

are generated by the passage through a dynamic Hopf bifurcation, similar to the one observed for the Olsen model in Section 6. This example illustrates how the methods described in this paper can be applied effectively to a system of higher dimension than just three or four.

8. Numerical methods for slow-fast systems. This section discusses numerical methods that we used to compute the two-dimensional slow manifolds shown in many of the figures, as well as stable and unstable manifolds of equilibrium points. The slow manifold computations choose an end point of each orbit segment on the critical manifold away from a fold; this approximation yields errors that are $O(\varepsilon)$ but decay quickly as one moves away from the end point. For stable or unstable manifolds of equilibria, orbit segments are chosen to lie in the linear eigenspace associated with the stable or unstable eigenvalues, respectively. The computational error associated with this approximation also decays quickly as one moves away from the endpoint; see [41, 130] for analysis of these approximation errors .

A simple and effective method for computing invariant manifolds as families of orbit segments is to use initial value solvers as the basic algorithm with initial conditions chosen on a mesh of points transverse to the flow in the invariant manifold; we call this the “sweeping” method. Despite its simplicity, this sweeping method fails to produce satisfactory results in some cases. In particular, strong convergence or divergence of trajectories toward one another makes the choice of the initial mesh problematic and can produce very non-uniform “coverage” of the desired manifold; see [59, 60]. In multiple-time-scale systems, the fast exponential instability of Fenichel manifolds that are not attracting makes initial value solvers incapable of tracking these manifolds by forward integration. These issues prompt the use of boundary value methods combined with continuation as an alternate strategy for computing invariant manifolds [131, 132] . We have used both strategies in this paper. This section presents more details of the techniques used to compute attracting and repelling slow manifolds of systems with one fast and two slow variables, as well as the continuation of canard orbits when a parameter is varied.

8.1. Sweeping invariant manifolds. The Fenichel manifolds of systems with a single fast variable are either attracting or repelling. As a result, forward trajectories with initial conditions on the critical manifold will converge quickly to an attracting Fenichel manifold and backward trajectories with initial conditions on the critical manifold will converge quickly to a repelling Fenichel manifold. Thus, one way to compute two-dimensional attracting and repelling Fenichel manifolds of a three-dimensional flow is to apply an initial value solver in the appropriate time direction to a mesh of initial conditions along a curve of the critical manifold transverse to the slow flow. We used this sweeping method to compute S_ε^T in Figure 11; see also [162] for an early use of this method to compute two-dimensional invariant manifolds and Wechselberger [233] and Guckenheimer and Haiduc [86] for an example involving folded nodes.

When incorporated into a continuation framework, the sweeping method can also be used if the critical manifold is not known in closed form and the mesh of initial conditions cannot be selected beforehand. Continuation methods [49] provide well-established algorithms that augment equation solvers like Newton’s method with strategies for choosing new starting points when solving under-determined systems of equations. More precisely, suppose $F : \mathbb{R}^{m+n} \rightarrow \mathbb{R}^m$ is a smooth function given by m equations of $m+n$ variables. The implicit function theorem states that the zeros of F form a smooth n -dimensional manifold M near points where the matrix DF of partial derivatives has full rank m . Moreover, the theorem gives a formula for the tangent space of M . Most continuation methods treat the case $n = 1$ where the set of solutions is a curve; see [101] for the case $n > 1$. In general, the methods are based on a predictor-corrector procedure: given a point on M , tangent (or higher-order) information is used to choose a new seed for the solver to find a new point on M . The sweep-

ing method described above selects the continuation step size based on equal increments of a specific coordinate or direction, but more sophisticated step size adaptations can be used as well. For example, *pseudo-arclength* continuation as implemented in AUTO [50] chooses points based on their distances from each other in \mathbb{R}^{m+n} .

We also used a sweeping method to compute the global unstable manifold $W^u(p)$ in Figure 11. The mesh of initial conditions was taken to lie along a ray in the tangent space of $W^u(p)$, with endpoints of the mesh at successive intersections of a trajectory of the linearized system with this ray. The sweeping method works well here, because the selected orbit segments provide adequate “coverage” of $W^u(p)$.

