Studies in Systems, Decision and Control 304

Oliver Junge · Oliver Schütze ·
Gary Froyland · Sina Ober-Blöbaum ·
Kathrin Padberg-Gehle *Editors*

Advances in Dynamics, Optimization and Computation

A volume dedicated to Michael Dellnitz on the occasion of his 60th birthday

Springer

Studies in Systems, Decision and Control

Volume 304

Series Editor

Janusz Kacprzyk, Systems Research Institute, Polish Academy of Sciences, Warsaw, Poland

The series "Studies in Systems, Decision and Control" (SSDC) covers both new developments and advances, as well as the state of the art, in the various areas of broadly perceived systems, decision making and control–quickly, up to date and with a high quality. The intent is to cover the theory, applications, and perspectives on the state of the art and future developments relevant to systems, decision making, control, complex processes and related areas, as embedded in the fields of engineering, computer science, physics, economics, social and life sciences, as well as the paradigms and methodologies behind them. The series contains monographs, textbooks, lecture notes and edited volumes in systems, decision making and control spanning the areas of Cyber-Physical Systems, Autonomous Systems, Sensor Networks, Control Systems, Energy Systems, Automotive Systems, Biological Systems, Vehicular Networking and Connected Vehicles, Aerospace Systems, Automation, Manufacturing, Smart Grids, Nonlinear Systems, Power Systems, Robotics, Social Systems, Economic Systems and other. Of particular value to both the contributors and the readership are the short publication timeframe and the world-wide distribution and exposure which enable both a wide and rapid dissemination of research output.

** Indexing: The books of this series are submitted to ISI, SCOPUS, DBLP, Ulrichs, MathSciNet, Current Mathematical Publications, Mathematical Reviews, Zentralblatt Math: MetaPress and Springerlink.

More information about this series at http://www.springer.com/series/13304

Oliver Junge · Oliver Schütze ·
Gary Froyland · Sina Ober-Blöbaum ·
Kathrin Padberg-Gehle
Editors

Advances in Dynamics, Optimization and Computation

A volume dedicated to Michael Dellnitz on the occasion of his 60th birthday

Springer

Editors
Oliver Junge
Department of Mathematics
Technical University of Munich
Munich, Germany

Gary Froyland
School of Mathematics and Statistics
University of New South Wales
Sydney, NSW, Australia

Kathrin Padberg-Gehle
Institute of Mathematics and its Didactics
Leuphana University Lüneburg
Lüneburg, Germany

Oliver Schütze
Computer Science Department
Cinvestav-IPN
Mexico City, Distrito Federal, Mexico

Sina Ober-Blöbaum
Department of Mathematics
Paderborn University
Paderborn, Germany

ISSN 2198-4182 ISSN 2198-4190 (electronic)
Studies in Systems, Decision and Control
ISBN 978-3-030-51263-7 ISBN 978-3-030-51264-4 (eBook)
https://doi.org/10.1007/978-3-030-51264-4

This Springer imprint is published by the registered company Springer Nature Switzerland AG
The registered company address is: Gewerbestrasse 11, 6330 Cham, Switzerland

Für Michael

Preface

This book is a collection of recent advances on problems in dynamical systems, optimal control and optimization. In many cases, computational aspects and techniques are central. We dedicate this volume to Michael Dellnitz on the occasion of his 60th birthday. In one way or the other, there is a connection to Michael's research in all of these contributions.

In Part I, we collect chapters related to problems in dynamical systems. We start with a new technique for computing highly degenerate periodic orbits in ordinary differential equations, so-called phase resetting curves which appear, e.g., in models of spiking neurons. The second contribution employs concepts from singularity theory in order to classify and compute homeostasis points, i.e., points in state space of, e.g., biochemical networks, in which some output variable is roughly constant while some input variable is changing. We continue with a review on recent developments on the set-oriented approximation of invariant sets, a technique that has been pioneered by Michael and which will reappear in several other chapters in this volume. In the following paper, this technique is employed to approximate the transfer operator in non-autonomous differential equations, enabling the computation of coherent behavior in otherwise turbulent fluid flows. In high-dimensional systems with high-dimensional invariant sets, other concepts for an approximation of the transfer operator have to be found—and this is the subject of the next chapter where empirical bases are employed to construct a finite-rank approximation of this operator. Eigenfunctions of the (approximate) transfer operator can be used in order to detect, e.g., rare events like transitions between almost–invariant, respectively, metastable subsets in state space—an observation which already appears in one of the earlier works of Michael from the late 1990s. This is built upon in the subsequent chapter where a new, weaker characterization of slowly changing coordinates in noisy dynamical systems is proposed. Another way to address the reliable detection of rare events is based on sampling techniques like importance sampling and this addressed in the next chapter, where several sampling algorithms are compared and validated. We close the first part of this book by a chapter which demonstrates the usefulness of concepts from dynamical systems for solving questions on the computational complexity of certain problems.

Part II is dedicated to optimal control problems. In the first contribution, symmetric optimal control problems are investigated. Lie group symmetries and associated motion primitives of mechanical systems are exploited to develop numerical methods for multi-objective model predictive optimal control problems. In the second contribution, we review the numerical treatment of a mixed-integer optimal control problem governed by linear convection–diffusion equations and binary control variables. Relaxation and sum-up rounding techniques are combined with model order reduction to make the numerical approximation computationally more efficient. In the third contribution, we review set-oriented methods for the construction of globally optimal controllers. Based on a discrete version of Bellman's optimality principle applied to a dynamic game, a discrete feedback is constructed which robustly stabilizes a given nonlinear control system. In the last contribution, we review and highlight some connections between the problem of nonlinear smoothing and optimal control problems involving control of probability densities.

Finally, in Part III, we present three contributions related to optimization. The first contribution deals with the occurrence of "dents" in Pareto fronts of continuous multi-objective optimization problems. This can be helpful to obtain information about the structure of the Pareto front without explicitly computing the entire Pareto set. The second contribution deals with equality constrained bi-level multi-objective optimization problems and proposes a novel set-oriented algorithm that aims for a well-distributed finite-size approximation of the Pareto front of the higher-level problem. The third contribution reviews the gradient subspace approximation which allows one to compute descent directions in a best-fit manner from given neighborhood information. The method works particularly well in combination with set-oriented searchers such as evolutionary algorithms.

April 2020

Oliver Junge
Oliver Schütze
Gary Froyland
Sina Ober-Blöbaum
Kathrin Padberg-Gehle

Acknowledgements

The editors are grateful to the authors for their contributions to this volume recognizing Michael's scientific achievements. All editors have been either graduate students or postdocs with Michael and would like to express their appreciation for Michael's mentoring and their collaborative endeavors, both within and beyond Michael's research group.

O. Schütze acknowledges support from Conacyt Basic Science project no. 285599 and SEP-Cinvestav project no. 231.

Contents

Optimal Control

Optimization

List of Contributors

Fernando Antoneli Escola Paulista de Medicina, Universidade Federal de São Paulo, São Paulo, SP, Brazil

Andreas Bittracher Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, Germany

Meagan Carney Max Planck Institute for the Physics of Complex Systems, Dresden, Germany

Michael Dellnitz Chair of Applied Mathematics, Paderborn University, Paderborn, Germany

Alessandro Dell'Aere Chair of Applied Mathematics, Institute for Mathematics University of Paderborn, Paderborn, Germany

Kathrin Flaßkamp Systems Modeling and Simulation, Saarland University, Saarbrücken, Germany

Gary Froyland School of Mathematics and Statistics, University of New South Wales, Sydney, NSW, Australia

Raphael Gerlach Paderborn University, Paderborn, Germany

Martin Golubitsky Mathematics Department, The Ohio State University, Columbus, OH, USA

Lars Grüne Mathematical Institute, University of Bayreuth, Bayreuth, Germany

Zhengyuan Huang The Ohio State University, Columbus, OH, USA

Oliver Junge Department of Mathematics, Technical University of Munich, Munich, Germany

Christian Jäkle Department of Mathematics and Statistics, University of Konstanz, Konstanz, Germany

Holger Kantz Max Planck Institute for the Physics of Complex Systems, Dresden, Germany

Jin Won Kim Coordinated Science Laboratory and Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA

Stefan Klus Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, Germany

Anna Klünker Institute of Mathematics and its Didactics, Leuphana Universität Lüneburg, Lüneburg, Germany

Bernd Krauskopf Department of Mathematics, The University of Auckland, Auckland, New Zealand

