Numerical continuation methods for large-scale dissipative dynamical systems

J. Sánchez Umbría[a] and M. Net[b]

Departament de Física, Universitat Politècnica de Catalunya

Abstract. A tutorial on continuation and bifurcation methods for the analysis of truncated dissipative partial differential equations is presented. It focuses on the computation of equilibria, periodic orbits, their loci of codimension-one bifurcations, and invariant tori. To make it more self-contained, it includes some definitions of basic concepts of dynamical systems, and some preliminaries on the general underlying techniques used to solve non-linear systems of equations by inexact Newton methods, and eigenvalue problems by means of subspace or Arnoldi iterations.

1 Introduction

Continuation methods have become the standard tools to analyze the dependence of the solutions of ordinary differential equations (ODEs) on parameters and to study their bifurcations. Equilibria, periodic and homoclinic orbits, loci of bifurcations, and boundary-value problems can be studied with the software packages developed during the past four decades. Among them AUTO [1] is probably the most widely known and used. The references and historical comments in [2–8] can be used to trace the history of these methods and the available software.

The applications of these techniques to partial differential equations (PDEs) is not so extended due essentially to the difficulties with the numerical linear algebra. There are two main fields in which they have been applied from the very beginning, structural and fluid mechanics. Two early examples are [9,10]. For these kind of problems, equilibria of the elasticity or Navier-Stokes equations were studied. This was possible, in most cases, thanks to the special structure of the matrices of the linear systems, obtained after discretizations by finite differences or elements, which allowed using adapted direct solvers. In the case of discretizations by spectral or pseudospectral methods, or in the case of periodic orbits, the Jacobians are no longer sparse and *matrix-free* methods are required. Their development for non-symmetric problems was impulsed by the appearance of GMRES [11], and by the development of Arnoldi methods to solve eigenvalue problems [12]. Several applications were presented in [13], and in the recent review article [14].

As far as we know there are only two freely available software packages for the computation of equilibria of PDEs which allow tracking steady solutions, and loci

---

[a] e-mail: `juan.j.sanchez@upc.edu`
[b] e-mail: `marta.net@upc.edu`

of, at least, some of the generic codimension-one bifurcations including symmetry breakings; LOCA [15], quite general and developed in C, and pde2path [16] for 2D problems and based on the FEM toolbox of Matlab, *pdetoolbox*.

The continuation of periodic orbits in large-scale dissipative systems has only been developed recently. Newton-Picard algorithms [17] were implemented in the package PDECONT, Broyden method were used in [18], and Newton-Krylov techniques in [19,20]. The computation of two-dimensional unstable manifolds of periodic orbits was developed in [21]. Methods for tracking invariant tori for large-scale applications were first considered in [22], and improved with a parallel algorithm in [23]. The continuation of codimension-one bifurcations of periodic orbits for high-dimensional systems has been developed only recently [24].

This tutorial article focuses on Newton-Krylov methods for the computation of steady states, periodic orbits, their bifurcation loci, and invariant tori, of systems of parabolic PDEs. It is therefore mainly devoted to computations based on time integrations. Some of these algorithms were already introduced with less details in reference [14]. The examples presented there were of equilibria or waves, which computation can be reduced to that of steady solutions. Since we only give references to examples, the description of the numerical methods and the underlying algorithms are given with more detail here. Moreover we include new types of continuations; invariant tori by a parallelizable method, much faster than that presented in [14], and codimension-one bifurcations of periodic orbits. Brief introductions to the numerical linear algebra tools needed to solve linear systems and eigenvalue problems, to continuation and inexact Newton's methods, and to some elementary concepts of stability in dynamical systems are provided for completeness. References to available software for these basic tools are also given. It is justified why the methods work efficiently, at a qualitatively level. With regard to efficiency, it is shown that once a particular problem has been discretized, giving rise to a system of ODEs of dimension n, the calculations of these invariant objects involve, at most, the time integration of systems of dimension $2n$ in the case of periodic orbits and tori, or $5n$ in order to track bifurcation loci of periodic orbits. This minimizes the computational work needed to calculate them. The last section contains references to recent publications using these techniques.

2 Equilibria and periodic orbits

Suppose
$$\dot{y} = f(y, p), \quad (y, p) \in \mathcal{U} \subset \mathbb{R}^n \times \mathbb{R} \tag{1}$$
is a large-scale ($n \gg 1$) autonomous system of ordinary differential equations (ODEs) obtained after the spatial discretization of a system of parabolic partial differential equations (PDEs) (reaction-diffusion or Navier-Stokes equations, for instance) and that
$$\varphi(t, x, p) \tag{2}$$
is its solution with initial condition x at $t = 0$ for a fixed value of p, i.e. $\varphi(0, x, p) = x$. In the case of discretizations by finite elements methods or for differential-algebraic equations (DAEs) it is possible that the system be of the form $M\dot{y} = f(y, p)$, with M a linear operator which might not be invertible. The form of Eq. (1) will be assumed for simplicity. We are interested in the computation of its equilibria, x, given by

$$f(x, p) = 0, \tag{3}$$

and its periodic regimes. A trajectory $\varphi(t, x, p)$ of (1), starting at x, is a periodic orbit if there is a minimal $T > 0$ (its period) such that $\varphi(T, x, p) = x$. Then a particular

snapshot of the periodic regime satisfies the system of equations

$$x - \varphi(T, x, p) = 0, \tag{4}$$
$$g(x, p) = 0, \tag{5}$$

where x is selected by a phase condition $g(x, p) = 0$, which can depend or not on p. The system determines simultaneously x and the period T for a fixed value of p. In many cases $g(x, p) = 0$ is the equation of a hyperplane cutting transversely the periodic orbits. It can be fixed or varied with the parameter p. The particular case $g(x, p) = \langle f(x_0, p_0), x - x_0 \rangle$, where x_0 is an initial approximation of x for $p = p_0$ and $\langle \cdot, \cdot \rangle$ is the Euclidean dot product, is known as the Poincaré condition.

The computation of both invariant objects (equilibria and periodic orbits) as functions of parameters, and the study of their stability imply solving large-scale nonlinear systems of equations and eigenvalue problems.

3 Continuation of zeros of a nonlinear system of equations

Consider a general system of nonlinear equations depending on a parameter p

$$H(x, p) = 0, \quad (x, p) \in \mathcal{U} \subset \mathbb{R}^m \times \mathbb{R} \tag{6}$$

with $m \gg 1$ for which one is interested in its solutions and their dependence on p. Parameter, pseudo-arclength, or local parameterization continuation methods are used to obtain the curves $(x(s), p(s))$ of solutions as functions, for instance, of the arclength s. They consist in two stages, a predictor step in which an initial approximation to a new point on the curve (x_0, p_0) is computed by extrapolation from the previous solutions, and a corrector step in which the prediction is refined, usually by Newton's method. The three continuation methods admit an unified formulation by adding to (6) an equation to determine simultaneously x and p

$$h(x, p) = 0. \tag{7}$$

If $h(x, p) = p - p_0$ the equation fixes the parameter p, leading to parameter continuation. It is not suitable for systems with strong variations with p or in the presence of folds of the curve of solutions. If $h(x, p) = h_x^\top(x - x_0) + h_p(p - p_0)$, (x_0, p_0) and (h_x, h_p) being the predictions of a new point and the tangent to the curve of solutions, respectively, the hyperplane $h(x, p) = 0$ will cut it transversely if the prediction is not far away from the previous point, and the algorithm allows passing turning points. The condition $h(x, p) = (x - x_0)^{(i)}$ or $h(x, p) = p - p_0$, where $(x - x_0)^{(i)}$ indicates here the i-th component of $x - x_0$, corresponds to the local parameterization method. It fixes one component of the vector (x, p), selected by looking at the variable which has the local largest variation. The index is that of the largest component of the tangent vector (h_x, h_p), or the largest component of the difference between the two last computed points on the curve of solutions. In any of the three cases the system that determines a unique solution, $(x, p) \in \mathbb{R}^{m+1}$, is

$$\widetilde{H}(x, p) = \begin{pmatrix} H(x, p) \\ h(x, p) \end{pmatrix} = 0.$$

Large-scale systems, $\widetilde{H}(x, p) = 0$, are usually solved by an inexact Newton's methods. Instead of solving the linear systems by means of direct algorithms, iterative methods are used. If the linear solvers are based on Krylov subspaces (GMRES,

BiCGStab, FOM, TFQMR, etc., see for instance [25]) they are also called Newton-Krylov methods. In the particular case of periodic orbits, iterations based on Newton-Picard [17] and Broyden [26] algorithms have also been proposed in the past. In inexact Newton's methods, a sequence of approximations are obtained, starting from the initial prediction (x_0, p_0), with

$$(x_{i+1}, p_{i+1}) = (x_i, p_i) + (\Delta x_i, \Delta p_i),$$

where $(\Delta x_i, \Delta p_i)$ satisfies the linear system

$$\begin{pmatrix} D_x H(x_i, p_i) & D_p H(x_i, p_i) \\ h_x^\top & h_p \end{pmatrix} \begin{pmatrix} \Delta x_i \\ \Delta p_i \end{pmatrix} = \begin{pmatrix} -H(x_i, p_i) \\ -h(x_i, p_i) \end{pmatrix} \qquad (8)$$

which is solved iteratively. The term *matrix-free methods* is used when the linear solver only requires the user to provide matrix products. In our case products of the form

$$\begin{pmatrix} D_x H(x_i, p_i) & D_p H(x_i, p_i) \\ h_x^\top & h_p \end{pmatrix} \begin{pmatrix} \delta x \\ \delta p \end{pmatrix},$$

with $(\delta x, \delta p)^\top$ an arbitrary vector. The linear solvers do not need to be convergent for an arbitrary system, and they hardly do when looking for equilibria of discretizations of systems of parabolic PDEs. It is necessary then to use preconditioners to improve the convergence.