8.2. Continuation of orbit segments with boundary value solvers. The core algorithms of AUTO [50] are boundary value solvers and continuation methods. The sweeping method described in the previous section can also be implemented in AUTO [50], so that the initial value problems are solved using a collocation method; see [49] for details. The techniques described in this section impose boundary conditions on both end points of the orbit segments, which makes the method more versatile and suitable in a wider context; see also [131]. We describe here how to formulate two-point boundary value problems (BVP) in order to compute slow manifolds and associated canard orbits.

We consider two-point boundary value problems of the form

$$\begin{cases} \dot{\mathbf{u}} &= T\mathbf{g}(\mathbf{u}, \lambda), \\ \mathbf{u}(0) &\in L, \\ \mathbf{u}(1) &\in \Sigma, \end{cases} \quad (8.1)$$

where $\mathbf{g} : \mathbb{R}^n \times \mathbb{R}^p \rightarrow \mathbb{R}^n$ is sufficiently smooth, $T \in \mathbb{R}$, $\lambda \in \mathbb{R}^p$ are parameters and L and Σ are submanifolds of \mathbb{R}^n . The parameter T rescales time so that the orbit segments always correspond to trajectories in the time interval $[0, 1]$. Hence, the boundary conditions at the two end points always apply to $\mathbf{u}(0)$ and $\mathbf{u}(1)$. In order to have a well-posed problem with isolated solutions, the number of boundary conditions should equal the number of equations (n , because $(\mathbf{u}) \in \mathbb{R}^n$) plus the number of free parameters (at most $p + 1$ for the parameter λ and the total integration time T). We are interested in one-parameter families of solutions of (8.1), which means that we allow one fewer boundary condition (or one additional free parameter). Note that T is typically unknown and we may view T as the extra free parameter.

Let us first consider the computation of two-dimensional attracting and repelling slow manifolds S_ε^a and S_ε^r . To simplify the explanation, we assume that we have a three-dimensional slow-fast system with two slow variables and a folded node. In this context, the parameter λ remains fixed, and we obtain a one-parameter family of orbit segments (with unknown total integration times T) by imposing a total of three boundary conditions. This means that the dimensions of L and Σ in (8.1) sum up to $n = 3$. Our approach is to choose L as a curve (or straight line) on the critical manifold, which requires two boundary conditions, and Σ as a surface (or plane), which requires one boundary condition, such that the associated one-parameter family of orbit segments covers the desired portion of the slow manifold. For example, in order for S_ε^a to come into the folded node region, we let L be a curve on the attracting sheet of the critical manifold transverse to the slow flow and Σ be a surface orthogonal to the fold curve F at the folded node. The same approach works for S_ε^r , where we choose L on the repelling sheet of the critical manifold; note that $T < 0$ for such a family of orbit segments. We remark that these choices can also be used with the sweeping method and an initial value solver that detects a “stopping condition” defined by the level set of a function. With the boundary value solvers, we can exchange the roles of L and Σ , which is more appropriate for finding canard orbits; see Section 8.3. The slow manifolds can be

extended by choosing cross-sections Σ orthogonal to F at points that lie beyond the folded node. Figures 6, 20 and 29 give examples of such visualizations; see also [40, 41, 43].

As with all continuation, an important issue is to find a first solution. When continuing solutions of a boundary value problem, explicit solutions may be known from which such a first solution may be constructed; see [41] for an example. However, in general no explicit solution is known and a first solution must be found in a different way. We use a homotopy method to generate an initial orbit segment; the main idea is to continue intermediate orbit segments via two auxiliary BVPs — the first to obtain an orbit segment from a point on the fold curve F to the section, and the second to move the end point on F along the critical manifold to a suitable distance from F ; see [40] for details.

