Celestial Mechanics and Dynamical Astronomy (2019) 131:13
https://doi.org/10.1007/s10569-019-9890-8

ORIGINAL ARTICLE

Homoclinic dynamics in a restricted four-body problem: transverse connections for the saddle-focus equilibrium solution set

Shane Kepley[1] · **J. D. Mireles James**[2]

Received: 5 August 2018 / Revised: 12 February 2019 / Accepted: 12 February 2019 /
Published online: 4 March 2019
© Springer Nature B.V. 2019

Abstract
We describe a method for computing an atlas for the stable or unstable manifold attached to an equilibrium point and implement the method for the saddle-focus libration points of the planar equilateral restricted four-body problem. We employ the method at the maximally symmetric case of equal masses, where we compute atlases for both the stable and unstable manifolds. The resulting atlases are comprised of thousands of individual chart maps, with each chart represented by a two-variable Taylor polynomial. Post-processing the atlas data yields approximate intersections of the invariant manifolds, which we refine via a shooting method for an appropriate two-point boundary value problem. Finally, we apply numerical continuation to some of the BVP problems. This breaks the symmetries and leads to connecting orbits for some nonequal values of the primary masses.

Keywords Gravitational 4-body problem · Invariant manifolds · High-order Taylor methods · Automatic differentiation · Numerical continuation

Mathematics Subject Classification 70K44 · 34C45 · 70F15

1 Introduction

Illuminating studies by Darwin, Strömgren, and Moulton in the first decades of the Twentieth Century established the importance of numerical calculations in the qualitative theory of Hamiltonian systems (Darwin 1897; Strömgren 1934; Moulton et al. 1920). In particular,

The second author was partially supported by NSF Grant DMS-1813501. Both authors were partially supported by NSF Grant DMS-1700154 and by the Alfred P. Sloan Foundation Grant G-2016-7320.

✉ J. D. Mireles James
jmirelesjames@fau.edu

Shane Kepley
sk2011@math.rutgers.edu

[1] Department of Mathematics, Rutgers University, New Brunswick, USA

[2] Department of Mathematical Sciences, Florida Atlantic University, Boca Raton, USA

their work gave new insights into the orbit structure of the circular restricted three-body problem (CRTBP), a problem already immortalized by Poincaré. Interest in the CRTBP was reinvigorated in the 1960s with the inauguration of the space race and a number of authors including Szebehely and Nacozy (1967), Szebehely and Flandern (1967) harnessed the newly available power of digital computing to settle some questions raised by Strömgren. The interested reader will find a delightful retelling of this story with many additional references in the book of Szebehely (1967).

Motivated by the works just mentioned, in 1973 Henrard proved a theorem settling a conjecture of Strömgren about the role of asymptotic orbits. More precisely, Henrard showed that the existence of a transverse homoclinic for a saddle-focus equilibrium in a two-degree-of-freedom Hamiltonian system implies the existence of a tube of periodic orbits parameterized by energy and accumulating to the homoclinic (Henrard 1973). In the same paper he showed that the period of the orbits in the family goes to infinity and their stability changes infinitely often as they accumulate to the homoclinic. This phenomenon was called the *blue sky catastrophe* by Abraham (1985) and has been studied by a number of authors including Shilnikov et al. (2014), Devaney (1977).

In 1976 it was further shown by Devaney that such a transverse homoclinic—again for a saddle-focus in a two-degree-of-freedom Hamiltonian system—implies the existence of chaotic dynamics in the energy level of the equilibrium (Devaney 1976). See also the works of Lerman (1991, 2000). Such theorems should be thought of as Hamiltonian versions of the homoclinic bifurcations studied by Shilńikov (1967, 1970a, b). Taken together the results cited so far paint a vivid picture of the rich dynamics near a transverse homoclinic connection in a two-degree-of-freedom Hamiltonian system.

The present study concerns asymptotic orbits in the planar equilateral restricted four-body problem, henceforth referred to as the circular restricted four-body problem (CRFBP). The problem has a rich literature dating at least back to the work of Pedersen (1944, 1952). Detailed numerical studies of the equilibrium set, as well as the planar and spatial Hill's regions, are found in Simó (1978), in Baltagiannis and Papadakis (2011a), and in Álvarez-Ramírez and Vidal (2009). Mathematically rigorous theorems about the equilibrium set and its bifurcations are proven by Leandro (2006), Barros and Leandro (2011, 2014) (with computer assistance). They show that for any value of the masses there are either 8, 9, or 10 equilibrium solutions with 6 outside the equilateral triangle formed by the primary bodies (see Fig. 1).

Fundamental families of periodic orbits are considered by in Papadakis (2016a, b), and by Burgos-García and Delgado (2013a), Burgos-García and Bengochea (2017). A study by Burgos-García, Lessard, and Mireles James proves the existence of some spatial periodic orbits for the CRFBP (Burgos-García et al. 2019) (again with computer assistance). An associated Hill's problem is derived, and its periodic orbits are studied by Burgos-García (2016), Burgos-García and Gidea (2015).

Regularization of collisions is studied by Alvarez-Ramírez et al. (2014). Chaotic motions were studied numerically by Gidea and Burgos (2003) and by Alvarez-Ramírez and Barrabés (2015). Perturbative proofs of the existence of chaotic motions are found in the work of Cheng and She (2017), She and Cheng (2014), She et al. (2013) and also in the work of Alvarez-Ramírez et al. (2018). Blue sky catastrophes in the CRFBP were previously studied by Burgos-García and Delgado (2013b) and by Kepley and Mireles James (2018). This last reference develops (computer-assisted) methods of proof for verifying the hypotheses of the theorems of Hernard and Devaney.

The main goal of the present work is to study orbits which are homoclinic to a saddle-focus equilibrium solution in the equilateral restricted four-body problem. We apply the parameterization method of Cabré, Fontich, and de la Llave to compute a chart for the

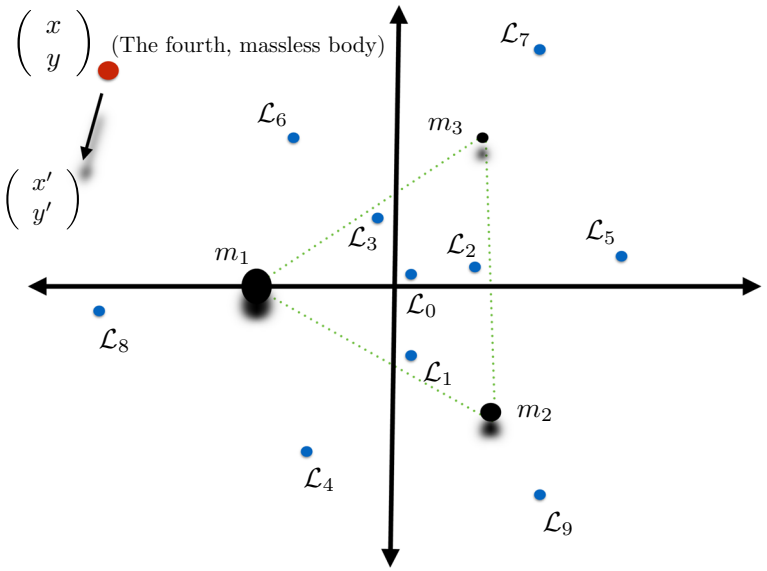

Fig. 1 Configuration space for the CRFBP: the three primary bodies with masses m_1, m_2, and m_3 are arranged in an equilateral triangle configuration of Lagrange, which is a relative equilibrium solution of the three-body problem. After transforming to a co-rotating frame, we consider the motion of a fourth massless body. The equations of motion have 8, 9, or 10 equilibrium solutions (libration points) denoted by $\mathcal{L}_j$ for $0 \leq j \leq 9$. The number of libration points, and their stability, varies depending on m_1, m_2, and m_3. In this work we study the points, $\mathcal{L}_{0,4,5,6}$, which are the only libration points which can have saddle-focus stability

stable or unstable manifold in a neighborhood of the equilibrium (Cabré et al. 2003a, b, 2005). Then, we implement the analytic continuation scheme for local invariant manifolds developed by Kalies et al. (2018), where it was applied to some two-dimensional manifolds in the Lorenz system. We adapt this scheme for the CRFBP and compute atlases for the local stable/unstable manifolds attached to a saddle-focus equilibrium. By an atlas, we mean a collection of analytic maps or *charts* of the form, $P\colon [-1, 1]^2 \to \mathbb{R}^4$, where the image of P lies in the stable or unstable manifold. The union of these charts is a piecewise approximation for a large portion of the manifold away from the equilibrium. For a more formal definition, see any standard text on differential geometry. The charts are computed using high-order polynomial approximations with algorithms that exploit automatic manipulations of formal series.

After computing the stable/unstable manifold atlases, we post-process to find approximate intersections. Once a potential intersection is located, we refine the approximation using a Newton scheme for a two-point boundary value problem as in the classical work of Doedel and Friedman (1989), Doedel et al. (1997). In the case of the CRFBP, our algorithm identifies a large collection of connecting orbits which are naturally ordered by connection time. We focus on the maximally symmetric case of equal masses, which we refer to as *the triple Copenhagen problem*. We prove that a rotational symmetry in this case reduces the complexity of the atlas computations by a factor of 3.

The algorithm for producing the atlases utilizes an adaptive subdivision routine to carefully control errors. This results in a large number of charts, on the order of tens of thousands, in only a few minutes of computation time. These computations are expensive in terms of memory usage, and it is impractical to recompute the atlases for a large number of parameter

values, at least given the resources of the present study, namely laptop/desktop computers running single threads. Instead, after computing an ensemble of connecting orbits for the triple Copenhagen problem, we apply numerical continuation to the boundary value problem describing the homoclinics. That is, we use the connections found for the equal mass case as a jumping off point for exploring nearby—but nonsymmetric—mass parameters. Continuation of the connecting orbits is much more efficient than continuing the entire invariant manifold atlas.

As is well known, the bifurcation structure of the homoclinic continuation problem in the Hamiltonian setting is rich. We do not attempt automatic tracking of new branches, nor do we follow folds. A more systematic study of the branching would make an excellent topic for future study, perhaps by combining our invariant manifold atlas data with powerful continuation software such as AUTO (Champneys et al. 1996).

We emphasize that our restriction to the equal masses case is due to convenience and is not a technical restriction on the method itself. Our atlas algorithm applies to any choice of parameters or even to other Hamiltonian systems. Thus, even though we abandon the branch whenever the homoclinic continuation algorithm fails, we always have the ability to dig deeper into the cause of failure by running the full atlas computation from scratch.

We remark that our method is deployed in the full phase space and does not require choosing a fixed surface of section in which to study intersections of the invariant manifolds. This is advantageous as many problems do not admit a single section for which the return map is topologically conjugate to the true dynamics. Considering the intersections of the stable/unstable manifolds in a particular section may not reveal all the connecting orbits. Moreover, the first intersections to appear in phase space may not be the first to appear in a given section. Indeed, projecting to a section can introduce discontinuities which make it impossible to precisely formulate notions like "first intersection." The great virtue of a surface of section (restricted to an energy level) is that it leads—at least in the case of a two-degree-of-freedom Hamiltonian—to a two-dimensional representation of the dynamics. We remark that the methods of the present work generalize to systems with three or more degrees of freedom, where considering surfaces of section is less fruitful.

2 Saddle-focus equilibrium solutions of the equilateral CRFBP

In this section, we review well known results about the set of equilibrium solutions in the CRFBP, focusing on material which informs the calculations carried out in the remainder of the work. We are especially interested in the number and location of saddle-foci and in how these depend on the mass ratios. First, we recall the mathematical formulation of the problem and some of its elementary properties.

2.1 The planar equilateral circular restricted four-body problem

Consider three particles with masses $0 < m_3 \leq m_2 \leq m_1 < 1$, normalized so that

$$m_1 + m_2 + m_3 = 1.$$

These massive particles are referred to as the "primaries." Suppose that the primaries are located at the vertices of a planar equilateral triangle, rotating with constant angular velocity. That is, we assume that the three massive bodies are in the triangular configuration of Lagrange. We choose a co-rotating coordinate frame which puts the triangle in the xy-plane

and fixes the center of mass at the origin. We orient the triangle so that the first primary is on the negative x-axis, the second body is in the lower right quadrant, and the smallest body is in the upper right quadrant. Once in co-rotating coordinates, we are interested in the dynamics of a fourth, massless particle with coordinates (x, y), moving in the gravitational field of the primaries. The situation is illustrated in Fig. 1.