Suppose that a generic system $Ax = b$ is to be solved, and that M is a matrix which approximates A in some sense ($M \approx A$) and which is easy to invert, i.e. it is easy to solve systems with matrix M. Left preconditioning consists in solving the system

$$M^{-1}Ax = M^{-1}b.$$

Its solution is that of $Ax = b$. In right preconditioning the system solved is

$$AM^{-1}y = b.$$

Then the solution of $Ax = b$ is $x = M^{-1}y$. This means that when applying a matrix-free method each matrix product by A is substituted by a matrix product by A followed by a system solve with matrix M in the case of left preconditioning, or by a system solve with matrix M followed by a matrix product by A for right preconditioning. If $M \approx A$ then $M^{-1}A$ or AM^{-1} will be closer to the identity than the original matrix A. Most iterative linear solvers based on Krylov subspaces converge quickly for matrices near the identity.

It is important to state a fundamental difference between the computation of equilibria and periodic orbits of discretized parabolic PDEs. Assume, to simplify, that the condition (7) is $p - p_0 = 0$ such that the last row of system (8) is $(0, \ldots, 0, 1)$, and that one has to solve a system with matrix $D_x H(x_i, p_0)$. In the case of equilibria $D_x H(x_i, p_0) = D_x f(x_i, p_0)$, and since $f(x, p)$ is the discretization of a system of elliptic equations, its spectrum will have eigenvalues of a wide range of negative real parts. This is the reason for the poor convergence of the iterative solvers. This will be seen in detail for the GMRES method later. In the case of periodic orbits the Jacobian with respect to x of Eq. (4) is $I - D_x \varphi(T, x, p)$. Since the flow of a parabolic differential equation (and of its consistent and stable discretizations) is strongly contractive, $D_x \varphi(T, x, p)$ has most of its spectrum tightly clustered around the origin, making $I - D_x \varphi(T, x, p)$ close to the identity. The addition of the phase condition, using a different continuation condition $h(x, p) = 0$, or the presence of some unstable multipliers of the periodic orbit (see below for definitions) out of the unit circle does not change this, except on a low-dimensional subspace. The case of

adding borders to the matrix (additional rows and columns as in system (8)) was considered in [27]. This means in practice that there is no need to use preconditioners in the case of periodic orbits. This is why this method has been used successfully by many authors, since its introduction in [19].

A consequence of the simplicity of computing periodic orbits, from the point of view of the linear solvers, is that it is also possible to find equilibria by solving $x - \varphi(T', x, p) = 0$, where now T' is an arbitrary characteristic time of the system, if a good time-stepper is available (it will be if periodic orbits are also being computed). It is clear that if $f(x, p) = 0$ then $x - \varphi(T', x, p) = 0$ for all $T' > 0$. The time T' must be as short as possible to save computing time, but large enough to have the required contractive properties of the flow $\varphi(t, x, p)$ just mentioned. This method can be used as a second option, for instance, when the continuation method applied to $f(x, p)$ fails due to the lack of good preconditioners for the linear systems or the eigenvalue problems.

When Newton-Krylov methods are applied to the computation of periodic orbits of (1) it is necesary to evaluate $x - \varphi(T, x, p)$, and the action of its Jacobian on vectors $(\delta T, \delta x, \delta p)$,

$$\delta x - D_x \varphi(T, x, p) \delta x - D_p \varphi(T, x, p) \delta p - D_t \varphi(T, x, p) \delta T, \tag{9}$$

where $D_t \varphi(T, x, p) = f(\varphi(T, x, p), p)$. The calculation of $\varphi(T, x, p)$ implies integrating (1) during a time T, with initial condition $y(0) = x$ and p fixed. The matrix product

$$D_x \varphi(T, x, p) \delta x + D_p \varphi(T, x, p) \delta p \tag{10}$$

can be computed by integrating a first variational equation. If

$$y(t) = \varphi(t, x, p), \tag{11}$$
$$y_1(t) = D_x \varphi(t, x, p) \delta x + D_p \varphi(t, x, p) \delta p, \tag{12}$$

are defined, y_1 satisfies, by interchanging the order of the derivatives and using that $\varphi(0, x, p) = x$,

$$\dot{y}_1 = D_y f(y, p) y_1 + D_p f(y, p) \delta p \quad \text{and} \quad y_1(0) = \delta x.$$

This equation must be solved coupled with that for y,

$$\dot{y} = f(y, p), \qquad\qquad y(0) = x, \tag{13}$$
$$\dot{y}_1 = D_y f(y, p) y_1 + D_p f(y, p) \delta p, \qquad\qquad y_1(0) = \delta x. \tag{14}$$

At time $t = T$

$$D_x \varphi(T, x, p) \delta x + D_p \varphi(T, x, p) \delta p = y_1(T). \tag{15}$$

Therefore the evaluation of (9) requires the integration of a system of $2n$ ODEs, in contrast to that of the full Jacobian $(D_x \varphi(T, x, p), D_p \varphi(T, x, p))$ which would require integrating, at least, $n(n + 2)$ equations.

The same holds if equilibria are found by solving, as suggested before, $x - \varphi(T', x, p) = 0$. The only difference is that T' is fixed and then the term $D_t \varphi(T, x, p) \delta T'$ is missing in (9).

4 The convergence of inexact Newton's methods

The efficient computation of the invariant objects considered in this paper requires that the number of iterations of Newton's method and the linear solvers be small. The

linear solvers can be preconditioned, or it might be unnecessary if the map for which the fixed points are searched involves the time evolution of parabolic equations, as seen before for periodic orbits. It remains to see what happens with Newton's method. Let us recall some results on its convergence [28].

Iterative methods can be classified by their rate of convergence which are defined as follows.

Definition 1 *Let $\{x_k\}_{k \in \mathbb{N}} \subset \mathbb{R}^n$ and $x^* \in \mathbb{R}^n$. Then*

- *$x_k \to x^*$ linearly with factor $\sigma \in (0,1)$ if $\|x_{k+1} - x^*\| \leq \sigma \|x_k - x^*\|$.*
- *$x_k \to x^*$ superlinearly if $\lim_{k \to \infty} \|x_{k+1} - x^*\| / \|x_k - x^*\| = 0$.*
- *$x_k \to x^*$ superlinearly with order $\alpha > 1$ if $x_k \to x^*$ and there is a $K > 0$ such that $\|x_{k+1} - x^*\| \leq K \|x_k - x^*\|^\alpha$.*
- *$x_k \to x^*$ quadratically if $x_k \to x^*$ and there is a $K > 0$ such that $\|x_{k+1} - x^*\| \leq K \|x_k - x^*\|^2$.*

The following results state the rate of convergence of Newton's and inexact Newton's methods. We refer to [28–30] for more details and proofs. Suppose one seeks to solve the system $F(x) = 0$ with $x, F(x) \in \mathbb{R}^n$, and assume the following conditions hold:

- $F(x) = 0$ has a solution x^*,
- F is continuously differentiable in a neighborhood of x^*,
- $DF(x^*)$ is nonsingular.

Theorem 1 *Under the above assumptions there is a $\delta > 0$ such that if $\|x_0 - x^*\| < \delta$ the Newton's iteration*

$$x_{k+1} = x_k + \Delta x_k, \quad with \quad DF(x_k)\Delta x_k = -F(x_k)$$

converges quadratically to x^*.

Newton's methods is therefore interesting because of its fast convergence, but it requires solving the linear systems $DF(x_k)\Delta x_k = -F(x_k)$ exactly. This is not possible, in general, for huge systems. Suppose that they are solved instead by an iterative method with stopping criteria

$$\|DF(x_k)\Delta x_k + F(x_k)\| \leq \eta_k \|F(x_k)\|, \tag{16}$$

for $\eta_k > 0$. The factors η_k are sometimes called the forcing terms in the literature after [31].

Theorem 2 *Let the above conditions hold. There exists a $\delta > 0$ such that if $\|x_0 - x^*\| < \delta$, and $\{\eta_k\} \subset [0, \eta]$ with $\eta < \bar{\eta} < 1$, then the inexact Newton iteration*

$$x_{k+1} = x_k + \Delta x_k, \quad with \quad \|DF(x_k)\Delta x_k + F(x_k)\| \leq \eta_k \|F(x_k)\|,$$

converges linearly to x^* with factor $\bar{\eta}$, with respect to the norm $\|\cdot\|_* = \|DF(x^*) \cdot \|$. Moreover

- *if $\eta_k \to 0$ the convergence is superlinear, and*
- *if $\eta_k \leq K_\eta \|F(x_k)\|^r$ for some $K_\eta > 0$ the convergence is superlinear with order $1+r$.*

Therefore if the linear systems are solved with the adequate tolerances for the residuals ($\eta_k \leq K_\eta \|F(x_k)\|$), quadratic convergence can still be obtained. In practice and for a particular problem, one can do some experiments with different stopping

criteria in order to minimize the total computational cost of the inexact Newton's method.

The expression inexact Newton's method was coined in [29] to refer to the case where the linear solver is stopped when the condition (16) holds. There are, of course, other sources of inexactness due to the errors in the evaluation of the function and the action by its Jacobian. They can degrade or event prevent the convergence of Newton's method (see [28]). This is specially important for invariant objects found as fixed points of maps that involve time integrations (periodic orbits and their bifurcations loci or invariant tori, for example). In these cases the approximation by finite differences of the action by the Jacobian is very dangerous because the error in the time evolution is amplified when it is divided by the step of the difference formula. Therefore, the integration of a variational equation as described at the end of section 3 is almost compulsory, unless the time stepper be very accurate, which is hardly the case when solving PDEs.

With regard to available software we will only mention NITSOL [32]. It includes standard and globalized Newton's methods, and three linear solvers (GMRES, BiCGSTAB and TFQMR). It is therefore a standalone implementation of inexact Newton's methods.