Adriana Lara ESFM del Instituto Politécnico Nacional, Mexico City, Mexico

Peter Langfield IHU Liryc, Electrophysiology and Heart Modeling Institute, Fondation Bordeaux Université, Pessac - Bordeaux, France; Institut de Mathématiques de Bordeaux UMR 5251, Université de Bordeaux, Talence, France

Prashant G. Mehta Coordinated Science Laboratory and Department of Mechanical Science and Engineering, University of Illinois at Urbana-Champaign, Urbana, IL, USA

Mattes Mollenhauer Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, Germany

Mirko Hessel-von Molo Faculty of Computer Science Fachhochschule Dortmund – University of Applied Sciences and Arts, Dortmund, Germany

Matthew Nicol Department of Mathematics, University of Houston, Houston, TX, USA

Sina Ober-Blöbaum Department of Mathematics, Paderborn University, Paderborn, Germany

Hinke M. Osinga Department of Mathematics, The University of Auckland, Auckland, New Zealand

Kathrin Padberg-Gehle Institute of Mathematics and its Didactics, Leuphana Universität Lüneburg, Lüneburg, Germany

Sebastian Peitz Department of Mathematics, Paderborn University, Paderborn, Germany

Tuhin Sahai Raytheon Technologies Research Center, Berkeley, CA, USA

Christiane Schneide Institute of Mathematics and its Didactics, Leuphana Universität Lüneburg, Lüneburg, Germany

Jörg Schumacher Department of Mechanical Engineering, Technische Universität Ilmenau, Ilmenau, Germany

Ingmar Schuster Zalando Research, Zalando SE, Berlin, Germany

Christof Schütte Department of Mathematics and Computer Science, Freie Universität Berlin, Berlin, Germany; Zuse Institute Berlin, Berlin, Germany

Oliver Schütze Computer Science Department, Cinvestav-IPN, Mexico City, Mexico

Ian Stewart Mathematics Institute, University of Warwick, Coventry, UK

Lourdes Uribe ESFM del Instituto Politécnico Nacional, Mexico City, Mexico

Stefan Volkwein Department of Mathematics and Statistics, University of Konstanz, Konstanz, Germany

Yangyang Wang Department of Mathematics, The University of Iowa, Iowa City, IA, USA

Katrin Witting dSPACE GmbH, Paderborn, Germany

Adrian Ziessler Paderborn University, Paderborn, Germany

Dynamics

A Continuation Approach to Computing Phase Resetting Curves

Peter Langfield[1,2], Bernd Krauskopf[3], and Hinke M. Osinga[3(✉)]

[1] IHU Liryc, Electrophysiology and Heart Modeling Institute, Fondation Bordeaux Université, 33600 Pessac - Bordeaux, France
peter.langfield@u-bordeaux.fr
[2] Institut de Mathématiques de Bordeaux UMR 5251, Université de Bordeaux, 33400 Talence, France
[3] Department of Mathematics, The University of Auckland, Private Bag 92019, Auckland 1142, New Zealand
{b.krauskopf,h.m.osinga}@auckland.ac.nz

Abstract. Phase resetting is a common experimental approach to investigating the behaviour of oscillating neurons. Assuming repeated spiking or bursting, a phase reset amounts to a brief perturbation that causes a shift in the phase of this periodic motion. The observed effects not only depend on the strength of the perturbation, but also on the phase at which it is applied. The relationship between the change in phase after the perturbation and the unperturbed old phase, the so-called phase resetting curve, provides information about the type of neuronal behaviour, although not all effects of the nature of the perturbation are well understood. In this chapter, we present a numerical method based on the continuation of a multi-segment boundary value problem that computes phase resetting curves in ODE models. Our method is able to deal effectively with phase sensitivity of a system, meaning that it is able to handle extreme variations in the phase resetting curve, including resets that are seemingly discontinuous. We illustrate the algorithm with two examples of planar systems, where we also demonstrate how qualitative changes of a phase resetting curve can be characterised and understood. A seven-dimensional example emphasises that our method is not restricted to planar systems, and illustrates how we can also deal with non-instantaneous, time-varying perturbations.

1 Introduction

Measuring phase resetting is a common approach for testing neuronal responses in experiments: a brief current injection perturbs the regular spiking behaviour of a neuron, resulting generally in a shifted phase as the neuron returns to its regular oscillating behaviour. This phase shift can be advanced or delayed—meaning that the next spike arrives earlier or later compared with the unperturbed spiking oscillation—and which effect occurs also depends on the moment when the

O. Junge et al. (Eds.): SON 2020, SSDC 304, pp. 3–30, 2020.
https://doi.org/10.1007/978-3-030-51264-4_1

current is applied; see [9] for more details. A plot of the shifted phase ϑ_{new} versus the original phase ϑ_{old} at which the current was applied is known as the *phase transition curve* (PTC). Experimentally, it is often easier to represent the reset in terms of the resulting phase difference $\vartheta_{\text{new}} - \vartheta_{\text{old}}$ as a function of ϑ_{old}, which can be measured as the time to the next spike; such a representation is called a *phase response curve* or *phase resetting curve* (PRC).

The shape of a PTC or PRC of a given system obviously depends on the size of the applied perturbation: already for quite small amplitudes, nonlinear effects can dramatically affect a PTC or PRC. The shape of the PTC or PRC has been used to classify neuronal behaviour [1,7,16], where the underlying assumption is that the size of the applied perturbation is sufficiently small. Hodgkin [17] distinguished between so-called Type-I and Type-II excitable membranes, where neurons with membranes of Type II are not able to fire at arbitrarily low frequencies. Note that transitions from Type-I to Type-II can occur when system parameters are changed [8]. Ermentrout [7] found that the PRC of a Type-I neuron always has the same sign, while that of a Type-II neuron changes sign; this means that the PTC is always entirely above or below the diagonal for Type-I neurons, while it intersects the diagonal for Type-II neurons. In either case, the PTC is invertible for sufficiently small perturbation amplitudes, since it can be viewed as a continuous and smooth deformation of the identity, which is the PTC in the limit of zero amplitude. Invertibility itself has also been used as a distinguishing property of PTCs: noninvertible PTCs are said to be of type 0 (or strong) and invertible PTCs are of type 1 (or weak) [9,37]. If an increasingly stronger perturbation is applied, for example, in the context of synchronisation, it is well known that PTCs can change from type 1 to type 0, that is, become noninvertible [11,37].

A motivation in recent work on phase resetting has been the idea of interpreting the PTC as defining a one-dimensional phase-reduction model that, hopefully, captures the essential dynamics of a possibly high-dimensional oscillating system. The main interest is in coupled systems, formed by two or more (planar) systems with known PRCs; for example, see [31,32] for mathematical as well as experimental perspectives. Unfortunately, the convergence back to the limit cycle after some perturbation can be quite slow for a coupled system, such that only (infinitesimally) small perturbations are accurately described. Furthermore, it makes physiological sense to assume a time-varying input, usually in the form of a short input pulse, rather than the instantaneous perturbation assumed for the theoretical phase reset. Moreover, the perturbation may be repeated at regular intervals. In this context, PTCs and PRCs can be useful for explaining the resetting behaviour, though strictly speaking, the theory is only valid at low firing rates [15,35]. More recently, the idea of a phase-amplitude description has led to a better understanding of the effects resulting from these kinds of repeated time-varying resets [2,3,26,30,34].

From a dynamical systems perspective, the key question of phase resetting is how the perturbed initial conditions relax back to an attracting periodic orbit Γ with period T_Γ of an underlying continuous-time model, which we take here to be a vector field on $\mathbb{R}^n$, that is, a system of n first-order autonomous ordinary

differential equations. All points in its basin $\mathcal{B}(\Gamma)$ converge to Γ, and they do so with a given asymptotic phase. The subset of all points in $\mathcal{B}(\Gamma)$ that converge to Γ in phase with the point $\gamma_\vartheta \in \Gamma$, where $\vartheta \in [0,1)$ by convention, is called the (forward-time) *isochron* of γ_ϑ, which we refer to as $I(\gamma_\vartheta)$. Isochrons were defined and named by Winfree [36]. Guckenheimer [12] showed that $I(\gamma_\vartheta)$ is, in fact, an $(n-1)$-dimensional invariant stable manifold of the attracting fixed point $\gamma_\vartheta \in \Gamma$ under the time-T_Γ map. In particular, it follows that $I(\gamma_\vartheta)$ is tangent to the attracting linear eigenspace of γ_ϑ and, hence, transverse to Γ. Moreover, the ϑ-dependent family of all isochrons $I(\gamma_\vartheta)$ foliates the basin $\mathcal{B}(\Gamma)$. In other words, any point in $\mathcal{B}(\Gamma)$ has a unique asymptotic phase determined by the isochron it lies on.