We write (x_1, y_1), (x_2, y_2) and (x_3, y_3) to denote the locations of the primary masses. Let

$$K = m_2(m_3 - m_2) + m_1(m_2 + 2m_3).$$

Taking into account the normalizations discussed above, the precise positions of the primary bodies are given by the formulas

$$x_1 = \frac{-|K|\sqrt{m_2^2 + m_2m_3 + m_3^2}}{K}, \qquad y_1 = 0,$$

$$x_2 = \frac{|K|\,[(m_2 - m_3)m_3 + m_1(2m_2 + m_3)]}{2K\sqrt{m_2^2 + m_2m_3 + m_3^2}} \qquad y_2 = \frac{-\sqrt{3}m_3}{2m_2^{3/2}}\sqrt{\frac{m_2^3}{m_2^2 + m_2m_3 + m_3^2}}$$

$$x_3 = \frac{|K|}{2\sqrt{m_2^2 + m_2m_3 + m_3^2}}, \qquad y_3 = \frac{\sqrt{3}}{2\sqrt{m_2}}\sqrt{\frac{m_2^3}{m_2^2 + m_2m_3 + m_3^2}}.$$

Define the potential function

$$\Omega(x, y) := \frac{1}{2}(x^2 + y^2) + \frac{m_1}{r_1(x, y)} + \frac{m_2}{r_2(x, y)} + \frac{m_3}{r_3(x, y)}, \tag{1}$$

where

$$r_j(x, y) := \sqrt{(x - x_j)^2 + (y - y_j)^2}, \qquad j = 1, 2, 3, \tag{2}$$

and let $\mathbf{x} = (x, \dot{x}, y, \dot{y}) \in \mathbb{R}^4$ denote the state of the system. The equations of motion in the rotating frame are

$$\mathbf{x}' = f(\mathbf{x}),$$

where

$$f(x, \dot{x}, y, \dot{y}) := \begin{pmatrix} \dot{x} \\ 2\dot{y} + \Omega_x(x, y) \\ \dot{y} \\ -2\dot{x} + \Omega_y(x, y) \end{pmatrix}. \tag{3}$$

The system conserves the quantity

$$E(x, \dot{x}, y, \dot{y}) = -\left(\dot{x}^2 + \dot{y}^2\right) + 2\Omega(x, y), \tag{4}$$

which is called the *Jacobi integral*. Note that E is smooth—in fact real analytic—away from the primaries. The zero velocity curves are defined by fixing a value of the energy and setting $\dot{x}, \dot{y}$ to zero. These curves are useful for understanding the structure of the phase space and are illustrated in Fig. 2.

As mentioned in the introduction, the CRFBP has exactly 8, 9 or 10 equilibrium solutions, depending on the values of the mass parameters m_1, m_2, and m_3. The equilibria are referred to as *libration points* in the dynamical astronomy literature, and we denote them by $\mathcal{L}_j$ for $0 \le j \le 9$. A typical configuration of these libration points is illustrated in Fig. 1, which also illustrates out naming convention. In the present work we are interested in the

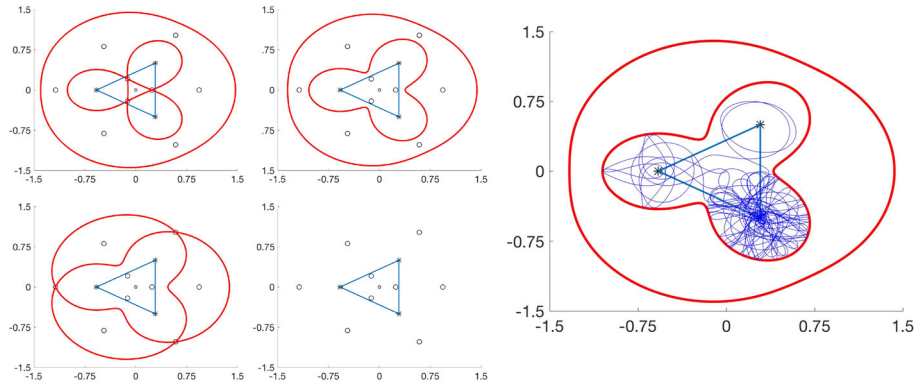

Fig. 2 Zero velocity curves for the triple Copenhagen problem: fixing a value of the Jacobi constant and setting velocity equal in Eq. (4) implicitly defines the zero velocity curves in the phase space of the CRFBP. An orbit which reaches one of these curves arrives with zero velocity and hence turns around immediately. These define natural boundaries which orbits at a given energy level may not cross. Left: the zero velocity curves associated with the energy levels of $\mathcal{L}_{1,2,3}$ (top left) $\mathcal{L}_0$ (top right), $\mathcal{L}_{4,5,6}$ (bottom left), and $\mathcal{L}_{7,8,9}$. Right: a typical orbit in the $\mathcal{L}_0$ energy level confined by the zero velocity curves

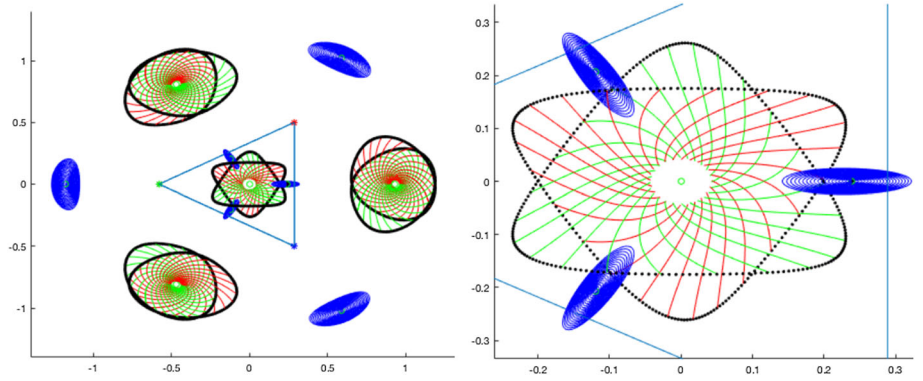

Fig. 3 Two-dimensional local invariant manifolds in the triple Copenhagen problem (CRFBP with equal masses): Left: all two-dimensional attached invariant manifolds for libration points in the equal mass case (one-dimensional manifolds not shown). In the case of equal masses, the libration points $\mathcal{L}_{0,4,5,6}$ have saddle-focus stability. Orbits are shown accumulating to the libration points in forward/backward time (green/red, respectively). The libration points $\mathcal{L}_{4,5,6,7,8,9}$ on the other hand have saddle $\times$ center stability. In this case each libration point has an attached center manifold foliated by periodic orbits—the so-called planar Lyapunov orbits. We make no systematic study the Lyapunov orbits in the present work and only remark that they appear to organize some of the homoclinic orbits in the discussion to follow. Right: closeup on the inner libration points and their invariant manifolds. All references to color refer to the online version

linear stability of the libration points. We are especially interested in determining the mass ratios where $\mathcal{L}_j$ with $j = 0, 4, 5, 6$ are saddle-focus—as opposed to real saddle or center $\times$ center—equilibria. This question is considered from a numerical point of view in Sect. 2.2.

We note that for all values of the masses, $\mathcal{L}_j$ with $j = 1, 2, 3, 7, 8, 9$ have either saddle $\times$ center, or center $\times$ center stability depending on the values of the masses. The local two-dimensional invariant manifolds attached to all ten libration points are illustrated in Fig. 3, for the case of equal masses.

2.2 Saddle-foci in parameter space

The CRFBP admits as many as four and as few as zero saddle-focus equilibrium points, depending on the mass ratios. We now consider briefly what happens in between these extremes as the masses are varied. The problem is normalized so that $m_1 + m_2 + m_3 = 1$, with $m_3 \leq m_2 \leq m_1$, so we have that $m_1 \in [1/3, 1]$, $m_2 \in [0, 1/2]$ and $m_3 \in [0, 1/3]$. Considering the 2-simplex in $\mathbb{R}^3$ satisfying these constraints, we see that when $m_1 \in [1/3, 1/2]$ we have

$$m_3 \in \left[-2m_1 + 1, \frac{-1}{2}m_1 + \frac{1}{2}\right],$$

while for $m_1 \in [1/2, 1]$ we have

$$m_3 \in \left[0, \frac{-1}{2}m_1 + \frac{1}{2}\right].$$

In either case, once we choose m_1 and m_3, the value of m_2 is determined by

$$m_2 = 1 - m_1 - m_3.$$

The question is, how does the stability of the libration points depend on the mass ratios? We address the question for each of the points, $\mathcal{L}_{0,4,5,6}$, as follows. Beginning with the case of equal masses, $m_1 = m_2 = m_3 = 1/3$, we numerically continue each equilibrium to the opposite boundary of the parameter simplex at $m_3 = 0$. Throughout the computation, we track the stability of each libration point and label a parameter point with a black dot whenever the stability is of saddle-focus type. The results are summarized in Fig. 4. We refer to the curve in the parameter simplex where the stability changes as the Routh–Gascheau curve.

Roughly speaking, we see that when $1/3 \leq m_1 \leq 0.42$ the libration point $\mathcal{L}_0$ is a saddle-focus for all allowable values of m_2, m_3. When $m_1 > 0.43$, the libration point $\mathcal{L}_0$ is no longer a saddle, no matter the values of m_2, m_3. The points $\mathcal{L}_{4,6}$ on the other hand have saddle-focus stability for most parameter values, and only bifurcate after $m_1 > 0.95$ (with $\mathcal{L}_6$ a little more robust than $\mathcal{L}_4$ except when $m_2 = m_3$). The libration point $\mathcal{L}_5$ is the most robust. It maintains saddle-focus stability until $m_1 \approx 0.99$. For $m_1 > 0.995$ there are no more saddle-foci at all. By reading parameter values off of the frames in Fig. 4, we can arrange that the CRFBP has 1, 2, 3 or 4 saddle-focus equilibria. In the sequel we are interested in homoclinic connections for such parameters.

2.3 Two ways to formulate a connecting orbit: phase space geometry and boundary value problems

There are two standard ways to think about connecting orbits and—while they are completely equivalent from a mathematical point of view—in practice they have different advantages and disadvantages. In the following let $f : \mathbb{R}^n \to \mathbb{R}^n$ denote a smooth vector field and let $\mathbf{x}_0 \in \mathbb{R}^n$ be an equilibrium solution for f. We write $W^s(\mathbf{x}_0)$ and $W^u(\mathbf{x}_0)$ to denote, respectively, the stable and unstable manifolds attached to $\mathbf{x}_0$.

– **Analytic definition** If $\mathbf{x} : \mathbb{R} \to \mathbb{R}^n$ satisfies

$$\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{x}(t) = f(\mathbf{x}(t)),$$

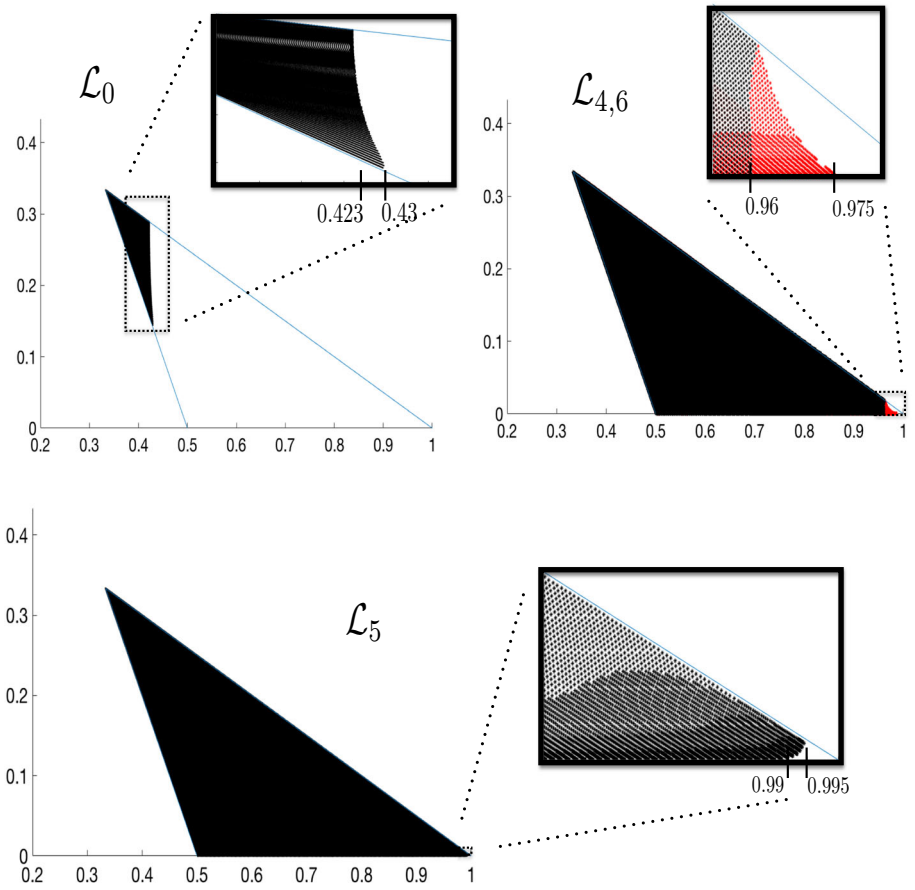

Fig. 4 Mass values and saddle-focus stability: results of a numerical search of the parameter space. Values of m_1 are on the horizontal axis and values of m_3 are on the vertical axis. These determine the remaining mass parameter through the relation $m_2 = 1 - m_1 - m_3$. In each frame a parameter pair is marked with a black or red dot if the libration point $\mathcal{L}_{0,4,5,6}$ has saddle-focus stability. The top left figure reports the results for $\mathcal{L}_0$, the top right for $\mathcal{L}_{4,6}$, and the bottom frame for $\mathcal{L}_5$. In each case the inlay zooms in on the Routh–Gascheau bifurcation curve. Note that these bifurcation curves are nonlinear, and that in the top right results for $\mathcal{L}_4$ are black and results for $\mathcal{L}_6$ are red. We remark that the changes in the dot pattern in the bottom right inlay are due to the use of an adaptive step size in our continuation algorithm

for all $t \in \mathbb{R}$, and satisfies the asymptotic boundary conditions

$$\lim_{t \to \pm\infty} \mathbf{x}(t) = \mathbf{x}_0,$$

then we say that $\mathbf{x}$ is a homoclinic connecting orbit for $\mathbf{x}_0$.