5 GMRES and its convergence

In order to understand the reasons for the fast or slow convergence of the linear solvers it is good to have some knowledge of at least a particular method. We will focus on GMRES since it is the most widely used linear solver for non-symmetric problems, and because the theory is quite simple in this case. Assume that a non-singular system $Ax = b$ of dimension $n \gg 1$ is to be solved. Projection methods produce, starting with an initial guess x_0, a sequence of approximations, x_k, to the solution $x^* = A^{-1}b$, in the affine subspace $x_k \in x_0 + \mathcal{K}_k$, which satisfy the Petrov-Galerkin condition $b - Ax_k \perp \mathcal{L}_k$, where $\mathcal{K}_k$ and $\mathcal{L}_k$ are two k-dimensional linear subspaces. If $\mathcal{L}_k = A\mathcal{K}_k$, then x_k satisfies

$$||b - Ax_k||_2 = \inf_{x \in x_0 + \mathcal{K}_k} ||b - Ax||_2,$$

i.e., x_k minimizes the Euclidean norm $||b - Ax||_2$ over $x \in x_0 + \mathcal{K}_k$.

In the particular case of GMRES, $\mathcal{L}_k = A\mathcal{K}_k$ and $\mathcal{K}_k$ is the Krylov subspace

$$\mathcal{K}_k = \{r_0, Ar_0, A^2 r_0, \dots, A^{k-1} r_0\},$$

$r_0 = b - Ax_0$ being the initial residual. It follows that if $x_k = x_o + z_k$ with $z_k \in \mathcal{K}_k$ then

$$
\begin{aligned}
r_k \equiv & b - Ax_k = b - A(x_0 + z_k) = r_0 - Az_k = \\
= & Ir_0 - A(\alpha_1 r_0 + \alpha_2 Ar_0 + \dots + \alpha_k A^{k-1} r_0) \\
= & (I - \alpha_1 A - \alpha_2 A^2 - \dots - \alpha_k A^k) r_0 = p_k(A) r_0,
\end{aligned}
$$

p_k being a polynomial of degree k, with $p_k(0) = 1$. This means that the residual r_k can be written as the action of a polynomial of the matrix A acting on the initial residual r_0. Now, by using that:

- if $A = V \Lambda V^{-1}$ then $A^l = V \Lambda^l V^{-1}$ and $p(A) = V p(\Lambda) V^{-1}$,
- if $\Lambda = diag(\lambda_1, \dots, \lambda_n)$ then $p(\Lambda) = diag(p(\lambda_1), \dots, p(\lambda_n))$,

- $\|p(A)\|_2 \leq \|p(\Lambda)\|_2 \|V\|_2 \|V^{-1}\|_2 = \kappa_2(V)\|p(\Lambda)\|_2$, with $\kappa_2(V) = \|V\|_2\|V^{-1}\|_2$ the norm-2 condition number of V, and
- if $\Lambda = diag(\lambda_1, \ldots, \lambda_n)$ then $\|p(\Lambda)\|_2 = \max_{i=1,\ldots,n} |p(\lambda_i)|$,

where $diag(a_1, \ldots, a_n)$ means the diagonal with entries the a_i, the following result is obtained.

Theorem 3 *(Saad and Schultz 1986) Assume that A is diagonalizable with $A = V\Lambda V^{-1}$, where $\Lambda = diag(\lambda_1, \cdots, \lambda_n)$ is the diagonal matrix of eigenvalues, P_k is the set of polynomials of degree at most k such that $p(0) = 1$, and $\kappa_2(V) = \|V^{-1}\|_2\|V\|_2$ is the norm-2 condition number of V. Then at the k-th step of GMRES*

$$\frac{\|b - Ax_k\|_2}{\|b - Ax_0\|_2} \leq \kappa_2(V) \inf_{p \in P_k} \max_{i=1,\ldots,n} |p(\lambda_i)|.$$

With the above preparation the proof is trivial. If $x_k = p_k(A)r_0$ minimizes $\|b - Ax\|_2$ then

$$\|b - Ax_k\|_2 = \inf_{p \in P_k} \|p(A)r_0\| \leq \kappa_2(V)\|b - Ax_0\|_2 \inf_{p \in P_k} \max_{i=1,\ldots,n} |p(\lambda_i)|.$$

It is very difficult to obtain estimates of $\kappa_2(V)$ for large-scale systems. It is known that $\kappa_2(V) \geq 1$, but one can be almost sure that it will be large for realistic problems. In any case, to conclude from this result that after a small number of iterations k, $\|b - Ax_k\|_2$ will be small it is necessary that there exists a polynomial of small degree k which is very small on the spectrum of A. This is accomplished if the spectrum consists of a small number of clusters. Estimates in this situation can be found in [33]. In the case of periodic orbits the matrix of the linear system is essentially $I - D_x\varphi(T, x, p)$ with most of the spectrum in a small cluster surrounding $+1$. This is almost the ideal situation and GMRES converges quickly. In the case of equilibria the matrix is $D_x f(x, p)$, and since this Jacobian will include diffusion terms its eigenvalues will spread to $-\infty$ in the complex plane. Therefore preconditioners M are required to reduce the diameter of the spectrum of $M^{-1}D_x f(x, p)$ (or $D_x f(x, p)M^{-1}$), in order to accelerate the convergence.

A problem is still pending; how to find the x_k solving the minimizing problem $\inf_{x \in x_0 + \mathcal{K}_k} \|b - Ax\|_2$. This is done by reducing it, using the Arnoldi's decomposition (see below), to one for a matrix of dimension k which can be solved very efficiently. We skip the details and recommend the references [25, 28, 11] for details and further reading. There are very good sources of software for Krylov linear solvers. See for instance [34, 35].

6 Stability of equilibria and periodic orbits

In this section the definitions of stability of equilibria and periodic orbits, and their characterizations in terms of eigenvalues are recalled. The dependence on the parameter p of the system of ODEs (1) is not made explicit. It is written as $\dot{y} = f(y)$, and $\varphi(t, x)$ for its solution with initial condition x. Let x^* be a fixed point (equilibrium) of the system, i.e. $f(x^*) = 0$.

Definition 2 *The fixed point is said to be Lyapunov stable if for every neighborhood N of x^* there is a neighborhood $M \subset N$ of x^* such that if $x \in M$, then $\varphi(t, x) \in N$ for all $t \geq 0$.*

An equilibrium that is not stable is said to be unstable.

Definition 3 *The fixed point is said to be asymptotically stable if it is Lyapunov stable and there is a neighborhood N of x^* such if $x \in N$ then $\lim_{t \to \infty} \|\varphi(t, x) - x^*\| = 0$.*

Definition 4 *The fixed point said to be exponentially stable if it is asymptotically stable and there exist $\alpha > 0$, and $\beta > 0$, and a neighborhood N of x^* such that if $x \in N$, then $\|\varphi(t, x) - x^*\| \le \alpha \|x - x^*\| e^{-\beta t}$, for $t \ge 0$.*

Theorem 4 *If f is of class C^1 and x^* is a fixed point such that all the eigenvalues of $Df(x^*)$ have strictly negative real parts, then x^* is exponentially stable (and hence asymptotically stable). If at least one eigenvalue has strictly positive real part, then x^* is unstable.*

This means that the eigenvalues of $Df(x^*)$ of largest real part must be computed to know the stability and detect bifurcations of fixed points.

Definition 5 *A set Λ is said to be invariant under the flow $\varphi(t, x)$ if $\varphi(t, \Lambda) = \Lambda$ for all t; that is, for each $x \in \Lambda$, $\varphi(t, x) \in \Lambda$ for any t.*

Definition 6 *The invariant set is said to be stable if for every neighborhood N of Λ there is a subset $M \subset N$ such that if $x \in M$, then $\varphi(t, x) \in N$ for all $t \ge 0$.*

An set that is not stable is called unstable.

Definition 7 *The invariant set is said to be asymptotically stable if it is stable and there is a neighborhood N of Λ such if $x \in N$ then, then $\lim_{t \to \infty} \rho(\varphi(t, x), \Lambda) = 0$, with $\rho(x, \Lambda) = \inf_{y \in \Lambda}(\|x - y\|)$.*

Consider the first variational equation $\dot{M} = Df(x(t))M$ about a periodic orbit $x(t) = \varphi(t, x)$ of period T ($\varphi(T, x) = x$), with initial condition $M(0) = I$. The solution at time T is called the monodromy matrix $M(T)$, and its eigenvalues the Floquet multipliers of the periodic orbit.

Theorem 5 *The monodromy matrix $M(T)$ always has a unit eigenvalue with eigenvector $\dot{x}(0) = \dot{x}(T) = f(x(0))$.*

This unit eigenvalue is named the trivial eigenvalue of the periodic orbit.

Theorem 6 *If $x(t)$ is a periodic orbit of a C^2 flow $\varphi(t, x)$ that is linearly asymptotically stable (its monodromy matrix has all the eigenvalue inside the unit circle except the trivial one), then it is asymptotically stable.*

Therefore the eigenvalues of $M(T)$ of largest modulus have to be computed to know the stability and detect bifurcations of periodic orbits.

7 Large-scale eigenvalue problems

Most algorithms to obtain the leading (largest magnitude) eigenvalues of a large-scale matrix are based on the Rayleigh-Ritz method [36]. Assume A is a real matrix of dimension $n \times n$ ($n \gg 1$), and that $V_m = [v_1, \ldots, v_m]$ is of dimension $n \times m$ ($m \ll n$) such that $V_m^\top V_m = I_m$ (I_m identity of dimension m) and

$$AV_m = V_m B_m,$$

with B_m of dimension $m \times m$. This expression says that the columns of V_m generate an invariant subspace of A of dimension m. Moreover, if (λ, u) is an eigenpair of B_m ($B_m u = \lambda u$) then $(\lambda, V_m u)$ is an eigenpair of A:

$$AV_m u = V_m B_m u = \lambda V_m u.$$

If the identity $AV_m = V_m B_m$ is only approximate, the pairs $(\lambda, V_m u)$ are called Ritz values and vectors of A, respectively. It is clear that since $V_m^\top V_m = I_m$ then $B_m = V_m^\top A V_m$, so obtaining B_m reduces to the computation of matrix vectors by A followed by dot products.

There are two main methods to obtain approximate bases, V_m, of the subspace corresponding to the leading eigenvalues of a large-scale matrix; subspace iteration and Arnoldi's method. They are described in next section.