For a given ϑ_{old}, consider now the perturbed point $\gamma_{\vartheta_{\text{old}}} + A\,\mathbf{d} \in \mathcal{B}(\Gamma)$, obtained from $\gamma_{\vartheta_{\text{old}}} \in \Gamma$ by applying the perturbation of strength A in the given direction $\mathbf{d}$. The asymptotic phase ϑ_{new} is, hence, uniquely determined by the isochron $I(\gamma_{\vartheta_{\text{new}}})$ on which this point lies. This defines a circle map $P : [0,1) \to [0,1)$ with $P(\vartheta_{\text{old}}) = \vartheta_{\text{new}}$. Therefore, finding the PTC is equivalent to determining how the perturbed cycle $\Gamma + A\,\mathbf{d} = \{\gamma_{\vartheta_{\text{old}}} + A\,\mathbf{d} \mid \vartheta_{\text{old}} \in [0,1)\}$ intersects the foliation of $\mathcal{B}(\Gamma)$ by the isochrons $I(\gamma_{\vartheta_{\text{new}}})$ for $\vartheta_{\text{new}} \in [0,1)$. Notice further that the PTC is the graph of the circle map P on the unit torus $\mathbb{T}^2$, represented by the unit square $[0,1) \times [0,1)$.

When considering the amplitude A of the perturbation as a parameter (while keeping the direction $\mathbf{d}$ fixed throughout), one can deduce some important properties of the associated PTC. Suppose that $0 < A_{\max}$ is such that $\Gamma_A := \Gamma + A\,\mathbf{d} \subset \mathcal{B}(\Gamma)$ for all $0 \le A < A_{\max}$. Then none of these perturbed cycles Γ_A intersects the boundary of the basin $\mathcal{B}(\Gamma)$ and the associated circle map $P = P_A$ is well defined for all $\vartheta_{\text{old}} \in [0,1)$. The map P_0 for zero perturbation amplitude is the identity on $\mathbb{T}^2$, which means that, as its graph, the PTC is the diagonal on $[0,1) \times [0,1)$ and a 1:1 torus knot on $\mathbb{T}^2$; in particular, P_0 is invertible, that is, it is injective and surjective. Because of smooth dependence on the amplitude A and the fact that P_A is a function over $[0,1)$, the PTC remains a 1:1 torus knot on $\mathbb{T}^2$ and P_A is surjective for all $0 \le A < A_{\max}$.

Since the isochrons are transverse to Γ, the circle map P_A is C^1-close to the identity, and hence, also injective, for sufficiently small A. As the graph of a near-identity transformation, the PTC is then strictly monotone, invertible, and hence, of type 1 (or weak) in the notation of [9,37]. While surjectivity is preserved, injectivity may be lost before $A = A_{\max}$ is reached. Indeed, the PTC is either invertible for all $0 \le A < A_{\max}$, or there is a maximal $0 < A_{\text{inv}} < A_{\max}$ such that P_A is invertible only for all $0 < A \le A_{\text{inv}}$. The loss of injectivity of P_A at $A = A_{\text{inv}}$ happens generically because of the emergence of an inflection point. For $0 \le A < A_{\max}$ this transition creates a local minimum and a local maximum of the PTC, which is now no longer invertible and so of type 0 (or strong) in the notation of [9,37]. As we will show, an inflection point of P_A corresponds to a cubic tangency between the perturbed cycle Γ_A and an isochron. Indeed, additional inflection points and, hence, local minima and maxima may appear at subsequent cubic isochron tangencies. Since P_A is a circle map, these must come

in pairs; hence, counting the number of its local maxima (or minima) would provide a further refinement of the notation of a type 0 (or strong) PTC.

The above discussion shows that, when the applied perturbation A is sufficiently weak, it suffices to consider only the linear approximation to the isochron family, which is given by the ϑ-family of stable eigenspaces of the time-T_Γ map for each ϑ. In practice, nonlinear effects are essential, especially when multiple time scales are present or the phase reset involves relatively strong perturbations. Isochrons are often highly nonlinear objects of possibly very complicated geometry [21,36]. While the geometric idea of isochrons determining the phase resets has been around since the mid 1970s, the practical implementation has proven rather elusive. In practice, it is not at all straightforward to compute the isochrons of a periodic orbit. In planar systems, when such isochrons are curves, three different approaches have been proposed, based on Fourier averages [23,25], a parametrisation formulated in terms of a functional equation [14,18], and continuation of solutions to a suitably posed two-point boundary value problem [21,29]. In principle, all three approaches generalise to higher-dimensional isochrons, but there are only few explicit examples [14,25].

From the knowledge of the isochron foliation of $\mathcal{B}(\Gamma)$, one can immediately deduce geometrically the phase resetting for perturbations of any strength and in any direction. However, already for planar and certainly for higher-dimensional systems, this is effectively too much information when one is after the PTC resulting from a perturbation in a fixed direction and with a specific amplitude. In essence, finding a PTC or PRC remains the one-dimensional problem of finding the asymptotic phase of all points on the perturbed cycle.

In this chapter, we show how this can be achieved with a multi-segment boundary value problem formulation. Specifically, we adapt the approach from [21,22] to set up the calculation of the circle map P_A by continuation, first in A from $A = 0$ for fixed ϑ_{old}, and then in $\vartheta_{\text{old}} \in [0,1)$ for fixed A. In this way, we obtain accurate numerical approximations of the PTC or PRC as continuous curves, even when the system shows strong phase sensitivity. The set-up is extremely versatile, and the direct computation of a PTC in this way does not require the system to be planar. We demonstrate our method with a constructed example going back to Winfree [37, Chapter 6], where we also show how injectivity is lost in a first cubic tangency of Γ_A with an isochron. The robustness of the method is then illustrated with the computation of a PTC of a perturbed cycle that cuts through a region of extreme phase sensitivity in the (planar) FithHugh–Nagumo system; in spite of very large derivatives due to this phase sensitivity, the PTC is computed accurately as a continuous curve. Our final example of a seven-dimensional system from [20] modelling a type of cardiac pacemaker cell shows that our approach also works in higher dimensions; this system also features phase sensitivity due to the existence of different time scales.

This chapter is organised as follows. In the next section, we provide precise details of the setting and explain the definitions used. Section 3 presents the numerical set-up for computing a resetting curve by continuation of a

multi-segment boundary value problem. We then discuss two planar examples in depth, which are both taken from [22]: a variation of Winfree's model in Sect. 4 and the FitzHugh–Nagumo system in Sect. 5. The third and higher-dimensional example from [20] is presented in Sect. 6. A summary of the results is given in Sect. 7, where we also discuss some consequences of our findings and directions of future research.

2 Basic Setting and Definitions

As mentioned in the introduction, we consider a dynamical system with an attracting periodic orbit Γ. For simplicity, we assume that the state space is $\mathbb{R}^n$ and consider the dynamical system

$$\dot{\mathbf{x}} = \mathbf{F}(\mathbf{x}), \tag{1}$$

where $\mathbf{F} : \mathbb{R}^n \to \mathbb{R}^n$ is at least once continuously differentiable. We assume that system (1) has an attracting periodic orbit Γ with period $T_\Gamma > 0$, that is,

$$\Gamma := \{\gamma(t) \in \mathbb{R}^n \mid 0 \le t \le T_\Gamma \text{ with } \gamma(T_\Gamma) = \gamma(0)\},$$

and T_Γ is minimal with this property. We associate a phase $\vartheta \in [0, 1)$ with each point $\gamma_\vartheta \in \Gamma$, defining $\gamma_\vartheta := \gamma(t)$ with $t = \vartheta T_\Gamma$. Here $\gamma_0 := \gamma(0)$ needs to be chosen, which is usually done by fixing it to correspond to a maximum in the first component. The (forward-time) isochron $I(\gamma_\vartheta)$ associated with $\gamma_\vartheta \in \Gamma$ is then defined in terms of initial conditions $\mathbf{x}(0)$ of forward trajectories $\mathbf{x} := \{\mathbf{x}(t) \in \mathbb{R}^n \mid t \in \mathbb{R}\}$ of system (1) that accumulate on Γ, namely, as

$$I(\gamma_\vartheta) := \{\mathbf{x}(0) \in \mathbb{R}^n \mid \lim_{k\to\infty} \mathbf{x}(k\, T_\Gamma) = \gamma_\vartheta \text{ with } k \in \mathbb{N}\}.$$

In other words, the trajectory $\mathbf{x}$ approaches Γ in phase with γ_ϑ. Note that $I(\gamma_\vartheta)$ is the stable manifold of the fixed point γ_ϑ of the time-T_Γ return map; in particular, this means that $I(\gamma_\vartheta)$ is of dimension $n - 1$ and tangent at γ_ϑ to the stable eigenspace $E(\gamma_\vartheta)$, which is part of the stable Floquet bundle of Γ [12,13]; we utilise this property when computing isochrons, and also when computing a PTC or PRC.

