HARRY DANKOWICZ • FRANK SCHILDER

Recipes for Continuation

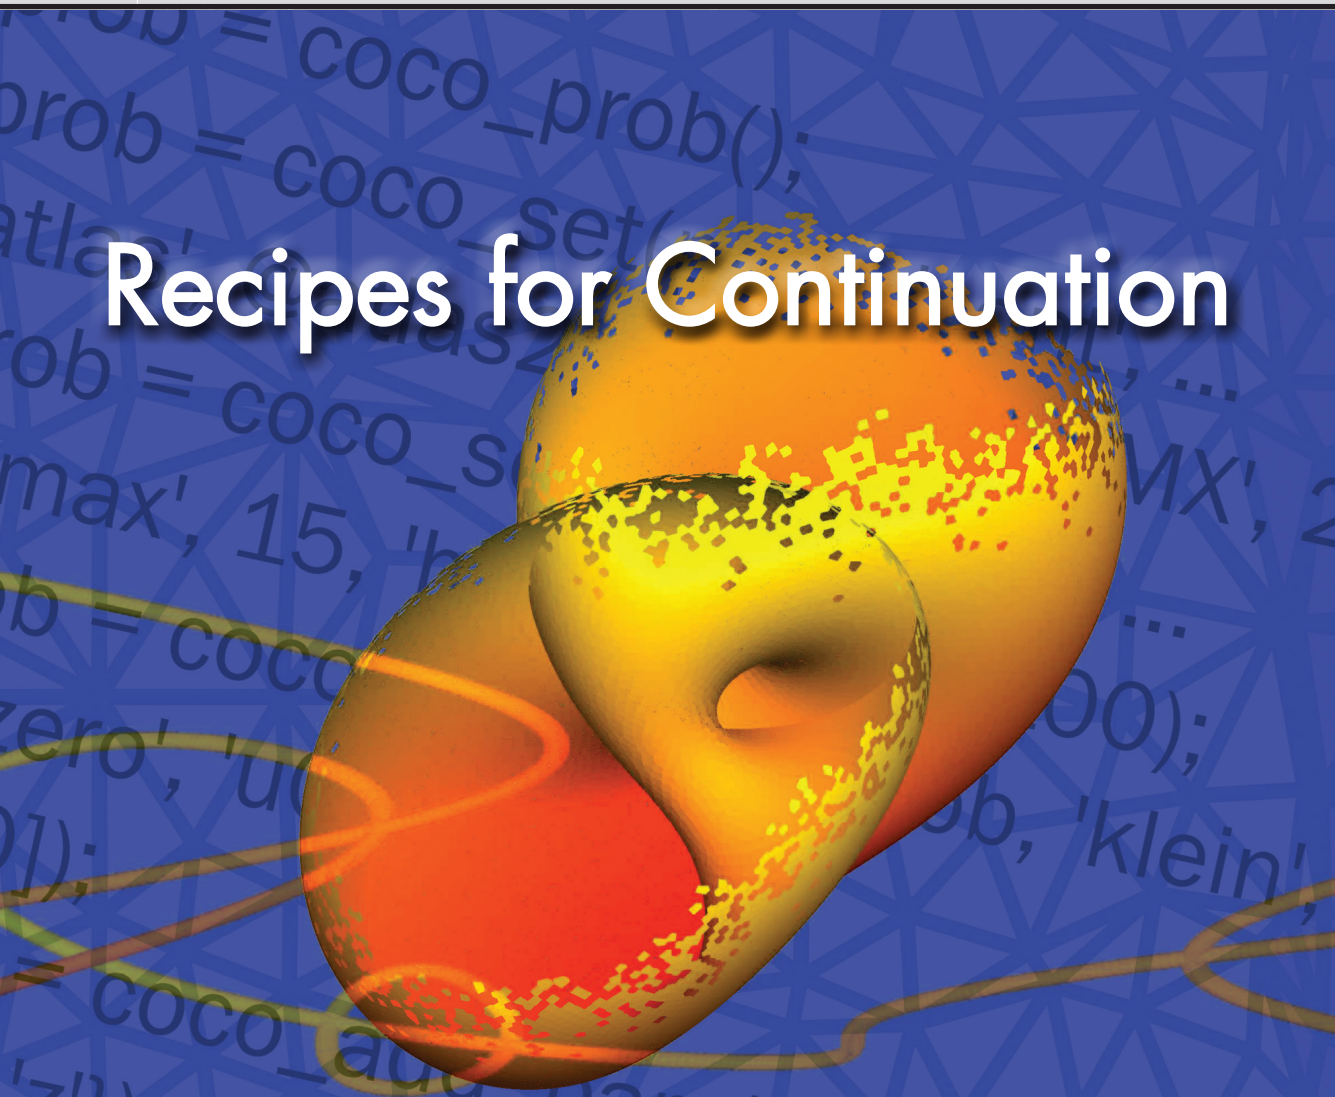

siam
Computational Science & Engineering

Recipes for Continuation

Computational Science & Engineering

The SIAM series on Computational Science and Engineering publishes research monographs, advanced undergraduate- or graduate-level textbooks, and other volumes of interest to an interdisciplinary CS&E community of computational mathematicians, computer scientists, scientists, and engineers. The series includes both introductory volumes aimed at a broad audience of mathematically motivated readers interested in understanding methods and applications within computational science and engineering and monographs reporting on the most recent developments in the field. The series also includes volumes addressed to specific groups of professionals whose work relies extensively on computational science and engineering.

SIAM created the CS&E series to support access to the rapid and far-ranging advances in computer modeling and simulation of complex problems in science and engineering, to promote the interdisciplinary culture required to meet these large-scale challenges, and to provide the means to the next generation of computational scientists and engineers.

Series Volumes

Dankowicz, Harry and Schilder, Frank, *Recipes for Continuation*

Mueller, Jennifer L. and Siltanen, Samuli, *Linear and Nonlinear Inverse Problems with Practical Applications*

Shapira, Yair, *Solving PDEs in C++: Numerical Methods in a Unified Object-Oriented Approach, Second Edition*

Borzì, Alfio and Schulz, Volker, *Computational Optimization of Systems Governed by Partial Differential Equations*

Ascher, Uri M. and Greif, Chen, *A First Course in Numerical Methods*

Layton, William, *Introduction to the Numerical Analysis of Incompressible Viscous Flows*

Ascher, Uri M., *Numerical Methods for Evolutionary Differential Equations*

Zohdi, T. I., *An Introduction to Modeling and Simulation of Particulate Flows*

Biegler, Lorenz T., Ghattas, Omar, Heinkenschloss, Matthias, Keyes, David, and van Bloemen Waanders, Bart, Editors, *Real-Time PDE-Constrained Optimization*

Chen, Zhangxin, Huan, Guanren, and Ma, Yuanle, *Computational Methods for Multiphase Flows in Porous Media*

Shapira, Yair, *Solving PDEs in C++: Numerical Methods in a Unified Object-Oriented Approach*

HARRY DANKOWICZ
University of Illinois at Urbana-Champaign
Urbana, Illinois

FRANK SCHILDER
Technical University of Denmark
Lyngby, Denmark

Recipes for Continuation

siam
Society for Industrial and Applied Mathematics
Philadelphia

10 9 8 7 6 5 4 3 2 1

This material is based on work supported by the National Science Foundation under grants 0237370, 0635469, and 1016467 (HD), by the Engineering and Physical Sciences Research Council under grant EP/D063906/1 (FS), and by the Danish Research Council FTP under project 09-065890/FTP (FS).

Library of Congress Cataloging-in-Publication Data
Dankowicz, Harry.
Recipes for continuation / Harry Dankowicz, University of Illinois at Urbana-Champaign, Urbana, Illinois, Frank Schilder, Technical University of Denmark, Lyngby, Denmark.
pages cm – (Computational science and engineering series)
Includes index.
ISBN 978-1-611972-56-6
1. Continuation methods. I. Schilder, Frank. II. Title.
QA377.D364 2013
515′.35–dc23
2012046499

Partial royalties from the sale of this book are placed in a fund to help students attend SIAM meetings and other SIAM-related activities. This fund is administered by SIAM, and qualified individuals are encouraged to write directly to SIAM for guidelines.

1 A Continuation Paradigm

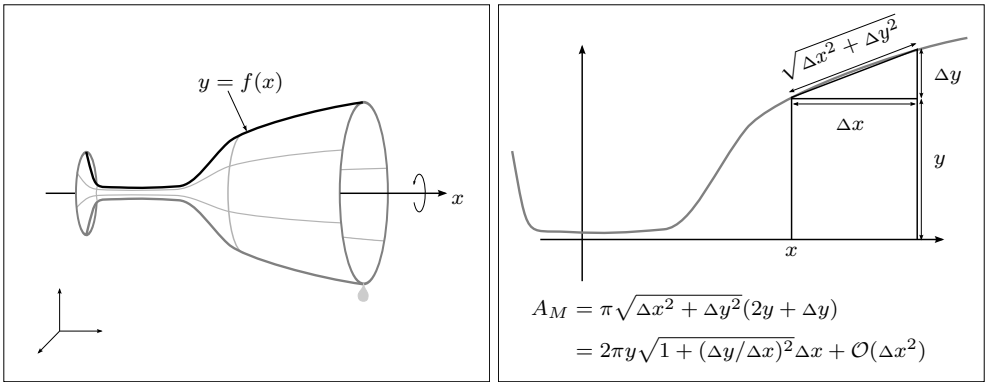

(a) Generating a surface of revolution. (b) Surface area of a cone section.

Figure 1.1. *A surface of revolution is generated by rotating a planar curve around an axis in the corresponding plane (a). If the curve is the graph $(x, f(x))$ of a function f, its surface area may be approximated by a* Riemann sum*, obtained by constructing a sequence of cone sections and summing their respective areas (b). The functional in Eq. (1.1) is obtained in the limit as $\max \Delta x \to 0$, where the maximum is taken over all cone sections.*

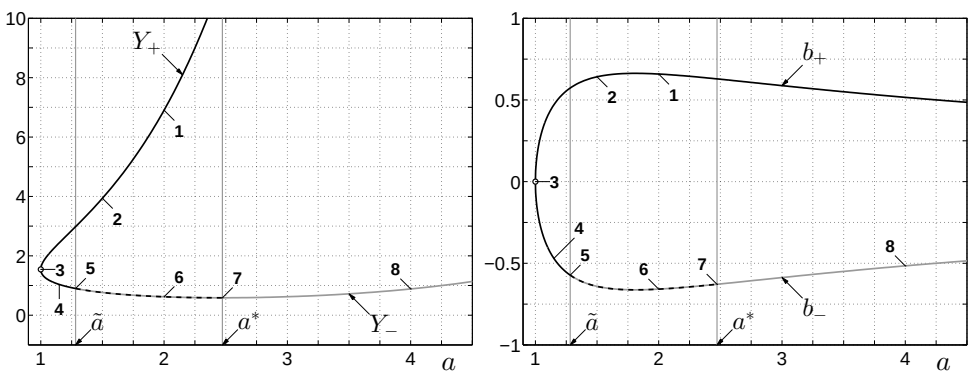

(a) $Y_{+/-}$ defined in Eqs. (1.16) and (1.17). (b) $b_{+/-}$ defined in Eqs. (1.14) and (1.15).

Figure 1.2. *Values of a, b, and Y corresponding to extremal curves of the integral functional in Eq. (1.3) with action integrand given in Eq. (1.2) and satisfying the given boundary conditions at $x = 0$ and $x = 1$. Points found along the solid black segment correspond to global minimizers, whereas those corresponding only to weak local minimizers are found along the dashed segment $(a, b_-(a), Y_-(a))$, for $a \in (\tilde{a}, a^*)$ (see Sects. 1.2.1 and 1.2.3 for definitions of a^* and $\tilde{a}$, respectively). Selected extremal curves for the points with labels 1-8 are shown in Fig. 1.3.*

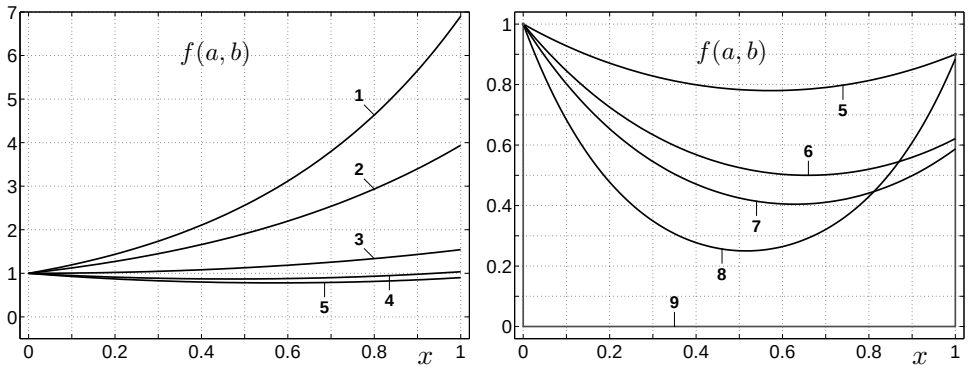

(a) Global (1-4) and local (5) minimizers.

(b) Local mini-mizers (5-7), stationary points (8), and a boundary curve (9).

Figure 1.3. *Extremal curves of the integral functional in Eq. (1.3) with action integrand given in Eq. (1.2) and satisfying the given boundary conditions at $x = 0$ and $x = 1$. The labels correspond to the points marked in Fig. 1.2. Global minimizers are shown in panel (a). Curves 5, 6, and 7 in panel (b) are weak local minimizers, while curve 8 is a stationary curve, but not a minimizer. The piecewise-straight curve 9 is the boundary curve defined in Sect. 1.2.3. The surfaces of revolution generated by curves 5 and 9 have the same surface area.*

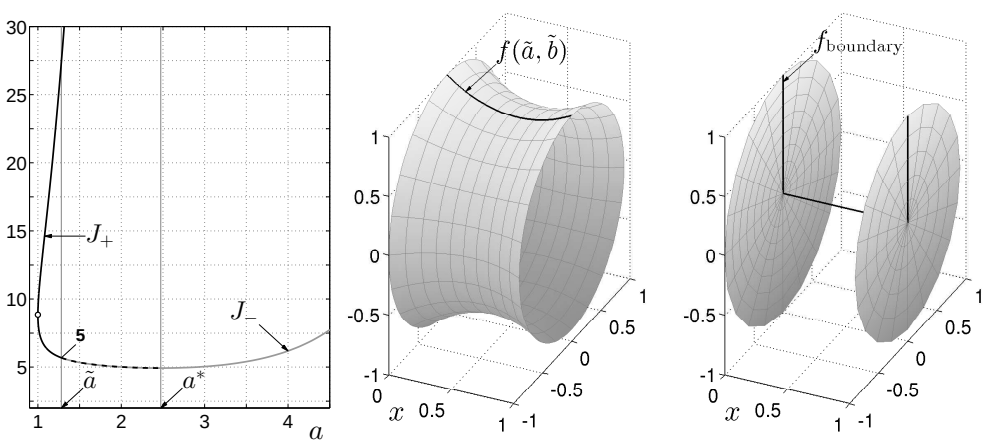

(a) Graphs of $J_\pm$.

(b) Surface generated by 5.

(c) Surface generated by 9.

Figure 1.4. *Graphs of the functions $J_\pm$ defined in* (1.32) *along the family of stationary solution curves (a). Points found along the solid black segment correspond to global minimizers, whereas those corresponding only to weak local minimizers are found along the dashed segment $(a, J_-(a))$, for $a \in (\tilde{a}, a^*)$ (see Sects. 1.2.1 and 1.2.3 for definitions of a^* and $\tilde{a}$, respectively). Panels (b)-(c) show the surfaces of revolution generated by curves 5 and 9 in Figure 1.3, respectively. The surface generated by the "boundary" curve defined in Sect. 1.2.3 consists of two circular disks.*

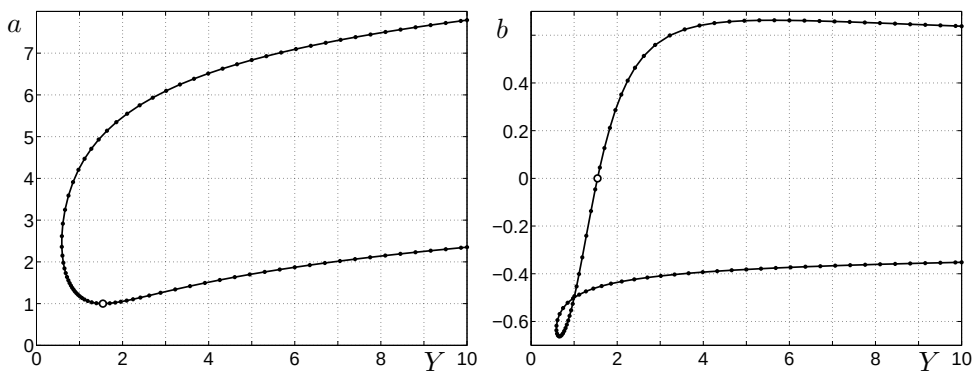

(a) Solution curve in Y-a-plane. (b) Solution curve in Y-b-plane.

Figure 1.5. *A branch of solutions of Eqs. (1.12-1.13) obtained using numerical continuation, as described in Sect. 1.3.1. Here, dots indicate computed points, which are connected with straight-line segments. The circle denotes the initial point at* $a = 1$, $b = 0$, *and* $Y = \cosh 1$. *Compare the numerical results with the analysis shown in Fig. 1.2.*

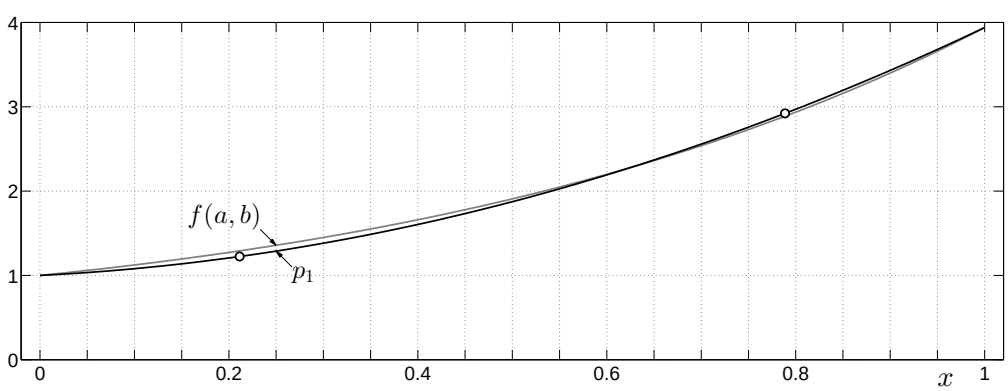

Figure 1.6. *Comparison of the exact solution (gray), obtained from Eq. (1.11), with a numerical solution of the form in Eq. (1.35) of degree* $m = 2$ *for the case when* $b = b_+(a)$, $Y = Y_+(a)$, *and* $a = 1.5$. *The two collocation nodes are marked with circles. For increasing order* m, *we observe spectral convergence; see Fig. 1.7.*

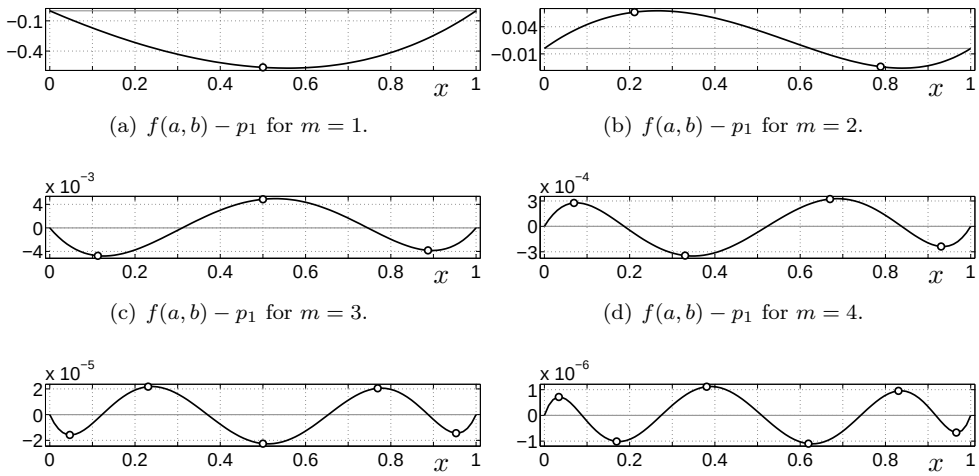

(a) $f(a,b) - p_1$ for $m = 1$.

(b) $f(a,b) - p_1$ for $m = 2$.

(c) $f(a,b) - p_1$ for $m = 3$.

(d) $f(a,b) - p_1$ for $m = 4$.

(e) $f(a,b) - p_1$ for $m = 5$.

(f) $f(a,b) - p_1$ for $m = 6$.

Figure 1.7. *Graphs of the approximation errors $\delta_m := f(a,b) - p_1$ of polynomial approximants p_1 of degree m as in Eq. (1.35), for the case when $b = b_+(a)$, $Y = Y_+(a)$, and $a = 1.5$. In each case, the m collocation nodes are indicated by circles. Panel (b) corresponds to the graphs in Fig. 1.6. We observe spectral convergence, i.e., $\|\delta_m\| < CM^{-m}$ for some $0 < C < \infty$ and $M > 1$.*

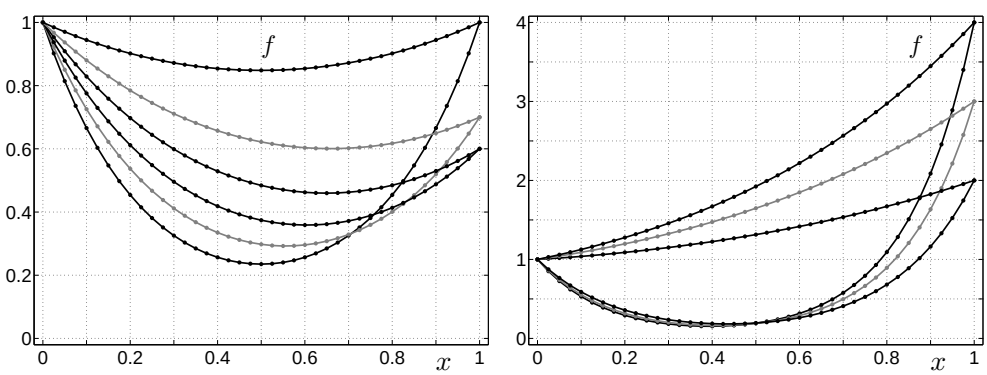

(a) Solutions of the boundary-value problem formulation.

(b) Solutions of the quadrature approximation.

Figure 1.8. *Sample approximate extremal curves of the integral functional in Eq. (1.3) with action integrand given in Eq. (1.2) and satisfying the given boundary conditions at $x = 0$ and $x = 1$. The results shown here were obtained using a collocation method with $N = 10$ and $m = 4$ applied to a two-point boundary value problem (a) as defined in Sect. 1.3.2, and to a quadrature approximation (b) as defined in Sect. 1.3.3. Each panel shows pairs of stationary curves for each given value of Y.*

2 Encapsulation

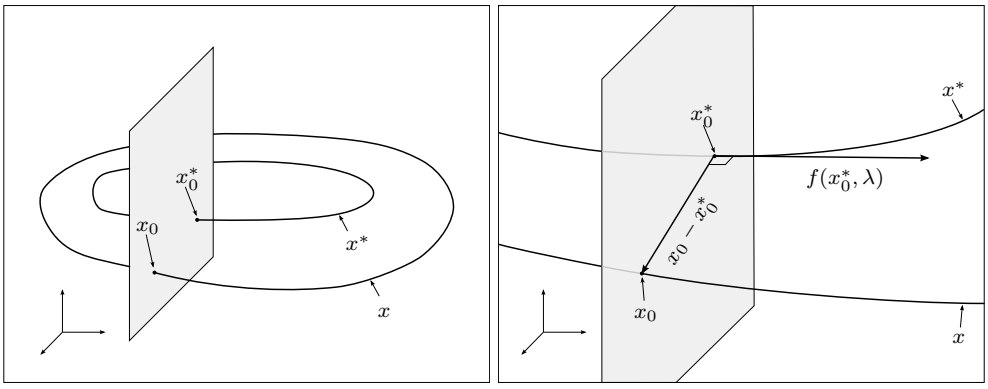

(a) Initial points in hyperplane transversal to x^* and x.

(b) Fixing x_0 in plane through x_0^* and normal to $f(x_0^*, \lambda)$.

Figure 2.1. *Due to translational invariance, a periodic solution of an autonomous ordinary differential equation has a free phase, i.e., the point $x_0 := x(0)$ is not defined uniquely. A typical way to obtain a unique initial point is to require that it lie in a hyperplane (a), transversal to the state-space representation of the solution. If the reference function x^* is sufficiently close to the sought-after orbit x, as seen in (b), one may use the hyperplane through x_0^* and perpendicular to the tangent vector $f(x_0^*, \lambda)$.*

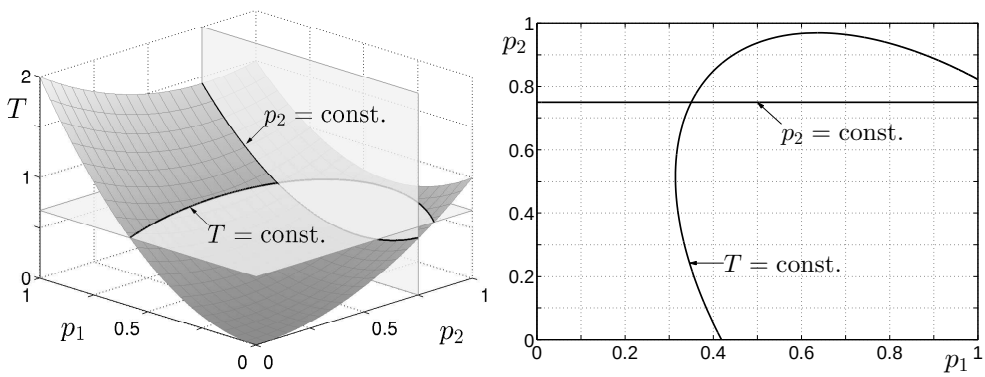

(a) 2-dimensional solution manifold.

(b) 1-dimensional solution curves.