– **Geometric definition** If

$$\hat{x} \in W^s(\mathbf{x}_0) \cap W^u(\mathbf{x}_0),$$

and $\mathbf{x} = \text{orbit}(\hat{x})$ denotes the orbit which passes through $\hat{x}$, then $\mathbf{x}$ is a homoclinic connecting orbit for $\mathbf{x}_0$. If the intersection of the manifolds is transverse, then we say that $\mathbf{x}$ is a transverse homoclinic connection.

The analytic definition is recast as a finite time boundary value problem by projecting the boundary conditions onto local stable/unstable manifolds. If P, Q are parameterizations of the local unstable and stable manifolds, respectively, then we look for $T > 0$ and $\mathbf{x}\colon [0, T] \to \mathbb{R}^n$, so that $\mathbf{x}$ solves the differential equation subject to the boundary conditions

$$\mathbf{x}(0) \in \text{image}(P), \qquad \text{and} \qquad \mathbf{x}(T) \in \text{image}(Q).$$

In applications one frequently replaces P and Q by their linear approximations. In Sect. 3 we review an approach called *the parameterization method* for computing high-order polynomial approximations of the local charts P, Q.

Remark 1 (Relative strengths and weaknesses) One great advantage of the analytic formulation is that, since it is equivalent to a two-point boundary value problem, we can utilize the Newton method to find very accurate solutions—often on the order of machine precision. The formulation as a boundary value problem also lends itself to numerical continuation schemes, which are very useful for exploring the parameter space. The disadvantages are twofold. First, in this formulation it is necessary to begin the Newton iteration with a fairly good approximate solution and this raises the question: *Where do the approximate solutions come from?* Second, it is difficult to rule out solutions using the BVP approach.

In the geometric approach, there is no need to make a guess. Instead, one moves along the stable and unstable manifolds and identifies connections by locating intersections in phase space. At the same time, the geometric approach allows one to rule out connecting orbits by showing that a particular region of phase space does not contain any intersections. The difficulty with the geometric perspective is that it provides information only as good as our knowledge of the embeddings of the stable/unstable manifolds. Computing embeddings of invariant manifolds is challenging, and methods tend to decrease in accuracy the farther from the equilibrium they are applied.

The important point, from the perspective of the present work, is that these two approaches complement one another. The geometric formulation is good for locating and ruling out connections, while the analytic formulation is good for refining approximations and for continuation with respect to parameters. This suggests the approach of the present work: namely that we use the two formulations in concert, playing the strengths of one against the weaknesses of the other as appropriate.

We remark that in many applications it is convenient to examine the intersections of the invariant manifolds in an intermediate surface of section. This is especially true for two-degree-of-freedom systems as the section intersected with the energy level leads to a two-dimensional image which is easy to visualize. Often an appropriate section is suggested by the geometry of the problem, or by the goals of a particular space mission. We refer the interested reader to the works (Koon et al. 2000; Canalias and Masdemont 2006; Barrabés et al. 2009) for examples and fuller discussion.

3 Numerical computation of the stable/unstable manifolds

The results of Sect. 2.2 show that *for most parameter values*, the CRFBP has either three or four saddle-focus equilibria—though for some parameters it may have only two, or one, or none. For a given saddle-focus equilibrium with fixed values of the mass parameters, we compute the invariant manifolds in two steps. First, we find a high-order expansion of an initial local chart containing the equilibrium solution. Then we use a high-order Taylor integration scheme to advect the boundary of the initial chart one subarc at a time. The second

step is repeated until a certain integration time has been reached, or until some error tolerance has been exceeded. Along the way, it is sometimes necessary to subdivide boundary arcs in order to manage the truncation errors.

Our computation of the initial chart employs the parameterization method, which is reviewed in Sect. 3.1. Advection of the boundary uses a Taylor integration scheme similar to the one developed in Kalies et al. (2018), but adapted to the problem at hand. Both procedures exploit differential-algebraic manipulations of formal power series, and these manipulations are delicate due to the presence of the minus two-thirds of power in the nonlinearity of the CRFBP vector field.

One technique for manipulating power series of several complex variables involves *automatic differentiation* combined with the radial gradient. This procedure is developed in Haro et al. (2016) and is reviewed in "Appendix B." Another technique involves appending additional variables and equations to the problem, so that the enlarged field is polynomial and equivalent to the original CRFBP on a certain submanifold. This option is discussed at length for the CRFBP in Kepley and Mireles James (2018) which also includes a more precise definition of what "equivalent" means here. See also Lessard et al. (2016) and Rabe (1961).

3.1 Parameterization method for the local invariant manifold

We now review the parameterization method adapted to the needs of the present work, namely for a stable/unstable manifold attached to a saddle-focus equilibrium in $\mathbb{R}^4$. Much more general treatment of the parameterization method is found in Cabré et al. (2003a, b, 2005). See also the book on this topic (Haro et al. 2016).

Let $\mathbf{x}_0 \in \mathbb{R}^4$ denote a saddle-focus equilibrium point. Specifically, we suppose $f(\mathbf{x}_0) = 0$,

$$\lambda_{1,2} = -\alpha \pm i\beta,$$

with $\alpha, \beta > 0$ denotes the stable eigenvalues for $Df(\mathbf{x}_0)$, and $\xi_{1,2} \in \mathbb{C}^4$ denotes a choice of associated complex conjugate eigenvectors.

Since the eigenvalues are complex, it is convenient to look for a complex parameterization of a local stable manifold. Let

$$D^2 = \left\{(z_1, z_2) \in \mathbb{C}^2 \,:\, |z_j| < 1,\ j = 1, 2\right\}$$

denote the unit complex polydisc. We look for a parameterization $P\colon D^2 \to \mathbb{C}^4$ satisfying the infinitesimal conjugacy given by

$$DP(\mathbf{z})\Lambda\mathbf{z} = f(P(\mathbf{z})), \tag{5}$$

where $\mathbf{z} = (z_1, z_2)^{\mathrm{T}}$, and

$$\Lambda = \begin{pmatrix} \lambda_1 & 0 \\ 0 & \lambda_2 \end{pmatrix}.$$

Equation (5) is subject to the first-order constraints

$$P(0,0) = \mathbf{x}_0, \quad \text{and} \quad \frac{\partial}{\partial z_{1,2}} P(0,0) = \xi_{1,2}. \tag{6}$$

Note that

$$DP(\mathbf{z})\Lambda\mathbf{z} = \lambda_1 z_1 \frac{\partial}{\partial z_1} P(z_1, z_2) + \lambda_2 z_2 \frac{\partial}{\partial z_2} P(z_1, z_2),$$

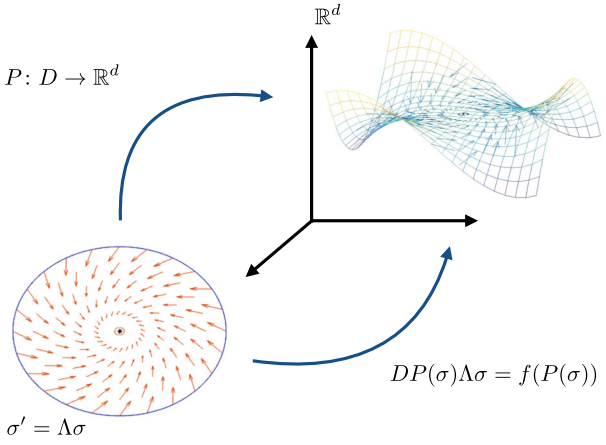

Fig. 5 Geometric interpretation of the parameterization method for differential equations: Eq. (5) requires that the push forward of the vector field Λ by P matches the vector field f on the image of P. A function satisfying this equation is a parameterization of a local stable manifold

is the push forward of the linear vector field by P. The geometric meaning of Eq. (5) is illustrated in Fig. 5.

Let Φ denote the flow generated by f. Any P satisfying Eq. (5) on D^2 also satisfies the flow conjugacy

$$\Phi(P(z_1, z_2), t) = P(e^{\lambda_1 t} z_1, e^{\lambda_2 t} z_2), \quad (z_1, z_2) \in D^2. \tag{7}$$

In particular, if P satisfies both Eq. (5) and the constraints of Eq. (6), then for any $(z_1, z_2) \in D^2$ it follows that

$$\begin{aligned} \lim_{t\to\infty} \Phi(P(z_1, z_2), t) &= \lim_{t\to\infty} P(e^{\lambda_1 t} z_1, e^{\lambda_2 t} z_2) \\ &= P(0, 0) \\ &= \mathbf{x}_0, \end{aligned}$$

so that $P(D^2) \subset W^s(\mathbf{x}_0)$. Combining this with the fact that the image of P contains $\mathbf{x}_0$ and is tangent to the stable eigenspace at $\mathbf{x}_0$ we see that P parameterizes a local stable manifold for $\mathbf{x}_0$. Moreover, we recover the dynamics on the manifold through the conjugacy.

When the vector field f is analytic near $\mathbf{x}_0$, then $W^s(\mathbf{x}_0)$ is an analytic manifold, and it makes sense to look for an analytic chart of the form

$$P(z_1, z_2) = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} p_{m,n} z_1^m z_2^n,$$

with $p_{m,n} \in \mathbb{C}^4$ for all $m, n \in \mathbb{N}$. Since we are interested in the real image of the chart, we look for a solution of Eq. (5) with

$$P(z, \bar{z}) \in \mathbb{R}^4,$$

for all $|z| < 1$. This is achieved whenever the power series coefficients of the solution satisfy

$$p_{n,m} = \overline{p_{m,n}}, \tag{8}$$

for all $(m, n) \in \mathbb{N}^2$. The real parameterization $\tilde{P}\colon B \to \mathbb{R}^4$ is recovered using complex conjugate variables

$$\tilde{P}(\sigma_1, \sigma_2) = P(\sigma_1 + i\sigma_2, \sigma_1 - i\sigma_2).$$

Elementary proofs of the facts discussed in this section are found, for example, in Kepley and Mireles James (2018).

3.2 Power series solution of Eq. (5)

We describe three methods for computing the power series coefficients of an analytic solution of the invariance equation given in Sect. 3.1. Combining these methods leads to very efficient numerical methods.