We are now ready to give formal definitions of the PTC and PRC; see also [9].

Definition 1 (Phase Transition Curve)
The phase transition curve or PTC associated with a perturbation of amplitude $A \ge 0$ in the direction $\mathbf{d} \in \mathbb{R}^n$ is the graph of the map $P : [0, 1) \to [0, 1)$ defined as follows. For $\vartheta \in [0, 1)$, the image $P(\vartheta)$ is the phase φ associated with the isochron $I(\gamma_\varphi)$ that contains the point $\gamma_\vartheta + A\,\mathbf{d}$ for $\gamma_\vartheta \in \Gamma$.

Definition 2 (Phase Response Curve)
The phase response curve or PRC associated with a perturbation of amplitude $A \ge 0$ in the direction $\mathbf{d} \in \mathbb{R}^n$ is the graph of the phase difference $\Delta(\vartheta) = P(\vartheta) - \vartheta \pmod 1$, where the map P is as above.

The definitions of the PTC and PRC are based on knowledge of the (forward-time) isochron $I(\gamma_\varphi)$ associated with a point $\gamma_\varphi \in \Gamma$. We previously designed an algorithm based on continuation of a two-point boundary value problem (BVP) that computes one-dimensional (forward-time and backward-time) isochrons of a planar system up to arbitrarily large arclengths [21,22,29]. Here, we briefly describe this algorithm in its simplest form, because this is useful for understanding the basic set-up, and for introducing some notation. The description is presented in the style that is used for implementation in the software package AUTO [4,5]. In particular, we consider a time-rescaled version of the vector field (1), which represents an orbit segment $\{\mathbf{x}(t) \mid 0 \le t \le T\}$ of (1) as the orbit segment $\{\mathbf{u}(t) \mid 0 \le t \le 1\}$ of the vector field

$$\dot{\mathbf{u}} = T\,\mathbf{F}(\mathbf{u}), \tag{2}$$

so that the total integration time T is now a parameter of the system.

We approximate $I(\gamma_0)$ as the set of initial points of orbit segments that end on the linear space $E(\gamma_0)$, the linearised isochron of $I(\gamma_0)$, close to γ_0 after integer multiples of the period T_Γ. These points are formulated as initial points $\mathbf{u}(0)$ of orbit segments $\mathbf{u}$ that end on $E(\gamma_0)$ at a distance η from γ_0; hence, η defines a one-parameter family of orbit segments. Each orbit segment in this family is a solution of system (2) with $T = k\,T_\Gamma$ for $k \in \mathbb{N}$; the corresponding boundary conditions are:

$$[\mathbf{u}(1) - \gamma_0] \cdot \mathbf{v}_0^\perp = 0, \tag{3}$$

$$[\mathbf{u}(1) - \gamma_0] \cdot \mathbf{v}_0 = \eta, \tag{4}$$

where $\mathbf{v}_0$ is the normalised vector that spans $E(\gamma_0)$ and $\mathbf{v}_0^\perp$ is perpendicular to it. Note that Γ itself, when starting from γ_0, is a solution to the two-point BVP (2)–(4) with $T = T_\Gamma$ and $\eta = 0$. This gives us a first solution to start the continuation for computing $I(\gamma_0)$. We fix $T = T_\Gamma$ and continue the orbit segment $\mathbf{u}$ in η up to a maximum prespecified tolerance $\eta = \eta_{\max}$. As the end point $\mathbf{u}(1)$ is pushed away from γ_0 along $E(\gamma_0)$, the initial point $\mathbf{u}(0)$ traces out a portion of $I(\gamma_0)$.

Once we reach $\eta = \eta_{\max}$, we can extend $I(\gamma_0)$ further by considering points that map to $E(\gamma_0)$ after one additional period, that is, after time $T = 2\,T_\Gamma$. We start the continuation with the orbit segment formed by concatenation of the final orbit segment with Γ; here, we rescale time such that this first orbit is again defined for $0 \le t \le 1$, we set $T = 2\,T_\Gamma$, and $\eta = 0$. Note that this orbit segment has a discontinuity at $t = \frac{1}{2}$, but it is very small and AUTO will automatically correct and close it as part of the first continuation step. This correction will cause a small shift in η away from 0, but η will still be much smaller than $\eta_{\max}$ (in absolute value). We can keep extending $I(\gamma_0)$ further in this way, by continuation with $T = k\,T_\Gamma$, for integers $k > 2$. See [21,29] for more details on the implementation and, in particular, see [19,29] for details on how to find $E(\gamma_0)$ represented by the first vector $\mathbf{v}_0$ in the stable Floquet bundle of Γ.

The computational set-up forms a well-posed two-point BVP with a one-parameter solution family that can be found by continuation, provided the equality $\mathsf{NDIM} - \mathsf{NBC} + \mathsf{NPAR} = 1$ holds for the dimension NDIM of the problem, the

number NBC of boundary conditions, and the number NPAR of free parameters. Indeed, for the computation of $I(\gamma_0)$, we have NDIM = 2, because we assumed that the system is planar; NBC = 2, namely, one condition to restrict $\mathbf{u}(1)$ to the linearised isochron of $I(\gamma_0)$, and one condition to fix its distance to γ_0; and NPAR = 1, because we free the parameter η.

To compute $I(\gamma_\varphi)$ for other $\varphi \in [0, 1)$, this same approach can be used, working with a shifted periodic orbit Γ so that its head point is γ_φ, and determining the associated direction vector $\mathbf{v}_\varphi$ that spans the eigenspace $E(\gamma_\varphi)$ to which $I(\gamma_\varphi)$ is tangent. In [29], approximations of γ_φ and $\mathbf{v}_\varphi$ are obtained by interpolation of the respective mesh discretisations from AUTO. We describe an alternative approach in [21], where we consider $I(\gamma_\varphi)$ as the set of initial points of orbit segments that end in the linear space $E(\gamma_0)$ of $I(\gamma_0)$ sufficiently close to γ_0 after total integration time $T = k\,T_\Gamma + (1 - \varphi)\,T_\Gamma$.

For the computation of a phase resetting curve, we use a combination of these two approaches, but rather than interpolation, we shift the periodic orbit by imposing a separate two-point BVP. More precisely, we set up a multi-segment BVP comprised of several subsystems of two-point BVPs; the set-up for this extended BVP is explained in detail in the next section.

3 Algorithm for Computing a Phase Resetting Curve

Based on the definition of PTC and PRC, one could now calculate a sufficiently large number of isochrons and determine the resetting curve numerically from data. We prefer to compute the PTC or PRC directly also with a BVP set-up and continuation. The major benefit of such a direct approach is that it avoids accuracy restrictions arising from the selection of computed isochrons; in particular, any phase sensitivity of the PTC or PRC will be dealt with automatically as part of the pseudo-arclength continuation with AUTO [4,5].

For ease of presentation, we will formulate and discuss our continuation set-up for the case of a planar system. We remark, however, that it can readily be extended for use in $\mathbb{R}^n$ with $n > 2$, because the dimensionality of the problem is not determined by the dimension $n - 1$ of the isochrons but by the dimension of the PTC or PRC, which is always one; see also the example in Sect. 6.