Figure 2.2. *A manifold of periodic orbits of a hypothetical dynamical system depending on two parameters p_1 and p_2, projected onto the hyperplane with coordinates p_1, p_2, and T. Codimension-1 intersections of the manifold with hyperplanes corresponding to constant values of the parameter p_2 or of the period T, respectively, are shown as curves in (a), and projected onto the (p_1, p_2)-parameter plane in (b). In the terminology of Sect. 2.2, each curve is obtained by applying continuation to a suitably restricted continuation problem that is initializable with a 1-dimensional atlas algorithm.*

3 Construction

$$\begin{tikzcd}
(v_1, v_s, v_2) \arrow[r, "\mathbb{K}_1"'] \arrow[d, "\Phi"] & (v_1, v_s) \arrow[d, "\widehat{\Phi}"] \\
\begin{bmatrix} f_1(v_1, v_s) \\ f_2(v_s, v_2) \end{bmatrix} \arrow[r, "\mathbb{L}_1"] & f_1(v_1, v_s)
\end{tikzcd}$$

(a) Embedding of $\widehat{E}$ in E.

$$\begin{tikzcd}
(v_1, v_s, v_2) \arrow[r, "\mathbb{K}_2"'] \arrow[d, "\Phi"] & (v_s, v_2) \arrow[d, "\widetilde{\Phi}"] \\
\begin{bmatrix} f_1(v_1, v_s) \\ f_2(v_s, v_2) \end{bmatrix} \arrow[r, "\mathbb{L}_2"] & f_2(v_s, v_2)
\end{tikzcd}$$

(b) Embedding of $\widetilde{E}$ in E.

Figure 3.1. *Let $E = (\Phi, \emptyset, u_0)$ denote the extended continuation problem with the combined zero problem $\Phi\big((v_1, v_s, v_2)\big) = \big(f_1(v_1, v_s), f_2(v_s, v_2)\big)^T$ and the initial solution guess $u_0 = (v_{1,0}, v_{s,0}, v_{2,0})$. Furthermore, let $\widehat{E} = (\widehat{\Phi}, \emptyset, \widehat{u}_0)$, $\widehat{\Phi}\big((v_1, v_s)\big) = f_1(v_1, v_s)$, $\widehat{u}_0 = (v_{1,0}, v_{s,0})$, and $\widetilde{E} = (\widetilde{\Phi}, \emptyset, \widetilde{u}_0)$, $\widetilde{\Phi}\big((v_s, v_2)\big) = f_2(v_s, v_2)$, $\widetilde{u}_0 = (v_{s,0}, v_{2,0})$, denote extended continuation problems constructed from the respective individual zero problems together with the index sets $\mathbb{K}_{1/2}$ and $\mathbb{L}_{1/2}$ as defined in Example 3.1. Then, the commutative diagrams (a) and (b) imply the embeddings $\widehat{E} \subseteq E$ and $\widetilde{E} \subseteq E$. The embedding $\widehat{E} \subseteq E$ is canonical.*

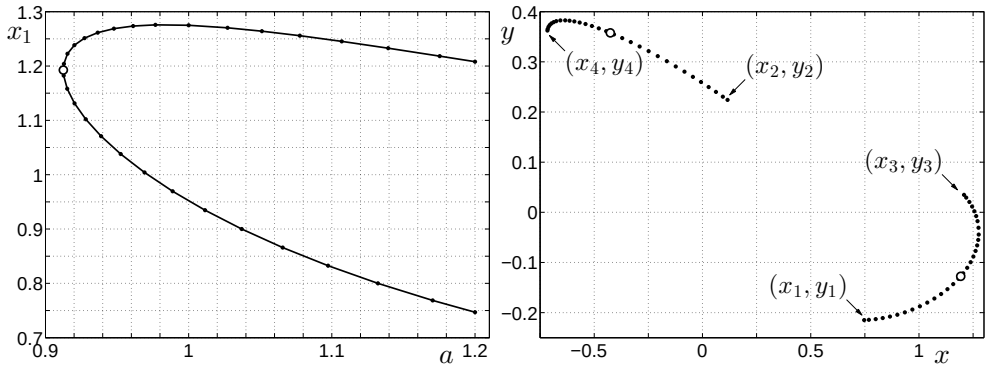

(a) Bifurcation diagram. (b) Period-4 points.

Figure 3.2. *A family of period-4 orbits of the Hénon map obtained in Example 3.13. The family passes through a period-halving point marked with an open circle, where the period-4 family intersects a period-2 family. A typical bifurcation diagram is shown in (a). Here, we graph a projection of the solution manifold onto the hyperplane with coordinates given by the parameter a and the x-component of the first point on the orbit. The full state-space orbits are shown in (b).*

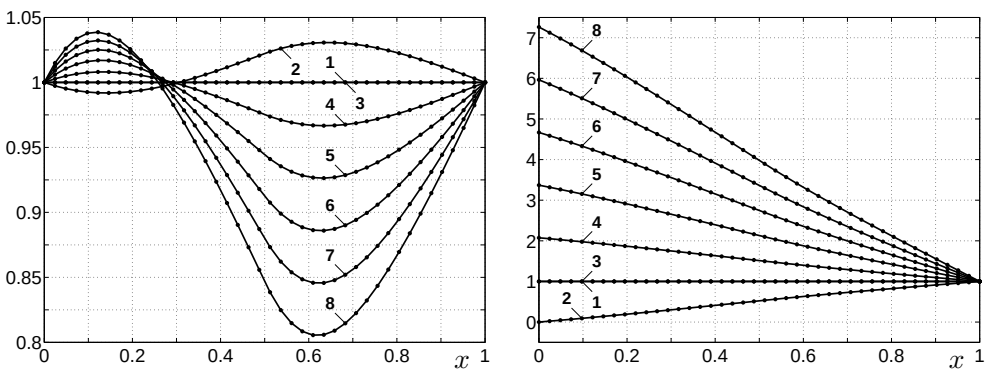

(a) Graphs of $f(x)$. (b) Graphs of $g(x)$.

Figure 3.3. *A set of profiles representative of the family of solutions to the Brusselator boundary-value problem in Example 3.14. The boundary conditions are $f(0) = f(1) = g(1) = 1$ and $g(0) = g_0$, where g_0 is the continuation parameter. We start with the exact solution $f(x) \equiv g(x) \equiv 1$ for $g_0 = 1$ and obtain solutions with nontrivial shape as the parameter g_0 varies. The labels correspond to the session output included in the text.*

5 Task Embedding

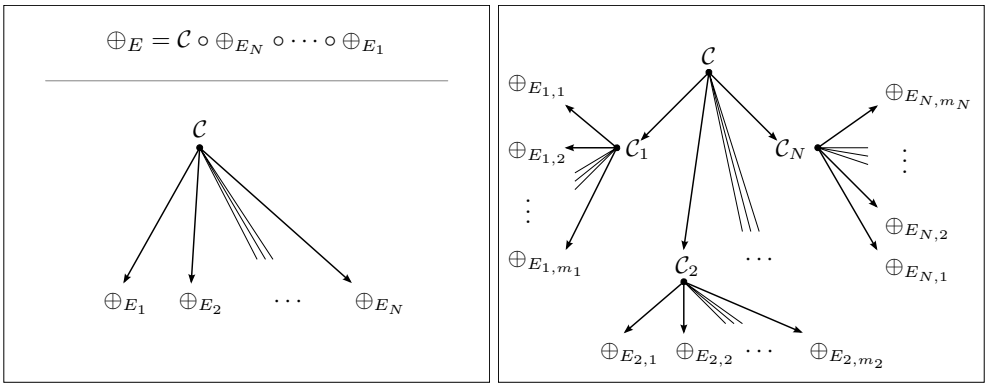

(a) Simple tree. (b) Recursive tree.

Figure 5.1. *The construction operator $\oplus_E$ appends E to an extended continuation problem. Following the decomposition in Eq. (5.7), it may be represented as a rooted tree with root at the closer $\mathcal{C}$, as illustrated in (a). Since the operation $\oplus$ is associative, a valid sequence of application is generated by a depth-first, bottom-up node traversal starting at $\mathcal{C}$. This concept can be extended to recursive rooted trees (b), where again a valid sequence of application of the operators may be obtained by performing a depth search starting at the root.*

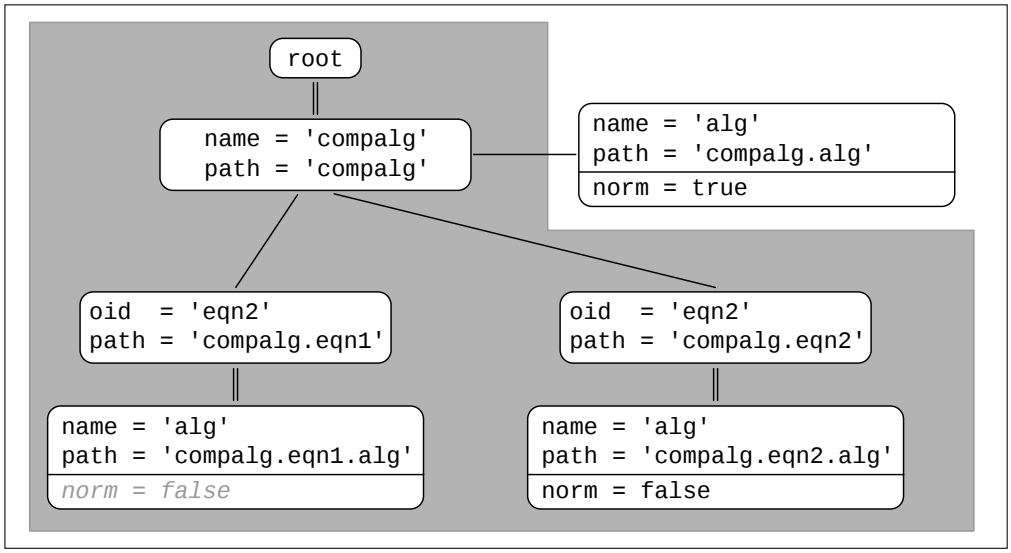

Figure 5.2. *The settings tree utility implemented in* COCO *facilitates the definition of toolbox options for individual or groups of toolbox instances. In its basic form it mimics the tree structure generated during the construction of an extended continuation problem. For example, the subtree inside the gray shaded area is a representation of the toolbox tree obtained in Example 5.6 if one interprets boxes connected by a double line as a single node. As one can see, each leaf has a unique path identifier, which allows the definition of toolbox options for every single instance individually. In the example, the* `'norm'` *property of the toolbox instance with path* `'compalg.eqn2.alg'` *has been set to* `false` *explicitly, while no such assignment has been made for the toolbox with path* `'compalg.eqn1.alg'`*. In this case, the toolbox default is used. In a slight deviation from the toolbox tree concept, the settings tree supports the addition of leafs that do not correspond to an actual object instance, but rather define options for groups of toolboxes. In the example, the assignment of a value to the* `'norm'` *property for the path* `'compalg.alg'` *adds the leaf to the settings tree that is outside the shaded area. This leaf defines options for all instances of a toolbox with name* `'alg'` *that reside on a lower level in the hierarchy than the node with path identifier* `'compalg'`*. Note that explicit assignments on lower levels have precedence over assignments on higher levels. In the example under consideration, extracting the property* `'norm'` *from* `'compalg.eqn1.alg'` *will, therefore, result in* `true`*, while the same operation for* `'compalg.eqn2.alg'` *will result in* `false`*.*

6 Discretization

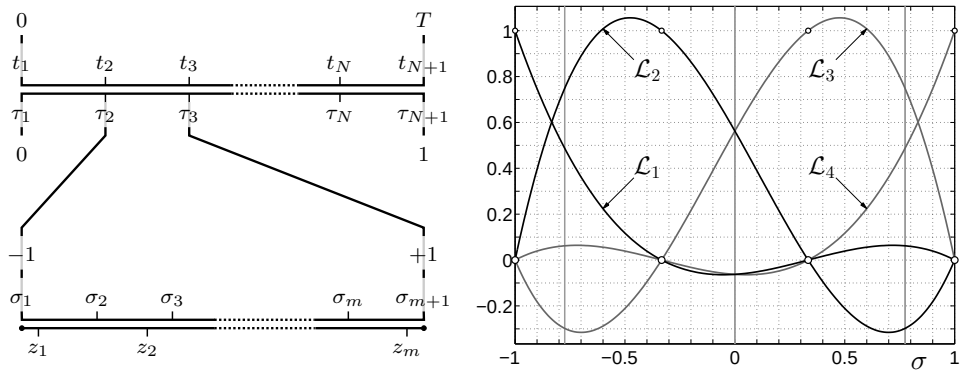

(a) Rescalings and reference interval. (b) Lagrange base polynomials.

Figure 6.1. *To construct a collocation zero problem, we first rescale the variable* $t \in [0, T]$ *to* $\tau \in [0, 1]$. *In a second step, we construct a uniform partition* $0 = \tau_1 < \tau_2 < \cdots < \tau_{N+1} = 1$ *of size* N *of the unit interval. In a third step, we map each subinterval* $[\tau_j, \tau_{j+1}]$ *onto the reference interval* $\sigma \in [-1, 1]$. *In a fourth step, we introduce a partition* $-1 = \sigma_1 < \sigma_2 < \cdots < \sigma_{m+1} = 1$ *of size* m *of the reference interval and represent the sought-after solution using Lagrange polynomial interpolation. In a last step, we require that the interpolating polynomial satisfy the ordinary differential equation* $dy/dt = f(y, p)$ *at* m *collocation nodes* z_l, $l = 1, \ldots, m$. *Panel (a) shows the hierarchy of scalings and partitions and panel (b) the family of Lagrange polynomials of degree 3 on the reference interval* $[-1, 1]$. *Values at the base points are shown as circles and the location of the three corresponding Gauss collocation nodes are indicated by vertical gray lines.*

7 The Collocation Continuation Problem

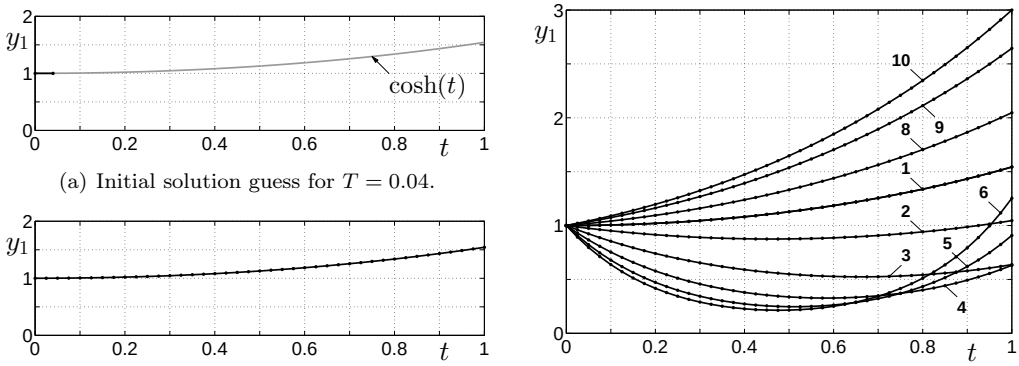

(a) Initial solution guess for $T = 0.04$.

(b) Solution for $T = 1$.

(c) Family of solutions obtained in continuation run with respect to parameter Y.

Figure 7.1. *A technique for constructing initial solutions to nonlinear boundary-value problems is continuation in the interval length T. Here, one starts with a short trajectory segment for small $T \ll 1$ and continues in T until reaching the desired value. This process is called* growing an initial orbit *and is illustrated in Sect. 7.3.1 in the context of the problem from calculus of variations, introduced in Chap. 1. Panel (a) shows an initial guess obtained with an Euler step at $t = 0$ with step size $h = 0.04$ and compares this approximation with the exact solution. The grown solution is shown in panel (b). From this solution one can, finally, start a continuation in Y. The labeled solution curves in (c) correspond to the labels in the session output included in the text.*

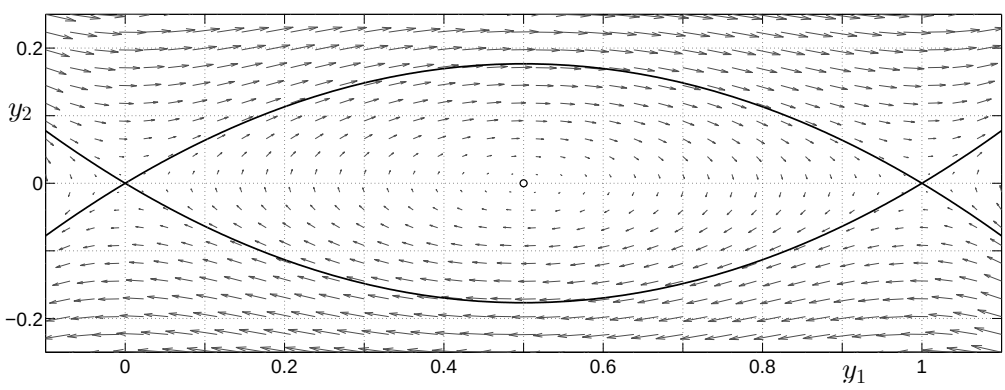

Figure 7.2. *Skeleton of the dynamics of the Huxley dynamical system given by the vector field in Eq. (7.26) and its direction field for $p_1 = 1/2$ and $p_2 = 0$. The system has three equilibria: a center at $(1/2, 0)$ and two saddles that are connected in a heteroclinic cycle. We demonstrate a continuation of the upper heteroclinic orbit in Sect. 7.3.2.*

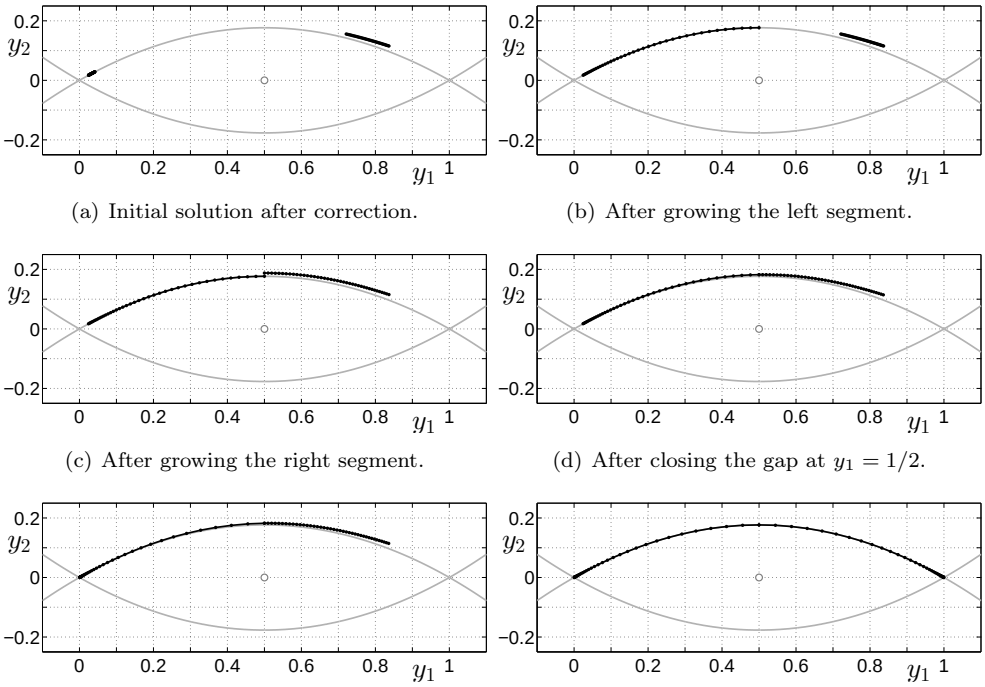

(a) Initial solution after correction.

(b) After growing the left segment.

(c) After growing the right segment.

(d) After closing the gap at $y_1 = 1/2$.

(e) After reducing ε_1.

(f) After reducing ε_2.

Figure 7.3. *In contrast to the example in Sect. 7.3.1, the construction of an initial solution for the continuation of a heteroclinic orbit in the Huxley dynamical system given by the vector field in Eq. (7.26) consists of a sequence of steps. We start with two short orbit segments close to each saddle, initial guesses to which can be obtained by Euler steps in the respective invariant eigenspace. The two segments after initial correction are shown in panel (a). In the subsequent step we grow both segments until they terminate on the hyperplane $y_1 = 0.5$, (b) and (c). At this point, the two end points are separated by a small gap, the so-called* Lin gap, *which we close in the next step, (d). In the last step, (e) and (f), we grow the two segments toward the saddle equilibria and obtain an approximation to a connecting orbit, which can be used as a starting point for subsequent continuation runs.*

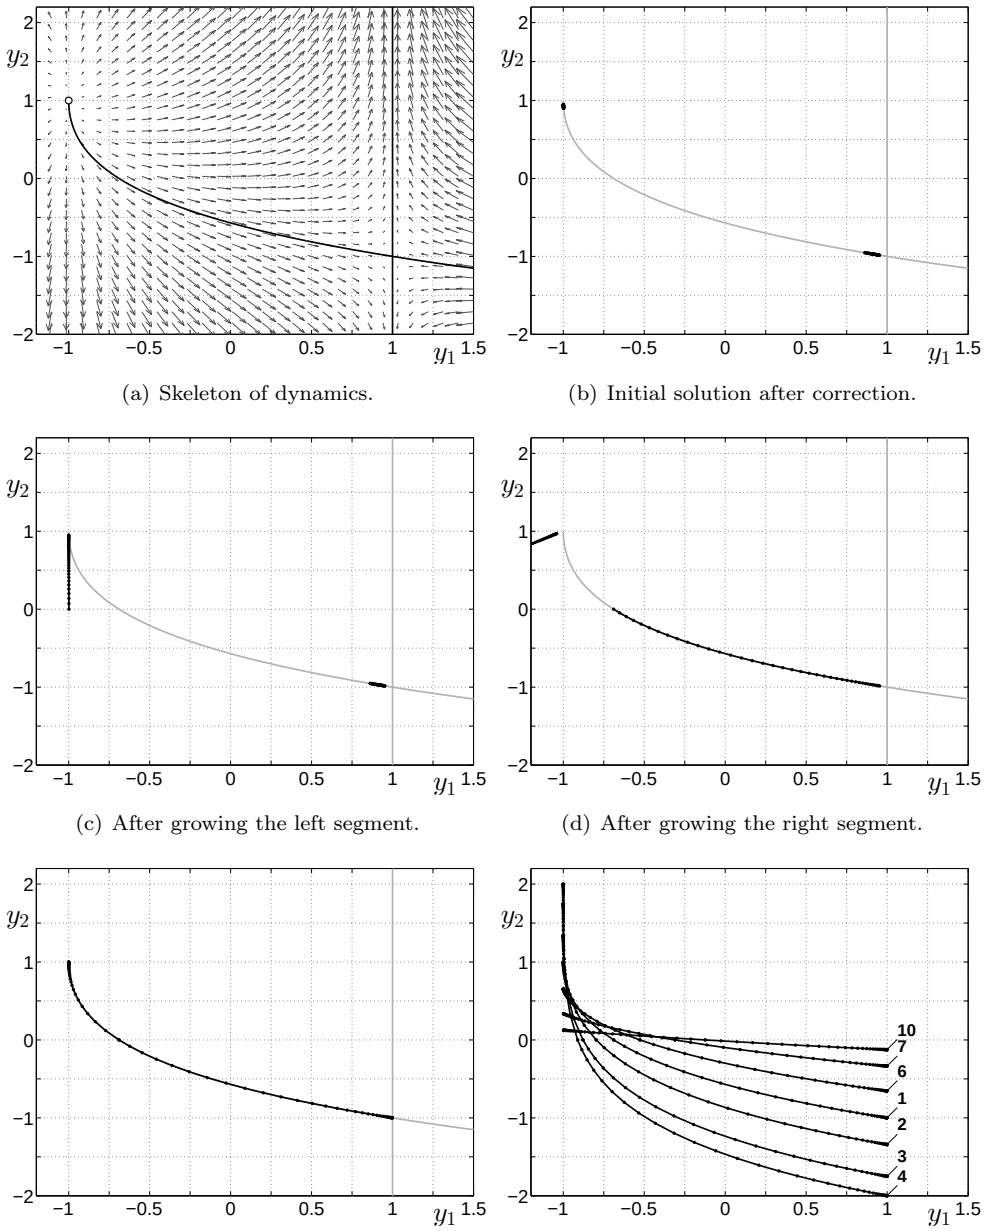

(a) Skeleton of dynamics.

(b) Initial solution after correction.

(c) After growing the left segment.

(d) After growing the right segment.

(e) After closing the gap at $y_2 = 0$ and reducing $\varepsilon_{1/2}$.

(f) Family of heteroclinic orbits.

Figure 7.4. *Computation of an initial approximation of a heteroclinic connection in the dynamical system with vector field given in Eq. (7.39), following the methodology described in Sect. 7.3.2; cf. Fig. 7.3. Here, we combine two instances of the* `'coll'` *toolbox with two instances of the* `'alg'` *toolbox, which are used to solve for the equilibria and an invariant eigenspace. A family of connecting orbits is shown in panel (f). The labels correspond to the session output of the last run included in the text.*

8 Single-Segment Continuation Problems

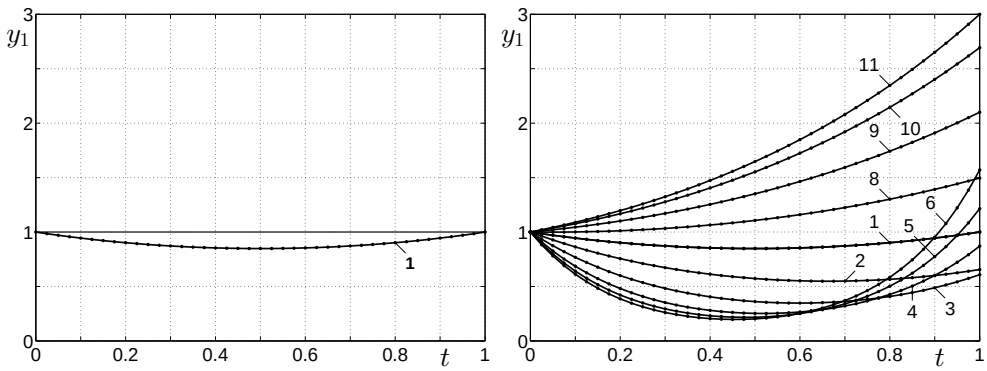

(a) Initial solution guess and first converged solution.

(b) Family of solutions.

Figure 8.1. *Solutions to the boundary-value problem given by the system of differential equations in Eqs. (8.6-8.7) and the given boundary conditions at $x = 0$ and $x = 1$. We start with the function $y_1 = 1$ shown in panel (a) as an initial guess. The initial correction converges to solution 1. A family of solutions obtained using numerical continuation is shown in panel (b). The labels correspond to the session output included in the text.*

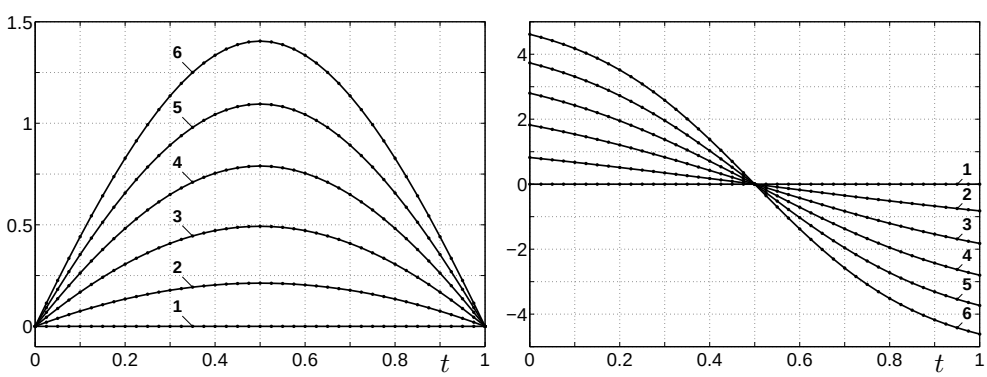

(a) Graphs of $z(t)$.