We now illustrate this method with the Koper model (4.1), which was also used for the case study in Section 4. We use the parameters $(\varepsilon_1, \varepsilon_2, \lambda, k) = (0.1, 1, 7, -10)$; note that $\lambda > 0$ as in [122], which is symmetrically related to the case with $\lambda = -7$ considered in Section 4. As shown in Section 4, there is a folded node in this model, which organizes the SAOs in some of the observed MMOs; in original coordinates it is at

$$\mathbf{p}_{\text{fn}} = \left(-1, \frac{2 + \lambda}{k}, \frac{2\lambda + 4 + k}{k} \right) = (-1, -0.9, -0.8). \quad (8.2)$$

We compute $S_{\varepsilon_1}^a$ and $S_{\varepsilon_1}^r$ as solutions to the BVPs given by (8.1), where \mathbf{g} is defined as the right-hand side of (4.1). As boundary conditions, we use the same section Σ for both $S_{\varepsilon_1}^a$ and $S_{\varepsilon_1}^r$ with respective lines $L = L^a$ and $L = L^r$ as follows

$$\Sigma_{\text{fn}} := \{(x, y, z) \in \mathbb{R}^3 \mid z = -0.8\}, \quad (8.3)$$

$$L^a := S \cap \{x = -1.5\}, \quad (8.4)$$

$$L^r := S \cap \{x = -0.2\}. \quad (8.5)$$

Figure 33 shows the result of the computations. We find a first orbit segment on $S_{\varepsilon_1}^a$ using two homotopy steps; this is illustrated in Figure 33(a). Starting from the trivial solution $\mathbf{u} = \{\mathbf{p}_{\text{fn}} \mid 0 \leq t \leq 1\}$, with total integration time $T = 0$, we continue the family of orbit segments that solves (4.1) subject to $\mathbf{u}(1) \in \Sigma_{\text{fn}}$ and $\mathbf{u}(0) \in F$. We stopped the computation, detected by a user-defined function in AUTO, as soon as

$$\mathbf{u}(0) \in \tilde{\Sigma}^a := \{(x, y, z) \in \mathbb{R}^3 \mid z = -0.76\}.$$

The orbit segment with its end point on F in Figure 33(a) is this last computed solution of the family. The second step of the homotopy moves $\mathbf{u}(0) \in S$ away from F (approximately) parallel to Σ , that is, we next continue the family of orbit segments that solves (4.1) subject to $\mathbf{u}(1) \in \Sigma_{\text{fn}}$ and $\mathbf{u}(0) \in \tilde{L}^a = S \cap \tilde{\Sigma}^a$. The continuation stops when L^a is reached, which is again detected by a user-defined function in AUTO. A selection of orbit segments in this family are shown in Figure 33(a) (red curves); only the last orbit segment \mathbf{u}^a (dark red) lies on $S_{\varepsilon_1}^a$ to good approximation. A similar computation was done to obtain a first orbit segment on $S_{\varepsilon_1}^r$, where we use the intermediate section $\tilde{\Sigma}^r := \{z = -0.87\}$; this is illustrated in Figure 33(b), where the orbit segment \mathbf{u}^r (cyan) serves as a first solution on $S_{\varepsilon_1}^r$.

Once the first orbit segments \mathbf{u}^a and \mathbf{u}^r have been found we start the continuation of (8.1) with (8.3) and (8.4) for the attracting slow manifold $S_{\varepsilon_1}^a$ and with (8.3) and (8.5) for the repelling slow manifold $S_{\varepsilon_1}^r$. The result is presented in Figure 33(c), and the intersection curves of $S_{\varepsilon_1}^a$ and $S_{\varepsilon_1}^r$ with Σ_{fn} are shown in Figure 33(d). The transverse intersection points of $S_{\varepsilon_1}^a \cap \Sigma_{\text{fn}}$ and $S_{\varepsilon_1}^r \cap \Sigma_{\text{fn}}$ in panel (d) correspond to secondary canard orbits; the three-dimensional view in panel (c) shows three of these, labeled ξ_1 , ξ_2 and ξ_3 . Precisely for the purpose of locating and continuing canard orbits it is necessary to choose the common cross-section Σ_{fn} for the calculations of $S_{\varepsilon_1}^a$ and $S_{\varepsilon_1}^r$; see also the next section.