It is important to realize that in the case of periodic orbits all that is needed is having one of these algorithms available. They provide precisely the eigenvalues of interest. In the case of equilibria, transformations of the matrix are required. This is explained below. Studying the stability of periodic orbits is easier than for equilibria, as in the case of solving the linear systems. It is also possible to study the stability of fixed points by using time evolution. If (λ, v) is an eigenpair of $Df(x^*)$ then (μ, v), with $\mu = \exp(\lambda T')$, is an eigenpair of $D_x \varphi(T', x)$. By taking logarithms, and using that $\lambda = \langle v, Av \rangle / \langle v, v \rangle$ to recover the undetermined imaginary part, λ can be obtained from μ. Since the exponential transforms the left half complex plane into the interior of the unit circle, the *multipliers* μ of larger modulus are those of interest.

7.1 Subspace iteration

Subspace iteration, for which several good implementation exist [37,38], works as follows.

1. Start: Choose an initial system of orthonormal vectors $V_m = [v_1, \ldots, v_m]$, $(m \ll n)$
2. Iterate: Until convergence do:
   (a) For l=1,..., k do
       i. Compute $Z_m = AV_m$
       ii. Orthonormalize Z_m by computing $Z_m = Q_m R_m$ with $Q_m^\top Q_m = I_m$ and R_m upper-triangular, and set $V_m = Q_m$
   (b) Form $B_m = V_m^\top A V_m$ and compute the eigenpairs (λ_i, z_i), $i = 1, \cdots, m$ of B_m by the QR method (with LAPACK [39] for instance)
   (c) Test for convergence of eigenvalues and/or eigenvectors
3. Stop: When satisfied, compute the approximate eigenvectors of A as $x_i = V_m z_i$, $i = 1, \cdots, m$. The λ_i, $i = 1, \cdots, m$ are the approximate eigenvalues.

The following results gives a complete description of the convergence of the method.

Theorem 7 *Suppose that the n eigenvalues of A are ordered by decreasing modulus* $|\lambda_1| \geq |\lambda_2| \geq \cdots |\lambda_m| > |\lambda_{m+1}| \geq \cdots \geq |\lambda_n|$. *If the initial set of vectors V_m is not deficient in the eigenvectors corresponding to $\lambda_1, \cdots, \lambda_m$, and if the total number of iterations taken, k, is large enough, then the previous algorithm computes approximations $\hat{\lambda}_{i,k}$ to λ_i $(i = 1, \ldots, m)$ with*

$$|\hat{\lambda}_{i,k} - \lambda_i| = O\left(\left|\frac{\lambda_{m+1}}{\lambda_i}\right| + \varepsilon_{i,k}\right)^k, \quad \lim_{k \to \infty} \varepsilon_{i,k} = 0.$$

Moreover, if λ_i is simple, then $\varepsilon_{i,k} = 0$.

It is clear that λ_1 will, generically, converge first, then λ_2, etc. Therefore in good implementations, as those already mentioned, care is taken to freeze the part of the approximate invariant base $\{v_1, \ldots, v_m\}$ already converged in order to save computing time. Moreover, the columns of the product AV_m can be computed in parallel since they are independent.

7.2 Arnoldi iterations

The Arnoldi's method is based on the Arnoldi's decomposition. The latter is also used in the iterative solution of linear systems (for instance in GMRES) and to find reduced order models.

Let A be a given matrix, and v_1 an unitary vector. The Arnoldi's decomposition finds an orthonormal basis of the Krylov subspace $\mathcal{K}_m = \{v_1, Av_1, A^2v_1, \ldots, A^{m-1}v_1\}$ by applying a modified Gram-Schmidt orthogonalization:
Iterate: for $j = 1, 2, \ldots, m$ compute

1. $w = Av_j$
2. for $i = 1, 2, \ldots, j$ do
   (a) $h_{i,j} = \langle w, v_i \rangle$
   (b) $w \leftarrow w - h_{i,j} v_i$
3. $h_{j+1,j} = \|w\|$, if $h_{j+1,j} = 0$ stop
4. $v_{j+1} = w / h_{j+1,j}$

If $V_m = [v_1, \ldots, v_m]$, its columns form an orthonormal basis of $\mathcal{K}_m$, and if H_m is the $m \times m$ upper Hessenberg matrix whose nonzero entries are the $h_{i,j}$ then

$$AV_m = V_m H_m + w_m e_m^\top, \quad \text{and} \quad V_m^\top AV_m = H_m,$$

with $w_m = h_{m+1,m} v_{m+1}$, and $e_m^\top = (0, \ldots, 0, 1) \in \mathbb{R}^m$. At line 3 of the above pseudocode, the stop indicates that an exact invariant subspace has been found. If more eigenvalues are needed the method can be restarted with a new v_1 linearly independent of the basis already found.

The Arnoldi's iteration to obtain Ritz eigenpairs, implemented for instance in the package ARPACK [12], is:

1. Start: Choose an initial unitary vector v_1.
2. Iterate: Until convergence do:
   (a) Compute the Arnoldi factorization $AV_m = V_m H_m + w_m e_m^\top$.
   (b) Compute the eigenpairs (λ_i, z_i), $i = 1, \cdots, m$ of $H_m = V_m^\top AV_m$ by the QR method (LAPACK).
   (c) Test for convergence of eigenvalues and/or eigenvectors. If not converged select a new initial vector v_1 from the Arnoldi factorization.
3. Stop: When satisfied, compute approximate eigenvectors of A as $x_i = V_m z_i$, $i = 1, \cdots, m$. The λ_i, $i = 1, \cdots, m$ are the approximate eigenvalues.

The size of the Krylov subspace is limited to m to save storage and computational cost. When this size is reached the method is restarted by choosing a new initial vector v_1 in step (c). This must be done trying to preserve as much information as possible of the eigenvalues of interest. The way this is done is out of the scope of this article. The reference [40] contains the algorithm implemented in ARPACK. The algorithm of the Arnoldi decomposition is sequential and therefore parallelism has to be implemented at the level of the computation of each matrix product.

The results on the convergence of this method are not as clean and elegant than for subspace iteration (see [36]). The practice shows that it is, in general, faster and therefore Arnoldi's method is now preferred.

During the computation of solutions by Newton-Krylov methods, Arnoldi's decompositions are generated (for instance inside GMRES). They can be used to obtain estimates of the eigenvectors, probably by increasing the size of the Krylov base. This possibility was studied in [41].

8 Eigenvalue Transformations

Two ways of obtaining eigenvalues of maximal modulus have been described. This is exactly what it is required in the analysis of the stability of periodic orbits. To find those of maximal real part, in the case of equilibria, matrix transformations have to be used to convert the eigenvalues of maximal real part to those of maximal modulus of the transformed matrix [36, 42]. This process is also known as eigenvalue preconditioning. Some possibilities are the following.

Shift-invert with real or complex shift σ:

$$Av = \lambda v \implies (A - \sigma I)^{-1} v = \mu v \quad \text{with} \quad \mu = 1/(\lambda - \sigma).$$

The eigenvalues of A close to the shift σ become those of largest modulus of $(A - \sigma I)^{-1}$. It implies solving systems with the matrix $A - \sigma I$. If the original problem is real, one can take a real $\sigma > 0$, compute a large number of eigenvalues μ of $(A - \sigma I)^{-1}$, transform them by $\lambda = \sigma + 1/\mu$ and reorder by the real part [43]. This can be useful if the imaginary parts are not too large. If a wide range of imaginary parts are present and it is possible to use complex arithmetics a swept of complex shifts σ can be used [44, 45].

Generalized Cayley transformation with shift σ and anti-shift τ:

$$Av = \lambda v \implies (A - \sigma I)^{-1}(A - \tau I) v = \mu v \quad \text{with} \quad \mu = (\lambda - \tau)/(\lambda - \sigma).$$

The line $\text{Re}(\lambda) = (\sigma + \tau)/2$ is mapped to the unit circle and $\text{Re}(\lambda) < (\sigma + \tau)/2$ $(\text{Re}(\lambda) > (\sigma + \tau)/2)$ is mapped to the interior (exterior) of the unit circle. This makes this transformation interesting [42, 46, 47]. The problem is that the choice of σ and τ can be very tricky because all the eigenvalues of A of large negative real part, which are of no interest, are transformed into μ near $+1$. This can make the convergence very slow. A possible strategy is starting with a filtering of the initial vector by applying several actions of $(A - \sigma I)^{-1}$ to remove the components along the unwanted eigenvectors. A similar technique was used recently in [47].

Double complex shift: If $\sigma = \rho + i\theta$,

$$Av = \lambda v \implies (A - \sigma I)^{-1}(A - \bar{\sigma} I)^{-1} v = \mu v \quad \text{with} \quad \mu = 1/((\lambda - \rho)^2 + \theta^2).$$

Systems with the matrix $(A - \sigma I)(A - \bar{\sigma} I) = (A - \rho I)^2 + \theta^2 I$ have to be solved. The advantage is that it is real if A does. This method was used, together with exponential transformations in [48] because there the matrix A was block-tridiagonal, then $(A - \rho I)^2 + \theta^2 I$ was block-pentadiagonal and a tailored LU decomposition could be used to solve the systems.

Exponential:

$$Av = \lambda v \implies \exp(TA) v = \mu v \quad \text{with} \quad \mu = \exp(\lambda T),$$

with an arbitrary time T. This is the ideal transformation. The lines $\text{Re}(\lambda) = a$ are mapped to the circles $|\mu| = \exp(a)$. The computation of the exponential requires the time integration of $\dot{y} = Ay$, but it is supposed that a time integrator for the original system is already available, which could be modified for this purpose. This method is expensive but very robust. It has ben used in [48, 49].

9 Continuation of bifurcation curves

Suppose now that the parameter p in (1) is two-dimensional ($p = (p_1, p_2)$). It is now described how to track curves of codimension-one bifurcations of equilibria and periodic orbits in system with or without symmetries. These curves separate the regions in the parameter space of different phase portraits. The extended systems in the case of equilibria are the same as for low-dimensional systems [6,50,7]. The only difference being that the linear systems have to be solved iteratively. The package LOCA incorporates the tracking of codimension-one bifurcations of steady solutions [15,51]. We will focus on the case of periodic orbits of large-scale systems because its developments is very recent [24]. The methods described here can also be applied to equilibria if they are found as fixed points of the flow as described above.

