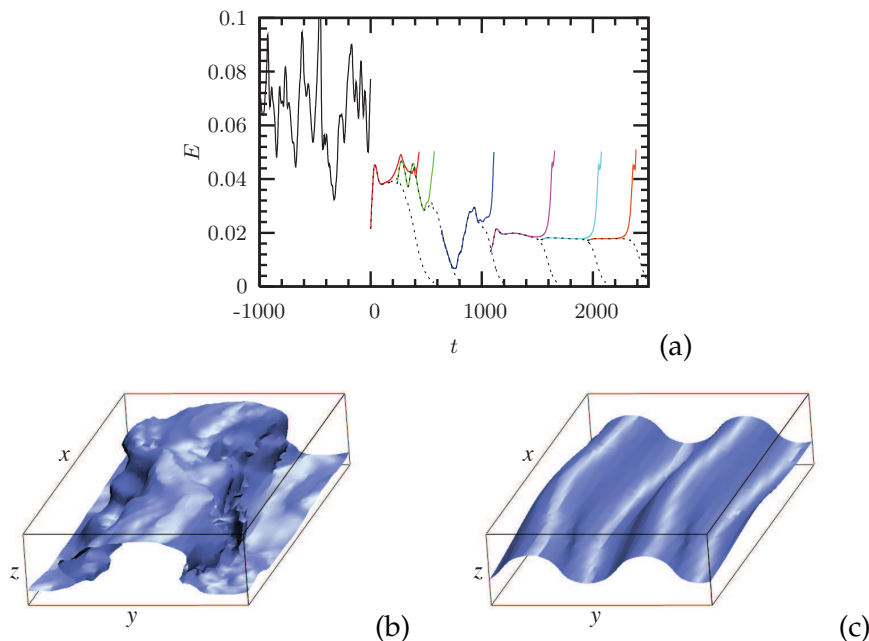


Figure 5: Edge tracking and coherent structures in plane Couette flow. Frame (a) shows the energy of the flow as a function of time. For negative times the time-evolution is unconstrained, and for positive time edge tracking is applied. The different colors indicate different segments that approximate the edge state. The flow fields represented by the surface where the downstream velocity component vanishes are shown in frames (b) and (c): for the initial turbulent state the surface is rough, indicating the small scale turbulence, but for the final state it is smooth (figure from [107]).

4.2 Thin film flows

Micro- and macrofluidic flows involving free surfaces, e.g., film flows along inclined walls or the motion of drops on homogeneous or heterogeneous substrates, can often

be described by the *thin film* or *lubrication* equations. These are obtained from the full Navier-Stokes description via a long-wave approximation [86]. The equations are usually combined into a single evolution equation for the fluid thickness profile $h(\mathbf{r}, t)$:

$$\partial_t h = \nabla \cdot \{ Q(h) \nabla p(h, \mathbf{r}) - \mu(h) \mathbf{e}_x \}, \quad (4.3)$$

where $\mathbf{r} = (x, y)$ is the position on the surface below the fluid, \mathbf{e}_x is the unit vector in x -direction and t is time. Furthermore, $p(h, \mathbf{r})$ is a generalized pressure and $Q(h)$ is the fluid flux, or mobility function. For Poiseuille flow without slip at the substrate, $Q(h) = h^3 / 3\eta$, where η is the dynamic viscosity of the fluid. In Eq. (4.3), $\mu(h)$ represents a general lateral driving force, e.g. $\mu(h) = \alpha g \rho Q(h)$ for a gravitational acceleration g driving the flow along an incline of small inclination angle α , and liquid density ρ .

Continuation techniques have been successfully applied to the class of equations (4.3) in cases when they can be reduced to ODEs, in which case the package AUTO can again be used. Non-uniform solutions emerge from the trivial flat film state via a Hopf bifurcation [123] and continuation is normally started there. This allows one to use AUTO to obtain various families of (i) periodic solutions that correspond to travelling waves, and arrays of sitting or sliding drops; (ii) homoclinic orbits that correspond to solitary waves, and individual drops; (iii) heteroclinic orbits that correspond to fronts; (iv) self-similar solutions related to film rupture; and to track the various occurring bifurcations [101–103, 122, 123, 125, 132]. Other continuation packages are used as well [23, 24, 61]. Some ready-to-run examples can be found in the supplementary material of [58]. A set of AUTO tutorials and ready-to-run files for sitting and sliding drop solutions of thin film equations, can be obtained via www.uwethiele.de.

At present, much less is known about the solution and bifurcation structure in the three-dimensional case, i.e., with two spatial dimensions as in Eq. (4.3). Here, most publications focus on studies of thin film dynamics by time integration, e.g., for heated films [11, 52, 84], and the dewetting of partially wetting films [10, 85, 110]. To our knowledge only [6] have developed suitable continuation tools, presenting a common framework to perform time-stepping and continuation tasks for thin film equations containing a bi-Laplacian. The time stepping is based on an exponential propagation scheme and requires the computation of the exponential of the Jacobian matrix. The exponential of the matrix is not directly computed, but its action on vectors is estimated using projections on small Krylov subspaces [93]. For the continuation, a tangent predictor/secant corrector scheme is used, leading in both steps to linear systems involving the Jacobian matrix. This system is reduced using the Cayley-Arnoldi method, as also employed in the exponentiation of the Jacobian [75, 141]. Furthermore, knowledge of the leading eigenvalues facilitates the stability analysis and allows one to detect bifurcations.

The package developed by [6] has up to now only been applied to three-dimensional liquid ridges and drops on an inclined heterogeneous substrate [4, 5]. Typical examples of a bifurcation diagram and corresponding steady profiles of stable and unstable drops are given in Fig. 6. In particular, Fig. 6 shows spanwise invariant ridges (thin red lines),




Figure 6: (a) Bifurcation diagram for steady ridges (SR, thin red lines), and steady drops (SD, black lines) pinned by a hydrophobic line defect of strength $\epsilon = 0.3$ for $L = 27.8$ showing the L^2 norm $\|\delta h\|$ of steady solutions as a function of the driving force μ . Solid (dashed) lines denote stable (unstable) solutions. The SD_3 branch becomes stable in a Hopf bifurcation at $\mu = 0.064$ (not shown). Downward and upward pointing triangles indicate stick-slipping ridge (SSR) and stick-slipping drop (SSD) solutions, respectively, as obtained by time integration. For the remaining parameters, explanation of the remaining symbols, details of the pressure term and further analysis of the system see [5]. (b) Steady states at locations indicated by open squares in Fig. 6a in terms of contours of constant $h(x, y)$, (left panel) at $\mu = 0.00082$ on the drop branch SD_1 , (left middle panel) on the drop branch SD_1 at $\mu = 0.00110$, (right middle panel) at the saddle-node $\mu_{sn2}^{3d} = 0.00582$ on SD_1 , and (right panel) at the sniper bifurcation at $\mu_{sn4}^{3d} = 0.00724$ on SD_2 (figures from [5]).

and spanwise modulated or drop solutions (heavy black lines). Stable and linearly unstable solutions are indicated by solid and dashed lines, respectively. Downward and upward pointing triangles indicate stick-slipping ridge (SSR) and drop (SSD) solutions, respectively [5].

The limitations of the presently available tools for thin film equations are also well-illustrated by the example in Fig. 6: Although the steady states can be continued, no tools exist to continue time-periodic solutions on the branches of stick-slipping drops and ridges. This leaves important open questions, such as the detailed bifurcation structure related to the emergence of the spanwise invariant SSR branch from the SR_1 branch (cf. [124]). At the chosen parameters the SR_1 becomes unstable through a subcritical Hopf bifurcation; therefore not all branches of time-dependent solutions can be obtained through time integration. At other parameter values the SSD branch emerges from the SSR branch through a subcritical bifurcation, whose details are also not known. Contin-

uation of the time-periodic solutions of large period would also allow for a better understanding of the observed transitions between sniper and homoclinic bifurcations, which is also relevant for depinning transitions in the context of deposition patterns.

4.3 Channel and shoal development in tidal embayments

Tide-dominated inlet systems, such as those located in the Wadden Sea, display complex channel and shoal systems [42]. To gain insight into the physical mechanisms responsible for these patterns, and their sensitivity to parameter values, a two-dimensional nonlinear idealized model is developed and analyzed using a bifurcation analysis [120]. The geometry of the tidal inlet is taken to be rectangular with width B and length L . The local water depth is given by $H - h + \zeta$, where H is the depth at the entrance, h the bed level and ζ the tidal elevation. The sea bottom of the basin consists of uniform fine sand.

The evolution of this coastal system is described by a coupled system of equations that describes the water motion, sediment transport and bed update. Focussing on transport of suspended sediment by diffusive processes only, and assuming that the bed profile only changes due to tidally averaged erosion and deposition, the resulting dimensionless equations are given by (see [120] for details)

$$\zeta_t + [(1-h)u]_x + [(1-h)v]_y = 0, \quad (4.4a)$$

$$\zeta_x = 0, \quad (4.4b)$$

$$\zeta_y = 0, \quad (4.4c)$$

$$v_{xt} - u_{yt} = - \left[\frac{ru}{1-h+h_0} \right]_x + \left[\frac{rv}{1-h+h_0} \right]_y, \quad (4.4d)$$

$$a \{ C_t - \kappa C_{xx} - \kappa C_{yy} \} = u^2 + v^2 - C, \quad (4.4e)$$

$$h_\tau + \langle C \rangle = - \langle u^2 + v^2 \rangle + \mu \langle [h - h_{eq}]_{xx} + [h - h_{eq}]_{yy} \rangle. \quad (4.4f)$$

The depth-averaged dimensionless velocities are represented by u and v in the longitudinal (x) and lateral (y) direction, respectively. The quantity C is the depth-integrated suspended sediment concentration and h is the bed level.

In this system of equations two time scales appear: the short time scale (denoted by t) with a typical period of 12.5 hours and the much larger morphodynamic time scale (denoted by τ) in the order of years. The parameter r denotes the dimensionless friction parameter, h_0 is introduced to normalize bottom friction when approaching zero water depth, a is the ratio of the deposition time scale over the tidal time scale, and κ is the horizontal eddy diffusivity coefficient. Tidal averaging of physical quantities is denoted by $\langle \cdot \rangle$.