The essential difference between calculating a resetting curve rather than an isochron is the following: for an isochron $I(\gamma_\vartheta)$, we compute orbit segments with total integration time $T = T_\Gamma$ (or integer multiples), where we move the end point $\mathbf{u}(1)$ along the linear approximation of $I(\gamma_\vartheta)$ to some distance η from Γ, while the initial point $\mathbf{u}(0)$ traces out a new portion of $I(\gamma_\vartheta)$; imagining the same set-up, if we move $\mathbf{u}(0)$ transverse to $I(\gamma_\vartheta)$, the end point $\mathbf{u}(1)$ will move to lie on the linearisation of an isochron $I(\gamma_\varphi)$ with a different phase φ. (Here, one should expect that the distance to Γ also changes, but we assume it is still less than $\eta_{\max}$). The key idea behind our approach is that we find a way to determine the different phase φ, or the phase shift $\varphi - \vartheta$, by allowing Γ and its corresponding stable Floquet bundle to rotate as part of an extended system. We ensure the head point of Γ moves with the phase-shifted point, that is, the first

point on Γ will be γ_φ. In this way, we can determine the shifted phase φ along any prescribed arc traced out by $\mathbf{u}(0)$, provided it lies in the basin of attraction of Γ. For the PTC or PRC associated with a perturbation of amplitude $A \geq 0$ in the direction $\mathbf{d} \in \mathbb{R}^n$, this prescribed arc is the perturbed cycle $\Gamma + A\,\mathbf{d}$, that is, $\mathbf{u}(0)$ traces out the closed curve $\{\gamma_\vartheta + A\,\mathbf{d} \mid \vartheta \in [0, 1)\}$.

3.1 Continuation Set-Up for Rotated Representation of Γ

We formulate an extended BVP that represents a rotated version of Γ with a particular phase, meaning that we automatically determine the phase of the head point relative to γ_0. To this end, we assume that the zero-phase point $\gamma_0 \in \Gamma$ and its associated linear vector $\mathbf{v}_0$, or more practical, its perpendicular $\mathbf{v}_0^\perp$, are readily accessible as stored parameters, or constants that do not change. Hence, even when Γ is rotated and its first point is γ_φ for some different $\varphi \in [0, 1)$, we can still access the coordinates of γ_0 and $\mathbf{v}_0^\perp$ from the parameter/constants list.

The extended BVP consists of three components, one to define Γ, one to define the associated (rotated) linear bundle, and one to define the associated phase. We start by representing Γ as a closed orbit segment $\mathbf{g}$ that solves system (2) for $T = T_\Gamma$. Hence, we define

$$\dot{\mathbf{g}} = T_\Gamma\, \mathbf{F}(\mathbf{g}), \tag{5}$$

with periodic boundary condition

$$\mathbf{g}(1) - \mathbf{g}(0) = 0. \tag{6}$$

The stable Floquet bundle of Γ is coupled with the BVP (5)–(6) via the first variational equation. More precisely, we consider a second orbit segment $\mathbf{v}_\mathbf{g}$, such that each point $\mathbf{v}_\mathbf{g}(t)$ represents a vector associated with points $\mathbf{g}(t)$ of the orbit segment that solves (5). The orbit segment $\mathbf{v}_\mathbf{g}$ is a solution to the linearised flow such that $\mathbf{v}_\mathbf{g}(0)$ is mapped to itself after one rotation around Γ. The length of $\mathbf{v}_\mathbf{g}(0)$ is contracted after one rotation by the factor $\exp(T_\Gamma\, \lambda_\mathrm{s})$, which is the stable Floquet multiplier of Γ. We prefer formulating this in logarithmic form, which introduces the stable Floquet exponent λ_s to the first variational equation, rather than affecting the length of $\mathbf{v}_\mathbf{g}(0)$. Therefore, the BVP (5)–(6) is extended with the following system of equations:

$$\dot{\mathbf{v}}_\mathbf{g} = T_\Gamma\, \left[\mathrm{D}_\mathbf{g}\mathbf{F}(\mathbf{g})\, \mathbf{v}_\mathbf{g} - \lambda_\mathrm{s}\, \mathbf{v}_\mathbf{g}\right], \tag{7}$$

$$\mathbf{v}_\mathbf{g}(1) - \mathbf{v}_\mathbf{g}(0) = 0, \tag{8}$$

$$\|\, \mathbf{v}_\mathbf{g}(0) \,\| = 1. \tag{9}$$

In particular, $\mathbf{v}_\mathbf{g}(0) = \mathbf{v}_\mathbf{g}(1)$ is the normalised vector that spans the local linearised isochron associated with $\mathbf{g}(0)$.

We have not specified a phase condition and, indeed, we allow $\mathbf{g}$ to shift and start at any point $\gamma_\vartheta \in \Gamma$. Consequently, the linear bundle $\mathbf{v}_\mathbf{g}$ will also shift such that $\mathbf{v}_\mathbf{g}(0)$ still spans the local linearised isochron associated with $\mathbf{g}(0)$.

Phase shifting the periodic orbit and its linear bundle by continuation in this way has been performed before [8]. However, the implementation in [8] requires accurate knowledge of the coordinates of the point γ_ϑ in order to decide when to stop shifting. Our approach uses another BVP set-up to monitor the phase shift, so that both γ_ϑ and $\mathbf{v}_\vartheta$ are determined up to the accuracy of the computation. To this end, we introduce a third orbit segment $\mathbf{w}$ that lies along Γ, with initial point $\mathbf{w}(0)$ equal to $\mathbf{g}(0)$, and end point $\mathbf{w}(1)$ equal to γ_0. The total integration time associated with this orbit segment $\mathbf{w}$ is the fraction of the period T_Γ that $\mathbf{g}(0)$ lies away from γ_0 along Γ; hence, it is directly related to the phase of $\mathbf{g}(0)$. We extend the BVP (5)–(9) with the following system of equations:

$$\dot{\mathbf{w}} = \nu\, T_\Gamma\, \mathbf{F}(\mathbf{w}), \tag{10}$$

$$\mathbf{w}(0) = \mathbf{g}(0), \tag{11}$$

$$[\mathbf{w}(1) - \gamma_0] \cdot \mathbf{v}_0^\perp = 0. \tag{12}$$

Here, we do not impose $\mathbf{w}(1) = \gamma_0$. Instead, condition (12) allows $\mathbf{w}(1)$ to move in the linearisation of $I(\gamma_0)$ at γ_0; this relaxation is necessary to ensure that the BVP remains well posed and the discretised problem has a solution. In practice, since $\mathbf{w}(0) \in \Gamma$, the difference between $\mathbf{w}(1)$ and γ_0 will be of the same order as the overall accuracy of the computation. Note that it is important to ensure $\nu \geq 0$ in Eq. (10), because $\mathbf{w}(1)$ may diverge from γ_0 along $E(\gamma_0)$ otherwise. We found it convenient to start the calculation with $\nu = 1$, which corresponds to the orbit segment $\mathbf{w} = \mathbf{g}$.

The combined solution $\{\mathbf{g}, \mathbf{v_g}, \mathbf{w}\}$ to the multi-segment BVP (5)–(12) represents a rotated version of Γ and its stable Floquet bundle so that the head point is γ_φ with phase $\varphi = 1 - \nu \,(\mathrm{mod}\ 1)$. We remark here that this extended set-up can also be used to compute $I(\gamma_\varphi)$, for any phase $0 < \varphi < 1$, with the method for $I(\gamma_0)$ described in Sect. 2; such a computation would approximate each isochron up to the same accuracy, without introducing an additional interpolation error.