(b) Graphs of $z'(t)$.

Figure 8.2. *Solutions to the Bratu boundary-value problem given by Eq. (8.11) and the given boundary conditions at $t = 0$ and $t = 1$. We start with the exact solution $z(t) \equiv 0$ at $p = 0$ and obtain a family of solutions with increasing amplitudes. The labels correspond to the session output included in the text.*

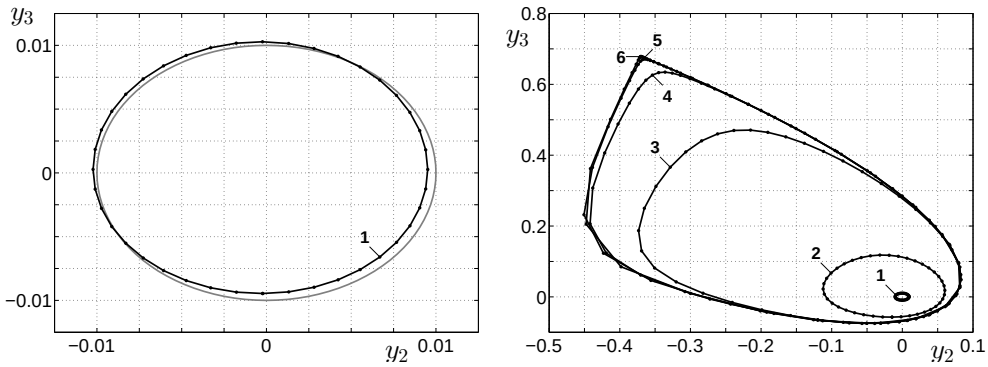

(a) Initial solutions. (b) Family of periodic orbits.

Figure 8.3. *We compute a family of periodic orbits emanating from a Hopf bifurcation point of the dynamical system given by the vector field in Eq. (8.24). We obtain an initial solution guess from normal form analysis, shown in gray in panel (a). The initial correction step converges to orbit 1. Panel (b) shows the family of periodic orbits of increasing amplitudes that seem to approach a homoclinic orbit, indicated by the corner that develops in the top left part of the plot and allocates many mesh points due to slow dynamics. The labels correspond to the session output included in the text.*

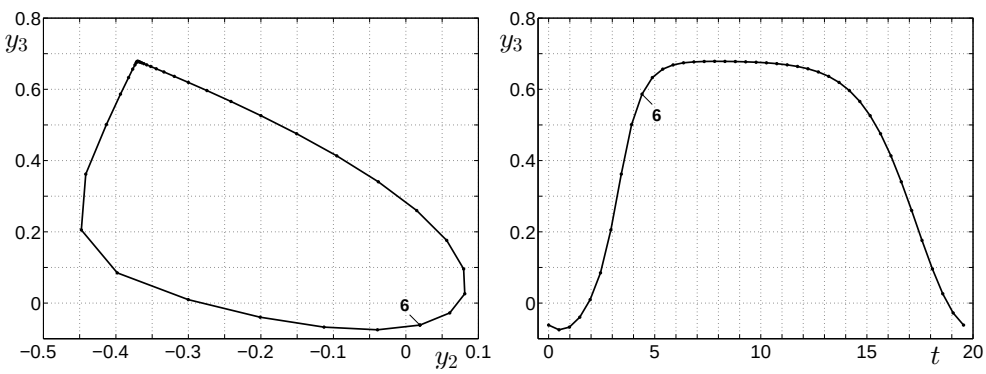

(a) Phase plot. (b) Time profile.

Figure 8.4. *The last periodic orbit 6 obtained in the continuation run shown in Fig. 8.3. The left panel shows the phase plot and the right panel the time profile. The time profile shows a phase of slow dynamics, indicating existence of a nearby equilibrium point, followed by a fast excursion. To confirm our hypothesis, we use this orbit and extend the part that seems to be close to an equilibrium, see Fig. 8.5.*

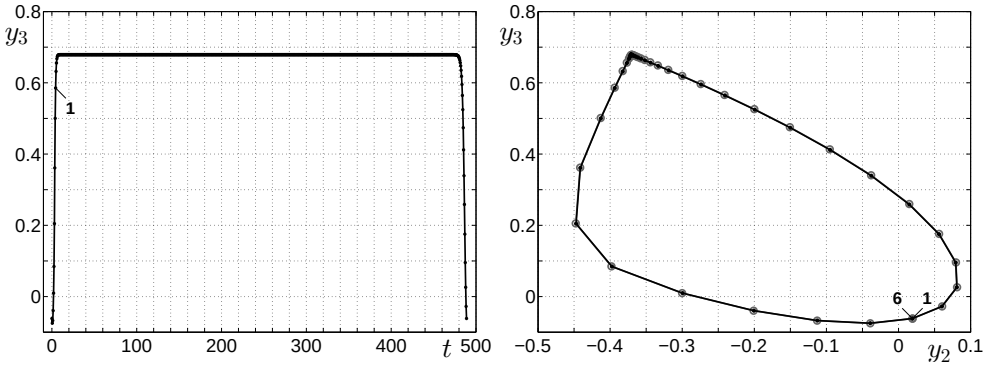

(a) Extended time profile. (b) Phase profiles.

Figure 8.5. *Starting with orbit 6 from Fig. 8.4, we insert a long segment of constant dynamics and rescale the period such that the shape of the orbit in phase space should be unchanged if there exists a nearby homoclinic orbit. The extended time profile after the initial correction step is shown in panel (a). We clearly observe an elongated phase of near-constant dynamics. We overlay this new solution 1 (black dot) on top of the previous orbit 6 (gray circle) in panel (b). The phase plots, including the distribution of mesh points, are virtually identical, which supports the assumption that a nearby homoclinic orbit exists. We continue a family of high-period orbits in Fig. 8.6.*

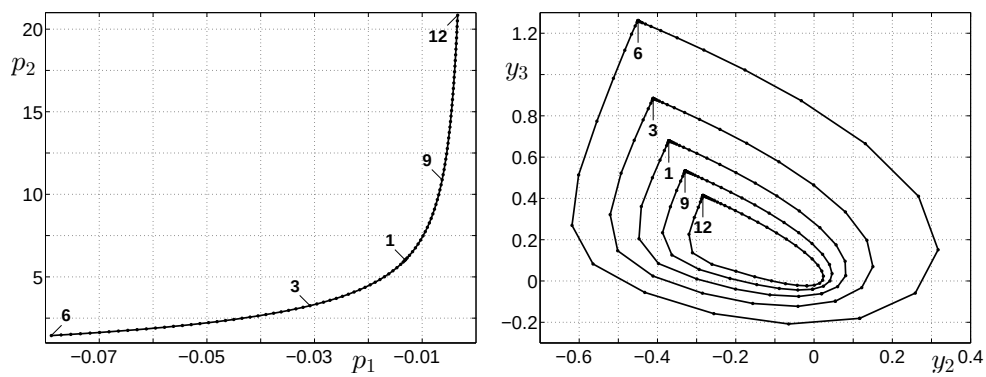

(a) Homoclinic bifurcation curve. (b) Homoclinic orbits approximated by high-period orbits.

Figure 8.6. *Continuation of the periodic orbit with high period, illustrated in Fig. 8.5, while keeping the period constant, resulting in an approximation to a* homoclinic bifurcation curve *(a). Each point on this curve corresponds to a terminal point along a family of periodic orbits emanating from a Hopf bifurcation under variations in p_2. Panel (b) shows selected members of the family of high-period orbits. The labels correspond to the session output included in the text.*

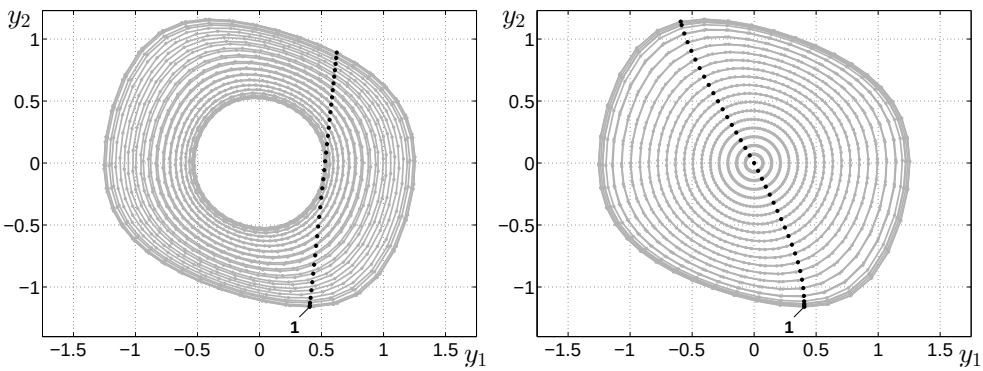

(a) Fixed phase condition.

(b) Phase condition with updates.

Figure 8.7. *Continuation of periodic orbits of the Lienard system in Example 8.5. Both continuation runs start at the same initial solution marked with label 1. The initial point $y(0)$ on each orbit is emphasized. With a fixed phase condition, all initial points must lie on a straight line (a), and we are unable to compute the full family due to a tangency. Updating the phase condition in each continuation step results in the full family; the initial points now lie on a curve passing through the Hopf bifurcation point at the origin within numerical accuracy (b).*

9 Multisegment Continuation Problems

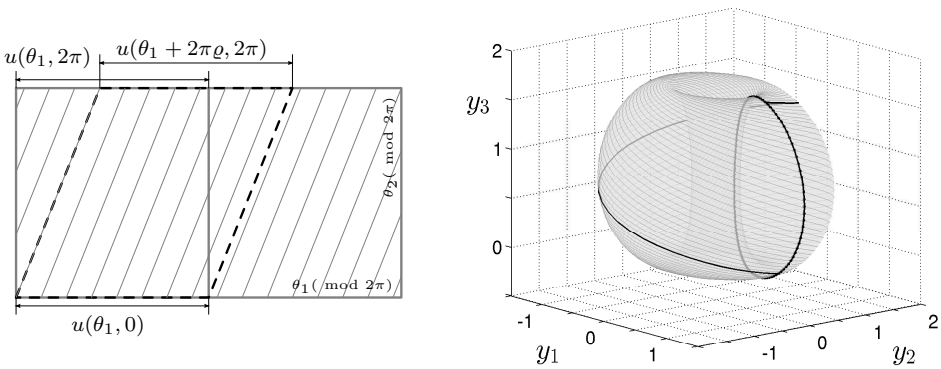

(a) Fundamental domain and characteristic field. (b) Image of characteristics in phase space.

Figure 9.1. *Two patches of a lift onto $\mathbb{R}^2$ of the fundamental domain of the transport equation in Eq. (9.13) next to each other with a sketch of the characteristic field (a). Applying the method of characteristics with initial condition $u(\theta_1, 0)$ maps the closed curve $u(\cdot, 0)$ onto the closed curve $u(\cdot, 2\pi)$, whereby the parametrization experiences a rotation by $2\pi\varrho$. In state space, this translates to a flow on an invariant torus for which there exists a curve $y(\varphi) = \upsilon(\varphi, 0) := u(2\pi\varphi, 0)$, for $\varphi \in [0, 1]$, that is mapped onto itself under the original flow after time $T = 2\pi/\omega_2$ and under a rotation of the parametrization by ϱ, as illustrated in panel (b).*

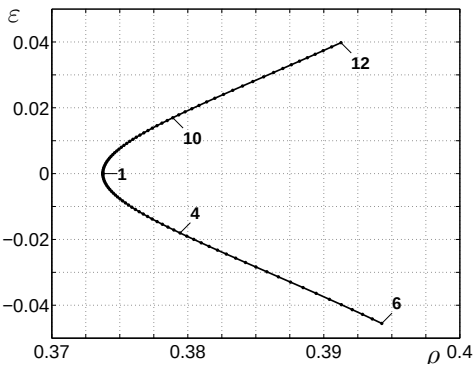

(a) Quasi-periodic arc.

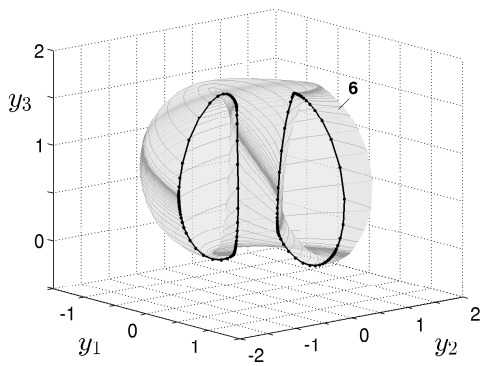

(b) Torus 6.

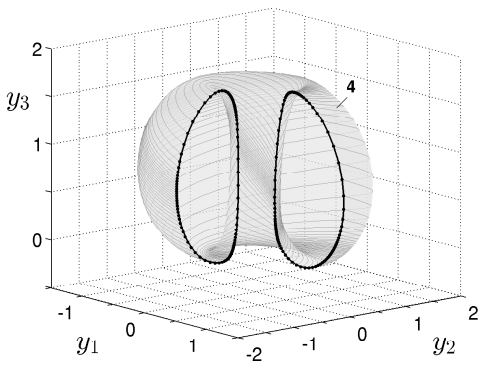

(c) Torus 4.

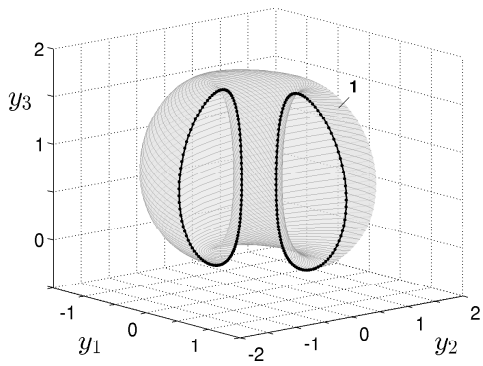

(d) Torus 1.

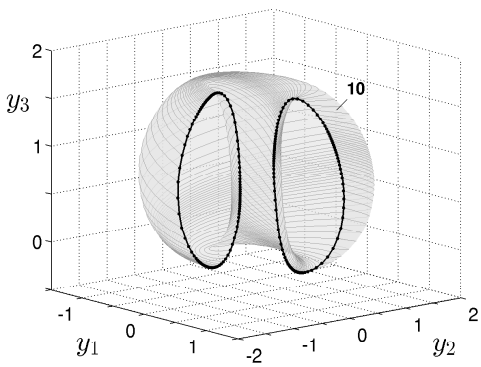

(e) Torus 10.

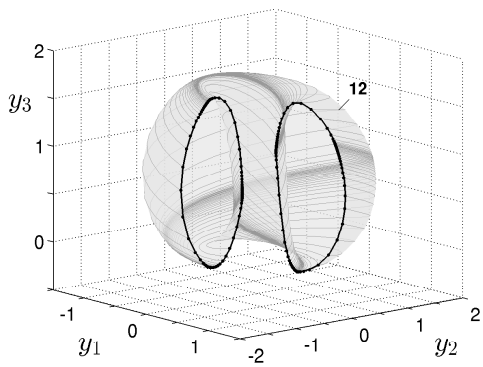

(f) Torus 12.

Figure 9.2. *A continuation of quasi-periodic invariant tori of the Langford system* (9.4) *results in the curve shown in panel (a), which is referred to as a* quasi-periodic arc *or a* quasi-periodic hair*. Selected members of this family are shown in panels (b) to (f). The accumulation of orbits on the torus suggests that this family approaches the vicinity of a 1:3 resonance at both ends of the arc. The labels correspond to the session output included in the text.*

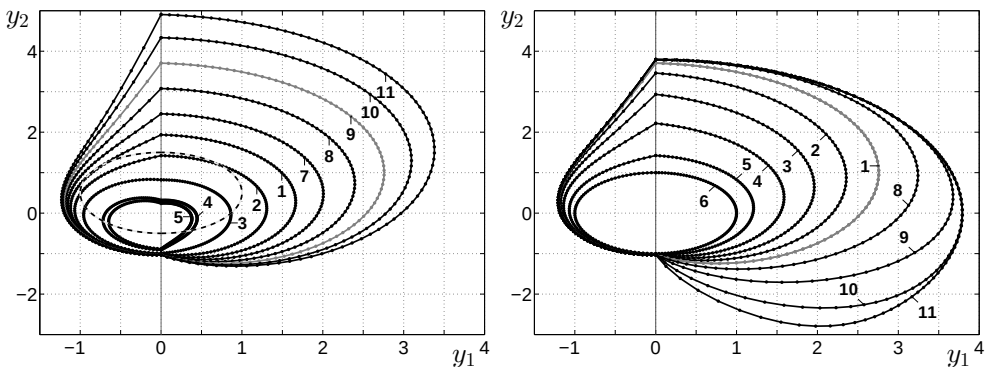

(a) Continuation in β.

(b) Continuation in α.

Figure 9.3. *Families of multisegment periodic orbits from Example 9.2. The initial guess is shown as the dashed closed curve in panel (a). The initial correction step converges to orbit 1. Selected members of the family of orbits resulting from a continuation in the parameter β are shown in (a). We restart a continuation in α from orbit 9 highlighted in gray. Some members of this family are shown in (b). The gray orbit is identical in both panels. The labels correspond to the session output included in the text.*

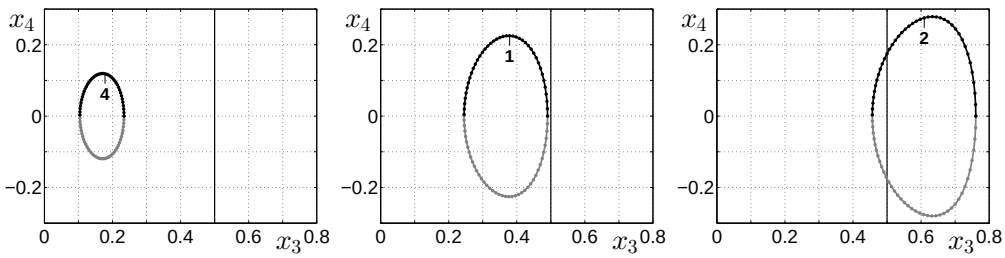

(a) Far from impact. (b) Close to impact. (c) Nonphysical orbit.

Figure 9.4. *Selected members of a family of nonimpacting two-segment periodic orbits of the stick-slip oscillator considered in Example 9.3 in the coordinates defined on page 237. Orbit 1 is close to impact and orbit 2 is nonphysical, because it penetrates through a surface of impact. Consequently, between 1 and 2 there must exist a* grazing orbit*; see Fig. 9.5. The labels correspond to the session output included in the text.*

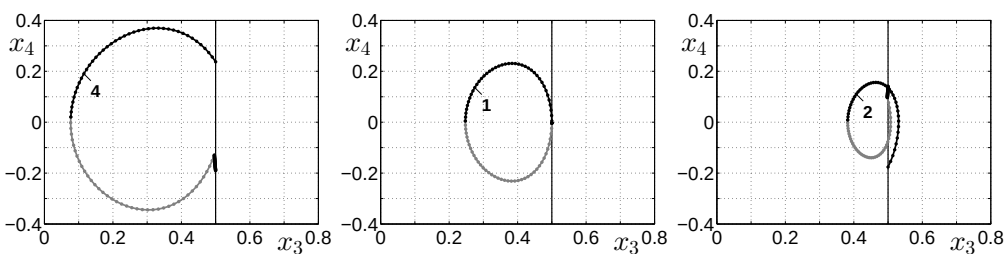

(a) Impacting orbit. (b) Grazing orbit. (c) Nonphysical orbit.

Figure 9.5. *Selected members of the family of impacting three-segment periodic orbits of the stick-slip oscillator considered in Example 9.4 in the coordinates defined on page 237. We initialize a computation of this family by resegmenting the grazing orbit from Example 9.3 and performing an initial correction under the additional constraint that the new segment has length 0; that is, it corresponds to the point of grazing contact. We again observe nonphysical orbits penetrating through the impact surface. The labels correspond to the session output included in the text.*

10 The Variational Collocation Problem

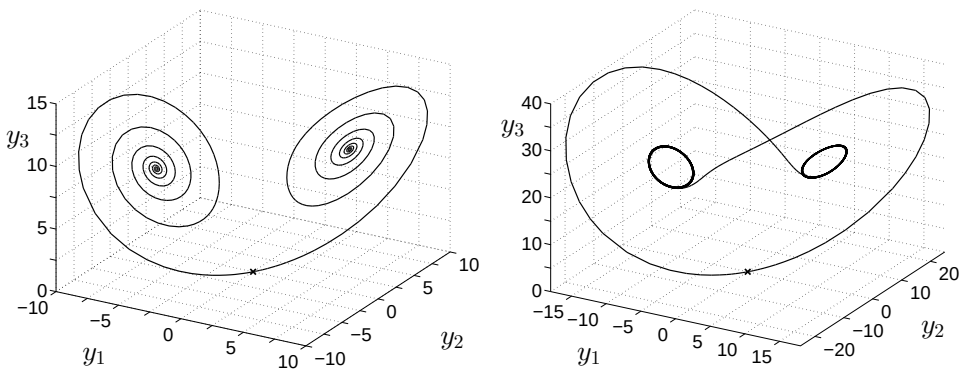

(a) Manifold for $r = 10$. (b) Manifold for $r = 24$.

Figure 10.1. *Orbits of the Lorentz system given by the vector field in Eq. (10.87) starting in the unstable eigenspace of the equilibrium at* 0*, tracing the unstable manifold. The orbits in (a) approach equilibria, while the orbits in (b) seem to approach periodic orbits. This approach is shown in more detail in Fig. 10.2.*

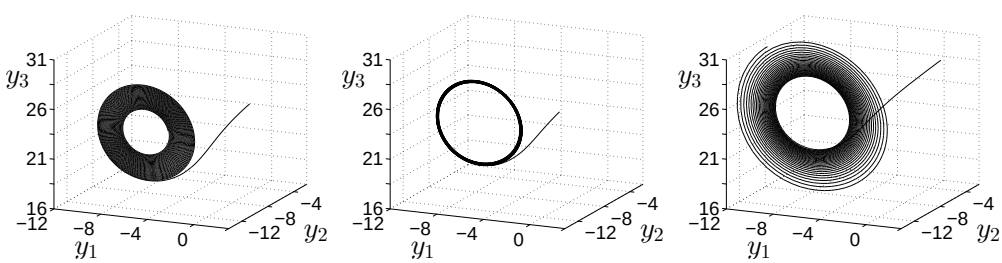

(a) Approach for $r = 23$. (b) Approach for $r = 24$. (c) Approach for $r = 25$.

Figure 10.2. *A close-up of the potential approach and exit of $\mathcal{W}_0^u$ to and from a periodic orbit, as observed in Fig. 10.1. In panel (a), the manifold exits spiraling inward, while the exit is outward in (c). This corresponds to a switch of approach from an orbit inside a stable manifold of a periodic orbit to an approach from an orbit outside. This suggests that, in between these parameter values, there exists an orbit approaching* on *the stable manifold—a heteroclinic connection between the equilibrium at* 0 *and a periodic orbit of saddle type.*

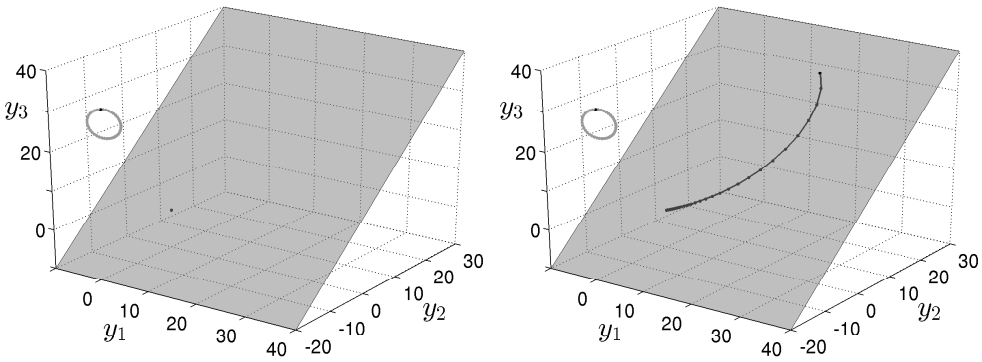

(a) Initial three-segment solution. (b) After growing the unstable manifold.

Figure 10.3. *Construction of an initial approximation of an orbit connecting an equilibrium and a periodic orbit in the Lorentz system given by the vector field in Eq. (10.87) following the homotopy approach described in Sect. 10.2.2. A state-space representation of the three-segment solution, consisting of a periodic orbit (gray) and two zero-length segments (black dots), that is used to initialize Stage II of the homotopy is shown in panel (a), together with the hyperplane Σ that separates the periodic orbit from the equilibrium. In Stage II, we grow an orbit in $\mathcal{W}_0^u$ until it terminates on Σ, as shown in (b). In the subsequent Stage III we grow an orbit in $\mathcal{W}_{\text{per}}^s$ in a similar way; see Fig. 10.4.*

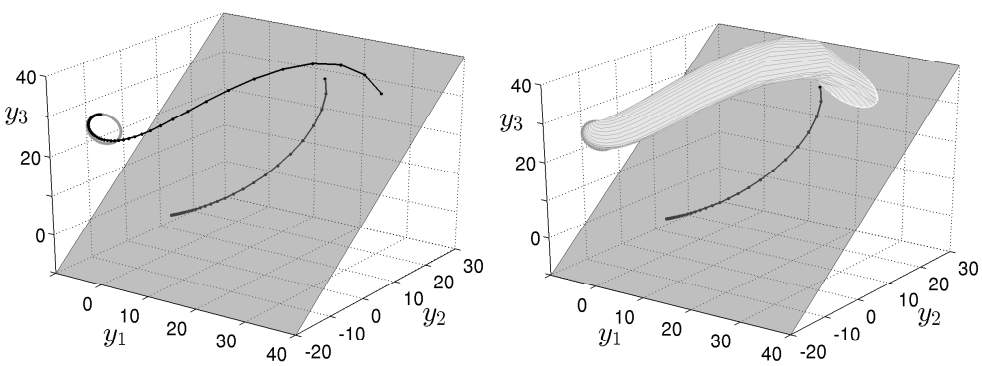

(a) Growing one orbit in the stable manifold. (b) After sweeping the stable manifold.