As the coastlines are fixed, there are no normal fluxes of water or sediment through these boundaries. The bed is assumed fixed at the entrance, and the water motion is driven by a prescribed vertical tide consisting only of the M_2 tidal constituent, i.e., $\zeta = \cos(t)$.

The temporal dependence of the variables on the fast time scale is discretized by using two tidal constituents and a residual component (a Fourier expansion in time, truncated after the second Fourier mode); their spatial dependence is discretized using a Galerkin approach (25 Fourier modes in the lateral direction and 30 Chebychev polynomials in the longitudinal direction). Hence \mathcal{M} and \mathcal{L} in (2.1) are of size 4500×4500 , with \mathcal{M} a singular matrix and \mathcal{L} a full matrix.

[108] and [121] investigated this model, and found semi-analytic width-invariant morphodynamic equilibria. When the bottom friction parameter is increased above a critical value, these equilibria become unstable. Using a simple continuation technique, they constructed bifurcation diagrams showing the number of nonlinear morphodynamic equilibria and their stability as a function of the bifurcation parameter r . An example is shown in Fig. 7, in which the equilibrium bed profiles are shown on the different branches.

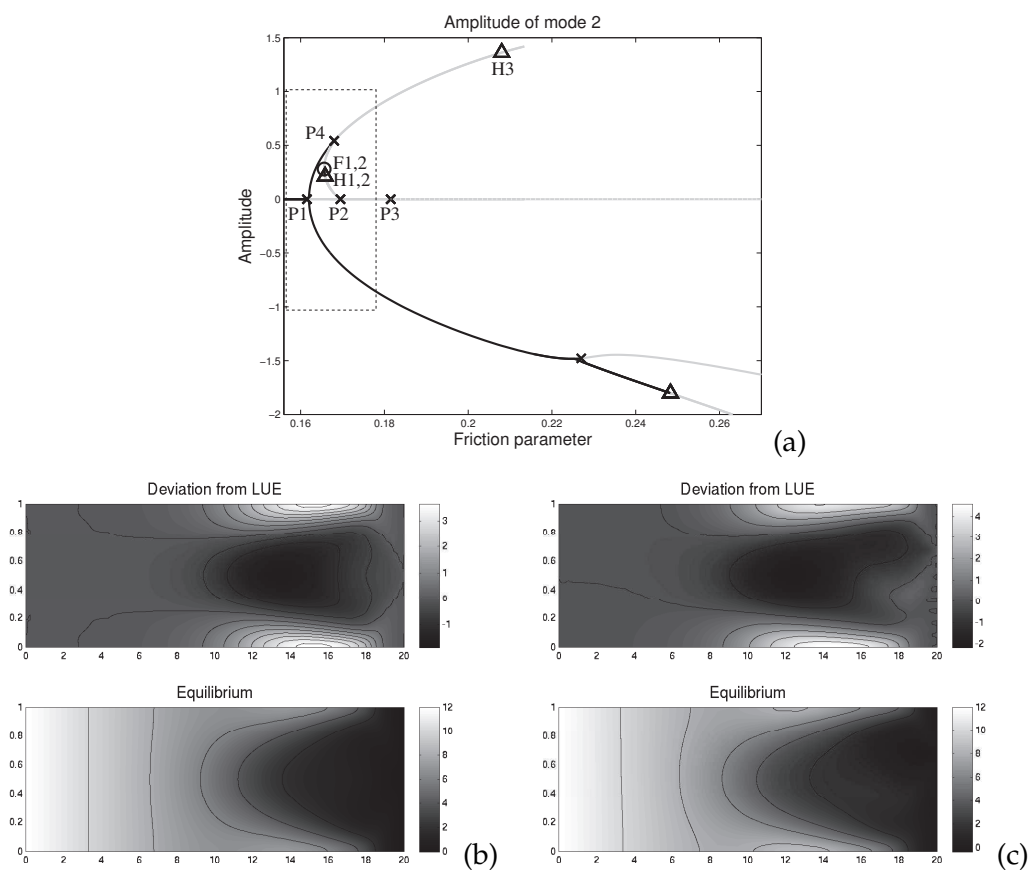


Figure 7: Bifurcation diagram for an embayment with a width of 1 kilometer. In (a) the amplitude corresponding to lateral mode 2 is plotted. The equilibrium bed level (bottom panel) and deviation from the constantly sloping trivial solution (top panel) for $r \sim 0.2$ and $r = 0.24$ are given in (b) and (c), respectively.

4.4 Rayleigh-Bénard convection

Rayleigh-Bénard convection is probably the earliest example of an analytical solution of a stability problem for viscous flow. The classical problem considers an infinite horizontal layer whose boundaries are maintained at constant temperatures, with the lower boundary considered hot and the upper one cold. The base state is motionless, with a linear temperature profile between the boundaries. When the temperature difference exceeds a critical value, the motionless state becomes unstable via a pitchfork bifurcation and slightly supercritical convective motion appears as pairs of counter-rotating two-dimensional convective rolls [62].

Buoyancy convection flows are usually described in the framework of the Boussinesq approximation [22]. This approximation consists of neglecting the temperature dependence of the fluid parameters, except for the density in the buoyancy term, which is assumed to vary linearly with temperature. The dimensionless equations are given by

$$Pr^{-1}\left(\frac{d\mathbf{v}}{dt} + \mathbf{v} \cdot \nabla \mathbf{v}\right) = -\nabla p + \nabla^2 \mathbf{v} + RaT\mathbf{e}_z, \quad (4.5a)$$

$$\nabla \cdot \mathbf{v} = 0, \quad (4.5b)$$

$$\frac{dT}{dt} + \mathbf{v} \cdot \nabla T = \nabla^2 T, \quad (4.5c)$$

where \mathbf{v} is the velocity field, T the temperature, p the pressure and \mathbf{e}_z the unit vector in vertical direction. The parameters in these equations are the Prandtl number $Pr = \nu/\kappa$, the Rayleigh number $Ra = g\beta\Delta TH^3/\nu\kappa$, and geometric relations, e.g., aspect and width ratios. Here ν is the kinematic viscosity, κ is the thermal diffusivity, g is the gravitational acceleration, β is the thermal expansion coefficient, ΔT is the characteristic temperature difference, and H is the characteristic length, e.g., the height of the horizontal fluid layer.

For convection in the Rayleigh-Bénard configuration, which allows for a quiescent base state, it can be shown that the instability threshold is defined only by the Rayleigh number. The onset of secondary instabilities, however, depends strongly on the Prandtl number and the geometry. This dependence is the most common objective of stability studies of convective flows. Several variants of the classic Rayleigh-Bénard problem in an infinite horizontal layer can be defined by the boundary conditions at the upper and lower bounding plates on the velocity, which can be either stress-free or no-slip.

A detailed analysis of the stability properties of the two-dimensional Rayleigh-Bénard convection problem can be found in [65]. The academic case of two stress-free boundaries has a simple analytic solution, which yields $Ra_{cr} = 27\pi^4/4 \approx 657.511$. The most unstable disturbance is two-dimensional with a critical spatial wavenumber of $\alpha_{cr} = \pi/\sqrt{2} \approx 2.221$. The physically achievable no-slip condition on either the lower or on both boundaries is slightly more complicated. The well-known results are $Ra_{cr} = 1100.657$, $\alpha_{cr} = 2.682$ for no-slip lower and stress-free upper boundaries and $Ra_{cr} = 1707.762$, $\alpha_{cr} = 3.117$ for two no-slip boundaries. These three pairs of critical values are the necessary preliminary benchmark results for any stability solver dealing with convective flows. The numerical problem

is formulated in a rectangle of width $2\pi/\alpha_{cr}$ and unit height with periodic horizontal boundary conditions. Alternatively, taking into account that the instability sets in as a pair of counter-rotating symmetric rolls, one can consider a smaller problem, with the width set to π/α_{cr} , and reflection-symmetric horizontal boundary conditions.

A large number of experimental and computational studies have been carried out of the primary and secondary instabilities. The reader is referred to books [46, 62, 70] and review papers [12, 16, 68, 69, 74]. Some benchmark-quality results on primary Rayleigh-Bénard instability for two-dimensional flows are available [45] as well as for three-dimensional flows [21, 45, 88]. Here, the system of PDEs are reduced to systems of ODEs through Galerkin expansions. Examples of three-dimensional spatial structures of the temperature perturbation at the first bifurcation point are shown in Fig. 8.

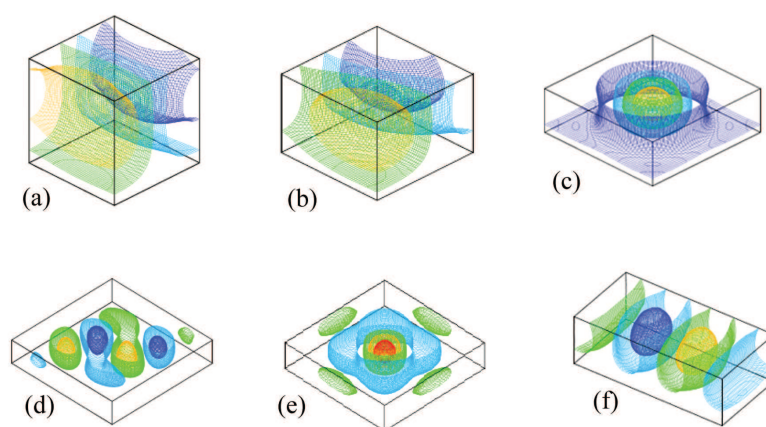


Figure 8: Isosurfaces of the temperature perturbation of primary Rayleigh-Bénard instability in three-dimensional rectangular containers with different height and width ratios. (a) $A_x = A_y = 1$; $Ra_{cr} = 4347$; (b) $A_x = A_y = 1.5$; $Ra_{cr} = 2506$; (c) $A_x = A_y = 3$; $Ra_{cr} = 1728$; (d) $A_x = A_y = 6.5$; $Ra_{cr} = 1597$; (e) $A_x = A_y = 6$; $Ra_{cr} = 1598$; (f) $A_x = 4$, $A_y = 1$; $Ra_{cr} = 2645$. The Rayleigh-Bénard instability turns a motionless non-uniformly heated fluid into a convective flow organized in ascending/descending motion called convective rolls. Depending on the aspect / width ratio and initial conditions the rolls can have different symmetries and can be aligned either parallel to a wall, parallel to a diagonal, or “centralized” with 90° rotational symmetry. The temperature perturbation isosurfaces shown in the figure repeat the flow symmetry (reprinted with permission from [45]).