3.2.1 Solution by power matching

Plugging the unknown power series expansion for P into Eq. (5) leads to

$$\sum_{m=0}^{\infty}\sum_{n=0}^{\infty}(m\lambda_1 + n\lambda_2)p_{m,n}z_1^m z_2^n = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty}[f \circ P]_{m,n}z_1^m z_2^n.$$

It is shown in Cabré et al. (2003a) (see also the discussion in Haro et al. 2016) that when we match like powers and isolate $p_{m,n}$ we are led to an expression of the form

$$\begin{aligned}(m\lambda_1 + n\lambda_2)p_{m,n} &= [f \circ P]_{m,n} \\ &= Df(p_{0,0})p_{m,n} + R(P)_{m,n},\end{aligned}$$

where $R(P)_{m,n}$ depends in a nonlinear way on coefficients $p_{j,k}$ with $0 \leq j + k < m + n$. Isolating the variable $p_{m,n}$ on the left leads to the *homological equations*

$$[Df(\mathbf{x}_0) - (m\lambda_1 + n\lambda_2)\mathrm{Id}]\, p_{m,n} = -R(P)_{m,n}. \tag{9}$$

Remark 2 (The formal solution is well defined) Observe that Eq. (9) is linear in $p_{m,n}$ and has a unique solution as long as $m\lambda_1 + n\lambda_2$ is not an eigenvalue of $Df(\mathbf{x}_0)$. But $\lambda_2 = \overline{\lambda_1}$, and since any remaining eigenvalues are assumed to be unstable, we have that $m\lambda_1 + n\lambda_2$ is never an eigenvalue of $Df(\mathbf{x}_0)$. Hence the matrix on the left-hand side of the homological equation (9) is invertible for all $m + n \geq 2$.

Given any first-order data as in the constraint Eq. (6), the homological equations are uniquely solvable to all orders and the corresponding formal series solution of Eq. (5) is well defined. Since each Taylor coefficient is uniquely determined by the homological equations (9), it follows that the formal series solution is unique up to the choice of the scalings of the eigenvectors in Eq. (6). Solving the homological equations recursively to order $N \geq 2$ provides a polynomial chart P^N which approximately parameterizes the local stable manifold.

Remark 3 (Reality of the parameterization) Taking complex conjugates in the homological equations (9) shows that the coefficients $p_{m,n}$ have the symmetry of Eq. (8).

3.2.2 A Newton scheme

A quadratic convergence scheme for Eq. (5) is obtained as follows. Define the nonlinear operator

$$\Psi[P](\sigma) = DP(\sigma)\Lambda\sigma - f(P(\sigma)),$$

where f is the CRFBP vector field, and note that a zero of Ψ is a solution of Eq. (5). Moreover, we note that, at least formally, the Fréchet derivative is given by

$$D\Psi[P]H(\sigma) = DH(\sigma)\Lambda\sigma - Df(P(\sigma))H(\sigma).$$

In fact this is the correct Fréchet derivative of Ψ when, for example, we consider Ψ defined on a Banach space of analytic functions, see Cabré et al. (2003a, 2005), de la Llave and Mireles James (2012).

Choose P_0 an approximate zero of Ψ, and define the sequence

$$P_{n+1} = P_n + \Delta_n,$$

where Δ_n is the formal series solution of the linear equation

$$D\Psi[P]\Delta = -\Psi[P]. \tag{10}$$

If P_0 is a good enough approximate solution of Eq. (5) we expect P_n to converge quadratically to a zero of Ψ. The linear operator $D\Psi[P]$ nonconstant coefficient, and Eq. (10) may be solved recursively via the following power matching scheme. Define

$$\Delta(\sigma_1, \sigma_2) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} \Delta_{m,n}\sigma_1^m\sigma_2^n,$$
$$Df(P(\sigma)) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} A_{m,n}\sigma_1^m\sigma_2^n,$$

and

$$-\Psi(P(\sigma)) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} q_{m,n}\sigma_1^m\sigma_2^n.$$

Here $\Delta_{m,n}, q_{m,n} \in \mathbb{C}^4$, and A_{mn} are 4×4 complex valued matrices for all $(m, n) \in \mathbb{N}^2$. Plugging these series expansions into Eq. (10) leads to

$$\sum_{m+n\geq 2}\left((m\lambda_1 + n\lambda_2)\Delta_{m,n} - \sum_{j=0}^{m}\sum_{k=0}^{n} A_{m-j,n-k}\Delta_{j,k}\right)\sigma_1^m\sigma_2^n = \sum_{m+n\geq 2} q_{m,n}\sigma_1^m\sigma_2^n,$$

or, upon matching like powers,

$$(m\lambda_1 + n\lambda_2)\Delta_{m,n} - \sum_{j=0}^{m}\sum_{k=0}^{n} A_{m-j,n-k}\Delta_{j,k} = q_{m,n},$$

for all $m + n \geq 2$. We note that the sum contains one term of order Δ_{mn}, appearing when $j = m$ and $k = n$. That is

$$\sum_{j=0}^{m}\sum_{k=0}^{n} A_{m-j,n-k}\Delta_{j,k} = A_{00}\Delta_{mn} + \text{“lower-order terms of } \Delta\text{”}.$$

Let

$$\tilde{\delta}_{j,k}^{m,n} = \begin{cases} 1 & j < m \text{ or } k < n \\ 0 & j = m \text{ and } k = n \end{cases}.$$

Then we use $\tilde{\delta}^{m,n}_{j,k}$ to extract terms of order (m,n) from the sum and write the equation for Δ_{mn} as

$$(m\lambda_1 + n\lambda_2)\Delta_{m,n} - A_{0,0}\Delta_{m,n} - \sum_{j=0}^{m}\sum_{k=0}^{n} \tilde{\delta}^{m,n}_{j,k} A_{m-j,n-k}\Delta_{j,k} = q_{m,n}.$$

Recall that $A_{0,0} = Dg(0) = Df(\mathbf{x}_0)$, so that rearranging terms leads to the linear equations

$$(Df(\mathbf{x}_0) - (m\lambda_1 + n\lambda_2)\mathrm{Id})\,\Delta_{m,n} = -q_{m,n} - \sum_{j=0}^{m}\sum_{k=0}^{n} \tilde{\delta}^{m,n}_{j,k} A_{m-j,n-k}\Delta_{j,k}, \tag{11}$$

for $m+n \geq 2$. Since the right-hand side of Eq. (11) is exactly the right-hand side appearing in the homological equations (9) of Sect. 3.2.1, arguing as in Remarks 2 and 3 shows that the equations of (11) are uniquely solvable for all $m+n \geq 2$ just as before, and that the resulting power series coefficients have the desired symmetry. Then this Newton scheme is well defined on the space of formal power series.

3.2.3 A pseudo-Newton scheme

While the Newton scheme of the previous section converges rapidly (in the sense of the number of necessary iterations), solving the required nonconstant coefficient linear equations is expensive. In this case the overall computation may be slow just because of the cost of computing the individual corrections. The iterations can be speeded up as follows.

First, we note that

$$D\Psi[P]\Delta(\sigma) = D\Delta(\sigma)\Lambda\sigma - Df(\mathbf{x}_0)\Delta(\sigma) + \text{“higher-order terms”},$$

and we define a new iterative scheme

$$P_{k+1}(\sigma) = P_k(\sigma) + \tilde{\Delta}_k(\sigma),$$

where $\tilde{\Delta}_k$ is a solution of the *constant coefficient linear equation*

$$D\tilde{\Delta}_k(\sigma)\Lambda\sigma - Df(\mathbf{x}_0)\tilde{\Delta}_k(\sigma) = -\Psi(H_k).$$

On the level of power series, this equation becomes

$$\sum_{m=0}^{\infty}\sum_{n=0}^{\infty}[(m\lambda_1 + n\lambda_2)\mathrm{Id} - Df(\mathbf{x}_0)]\,\tilde{\Delta}_{m,n}\sigma_1^m\sigma_2^n = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} q_{m,n}\sigma_1^m\sigma_2^n,$$

and matching like powers yields the linear equations

$$[Df(\mathbf{x}_0) - (m\lambda_1 + n\lambda_2)\mathrm{Id}]\,\tilde{\Delta}_{m,n} = -q_{m,n}.$$

These homological equations uniquely determine the coefficients $\tilde{\Delta}_{m,n}$ and have the virtue of being “diagonal” in Taylor coefficient space. In practice we find that the pseudo-Newton scheme requires more iterates than the Newton method to converge. However, a single iteration step is much faster and for reasonable values of N the pseudo-Newton method is faster overall. We discuss this further below.

Remark 4 In practice the linear approximation of P by the eigenvectors provides a good initial guess for the Newton and pseudo-Newton schemes, especially when computations are started “from scratch.” However, within the context of calculations based on parameter

Table 1 Runtime data for the parameterization method: here the manifolds are first computed to order N_0 in order to measure the exponential decay rate associated with the Taylor coefficients

N	N_0	τ	$\max_{m+n=N} \|p_{m,n}\|$	Recursion (s)	Newton (s)	Pseudo-Newton (s)
10	5	0.024	8.4×10^{-15}	3.1	0.49	0.45
20	5	0.13	2.9×10^{-12}	3.3	0.94	0.73
20	10	0.09	9.9×10^{-16}	3	0.7	0.62
30	10	0.15	2.9×10^{-15}	3.5	2.1	1.2
30	15	0.15	1.2×10^{-15}	3.8	1.6	1.1
30	20	0.15	1.2×10^{-15}	4.4	1.7	0.93
40	20	0.21	4.5×10^{-15}	4.6	4.0	2.01
70	30	0.27	4.1×10^{-15}	9.9	28.1	14.9
70	50	0.27	1.2×10^{-15}	10.9	30.8	12.2

These data are used to determine the optimal eigenvector scaling, and then the coefficients are computed to order N in a "production run." The initial computation is always computed to order N_0 by recursion. Then the production run is computed either by recursion, by Newton, or by the pseudo-Newton method. The computations were performed on a MacBook Air with a 1.8 GHz Intel Core i5 processor and 8 GB of 1600 MHz RAM running the version 10.12.6 of the Sierra operating system with MATLAB version 2017*b*. The same computations run about twenty percent faster on a Mac Pro desktop with a 3.7 GHz quad-core Intel Xenon E5 processor and the same version of MATLAB

continuation, we will take P_0 as the high-order parameterization from the previous mass values.

Indeed, it seems that the best results are obtained by a "hybrid" approach. That is, we compute an initial guess P_0 by recursively solving Eq. (9) to some fixed order, N_0. Then, we refine this approximation via the Newton or pseudo-Newton scheme to obtain a polynomial approximation to order, $N > N_0$. The runtime performance for this hybrid approach is recorded in Table 1.

Remark 5 (Quantifying the errors) Suppose that the polynomial

$$P^N(z_1, z_2) = \sum_{0 \le m+n \le N} p_{m,n} z_1^m z_2^n,$$

is an approximation solution of Eq. (5). One way to measure the quality of the approximation is to measure the *defect* associated with P^N defined by the quantity

$$\text{defect}(P^N) = \sup_{\mathbf{z} \in D^2} \left\| DP^N(\mathbf{z}) \Lambda \mathbf{z} - f(P(\mathbf{z})) \right\|_{\mathbb{C}^4}.$$

This quantity could be approximated by evaluating on a mesh of points in D. On the other hand, we can use the fact that for power series on the unit disk we have the bound

$$\sup_{\mathbf{z} \in D^2} \|g(\mathbf{z})\|_{\mathbb{C}^4} \le \sum_{m+n=0}^{\infty} \|a_{m,n}\|_{\mathbb{C}^4},$$

where the infinite sum can be approximated by a finite sum. Then another useful a-posteriori indicator is obtained by choosing an $N' > N$ and computing the quantity

$$\varepsilon_{\text{a-posteriori}} = \sum_{0 \le m+n \le N'} \left\| (m\lambda_1 + n\lambda_2) p_{m,n}^N - [f \circ P^N]_{m,n} \right\|_{\mathbb{C}^4},$$

Table 2 Taylor order, scaling, and error bounds for the parameterization method: table reports the numerical defect and numerical conjugacy error associated with the local stable/unstable manifold parameterization for a number of example computations, as functions of the polynomial order and eigenvector scalings

N	$\|\xi\|$	$\max_N \|p_{m,n}\|$	Defect	Conj error	# Machine eps
1	10^{-8}	10^{-8}	2.3×10^{-15}	3.9×10^{-15}	18
1	10^{-6}	10^{-6}	1.4×10^{-11}	1.3×10^{-11}	5,855
1	10^{-4}	10^{-4}	1.39×10^{-7}	1.3×10^{-7}	5.9×10^{8}
1	1	1	13.9	4.1	2.3×10^{16}
2	10^{-6}	1.4×10^{-12}	9.2×10^{-16}	1.5×10^{-14}	68
3	10^{-4}	2.9×10^{-12}	6.5×10^{-15}	1.6×10^{-14}	73
4	10^{-3}	3.9×10^{-12}	7.9×10^{-14}	2.8×10^{-14}	356
5	10^{-3}	2.5×10^{-15}	1.3×10^{-15}	1.6×10^{-14}	73
7	10^{-2}	3.6×10^{-13}	2.9×10^{-13}	8.6×10^{-14}	1036
10	10^{-2}	1.2×10^{-18}	9.3×10^{-16}	1.6×10^{-14}	73
15	10^{-1}	6.2×10^{-12}	3.9×10^{-10}	5.3×10^{-11}	1.7×10^{6}
20	10^{-1}	4.1×10^{-15}	3.9×10^{-13}	4.9×10^{-14}	1, 756
25	10^{-1}	2.9×10^{-18}	1.6×10^{-15}	5.6×10^{-14}	253
35	1.5×10^{-1}	2.2×10^{-18}	2.4×10^{-15}	1.1×10^{-13}	495
45	2×10^{-1}	3.1×10^{-17}	3.3×10^{-14}	1.0×10^{-13}	450
65	2.5×10^{-1}	2.6×10^{-17}	7.7×10^{-14}	1.1×10^{-13}	495

The first column records the polynomial order N of the approximation P^N. The second column is the magnitude $\|\xi\|$ of the eigenvector (the scaling of the parameterization). These are the inputs which must be specified by the user in any computation using the parameterization method. The third column reports the corresponding bound on magnitude of the highest order Taylor coefficients. The fourth and fifth columns record the numerical defect and conjugacy error, respectively. The sixth column reports the worst of these two quantities measured in multiples of machine epsilon

where $p^N_{m,n}$ are the power series coefficients of P^N, and $[f \circ P^N]_{m,n}$ are the coefficients of $f(P^N(\mathbf{z}))$. Of course this bounds also the real image of P^N.