3.2 Continuation Set-Up for the Phase Reset

Recall the set-up for computing a phase reset by moving $\mathbf{u}(0)$ transversely to $I(\gamma_\vartheta)$, so that the end point $\mathbf{u}(1)$ will move and lie on the linearisation of an isochron $I(\gamma_\varphi)$ with a different phase φ. Here, the orbit segment $\mathbf{u}$ is a solution of

$$\dot{\mathbf{u}} = k\, T_\Gamma\, \mathbf{F}(\mathbf{u}), \tag{13}$$

for some $k \in \mathbb{N}$. The end point $\mathbf{u}(1)$ should lie close to Γ on the linearisation of $I(\gamma_\varphi)$, for some $\varphi \in [0, 1)$. We stipulate that the rotated version of Γ is shifted such that $\mathbf{u}(1)$ lies close to $\mathbf{g}(0)$ along the direction $\mathbf{v_g}(0)$. Hence, we require the two boundary conditions

$$[\mathbf{u}(1) - \mathbf{g}(0)] \cdot \mathbf{v_g}(0) = \eta, \tag{14}$$

$$[\mathbf{u}(1) - \mathbf{g}(0)] \cdot \mathbf{v_g}(0)^\perp = 0, \tag{15}$$

where $\mathbf{v_g}(0)^\perp$ is the vector perpendicular to $\mathbf{v_g}(0)$. Here, η measures the (signed) distance between $\mathbf{u}(1)$ and $\mathbf{g}(0)$, which is along $\mathbf{v_g}(0)$. Since $\mathbf{u}$ is a solution of (13) and $k \in \mathbb{N}$, the initial point $\mathbf{u}(0)$ has the same phase as the last point $\mathbf{u}(1)$, and the combined multi-segment BVP (5)–(15) ensures that $\mathbf{u}(1)$ has (approximate) phase $1 - \nu \,(\mathrm{mod}\, 1)$. In practice, we should choose $k \in \mathbb{N}$ large enough such that $\eta < \eta_{\max}$. If $\mathbf{u}(0)$ lies close to Γ, it will be sufficient to set $k = 1$. In order to consider phase resets of large perturbations, for which $\mathbf{u}(0)$ starts relatively far away, we need $k > 1$, to allow for sufficient time to let $\mathbf{u}$ converge and have $\mathbf{u}(1)$ lie close to Γ.

At this stage, the multi-segment BVP (5)–(15) is a system of $\mathsf{NDIM} = 8$ ordinary differential equations (for the case of a planar system), with $\mathsf{NBC} = 10$ boundary conditions, and $\mathsf{NPAR} = 4$ free parameters, namely, T_Γ, λ_s, ν, and η; the period T_Γ and stable Floquet exponent λ_s must remain free parameters to ensure that the discretised problem has a solution, but their variation will be almost zero. Hence, $\mathsf{NDIM} - \mathsf{NBC} + \mathsf{NPAR} = 2 \neq 1$, and one more condition is needed to obtain a one-parameter family of solutions.

The final step in the set-up is to impose an extra condition that specifies how $\mathbf{u}(0)$ moves along an arc or closed curve in the phase plane. Consequently, since $k\,T_\Gamma$ is fixed, the orbit segment $\mathbf{u}$ changes, so that $\mathbf{u}(1)$ will move as well, and $\mathbf{g}(0)$, along with $\mathbf{v_g}(0)$ will shift accordingly. This causes a variation in ν to maintain $\mathbf{w}(0) = \mathbf{g}(0)$, and these ν-values precisely define the new phase in the continuation run as a function of the position along the chosen arc or closed curve.

To compute the PRC, we need to let $\mathbf{u}(0)$ traverse the closed curve $\{\gamma_\vartheta + A\,\mathbf{d} \mid \vartheta \in [0, 1)\}$ obtained by the (instantaneous) perturbation of Γ in the direction $\mathbf{d}$ for distance A. We can impose this relatively complicated path on $\mathbf{u}(0)$ by including another system of equations to the multi-segment BVP, namely, the BVP that defines Γ in terms of another rotated orbit segment $\mathbf{g_u}$. Furthermore, in order to keep track of the phase ϑ along this path, we introduce another segment $\mathbf{w_u}$ that plays the same role as $\mathbf{w}$ in Sect. 3.1; compare with equations (5)–(6) and (10)–(12). In other words, we extend the BVP (5)–(15) by the following system of equations

$$\dot{\mathbf{g}}_\mathbf{u} = \widehat{T}_\Gamma\, \mathbf{F}(\mathbf{g_u}), \tag{16}$$

$$\mathbf{g_u}(1) - \mathbf{g_u}(0) = 0. \tag{17}$$

$$\dot{\mathbf{w}}_\mathbf{u} = (1 - \vartheta)\, \widehat{T}_\Gamma\, \mathbf{F}(\mathbf{w_u}), \tag{18}$$

$$\mathbf{w_u}(0) = \mathbf{g_u}(0), \tag{19}$$

$$[\mathbf{w_u}(1) - \gamma_0] \cdot \mathbf{v}_0^\perp = 0. \tag{20}$$

Here, we decrease ϑ from 1 to 0, during which $\mathbf{w_u}$ grows and $\mathbf{g_u}$ tracks γ_ϑ. In order for a solution to exist, the periods T_Γ and $\widehat{T}_\Gamma$ must be two different free parameters, although they remain constant (and equal) to within the accuracy of the computation. The phase reset is now obtained by imposing

$$\mathbf{u}(0) = \mathbf{g_u}(0) + A\,\mathbf{d}. \tag{21}$$

The multi-segment BVP (5)–(21) is now a system of dimension NDIM = 12, with NBC = 17 boundary conditions, and NPAR = 6 free parameters, which are T_Γ, λ_s, ν, η, $\widehat{T}_\Gamma$, and either ϑ or A. Since, NDIM − NBC + NPAR = 1, we obtain a one-parameter solution family by continuation. As the first solution in the continuation, we use the known solution $\mathbf{g} = \mathbf{w} = \mathbf{u} = \mathbf{g_u} = \Gamma$, which starts with the head point γ_0, the associated stable linear bundle $\mathbf{v}_0$ that we assumed has been pre-computed, and $\mathbf{w_u} = \gamma_0$; then $T_\Gamma = \widehat{T}_\Gamma$ and λ_s are set to their known computed values, $\eta = A = 0$, and $\nu = \vartheta = 1$. Initially, $k = 1$, and one should monitor η to make sure it does not exceed $\eta_{\max}$.

To obtain the PTC or PRC we first perform a homotopy step, where we fix $\vartheta = 1$ and vary the amplitude A until the required value is reached. This continuation run produces a one-parameter family of solutions representing the effect of a reset of varying amplitude A from the point γ_0. In the main continuation run, we then fix A and decrease ϑ until $\vartheta = 0$, so that it covers the unit interval; the associated solution family of the multi-segment BVP (5)–(21), hence, provides the resulting phase $\vartheta_{\text{new}} := 1 - \nu \,(\text{mod } 1)$ as a function of the phase $\vartheta_{\text{old}} := \vartheta$ along the perturbed periodic orbit.

4 Illustration of the Method with a Model Example

We illustrate our method for computing a PTC with a constructed example, namely, a parametrised version of the model introduced by Winfree [37, Chapter 6], which we also used in [22]; it is given in polar coordinates as

$$\begin{cases} \dot{r} = (1-r)\,(r-a)\,r, \\ \dot{\psi} = -1 - \omega\,(1-r). \end{cases}$$

In Euclidean coordinates, the system becomes

$$\begin{cases} \dot{x} = (1-\sqrt{x^2+y^2})\,\left(x\,(\sqrt{x^2+y^2}-a)+\omega y\right)+y, \\ \dot{y} = (1-\sqrt{x^2+y^2})\,\left(y\,(\sqrt{x^2+y^2}-a)-\omega x\right)-x. \end{cases} \tag{22}$$

Note that this system is invariant under any rotation about the origin; moreover, its frequency of rotation only depends on $r = \sqrt{x^2+y^2}$; see [22] for details. We now fix the parameters to $a = 0$ and $\omega = -0.5$, as in [22]. Then the unit circle is an attracting periodic orbit Γ with period $T_\Gamma = 2\pi$ and the origin is an unstable equilibrium $\mathbf{x}^*$.