The importance of convective flows in many scientific and technological areas has led to their intensive study in various configurations, under many different heating conditions, and sometimes with additional forcing, such as an applied magnetic field. An impressive collection of results, including state-of-the-art stability and path-following studies, can be found in the recent book of [70].

4.5 Convection mitigated thermal runaway

In this application, a chemical reactor is considered with an exothermic chemical reaction. To remove the heat, the walls of the reactor are cooled and the fluid is stirred. If

the stirring mechanism breaks down, then the fluid will develop a hot spot away from the cooled walls. Due to the exponential increase in reaction rate with respect to temperature, thermal runaway can occur and lead to explosion. Natural convection due to the temperature differences can induce a flow that will circulate the fluid near the cooled walls and delay or prevent the explosion compared to a situation where there is no flow. In this problem, bifurcation analysis can provide guidance on the design of the reactor by delineating regimes in parameter space where the un-stirred reactor will explode and where the passive convection behavior will be enough to prevent explosion.

The model often used comes from previous studies of explosion and natural convection [78]. The full model includes the Navier-Stokes equations for the flow and coupled heat and chemical species balances. To reduce the number of dimensionless parameters that govern the problem, it was assumed that reactant depletion is negligible (eliminating the species balance) and the Frank-Kamenetskii representation of Arrhenius kinetics [44] as an exponential dependence on temperature (removing the dependence of the chemical reaction activation energy) is used. The equations are the same as (4.5) except for an extra term proportional to the Frank-Kamenetskii number F_K in the right hand side of (4.5c).

With the Boussinesq approximation for the density dependence on temperature, and assuming other physical properties are independent of temperature, focus is on the behavior of the flow while varying two dimensionless parameters: the Rayleigh number Ra representing the ratio of buoyancy forces to stabilizing viscous and conduction effects, and F_K which is the ratio of the heat generated by reactions to the dissipation of conduction.

The problem was studied for a cylinder with unit aspect ratio, no-slip boundary conditions for the velocity on all sides, a cold outer wall held at dimensionless temperature $T = 0$, and adiabatic top and bottom surfaces. The problem was discretized with a pressure-stabilized Galerkin Finite Element methods, a mesh of hexahedral elements, and with trilinear basis functions for all unknowns. The Jacobian matrix was constructed analytically using Automatic Differentiation, specifically the Sacado package in Trilinos. The parallel data structures, linear solver, preconditioner, nonlinear solver, eigensolver, continuation, and bifurcation tracking algorithms all come from the packages within the Trilinos solver framework.

Fig. 9a shows vertical velocity contours on half the domain, with a hot spot at the center and cooling at the walls generating a recirculation. For a fixed value of $Ra = 10^6$, a continuation run was done for the reaction rate parameter F_K , and shown in Fig. 9b. A fold point located near $F_K = 16$ shows that, for a fluid and geometry corresponding to this Ra , a reaction rate over this critical value would result in thermal runaway. More design information comes from the fold-tracking calculation in Fig. 9c, where the curve of fold points is calculated as a function of the Rayleigh number. The curve intercepts the $Ra = 0$ axis at the known analytic solution of a critical value of $F_K = 2$, below which conduction alone is enough to remove the heat and prevent explosion. Here the regions in parameter space with qualitatively different behaviors (whether or not the reactor explodes) are delineated by a single bifurcation-tracking continuation calculation. This is exactly the




Figure 9: (a) Contour plot of vertical velocity variable for the convection mitigated thermal runaway problem. This natural convection pattern is driven by heat generated by a chemical reaction and cooling at the outer walls. (b) The result of a pseudo-arclength continuation run for the convection mitigated thermal runaway problem. For a fixed value of $Ra=10^6$, the solution curve, parameterized by the dimensionless reaction rate F_K , shows a fold point corresponding to thermal runaway. (c) A fold-tracking run for this same problem directly delineates the regions in two-parameter space between safe and unsafe designs.

type of plot needed by engineers to know how to design and operate the reactors so that, in the event of a failure in the stirring mechanism, the system will avoid thermal runaway.

This relatively straightforward calculation required a few hours on 128 processors for a discretization of over one million unknowns, and little human intervention. This is in stark contrast with the human and computational resources that would be needed to extract this same information by repeated transient simulations covering the two-parameter space.

4.6 Swirling flows in Francis turbines

Due to the variable demand on the energy market and the limited storage capabilities, modern turbines are often forced to operate far from optimal efficiency point. The swirling flow exiting the runner of the turbine is further ingested by the draft tube where

it is decelerated in order to convert as much as possible the excess of dynamic pressure into static pressure. The deceleration of the swirling flow in the draft tube may lead to the occurrence of vortex breakdown with the development of a central quasi-stagnation region or severe pressure fluctuations associated with a precessing helical vortex. The onset of vortex breakdown is characterized by a transition between a jet-like profile and a wake-like profile with a local minimum on the axis [83]. Understanding the mechanism of vortex breakdown is extremely important in developing flow control techniques.

Numerical simulations of three-dimensional turbulent swirling flow in the particular case of a flow downstream from a Francis turbine runner, using very powerful available computer resources, take about two months for the investigation of only one operating point. One way to help interpret the significant amount of data is the application of stability and bifurcation analysis in order to recover the most relevant information using much less computation time for each set of parameters.

The equations of motion in the primitive variable formulation for an incompressible flow are the Navier-Stokes equations in cylindrical coordinates. The computational domain corresponds to the diffuser shape (Fig. 10a) with the wall radius r_{wall} given with respect to the inlet radius, $R_{inlet} = \sqrt{2}$, as [119]

$$\frac{r_{wall}(z)}{R_{inlet}} = 1 + \frac{1}{2} \left(1 - \cos\left(\frac{\pi z}{6}\right) \right), \quad \text{for } 0 \leq z \leq 6. \quad (4.6)$$

There are two cylindrical segments of radius R_{inlet} for $-2 \leq z \leq 0$ and of radius $2R_{inlet}$ for $6 \leq z \leq 8$, added upstream and downstream of the diffuser. The inlet swirl corresponds to a Burgers vortex, with velocity components given by

$$V_{zin} = 1; \quad V_{rin} = 0; \quad V_{\theta in}(r) = \frac{\Omega R^2}{r} \left(1 - e^{-\frac{r^2}{R^2}} \right), \quad (4.7)$$

where Ω is the angular velocity at the axis, R is the vortex characteristic radius and V_{zin} , V_{rin} , $V_{\theta in}$ are the inlet axial, radial and swirl components of the velocity. The asymptotic behavior of the swirl velocity is

$$V_{\theta} \approx \Omega R \frac{r}{R} \quad \text{if } r \leq R \quad \text{and} \quad V_{\theta} \approx \Omega R \frac{R}{r} \quad \text{if } r \geq R, \quad (4.8)$$

corresponding to the forced vortex and free vortex, respectively. The characteristic radius is taken as constant $R = 0.4$, but the swirl intensity is increased as $\xi = 2\Omega R_{inlet} / V_{zin}$ and for our purpose takes the values 1, 2 or 3. A simple radial pressure equilibrium condition on the outflow section, $\partial p / \partial r = \rho V_{\theta}^2 / r$, is used in the numerical computations, although its accuracy deteriorates when a vortex ring is convected downstream through the outlet section.

Steady axisymmetric flows in Francis turbines and their linear stability were investigated in [36] and [119]. When the swirl intensity is $\xi = 1$, the streamlines in a meridional half-plane for a low swirl intensity (Fig. 10b) reveal a steady flow: no vortex breakdown