9.1 Saddle-node and period-doubling bifurcations

The saddle-node and period-doubling bifurcations can be treated together. A point x on a periodic orbit of period T with a non-trivial multiplier $\lambda = 1$ (saddle-node), or $\lambda = -1$ (period-doubling) satisfies the system of equations $H(x, u, T, p) = 0$ given by

$$x - \varphi(T, x, p) = 0,$$
$$g(x, p) = 0,$$
$$\lambda u - \left(D_x \varphi(T, x, p)u - \frac{1}{2}(1 + \lambda)\frac{\langle w, u \rangle}{\langle w, w \rangle}w \right) = 0,$$
$$\langle u_r, u \rangle = 1,$$

with $w = f(x, p)$. The first two equations are the same as (4)-(5) and determine a single point on a periodic orbit and its period. The third equation says that there is an eigenvector u of the monodromy matrix of eigenvalue λ, normalized by the fourth equation, u_r being a reference vector. It can be $u_r = u$. The last term of the third equation is missing in the case of a period-doubling. For a saddle-node point it becomes $(\langle w, u \rangle / \langle w, w \rangle)w$. It is necessary to remove the trivial $+1$ multiplier in order to have a regular system of equations. This is done by deflating the trivial eigenvector which is $f(x, p)$. This prodedure is known as Wieland's deflation [36]. It shifts the trivial $+1$ multiplier to zero. It is clear that the vector of unknowns $X = (x, u, T, p)$ has dimension $2n + 3$ (remember that p is two-dimensional), and that the $2n + 2$ equations define the curve of bifurcations.

The terms $\varphi(T, x, p)$ and $D_x \varphi(T, x, p)u$ can be computed as in the case of evaluating the product (10), but with $\delta p = 0$, by using the auxiliary functions $y(t)$ and $y_1(t)$ (see Eqs. (11),(12)). The evaluation of the deflation term is trivial. The action of $D_X H(x, u, T, p)$ on $(\delta x, \delta u, \delta T, \delta p)$ needed in the matrix-free methods to solve linear systems is

$$\delta x - D_t \varphi(T, x, p)\delta T - D_x \varphi(T, x, p)\delta x - D_p \varphi(T, x, p)\delta p, \tag{17}$$

$$D_x g(x, p)\delta x + D_p g(x, p)\delta p, \tag{18}$$

$$\lambda \delta u - D_{tx}^2 \varphi(T, x, p)(u, \delta T) - D_{xx}^2 \varphi(T, x, p)(u, \delta x) - D_{xp}^2 \varphi(T, x, p)(u, \delta p)$$
$$- D_x \varphi(T, x, p)\delta u$$
$$+ \frac{1 + \lambda}{2\langle w, w \rangle} \left(\langle w, u \rangle z + \left(\langle z, u \rangle + \langle w, \delta u \rangle - \frac{2\langle w, z \rangle}{\langle w, w \rangle}\langle w, u \rangle \right) w \right), \tag{19}$$

$$\langle u_r, \delta u \rangle, \tag{20}$$

where $w = f(x, p)$ and $z = D_y f(x, p)\delta x + D_p f(x, p)\delta p$. Let us define

$$
\begin{aligned}
y(t) &= \varphi(t, x, p), \\
y_1(t) &= D_x \varphi(t, x, p) u, \\
y_2(t) &= D_x \varphi(t, x, p)\delta x + D_p \varphi(t, x, p)\delta p, \\
y_3(t) &= D_{xx}^2 \varphi(t, x, p)(u, \delta x) + D_{xp}^2 \varphi(t, x, p)(u, \delta p), \\
y_4(t) &= D_x \varphi(t, x, p)\delta u.
\end{aligned}
$$

It is easy to check that all the terms in the expressions (17)-(20) are contained in the vectors $y_i(T)$ $i = 1, \ldots, 4$, or can be reduced to evaluations of the vector field $f(x, p)$ and its first order derivatives. For instance,

$$
D_t \varphi(T, x, p)\delta T = f(y(T), p)\delta T,
$$
$$
D_{tx}^2 \varphi(T, x, p)(u, \delta T) = \delta T \, D_y f(\varphi(T, x, p), p) D_x \varphi(T, x, p) u = \delta T \, D_y f(y(T), p) y_1(T).
$$

The system which must be integrated to obtain $y(T)$, $y_i(T)$, $i = 1, \ldots, 4$ is

$$
\begin{array}{lll}
\dot{y} = f(y, p), & y(0) = x & (21) \\
\dot{y}_1 = D_y f(y, p) y_1, & y_1(0) = u & (22) \\
\dot{y}_2 = D_y f(y, p) y_2 + D_p f(y, p)\delta p, & y_2(0) = \delta x & (23) \\
\dot{y}_3 = D_y f(y, p) y_3 + D_{yy}^2 f(y, p)(y_1, y_2) + D_{yp}^2 f(y, p)(y_1, \delta p), & y_3(0) = 0 & (24) \\
\dot{y}_4 = D_y f(y, p) y_4, & y_4(0) = \delta u. & (25)
\end{array}
$$

It can be decoupled into two subsystems which can be solved in parallel; (21) and (25) of dimension $2n$, and (21)-(24) of dimension $4n$.

9.2 Neimark-Sacker bifurcations

The Hopf bifurcations of periodic orbits with multiplier $e^{i\theta}$ and associated eigenvector $u + iv$ are solutions of the system $H(x, u, v, T, \theta, p) = 0$ given by

$$
\begin{aligned}
x - \varphi(T, x, p) &= 0, \\
g(x, p) &= 0, \\
u\cos\theta - v\sin\theta - D_x\varphi(T, x, p)u &= 0, \\
u\sin\theta + v\cos\theta - D_x\varphi(T, x, p)v &= 0, \\
\langle u, u\rangle + \langle v, v\rangle &= 1, \\
\langle u, v\rangle &= 0.
\end{aligned}
$$

The first two equations determine a single point on a periodic orbit and its period, the next two are the real and imaginary parts of $e^{i\theta}(u+iv) - D_x\varphi(T, x, p)(u+iv) = 0$, and the last two uniquely determine the eigenvector $u + iv$. Now $X = (x, u, v, T, \theta, p)$ has dimension $3n + 4$, and the $3n + 3$ equations define the curve of Neimark-Sacker points. The work required to evaluate $H(X)$ and $D_X H(X)\delta X$ doubles since there are now two products by $D_x\varphi(T, x, p)$ in the system. They can obviously be computed in parallel. All the terms required to compute H, or the action by its Jacobian have already been considered in the case of saddle-node bifurcations.

The case of pitchfork bifurcations, when simmetries are present, was also considered in [24], where an example of thermal convection of a binary fluid was presented,

having all the types of codimension-one bifurcation, and where the reasons for the convergence of these methods were explained.

The bifurcations of steady solutions can also be continued with the preceding systems by taking as T a fixed characteristic time of the problem, which is no longer an unknown, and removing the phase condition. This is not, in general, the most efficient way because it involves time integrations, but as stated before there is no need to find adequate preconditioners for the extended systems.

10 Computation of invariant tori

After a Neimark-Sacker bifurcation of a periodic orbit, and under very general conditions, it changes its stability and a new branch of two-dimensional invariant tori appears. Their study is simplified if a Poincaré section is taken, cutting the periodic orbit at a point and the torus at an invariant curve. The methods used to compute invariant tori for low-dimensionl systems consist in discretizing the invariant curve or the full tori expanding them in Fourier series (simple or double), using finite differences to approximate the invariance conditions, or finding a single point of the tori by adding a second section [52–57]. In the case of high-dimensional systems, tori are usually computed by time evolution of the equations. Only attracting tori are obtained by this method. It is a very expensive way of studying their dependence on parameters, specially near bifurcations, where due to the weak attraction the transients are very long. We summarize here the methods described in [22,23], which use all the previous tools in this article, are much more efficient, and allow computing weakly unstable tori. Since discretizing a full invariant curve in a large phase space would be too expensive, the two methods presented try to find a single point or an small arc of it.

10.1 Poincaré maps and its derivatives

Consider again system (1) with a one-dimensional parameter p, and suppose there is a periodic orbit which cuts transversely a hyperplane Σ_1 given by

$$\Sigma_1 = \left\{ y \in \mathbb{R}^n \,/\, \langle v_1, y - x_1^\Sigma \rangle = 0 \right\},$$

x_1^Σ and v_1 being a point on the hyperplane and its normal unitary vector, respectively. Let $\mathcal{V} \subset \Sigma_1$ be a neighborhood of the intersection of the periodic orbit with Σ_1 at which the Poincaré map $P : \mathcal{V} \subset \Sigma_1 \to \Sigma_1$ is defined as

$$P(x,p) = \varphi(t(x,p), x, p), \tag{26}$$

$t(x,p) > 0$ being the first time for which $\varphi(t(x,p), x, p) \in \Sigma_1$ with $\mathrm{sign}\langle v_1, f(x,p)\rangle = \mathrm{sign}\langle v_1, f(P(x,p), p)\rangle$ (or the second after a period doubling, etc.). Computing $P(x,p)$ requires a time integrator able to stop when the condition $\varphi(t,x,p) \in \Sigma_1$ holds. This can be done by tracking the sign of $\langle v_1, \varphi(t, x, p) - x_1^\Sigma \rangle$ and apply a root-finding method to compute the intersection. This is implemented in many libraries of numerical integrators of ODEs.