4.1 Computing the PTC

We choose $\gamma_0 = (1, 0)$ and compute the normalised linear direction associated with its isochron as $\mathbf{v}_0 \approx (-0.83, -0.55)$. As was explained in Sect. 3.2, the computation is performed in two separate continuation runs: first, we apply a perturbation to the point γ_0 in a fixed direction $\mathbf{d}$, where we vary the amplitude A from 0 to 0.75 during the homotopy step. Next, we fix $A = 0.75$ and apply

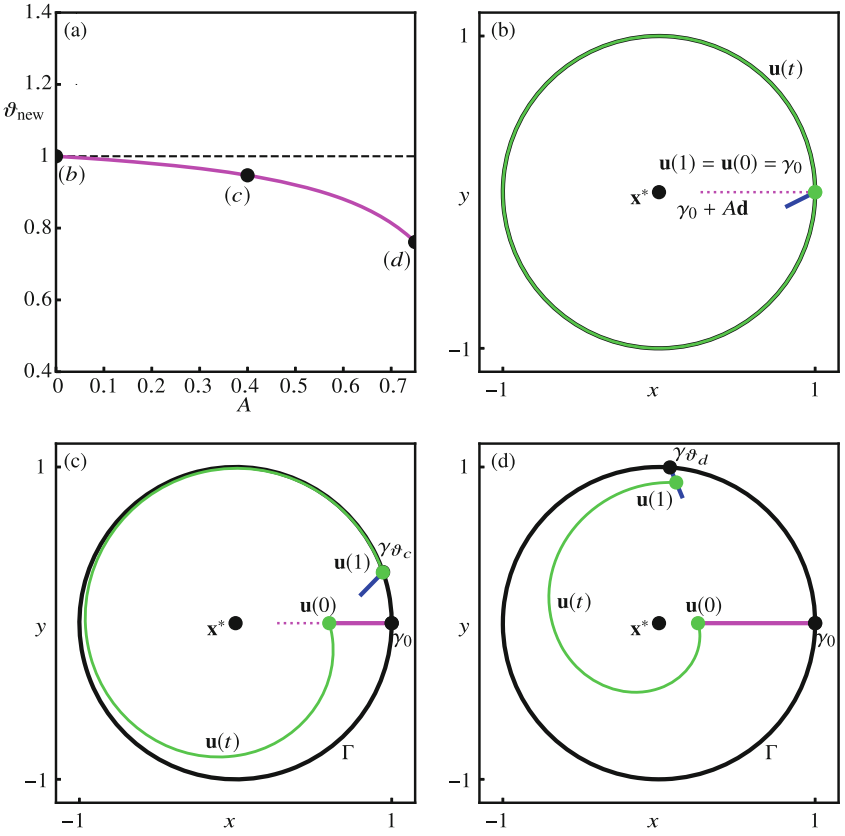

Fig. 1. Phase reset of system (22) at fixed γ_0 in the direction $\mathbf{d} = (-1, 0)$ with amplitude $A \in [0, 0.75]$ (a), and continuation set-up at the three labelled points (b), (c), and (d). Panel (b) shows the initial set-up when $A = 0$ and $\vartheta_{\text{new}} = 1$, in panel (c) the continuation has progressed to $A = 0.4$ and $\vartheta_{\text{new}} = \vartheta_c \approx 0.96$, and in panel (d) $A = 0.75$ has been reached and $\vartheta_{\text{new}} = \vartheta_d \approx 0.76$.

the same perturbation to each point $\gamma_\vartheta \in \Gamma$. For the purpose of visualising the computational set-up, we choose the (somewhat unusual) direction $\mathbf{d} = (-1, 0)$ and set the maximum distance along the linearised isochron to the relatively large value of $\eta_{\max} = 0.2$.

The first continuation run of the multi-segment BVP (5)–(21) is illustrated in Fig. 1. Here, the free amplitude A increases while $\vartheta = 1 = 0 \,(\text{mod } 1)$ is fixed and, hence, the perturbation is always applied at γ_0 and grows in size. Figure 1(a) shows the resulting phase ϑ_{new} as a function of A. Three points are labelled, indicating the three stages during the continuation that are illustrated in panels (b), (c) and (d). In each of these panels we show the periodic orbit Γ in black, and the current orbit segment $\mathbf{u}$ of the continuation run in green. Note that Γ is rotated here and its head point $\mathbf{g}(0)$ lies at the point on Γ with phase ϑ_{new}. A

short segment of the associated linearisation of the isochron of $\gamma_{\vartheta_{\text{new}}}$ is shown in blue. We do not plot the orbit segment $\mathbf{w}$ that determines the value of ϑ_{new}, but it follows Γ from $\mathbf{g}(0)$ back to $\mathbf{g}(0)$ and then extends (approximately) along Γ to γ_0. Indeed, notice in Fig. 1(a) that ϑ_{new} is decreasing, which means that $\nu > 1$ is increasing so that $\mathbf{w}$ becomes longer. We also do not show the orbit segments $\mathbf{g_u}$ and $\mathbf{w_u}$ that determine the phase $\vartheta = \vartheta_{\text{old}}$ at which the perturbation is applied, because $\vartheta_{\text{old}} = 1$ is fixed in this continuation run.

Figure 1(b) shows the initial set-up, with $\mathbf{g} = \mathbf{w} = \mathbf{u} = \mathbf{g_u} = \Gamma$, $\mathbf{w_u} = \gamma_0$, $T_\Gamma = \widehat{T}_\Gamma$ and λ_s set to their known values, and $\nu = 1$, $\eta = A = 0$, with $k = 1$ and $\vartheta = 1$. The dotted line segment in Fig. 1(b) indicates the direction $\mathbf{d}$ of the intended perturbation away from γ_0; its length is the maximal intended amplitude $A = 0.75$. An intermediate continuation step when $A = 0.4$ is shown in Fig. 1(c). The perturbation has pushed $\mathbf{u}(0)$ out along $\mathbf{d}$, such that $\mathbf{u}(1)$ now lies (approximately) on the linearised isochron, parametrised as $\mathbf{g}(0) + \eta\, \mathbf{v_g}(0)$ with $0 < \eta \leq \eta_{\max}$, associated with the rotated head point $\mathbf{g}(0) = \gamma_{\vartheta_c}$, where $\vartheta_c \approx 0.96$. Note that the orbit segment $\mathbf{w}$ (not shown) has now changed from its initialisation to match the solution to subsystem (10)–(12) with $\nu \approx 1.04$. Figure 1(d) illustrates the last step of the first continuation run, when $A = 0.75$. The head point $\mathbf{g}(0) \in \Gamma$ has rotated further to γ_{ϑ_d} with $\vartheta_d = 1 - \nu \approx -0.24 = 0.76 \,(\text{mod } 1)$. Notice that $\mathbf{u}(1)$ lies quite far along the linearised isochron, because we allow a relatively large distance η. The corresponding orbit segment $\mathbf{u}$ is determined for an integration time of only one period, that is, for $k = 1$. We show this case for illustration purposes, but in practice, it would be worth choosing a smaller value for $\eta_{\max}$, so that $\mathbf{u}$ would be extended, and the integer multiple of T_Γ set to $k = 2$, before reaching $A = 0.75$.

The second continuation run uses the fixed perturbation of size $A = 0.75$ along $\mathbf{d} = (-1, 0)$, and varies the phase ϑ at which it is applied. Since ϑ controls the integration time associated with the orbit segment $\mathbf{w_u}$, the multi-segment BVP (16)–(20) with solution $\{\mathbf{g_u}, \mathbf{w_u}\}$ and parameter $\widehat{T}_\Gamma$ now plays an important role. For each ϑ, the head point $\mathbf{g_u}(0)$ of $\mathbf{g_u}$ lies (approximately) at $\gamma_\vartheta \in \Gamma$, and $\mathbf{w_u}$ represents the remaining part of Γ from γ_ϑ to γ_0; hence, the total integration time of $\mathbf{w_u}$ is the fraction $1 - \vartheta$ of $\widehat{T}_\Gamma$, which is equal, up to the computational accuracy, to the period T_Γ of Γ.

Figure 2 illustrates different aspects of this continuation run. As $\vartheta_{\text{old}} = \vartheta$ decreases from 1, the multi-segment BVP (5)–(21) determines the orbit segment $\mathbf{u}$ with $\mathbf{u}(0) = \gamma_\vartheta + A\,\mathbf{d}$ and uses the rotated orbit segment $\mathbf{g}$ and $\mathbf{w}$ to establish the resulting phase $\vartheta_{\text{new}} = 1 - \nu \,(\text{mod } 1)$ of $\mathbf{u}(1)$. Panel (a) shows the PTC computed for $A = 0.75$. Note that ν takes values in the covering space $\mathbb{R}$; the output is then folded onto the unit torus by taking $\vartheta_{\text{new}} = 1 - \nu \,(\text{mod } 1)$, giving the solid curve in Fig. 2. The points labelled (b) and (c) in this panel correspond to $\vartheta_{\text{old}} = 0.9$ and $\vartheta_{\text{old}} = 0.1$, respectively. The continuation set-up for these two cases is shown in the corresponding panels (b) and (c). As in Fig. 1, the periodic orbit Γ is black and $\mathbf{u}$ is green. The path traced by the initial point $\mathbf{u}(0)$ is the magenta dotted circle, which is Γ shifted by $A = 0.75$ in the direction $\mathbf{d} = (-1, 0)$; hence, for fixed ϑ, the point $\mathbf{u}(0)$ corresponds to the perturbation of the point $\gamma_\vartheta \in \Gamma$