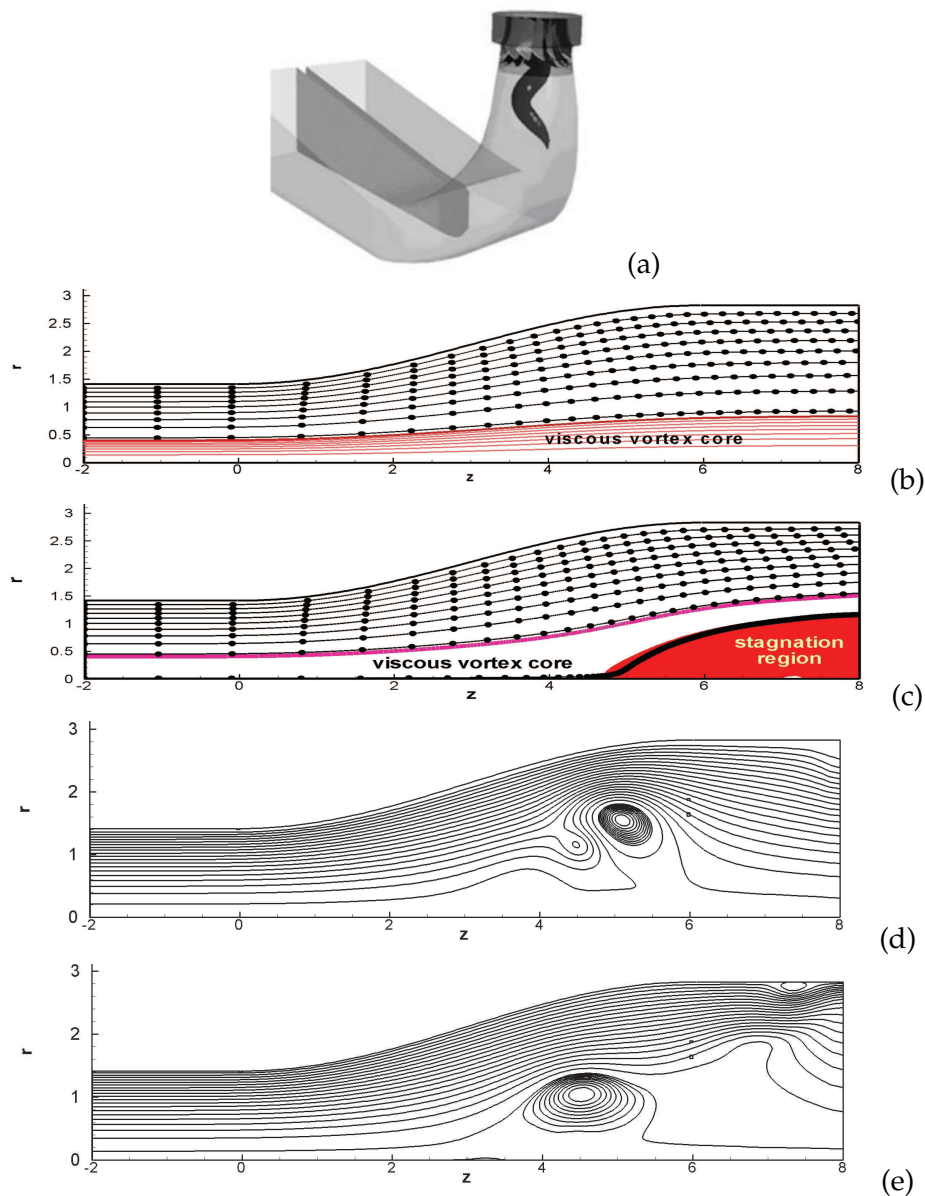


Figure 10: (a) Schematic of the vortex flow leaving the runner of the Francis turbine [77]. (b) Steady swirling flow without vortex breakdown for the swirl intensity $\zeta = 2\Omega R_{inlet}/V_{zin} = 1$. (c) Steady swirling flow with central stagnation region for $\zeta = 2$. (d-e) Two snapshots of the streamlines for unsteady swirling flow showing distinct vortex rings for $\zeta = 3$ (figure from [119]).

occurs. As the swirl intensity increases to $\zeta = 2$, steady axisymmetric vortex breakdown occurs at the end of the diffuser, with the development of a central stagnation region (Fig. 10c). We recognize vortex breakdown as an internal separation away from a solid

boundary. The streamline originating close to the axis in Fig. 10c displays severe flow deceleration as the fluid particles approach the tip of the stagnation region, then the velocity practically vanishes. The streamline originating at the vortex characteristic radius on the inlet section, separating the inner so-called viscous core from the outer inviscid-like flow, is shown as a thick red solid line. When the swirl intensity is further increased, a Hopf bifurcation can occur and the flow becomes highly unsteady, with periodic vortex rings development and convection downstream. The streamlines in the meridional half plane (Fig. 10d-e) for $\zeta=3$ at consecutive moments in time show the formation of periodic vortex rings. The phenomenon is also associated with a significant shifting in time of the radial discharge distribution on the outlet section.

The diagrams reveal the swirling flow characteristic for increasing intensity of the swirl very distinctly. These secondary structures are of great practical and theoretical importance for Francis turbines operating at partial discharge, where the bifurcation of the solution leads to the formation of an unsteady helical vortex breakdown in the draft tube cone. Vortex breakdown, which appears as a jump bifurcation due to structural instability of swirling flows when the solution locally fails to exist, is associated with severe pressure fluctuations and mechanical vibrations. In the presence of viscosity, these structures are related to global bifurcations for dynamical systems. Control of such recirculation zones is therefore necessary and suitable methodologies to control their generation are crucial for the design of new devices.

4.7 Flow in a diverging channel

We consider the stability of the flow in a channel with non-parallel walls, which has obvious applications to the internal transport of fluids in regions where the containing vessels exhibit variations in internal dimension. That said, our principal interest is in a more fundamental question: are (mathematically) simple similarity solutions useful predictors of nonlinear behaviour?

The classic Jeffery-Hamel similarity solution describes the flow of an incompressible, Newtonian fluid through a (two-dimensional) diverging wedge of infinite extent with fixed separation angle 2α . The bifurcation structure of this flow is well understood and the first bifurcation with increasing Reynolds number is a subcritical pitchfork that breaks the mid-plane symmetry [118]. The bifurcation structure of the flow through diverging channels of finite extent is not as well understood and there is contradictory evidence in the literature, albeit often inferred from IVP methods. Although the critical Reynolds number in finite domains is close to the predictions of the similarity solution, the criticality and even the nature of the bifurcation is disputed. The aim of the work reported in detail in [54] was to use bifurcation detection and tracking techniques to perform a comprehensive analysis of the bifurcation structure for the flow in a diverging channel of finite extent and to determine its connection, if any, to that of the Jeffery-Hamel similarity solution. The work is part of a larger research programme with the aim of understanding whether the complex behaviour found in similarity solutions of the Navier-Stokes

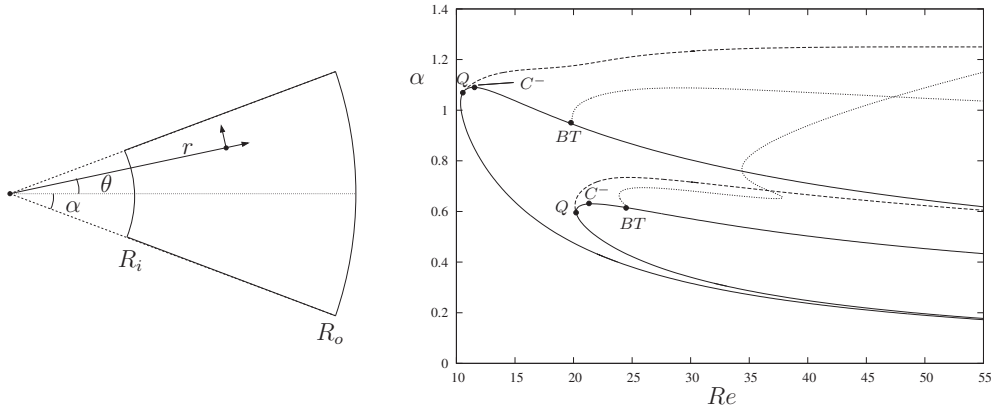


Figure 11: (a) Two parallel walls separated by a fixed angle 2α for which $R_i \leq r \leq R_o$, where r is the radial coordinate of a set of polar coordinates centred at the (virtual) intersection of the two walls. (b) Location of pitchfork (solid), limit-point (dashed) and Hopf (dotted) bifurcations in the (Re, α) plane. We conjecture that there is an infinite nested sequence of similar curves. Here C^- denotes a coalescence point, Q denotes a quartic point and BT denotes a Bogdanov-Takens point. Results are for $R_o/R_i=100$, with a quadratic inlet profile and a pseudo-traction free outlet (reprinted with permission from [54]).

equations, which are typically defined over semi-infinite domains, is ever relevant in physically-realizable finite domains.

The Navier-Stokes equations were solved in a wedge of finite extent via a plane polar coordinate system (Fig. 11a), and using a second-order accurate Galerkin method. No slip, zero velocity, boundary conditions were imposed on the wedge sidewalls and a variety of different inlet and outlet conditions were considered. For the results presented here, we imposed an ad hoc quadratic profile as a Dirichlet inlet condition and allowed the outlet to remain pseudo-traction free (the natural boundary condition for our finite element formulation). The discrete form of the governing equations was assembled and solved by `oomph-lib` [55]. Bifurcations were detected by solving the associated generalized eigenproblem with an Arnoldi method (ARPACK) and then tracked by forming the appropriate extended system and continuing in a second parameter.

A sparse direct solver (SuperLU) was used to solve the linear system that arises at each step of the global Newton method and our standard resolution was approximately 45,000 degrees of freedom for the base flow. The number of degrees of freedom is approximately doubled when continuing limit points and pitchfork bifurcations and approximately tripled when continuing Hopf bifurcations. Non-uniform spatial adaptivity was used to minimize the number of degrees of freedom required for an accurate solution. Typical memory usage was 200Mb–1.2Gb depending on the overall problem size (the greatest memory usage occurs when continuing Hopf bifurcations) and the bifurcation diagram shown in Fig. 11b can be generated in a few hours on a standard commodity PC. For the user, the most time-consuming operation is validating the computations and designing the logic to detect and switch the tracking for the different bifurcation types, which is not done automatically.

The bifurcation structure in the finite domain is complex, exhibiting an entire family of nested curves, each corresponding to a *supercritical* pitchfork bifurcation. Nonetheless, the outermost such curve approaches the path of the Jeffery-Hamel *subcritical* pitchfork bifurcation, $Re \sim 3\pi/\alpha$, as $\alpha \rightarrow 0$. In addition, isolated solution branches (not shown in Fig. 11b) were also found that correspond to wave-like solutions found in the infinite domain by [60]. Asymptotic analysis revealed that in the limit of small α , the critical eigenfunction in the finite domain can be constructed by the superposition of two infinite-domain eigenfunctions, which is enough to ensure that the inlet and outlet boundary conditions are satisfied. This result provides the connection between the finite and infinite domains and explains the change in criticality of the bifurcation. Moreover, the supercritical bifurcation was found to be the generic behaviour, unless the boundary conditions were specifically chosen to force the similarity solution. In summary, the use of numerical bifurcation methods allowed us to bring together and explain previous discrepancies in the stability of the flow through a diverging channel within a consistent framework.

4.8 The Atlantic ocean circulation

Over the last decade, the application of dynamical systems methods to a hierarchy of models of the ocean circulation has become an important complementary approach for understanding the origin of spatio-temporal variability of mid-latitude three-dimensional ocean flows. Canonical situations are flows in a single-hemispheric spherical sector representing the North Atlantic, forced by an idealized steady double-gyre wind-stress field (with a typical amplitude τ), a restoring heat flux (with a typical pole-to-equator temperature difference ΔT), and a prescribed freshwater flux (with a typical amplitude σ). In the context of the decadal-to-interdecadal variability, two different limits have been well studied [30]. One of these limits, $\tau=0$, represents purely buoyancy-forced flows with focus on instabilities and transitions (multiple equilibria) of the meridional overturning circulation [32, 90, 91, 127]. The other well-studied limit, $\Delta T = \sigma = 0$, is that of purely wind-driven flows with a prescribed density field [31, 113, 114]. In [7-9], Jacobian matrix-free methods are used to compute bifurcation diagrams with existing ocean models.