Figure 10.4. *Panel (a) shows the three-segment solution after completing Stage III, i.e., growing an orbit in $\mathcal{W}_{\text{per}}^s$ starting at the solution shown in Fig. 10.3(b). Here, the end point of the orbit segment in $\mathcal{W}_0^u$ and the starting point of the orbit segment in $\mathcal{W}_{\text{per}}^s$ both lie in Σ. Although $\mathcal{W}_{\text{per}}^s$ is 2-dimensional, the connecting orbit is unique. To obtain an initial approximation of the connecting orbit, we first sweep $\mathcal{W}_{\text{per}}^s$ in Stage IV and compute a set of orbit segments that cover the manifold sufficiently densely (b). From this family of orbits we select the one that terminates closest to the end-point of the segment in $\mathcal{W}_0^u$; see Fig. 10.5.*

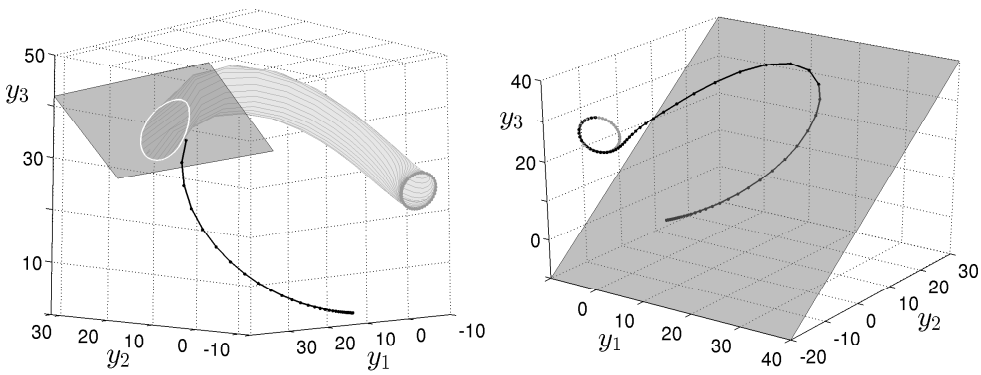

(a) A different view of the sweep.

(b) After closing the gap.

Figure 10.5. *Panel (a) shows a different view of Fig. 10.4(b), the result of a sweep of* $\mathcal{W}^s_{\text{per}}$*. The intersection with* Σ *is highlighted. We compute the point of the intersecting curve that is closest to the end point of the segment in* $\mathcal{W}^u_0$ *and initialize Stage V of the homotopy, i.e., the closing of the Lin gap. The resulting connecting orbit after closing the gap is shown in panel (b).*

11 Covering Manifolds

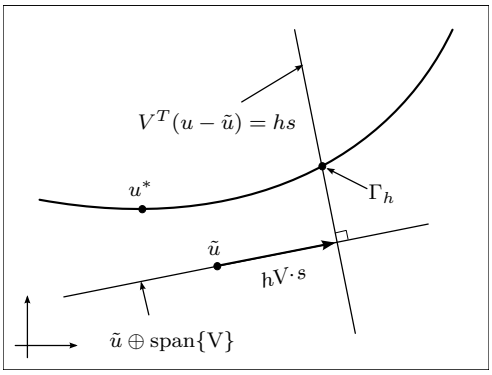

(a) General condition.

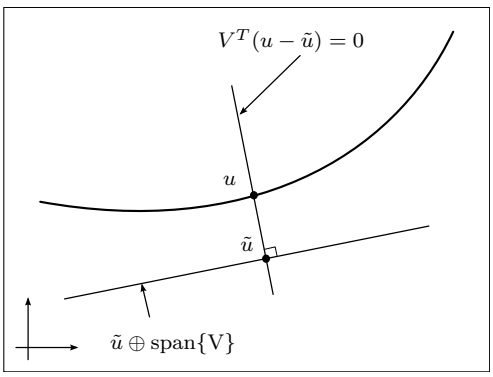

(b) Initial projection condition.

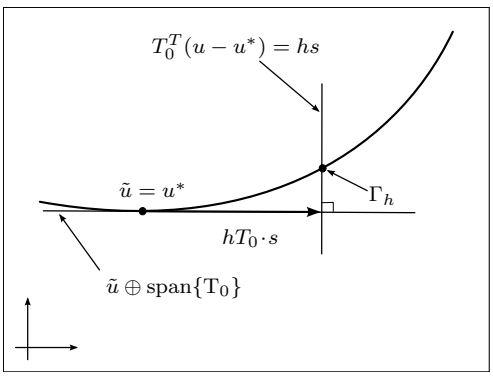

(c) Pseudo-arclength condition.

Figure 11.1. *Geometric interpretation of the projection condition in Eq. (11.46). The point Γ_h on the manifold projects orthogonally, with respect to the affine subspace $\tilde{u} \oplus \mathrm{span}\{V\}$, onto the point $\tilde{u} + hV \cdot s$, as illustrated in (a). The same condition for $hV \cdot s = 0$ reduces to Eq. (11.25) and allows for computing a first solution point on the manifold from an initial guess $\tilde{u}$ if a suitable V is available (b). For $\tilde{u} = u^*$ and $V = T_0$, we obtain the well-known pseudo-arclength condition (c).*

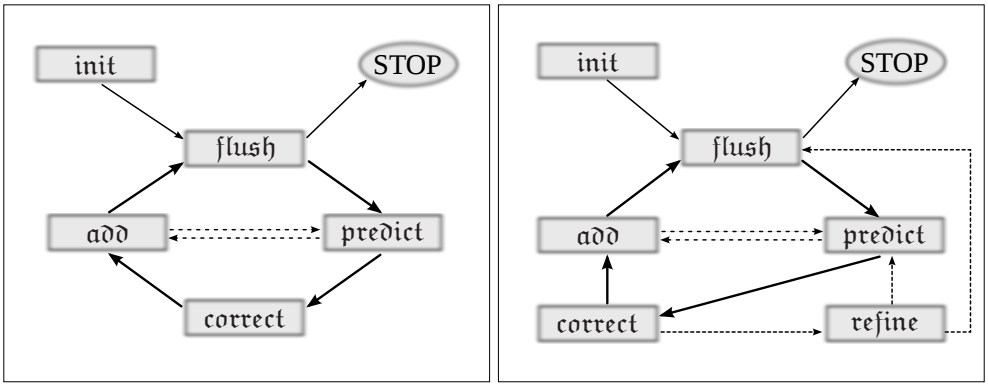

(a) FSM with basic loop. (b) FSM with additional refine state.

Figure 11.2. *In its simplest form an* atlas algorithm *may be implemented as an event-driven finite-state machine (FSM) consisting of five states (a), as explained in Sect. 11.2.1. The loop* flush, predict, correct, *and* add *is entered from* init *via* flush. *A terminal exit of the loop may be triggered in the* flush *state upon meeting some stop condition. The variant described in Sect. 11.2.3 and shown in (b) contains the additional state* refine *in order to allow for handling of a failure of convergence in* correct.

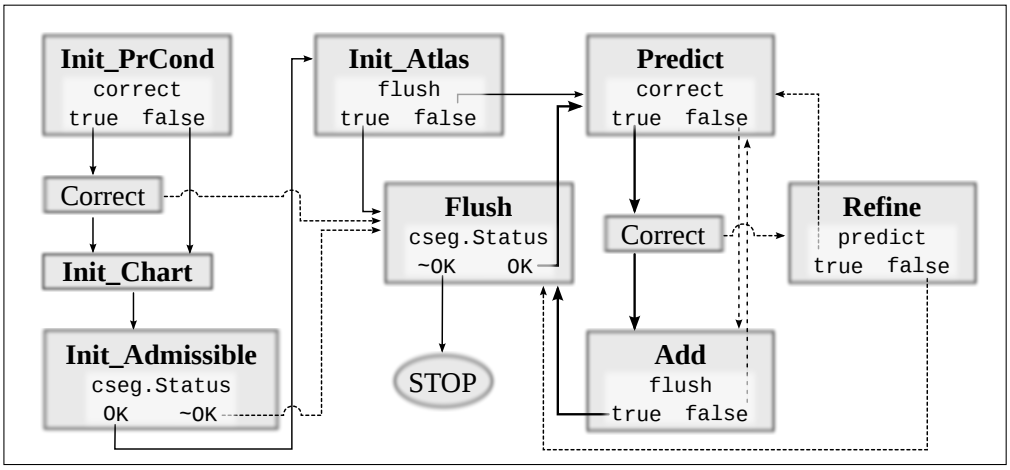

Figure 11.3. *The atlas algorithm implemented in* COCO *has a five-step initialization sequence described in Sect. 11.2.3 and is otherwise identical to the finite-state machine shown in Fig. 11.2(b). The names of states that interface with class methods of the* `AtlasBase` *class, introduced in Sect. 11.3, are typeset in bold face. This algorithm allows an* `AtlasBase` *subclass to alter the execution path via a boolean flag included in the class methods' return argument lists or by the content of* `cseg.Status` *in the case of the* `init_admissible` *and* `flush` *class methods, as indicated in the boxes.*

12 Single-Dimensional Atlas Algorithms

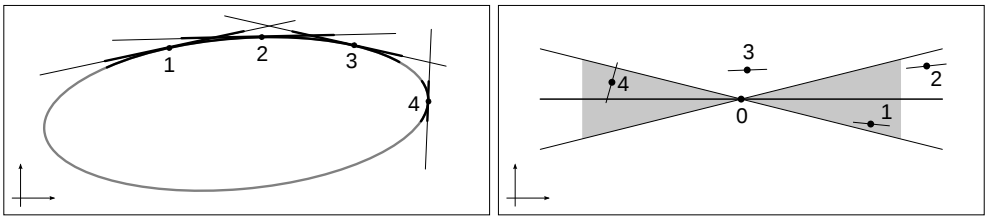

(a) Faithful covering and gaps. (b) Neighbors and nonneighbors.

Figure 12.1. *Every chart based at $\tilde{u}$ and with tangent matrix V defines an affine subspace $A := \tilde{u} \oplus \mathrm{span}\{V\}$ that approximates the manifold M locally. We define the local cover $C := \{v \in M \mid \mathrm{err}(v, A) < \mathrm{TOL}\}$, where* err *is some measure of the approximation error and* TOL *is a required accuracy. An important requirement of an atlas algorithm is that the intersection of local covers of neighboring charts be nonempty, i.e., that the respective local covers overlap. In panel (a) the local covers of charts 1-3 form a sequence of overlapping covers, while the local cover of chart 4 is disconnected. We say that the covering defined by charts 1-4 has a* gap. *A simple heuristic to prevent the occurrence of gaps is a test that ensures that the approximation error between neighboring charts is sufficiently small on a local cover with radius greater than $R/2$ such that overlapping local covers result. Panel (b) illustrates the conditions in Eqs. (12.22-12.26) that require that the base point of a neighboring chart lie within the shaded area, and that the angle between the tangent spaces be bounded. According to this condition, chart 1 is a neighbor of chart 0, while charts 2-4 are not.*

13 Multi-Dimensional Manifolds

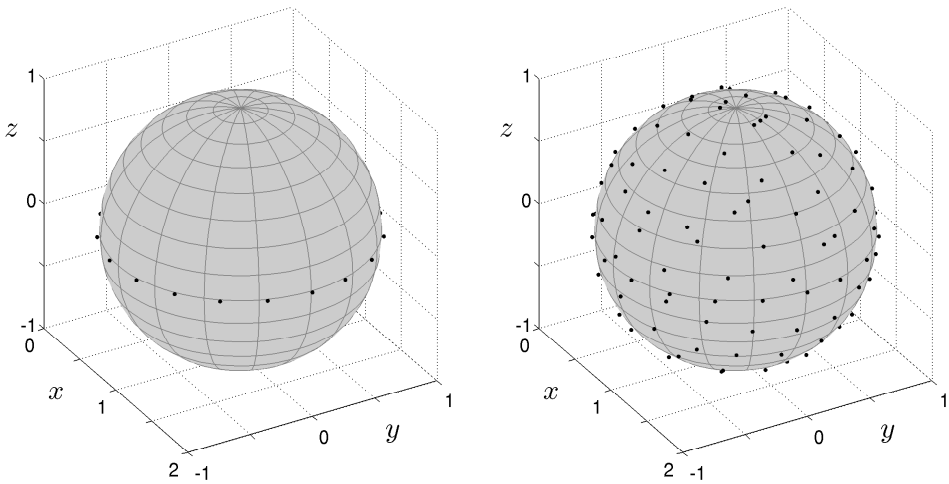

(a) Incomplete covering.

(b) Complete covering.

Figure 13.1. *Coverings of a unit sphere obtained in Example 13.1 with the modified* `atlas_2d_min` *atlas algorithm from Sect. 13.2. The first run results in an incomplete covering shown in panel (a): the algorithm terminates prematurely after covering a strip around the equator. The second run, with a different choice for the number of continuation directions for each chart, results in a complete cover (b). The dots mark the position of the base points of the charts in the final atlas.*

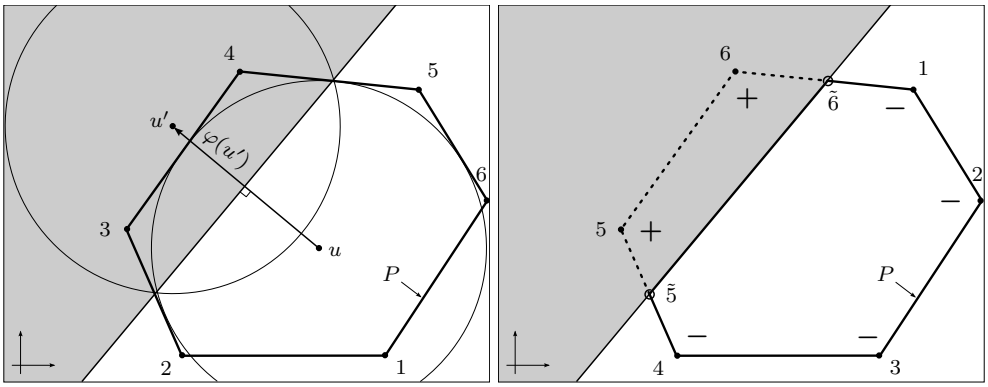

(a) Two overlapping charts. (b) Subtracting a half-space from P.

Figure 13.2. *Updating a polygon P associated with a chart $\{u, T, \Sigma, R\}$ during consolidation. The overlapping chart $\{u', T', \Sigma', R\}$ is projected onto the affine subspace $u \oplus \operatorname{span}\{T\}$, which is divided into two half-spaces along a straight line normal to the vector $\varphi(u')$ and separating u and u', as shown in panel (a). In a subsequent step, the polygon is updated to the intersection of P with the half-space containing u according to Eqs. (13.12-13.15), as illustrated in panel (b) after renumbering. Here, the sign of the expression $r^T \cdot \varphi(u') - \frac{1}{2}\|\varphi(u')\|^2$ is indicated next to the vertices.*

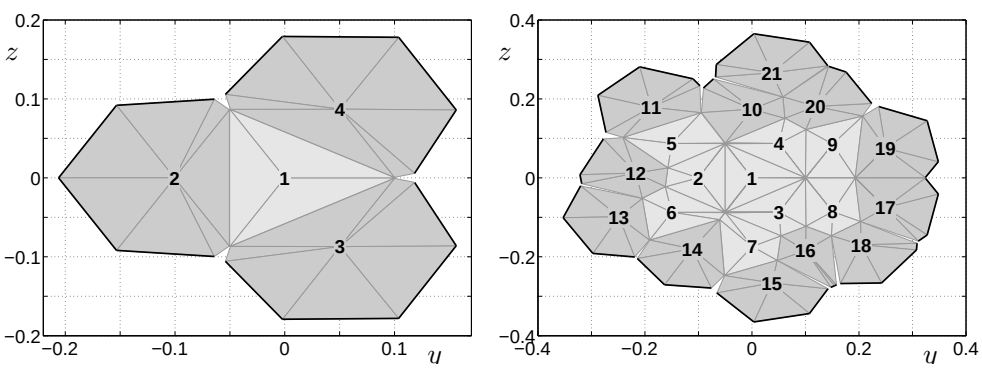

(a) After 4 steps. (b) After 21 steps.

Figure 13.3. *Initial steps of covering the manifold of Example 13.2 with the 2-dimensional atlas algorithm from Sect. 13.3. After four steps we obtain an atlas with one interior and three boundary charts (a), where each chart has all other charts as neighbors. The piecewise-linear approximation of the atlas boundary is the union of the bold line segments. The algorithm continues to grow the atlas and expands the boundary outwards (b).*

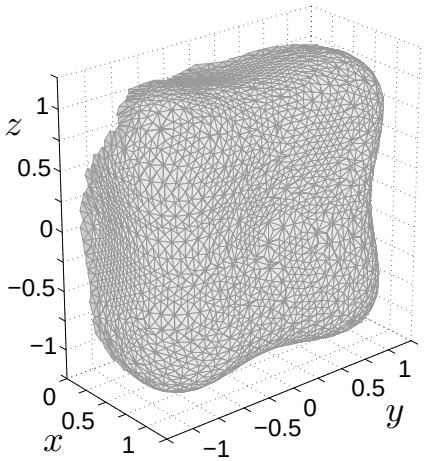

(a) Partial cover of pillow.

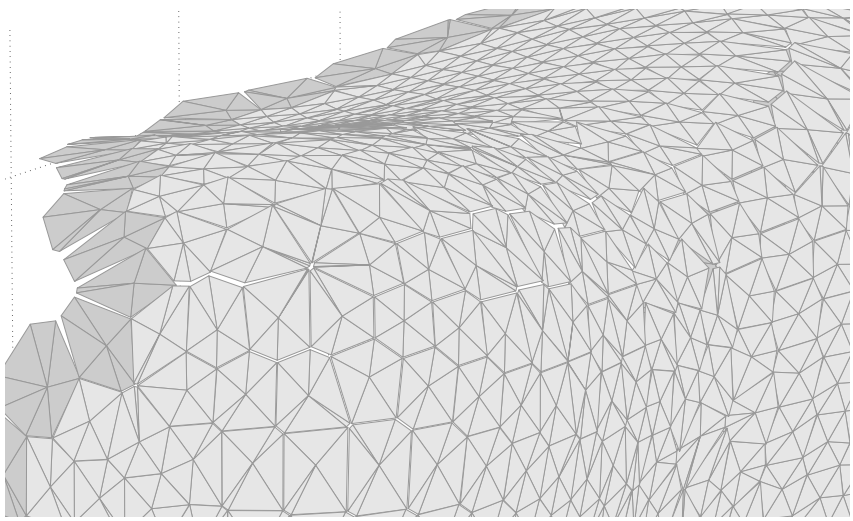

(b) Close up of top-left corner.

Figure 13.4. *After executing 1,000 steps of the algorithm from Sect. 13.3, we obtain the partial cover of the manifold of Example 13.2 shown in panel (a). The local geometry of this partial cover is visualized in the close-up of the top-left corner, shown in panel (b).*

14 Computational Domains

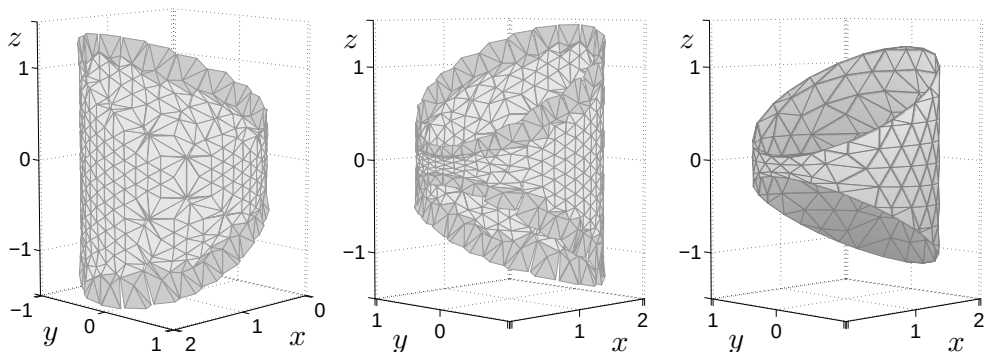

(a) Manifold cover. (b) Different view. (c) Triangulation of surface.

Figure 14.1. *Applying the modified 2-dimensional covering algorithm from Sect. 14.2 to the constrained cylinder in Example 14.4 results in the covering shown in different views in panels (a) and (b). A triangulation of the surface can be constructed using the neighbor information stored with each chart. A resulting triangulated surface is shown in panel (c) in the same view as in panel (b).*

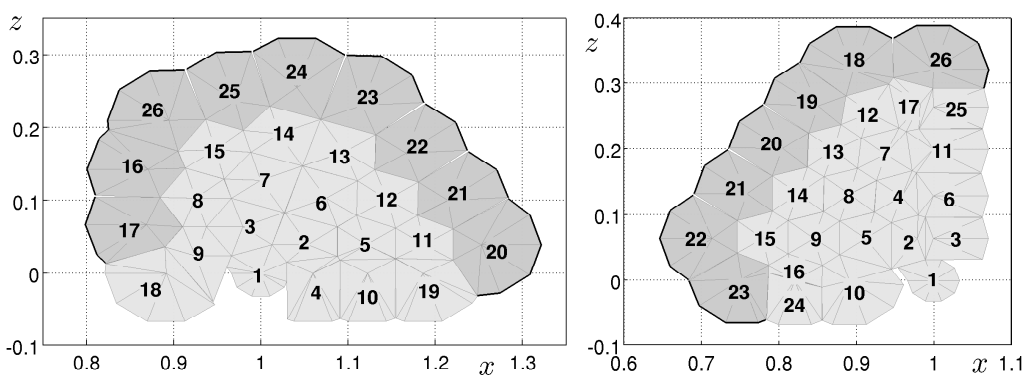

(a) One active constraint $z \geq 0$. (b) Two active constraints $z \geq 0$, $x \leq 1$.

Figure 14.2. *The 2-dimensional atlas algorithm discussed at the end of Sect. 14.2 supports starting on the boundary of the computational domain. This is illustrated with the manifold of Example 14.5, where we start on a boundary defined by one active constraint in panel (a), and at a corner defined by two active constraints in panel (b). In both cases the initial solution guess is given by $(x, y, z) = (1, -1, 0)$ and we show the atlas obtained after 25 continuation steps.*

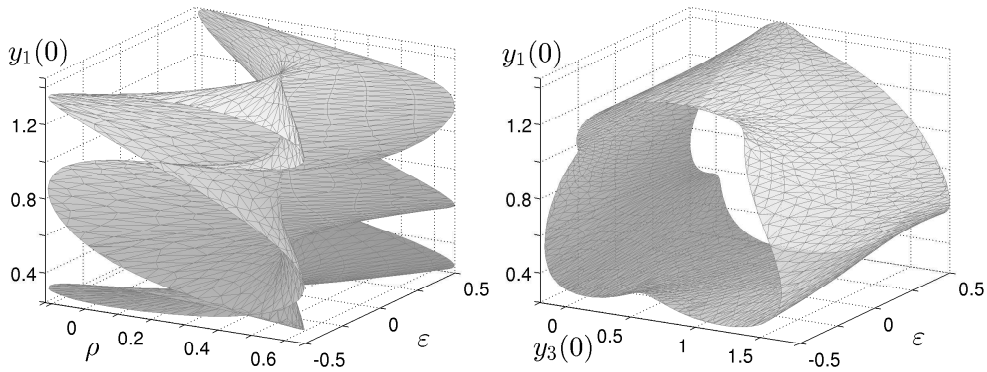

(a) 1:3 resonance surface in ρ-ε-$y_1(0)$- ... (b) ... and in $y_3(0)$-ε-$y_1(0)$-space.

Figure 14.3. *A so-called* resonance surface *obtained by applying the 2-dimensional covering algorithm from Sect. 14.2 to a boundary value problem for resonant periodic orbits of the Langford dynamical system, given by the vector field in Eq. (14.4). A projection of this surface onto the ρ-ε-$y_1(0)$-space is shown in panel (a). The striking cyclic triple-S shape is generic for 1:3-resonance surfaces. Similar shapes are observed for all m:n resonance surfaces. The apparent self-intersection is due to projection; the surface is locally isomorphic to a cylinder (b). A more familiar projection, a so-called* Arnol′d tongue, *is shown in Fig. 14.4(a).*

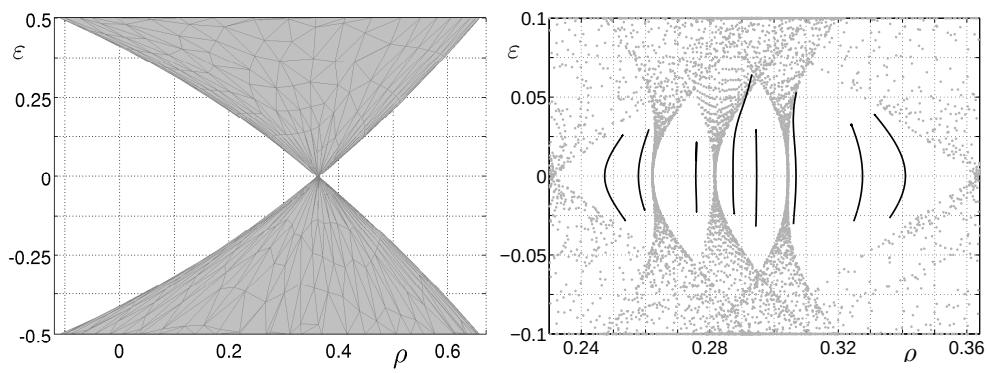

(a) 1:3 Arnol′d tongue. (b) Arnol′d tongue scenario.