Since the map is restricted to Σ_1 the action of the Jacobian of $P(x,p)$ has to be restricted to vectors $(\delta x, \delta p)$ with $\langle v_1, \delta x \rangle = 0$. The action is

$$w \equiv D_x P(x,p)\delta x + D_p P(x,p)\delta p = y_1 - \frac{\langle v_1, y_1 \rangle}{\langle v_1, z \rangle} z, \tag{27}$$
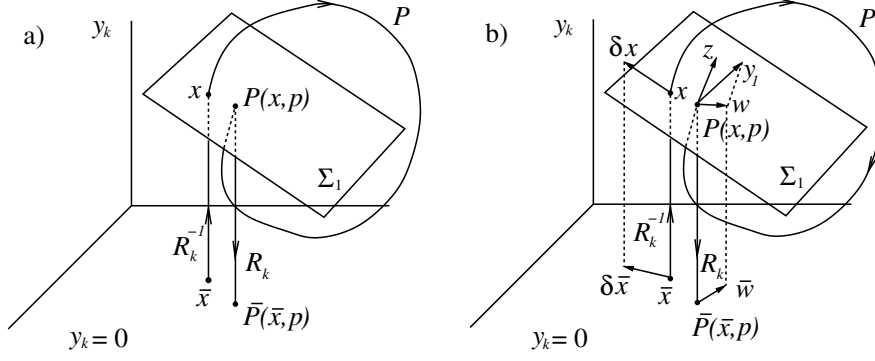
Fig. 1. Diagrams of a) the parameterized Poincaré map and b) the action by its Jacobian.

where $z = f(P(x,p), p)$, and y_1 is the solution, at the arrival time, $t(x,p)$, of the first variational equation (14). To see this, $P(x,p) = \varphi(t(x,p), x, p)$ is first differentiated to obtain

$$D_x P(x,p)\delta x + D_p P(x,p)\delta p = D_x\varphi(t(x,p), x, p)\delta x + D_p\varphi(t(x,p), x, p)\delta p$$
$$+ f(\varphi(t(x,p), x, p), p)\big(D_x t(x,p)\delta x + D_p t(x,p)\delta p\big) = y_1(t(x,p)) + cz,$$

with

$$y_1(t) = D_x\varphi(t, x, p)\delta x + D_p\varphi(t, x, p)\delta p,$$
$$c = D_x t(x,p)\delta x + D_p t(x,p)\delta p \in \mathbb{R},$$
$$z = f(\varphi(t(x,p), x, p), p).$$

Since $h_1(x,p) \equiv \langle v_1, P(x,p) - x_1^\Sigma \rangle = 0 \ \forall x \in \mathcal{V}$ and $\forall p$, its derivative must vanish identically for all δx satisfying $\langle v_1, \delta x \rangle = 0$, then

$$D_x h_1(x,p)\delta x + D_p h_1(x,p)\delta p = \langle v_1, y_1 \rangle + c\langle v_1, z \rangle = 0,$$

and $c = -\langle v_1, y_1 \rangle / \langle v_1, z \rangle$.

Since Σ_1 has dimension $n-1$ it is convenient to parameterize it and the Poincaré map. If v_{1_k} is the largest component of v_1, let us define R_k as the orthogonal projection from Σ_1 onto the hyperplane $y_k = 0$, where the subscript indicates here the k-th component of y. The inverse of this map is a regular parameterization of Σ_1. The map R_k simply removes the component of index k of a point x on Σ_1 so that $\bar{x} = R_k(x) \in \mathbb{R}^{n-1}$, and its inverse R_k^{-1} inserts a new component at the position k of a point $\bar{x} \in \mathbb{R}^{n-1}$ such that $x = R_k^{-1}(\bar{x}) \in \Sigma_1$ (see Fig.1). Then the map

$$\bar{P}(\bar{x}, p) = R_k(P(R_k^{-1}(\bar{x}), p))$$

is a projected Poincaré map defined on $R_k(\mathcal{V}) \subset \{y \in \mathbb{R}^n \ / \ y_k = 0\}$. By applying the chain rule to $\bar{P}(\bar{x}, p) = R_k(P(R_k^{-1}(\bar{x}), p))$

$$D_{\bar{x}}\bar{P}(\bar{x}, p)\delta\bar{x} + D_p\bar{P}(\bar{x}, p)\delta p = R_k\big(D_x P(x,p)DR_k^{-1}(\bar{x})\delta\bar{x} + D_p P(x,p)\delta p\big) = R_k w$$

where $\delta x = DR_k^{-1}(\bar{x})\delta\bar{x}$, and w is defined by Eq. (27).

The periodic orbits of (1) can be computed as fixed points of the parameterized version of the Poincaré map. Its fixed points, verifying $\bar{x} - \bar{P}(\bar{x}, p) = 0, \quad \bar{x} \in \mathbb{R}^{n-1}$, are in one-to-one correspondence with those of P by the map $x = R_k^{-1}(\bar{x})$. This was the method used originally in [19]. All the computations required are represented schematically in Fig. 1. The multiple shooting version using parallelism was studied in [20]. The extension is not trivial if some speedup is expected.
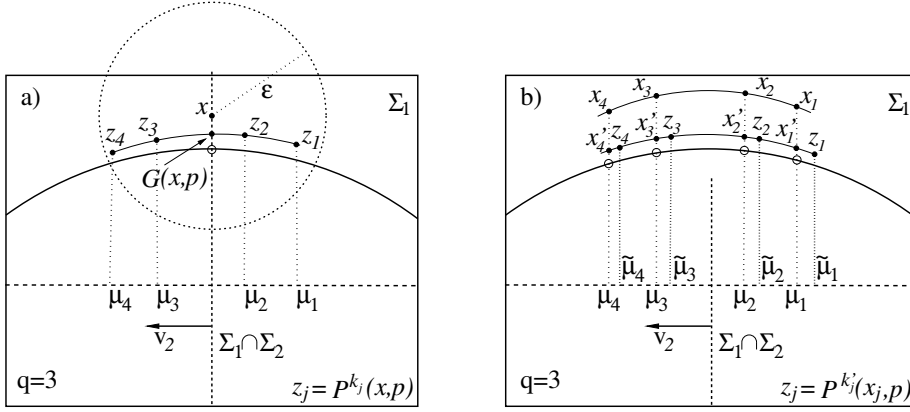
Fig. 2. Diagrams of the two generalized Poincaré maps used to compute invariant tori when the degree of interpolation is three.

10.2 Invariant tori

The first method to compute invariant tori tries to find a single point of the invariant curve. Similar methods were used in the past for low-dimensional dissipative [53] and conservative [54] problems. Assume that a Poincaré map has been defined as before on a hyperplane Σ_1, and let Σ_2 be another hyperplane, given by $\langle v_2, x - x_2^\Sigma \rangle = 0$, cutting transversely both Σ_1 and the invariant tori, i.e., it cuts the invariant curve at the point one wants to obtain. It is found as the fixed point of a generalized Poincaré map $G(\bar{x}, p) : \mathcal{U} \subset \mathbb{R}^{n-2} \times \mathbb{R} \to \mathbb{R}^{n-2}$ defined as follows. Let us fix $\varepsilon > 0$, and let $R : \Sigma_1 \cap \Sigma_2 \to \mathbb{R}^{n-2}$ be the orthogonal projection of $\Sigma_1 \cap \Sigma_2$ onto the $n-2$ manifold $\{y \in \mathbb{R}^n \ / \ y_{k_1} = 0, \ y_{k_2} = 0\}$, k_1 and k_2 being the indices of the components of maximal modulus of v_1 and v_2, respetively. If $\bar{x} \in \mathcal{U}$ let $x = R^{-1}(\bar{x}) \in \Sigma_1 \cap \Sigma_2$ and $z_j = P^{k_j}(x, p)$, $j = 1, \dots, q+1$ be the first $q+1$ powers of P applied to x such that $\|P^{k_j}(x, p) - x\| < \varepsilon$. Then

$$G(\bar{x}, p) = R\left(\sum_{j=1}^{q+1} l_j(0) P^{k_j}(R^{-1}(\bar{x}), p)\right) = R\left(\sum_{j=1}^{q+1} l_j(0) z_j\right), \qquad (28)$$

where the l_j are the Lagrange interpolation polynomials of degree q, associated to the mesh of points μ_j, and evaluated at $\mu = 0$ with

$$l_j(\mu) = \prod_{\substack{i=1 \\ i \neq j}}^{q+1} \frac{\mu - \mu_i}{\mu_j - \mu_i}, \quad \mu_j = \langle v_2, z_j - x_2^\Sigma \rangle, \quad j = 1, \cdots, q+1.$$

The variable μ is a local coordinate, perpendicular to Σ_2, which parameterizes the invariant curve close to the initial point x, and with origin $\mu = 0$ at Σ_2. What the maps does is finding the intersection of the curve interpolating the z_j with Σ_2. Since $z_j \in \Sigma_1$, $j = 1, \cdots, q+1$, the intersection will be in $\Sigma_1 \cap \Sigma_2$. Fig 2a shows a scheme of the process. The solutions of $\bar{x} - G(\bar{x}, p) = 0$ correspond, via $x = R^{-1}(\bar{x})$, to an approximation of the point of the invariant tori in $\Sigma_1 \cap \Sigma_2$.

The map $G(\bar{x}, p)$ depends on two parameters, the radius ε and the degree of the interpolation polynomial q. At least ε has to be varied adaptively when the parameter p is moved in a continuation process. When resonance regions are crossed a point on

the manifolds connecting stable and unstable periodic orbits on the torus will be found provided the tori persists and the returning points z_j can still be computed.

To apply Newton-Krylov methods to the fixed point problem $\bar{x} - G(\bar{x}, p) = 0$ it is necessary to compute the action of the Jacobian of $G(\bar{x}, p)$ on vectors $(\delta\bar{x}, \delta p)$. It reduces to the case of the differential of the Poincaré map:

$$D_{\bar{x}} G(\bar{x}, p)\delta\bar{x} + D_p G(\bar{x}, p)\delta p = R \sum_{j=1}^{q+1} \Big[l_j(0)\big(D_x P^{k_j}(x, p)\delta x + D_p P^{k_j}(x, p)\delta p\big)$$

$$+ P^{k_j}(x, p) \sum_{i=1}^{q+1} \partial_{\mu_i} l_j(0)\langle v_2, D_x P^{k_i}(x, p)\delta x + D_p P^{k_i}(x, p)\delta p\rangle \Big],$$

with $x = R^{-1}(\bar{x})$ and $\delta x = D_{\bar{x}} R^{-1}(\bar{x})\delta\bar{x}$, and where

$$\langle v_2, D_x P^{k_i}(x, p)\delta x + D_p P^{k_i}(x, p)\delta p\rangle = D_x \mu_i \delta x + D_p \mu_i \delta p.$$

The derivatives $\partial_{\mu_i} l_j(0)$ of $l_j(0) = \prod_{i=1, i \neq j}^{q+1} \mu_i/(\mu_i - \mu_j)$ are trivial.