In the matrix-based THCM ocean model [28] flows in a spherical domain bounded by longitudes ϕ_w and ϕ_e and by latitudes θ_s and θ_n with continental boundaries introduced by a land mask are considered. The ocean basin has a bottom topography h_b and is hence bounded vertically by $z = -D + h_b(\phi, \theta)$ and a nondeformable ocean-atmosphere boundary at $z = 0$. The flows in this domain are forced by a heat flux Q_H (in Wm^{-2}), a zonal wind stress field (τ^ϕ, τ^θ) (in Pa) and a virtual salt flux Q_S (in ms^{-1}). Both the fluxes Q_H and Q_S are of restoring type where the surface temperature and salinity are restored to prescribed functions T_S and S_S , using restoring time scales τ_T and τ_S , respectively. The wind-stress forcing (τ^ϕ, τ^θ) is prescribed from data.

Temperature and salinity differences in the ocean cause density differences according to the linear equation of state $\rho = \rho_0(1 - \alpha_T(T - T_0) + \alpha_S(S - S_0))$, where α_T and α_S are

the volumetric expansion coefficients and T_0 , S_0 and ρ_0 are reference quantities. The full (dimensional) equations of the primitive equation ocean model can be found in [28]. In the model, Laplacian friction is used where A_H and A_V are the constant horizontal and vertical momentum (eddy) viscosities, respectively. Although the general tracer mixing equations (including isoneutral mixing and the Gent-McWilliams representation of eddy mixing) were formulated in [28], in most cases constant horizontal and vertical diffusivities K_H and K_V , respectively, are used (the case $\eta_M = \eta_G = 0$ in [28]). Slip conditions are assumed at the bottom boundary, while no-slip conditions are applied at all lateral boundaries.

To compute steady flows as a function of parameters in a realistic Atlantic geometry, a domain $\phi \in [262^\circ, 350^\circ]$, $\theta \in [10^\circ, 74^\circ]$ N and $z \in [-4000, 0]$ m is chosen, with full bathymetry from the ETOPO10 data set. Annual mean wind-stress forcing is again taken from data of [131]. Furthermore, the restoring profiles of T_S and S_S are taken from the [71] data set and a restoring time scale of 75 days is used. The restoring temperature and salinity profiles, as well as the wind stress amplitude are multiplied by a homotopy parameter η .

For these computations, a $176 \times 128 \times 16$ grid (corresponding to $1/2^\circ$ horizontal resolution) is used [126] giving a dynamical system with 2,162,688 degrees of freedom. The homotopy parameter η is increased from $\eta = 0$ (zero forcing) to $\eta = 1$ (realistic forcing) giving the results in Fig. 12. To visualize the flow, the streamfunction of the zonally averaged (the meridional overturning stream function, Ψ_M) and vertically averaged (the barotropic stream function, Ψ_B) are plotted. The maximum of the meridional overturning increases with η following a power law (Fig. 12a). The pattern of the meridional overturning streamfunction in Fig. 12b indicates a single overturning cell, but there are regions of large gradients due to a combination of the buoyancy forcing and the bottom topography.

Two large gyres are seen in the barotropic stream function (Fig. 12c); the sub polar gyre with minimum Ψ_P and the subtropical gyre with a maximum Ψ_T . While the strength of the subtropical gyre still increases approximately linearly with η , the subpolar gyre strength depends nonlinearly on η and appears to saturate at a value just above 10 Sv (Fig. 12a). The pattern of the barotropic streamfunction Ψ_B in Fig. 12c indeed shows a very confined subpolar gyre and a much larger subtropical gyre with a flow near the western boundary qualitatively resembling the Gulf Stream.

5 Summary and outlook

In this paper, we have given an overview of current techniques of numerical bifurcation analysis (NBT) and their applications to fluid dynamics. As illustrated by the highlights, one can obtain a much more detailed picture of the different solutions regimes and transitions in fluid flows than one can obtain from transient simulations (obtained through initial value problem (IVP) methods) alone. IVP methods are, by their nature, limited to finding only stable attracting sets. Quite often in bifurcation problems, there are more

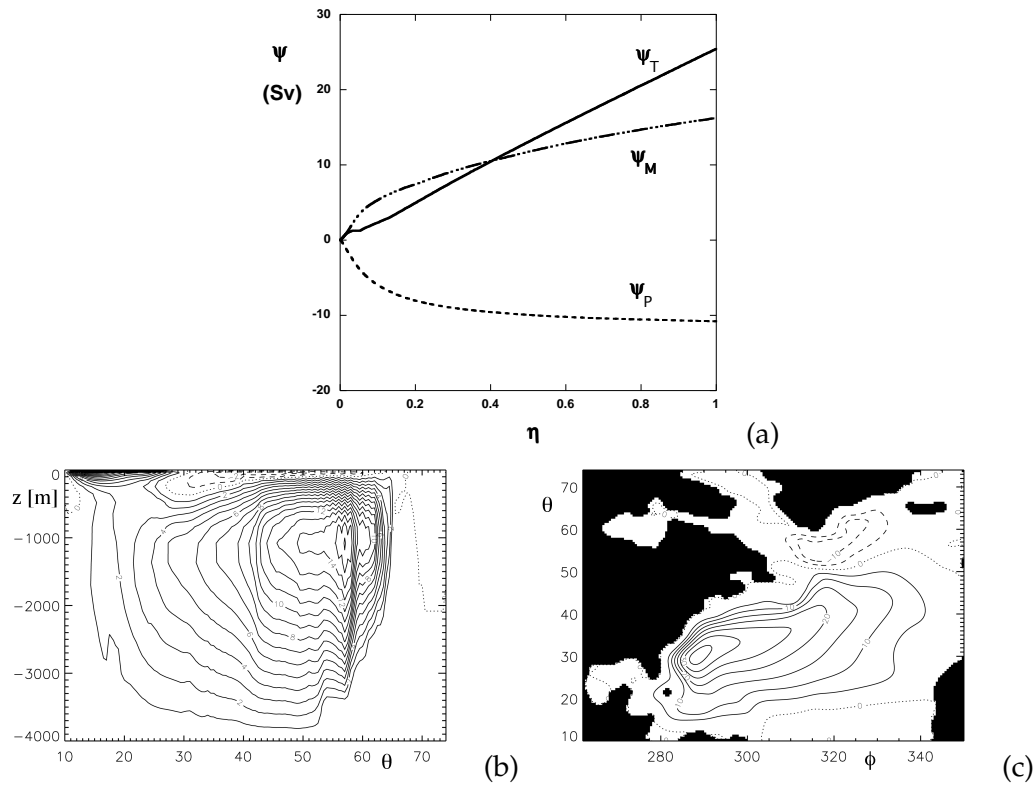


Figure 12: (a) Values of the maximum of the meridional overturning stream function ψ_M , the subtropical gyre ψ_T and the minimum of the subpolar gyre ψ_P versus η . (b) Contour plot of the meridional overturning streamfunction for $\eta=1$. (c) Contour plot of the barotropic streamfunction for $\eta=1$ (figures from [126]).

unstable solutions than stable ones and understanding the structure of all solutions can give a clearer picture of what is really going on. This is not so much an argument for NBT methods as it is for methods that are capable of finding unstable solutions. However, one of the major advantages of NBT methods is that they can be used to partition parameter space into disjoint regions (so-called regimes) in each of which the bifurcation diagrams are qualitatively the same. This is a powerful idea that can save huge amounts of computing time if it can be carried out successfully. Hence numerical bifurcation analysis techniques are an important complementary research tool.

Both matrix-free and matrix-based techniques are used in applications. For both, computations for substantial spatial resolutions can be performed on modern multicore hardware. The choice of either of these methods for a certain application depends on the availability of the Jacobian matrix and on the preconditioning strategy that is to be applied to solve the linear systems of equations. It is advantageous to use a numerical toolbox, for example Trilinos, to express the problem. Using such a toolbox, a Jacobian can be constructed by automatic differentiation, one linear system solver can easily be replaced by another, and the platform can also easily be changed. For matrix-based cases, efficient

preconditioners are crucial for the application of the techniques of numerical bifurcation analysis to fluid dynamics problems. Because of large fill-in, for example, memory usage may become a problem. In general, the stability study of a certain steady flow requires a good numerical resolution not only for the flow, but also for several leading eigenmodes. This can make requirements to numerical accuracy significantly more demanding than for a calculation of the steady state itself.

In the examples shown in Section 4 only the continuation of steady-states has been considered. In Section 3 several references were given to recent applications of the methods to compute periodic orbits described here. Periodic orbits show a much richer dynamics than fixed points and traveling waves, and are important because they contain information about the connection between the different fixed points and provide a more natural basis in which to describe turbulent dynamics. The computational requirements to obtain them is higher than for fixed points because each evaluation of the corresponding map involves a time-dependent computation. This can be partially overcome by the application of multiple shooting techniques. We hope to see more researchers in fluid mechanics to use these methods in the future. Given the much higher cost of computing invariant tori, only applications to two-dimensional problems have been considered so far. Improvements to the methods now employed, by using parallelism, are on their way.

Finally, there are many more areas in which continuation methods are used for low-dimensional models, such as in studies on optimization and predictability, which we have not addressed in this paper. The numerical techniques described here may be able to address important questions on control and predictability in high-dimensional models more efficiently.