Figure 14.4. *Projecting the resonance surface shown in Fig. 14.3 onto the ρ-ε plane results in an Arnol′d tongue (a). Combining the computation of resonance surfaces with the computation of quasi-periodic arcs, as shown in Sect. 9.2, enables the investigation of the so-called* Arnol′d tongue scenario. *This consists of a countable collection of Arnol′d tongues and a complementary Cantor-like set of quasi-periodic arcs. Panel (b) shows, from left to right, the 1:4, 3:11, 2:7, 3:10 and 1:3 tongues together with quasi-periodic arcs for rotation numbers obtained by a two-level recursive golden-mean subdivision of the interval* $[1/4, 1/3]$. *The tongues are represented by a dot for each of the 1500 resonant orbits computed on each resonance surface. We observe that one of the quasi-periodic arcs enters the 2:7 tongue. This violates theory and is an artifact caused by the absence of control for discretization errors, a task undertaken in Part V.*

15 Special Points and Events

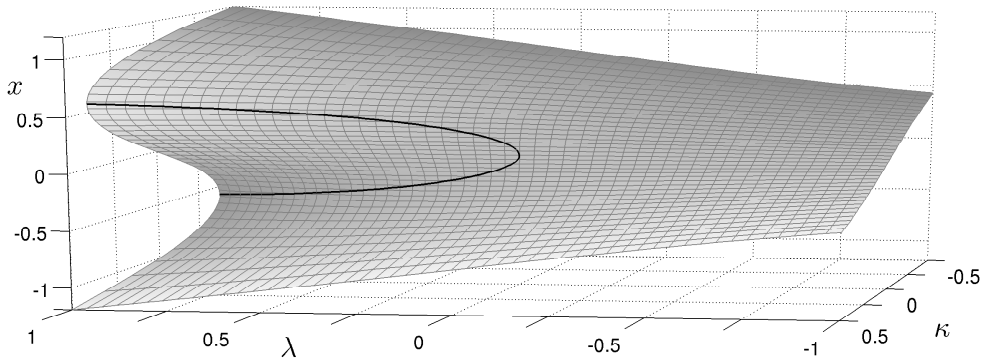

Figure 15.1. *Solution manifold of the zero problem* $\Phi(u) = 0$*, where* Φ *is given in Eq. (15.19). The solid curve marks the locus of vertical tangencies, implicitly defined by Eq. (15.22) and given by special points associated with the function* ψ *in Eq. (15.30) and the value* 0*.*

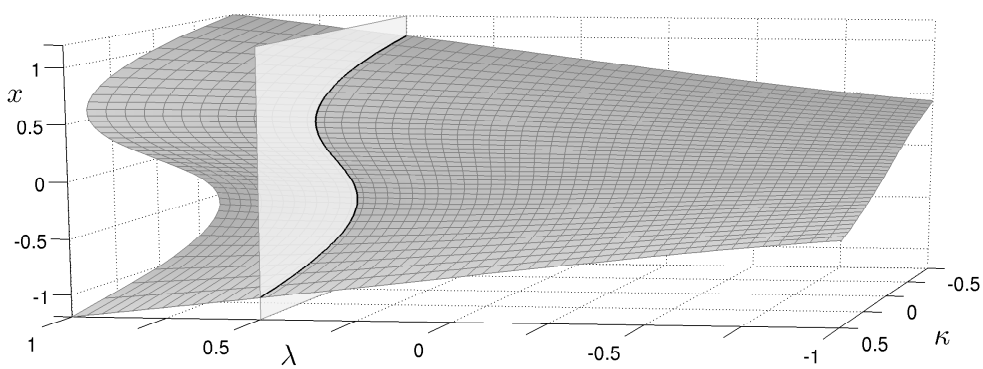

Figure 15.2. *Restricted continuation problem defined by the zero function* Φ *in Eq. (15.19) and the monitor functions* Ψ *in Eq. (15.39), where* $\mathbb{I} = \{2\}$*. The solution curve is the intersection of the solution manifold of* $\Phi(u) = 0$ *and the hyperplane* $\lambda = \text{const}$*. For* $\lambda = 0.5$*, we find two fold points along the highlighted curve.*

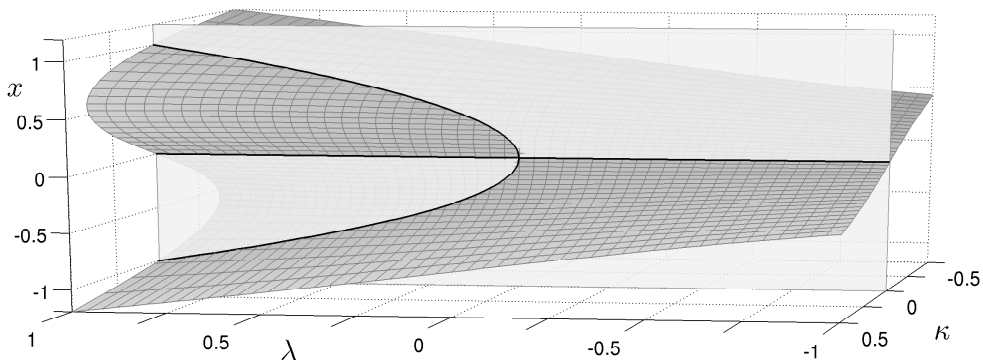

Figure 15.3. *Restricted continuation problem defined by the zero function* Φ *in Eq. (15.49) and the monitor functions* Ψ *in Eq. (15.50), where* $\mathbb{I} = \{1\}$*. The solution set is the intersection of the solution manifold of* $\Phi(u) = 0$ *and the hyperplane* $\kappa = \text{const.}$*. For* $\kappa = 0$*, this is the union of two curves that intersect at* $(0, 0, 0)$ *in a branch point.*

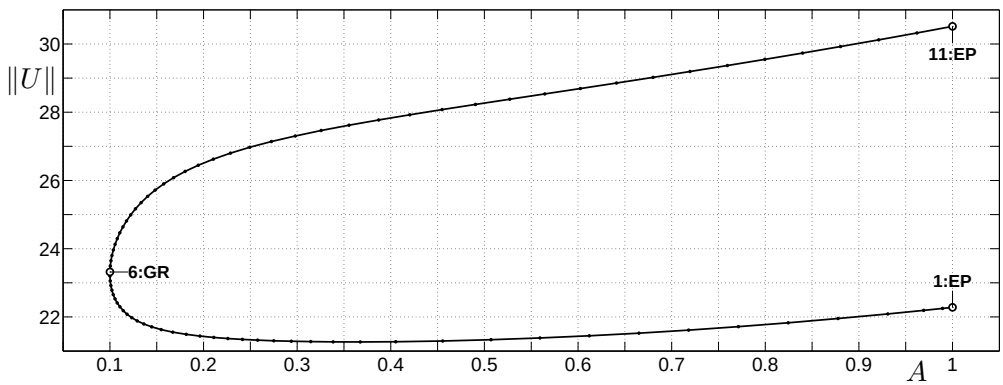

Figure 15.4. *Bifurcation diagram of the impact oscillator from Example 15.2 with* $\|U\| := \|(u, \mu_{\mathbb{J}})\|$. *Continuation starts with an impacting orbit at label 1 and passes through a grazing point. The part of the curve between labels 6 and 11 corresponds to nonphysical orbits penetrating through the impact surface. Some representative orbits are shown in Fig. 15.5.*

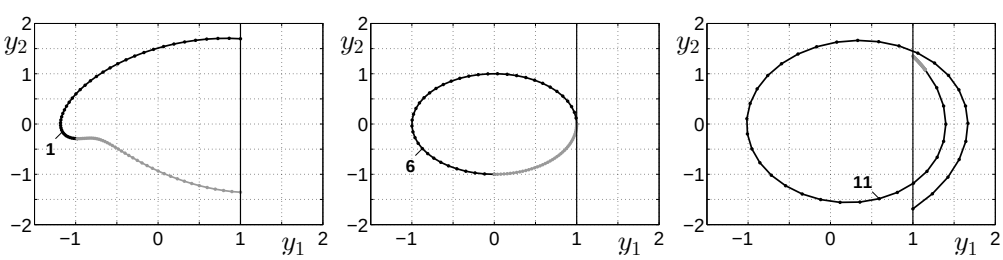

(a) Impacting orbit. (b) Grazing orbit. (c) Nonphysical orbit.

Figure 15.5. *Some orbits representative of the continuation illustrated in Fig. 15.4. Continuation starts with the impacting orbit shown in (a), passes through a grazing point (b) after which nonphysical orbits result (c).*

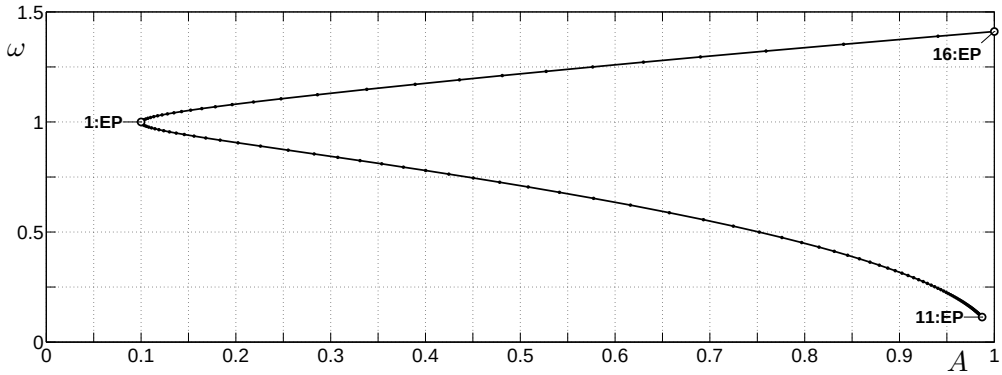

Figure 15.6. *Grazing curve in the (A, ω)-parameter plane computed for the impact oscillator in Example 15.2. For the system under consideration, this curve subdivides the parameter plane locally into a region of parameter values for which nonimpacting orbits exist (left-hand side) from parameter values for which orbits have at least one impact (right-hand side). Some representative grazing orbits are shown in Fig. 15.7.*

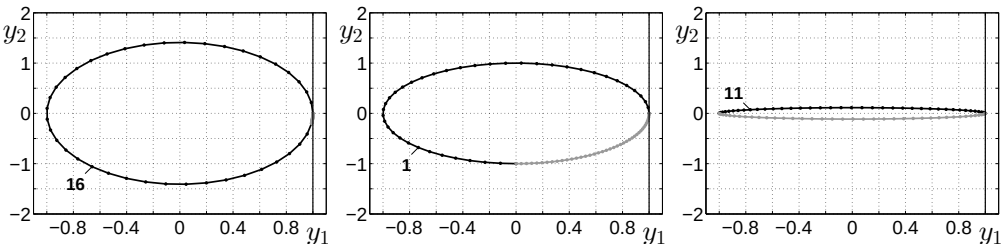

(a) Grazing orbit 16. (b) Grazing orbit 1. (c) Grazing orbit 11.

Figure 15.7. *Some grazing orbits representative of the continuation illustrated in Fig. 15.6. All orbits have the same amplitude, but different periods.*

16 Atlas Events and Toolbox Integration

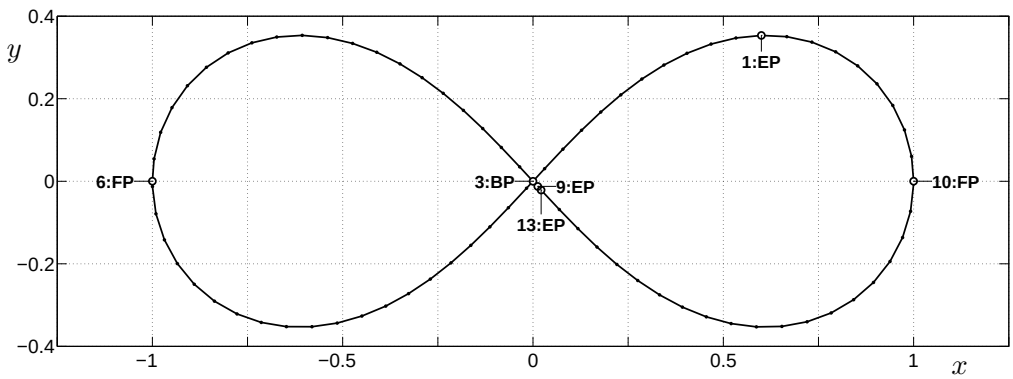

Figure 16.1. *Computation of the lemniscate curve in Example 16.1 with detection of atlas events. Two fold points at labels 6 and 10 and a branch point at label 3 are located. The covering closes at the end points 9 and 13.*

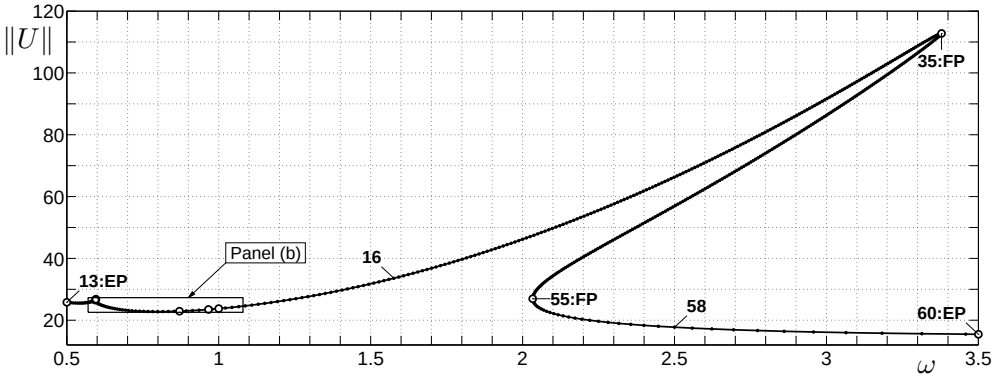

(a) Frequency response curve of Duffing oscillator. A number of atlas events are detected.

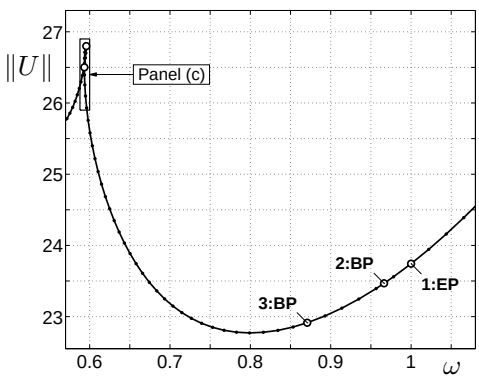

(b) Enlargement with two branch points.

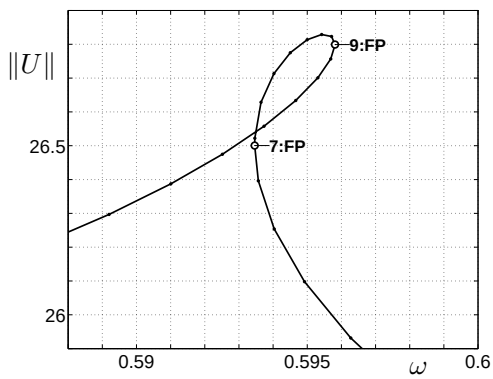

(c) Enlargement with two fold points.

Figure 16.2. *Frequency response curve of the Duffing oscillator given by the vector field in Eq. (16.10) for $A = 2.5$, $\lambda = 0.2$, and $\alpha = \varepsilon = 1$, as computed in Example 16.2 with $\|U\| := \|(u, \mu_{\mathbb{J}})\|$. A number of atlas events is detected along this curve, as shown in Panel (a) and the enlargements (b) and (c). For sample solutions at labels 16, 35, and 58 along the curve we observe a characteristic phase shift, as illustrated in Figure 16.3.*

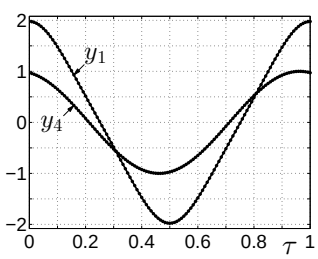

(a) Solution at label 16.

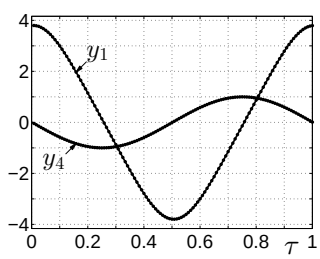

(b) Solution at label 35.

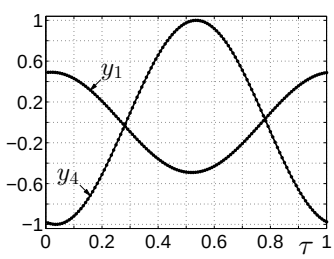

(c) Solution at label 58.

Figure 16.3. *For oscillations obtained at both sides of the main resonance peak of the frequency response curve shown in Fig. 16.2, we observe a characteristic phase shift between the forcing $y_4(t)$ and the response $y_1(t)$, here shown over the normalized time $\tau = t/T$. To the left, both oscillations are in phase (a), while to the right we observe an anti-phase response (c). The transition occurs close to the fold point at label 35 (b).*

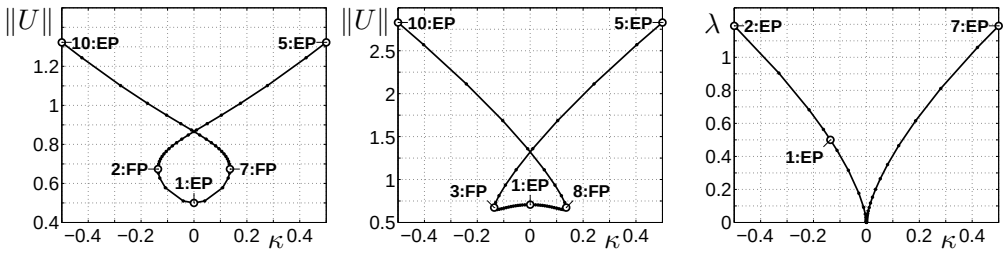

(a) Detection of fold points. (b) Different projection. (c) Locus of fold points.

Figure 16.4. *Fold-point detection and continuation for the cusp normal form considered in Sect. 16.2.3 using the monitor function defined in Sect. 16.2.1 embedded in the* `alg` *toolbox. Panels (a) and (b) show the result of computing the solution to the restricted continuation problem illustrated in Fig. 15.2 with the monitor function added as either* regular *or* active*, respectively, and with* $\|U\| := \|(u, \mu_{\mathbb{J}})\|$*. The difference in appearance is due to the exclusion or inclusion of the value of the monitor function in the vector of continuation parameters* μ*. Two fold points are located in each case. Restarting at one of these fold points while restricting the monitor function to 0 results in a covering of the locus of fold points (c); see also Fig. 15.1.*

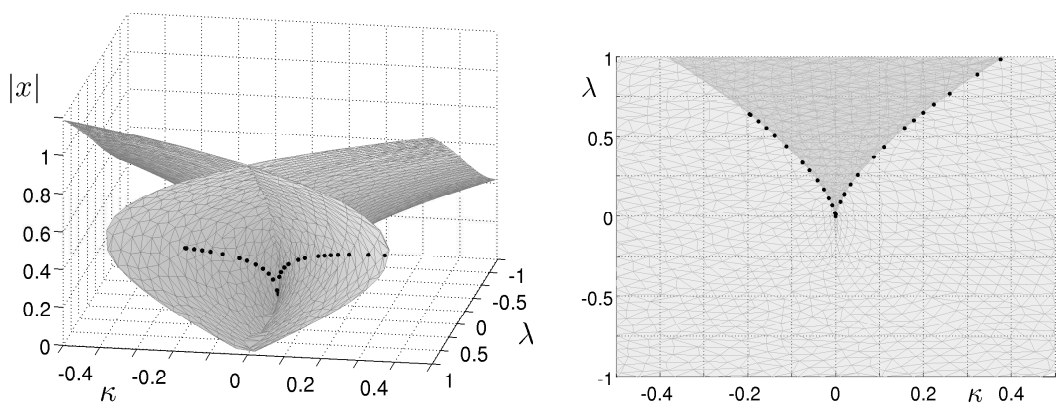

(a) 2-dimensional covering of cusp surface. (b) Projection onto (κ, λ)-parameter plane.

Figure 16.5. *The covering of the cusp surface obtained in the last run of Sect. 16.2.3; fold points are marked as black dots. The 2-dimensional atlas algorithm locates a fold point whenever a curve segment intersects an event surface transversally. In panel (a), one can see that events between neighboring charts may remain undetected if no connecting curve segment was constructed by the covering algorithm. Projecting the manifold onto the* (κ, λ)*-parameter plane again reveals the fold curve as the border curve between different shades of gray (b).*

17 Event Handlers and Branch Switching

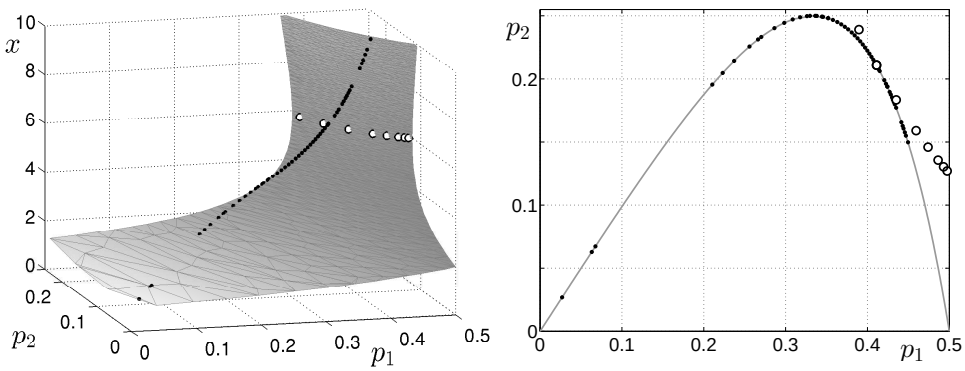

(a) Manifold of equilibrium points.

(b) Hopf- and fold points.

Figure 17.1. *Covering of a manifold of equilibrium points of Example 17.1. The black dots mark events associated with* ψ_{HB} *and the value* 0*, which are either Hopf points or neutral saddles. The circles mark fold points. From panel (a) it is evident that the loci of Hopf and fold points intersect each other transversally. When projected onto the* (p_1, p_2)*-parameter plane, these curves have a tangency (b). The gray curve in panel (b) is the locus of zeros of* ψ_{HB} *given by Eq. (17.4). The distinction between different events associated with the same monitor function is afforded by event handlers; see Fig. 17.3.*

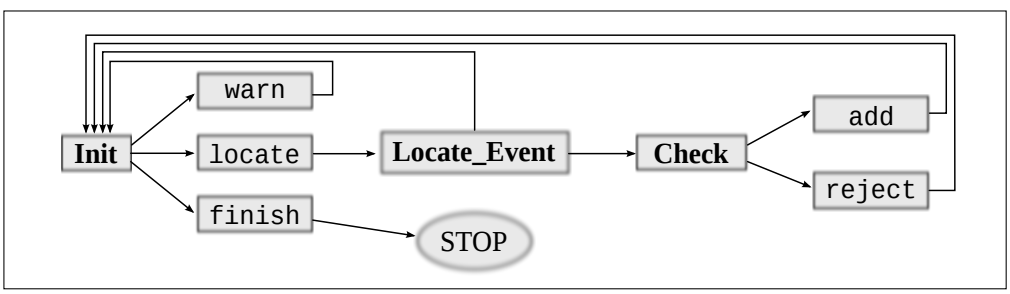

Figure 17.2. *The reverse communication protocol implemented in* COCO *is a three-state finite-state machine. An event handler triggers a state transition from either* **Init** *or* **Check** *and determines the subsequent flow through the finite-state machine by setting the* `action` *field of the* `msg` *message structure to one of the values following an outgoing edge of the current state. Since the lifetime of the message structure spans from* **Init** *until* STOP*, transitions triggered by an event handler can depend on (parts of) the complete history of state transitions. As a typical case, this allows an implementation to determine the state prior to* **Init***, which in turn supports a distinct response depending on whether or not* **Locate_Event** *was successful.*

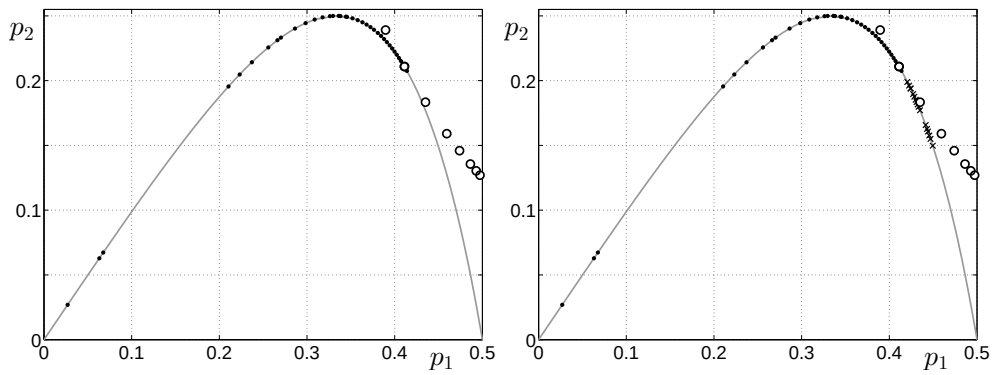

(a) Hopf points. (b) Hopf- and neutral saddle points.

Figure 17.3. *Repeating the computation of Example 17.1 using an event handler, as shown in Sect. 17.1.3, allows for a distinction between Hopf and neutral saddle points; compare with Fig. 17.1(b). Panel (a) shows only Hopf and fold points, while panel (b) shows also neutral saddle points marked with* $\times$*. The gray curve is the locus of zeros of* ψ_{HB} *given by Eq. (17.4).*

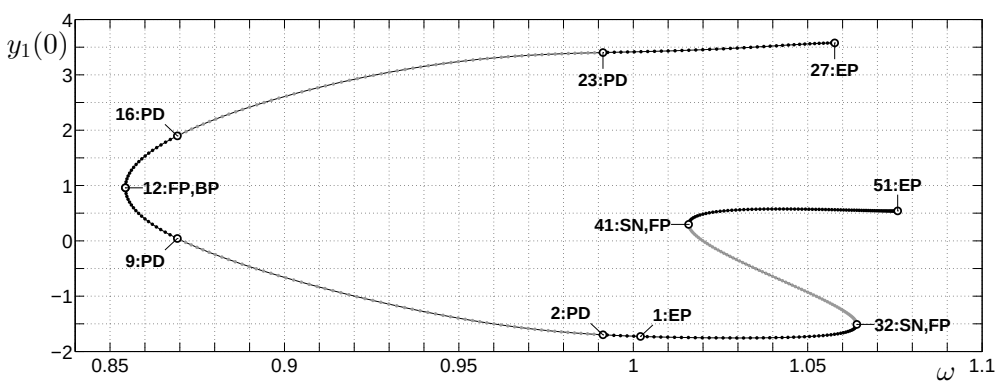

Figure 17.4. *Frequency response curve of the Duffing oscillator from Example 17.3 for forcing amplitude* $A = 26$ *under variations in the forcing frequency* ω*, while the remaining parameters are set to* $\lambda = 0.2$ *and* $\alpha = \varepsilon = 1$*. Here, black dots correspond to periodic orbits with all Floquet multipliers within the unit circle, whereas gray dots correspond to periodic orbits with at least one Floquet multiplier outside of the unit circle. A number of toolbox and atlas events is detected along this curve. As a special case, the branch point at label 12 is a pitchfork bifurcation point that marks a characteristic transition from nonsymmetric responses through a symmetric one to symmetry-conjugate responses. This transition is illustrated with sample solutions at labels 9, 12 and 16 in Fig. 17.5.*

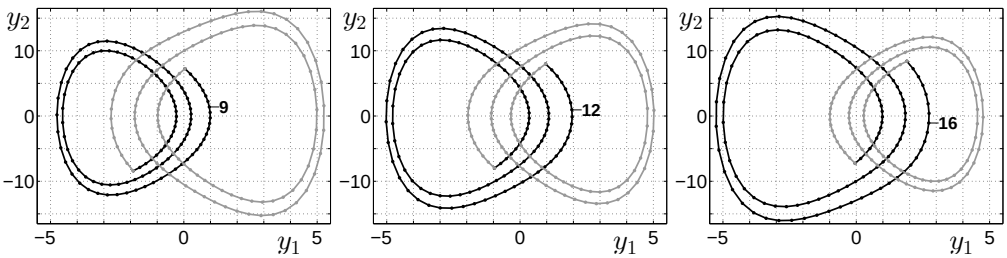

(a) Non-symmetric orbit. (b) Symmetric orbit. (c) Conjugate orbit of (a).