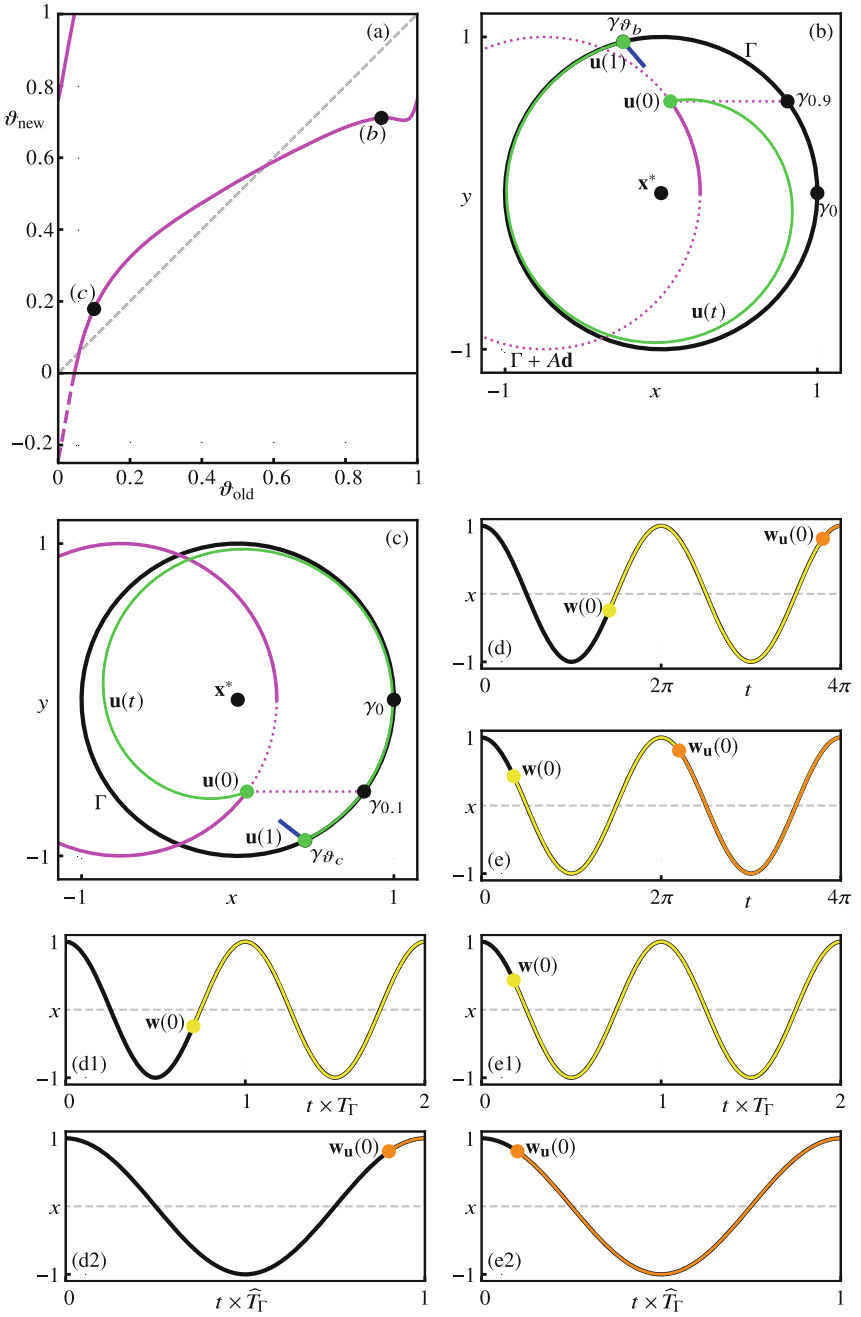

Fig. 2. PTC of Γ in system (22) for $\mathbf{d} = (-1, 0)$ and $A = 0.75$ (a), and continuation set-up at $\vartheta_{\text{old}} = 0.9$ (b) and at $\vartheta_{\text{old}} = 0.1$ (c) with $\mathbf{w}$ and $\mathbf{w}_{\mathbf{u}}$ in (d), (d1), (d2) and (e), (e1), (e2), respectively.

that lies horizontally to the right of $\mathbf{u}(0)$, as indicated by the magenta dotted line segment. The end point $\mathbf{u}(1)$ lies on the linearised isochron, parametrised as $\mathbf{g}(0) + \eta\,\mathbf{v_g}(0)$ with $0 < \eta \le \eta_{\max}$, associated with the rotated head point of $\mathbf{g}$, which is determined by subsystem (5)–(9). The phase of this head point is given by $\vartheta_{\text{new}} = 1 - \nu \,(\text{mod } 1)$, where ν is determined from subsystem (10)–(12) that defines the orbit segment $\mathbf{w}$.

Hence, the two orbit segments $\mathbf{w}$ and $\mathbf{w_u}$ essentially determine the PTC, that is, the map $P : \vartheta_{\text{old}} \mapsto \vartheta_{\text{new}}$. Their x-coordinates are plotted versus time in panel (d) for $\vartheta_{\text{old}} = 0.9$ and in panel (e) for $\vartheta_{\text{old}} = 0.1$, respectively, overlaid on two copies of Γ (black curve), that is, time t runs from 0 to 4π. The further panels (d1) and (d2) for $\vartheta_{\text{old}} = 0.9$ and panels (e1) and (e2) for $\vartheta_{\text{old}} = 0.1$ show $\mathbf{w}$ (yellow curve) and $\mathbf{w_u}$ (orange curve) individually, relative to the periods T_Γ and $\widehat{T}_\Gamma$, respectively. Note that both $\mathbf{w}$ and $\mathbf{w_u}$ end at $x = 1$, for $t = 4\pi$ and $t = 2\pi$, respectively, as required, but their initial points differ. As ϑ decreases from 1 to 0 during the continuation, the orbit segment $\mathbf{w_u}$ lengthens as expected, but note that $\mathbf{w}$ lengthens as well; this is due to the (near-)monotonically increasing nature of the PTC for this example.

4.2 Loss of Invertibility

Recall that any PTC is the identity for $A = 0$, and invertible for sufficiently small amplitude A of the perturbation, because its graph remains a 1:1 torus knot on the torus parametrised by the two periodic variables ϑ_{old} and ϑ_{new}. However, the PTC in Fig. 2(a) for $A = 0.75$ is no longer near the identity: it is not injective and, hence, not invertible.

To show how injectivity of the PTC is lost as A is increased, we consider again model (22), but now with $a = 0.25$; see also [22]. Apart from the attracting unit circle $\Gamma_{\text{s}} = \Gamma$ with period $T_\Gamma = 2\pi$, there exists then also a repelling circle Γ_{u} with radius $r = a = 0.25$ and period $2\pi/(1 + \omega\,(1 - a)) = 3.2\,\pi$; note that Γ_{u} forms the boundary of the basins of attraction of both Γ_{s} and the equilibrium $\mathbf{x}^*$ at the origin, which is now attracting.

We consider three phase resets of Γ_{s} of the form $\Gamma_{\text{s}} + A\,\mathbf{d}$ in the positive direction $\mathbf{d} = (1, 0)$ and with $A = 0.54$, $A = 0.59$, and $A = 0.64$. Figure 3 shows the three corresponding PTCs, the corresponding PRCs, and the perturbed cycles $\Gamma_{\text{s}} + A\,\mathbf{d}$ in increasingly darker shades of magenta as A increases in panels (a), (b), and (c), respectively. Panel (a) shows that the first PTC for $A = 0.54$ is injective and invertible. As A is increased to approximately $A = 0.59$, the graph has a cubic tangency near $(\vartheta_{\text{old}}, \vartheta_{\text{new}}) = (0.45, 0.24)$, because the associated map P has an inflection point at $\vartheta_{\text{old}} \approx 0.45$. For larger values of A, such as for $A = 0.64$, the PTC has a local maximum followed by a local minimum and is, hence, no longer invertible. Note from Fig. 3(b) that this qualitative change of the PTC does not lead to a corresponding qualitative change of the PRC.

Figure 3(c) and the enlargement near the basin boundary Γ_{u} in panel (d) show that the loss of injectivity of the PTC is due to a cubic tangency between the perturbed cycle $\Gamma_{\text{s}} + A\mathbf{d}$ and the foliation of the basin of Γ_{s} by (forward-time) isochrons; ten isochrons are shown in panel (c) and one hundred in panel (d),