This first method is clearly sequential. The computation of z_j must be preceded by that of z_{j+1}. The second method is a parallelizable version (see Fig. 2b). Instead of finding a single point on the invariant curve, several points are found which approximate a segment. With the same meaning for P, Σ_1 and Σ_2, let $\mu_1, \cdots, \mu_{q+1}$ be $q+1$ fixed coordinates along the line $x = x_2^\Sigma + \mu v_2$, let us fix $\varepsilon > 0$ and define the map $G(X, p) : \mathcal{U} \subset \mathbb{R}^{(n-1)(q+1)} \times \mathbb{R} \to \mathbb{R}^{(n-1)(q+1)}$ as follows. If $X = [x_1, \cdots, x_{q+1}] \in \mathcal{U}$ is a matrix with columns the coordinates of the points $x_1, \ldots, x_{q+1}$, let $z_j = P^{k'_j}(x_j, p)$ be the first power of P applied to x_j such that $\|P^{k'_j}(x_j, p) - x_j\| < \varepsilon$, $j = 1, \ldots, q+1$. Then

$$G(X, p) = X' = Z(X, p)\tilde{V}(X, p)^{-1}V, \tag{29}$$

with $X' = (x'_1, \cdots, x'_{q+1})$, $Z(X, p) = (z_1, \cdots, z_{q+1})$, and V and $\tilde{V}$ being the Vandermonde matrices

$$V = \begin{pmatrix} 1 & \cdots & 1 \\ \mu_1 & \cdots & \mu_{q+1} \\ \cdot & \cdot & \cdot \\ \mu_1^q & \cdots & \mu_{q+1}^q \end{pmatrix}, \quad \tilde{V} = \begin{pmatrix} 1 & \cdots & 1 \\ \tilde{\mu}_1 & \cdots & \tilde{\mu}_{q+1} \\ \cdot & \cdot & \cdot \\ \tilde{\mu}_1^q & \cdots & \tilde{\mu}_{q+1}^q \end{pmatrix},$$

associated with the μ_j and the $\tilde{\mu}_j = \langle v_2, P^{k'_j}(x_j, p) - x_2^\Sigma\rangle$, $\quad j = 1, \cdots, q+1$ respectively.

The map takes $q+1$ initial conditions x_j approximating a segment (see Fig. 2b), finds the first return point z_j of each one, interpolates them, and evaluates the interpolating polynomial at the original coordinates μ_j. The product $Z(X, p)\tilde{V}(X, p)^{-1}$ contains the vectorial coefficients of the monomials $1, \mu, \mu^2, \ldots, \mu^{q+1}$. Multiplying them by V is evaluating the interpolation polynomial at the nodes μ_j. A fixed point of $G(X, p)$ provides an approximation of an arc of the invariant curve in Σ_1.

In this case the parameterizations of the intersections of Σ_1 with the planes $\mu = \mu_j$ have not been made explicit. They are exactly the same as for the first method. The radius ε, and the position of the μ_j defining G should be varied adaptively during the continuation process. Finally, the action by the Jacobian of this second map, which can be written as $G = Z\tilde{V}^{-1}V$ to simplify, can also be reduced to that of the Poincaré map. If $\delta X = [\delta x_1, \cdots, \delta x_{q+1}]$ then

$$DG(X, p)(\delta X, \delta p) =$$

$$\Big[DZ(X, p)(\delta X, \delta p) - Z(X, p)\tilde{V}(X, p)^{-1} D\tilde{V}(X, p)(\delta X, \delta p) \Big] \tilde{V}(X, p)^{-1}V$$

where $DZ(X,p)(\delta X,\delta p) = [DP^{k'_1}(x_1,p)(\delta x_1,\delta p),\cdots,DP^{k'_{q+1}}(x_{q+1},p)(\delta x_{q+1},\delta p)]$,

$$
D\tilde{V}(X,p)(\delta X,\delta p) = \begin{pmatrix} 0 & \cdots & 0 \\ 1 & \cdots & 1 \\ 2\tilde{\mu}_1 & \cdots & 2\tilde{\mu}_{q+1} \\ \cdots\cdots\cdots\cdots\cdots \\ q\tilde{\mu}_1^{q-1} & \cdots & q\tilde{\mu}_{q+1}^{q-1} \end{pmatrix} \begin{pmatrix} \eta_1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & \eta_{q+1} \end{pmatrix},
$$

and $\eta_j = \langle v_2, DP^{k'_j}(x_j,p)(\delta x_j,\delta p)\rangle$. In short, $DG = \left[DZ - Z\tilde{V}^{-1}D\tilde{V} \right] \tilde{V}^{-1}V$.

In [23] the two methods were compared. The main conclusion is that the second should be preferred. It is trivially parallelizable because finding each z_j is now an independent task, and, if the x_j are close together, the number of iterations of the linear solver is quite independent of the number of points, $q+1$. Moreover the integration time required to find the x_j is much lower because only the first return point is needed. This allows also to find more unstable invariant curves than with the first method. The final result is that good speedups can be achieved with parallel implementations. Other formulations are possible based on splines interpolation or least squares fitting. Both were also considered in [23].

11 Applications and conclusions

The tools required to analyze efficiently systems of parabolic nonlinear PDEs, from the point of view of dynamical systems, have been introduced in this tutorial. No examples have been provided to reduce the length of the article. Instead, some references will be given representing recent uses of these algorithms, without attempting to be exhaustive. The reviews [58,14] present the current status of the application of dynamical systems techniques in Fluid Dynamics, including several examples.

Since computing equilibria and their stability is now quite common, we will only mention [47] as typical study, because it is for an open flow at high Reynolds number, using a model for the turbulence. The authors compute the flow around four-digits symmetric NACA airfoils of several thickness. Continuation methods allow them to pass the stall point.

Waves and modulated waves can be computed by writing the governing equations in the frame at which they become steady and periodic solutions, respectively. This approach has been used in [59] for the study of localized structures in binary convection of a two-dimensional layer, and in [60,61] for the thermal convection in rotating spherical shells.

Several studies of periodic regimes in fluid flow problems, which are not just waves, have appeared after the publication of [19]. The problems treated included the plane Couette flow [62], the pipe flow [63], thermal convection in a cubical cavity [64], the thermoacoustics of flames [65,66], and the continuation of modulated waves for the thermal convection in rotating spherical shells in [67], among others. This multiplicity of works is due to the fact that the continuation and eigenvalue computation algorithms are functionally separated from that of the time evolution, and can therefore be developed pretty much independently allowing the use of adapted legacy simulation codes.

As far as we know, the methods to compute invariant tori introduced in [22,23] have not yet been used by other authors, probably because they might seem too complicated, or simply because there has been no need to study in detail extended systems in quasiperiodic regimes.

The possible appearance in the future of general software packages implementing all these tools, will relieve the users from the burden of having to program everything from scratch and promote their use, as happened for low-dimensional systems.

12 Acknowledgments

References

1. E. Doedel, AUTO: Software for Continuation and Bifurcation Problems in Ordinary Differential Equations, Tech. report, Applied Mathematics, California Institute of Technology, Pasadena CA (1986).
2. M. Kubíček, M. Marek, Computational Methods in Bifurcation Theory and Dissipative Structures, Springer-Verlag, 1983.
3. W. C. Rheinboldt, Numerical Analysis of Parametrized Nonlinear Equations, J. Wiley, 1986.
4. E. L. Allgower, K. Georg, Numerical Continuation Methods: An Introduction, Vol. 13 of Computational Mathematics, Springer, 1990.
5. R. Seydel, Practical bifurcation and stability analysis. From equilibrium to chaos, Springer, New York, 1994.
6. Y. A. Kuznetsov, Elements of Applied Bifurcation Theory, Springer, Berlin, 1998.
7. K. A. Cliffe, A. Spence, S. Taverner, The numerical analysis of bifurcation problems with applications to fluid mechanics, Acta Numer. (2000) 39–131.
8. H. Dankowicz, F. Schilder, Recipes for Continuation, Computational Science and Engineering, SIAM, 2013.
9. E. Riks, The application of Newton's method to the problem of elastic stability, ASME J. Appl. Mech. 39 (4) (1971) 1060–1065.
10. R. Meyer-Spasche, H. B. Keller, Computation of the axisymmetric flow between rotating cylinders, J. Comput. Phys. 35 (1980) 100–109.
11. Y. Saad, M. H. Schultz, GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems, SIAM J. Sci. Stat. Comput. 7 (1986) 856–869.
12. R. B. Lehoucq, D. C. Sorensen, C. Yang, ARPACK User's Guide: Solution of Large-Scale Eigenvalue Problems with Implicitly Restarted Arnoldi Methods, Software, Environments, Tools, SIAM, 1998.
13. E. Doedel, L. S. Tuckerman (Eds.), Numerical Methods for Bifurcation Problems and Large-Scale Dynamical Systems, Vol. 119 of IMA Volumes in Mathematics and its Applications, Springer–Verlag, 2000.
14. H. A. Dijkstra, F. W. Wubs, A. K. Cliffe, E. Doedel, I. F. Dragomirescu, B. Eckhardt, A. Gelfgat, A. Hazel, V. Lucarini, A. Salinger, J. Sánchez, H. Schuttelaars, L. Tuckerman, U. Thiele, Numerical bifurcation methods and their application to fluid dynamics: Analysis beyond simulation, Commun. Comput. Phys. 15 (1) (2014) 1–45.
15. A. G. Salinger, N. M. Bou-Rabee, R. P. Pawlowsky, E. D. Wilkes, E. A. Burroughs, R. B. Lehoucq, L. A. Romero, LOCA 1.1. Library of Continuation Algorithms: Theory and Implementation Manual, Sandia National Laboratories, Albuquerque, NM (2002).
16. H. Uecker, D. Wetzel, J. Rademacher, pde2path - a matlab package for continuation and bifurcation in 2d elliptic systems, Num. Math.: Th. Meth. Appl. 7 (2014) 58–106.
17. K. Lust, D. Roose, A. Spence, A. Champneys, An adaptive Newton–Picard algorithm with subspace iteration for computing periodic solutions, SIAM J. Sci. Comput. 19 (4) (1998) 1188–1209.
18. T. L. van Noorden, S. M. Verduyn Lunel, A. Bliek, A Broyden rank p+1 update continuation method with subspace iteration, SIAM J. Sci. Comput.