Figure 17.5. *The Duffing equation with bang-bang forcing in Example 17.3 is equivariant under rotations of the (y_1, y_2) plane by π, combined with a phase shift by π. As a consequence, we expect to observe symmetric, as well as nonsymmetric oscillations, where existence of a nonsymmetric oscillation implies existence of the symmetry-conjugate oscillation. In the present case, continuation starts at a family of nonsymmetric oscillations, passes through a symmetry-increasing bifurcation point, and continues to produce symmetry-conjugate oscillations, as shown in the sequence of plots in panels (a)-(c). The solutions in panels (a) and (c) are obtained for the same value of ω, but on different sides of the branch point. Clearly, one can transform (a) into (c) by a rotation by π and exchanging black for gray.*

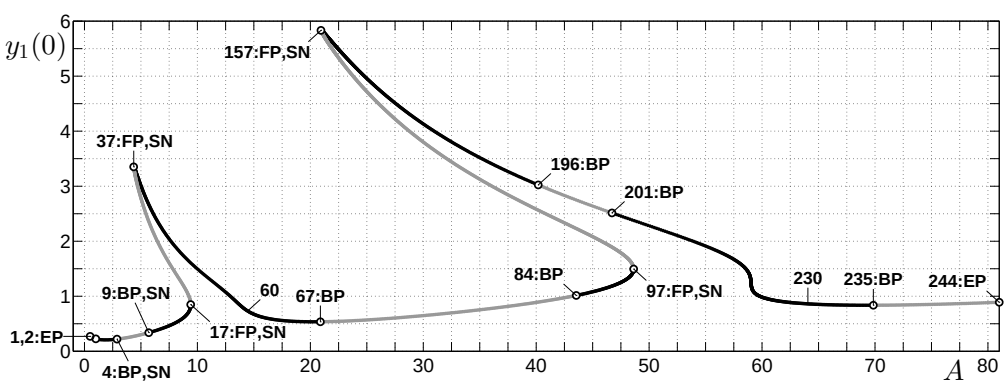

Figure 17.6. *Amplitude response curve of the Duffing oscillator from Example 17.4 for $\omega = 1$ computed with the atlas code developed in Sect. 16.1. Again, $\lambda = 0.2$ and $\alpha = \varepsilon = 1$. Along this curve we indicate the location of atlas and toolbox events. In addition, we encode stability information provided by the toolbox monitor function. Here, black corresponds to stable and gray to unstable oscillations. We observe two resonance peaks, where passage through each peak corresponds to the addition of an oscillation during one forcing half-cycle as illustrated with sample solutions at labels 1, 60, and 230 in Fig. 17.7. Note that all solutions along this response curve are symmetric. The analysis of this example continues in Fig. 17.8.*

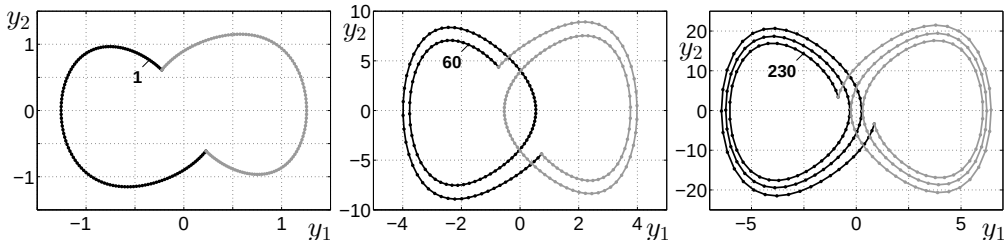

(a) One-loop orbit. (b) Two-loop orbit. (c) Three-loop orbit.

Figure 17.7. *After each passage through a resonance peak along the amplitude response curve shown in Fig. 17.6, we observe the addition of an oscillation during one half-cycle of the forcing.*

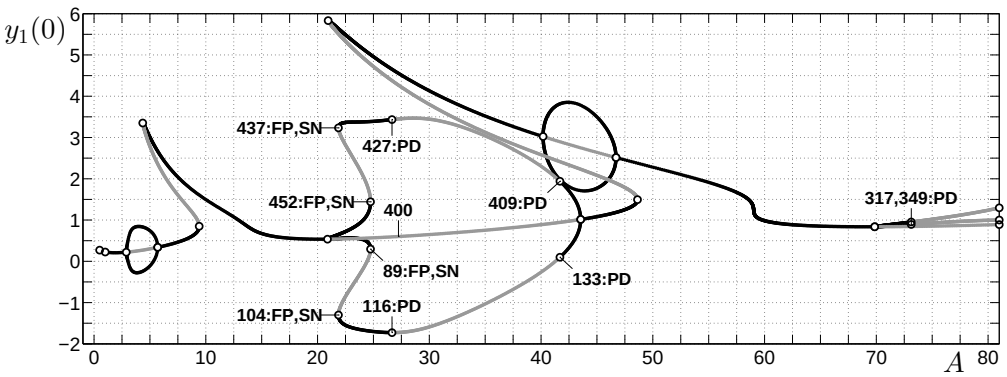

Figure 17.8. *Amplitude response curve of the Duffing oscillator from Example 17.4 for $\omega = 1$ computed with the atlas code developed in Sect. 17.3, which implements automatic branch switching at branch points. Again, $\lambda = 0.2$ and $\alpha = \varepsilon = 1$. In addition to the results shown in Fig. 17.6, we obtain a number of families of nonsymmetric oscillations emerging from the branch points located during continuation; see also Fig. 17.9. Three of these form closed families and the others terminate at the computational boundary. Along the closed family in the center of the figure, a number of toolbox events are located, including period-doubling bifurcation points. The analysis of this example continues in Fig. 17.10. The overlap of two curves at the period-doubling point 409 is an artifact of projection.*

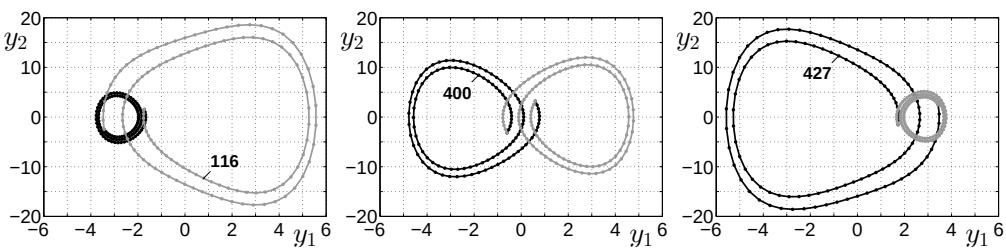

(a) Orbit 116. (b) Orbit 400 (c) Orbit 427

Figure 17.9. *A sample of periodic orbits for $A \approx 27$ on the family of symmetric (b), and along the closed family of nonsymmetric (a) and (c), two-loop oscillations located in the center of Fig. 17.8. The latter two orbits are again symmetry-conjugate and lie at period-doubling points.*

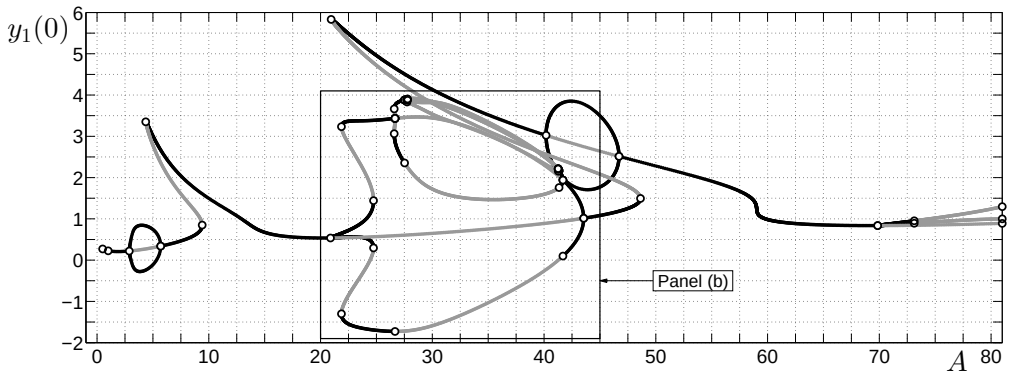

(a) Bifurcation diagram of Duffing oscillator with families of nonsymmetric and period-doubled solutions. A sequence of enlargements of details is shown in panels (b) to (d).

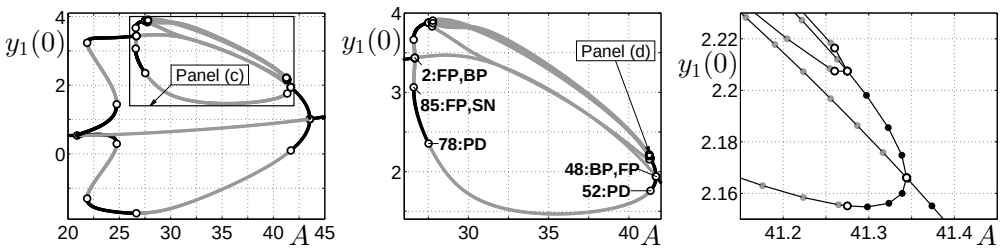

(b) Nonsymmetric family. (c) Period-doubled families. (d) Period-doubling sequence.

Figure 17.10. *Continuation of the analysis from Fig. 17.8. Using the algorithm for branch switching at period-doubling points developed in Sect. 17.3.3, we can further complete the bifurcation diagram (a). Panel (b) shows a zoom into the closed family, along which period-doubling points were detected previously. A continuation of the emerging period-doubled orbits results in another closed family, along which period-doubling points are again detected (c). Repetition of this procedure results in evidence of a sequence of period-doubling bifurcations. The onset of this sequence is shown in the enlargement (d), which includes local families of period-2, -4 and -8 orbits; see also Fig. 17.11. Note that the families of period-doubled solutions were computed with the atlas algorithm without automatic branch switching to prevent redundant coverage.*

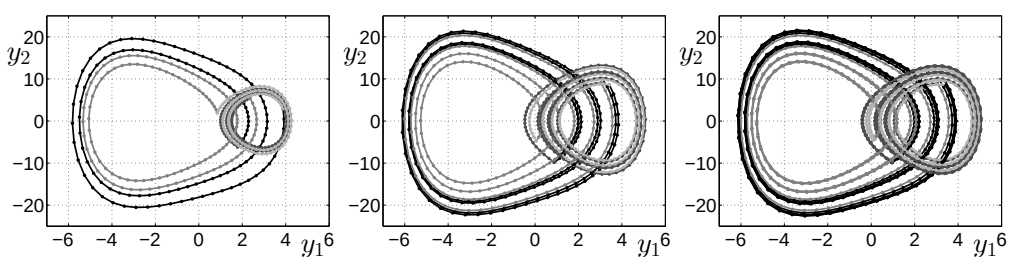

(a) Period-2 orbit. (b) Period-4 orbit. (c) Period-8 orbit.

Figure 17.11. *A sequence of period-doubled orbits along the families shown in Fig. 17.10(d).*

18 Pointwise Adaptation and Comoving Meshes

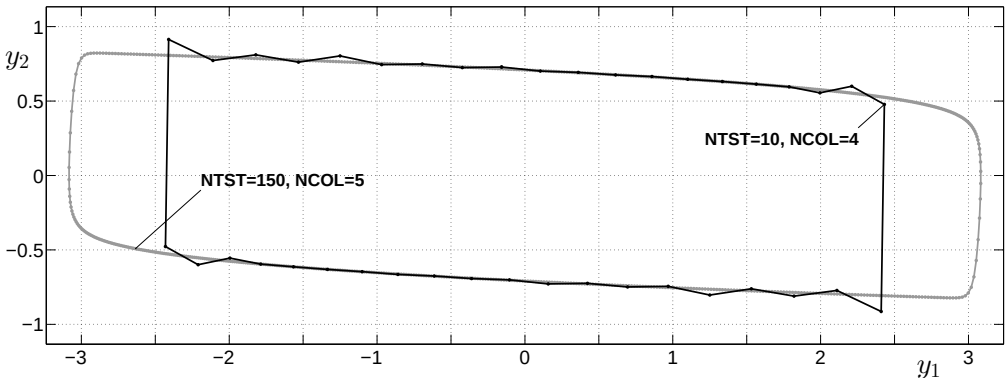

Figure 18.1. *Comparison of approximate periodic orbits of the dynamical system with vector field given by Eq. (18.5) at $\varepsilon = 20$, obtained in Example 18.1 using continuation for different uniform meshes. For the default settings with $N = 10$ and $m = 4$, we obtain a poor approximation. Increasing the discretization parameters to $N = 150$ and $m = 5$ results in a solution curve that seems acceptable, at least by visual inspection. A large subset of the additional mesh points contributes little to the improvement, however; only a small number is allocated along the problematic vertical part of the orbit.*

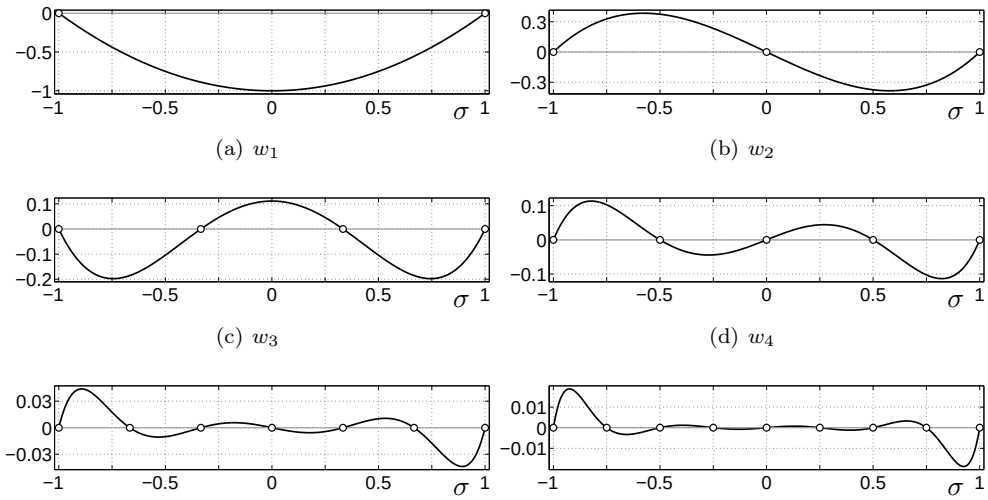

(a) w_1 (b) w_2 (c) w_3 (d) w_4 (e) w_6 (f) w_8

Figure 18.2. *Graphs of the function $w_m(\sigma) := \prod_{i=1}^{m+1}(\sigma - \sigma_i)$ for typical choices of $m = 1, 2, 3, 4, 6, 8$. We observe that, for all these functions, $|w_m|$ assumes its global maximum between the first two zeros located at σ_1 and σ_2, as claimed in Example 18.2.*

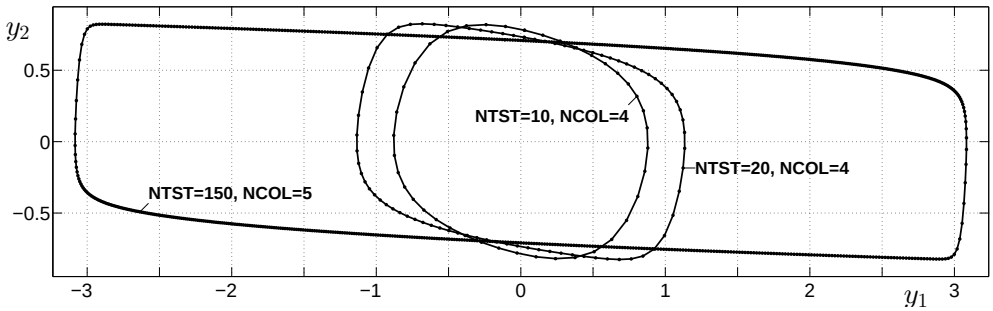

(a) Terminal solutions produced in each of the three runs in Example 18.4. The orbit with NTST equal to 150 is at $\varepsilon = 20$.

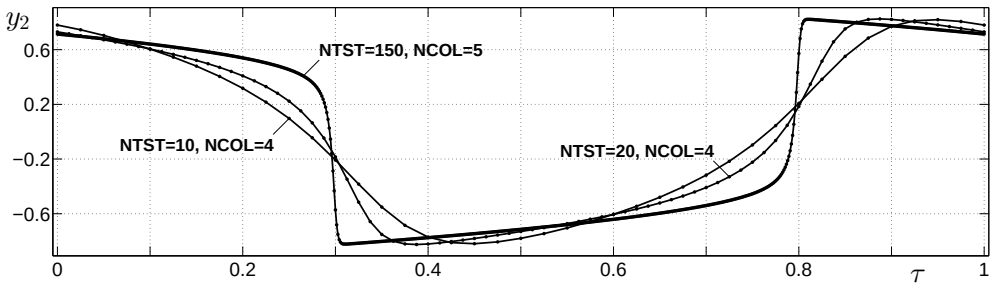

(b) Time profiles of y_2. The two sharp fronts give rise to a large, but localized, error.

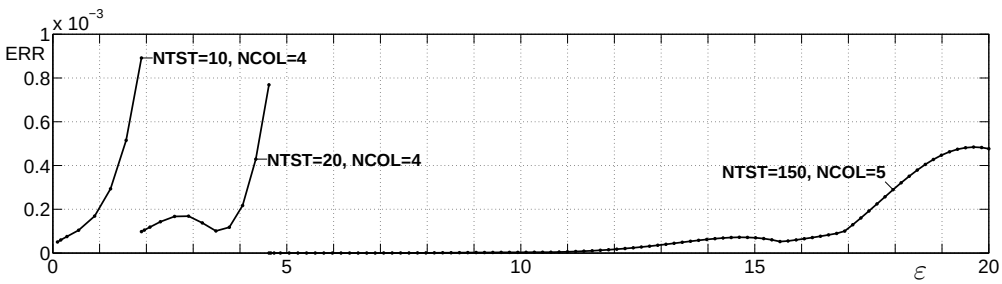

(c) The estimated error for different meshes plotted against variations in ε. The tolerance is 10^{-3}, and the first two runs, with NTST equal to 10 and 20, eventually exceed this tolerance and terminate. The last run, with NTST equal to 150 and NCOL equal to 5, is successful up to $\varepsilon = 20$.

Figure 18.3. *Continuation of periodic orbits of the dynamical system given by the vector field in Eq. (18.5) using pointwise adaptive changes to the discretization order, as described in Sect. 18.1.2. Here, a terminal event is associated with a monitor function that computes an estimate of the discretization error. Continuation stops automatically whenever the estimated error exceeds a predefined limit. As shown in Example 18.4, this approach allows one to compute an approximate solution family within a desired tolerance. Here, we need to restart twice with a finer mesh at terminal solution points, represented by the corresponding approximate periodic orbits in panel (a) together with the solution for* $\varepsilon = 20$. *The reason for the failure to remain within the desired tolerance with low discretization order is explained by the sharp fronts evident in the corresponding time profiles in panel (b). Panel (c) shows the estimated error plotted against variations in* ε.

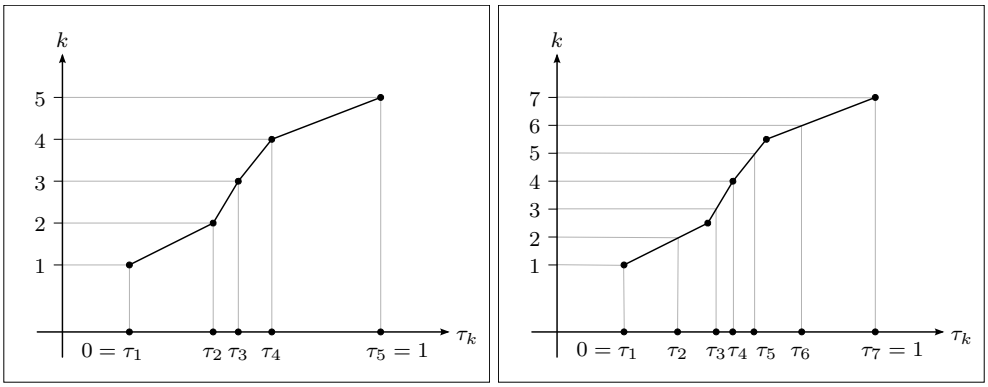

(a) $N = M = 4$.

(b) $N = 6$, $M = 4$.

Figure 18.4. *Initialization of the κ array, as shown in Example 18.8. The piecewise-linear interpolant $\hat{t}^{-1}$, here shown as a function of normalized time $\tau = t/T$, is constructed such that $\hat{t}\left((k-1)\frac{N}{M}+1\right) = \tau_k$ for $k = 1, \ldots, M+1$. As is evident from this construction visualized in (a), when $N = M$, the inverse of the sequence $\{1, \ldots, N+1\}$ under $\hat{t}$ is the sequence $\{\tau_1, \ldots, \tau_{N+1}\}$ itself. An initialization following the outlined procedure thus preserves the initial time distribution of mesh points when $N = M$. If the number of mesh points is changed, the new mesh will have a time distribution that incorporates adaptation information provided with an initial solution (b).*

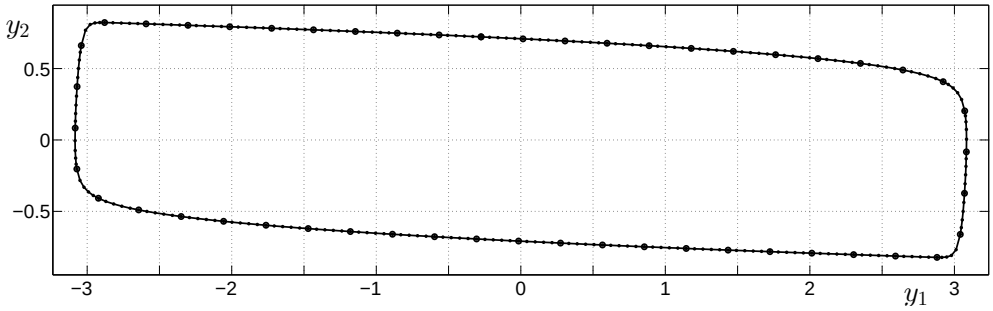

(a) Approximate periodic orbit for $\varepsilon = 20$. The fat dots mark the end points of mesh intervals. The arclengths of all mesh intervals are approximately equal, while points inside a mesh interval are equidistributed with respect to time.

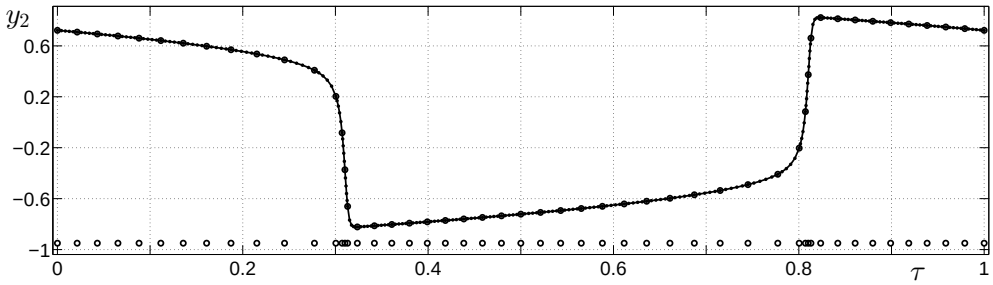

(b) Time profile of y_2. The circles at the bottom mark the end points of the mesh intervals.

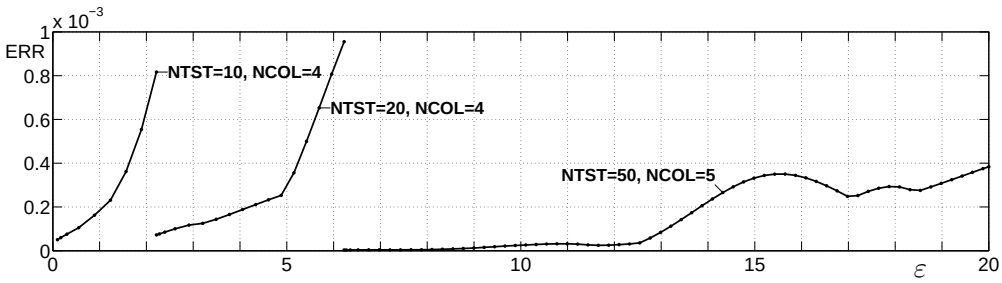

(c) The estimated error for different meshes plotted against variations in ε. The first two runs, with NTST equal to 10 and 20, respectively, eventually exceed the desired tolerance of 10^{-3} and terminate. The subsequent run, with NTST and NCOL equal to 50 and 5, respectively, is successful up to $\varepsilon = 20$.

Figure 18.5. *Continuation of periodic orbits of the dynamical system given by the vector field in Eq. (18.5) using comoving-mesh adaptation aimed at equalizing the arclength across all collocation intervals in state space, as described in Sect. 18.3.1. Here, the distribution of mesh points is uniquely determined along the solution manifold by appending a family of zero functions to the continuation problem, corresponding to a mixed Euler method applied to a gradient vector field. In contrast to Examples 18.1 and 18.4, the comoving mesh algorithm allows one to compute the solution family with the same desired tolerance as in Fig. 18.3, but with only* $1/3$ *of the number of mesh points. Again we need to restart twice with a finer mesh at terminal solution points. The solution obtained for* $\varepsilon = 20$ *is shown in panel (a). Panel (b) shows the corresponding time profile and illustrates the adaptive time mesh. The estimated error plotted against variations in* ε *is shown in panel (c).*

19 A Spectral Toolbox

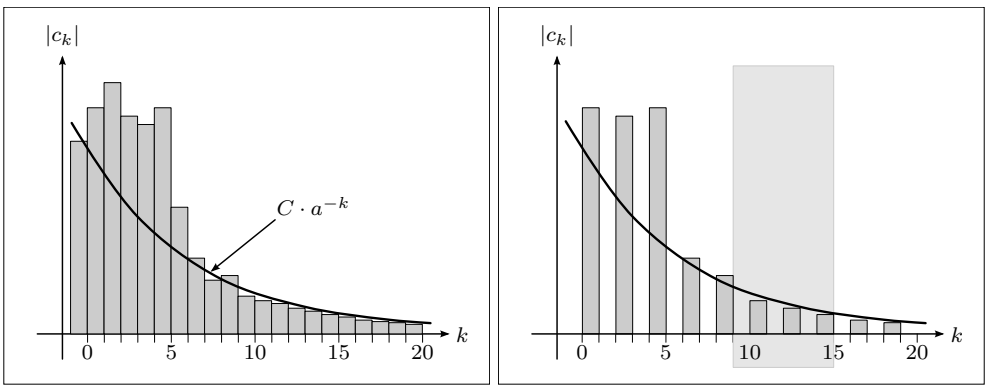

(a) Exponential decay. (b) Decay for odd function.