Acknowledgments

This review was initiated during a workshop on “Tipping points in Complex Flows” during October 31–November 4, 2011. We thank the staff of the Lorentz Center in Leiden, the Netherlands, for making this workshop such a pleasant and interesting event. The workshop and the work of F.W. Wubs and H.A. Dijkstra was partially sponsored by the Netherlands Organization of Scientific Research (NWO) through the NWO-COMPLEXITY project PreKurs. The participation of F.I. Dragomirescu to the workshop was partially supported by a Grant of the Romanian National Authority for Scientific Research, CNCS-UEFISCDI, project number PN-II-RU-PD-2011-3-0153, 31/5.10.2011. Sandia National Laboratory is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy’s National Nuclear Security Administration under contract DE-AC04-94AL85000. The work of J. Sanchez-Umbria was supported by projects MTM2010-16930 and 2009-SGR-67. Valerio Lucarini acknowledges the financial support received from the EU through the FP7/2007-2013/ERC Grant agreement No. 257106, project Thermodynamics of the Climate System–NAMASTE.

References

- [1] Assemat, P., A. Bergeon, and E. Knobloch, 2008: Spatially localized states in Marangoni convection in binary mixtures. *Fluid Dyn. Res.*, 40, 852–876.
- [2] Barkley, D. and R. D. Henderson, 1996: Floquet stability analysis of the periodic wake of a circular cylinder. *J. Fluid Mech.*, 322, 215–241.
- [3] Batiste, O., E. Knobloch, A. Alonso, and I. Mercader, 2006: Spatially localized binary fluid convection. *J. Fluid Mech.*, 560, 149–158.
- [4] Beltrame, P., P. Hänggi, and U. Thiele, 2009: Depinning of three-dimensional drops from wettability defects. *Europhys. Lett.*, 86, 24 006, doi:10.1209/0295-5075/86/24006.
- [5] Beltrame, P., E. Knobloch, P. Hänggi, and U. Thiele, 2011: Rayleigh and depinning instabilities of forced liquid ridges on heterogeneous substrates. *Phys. Rev. E*, 83, 016305, doi:10.1103/PhysRevE.83.016305.
- [6] Beltrame, P. and U. Thiele, 2010: Time integration and steady-state continuation method for lubrication equations. *SIAM J. Appl. Dyn. Syst.*, 9, 484–518, doi:10.1137/080718619.
- [7] Bernsen, E., H. A. Dijkstra, J. Thies, and F. W. Wubs, 2010: The application of Jacobian-free Newton-Krylov methods to reduce the spin-up time of ocean general circulation models. *Journal of Computational Physics*, 229 (21), 8167–8179.
- [8] Bernsen, E., H. A. Dijkstra, and F. W. Wubs, 2008: A method to reduce the spin-up time of ocean models. *Ocean Modelling*, 20 (4), 380–392.
- [9] Bernsen, E., H. A. Dijkstra, and F. W. Wubs, 2009: Bifurcation analysis of wind-driven flows with MOM4. *Ocean Modelling*, 30 (2-3), 95–105.
- [10] Bestehorn, M. and K. Neuffer, 2001: Surface patterns of laterally extended thin liquid films in three dimensions. *Phys. Rev. Lett.*, 87, 046 101, doi:10.1103/PhysRevLett.87.046101.
- [11] Bestehorn, M., A. Pototsky, and U. Thiele, 2003: 3D large scale Marangoni convection in liquid films. *Eur. Phys. J. B*, 33, 457–467.
- [12] Bodenschatz, E., W. Pesch, and G. Ahlers, 2000: Recent developments in Rayleigh-Bénard convection. *Annu. Rev. Fluid Mech.*, 32, 709–778.
- [13] Bohr, T., M. H. Jensen, G. Paladin, and A. Vulpiani, 1998: *Dynamical Systems Approach to Turbulence*. Cambridge University Press, Cambridge.
- [14] Boronska, K. and L. Tuckerman, 2010: Extreme multiplicity in cylindrical Rayleigh-Bénard convection. II. Bifurcation diagram and symmetry classification. *Phys. Rev.*, E81, 036321.
- [15] Brown, P. N. and Y. Saad, 1990: Hybrid Krylov methods for nonlinear systems of equations. *SIAM J. Sci. Stat. Comput.*, 11 (3), 450–481.
- [16] Busse, F., 1981: Transition to turbulence in thermal convection. *Topics in Applied Physics* (eds. H.L. Swinney and J.P. Gollub), 45, 81–137.
- [17] Canuto, C., M. Y. Hussaini, A. Quarteroni, and T. A. Zang, 2006: *Spectral Methods. Fundamentals in single domains*. Springer.
- [18] Canuto, C., M. Y. Hussaini, A. Quarteroni, and T. A. Zang, 2007: *Spectral Methods. Evolution to complex geometries and applications to Fluid Dynamics*. Springer.
- [19] Carey, G. F., K. Wang, and W. Joubert, 1989: Performance of iterative methods for Newtonian and generalized Newtonian flows. *Int. J. Num. Meth. Fluids*, 9, 127–150.
- [20] Cassak, P. A., J. F. Drake, M. A. Shay, and B. Eckhardt, 2007: Onset of fast magnetic reconnection. *Physical Review Letters*, 98, 215 001 (4 pages).
- [21] Catton, I., 1972: The effect of insulating vertical walls on the onset of motion in a fluid heated from below. *Int. J. Heat Mass Transfer*, 15, 665–672.
- [22] Chandrasekhar, S., 1981: *Hydrodynamic and Hydromagnetic Stability*. N.Y.

- [23] Chang, H.-C., 1994: Wave evolution on a falling film. *Ann. Rev. Fluid Mech.*, 26, 103–136.
- [24] Chang, H. C., E. A. Demekhin, and D. I. Kopelevich, 1993: Nonlinear evolution of waves on a vertically falling film. *J. Fluid Mech.*, 250, 433–480.
- [25] Cherubini, S., P. De Palma, J. C. Robinet, and A. Bottaro, 2011: Edge states in a boundary layer. *Physics Of Fluids*, 23 (5), 051705.
- [26] Clever, R. M. and F. H. Busse, 1997: Tertiary and quaternary solutions for plane Couette flow. *J. Fluid Mech.*, 344, 137–153.
- [27] Coleman, T. F., B. S. Garbow, and J. J. Moré, 1984: Software for Estimating Sparse Jacobian Matrices. *ACM Transactions on Mathematical Software*, 10 (3), 329–345.
- [28] De Niet, A. C., F. W. Wubs, A. D. Terwisscha van Scheltinga, and H. A. Dijkstra, 2007: A tailored solver for bifurcation studies of ocean-climate models. *J. Comp. Physics*, 277, 654–679.
- [29] Dhooge, A., W. J. F. Govaerts, and Y. A. Kuznetsov, 2003: MatCont: A MATLAB package for numerical bifurcation analysis of ODEs. *ACM Trans. Math. Soft.*, 29, 141–164.
- [30] Dijkstra, H. A., 2005: *Nonlinear Physical Oceanography: A Dynamical Systems Approach to the Large Scale Ocean Circulation and El Niño*, 2nd Revised and Enlarged edition. Springer, New York, 532 pp.
- [31] Dijkstra, H. A. and C. A. Katsman, 1997: Temporal variability of the wind-driven quasi-geostrophic double gyre ocean circulation: Basic bifurcation diagrams. *Geophys. Astrophys. Fluid Dyn.*, 85, 195–232.
- [32] Dijkstra, H. A. and W. Weijer, 2003: Stability of the global ocean circulation: The connection of equilibria in a hierarchy of models. *J. Mar. Res.*, 61, 725–743.
- [33] Doedel, E., 1986: AUTO: Software for continuation and bifurcation problems in ordinary differential equations. Pasadena, USA, Report Applied Mathematics, California Institute of Technology.
- [34] Doedel, E. J., 1980: AUTO: A program for the automatic bifurcation analysis of autonomous systems. *Proc. 10th Manitoba Conf. on Numerical Math. and Comp.*, Vol. 30, 265–274.
- [35] Doedel, E. J., H. B. Keller, and J. P. Kernévez, 1991: Numerical analysis and control of bifurcation problems, Part I : Bifurcation in finite dimensions. *Int. J. Bifurcation and Chaos*, 1 (3), 493–520.
- [36] Dragomirescu, F. I., D. A. Bistrián, S. Muntean, and R. Susan-Resiga, 2009: The stability of the swirling flows with applications to hydraulic turbines, *Proceedings of the 3rd IAHR WG on Cavitation and Dynamic Problems in Hydraulic Mach. and Systems*, Brno - Czech Rep., Oct., 14-16, 2009, 15-24.
- [37] Duguet, Y., C. C. T. Pringle, and R. R. Kerswell, 2008a: Relative periodic orbits in transitional pipe flow. *Phys. Fluids*, 20 (11), 114102.
- [38] Duguet, Y., P. Schlatter, D. Henningson, and B. Eckhardt, 2012: Self-sustained localized structures in a boundary layer-flow. *Phys. Rev. Lett.*, 108, 044501.
- [39] Duguet, Y., A. P. Willis, and R. R. Kerswell, 2008b: Transition in pipe flow: the saddle structure on the boundary of turbulence. *Journal Of Fluid Mechanics*, 613, 255–274.
- [40] Eckhardt, B., H. Faisst, A. Schmiegél, and T. M. Schneider, 2008: Dynamical systems and the transition to turbulence in linearly stable shear flows. *Philosophical Transactions Of The Royal Society A-Mathematical Physical And Engineering Sciences*, 366 (1868), 1297–1315.
- [41] Eckhardt, B., T. M. Schneider, B. Hof, and J. Westerweel, 2007: Turbulence transition in pipe flow. *Annual Review Of Fluid Mechanics*, Vol 43, 39, 447–468.
- [42] Ehlers, J., 1988: *The Morphodynamics of the Wadden Sea*. Balkema, Rotterdam, 397 pp.