FIG. 33. Computation of the slow manifolds $S_{\varepsilon_1}^a$ and $S_{\varepsilon_1}^r$ of the Koper model (4.1) with $(\varepsilon_1, \varepsilon_2, \lambda, k) = (0.1, 1, 7, -10)$. Panels (a) and (b) show the homotopy steps to construct first orbit segments \mathbf{u}^a (dark red) on $S_{\varepsilon_1}^a$ (red) and \mathbf{u}^r (cyan) on $S_{\varepsilon_1}^r$ (blue) that connect the section Σ_{fn} with curves L^a and L^r on the critical manifold S (grey), respectively. The red and blue families are generated during the second homotopy step, which starts from solutions that have one of their end points on the fold curve F of S . Panel (c) shows $S_{\varepsilon_1}^a$ and $S_{\varepsilon_1}^r$ together with three secondary canards ξ_1 , ξ_2 and ξ_3 . Panel (d) shows the intersection curves of $S_{\varepsilon_1}^a$ and $S_{\varepsilon_1}^r$ in Σ_{fn} that are used to detect canard orbits.

8.3. Finding and following canard orbits. Maximal canards near a folded node are transverse intersection curves of the two-dimensional attracting and repelling slow manifolds S_{ε}^a and S_{ε}^r . We briefly discuss here how to detect the canard orbits and subsequently continue them in a system parameter; see also [40, 41, 43]. To represent a maximal canard we must compute S_{ε}^a and S_{ε}^r using a common cross-section Σ of the fold curve at or near the folded node. The common cross-section allows us to obtain a representation of the canard orbit as the concatenation \mathbf{u}^c of an orbit segment $\mathbf{u}^a \subset S_{\varepsilon}^a$ with an orbit segment $\mathbf{u}^r \subset S_{\varepsilon}^r$, where \mathbf{u}^a and \mathbf{u}^r are chosen such that $\mathbf{u}^a \cap \Sigma = \mathbf{u}^r \cap \Sigma$. The concatenated orbit \mathbf{u}^c located with