Figure 19.1. *For real-analytical 2π-periodic functions, the sequence $\{|c_k|\}_{k=0}^{\infty}$ of the absolute values of the Fourier coefficients decays asymptotically with an exponential rate. There thus exist constants $C > 0$ and $a > 1$, and a mode number $K \geq 0$, such that $|c_k| \leq C \cdot a^{-k}$ holds for all $k \geq K$. If one interprets the sequence $\{|c_k|\}_{k=0}^{\infty}$ as a piecewise-constant function, as indicated in panel (a), one obtains the estimate $S_{K+1} := \sum_{k=K+1}^{\infty} |c_k| \leq \int_K^{\infty} C \cdot a^{-x}\, dx = C \cdot a^{-K}/\ln a$. Making the simplifying assumptions that $|c_K| \approx C \cdot a^{-K}$ and $\ln a = 1$ gives $S_{K+1} \leq |c_K|$. The tail S_{K+1} can thus be estimated by the first term omitted from the sum S_K. Some caution is necessary when using this idea for estimating an approximation error. For example, for odd functions it holds that $c_{2k} = 0$, for $k = 0, 1, 2, \ldots$. A robust error estimation should therefore use the values of $|c_k|$ from a window $k_0 \leq k \leq k_1$ of size $k_1 - k_0 \geq 1$ as indicated in panel (b) with $k_0 = 10$ and $k_1 = 15$. A related idea of error estimation is implemented in Sect. 19.1.5 in the context of the spectral zero problem.*

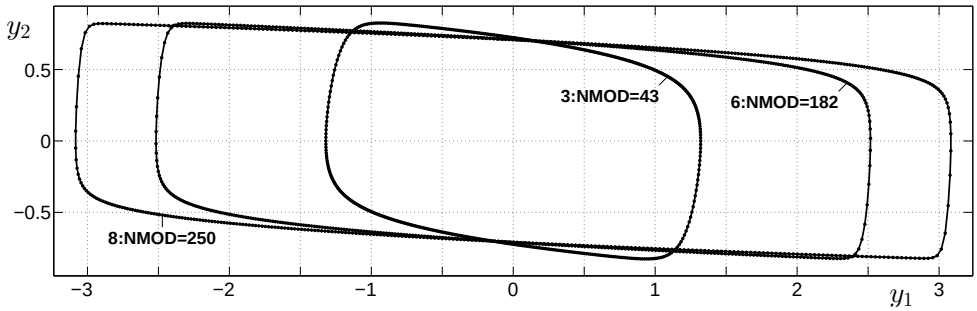

(a) Approximate periodic orbits obtained during continuation for increasing ε. The number of modes used for each approximation is indicated by the label.

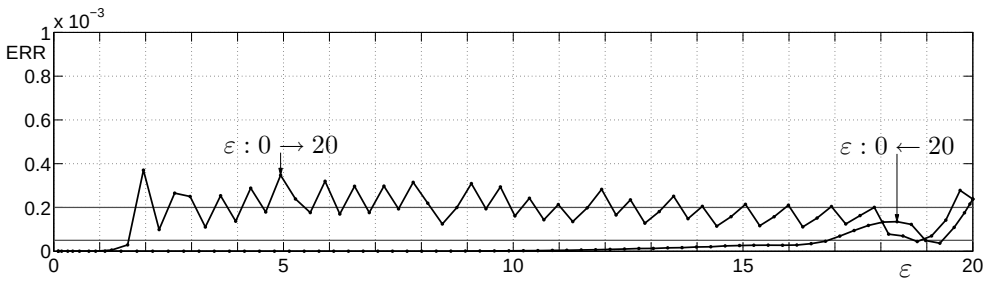

(b) Estimated error during continuation for increasing and decreasing values of ε. The tolerance is 10^{-3}, and the adaptation window is indicated with two horizontal lines.

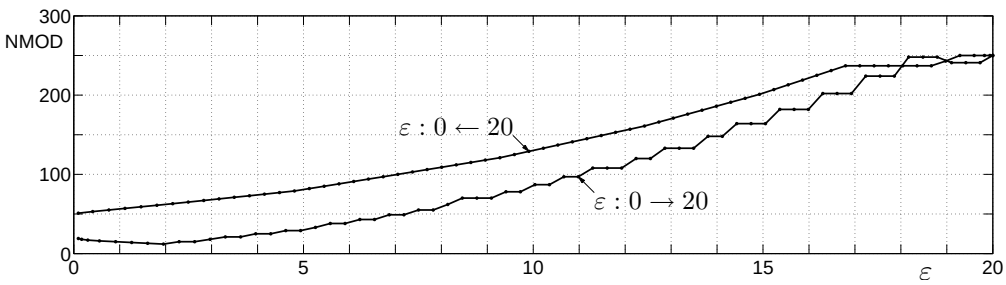

(c) Number of Fourier modes during continuation with `'NMAX'` equal to 250.

Figure 19.2. *Continuation of periodic orbits of the dynamical system given by the vector field in Eq. (18.5) using the adaptive spectral toolbox, as implemented in Sects. 19.2 and 19.3. The rapid decay of Fourier coefficients illustrated in Fig. 19.1 enables adaptation by exchanging equations that either fix a variable at* 0 *or make it equal to a Fourier mode. In contrast to the results shown in Figs. 18.3 and 18.5, continuation here computes the complete family in a single run. Panel (a) shows sample orbits. The effects of adaptation are illustrated in panels (b) and (c).*

20 Integrating Adaptation in Atlas Algorithms

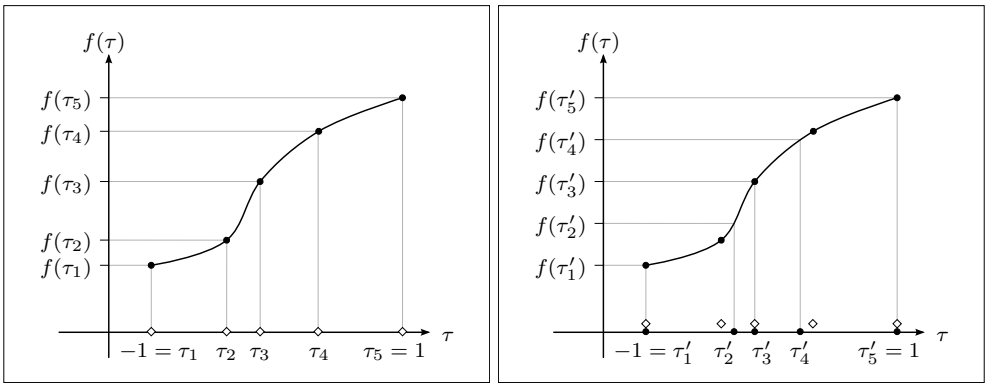

(a) Interpolation of f. (b) Inversion of equidistribution.

Figure 20.1. *The action of the remeshing map Υ defined in Eq. (20.6) is computed in a two-step algorithm. In a first step, an interpolation of an increasing function of τ is constructed for a given sequence $\{\tau_k\}_{k=1}^N$, as indicated with diamonds in panel (a). In a second step, one computes the inverse of an equidistributed sequence of f-values to obtain the sequence $\{\tau'_k\}_{k=1}^N$, which is equidistributed with respect to f and represented by the black dots in panel (b).*

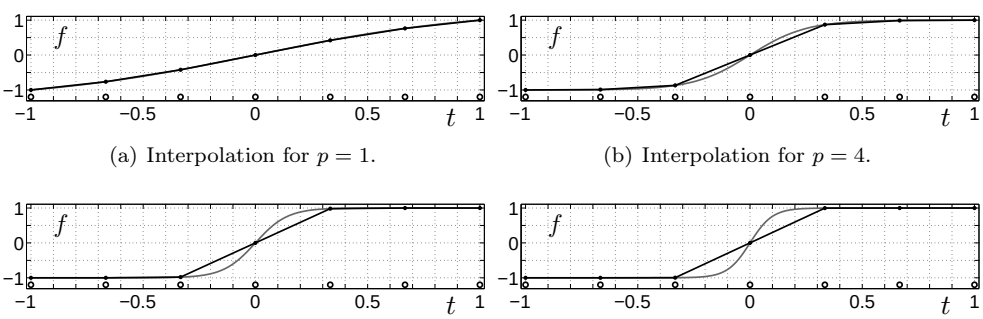

(a) Interpolation for $p = 1$. (b) Interpolation for $p = 4$.

(c) Interpolation for $p = 7$. (d) Interpolation for $p = 10$.

Figure 20.2. *Interpolation of the increasing test function $f(t,p) = \tanh(pt)/\tanh(p)$ in Example 20.4 on a uniform mesh with seven mesh points for increasing values of p, as described in Sect. 20.1.3. As seen in the sequence (a)-(d), a steep front develops in the center of the graph. This is clearly poorly resolved with a uniform mesh. The circles at the bottom indicate the mesh in t. The result of using a mesh that is equidistributed with respect to f is shown in Fig. 20.3.*

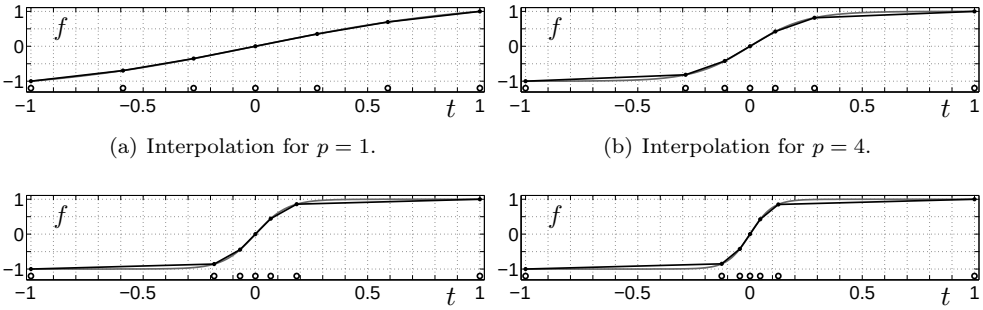

(a) Interpolation for $p = 1$.

(b) Interpolation for $p = 4$.

(c) Interpolation for $p = 7$.

(d) Interpolation for $p = 10$.

Figure 20.3. *Interpolation of the increasing test function* $f(t, p) = \tanh(pt)/\tanh(p)$ *using an equidistributed moving mesh with seven mesh points for increasing values of* p*, as described in Sect. 20.1.3. The circles at the bottom indicate the mesh in* t*. The quality of interpolation of the steep front improves compared with the graphs shown in Fig. 20.2. The result of using an equidistributed moving mesh with variable number of mesh points is shown in Fig. 20.4.*

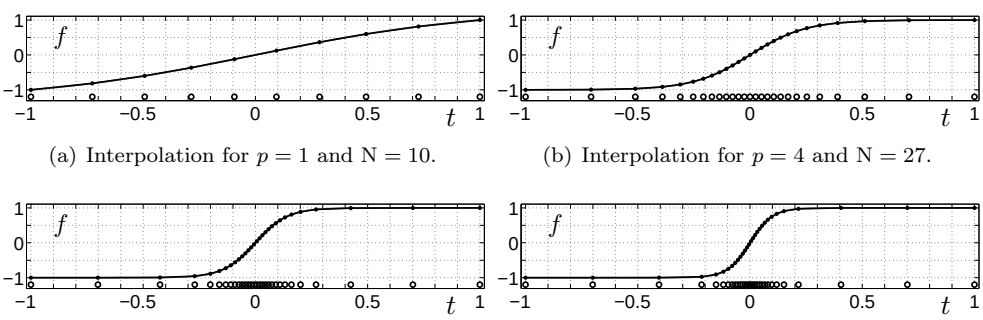

(a) Interpolation for $p = 1$ and N = 10.

(b) Interpolation for $p = 4$ and N = 27.

(c) Interpolation for $p = 7$ and N = 30.

(d) Interpolation for $p = 10$ and N = 30.

Figure 20.4. *Interpolation of the increasing test function* $f(t, p) = \tanh(pt)/\tanh(p)$ *using an equidistributed moving mesh with varying number of mesh points for increasing values of* p*, as described in Sect. 20.1.3. The circles at the bottom indicate the mesh in* t*. The overall quality of interpolation improves compared with Fig. 20.3, albeit at the expense of adding mesh points.*

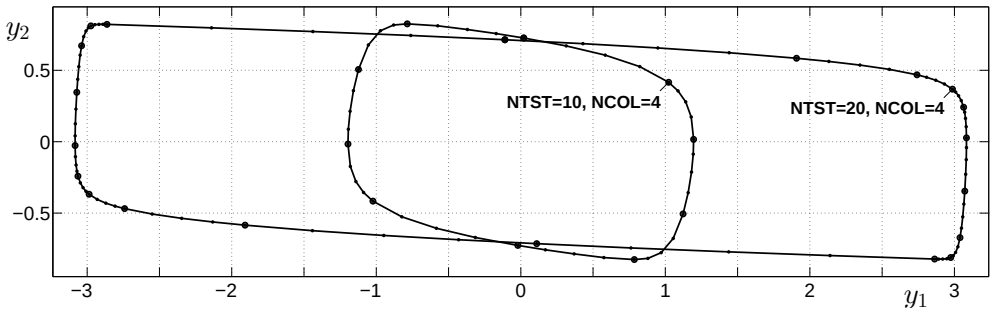

(a) Selected approximate periodic orbits obtained during continuation. The solution with 20 mesh intervals is obtained at $\varepsilon = 20$. The fat dots mark the end points of mesh intervals.

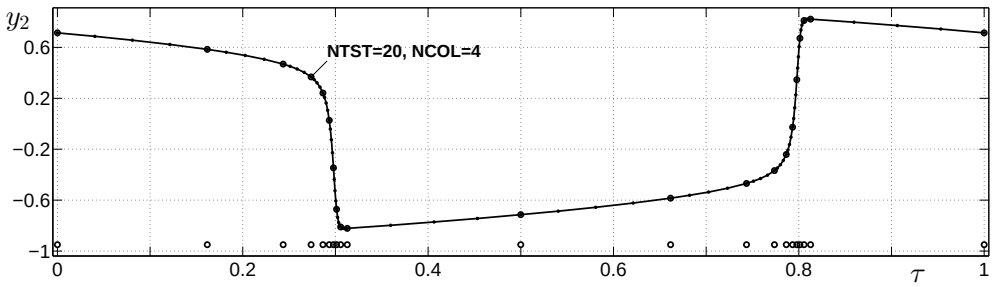

(b) Time profile of y_2 for $\varepsilon = 20$. The circles at the bottom mark the end points of the mesh intervals.

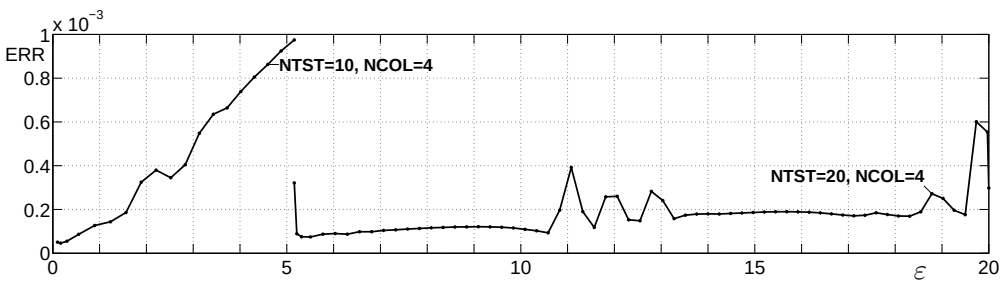

(c) The estimated error for different meshes plotted against variations in the parameter ε. The run with the `'NTST'` setting equal to 10 eventually exceeds the desired tolerance of 10^{-3} and terminates.

Figure 20.5. *Continuation of periodic orbits of the dynamical system given by the vector field in Eq. (18.5) using an equidistributed moving mesh without adaptive changes to the discretization order, as described in Sect. 20.2.1. The algorithm allows one to compute the solution family with the same desired tolerance as in Fig. 18.3 with* $1/9$ *of the number of mesh points. Panel (a) shows sample orbits obtained during continuation, and panels (b) and (c) illustrate the effect of adaptation. A refined strategy with varying discretization order is illustrated in Fig. 20.6.*

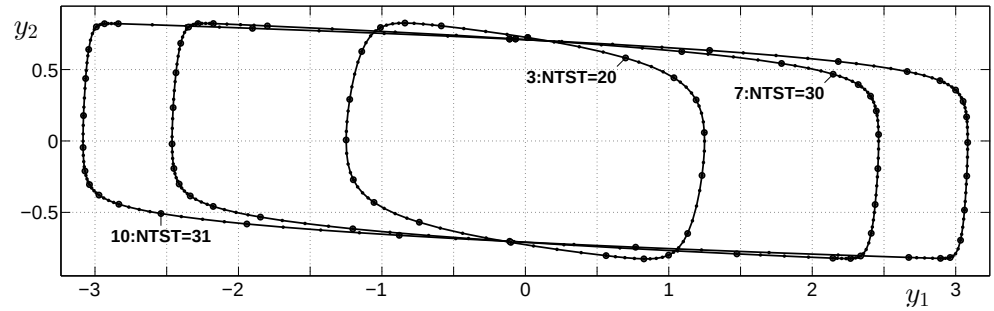

(a) Selected approximate periodic orbits obtained during continuation for increasing ε. The number NTST of mesh intervals is indicated with the label. The fat dots mark the end points of mesh intervals.

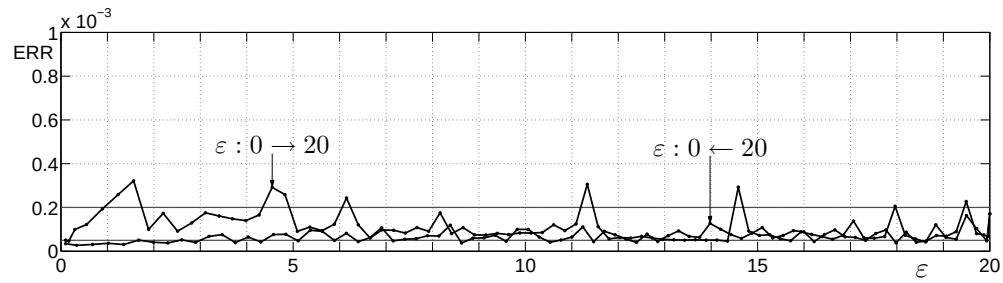

(b) Estimated error during continuation for increasing and decreasing ε. The tolerance is 10^{-3}, and the adaptation window is indicated by two horizontal lines.

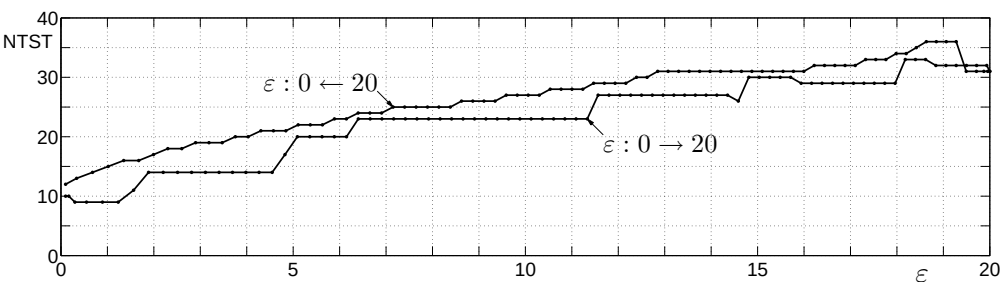

(c) Number of mesh intervals during continuation for increasing and decreasing ε.

Figure 20.6. *Continuation of periodic orbits of the dynamical system given by the vector field in Eq. (18.5) using an equidistributed moving mesh with adaptive changes to the discretization order, as described in Sect. 20.2.2. Here, a single continuation run suffices to compute the entire solution family. Due to the conservative definition of the adaptation window, the algorithm uses about 50% more mesh points than the analysis shown in Fig. 20.5. Panel (a) shows sample orbits obtained during continuation, and panels (b) and (c) illustrate the effect of adaptation.*

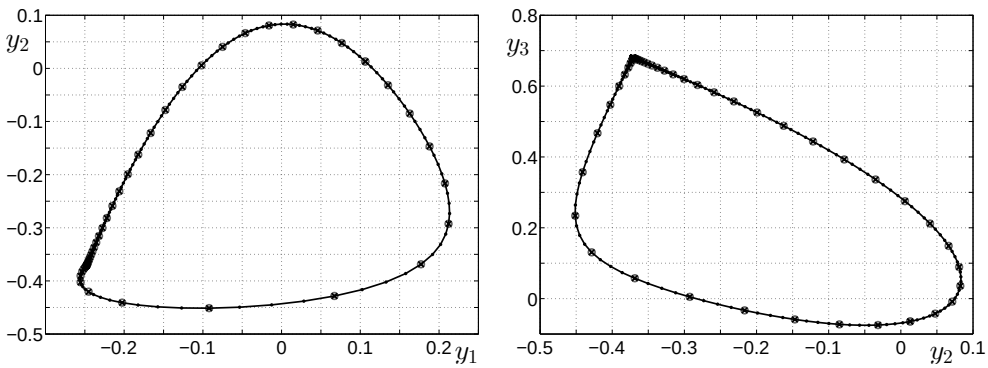

(a) In (y_1, y_2) plane. (b) In (y_2, y_3) plane.

Figure 20.7. *State-space reprentations of near-homoclinic approximate periodic orbits of the dynamical system given by the vector field in Eq. (20.44), obtained using continuation with the nonadaptive* `'coll'` *toolbox from Sect. 18.1.2 with* `'NTST'` *equal to* 50 *and* `'TOL'` *equal to* 10^{-4}. *Panels (a) and (b) show the terminal low-period solution (gray circles) with 50 mesh intervals found using continuation with a* 1*-dimensional atlas algorithm starting at a Hopf bifurcation point, as well as the corrected, high-period solution (black curve and end points of mesh intervals denoted by* $\times$*'s) found using continuation with a* 0*-dimensional atlas algorithm applied to a reconstructed periodic orbit with 1,250 mesh intervals. The estimated error is* 4.2532×10^{-5} *for the low-period orbit and* 4.2531×10^{-5} *for the high-period orbit. These results are compared in Figs. 20.9 and 20.11 with computations using moving-mesh adaptation.*

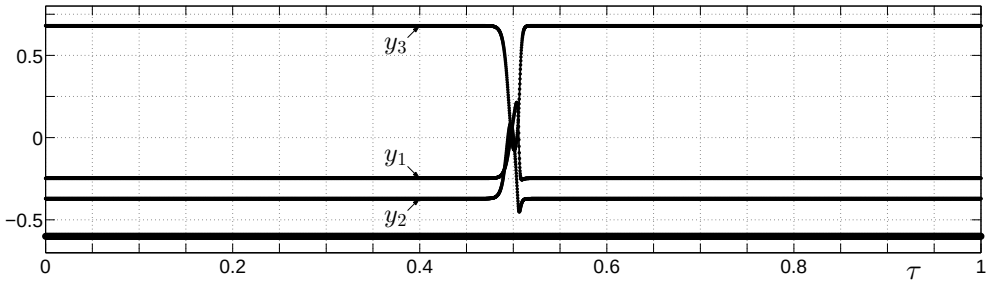

(a) Time profile of high-period orbit.

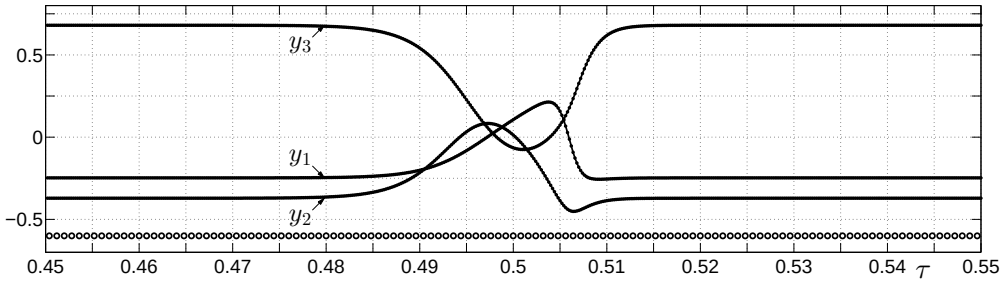

(b) Zoom into excursion from equilibrium.

Figure 20.8. *Time profiles of the high-period orbit constructed in Fig. 20.7 on a uniform mesh shown over the full period (a) and in a zoom into the excursion from the equilibrium (b). The distribution of mesh points is indicated with circles at the bottom. In panel (a), the mesh is so dense that it is not possible to see individual mesh intervals.*

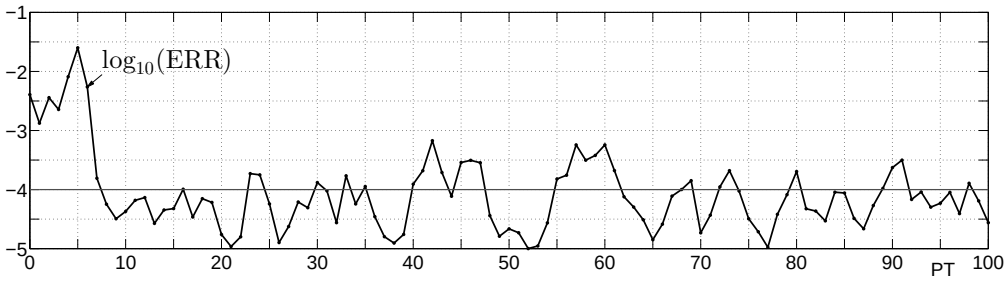

(a) The logarithm of the estimated discretization error during 100 iterations of the remesh-correct cycle of a 0-dimensional atlas algorithm. A mesh with the desired tolerance of 10^{-4} (gray horizontal line) is arrived at within 10 iterations.

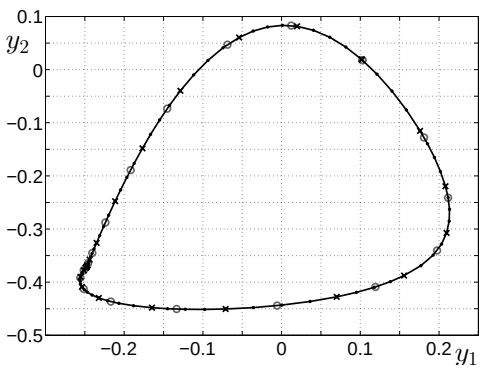

(b) In (y_1, y_2) plane.

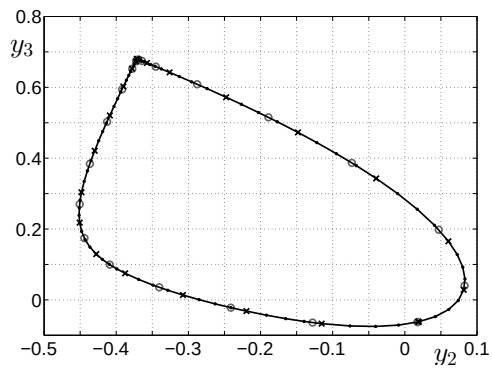

(c) In (y_2, y_3) plane.