If f is a polynomial of order K, then we take $N' = KN$. If f is not a polynomial, then the power series for $f \circ P^N$ has infinitely many terms even though P^N is polynomial. Then we choose $N' > N$ somewhat arbitrarily. Note that $p^N_{m,n}$ are zero when $m + n > N$, so that eventually the sum involves only the coefficients of the composition.

Yet another useful error indicator is obtained by considering the dynamical conjugacy of Eq. (7). Since the true solution satisfies the dynamical conjugacy exactly, we consider also the quantity defined by

$$\text{conjugacyDefect}(P^N) = \sup_{\mathbf{z}\in D^2} \sup_{t>0} \left|\left|\Phi(P(z_1, z_2), t) - P(e^{\lambda_1 t} z_1, e^{\lambda_2 t} z_2)\right|\right|_{\mathbb{C}^4}.$$

To approximate this quantity, we fix $\tau > 0$ and let Φ_{num} denote a numerical integrator and z_k, $1 \le k \le K$ be a mesh of the complex circle so $|z_k| = 1$. Define the indicator

$$\varepsilon_{\text{conjugacy}} = \max_{1\le k\le K} \left\| \Phi_{\text{num}}(P^N(z_k, \overline{z_k}), \tau) - P^N(e^{\lambda_1 \tau} z_k, e^{\lambda_2 \tau}\overline{z_k}) \right\|_{\mathbb{C}^4}.$$

Error bounds for a number of example computations are recorded in Table 2.

Remark 6 (Eigenvector scaling and coefficient decay) Solutions of Eq. (5) are only unique up to the choice of the scalings of the eigenvectors and this freedom is exploited in our numerical algorithms. Indeed, this is the reason we can always take our domain to be the unit disk. The results in Table 2 describe the dependence of the numerical errors on the approximation order and the eigenvector scalings. These numerical experiments lead to the following heuristic. If we scale the eigenvectors so that the final coefficients—that is the N-th-order coefficients of P^N—are on the order of machine epsilon, then we obtain a-posteriori errors on the order of machine epsilon.

3.3 Integration of analytic arcs

In Sect. 4 we present a scheme for computing an atlas for the stable/unstable manifolds which relies on integrating analytic arcs of initial conditions by the flow generated by f. We describe this integrator in terms of power series expansions. Let us assume that $\gamma : (-1, 1) \to \mathbb{R}^4$ is an analytic arc with power series expansion

$$\gamma(s) = \sum_{n=0}^{\infty} \gamma_n s^n \qquad \gamma_n \in \mathbb{R}^4.$$

Denote the formal series expansion

$$\Gamma(s, t) = \Phi(\gamma(s), t) = \sum_{m=0}^{\infty} \sum_{n=0}^{\infty} a_{m,n} s^n t^m \qquad a_{m,n} \in \mathbb{R}^4.$$

Here, we use the variables (s, t) in place of (z_1, z_2) to emphasize the intuition that s corresponds to the "spatial" parameterization along the initial data, and t corresponds to the "time" parameterization along the flow. In other words, we consider Γ as the solution of the parameterized family of initial value problems

$$\frac{\mathrm{d}}{\mathrm{d}t}\Gamma(s, t) = f(\Gamma(s, t)), \quad \Gamma(s, 0) = \gamma(s), \quad s \in (-1, 1).$$

Substituting the formal series into this IVP and matching like powers leads to the recursion relations

$$a_{m+1,n} = \frac{1}{m+1}[f \circ \Gamma]_{m,n}, \qquad a_{0,n} = \gamma_n,$$

which allow us to compute the coefficients of Γ to arbitrary order using the same methods described in Sect. 3.2. We also note that the precision of these formal series computations depend on convergence and domain decomposition of these series expansions which has not been addressed and will also be taken up in the following section.

4 Building an atlas for the local stable/unstable manifold

In this section, let $W^*(\mathbf{x}_0)$ denote an invariant stable/unstable manifold for a saddle-focus equilibrium, $\mathbf{x}_0$. Our goal is to describe an algorithm for producing an atlas of chart maps which parameterizes a large portion of the invariant manifold. The union of the images of these maps is a piecewise parameterization of a two-dimensional subset of $W^*(\mathbf{x_0})$. Our procedure is iterative and at each step outputs a (strictly) larger piecewise parameterization.

It is important to emphasize that our computations are carried out only to finite order. In particular, the charts described in this section are analytic functions of two complex variables. However, in practice we fix $(M, N) \in \mathbb{N}^2$, and for each chart we compute a finite polynomial approximation of order (M, N). Nevertheless, throughout this section we denote these analytic charts and their polynomial approximations using the same notation. We end this section by outlining methods for reliably, efficiently, and automatically computing these atlases. This includes algorithms for estimating and controlling truncation errors, identifying Taylor series blowup, domain decomposition, and stiffness.

4.1 Iterative method for computing charts

Before elaborating on the technical details of our method, we briefly describe the overall strategy. Starting from the parameterized local invariant manifolds obtained via methods described in Sect. 3.1, we want to build an even larger representation of the manifold. There are many ways to grow such a representation. We could, for example, simply integrate a collection of initial conditions meshing the boundary of the parameterization. However, as is well known, the exponential separation of initial conditions will force these orbits apart and eventually degrade the description of the manifold. Instead, we mesh the boundary into a collection of one-dimensional arcs and advect each of these under the flow. Propagating these arcs maintains the fidelity of the representation, and leads to new "patches" of the manifold.

Since the initial chart is parameterized by a high-order polynomial, we would like the same representation for new charts. To this end we develop a high-order Taylor integration scheme which applies to analytic arcs of initial conditions. This results in a power series representation of the flow of a boundary arc, and we take this as our next chart. After advecting each one of the boundary arcs, we have a new and strictly larger representation of local stable/unstable manifold. The idea is illustrated in Fig. 4.

After one step of this procedure, we have moved the boundary of the local invariant manifold. In some cases, the image of the advected arc undergoes excessive stretching due to the exponential separation of initial conditions. This stretching in phase space is matched by a corresponding blow up in coefficients of the Taylor expansion, and the computations become numerically unstable.

This problem is overcome by occasionally remeshing the boundary of the atlas. This comes at a cost of increasing the number of charts in the next step of the algorithm. Hence, efficiently computing large atlases while controlling numerical error requires automatic algorithms for managing the growth of the power series coefficients, deciding how long to integrate each individual arc, and deciding when and how to subdivide the new boundaries. These topics account for much of the technical details which follow (Fig. 6).

4.1.1 The initial local manifold

The first step in our algorithm is to compute a polynomial approximation of the local parameterization, either by directly solving the homological equations or by iterating the Newton or pseudo-Newton schemes described in Sect. 3.1. Let Γ_0 be a solution of Eq. (5), and D^2 denote the unit polydisc in $\mathbb{C}^2$. Recall that $\Gamma_0 : D^2 \to W^*_{\text{loc}}(\mathbf{x}_0)$ is analytic, and that $\Gamma_0(\partial D^2)$ is flow transverse. In particular, Γ_0 serves as our initial local parameterization, and we refer to it as a zeroth-generation *interior* chart and we write $\Gamma_0(D^2) = W^*_0(\mathbf{x}_0)$.

In practice, we compute Γ_0 to order (N, N) with $N \in \mathbb{N}$ chosen by applying the heuristic methods discussed in Sect. 3.2. This chart is represented in the computer as a polynomial

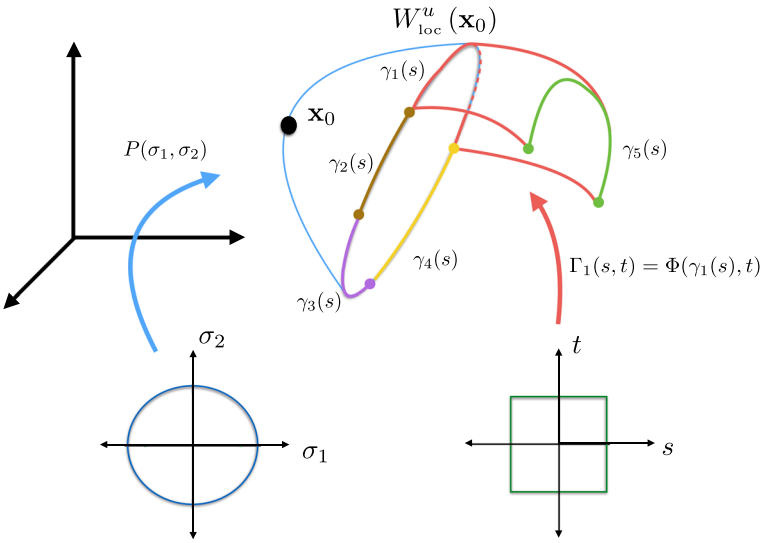

Fig. 6 Building an atlas: here P is a chart for a neighborhood of the equilibrium $\mathbf{x}_0$ computed using the parameterization method. To grow the atlas, we mesh the boundary of the P using a collection of analytic arcs $\gamma_j(s)$. Each of these arcs is advected under the flow Φ to produce a new chart $\Gamma_j(s,t) = \Phi(\gamma_j(s), t)$. The union of P with all the Γ_j is an atlas for a larger local stable/unstable manifold

in two complex variables of total degree $\deg(\Gamma_0) = (N-1)^2$. The truncation error of this approximation is controlled directly by choosing the eigenvector scaling as described in Remark 6, and in practice, is on the order of machine epsilon.

4.1.2 The initial manifold boundary

With Γ_0 in hand, we fix $K_0 \in \mathbb{N}$ and subdivide ∂D into K_0-many analytic segments, each of which has the form, $c_j : [-1,1] \to \partial D$, for $1 \le j \le K_0$. We parameterize $\partial W_0^*(\mathbf{x}_0)$ by defining $\gamma_j(s) = \Gamma_0 \circ c_j(s)$ and we refer to γ_j as a *lifted boundary*. Note that for each $1 \le j \le K_0$, $\gamma_j : [-1,1] \to \partial W^*_{\text{loc}}(\mathbf{x}_0)$ and $\gamma_j([-1,1])$ is a flow transverse arc since Γ_0 is a dynamical conjugacy and the image of c_j is transverse to the linear flow. Now, we define the zeroth-generation boundary to be

$$\partial W_0^*(\mathbf{x}_0) = \bigcup_{j=1}^{K_0} \gamma_j([-1,1]),$$

and refer to each γ_j as a zeroth-generation *boundary* chart.

4.1.3 The next generation

Now, we apply the high-order Taylor advection described in Sect. 3.3 to grow a larger local manifold denoted by $W_1^*(\mathbf{x}_0)$. Specifically, for $1 \le j \le K_0$, we choose $|\tau_j| > 0$, and our advection algorithm takes γ_j, τ_j as input and produces a chart, $\Gamma_{1,j} : D \to W^*(\mathbf{x}_0)$ which satisfies

$$\Gamma_{1,j}(s,t) = \Phi\left(\gamma_j(s), \frac{t}{\tau_j}\right) \qquad \text{for} \quad (s,t) \in [-1,1]^2.$$

In other words, $\Gamma_{1,j}$ parameterizes the advected image of γ_j under the flow over the time interval $[0, \tau_j]$. These new charts are referred to as first-generation interior charts which we add to our atlas to obtain the first-generation local parameterization

$$W_1^*(\mathbf{x}_0) = W_0^*(\mathbf{x}_0) \cup \bigcup_{j=1}^{K_0} \Gamma_{1,j}(D).$$

Note that $\tau_j \neq 0$ and since γ_j is flow transverse, we have $W_0^*(\mathbf{x}_0) \subsetneq W_1^*(\mathbf{x}_1)$ is a strict subset. In fact, transversality of γ_j implies the stronger condition that $\partial W_0^*(\mathbf{x}_0) \subset \mathrm{Int}(W_1^*(\mathbf{x}_0))$, i.e., the manifold has grown through *every* point on the previous boundary.