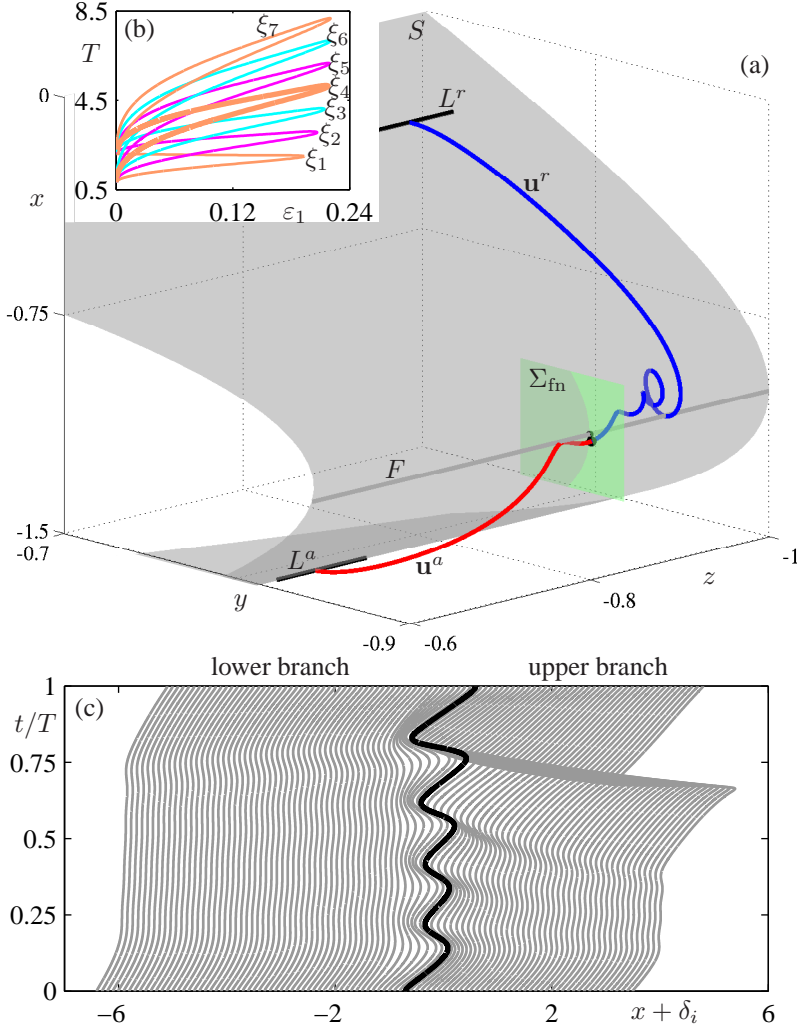


FIG. 34. Continuation of secondary canards of the Koper model (4.1) with $(\varepsilon_2, \lambda, k) = (1, 7, -10)$ starting from $\varepsilon_1 = 0.1$. Panel (a) shows the canard orbit ξ_4 represented by the concatenation \mathbf{u}^c of two orbit segments \mathbf{u}^a and \mathbf{u}^r that match up in Σ_{fn} . Panel (b) shows the continuation of the canard orbits ξ_1 – ξ_7 in ε_1 ; plotted as total integration time T versus ε_1 . Panel (c) shows a two-dimensional “waterfall diagram” of the time profiles of the fast variable x (subject to an offset δ_i) of computed orbit segments along the branch ξ_4 . The bold black curve in panel (c) is the canard orbit ξ_4 at the fold point of the (boldfaced) branch in panel (b).

this method can be continued in a system parameter without the need to recompute the slow manifolds at each step. Recall that AUTO always scales boundary value problems to the time interval $[0, 1]$, so we rescale time on \mathbf{u}^c appropriately and set $T = T^a + T^r$ in (8.1). We can then start the continuation (in a system parameter) subject to the boundary conditions

$$\mathbf{u}^c(0) \in L^a, \quad (8.6)$$

$$\mathbf{u}^c(1) \in L^r, \quad (8.7)$$

which determine \mathbf{u}^c as an isolated solution. In fact, such a continuation typically starts already provided that $\mathbf{u}^a \cap \Sigma \approx \mathbf{u}^r \cap \Sigma$; any small gap in Σ is forced to close by the first Newton

step. These two boundary conditions (8.6) and (8.7) force the orbit segment \mathbf{u}^c to stay very close to the attracting sheet of the critical manifold S until near the fold curve F , and then stay close to the repelling sheet of S up to L^T .

Figure 34 illustrates canard continuation with the Koper model (4.1), where we used ε_1 as the second free parameter (together with T) and kept $(\varepsilon_2, \lambda, k) = (1, 7, -10)$ fixed. Figure 34(a) shows the two orbit segments \mathbf{u}^a and \mathbf{u}^r with (almost) equal end points in the section $\Sigma = \Sigma_{\text{fn}}$; they have been detected as a good approximation of the maximal secondary canard orbit ξ_4 , which is then represented by the concatenated orbit \mathbf{u}^c . We continued ξ_4 , along with six other maximal secondary canards, for increasing and decreasing ε_1 ; see also Figure 33. Figure 34(b) shows these seven branches, labeled ξ_1 – ξ_7 ; here, the vertical axis shows the total integration time T because it clearly distinguishes the branches. When ξ_1 – ξ_7 are continued in the direction of increasing ε_1 , a fold in ε_1 is detected for each branch; we have already seen this in Section 5 and it has also been observed in other systems [43]. Figure 34(c) is a “waterfall diagram” that shows how the maximal secondary canard orbit ξ_4 evolves along the branch as ε_1 is varied; specifically, the time profile of the fast variable x of consecutively computed orbit segments along the branch ξ_4 are plotted with a suitable off-set δ_i . The orbit segment that corresponds to the fold of ξ_4 is highlighted in bold black. Observe that the orbit segments to the left of the fold have four SAOs, whereas past the fold there are only three SAOs followed by a fast segment. Hence the canard orbits past the fold are no longer maximal canards; see also Section 5.

9. Discussion. We described several mechanisms in slow-fast systems that produce mixed-mode oscillations, namely the twisting of slow invariant manifolds near a folded node, oscillations that follow the two-dimensional unstable manifold of a saddle-focus equilibrium near a singular Hopf bifurcation, and the tourbillion mechanism of a dynamic Hopf bifurcation. Geometric singular perturbation theory provides tools to identify the geometry associated with each mechanism, to quantify the MMO signatures, and to describe associated bifurcations. Analysis of the folded node case is more complete than the other cases. Recent work on singular Hopf bifurcation [85] and the transition from singular Hopf to folded nodes [143] provides substantial detail on the second case, but much remains to be discovered about the unfolding of a singular Hopf bifurcation that is relevant to MMOs. Historically, the dynamic Hopf bifurcation was discovered first, and detailed analysis exists for the case of a delayed Hopf bifurcation of the layer equations [168]. Together, these mechanisms constitute a partial framework for classifying MMOs in multiple-time-scale systems that can be further extended. Perhaps the most surprising aspect of the theory we have described is that oscillations can appear from the interaction of fast and slow time scales even when neither of these time scales individually displays oscillations.