Figure 20.9. *Mesh adaptation and an approximate near-homoclinic periodic orbit of the dynamical system given by the vector field in Eq. (20.44), obtained using the collocation method with a moving mesh of fixed order from Sect. 20.2.1 with* `'TOL'` *equal to* 10^{-4}. *Panel (a) shows variations in the discretization error during 100 iterations of a* 0-*dimensional remesh-correct cycle applied to the reconstructed initial solution guess obtained by a 5,000-fold increase in the period* T *and a* 6-*fold increase in discretization order* N. *In panels (b) and (c), we compare the corrected high-period orbit (black) with the low-period orbit (gray circles) obtained using continuation from the Hopf bifurcation point. The number of mesh points along the excursion from the equilibrium is nearly identical for both orbits, although the mesh points move somewhat. The estimated error is* 2.2619×10^{-5} *for the low-period orbit and* 4.2620×10^{-5} *for the high-period orbit. The adapted mesh and time profile of the high-period orbit are shown in Fig. 20.10. Figures 20.11-20.13 repeat the analysis with a moving mesh with variable discretization order.*

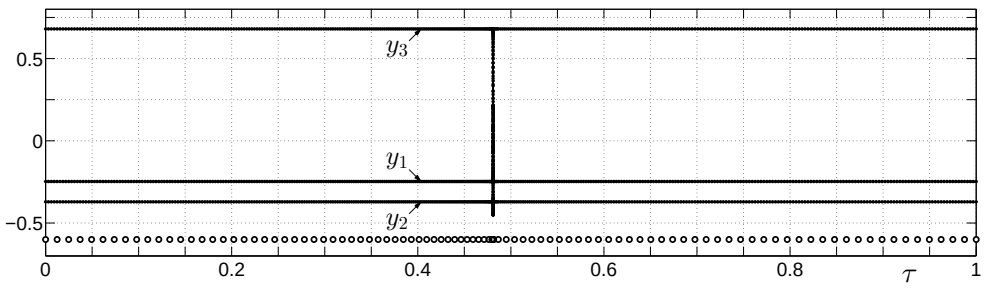

(a) Time profile of the high-period orbit.

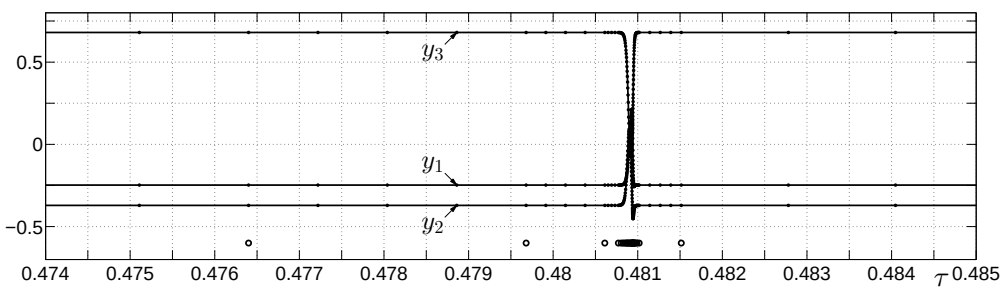

(b) First zoom into excursion from equilibrium.

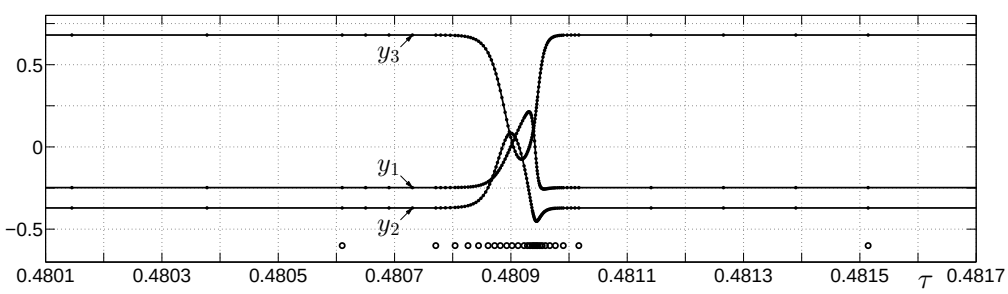

(c) Second zoom into excursion from equilibrium.

Figure 20.10. *Time profiles of the high-period orbit constructed in Fig. 20.9 on a moving mesh shown over the full period (a) and in two subsequent zooms into the excursion from the equilibrium (b) and (c). The distribution of mesh points is indicated with circles at the bottom.*

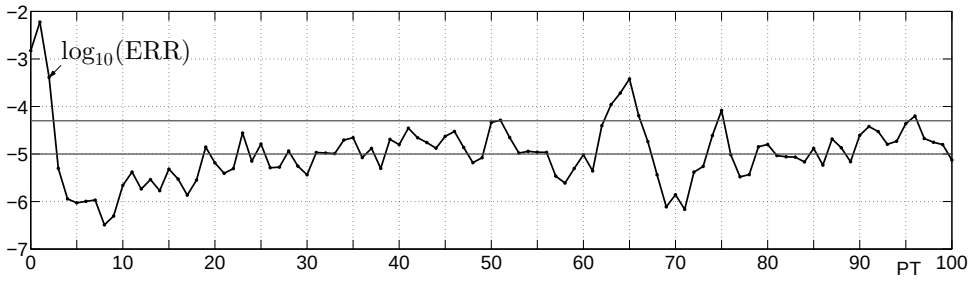

(a) The logarithm of the estimated discretization error during 100 iterations of the remesh-correct cycle of a 0-dimensional atlas algorithm.

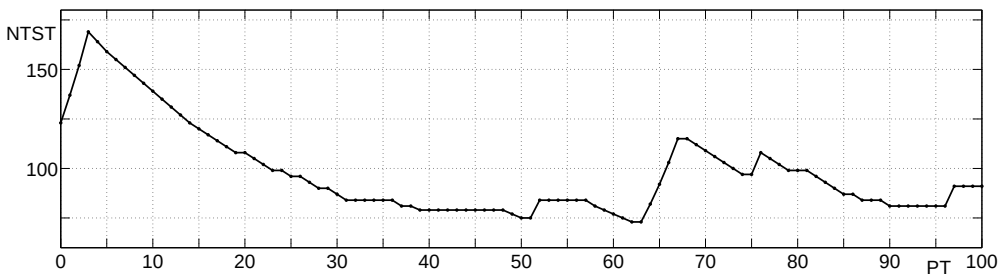

(b) The number of mesh intervals plotted against the number of iterations of the remesh-correct cycle of a 0-dimensional atlas algorithm.

Figure 20.11. *Mesh adaptation for an approximate near-homoclinic periodic orbit of the dynamical system given by the vector field in Eq. (20.44), obtained using the collocation method with a moving mesh of variable discretization order from Sect. 20.2.2 with* `'TOL'` *equal to the default tolerance of* 10^{-4}. *Panels (a) and (b) show variations in the discretization error and discretization order during 100 iterations of a* 0*-dimensional remesh-correct cycle applied to the reconstructed initial solution guess obtained by a 5,000-fold increase in the period* T *and a* 6*-fold increase in discretization order* N. *Again, we observe that the iteration settles onto an acceptable mesh after a somewhat longer transient phase. Although the approximation error drops much faster than in Fig. 20.9(a), the dynamics is dominated by the slow reduction of the number of mesh intervals. After 40 steps, the solution settles onto a mesh of size* $N = 79$. *A comparison of the resulting solutions is shown in Fig. 20.12. Figure 20.13 shows the adapted mesh and time profile of the high-period orbit.*

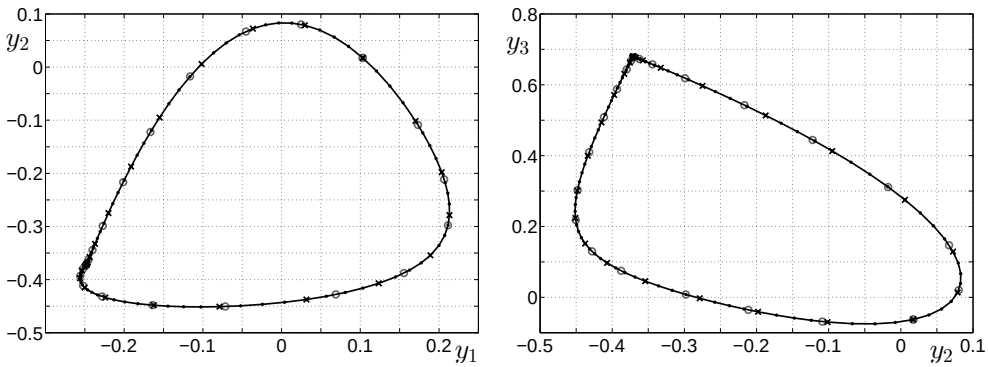

(a) In (y_1, y_2) plane. (b) In (y_2, y_3) plane.

Figure 20.12. *Comparison of the high-period orbit (black) obtained in Fig. 20.11, with the low-period orbit (gray circles) used for constructing the initial solution guess, in two different projections (a) and (b). Again, the number of mesh points along the excursion from the equilibrium is nearly identical for both orbits, although the mesh points move somewhat. The estimated error is* 1.2506×10^{-5} *for the low-period orbit and* 1.5747×10^{-5} *for the high-period orbit. Figure 20.13 shows the adapted mesh and time profile of the high-period orbit.*

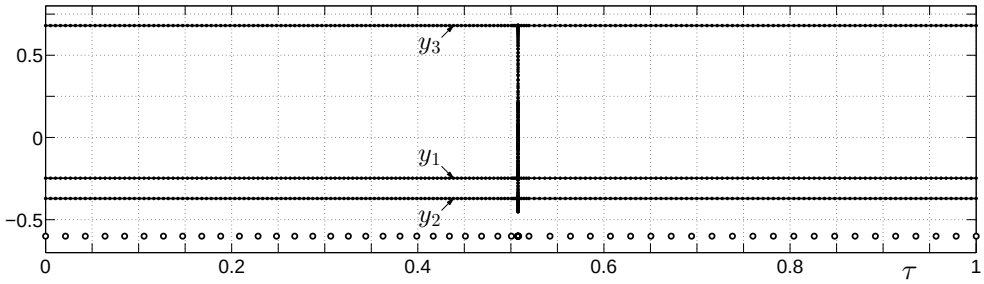

(a) Time profile of high-period orbit.

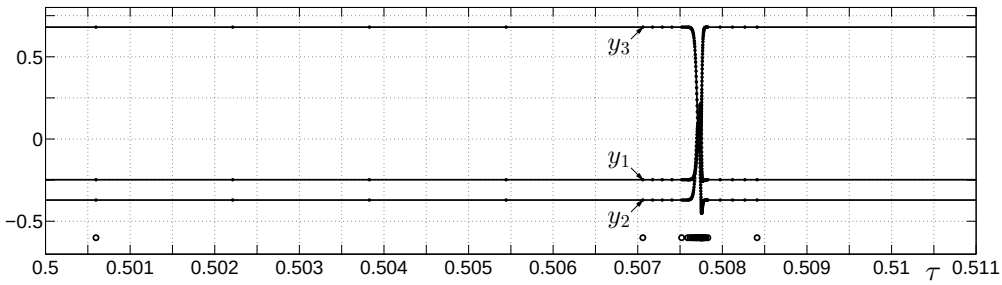

(b) First zoom into excursion from equilibrium.

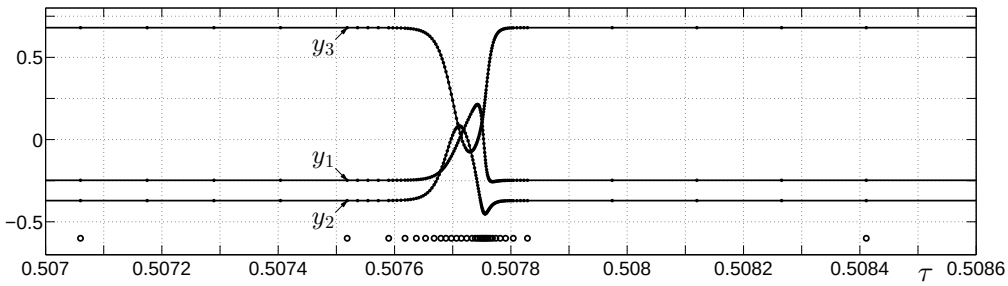

(c) Second zoom into excursion from equilibrium.

Figure 20.13. *Time profiles of the high-period orbit constructed in Fig. 20.11 on a moving mesh with variable discretization order shown over the full period (a) and in two subsequent zooms into the excursion from the equilibrium (b) and (c). The distribution of mesh points is indicated with circles at the bottom. A comparison between panel (a) and Fig. 20.10(a) indicates that the difference in discretization order of* 41 *mesh intervals is due mainly to a coarser mesh along the part with near-constant dynamics. The ratio between the largest and smallest mesh intervals,* $\max\{\kappa\}/\min\{\kappa\} = 9245$*, is here almost twice as large as for the moving mesh with fixed discretization order used in Fig. 20.10.*

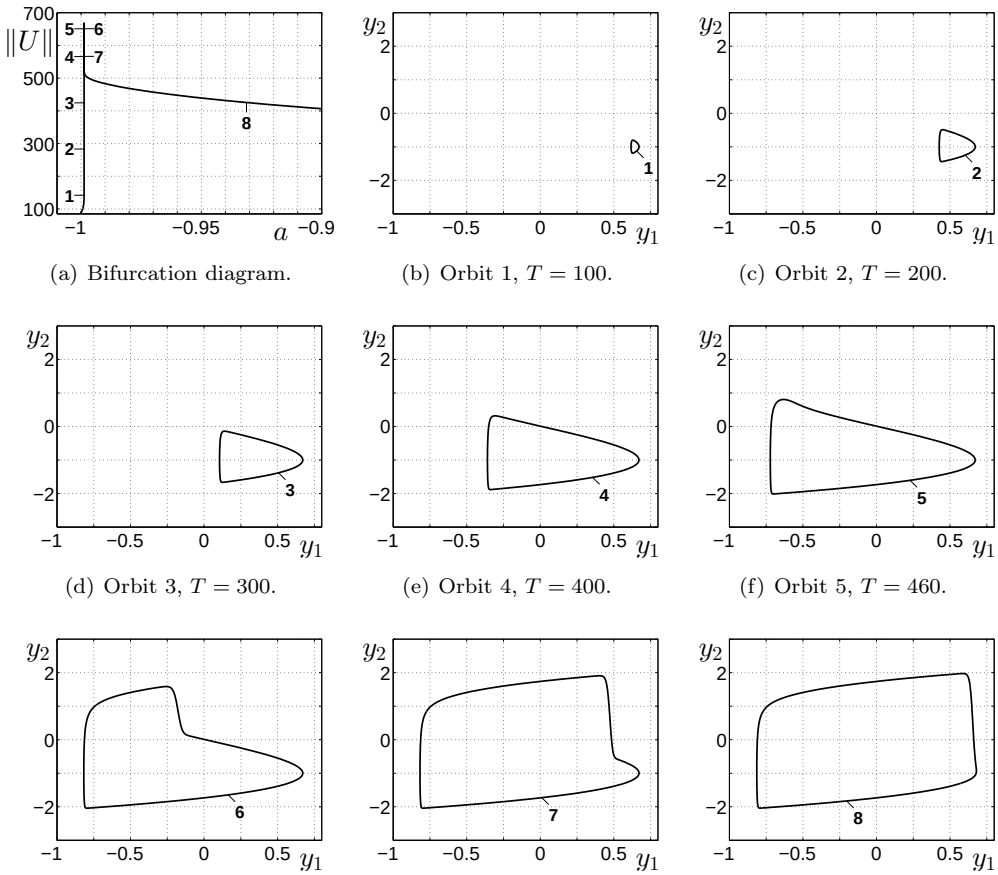

(a) Bifurcation diagram. (b) Orbit 1, $T = 100$. (c) Orbit 2, $T = 200$.

(d) Orbit 3, $T = 300$. (e) Orbit 4, $T = 400$. (f) Orbit 5, $T = 460$.

(g) Orbit 6, $T = 460$. (h) Orbit 7, $T = 400$. (i) Orbit 8, $T = 300$.

Figure 20.14. *A* canard explosion *in the Van der Pol dynamical system given by the vector field in Eq. (20.45). Here, a dramatic increase of amplitude (represented by the norm $\|U\| := \|(u, \mu_{\mathbb{J}})\|$) along a family of periodic orbits over an exceedingly small variation in a system parameter shows up as a virtually vertical branch in a bifurcation diagram (a). Panels (b) to (h) show selected members of the canard family. The orbits along the family were selected according to their period, indicated in the caption. We observe evidence for a fold point with respect to period between orbits 5 and 6.*

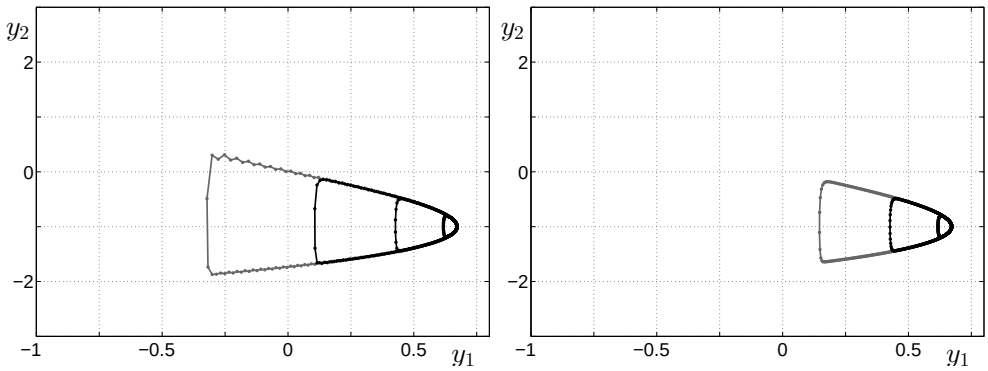

(a) Spectral method with adaptive order. (b) Collocation with uniform mesh.

Figure 20.15. *Selected approximate periodic orbits obtained during continuation along the canard family of the dynamical system given by the vector field in Eq. (20.45). Both the adaptive spectral method (a) from Sect. 19.3 with* `'NMAX'` *equal to* 100 *and the nonadaptive collocation method (b) from Sect. 18.1.2 with* `'NTST'` *and* `'NCOL'` *equal to* 100 *and* 4*, respectively, fail early on the canard family, given the desired error tolerance of* 10^{-2}.

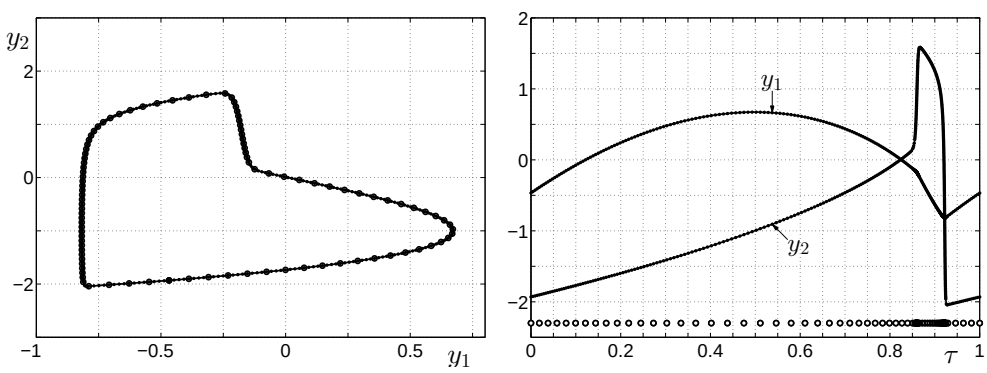

(a) State-space trajectory. (b) Time profile.

Figure 20.16. *Discretization of the canard orbit in Fig. 20.14(g) obtained using continuation with the comoving-mesh method from Sect. 18.3.1 on a mesh with* `'NTST'` *equal to* 100*,* `'NCOL'` *equal to 4,* `'TOL'` *equal to* 10^{-2}*, and* `'hfac'` *equal to 5. The fat dots in (a) and the open circles in (b) mark the end points of mesh intervals. The estimated discretization error of this solution is* 3.5326×10^{-3}.

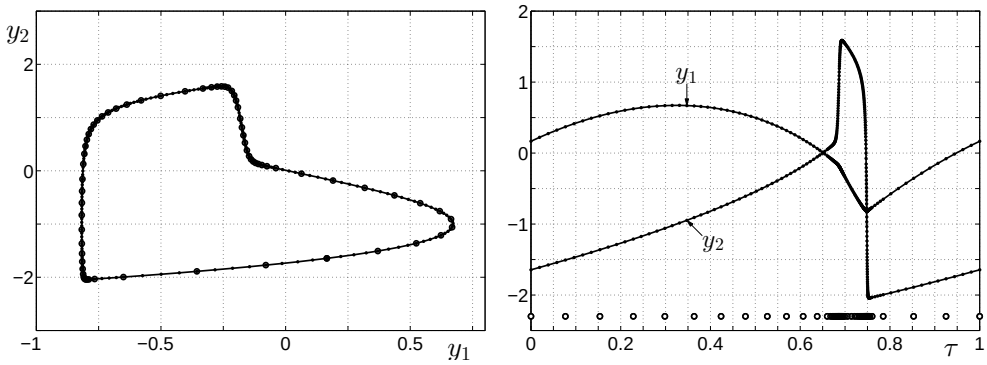

(a) State-space trajectory. (b) Time profile.

Figure 20.17. *Discretization of the canard orbit in Fig. 20.14(g) obtained using continuation with the moving-mesh method from Sect. 20.2.1 on a mesh with* `'NTST'` *equal to* 70, `'NCOL'` *equal to* 4, *and* `'TOL'` *equal to* 10^{-4}. *The fat dots in (a) and the open circles in (b) mark the end points of mesh intervals. The estimated error of this solution is* 5.1787×10^{-6}.

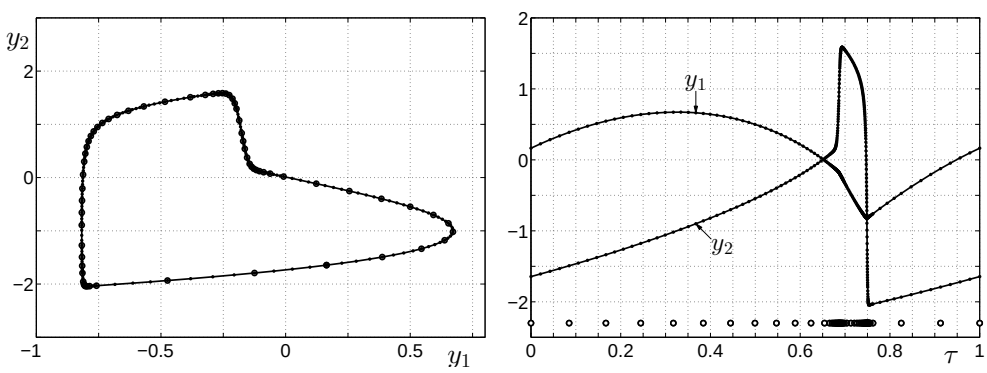

(a) State-space trajectory. (b) Time profile.

Figure 20.18. *Discretization of the canard orbit in Fig. 20.14(g) obtained using continuation with the moving-mesh method with adaptive discretization order from Sect. 20.2.2 on a mesh with* `'NTST'` *equal to* 66, `'NCOL'` *equal to* 4, *and* `'TOL'` *equal to* 10^{-4}. *The fat dots in (a) and the open circles in (b) mark the end points of mesh intervals. The estimated error of this solution is* 1.1882×10^{-5}.

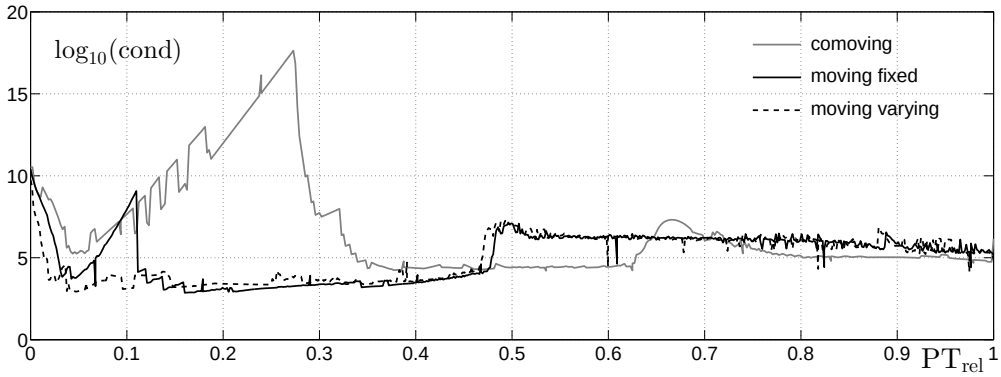

Figure 20.19. *The logarithm of the condition number of the Jacobian of the restricted continuation problem plotted against the point number normalized by the maximum point number during continuation of the canard family of the dynamical system given by the vector field in Eq. (20.45) using the (co)moving-mesh adaptive discretization strategies in Sects. 18.3.1, 20.2.1, and 20.2.2, respectively.*

This book provides a comprehensive introduction to the mathematical methodology of parameter continuation, the computational analysis of families of solutions to nonlinear mathematical equations. It develops a systematic formalism for constructing abstract representations of continuation problems and for implementing these in an existing computational platform.

Recipes for Continuation

- lends equal importance to theoretical rigor, algorithm development, and software engineering;
- demonstrates the use of fully developed toolbox templates for single- and multisegment boundary-value problems to the analysis of periodic orbits in smooth and hybrid dynamical systems, quasi-periodic invariant tori, and homoclinic and heteroclinic connecting orbits between equilibria and/or periodic orbits;
- shows the use of vectorization for optimal computational efficiency, an object-oriented paradigm for the modular construction of continuation problems, and adaptive discretization algorithms for guaranteed bounds on estimated errors; and
- contains extensive and fully worked examples that illustrate the application of the MATLAB®-based *Computational Continuation Core* (COCO) to problems from recent research literature that are relevant to dynamical system models from mechanics, electronics, biology, economics, and neuroscience.

A large number of exercises at the end of each chapter can be used as self-study or for course assignments that range from reflections on theoretical content to implementations in code of algorithms and toolboxes that generalize the discussion in the book or the literature. Open-ended projects throughout the book provide opportunities for summative assessments.

The book is intended for students and teachers of nonlinear dynamics and engineering, as well as engineers and scientists engaged in modeling and simulation, and is valuable to potential developers of computational tools for analysis of nonlinear dynamical systems. It assumes some familiarity with MATLAB programming and a theoretical sophistication expected of upper-level undergraduate or first-year graduate students in applied mathematics and/or computational science and engineering.

Harry Dankowicz is Professor of Mechanical Science and Engineering at the University of Illinois at Urbana-Champaign. He is author of a research monograph on chaos in Hamiltonian systems and a textbook on multibody mechanics, and serves as an Associate Editor of the *SIAM Journal on Applied Dynamical Systems*.

Frank Schilder has held post-doctoral research and teaching positions at the University of Bristol, the University of Surrey, and the Technical University of Denmark. In addition to COCO, he is the author of TORCONT and RAUTO and co-author, with Claudia Wulff and Andreas Schebesch, of SYMPERCON.

siam®

Society for Industrial and Applied Mathematics
3600 Market Street, 6th Floor
Philadelphia, PA 19104-2688 USA
+1-215-382-9800 • Fax: +1-215-386-7999
siam@siam.org • www.siam.org

CS11

ISBN 978-1-611972-56-6

90000

9781611972566