19. J. Sánchez, M. Net, B. García-Archilla, C. Simó, Newton-Krylov continuation of periodic orbits for Navier-Stokes flows, J. Comput. Phys. 201 (1) (2004) 13–33.
20. J. Sánchez, M. Net, On the multiple shooting continuation of periodic orbits by Newton-Krylov methods, Int. J. Bifur. Chaos 20 (1) (2010) 1–19.
21. L. van Veen, G. Kawahara, M. Atsushi, On matrix-free computation of 2D unstable manifolds, SIAM J. Sci. Comput. 33 (1) (2011) 25–44.
22. J. Sánchez, M. Net, C. Simó, Computation of invariant tori by Newton-Krylov methods in large-scale dissipative systems, Physica D 239 (2010) 123–133.
23. J. Sánchez, M. Net, A parallel algorithm for the computation of invariant tori in large-scale dissipative systems, Physica D 252 (1) (2013) 22–33.
24. M. Net, J. Sánchez, Continuation of bifurcations of periodic orbits for large-scale systems, SIAM J. Appl. Dynam. Systems 14 (2) (2015) 674–698.
25. Y. Saad, Iterative methods for sparse linear systems, PWS pub. company, New York, 1996.
26. T. L. van Noorden, S. M. Verduyn Lunel, A. Bliek, The efficient computation of periodic states of cyclically operated chemical processes, IMA J. Appl. Math. 68 (2003) 149–166.
27. K. I. Dickson, C. T. Kelley, I. C. F. Ipsen, I. G. Kevrekidis, Condition estimates for pseudo-arclength continuation, SIAM J. Numer. Anal. 45 (1) (2007) 263–276.
28. C. T. Kelley, Iterative methods for linear and nonlinear equations, Frontiers in applied mathematics, SIAM, 1995.
29. R. S. Dembo, S. C. Eisenstat, T. Steihaug, Inexact Newton methods, SIAM J. Numer. Anal. 19 (2) (1982) 400–408.
30. J. E. Dennis, Jr., R. B. Schnabel, Numerical Methods for Unconstrained Optimization and Nonlinear Equations (Classics in Applied Mathematics, 16), SIAM, 1996.
31. S. C. Eisenstat, H. F. Walker, Choosing the forcing terms in an inexact newton method, SIAM J. Sci. Comput. 17 (1) (1996) 16–32.
32. M. Pernice, H. F. Walker, NITSOL: A Newton iterative solver for nonlinear systems, SIAM J. Sci. Comput. 19 (1) (1998) 302–318.
33. S. L. Campbell, I. C. F. Ipsen, C. T. Kelley, C. D. Meyer, GMRES and the Minimal Polynomial, BIT Numerical Mathematics 36 (4) (1996) 664–675.
34. R. Barrett, M. Berry, T. F. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, H. V. der Vorst, Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods, SIAM, Philadelphia, 1994.
35. V. Frayssé, L. Giraud, S. Gratton, J. Langou, Algorithm 842: A set of GMRES routines for real and complex arithmetics on high performance computers, ACM Trans. Math. Softw. 31 (2) (2005) 228–238.
36. Y. Saad, Numerical Methods for Large Eigenvalue Problems, Manchester University Press, Manchester, 1992.
37. W. Stewart, A. Jennings, Algorithm 570: LOPSI: A simultaneous iteration algorithm for real matrices, ACM Trans. Math. Soft. 7 (2) (1981) 230–232.
38. Z. Bai, G. W. Stewart, Algorithm 776: SRRIT: A fortran subroutine to calculate the dominant invariant subspace of a nonsymmetric matrix, ACM Trans. Math. Softw. 23 (4) (1997) 494–513.
39. E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, LAPACK Users' Guide, Third Edition, SIAM, 1999.
40. D. C. Sorensen, Implicit application of polynomial filters in a k-step Arnoldi method, SIAM J. Matrix Anal. Appl. 13 (1) (1992) 357–385.
41. B. García-Archilla, J. Sánchez, C. Simó, Krylov methods and test functions for detecting bifurcations in one parameter-dependent partial differential equations, BIT 46 (4) (2006) 731–757.
42. K. Meerbergen, D. Roose, Matrix transformations for computing rightmost eigenvalues of large sparse non-symmetric eigenvalue problems, IMA J. Numer. Anal. 16 (3) (1996) 297–346.
43. A. Fortin, M. Jardak, J. J. Gervais, R. Pierre, Localization of Hopf bifurcations in fluid flow problems, Int. J. Numer. Meth. Fluids 24 (11) (1997) 1185–1210.

44. M. Rieutord, L. Valdettaro, Inertial waves in a rotating spherical shell, J. Fluid Mech. 341 (1997) 77–99.
45. J. Sánchez, F. Garcia, M. Net, Radial collocation methods for the onset of convection in rotating spheres, J. Comput. Phys. 308 (2016) 273 – 288.
46. R. B. Lehoucq, A. G. Salinger, Large-scale eigenvalue calculations for stability analysis of steady flows on massively parallel computers, Int. J. Numer. Meth. Fluids 36 (3) (2001) 309–327.
47. C. Wales, A. L. Gaitonde, D. P. Jones, D. Avitabile, A. R. Champneys, Numerical continuation of high reynolds number external flows, International Journal for Numerical Methods in Fluids 68 (2) (2012) 135–159.
48. M. Net, F. Garcia, J. Sánchez, On the onset of low-Prandtl-number convection in rotating spherical shells: non-slip boundary conditions, J. Fluid Mech. 601 (2008) 317–337.
49. D. Barkley, R. D. Henderson, Floquet stability analysis of the periodic wake of a circular cylinder, J. Fluid Mech. 322 (1996) 215–241.
50. W.-J. Beyn, A. Champneys, E. Doedel, W. Govaerts, Y. A. Kuznetsov, B. Sandstede, Chapter 4, Numerical continuation, and computation of normal forms, Vol. 2 of Handbook of Dynamical Systems, Elsevier Science, 2002, pp. 149 – 219.
51. A. G. Salinger, E. A. Burroughs, R. P. Pawlowski, E. T. Phipps, L. A. Romero, Bifurcation tracking algorithms and software for large scale applications, Int. J. Bifur. Chaos 15 (03) (2005) 1015–1032.
52. L. O. Chua, A. Ushida, Algorithms for computing almost periodic steady-state response of nonlinear systems to multiple input frequencies, IEEE Tran. Circuits Sys. 28 (10) (1981) 953–971.
53. C. Kaas-Petersen, Computation, continuation, and bifurcation of torus solutions for dissipative maps and ordinary differential equations, Physica D 25 (1–3) (1987) 288–306.
54. C. Simó, Effective Computations in Hamiltonian Dynamics, in: Cent ans après les Méthodes Nouvelles de H. Poincaré, Société Mathématique de France, 1996, pp. 1–23.
55. C. Simó, Effective computations in celestial mechanics and astrodynamics, in: V. Rumyantsev, A. Karapetyan (Eds.), Modern Methods of Analytical Mechanics and their Applications (CISM Courses and Lectures 387), Springer, 1998, pp. 55–102.
56. À. Jorba, Numerical computation of the normal behaviour of invariant curves on n-dimensional maps, Nonlinearity 14 (2001) 943–976.
57. F. Schilder, H. M. Osinga, W. Vogt, Continuation of Quasi-periodic Invariant Tori, SIAM J. Appl. Dynam. Systems 4 (3) (2005) 459–488.
58. G. Kawahara, M. Uhlmann, L. van Veen, The significance of simple invariant solutions in turbulent flows, Ann. Rev. Fluid Mech. 44 (1) (2012) 203–225.
59. T. Watanabe, M. Iima, Y. Y. Nishiura, Spontaneous formation of travelling localized structures and their asymptotic behaviour in binary fluid convection, J. Fluid Mech. 712 (2012) 219–243.
60. J. Sánchez, F. Garcia, M. Net, Computation of azimuthal waves and their stability in thermal convection in rotating spherical shells with application to the study of a double-Hopf bifurcation, Phys. Rev. E 87 (2013) 033014.
61. F. Feudel, L. S. Tuckerman, M. Gellert, N. Seehafer, Bifurcations of rotating waves in rotating spherical shell convection, Phys. Rev. E 92 (2015) 053015.
62. D. Viswanath, Recurrent motions within plane Couette turbulence, J. Fluid Mech. 580 (2007) 339–358.
63. Y. Duguet, C. C. T. Pringle, R. R. Kerswell, Relative periodic orbits in transitional pipe flow, Phys. Fluids 20 (11) (2008) 114102–.
64. D. Puigjaner, J. Herrero, C. Simó, F. Giralt, From steady solutions to chaotic flows in a Rayleigh–Bénard problem at moderate Rayleigh numbers, Phys. D 240 (2011) 920–934.
65. I. Waugh, S. Illingworth, M. Juniper, Matrix-free continuation of limit cycles for bifurcation analysis of large thermoacoustic systems, J. Comput. Phys. 240 (2013) 225–247.
66. I. C. Waugh, K. Kashinath, M. P. Juniper, Matrix-free continuation of limit cycles and their bifurcations for a ducted premixed flame, J. Fluid Mech. 759 (2014) 1–27.
67. F. Garcia, M. Net, J. Sánchez, Continuation and stability of convective modulated rotating waves in spherical shells, Phys. Rev. E 93 (2016) 013119.