We have used four case studies to illustrate theoretical concepts and they serve as a test-bed for the development of numerical methods. The MMOs in the Koper model and the three-dimensional reduction of the Hodgkin–Huxley equations have SAOs that occur on intermediate time scales due to folded nodes and singular Hopf bifurcations. In the folded-node mechanism, three parameters play key roles in determining the geometry of the small oscillations: the ratio ε of time scales, the eigenvalue ratio μ of the folded node in the desingularized reduced system, and the distance δ of global return trajectories from certain invariant manifolds. Intersections of invariant manifolds are prerequisite to global returns that produce MMOs in these examples, and tangencies between these manifolds constitutes a new type of bifurcation that is found on the boundaries of parameter regions yielding MMOs. We found fast oscillations of the layer equations in the Olsen and Showalter–Noyes–Bar-Eli models of chemical reactions. Both models exhibit MMOs due to the dynamic Hopf mechanism. These two case studies also illustrate how the theory applies in higher dimensions and how numer-

System / Reaction	References
Belousov-Zhabotinskii (BZ) reaction	
- Virginia	[83, 103, 104, 105, 202]
- Texas	[156, 157, 158, 194, 195]
- Bordeaux	[8, 9, 183, 193, 228]
- Other groups	[107, 155, 184, 185, 206]
Briggs-Rauscher (BR) reaction	[28, 73, 171, 231]
peroxidase-oxidase (PO) reaction	[76, 97, 98, 99, 100, 106, 173, 207]
HPTCu reaction	[15, 137, 175, 176, 227]
Bray-Liebhafsky (BL) reaction	[73, 149, 230]
copper and phosphoric acid	[6, 200]
indium/thiocyanate (IT) reaction	[125, 126]
BSFA-system	[128]
p-CuInSe ₂ /H ₂ O ₂ -system	[167, 182]
spin-wave experiment	[5]
rhythm neural network (PreBötC)	[39]
stellate cells	[45, 46, 61]
pituitary cells	[225, 229]
combustion oscillations	[82]
dusty plasmas	[160]
semiconductor lasers	[7, 81, 226]
CO oxidation	[57, 58, 136]

TABLE 9.1

References for experimental investigations of MMOs.

ical tools can be extended to investigate and identify the mechanisms for generating MMOs in higher-dimensional systems.

One of our goals for this paper is to facilitate fitting dynamical models to data. In the case of MMOs, this task has been less successful than with many other nonlinear dynamical phenomena. On the one hand, MMOs are a complex phenomenon, and on the other hand, numerical studies of models have yielded puzzling and sometimes paradoxical results. The theory that has been developed thus far deals best with circumstances where the SAOs have amplitudes that are far too small to be observed even in numerical simulations, but model studies frequently show MMOs with SAOs that are readily visible. Thus, numerical methods that identify the geometric objects highlighted by the theory are essential for bringing theory, models and empirical data together. We have reviewed recent advances in computing two-dimensional invariant manifolds and their intersections that are especially important in three-dimensional models. Extension of these methods to higher dimensions is one of the challenges for further advances in this subject.

We conclude this survey with a brief review of the MMO literature, and a short discussion of other mechanisms for MMOs in ODEs and beyond.

9.1. MMO literature review. This section provides an overview, in the form of three tables, of references where examples of MMOs have been studied experimentally or in model systems. We do not claim that this overview is complete; rather, these tables are intended as an entry point into the extensive literature on the subject. Table 9.1 lists experimental work on MMOs. The majority of these experiments have been carried out for chemical reactions. As suggested in [8], we subdivided the large number of references on the Belousov-Zhabotinskii