- [43] Faisst, H. and B. Eckhardt, 2003: Traveling waves in pipe flow. *Physical Review Letters*, 91 (22), 224 502.
- [44] Frank-Kamenetskii, D. A., 1939: Calculation of thermal explosion limits. *Acta. Phys.-Chim USSR*, 10, 365.
- [45] Gelfgat, A. Y., 1999: Different modes of Rayleigh-Bénard instability in two- and three dimensional rectangular enclosures. *J. Comp. Physics*, 156, 300–324.
- [46] Gershuni, G. Z. and E. Zhukhovitskii, 1976: *Convective Stability of Incompressible Fluids*. Keter, Jerusalem.
- [47] Gibson, J. F., J. Halcrow, and P. Cvitanovic, 2009: Equilibrium and traveling-wave solutions of plane Couette flow. *J. Fluid Mech.*, 638, 243–266.
- [48] Gladwell, M., 2000: *The Tipping Point*. Little Brown.
- [49] Golub, G. H. and C. F. Van Loan, 1983: *Matrix Computations*. The Johns Hopkins University Press, Baltimore, U.S.A.
- [50] Golubitsky, M., I. Stewart, and D. G. Schaeffer, 1988: *Singularities and Groups in Bifurcation Theory, Vol. II*. Springer-Verlag, New York, U.S.A.
- [51] Grossmann, S., 2000: The onset of shear flow turbulence. *Reviews Of Modern Physics*, 72 (2), 603.
- [52] Grün, G. and M. Rumpf, 2001: Simulation of singularities and instabilities arising in thin film flow. *Euro. Jnl of Applied Mathematics*, 12, 293–320.
- [53] Guckenheimer, J. and P. Holmes, 1990: *Nonlinear Oscillations, Dynamical Systems and Bifurcations of Vector Fields*, 2nd edition. Springer-Verlag, Berlin/Heidelberg.
- [54] Haines, P. E., R. E. Hewitt, and A. L. Hazel, 2011: The Jeffery-Hamel similarity solution and its relation to flow in a diverging channel. *Journal of Fluid Mechanics*, 687, 404–430.
- [55] Heil, M. and A. Hazel, 2006: *oomph-lib – An Object-Oriented Multi-Physics Finite-Element Library*. Fluid-Structure Interaction, Lecture Notes on Computational Science and Engineering, Editors: M. Schafer und H.-J. Bungartz.
- [56] Hof, B., et al., 2004: Experimental observation of nonlinear traveling waves in turbulent pipe flow. *Science*, 305 (5690), 1594–1598.
- [57] Huepe, C., L. S. Tuckerman, S. Métens, and M. E. Brachet, 2003: Stability and Decay Rates of Non-Isotropic Attractive Bose-Einstein Condensates. *Phys. Rev. A*, 68, 023 609.
- [58] Kalliadasis, S., C. Ruyer-Quil, B. Scheid, and M. G. Velarde, 2012: *Falling Liquid Films*, Applied Mathematical Sciences, Vol. 176. Springer Verlag.
- [59] Keller, H. B., 1977: Numerical solution of bifurcation and nonlinear eigenvalue problems. *Applications of Bifurcation Theory*, P. H. Rabinowitz, Ed., Academic Press, New York, U.S.A.
- [60] Kerswell, R., O. Tutty, and P. Drazin, 2004: Steady nonlinear waves in diverging channel flow. *Journal of Fluid Mechanics*, 501, 231–250.
- [61] Kevrekidis, I. G., B. Nicolaenko, and J. C. Scovel, 1990: Back in the saddle again - a computer-assisted study of the Kuramoto-Sivashinsky equation. *SIAM J. Appl. Math.*, 50, 760–790.
- [62] Koschmieder, E. L., 1993: *Bénard Cells and Taylor Vortices*. Cambridge University Press, Cambridge, UK.
- [63] Krauskopf, B. and H. Osinga, 2007: Computing invariant manifolds via the continuation of orbit segments. *Numerical Continuation Methods for Dynamical Systems: Path following and boundary value problems*, O. H. Krauskopf B. and G.-V. J., Eds., Springer-Verlag, *Understanding Complex Systems*, 117–154.
- [64] Krauskopf, B., H. M. Osinga, E. J. Doedel, M. E. Henderson, J. Guckenheimer, M. Dellnitz,

- and O. Junge, 2005: A survey of methods for computing (un)stable manifolds of vector fields. *Int. J. Bifurcation and Chaos*, 15, 763–791.
- [65] Kundu, P. J. and I. M. Cohen, 2002: *Fluid Mechanics*. Academic Press, New York.
- [66] Kuznetsov, Y. A. and V. V. Levitin, 1996: *CONTENT*, a multiplatform continuation environment. Tech. rep., CWI, Amsterdam, The Netherlands.
- [67] Landau, L. D. and I. M. Lifshitz, 1987: *Fluid Mechanics*. Butterworth - Heinemann, New York.
- [68] Lappa, M., 2005: On the nature and structure of possible three-dimensional steady flows in closed and open parallelepipedic and cubical containers under different heating conditions and driving forces. *Fluid Mechanics and Material Processing*, 1, 1–19.
- [69] Lappa, M., 2007: Secondary and oscillatory gravitational instabilities in canonical three-dimensional models of crystal growth from the melt. Part 1: Rayleigh-Bénard systems. *C.R. Mecanique*, 335, 253–260.
- [70] Lappa, M., 2010: *Thermal Convection: Patterns, Evolution and Stability*. John Wiley & Sons.
- [71] Levitus, S., 1994: *World Ocean Atlas 1994, Volume 4: Temperature*. NOAA/NESDIS E, US Department of Commerce, Washington DC, OC21, 1–117.
- [72] Lust, K., D. Roose, A. Spence, and A. Champneys, 1998: An adaptive Newton-Picard algorithm with subspace iteration for computing periodic solutions. *SIAM J. Sci. Comput.*, 19 (4), 1188–1209.
- [73] Mamun, C. K. and L. S. Tuckerman, 1995: Asymmetry and Hopf bifurcation in spherical Couette flow. *Phys. Fluids*, 7, 80–91.
- [74] Manneville, P., 2005: Rayleigh-Bénard convection, thirty years of experimental, theoretical, and modeling work. In: *Dynamics of Spatio-Temporal Cellular Structures. Henri Bénard Centenary Review* (eds. I. Mutabazi, J.E. Guyin, and J.E. Wesfreid), Springer.
- [75] Meerbergen, K. and D. Roose, 1996: Matrix transformations for computing rightmost eigenvalues of large sparse non-symmetric eigenvalue problems. *IMA J. Numer. Anal.*, 16, 297–346.
- [76] Meerbergen, K., A. Spence, and D. Roose, 1994: Shift-invert and Cayley transforms for the detection of rightmost eigenvalues of nonsymmetric matrices. *BIT*, 34, 409–423.
- [77] Menter, F. R., R. Langtry, and T. Hansen, 2004: *CFD simulation of turbomachinery flow-verification, validation and modelling*, European Congress on Computational Methods in Applied Sciences and Engineering, Jyväskylä, Vol. 1. 1-14 pp.
- [78] Merzhanov, A. G. and E. A. Shtessel, 1973: Free convection and thermal explosion in reactive systems. *Acta Astronautica*, 18, 191.
- [79] Molemaker, M. J. and H. A. Dijkstra, 2000: Multiple equilibria and stability of the North-Atlantic wind-driven ocean circulation. *Numerical methods for bifurcation problems and large-scale dynamical systems*, E. Doedel and L. S. Tuckerman, Eds., Springer, The IMA volumes in Mathematics and its applications, Vol. 119, 35–65.
- [80] Nagata, M., 1997: Three-dimensional traveling-wave solutions in plane Couette flow. *Physical Review E*, 55 (2), 2023–2025.
- [81] Net, M., A. Alonso, and J. Sánchez, 2003: From stationary to complex time-dependent flows in two-dimensional annular thermal convection at moderate Rayleigh numbers. *Phys. Fluids*, 15 (5), 1314–1326.
- [82] Net, M., F. Garcia, and J. Sánchez, 2008: On the onset of low-Prandtl-number convection in rotating spherical shells: non-slip boundary conditions. *J. Fluid Mech.*, 601, 317–337.
- [83] Olendraru, C. and A. Sellier, 2002: Absolute-convective instabilities of the Batchelor vortex

- in the viscous case. *J. Fluid Mech.*, 459, 371–396.
- [84] Oron, A., 2000a: Nonlinear dynamics of three-dimensional long-wave Marangoni instability in thin liquid films. *Phys. Fluids*, 12, 1633–1645.
- [85] Oron, A., 2000b: Three-dimensional nonlinear dynamics of thin liquid films. *Phys. Rev. Lett.*, 85, 2108–2111, doi:10.1103/PhysRevLett.85.2108.
- [86] Oron, A., S. H. Davis, and S. G. Bankoff, 1997: Long-scale evolution of thin liquid films. *Rev. Mod. Phys.*, 69, 931–980, doi:10.1103/RevModPhys.69.931.
- [87] Pawlowski, R. P., J. N. Shadid, J. P. Simonis, and H. F. Walker, 2006: Globalization techniques for Newton-Krylov methods and applications to the fully coupled solution of the Navier-Stokes equations. *SIAM Rev.*, 48, 700–721.
- [88] Puigjaner, D., J. Herrero, F. Giralt, and C. Simó, 2004: Stability analysis of the flow in a cubical cavity heated from below. *Phys. Fluids*, 16, 3639–3655.
- [89] Puigjaner, D., J. Herrero, C. Simó, and F. Giralt, 2011: From steady solutions to chaotic flows in a Rayleigh-Bénard problem at moderate Rayleigh numbers. *Physica D*, 240, 920–934.
- [90] Quon, C. and M. Ghil, 1992: Multiple equilibria in thermosolutal convection due to salt-flux boundary conditions. *J. Fluid Mech.*, 245, 449–484.
- [91] Quon, C. and M. Ghil, 1995: Multiple equilibria and stable oscillations in thermosolutal convection at small aspect ratio. *J. Fluid Mech.*, 291, 33–56.
- [92] Saad, Y., 1980: Variations on Arnoldi’s method for computing eigenelements of large unsymmetric matrices. *Lin. Alg. & Appl.*, 34, 269–295.
- [93] Saad, Y., 1992: Analysis of some Krylov subspace approximations to the matrix exponential operator. *SIAM J. Numer. Anal.*, 29 (1), 209–228.
- [94] Saad, Y., 1996: *Iterative Methods for Sparse Linear Systems*. PWS.
- [95] Sánchez, J., F. Marqués, and J. M. López, 2002: A continuation and bifurcation technique for Navier-Stokes flows. *J. Comput. Phys.*, 180, 78–98.
- [96] Sánchez, J. and M. Net, 2010: On the multiple shooting continuation of periodic orbits by Newton-Krylov methods. *Int. J. Bifurcation and Chaos*, 20 (1), 1–19.
- [97] Sánchez, J. and M. Net, 2013: A parallel algorithm for the computation of invariant tori in large-scale dissipative systems. *Physica D*, 252, 22–33.
- [98] Sánchez, J., M. Net, B. García-Archilla, and C. Simó, 2004: Newton-Krylov continuation of periodic orbits for Navier-Stokes flows. *J. Comput. Phys.*, 201 (1), 13–33.
- [99] Sánchez, J., M. Net, and C. Simó, 2010: Computation of invariant tori by Newton-Krylov methods in large-scale dissipative systems. *Physica D*, 239, 123–133.
- [100] Sánchez, J., M. Net, and J. Vega, 2006: Amplitude equations close to a triple-(+1) bifurcation point of D_4 -symmetric periodic orbits in $O(2)$ -equivariant systems. *Discrete and Continuous Dynamical Systems-Series B*, 6 (6), 1357–1380.
- [101] Scheid, B., A. Oron, P. Colinet, U. Thiele, and J. C. Legros, 2002: Nonlinear evolution of nonuniformly heated falling liquid films. *Phys. Fluids*, 14, 4130–4151.
- [102] Scheid, B., C. Ruyer-Quil, and P. Manneville, 2006: Wave patterns in film flows: modelling and three-dimensional waves. *J. Fluid Mech.*, 562, 183–222, doi:10.1017/S0022112006000978.
- [103] Scheid, B., C. Ruyer-Quil, U. Thiele, O. A. Kabov, J. C. Legros, and P. Colinet, 2005: Validity domain of the Benney equation including Marangoni effect for closed and open flows. *J. Fluid Mech.*, 527, 303–335, doi:10.1017/S0022112004003179.
- [104] Schneider, T. M. and B. Eckhardt, 2009: Edge states intermediate between laminar and turbulent dynamics in pipe flow. *Philosophical Transactions Of The Royal Society A-Mathematical Physical And Engineering Sciences*, 367 (1888), 577–587.