Remark 7 (Time rescaling) In this description, τ_j serves as a time rescaling of the flow. This allows direct control over the truncation error (in the time direction) and is analogous to the eigenvector scaling for the initial parameterization described in Remark 6. However, choosing this time rescaling is typically more difficult than choosing the eigenvector scaling and we postpone the discussion of this problem to Sect. 4.2.1.

Once the first-generation interior charts are computed by advection, the first-generation boundary arcs are now obtained by evaluation of the time variable. In particular, for $1 \leq j \leq K_0$, the evaluation, $\Gamma_{1,j}([-1, 1], 1) \subset \partial W_1^*(\mathbf{x}_0)$ is a flow transverse arc segment. We perform spatial rescaling as needed (see Remark 8 below) to obtain the next-generation boundary arcs, $\gamma_{1,j} : [-1, 1] \to \partial W_1^*(\mathbf{x}_0)$ where $1 \leq j \leq K_1$ for some $K_1 \geq K_0$ and

$$\gamma_{1,j}([-1, 1]) \subset \Gamma_{1,j'}([-1, 1], 1) \qquad \text{for some} \quad 1 \leq j' \leq K_0$$

is flow transverse. The advection and evaluation algorithms are then iterated to increase the number of charts in the atlas. The L^{th} step in the iteration chain has the form

$$\overset{\cdots}{\mapsto} \partial W_{L-1}^*(\mathbf{x}_0) \overset{\text{advection}}{\longmapsto} W_L^*(\mathbf{x}_0) \overset{\text{evaluation}}{\longmapsto} \partial W_L^*(\mathbf{x}_0) \overset{\cdots}{\mapsto}$$

where $W_L^*(\mathbf{x}_0)$ is parameterized by K_{L-1}-many interior charts (polynomials in both the space and time variables), $\partial W_L^*(\mathbf{x}_0)$ is parameterized by K_L-many boundary charts (polynomials in the space variable only), and $K_{L-1} \leq K_L$.

If we stop iteration, say at the Lth step, then the final atlas,

$$\mathcal{A} = \left\{ \Gamma_0, \bigcup_{j=1}^{K_0} \Gamma_{1,j}, \bigcup_{j=1}^{K_1} \Gamma_{2,j}, \ldots, \bigcup_{j=1}^{K_L} \Gamma_{L,j} \right\},$$

is a collection of $|\mathcal{A}| = 1 + \sum_{l=1}^{L} K_l$-many analytic charts is a piecewise parameterization a portion of the invariant manifold.

Remark 8 (Spatial rescaling) The parameters, $K_0, \ldots, K_L$, control the number of boundary subdivisions, and therefore, allow direct control over scaling in the spatial direction. As in the time rescaling problem, choosing these parameters effectively is a nontrivial problem which we take up in Sect. 4.2.2.

4.2 Convergence, manifold subdivision, and numerical integration

Thus far, we have ignored the issue of convergence for our formal power series computations. The best method for studying this issue is to combine rigorous numerical computations with

a-posteriori analysis and obtain a proof of the existence of an analytic solution and explicit error bounds on the polynomial approximation. Rigorously validated numerical methods for invariant manifold atlases are described in detail in Kalies et al. (2018), Kepley and Mireles James (2018). In the present work we explore the utility of invariant manifold atlases as a purely numerical tool and trade the computer-assisted proof of rigorous error bounds for improved runtime performance.

In the absence of a rigorous validation scheme, we develop more heuristic checks to insure the reliability of the computations. More precisely, we must automatically identify and fix numerical accuracy issues related to numerical Taylor integration. This amounts to rescaling our Taylor coefficients whenever the decay in either space or time becomes too slow. However, this is less straightforward than the eigenvector rescaling for the initial local parameterization described in Remark 6. In particular, it is helpful to consider the rescaling in space and time "directions" separately.

4.2.1 Time-stepping

Recall that at the saddle-focus equilibrium, the stable/unstable eigenvalues occur in complex conjugate pairs. In particular, both eigenvalues in each pair have equal real parts. It follows that identically rescaling each pair of eigenvectors is the ideal strategy. In fact, this strategy is also necessary and sufficient to ensure that the initial parameterization is real-valued, see Van den Berg et al. (2016). Moreover, in the general case of a hyperbolic equilibrium, the real part of each eigenvalue is a measure of the expansion or contraction rate in the direction of its associated eigenvector. Thus, in cases for which they are not equal, the real parts are still explicitly known and the eigenvectors are scaled proportional to these rates.

On the other hand, all but the initial chart in our atlas is obtained via our advection scheme. In this case, neither the expansion/contraction rates, nor their directions are explicitly known. Obtaining these estimates would require solving for the (spatial) derivative of the flow on each chart. For a general vector field defined on $\mathbb{R}^n$, this amounts to increasing the phase space dimension of our ODE solver from n, to $n + n^2$, which would significantly reduce the size of each manifold which is computationally feasible to produce.

Instead, we take an approach similar to Kalies et al. (2018), which describes heuristics for rescaling time and space independent of one another. Specifically, we adopt a time rescaling which ensures that the norm of the M^{th} "coefficient" (with respect to t) for each chart, is less than machine epsilon. Note that for a classical IVP this coefficient is of course just a scalar. However, in our case the coefficient is actually an analytic function of the spatial variable, represented as a power series and the norm of this coefficient is measured using the ℓ^1 norm. This is made more precise in the following section.

This choice is highly conservative, which gives us tight control over the truncation error in the time direction. On the other hand, the spatial rescaling in the present work deviates from the scheme presented in Kalies et al. (2018) and is detailed in Sect. 4.2.2.

4.2.2 Manifold subdivision

Next, we describe the spatial rescaling scheme which we refer to as *manifold subdivision*. We assume that the time rescaling described in the previous section has been carried out on each chart, and our interest is in rescaling each boundary arc to control truncation errors accumulating in the "space direction." This is equivalent to subdividing a manifold since it is reasonable to assume the rescaling will always shrink the domain. Thus, a single boundary arc will give rise to multiple subarcs defined on reduced domains.

To be more precise, we let C^ω denote the collection of real-valued, analytic functions defined on $(-1, 1)$, and let $\mathcal{S}$ denote the collection of real-valued sequences. We define the *Taylor transform*, $\mathcal{T} : C^\omega \to \mathcal{S}$, to be the mapping which sends an analytic function to its sequence of Taylor coefficients centered at $z = 0$. Specifically, if $g \in C^\omega$ has the Taylor expansion,

$$g(z) = \sum_{n=0}^{\infty} a_n z^n \qquad a_n \in \mathbb{R}, \quad z \in (-1, 1),$$

then $\mathcal{T}(g) = \{a_n\} = a \in \mathcal{S}$. Now, we equip $\mathcal{S}$ with the ℓ_1-norm defined by

$$||a||_1 = \sum_{n=0}^{\infty} |a_n|,$$

and we note that elements of $\mathcal{S}$ with finite norm form a closed subalgebra denoted as

$$\ell_1 = \{x \in \mathcal{S} : ||x||_1 < \infty\},$$

and we write $||a||_{\ell_1}$ when we want to emphasize that $a \in \ell_1$ (i.e., we write $||a||_{\ell_1}$ for the norm $||a||_1$ when $||a||_1$ is finite).

We remark that our error analysis is carried out using the ℓ_1-norm due to the efficiency of computing this norm for polynomials. However, if $\bar{g} \approx g$ is a numerical approximation, then the errors we are interested in are of the form

$$||\bar{g} - g||_\infty = \sup_{z \in [-1,1]} \{|\bar{g}(z) - g(z)|\}.$$

We are justified in using the ℓ_1 norm due to the well known result that $||\bar{g} - g||_\infty \leq ||\bar{g} - g||_{\ell_1}$.

Now, suppose $\gamma \in C^\omega$ and assume that $\mathcal{T}(\gamma) = a \in \ell_1$. Since Φ is a nonlinear flow, a typical arc segment undergoes rapid deformation and stretching when advected. This implies that for a single step in our algorithm with the general form,

$$\overset{\cdots}{\mapsto} \gamma \overset{\text{advection}}{\longmapsto} \Gamma \overset{\text{evaluation}}{\longmapsto} \gamma' \overset{\cdots}{\mapsto},$$

we expect both the arc length and curvature of γ' to be larger than for γ. On the level of Taylor coefficients, this statement about deformation/stretching says that if $b = \mathcal{T}(\gamma')$, then in general we expect $||a||_{\ell_1} \leq ||b||_{\ell_1}$. The relationship between this norm and the truncation error implies that advecting an arc adversely impacts the propagation error.

To see this, we recall that in practice our computation stores a truncated polynomial approximation for γ' in the form $\bar{b} = (b_0, \ldots, b_{N-1})$. In order that $\bar{b} \approx b$ is a "good" approximation (in the ℓ_1 topology), $|b_n|$ must be "small" for each $n \geq N$. These higher-order terms correspond to the truncation error for γ' and primarily arise from two sources. One source which we can not control (once N is fixed) is the truncation error associated with γ. However, by inspection of the Cauchy product formula in Eq. (19), it is clear that the polynomial coefficients stored for γ also contribute to this truncation error for γ' after applying the nonlinearity. We refer to these contributions as *spillover* terms.

This observation implies that for $\bar{b} \approx b$ to be a good approximation, we must also require that $|a_n|$ is "small" for each $n > N'$ where $N' < N$ depends on the degree of the nonlinearity. This motivates the following heuristic method for controlling truncation error for propagated arcs. We begin by assuming that a has approximately geometric decay. Specifically, we expect that there exists some $r < 1$ such that the tail of the series defined by γ decays faster than

the geometric series with ratio r. In this case, the truncation error is of order $\mathcal{O}(r^N)$. Now, fix $0 < N' < N$, and we define the tail ratio for a by

$$T_{N'}(a) := \frac{\sum_{n=N'}^{N-1} |a_n|}{\sum_{n=0}^{N-1} |a_n|} = \frac{\left|\left|a - a^{N'}\right|\right|_{\ell_1}}{||a||_{\ell_1}}. \tag{12}$$

Evidently, $T_{N'}(a)$ is small whenever "most" of the ℓ_1 weight of a is carried in the first N'-many coefficients. It follows that if $T_{N'}(a)$ is sufficiently small, then under the action of a nonlinear function, $f : \ell_1 \to \ell_1$, the spillover terms for $f(a)$ remain small. Of course, small is dependent on context and in particular, choices for N' as well as thresholding values for $T_{N'}$ are problem specific. In the present work, we prove it is always possible to control $T_{N'}$.

Remark 9 Strictly speaking, for the CRFBP we have $\gamma = \left(\gamma^{(1)}, \ldots, \gamma^{(4)}\right)$ where each $\gamma^{(j)} \in C^\omega$ is a coordinate for the boundary chart. Similarly, $\mathcal{T}(\gamma) = \left(a^{(1)}, \ldots, a^{(4)}\right) \in \ell_1^4$, and thus the discussion in Sect. 4.2.2 thus far is technically not applicable. However, our restriction to scalar-valued functions is justified by the fact that if $a \in \ell_1^4$, then defining

$$||a||_{\ell_1^4} = \max\left\{\left|\left|a^{(1)}\right|\right|_{\ell_1}, \ldots, \left|\left|a^{(4)}\right|\right|_{\ell_1}\right\}$$

makes ℓ_1^4 into a normed vector space. This choice of norm gives us the freedom to restrict the discussion of remeshing and tail ratios to scalar-valued functions.

Next, we describe our scheme for controlling the tail ratio. This algorithm takes a polynomial representation for γ, defined on $[-1, 1]$ as input, and returns a list of polynomials, $\{\gamma_1, \ldots, \gamma_K\}$, as outputs. The key point is that these polynomials are also defined on $[-1, 1]$, and they can be chosen such that $T_{N'}(\gamma_j)$ is arbitrarily small for $1 \le j \le K$. In this work, we assume the output polynomials are specified as coefficient vectors of length N (i.e., the same degree as the input); however, this is not required.

This gives rise to an additional *remeshing* step in our algorithm which is performed as needed after an evaluation step and prior to an advection step leading to an updated schematic

$$\overset{\cdots}{\longmapsto} \gamma \xrightarrow{\text{remeshing}} \{\gamma_j\}_{1\le j\le K} \xrightarrow{\text{advection}} \{\Gamma_j\}_{1\le j\le K} \xrightarrow{\text{evaluation}} \{\gamma'\}_{1\le j\le K} \overset{\cdots}{\longmapsto}$$

In the remeshing step, the tail ratio for each boundary arc from the previous step is computed and checked against a threshold. Boundary arcs which exceed this threshold are flagged as poorly conditioned, and subdivided into smaller subarcs which satisfy the threshold. The collection of resulting subarcs and well-conditioned arcs from the previous step is passed to the advection step where each results in a separate chart.

Before proving this threshold can always be satisfied, we describe the subdivision algorithm. As noted in Remark 9, it suffices to consider a single coordinate for a parameterized boundary arc. Thus, we assume $\gamma(s) : [-1, 1] \to \mathbb{R}$ is analytic with Taylor series

$$\gamma(s) = \sum_{n=0}^{\infty} a_n s^n,$$

and fix a subinterval, $[s_1, s_2] \subset [-1, 1]$. Define the constants

$$\hat{s} := \frac{s_1 + s_2}{2} \qquad \delta := \frac{s_2 - s_1}{2} \tag{13}$$

and define $\hat{\gamma} : [-1, 1] \to \mathbb{R}$ by

$$\hat{\gamma}(s) = \sum_{n=0}^{\infty} c_n s^n \qquad \text{where} \quad c_n = \delta^n \sum_{k=n}^{\infty} a_k \binom{k}{n} \hat{s}^{k-n}. \tag{14}$$

Then $\hat{\gamma}$ is a parameterization for the arc segment parameterized by γ restricted to $[s_1, s_2]$. In fact, $\hat{\gamma}$ is the Taylor series for γ after recentering at $\hat{s}$ and rescaling by δ which satisfies the functional equation

$$\hat{\gamma}(s) = \gamma(\hat{s} + \delta s) \qquad s \in [-1, 1]. \tag{15}$$

Moreover, the mapping $a \mapsto c$ is a linear transformation on $\mathcal{S}$, and in particular, if $a_n = 0$ for all $n \geq N$, then $c_n = 0$ for all $n \geq N$ also. Now, we prove that we have explicit control over the tail ratio for $\hat{\gamma}$.

Proposition 1 (Controlling tail ratios) *Suppose $\gamma : [-1, 1] \to \mathbb{R}$ is analytic, fix $\hat{s} \in (-1, 1)$, $1 \leq N' \leq N$, and let $\epsilon > 0$. Then there exists $\delta > 0$ such that $T_{N'}(c) < \epsilon$ where c is the truncation to order N for $\hat{\gamma} : [-1, 1] \to \mathbb{R}$ defined by $\hat{s}$, δ as in Eq. (14).*

Proof Define $\gamma^N : [-1, 1] \to \mathbb{R}$ to be the Taylor polynomial obtained by truncating the Taylor series for γ to order N. For $k \in \mathbb{N}$, define the usual C^k-norm on $[-1, 1]$ to be

$$||g||_{C^k} = \max_{0 \leq j \leq k} \left\{ \left|\left| g^{(j)} \right|\right|_{\infty} \right\}.$$

Since γ^N is a polynomial, we have the bound

$$\left|\left| \gamma^N \right|\right|_{C^k} \leq M := \left|\left| \gamma^N \right|\right|_{C^{N-1}} \qquad \text{for all} \quad k \in \mathbb{N}.$$

In particular, for any $\hat{s} \in (-1, 1)$, we have $\left|\gamma^{(n)}(\hat{s})\right| \leq M$, for $0 \leq n \leq (N-1)$, and we define

$$\delta := \min_{N' \leq n \leq N} \left\{ \left(\frac{\epsilon \left|\gamma(\hat{s})\right|}{M(N - N')} \right)^{\frac{1}{n}} \right\}.$$

It follows that

$$\delta^n \left|\gamma^{(n)}(\hat{s})\right| \leq \frac{\epsilon \gamma(\hat{s})}{N - N'} \quad \text{for all} \quad N' \leq n \leq N.$$

Now, let $\hat{\gamma}$ be defined as in Eq. (14). Recall that $\hat{\gamma}$ is also analytic on $[-1, 1]$, and by differentiating Eq. (15) we have the derivative formula, $\hat{\gamma}^{(n)}(s) = \delta^n \gamma^{(n)}\left(\hat{s} + \delta s\right)$, for all $n \in \mathbb{N}$. By Taylor's theorem, we obtain another explicit formula for c_n given by

$$c_n = \frac{\hat{\gamma}^{(n)}(0)}{n!} = \frac{\delta^n \gamma^{(n)}(\hat{s})}{n!},$$

and we note that $c_0 = \hat{\gamma}(0) = \gamma(\hat{s})$ does not depend on δ. We have the estimate for the tail ratio of $\hat{\gamma}$:

$$\begin{aligned} T_{N'}(c) &= \frac{1}{||c||_{\ell_1}} \sum_{n=N'}^{N-1} |c_n| \\ &= \frac{1}{||c||_{\ell_1}} \sum_{n=N'}^{N-1} \frac{\delta^n \left|\gamma^{(n)}(\hat{s})\right|}{n!} \end{aligned}$$

$$\leq \frac{1}{|c_0|} \sum_{n=N'}^{N-1} \frac{\epsilon \left|\gamma(\hat{s})\right|}{N - N'}$$

$$= \epsilon$$

which completes the proof. □

Proposition 1 establishes the fact that we may reparameterize γ on subintervals of $[-1, 1]$ with width, 2δ, and that as $\delta \to 0$ the tail ratio also approaches zero. We note that δ does not depend on the subinterval, and therefore, for a fixed ϵ the number of required subarcs is finite. In particular, no more than $K = \lceil \frac{2}{\delta} \rceil$ subarcs are required. To summarize the usefulness of this result, we present the following algorithm for controlling the spatial truncation error which was implemented for the atlases in this work.

1. Fix a threshold $0 < \epsilon \ll 1$, a cutoff $1 \leq N' < N$, and $K \in \mathbb{N}$. The threshold and cutoff are both chosen based on the alignment of γ with the flow, the degree of the nonlinearity in f, and the truncation size. In practice, these are problem specific choices which require some ad hoc experimentation in order to balance computational efficiency and truncation error.
2. Following each evaluation step in our algorithm, a boundary arc has the form $\gamma : [-1, 1] \to \mathbb{R}$ which is stored in the computer as a polynomial approximation, $\bar{a} = (a_0, \ldots, a_{N-1})$. If $T_{N'}(\bar{a}) < \epsilon$, continue to the advection step.
3. If $T_{N'}(\bar{a}) \geq \epsilon$, specify a partition of $[-1, 1]$ into K-many subintervals by choosing their endpoints, $\{s_0, s_1, \ldots, s_K\}$. Apply the formula in Eq. (14) to obtain $\{\gamma_1, \ldots, \gamma_K\}$ where for $1 \leq j \leq K$, $\gamma_j(s) = \gamma(\hat{s}_j + \delta_j s)$ where $\hat{s}_j = \frac{s_j + s_{j-1}}{2}$ and $\delta_j = \frac{s_j - s_{j-1}}{2}$.
4. Each resulting subarc which satisfies the tail ratio threshold passes to the advection step. Subarcs which violate the threshold are subdivided again by repeating step 3. By Proposition 1, this condition is eventually met for every subarc and the algorithm proceeds to the advection step.

4.2.3 Stiffness

The final numerical consideration which we address is the stiffness problem. We recall that the CRFBP vector field is analytic away from the primary masses which correspond to singularities of Eq. (3). Since this system is Hamiltonian, any trajectory which collides with one of these primaries must blow up in finite time. However, smooth trajectories may pass arbitrarily close to these primaries and as they do, the velocity coordinates, $\dot{x}, \dot{y}$, become arbitrarily large.

Recall that a single boundary arc, $\gamma : [-1, 1] \to \mathbb{R}^4$, is a parameterized manifold of initial data. Then its advected image, $\Gamma : [-1, 1] \times [0, 1] \to \mathbb{R}^4$, is a parameterized bundle of trajectory segments. For any $s_0 \in [-1, 1]$, $\Gamma(s_0, t)$ parameterizes the trajectory passing through $\gamma(s)$ over the (nonscaled) time interval, $[0, \tau]$.

Now, suppose that for $s_0 \in [-1, 1]$, the trajectory through $\gamma(s_0)$ passes "close" to a primary at time $t = t_0$. Then, we have

$$||f(\Gamma(s_0, t_0))||_{\mathbb{C}^4} \gg 1.$$

Recalling our time rescaling algorithm described in Sect. 4.2.1, it is clear that controlling truncation in the time direction will require taking increasingly shorter time-steps. Of course, this is not surprising; however, the difficulty arises from the fact that other choices of $s \in [-1, 1]$ often correspond to trajectory segments which remain far away from the primary and our time

rescaling is applies uniformly on $[-1, 1]$. Hence, the advection of the entire boundary chart is slowed dramatically whenever any portion of its image approaches a primary. We refer to these charts as *stiff*. Obviously, this is a major problem for our "breadth-first" approach for computing the manifold atlas. Namely, the integrator gets stuck on the stiff charts causing the computation to stall.

A naive method for dealing with this is to define the *speed* for a boundary chart which is a parameterized curve of the form, $\gamma(s) = (x(s), \dot{x}(s), y(s), \dot{y}(s))$, by

$$S(\gamma) = \sup_{s\in[-1,1]} \left\{ \sqrt{\dot{x}(s)^2 + \dot{y}(s)^2} \right\}, \tag{16}$$

set a threshold, κ, and cease advection of γ whenever $S(\gamma) > \kappa$. While this fixes the problem of computational efficiency, we also lose large portions of the manifold which remain far from the primaries. Instead, we leverage the manifold subdivision procedure which is already introduced in Sect. 4.2.2 to modify the naive algorithm in order to retain these portions of the manifold as follows.

1. Fix a maximum speed threshold, $\kappa > 0$. For each boundary chart, γ, present after the evaluation step, check that $S(\gamma) \leq \kappa$ and if so, continue to the remeshing step.
2. If $S(\gamma) > \kappa$, write $\gamma(s) = (x(s), \dot{x}(s), y(s), \dot{y}(s))$ and compute
$$\left\{ s \in [-1, 1] : \dot{x}(s)^2 + \dot{y}(s)^2 - \kappa^2 = 0 \right\}.$$
Since $\dot{x}, \dot{y}$ are polynomial approximations, this set is a finite collection of roots of a polynomial which we denote by, $\{s_0, \ldots, s_K\}$.
3. For $1 \leq j \leq K$, check that $\dot{x}(s)^2 + \dot{y}(s)^2 - \kappa^2 < 0$ holds on $[s_j, s_{j+1}]$ and if so, compute $\hat{\gamma}_j$ as in Eq. (14) and continue to the remeshing step. Subintervals which fail this check are discarded.

To summarize, our algorithm identifies regions of the manifold boundary which pass close to a primary by checking the maximum speed. Regions which exceed a threshold are cut away, while regions of the nearby boundary continue to be advected. The cut regions cause the apparent holes punched out around each primary in the manifold plots, as in Figs. 7 and 8.

4.3 Computational results: manifold atlases for the triple Copenhagen problem

Performance results for atlas computations at the libration points $\mathcal{L}_0$ and $\mathcal{L}_5$ are given in Tables 3 and 4, respectively. The computations are performed for the case of equal masses, that is for the triple Copenhagen problem. The tables report the advection time—that is the number of time units the boundary of the local parameterizations are integrated—as well as the time required to complete the computations and the number of polynomial charts comprising the atlas. All computations were performed on a MacBook Air laptop running Sierra version 10.12.6, on a 1.8 GHz Intel Core i5, with 8 GB of 1600 MHz DDR3 memory.

The resulting atlases for $\mathcal{L}_0$ and $\mathcal{L}_5$ are illustrated in Figs. 7 and 8 for various integration times. The boundaries for the charts are also shown, making it clear that the computational effort goes up dramatically near the primaries. Note that the chart boundary lines running out of the local parameterizations are actual orbits of the system and hence give a sense of the dynamics on the manifold. The pictures provide some insight into the dynamics of the problem; however, their complexity illustrates the need for more sophisticated search techniques in order to extract further useful qualitative information from the atlases.