- [105] Schneider, T. M., B. Eckhardt, and J. Vollmer, 2007a: Statistical analysis of coherent structures in transitional pipe flow. *Physical Review E*, 75 (6), 066313.
- [106] Schneider, T. M., B. Eckhardt, and J. A. Yorke, 2007b: Turbulence transition and the edge of chaos in pipe flow. *Physical Review Letters*, 99 (3), 034502.
- [107] Schneider, T. M., J. F. Gibson, M. Lagha, F. De Lillo, and B. Eckhardt, 2008: Laminar-turbulent boundary in plane Couette flow. *Physical Review E*, 78 (3), 037301.
- [108] Schuttelaars, H. and H. De Swart, 1996: An idealized long-term morphodynamic model of a tidal embayment. *Eur. J. Mech., B/Fluids*, 15 (1), 55–80.
- [109] Seydel, R., 1994: *Practical Bifurcation and Stability Analysis: From Equilibrium to Chaos*. Springer-Verlag, New York, U.S.A.
- [110] Sharma, A. and R. Khanna, 1998: Pattern formation in unstable thin liquid films. *Phys. Rev. Lett.*, 81, 3463–3466, doi:10.1103/PhysRevLett.81.3463.
- [111] Shroff, G. M. and H. B. Keller, 1993: Stabilization of unstable procedures: the recursive projection method. *SIAM J. Numer. Anal.*, 30 (4), 1099–1120.
- [112] Simó, C., 1990: Analytical and numerical computation of invariant manifolds. *Modern methods in celestial mechanics*, C. Benest and C. Froeschlé, Eds., Editions Frontières, 285–330. Also at <http://www.maia.ub.es/dsg/2004/>.
- [113] Simonnet, E., M. Ghil, and H. A. Dijkstra, 2005: Homoclinic bifurcations of barotropic QG double-gyre circulation. *J. Mar. Res.*, 63, 931–956.
- [114] Simonnet, E., M. Ghil, K. Ide, R. Temam, and S. Wang, 2003: Low-frequency variability in shallow-water models of the wind-driven ocean circulation. Part I: Steady-state solutions. *J. Phys. Oceanogr.*, 33, 712–728.
- [115] Skufca, J., J. A. Yorke, and B. Eckhardt, 2006: The edge of chaos in a model for a parallel shear flow. *Phys. Rev. Lett.*, 96, 174101.
- [116] Sleijpen, G. and H. van der Vorst, 1996: A Jacobi-Davidson iteration method for linear eigenvalue problems. *SIAM J. Matrix Anal. Appl.*, 17, 401–425.
- [117] Sleijpen, G., H. van der Vorst, and F. Wubs, 1999: Preconditioning for eigenproblems. *Recent Advances in Numerical Methods and Applications II*, Proceedings of the fourth International Conference, NMA '98, Sofia Bulgaria, 19-23 August 1998, O. Iliev, M. Kasciev, S. Margenov, B. Sendov, and F. Vassilevski, Eds., World Scientific, 170–177.
- [118] Sobey, I. and P. Drazin, 1986: Bifurcations of two-dimensional channel flows. *Journal of Fluid Mechanics*, 171, 263–287.
- [119] Susan-Resiga, R., G. Ciocan, S. Muntean, I. Anton, and F. Avellan, 2006: Numerical simulation and analysis of swirling flow in the draft tube cone of a Francis turbine: 23rd IAHR Symposium-Yokohama., Vol. 15. 1-15 pp.
- [120] ter Brake, M., 2011: Tidal embayments: modelling and understanding their morphodynamics. Ph.D. thesis, Delft University of Technology.
- [121] ter Brake, M. and H. Schuttelaars, 2010: Modeling equilibrium bed profiles of short tidal embayments. on the effect of the vertical distribution of suspended sediment and the influence of the boundary conditions. *Ocean Dynamics*, 60, 183–204.
- [122] Thiele, U., 2010: Thin film evolution equations from (evaporating) dewetting liquid layers to epitaxial growth. *J. Phys.: Condens. Matter*, 22, 084019, doi:10.1088/0953-8984/22/8/084019.
- [123] Thiele, U. and E. Knobloch, 2004: Thin liquid films on a slightly inclined heated plate. *Physica D*, 190, 213–248.
- [124] Thiele, U. and E. Knobloch, 2006: Driven drops on heterogeneous substrates: Onset of sliding motion. *Phys. Rev. Lett.*, 97, 204501, doi:10.1103/PhysRevLett.97.204501.

- [125] Thiele, U., M. G. Velarde, K. Neuffer, M. Bestehorn, and Y. Pomeau, 2001: Sliding drops in the diffuse interface model coupled to hydrodynamics. *Phys. Rev. E*, 64, 061 601, doi: 10.1103/PhysRevE.64.061601.
- [126] Thies, J., F. Wubs, and H. A. Dijkstra, 2009: Bifurcation analysis of 3D ocean flows using a parallel fully-implicit ocean model. *Ocean Modelling*, 30 (4), 287–297.
- [127] Thual, O. and J. C. McWilliams, 1992: The catastrophe structure of thermohaline convection in a two-dimensional fluid model and a comparison with low-order box models. *Geophys. Astrophys. Fluid Dyn.*, 64, 67–95.
- [128] Tiesinga, G., F. Wubs, and A. Veldman, 2002: Bifurcation analysis of incompressible flow in a driven cavity by the Newton-Picard method. *J. Comput. Appl. Math.*, 140 (1–2), 751–772.
- [129] Toh, S. and T. Itano, 2003: A periodic-like solution in channel flow. *J. Fluid Mech.*, 481, 67–76.
- [130] Trefethen, L. and D. Bau, 1997: *Numerical Linear Algebra*. SIAM, Philadelphia.
- [131] Trenberth, K. E., J. G. Olson, and W. G. Large, 1989: A global ocean wind stress climatology based on ECMWF analyses. Tech. rep., National Center for Atmospheric Research, Boulder, CO, U.S.A.
- [132] Tseluiko, D., J. Baxter, and U. Thiele, 2013: A homotopy continuation approach for analysing finite-time singularities in thin liquid films. *IMA J. Appl. Math.*, doi: 10.1093/imamat/hxt021, (online).
- [133] Tuckerman, L. and D. Barkley, 1988: Global bifurcation to travelling waves in axisymmetric convection. *Phys. Rev. Lett.*, 61, 408–411.
- [134] Tuckerman, L. S. and D. Barkley, 2000: Bifurcation analysis for timesteppers. *Numerical Methods for Bifurcation Problems and Large-Scale Dynamical Systems*, E. Doedel and L. S. Tuckerman, Eds., Springer-Verlag, IMA Volumes in Mathematics and its Applications, Vol. 119, 453–466.
- [135] van Veen, L., G. Kawahara, and M. Atsushi, 2011: On Matrix-Free Computation of 2D Unstable Manifolds. *SIAM Journal on Scientific Computing*, 33 (1), 25–44.
- [136] Viswanath, D., 2007: Recurrent motions within plane Couette turbulence. *Journal Of Fluid Mechanics*, 580, 339–358.
- [137] Viswanath, D., 2009: The critical layer in pipe flow at high Reynolds number. *Philosophical Transactions Of The Royal Society A-Mathematical Physical And Engineering Sciences*, 367 (1888), 561–576.
- [138] Waleffe, F., 2003: Homotopy of exact coherent structures in plane shear flows. *Phys. Fluids*, 15, 1517–1534.
- [139] Walker, H. F., 1999: An adaptation of Krylov subspace methods to path following problems. *SIAM J. Sci. Comput.*, 21, 1191–1198.
- [140] Wedin, H. and R. R. Kerswell, 2004: Exact coherent structures in pipe flow: travelling wave solutions. *J. Fluid Mech.*, 508, 333–371.
- [141] Wheeler, P. and D. Barkley, 2006: Computation of spiral spectra. *SIAM J. Appl. Dyn. Syst.*, 5, 157–177.