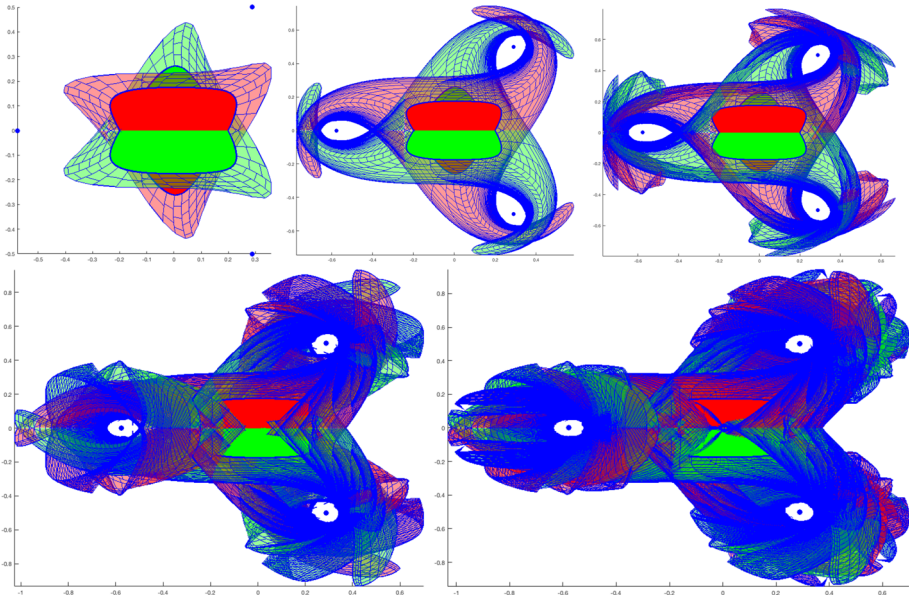

Fig. 7 Atlases at $\mathcal{L}_0$ in the Triple Copenhagen problem: the center of each frame shows the initial local stable chart (green), and unstable chart (red), computed to order 45 using the parameterization method. The three blue dots in each frame represent the location of the primaries. One-third of the boundary of each local manifold is meshed into ten analytic arcs and propagated in time with the boundary of each chart illustrated in blue. By Lemma 1, the rest of the atlas is obtained via $\pm 120^\circ$ rotations. The five frames illustrate the complete atlases obtained after advecting the boundary arcs for ± 0.25, ± 1.0, ± 1.5 time units (top row) and ± 2.5 and ± 4.0 time units (bottom row). After a fairly short integration time, the resulting atlases become complicated enough that visual analysis is difficult or impossible. This complexity motivates development of the post-processing schemes described in Sect. 5.2. Each chart is approximated using Taylor order 20 in space and 40 in time. Runtime and number of charts are given in Table 3

5 Homoclinic dynamics in the CRFBP

In this section, we discuss connecting orbits found for the symmetric $m_1 = m_2 = m_3 = 1/3$ case by searching the manifold atlases computed in the previous section.

5.1 Mining the atlases

Assume we have computed atlases, $\mathcal{A}^{s,u}$, for the stable/unstable manifolds of $\mathbf{x_0}$. We are interested in "mining" the chart data to find transverse connections. Since each atlas is stored as a collection of polynomial charts, it suffices to identify pairwise intersections between stable and unstable charts. Thus, throughout we assume $\Gamma^{s,u} : [-1, 1]^2 \to W^{s,u}(\mathbf{x}_0)$ is a pair of charts which parameterize a portion of the stable/unstable manifold. We write $\Gamma^{s,u}_{1,2,3,4}$ denote the scalar coordinates of each chart. The following theorem whose proof can be found in Kepley and Mireles James (2018) provides a computable condition for verifying transverse intersection of a pair of charts.

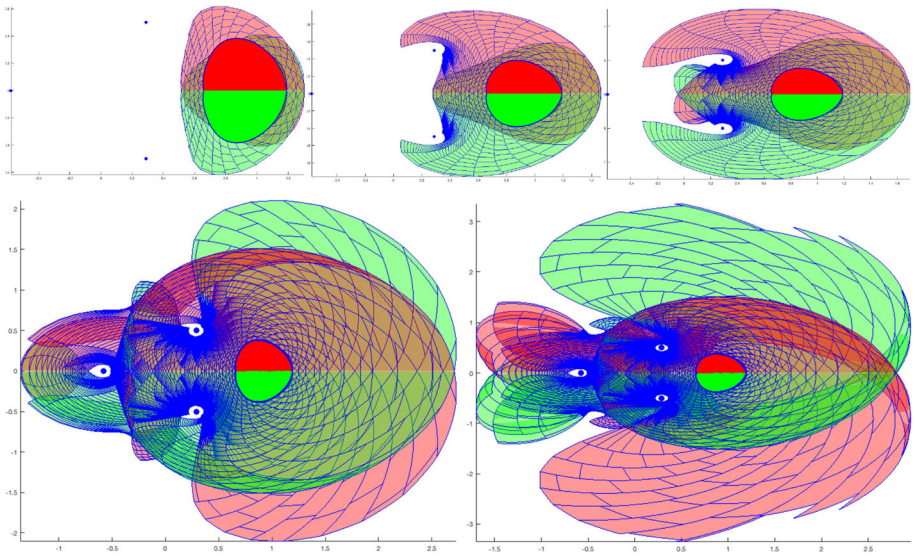

Fig. 8 Atlases at $\mathcal{L}_5$ in the Triple Copenhagen Problem: center of each frame shows the local stable/unstable charts computed using the parameterization method (red and green, respectively). The locations of the primaries are denoted by the blue dots in each frame. Parameterizations approximated to polynomial order 45. The boundary of the local stable/unstable manifold is meshed into thirty analytic arcs. The five frames illustrate the atlases obtained by advecting the boundary arcs by ± 0.75, ± 1.15, ± 1.5 time units (top row) and by ± 3.0 and ± 4.0 time units (bottom row). Again it is difficult to analyze the results by eye, and some post-processing is necessary. Each chart is approximated using Taylor order 20 in space and 40 in time. Runtime and number of charts are given in Table 4

Table 3 Atlas computations at $\mathcal{L}_0$ in the triple Copenhagen problem: each chart is computed to polynomial order 20 in space and order 40 in time

Integration time	Runtime (both manifolds)	# Stable charts	# Unstable charts
± 0.25	17.07 s	39	39
± 0.5	37.9 s	146	146
± 0.75	147 s	497	497
± 1.0	4.75 min	700	700
± 1.5	8.3 min	1579	1579
± 2.5	21.8 min	3530	3493
± 4.0	60.8 min	9372	9295

Unlike the later computations where we have used a speed threshold of 2, here we set the threshold at 3 to better illustrate how the atlas size and computation time grow when propagating the stable/unstable manifolds. Additionally, by Lemma 1 we must only consider one-third of the boundary of each local manifold, so our initial subdivision into 10 sub arcs actually corresponds to mesh which is 3 times finer. Moreover, the computation time is 3 times faster and the final atlas size is 3 times smaller than for the CRFBP with nonequal masses

Theorem 1 *Define* $G\colon [-1,1]^3 \to \mathbb{R}^3$ *by*

$$G(s,t,\sigma) := \begin{pmatrix} \Gamma_1^u(s,t) - \Gamma_1^s(\sigma,0) \\ \Gamma_2^u(s,t) - \Gamma_2^s(\sigma,0) \\ \Gamma_3^u(s,t) - \Gamma_3^s(\sigma,0) \end{pmatrix},$$

Table 4 Atlas computations at $\mathcal{L}_5$ in the triple Copenhagen problem: each chart is computed to polynomial order 20 in space and order 40 in time

Integration time	Runtime (both manifolds)	# Stable charts	# Unstable charts
± 0.5	40.5 s	124	124
± 0.75	57.7 s	216	216
± 1.0	2.3 min	487	487
± 1.5	7.2 min	634	634
± 2.0	15.3 min	1466	1466
± 3.0	32.9 min	2899	2899
± 4.0	53 min	4983	4751

Velocities greater than 2.5 are discarded. We consider the entire boundary of the local stable/unstable manifolds, and we initially divide into 30 sub arcs. Because of this, the computations are roughly 3 times longer than at $\mathcal{L}_0$. But we obtain the manifolds at $\mathcal{L}_{4,6}$ by rotational symmetry

and suppose $(\hat{s}, \hat{t}, \hat{\sigma}) \in [-1, 1]^3$ *satisfies* $G(\hat{s}, \hat{t}, \hat{\sigma}) = 0$. *If* $\Gamma_4^u(\hat{s}, \hat{t})$ *and* $\Gamma_4^s(\hat{\sigma}, 0)$ *have the same sign, then* $\hat{\mathbf{x}} := \Gamma^u(\hat{s}, \hat{t})$ *is homoclinic to* $\mathbf{x}_0$. *Moreover, if* $DG(\hat{s}, \hat{t}, \hat{\sigma})$ *is nonsingular and if* $\nabla E(\hat{\mathbf{x}}) \neq 0$ *(where* E *is the CRFBP energy), then the energy level set is a smooth* 3*-manifold near* $\hat{\mathbf{x}}$ *and the stable/unstable manifolds of* $\mathbf{x}_0$ *intersect transversally in the energy manifold.*

We emphasize that Theorem 1 provides a computable condition for verifying a transverse intersection using rigorous numerics. However, we will use the same theorem to detect transverse intersections in the purely numerical setting of this paper. This is made explicit in the following algorithm utilized in the mining scheme for all results in the present work.

Assume $\Gamma^{s,u}$, G are as defined in Theorem 1. Apply Newton's method to find an approximate root of G. Let $\hat{v} = (\hat{s}, \hat{t}, \hat{\sigma})$ denote an approximate solution with $G(\hat{v}) \approx 0$, and check the following conditions:

1. $\Gamma_4^u(\hat{s}, \hat{t})$, and $\Gamma_4^s(\hat{\sigma}, 0)$ are both "far" from 0.
2. $\Gamma_4^u(\hat{s}, \hat{t})$, and $\Gamma_4^s(\hat{\sigma}, 0)$ have the same sign.

If condition 1 holds without condition 2, then these charts are nonintersecting. In this case, these charts lie on separated portions of the stable/unstable manifolds which are symmetric with respect to the fourth coordinate. We refer to these as "pseudo-intersections." On the other hand, if both conditions hold, then we conclude from Theorem 1 that we have numerically found a transverse homoclinic for $\mathbf{x}_0$ passing through $\Gamma^u(\hat{s}, \hat{t}) = \hat{\mathbf{x}}$.

Note that condition 1 serves two purposes in this setting. First, it serves as an easily computable condition for checking that $\nabla E(\hat{\mathbf{x}}) \neq 0$ as required in the theorem. This follows by noting that

$$\pi_4 \circ \nabla E(\hat{\mathbf{x}}) = \hat{\mathbf{x}}_4 = \Gamma_4^u(\hat{s}, \hat{t}),$$

so it follows that $\nabla E(\hat{\mathbf{x}}) \neq 0$ is satisfied automatically whenever condition 1 is satisfied.

In addition, condition 1 gives us some confidence that the sign difference from condition 2 holds due to transversality of the homoclinic, as opposed to numerical error. Indeed, if condition 1 is not satisfied, then $\Gamma_4^u(\hat{s}, \hat{t})$, and $\Gamma_4^s(\hat{\sigma}, 0)$ take values near zero in which case sign errors for either coordinate are likely due to integration errors. In this case, even if condition 2 is satisfied we are unable to trust the result, and hence unable to conclude whether the zero of G corresponds to a transverse intersection or a pseudo-intersection. Fortunately,

this situation can be remedied as discussed in Remark 11. As a result, we are free to choose our threshold for what is meant by "far" in the statement of condition 1 very conservatively which leads to a great deal of confidence that our mining algorithm returns only transverse homoclinic orbits.

We further increase our confidence in the approximate connection by using it as the input for a BVP solver based on Newton's method, which allows us to refine our approximation to nearly machine precision, and it is the BVP formulation to which we then apply continuation methods. Every connection reported in this section has been so certified, and none of the connections identified from the mining algorithm had a BVP which failed to converge. In other words, the mining algorithm did not return any false homoclinics.

5.2 Efficient atlas mining

It is not desirable to check every pair of charts from each atlas using the above procedure, and we introduce two methods which significantly reduce the number of chart pairs which must be checked via the Newton intersection scheme based on Theorem 1.

5.2.1 The ℓ_1 box approximation

The first method for improving the mining efficiency is to apply a coarse preprocessing step to each pair of charts which must be compared. The main idea is based on the fact that for most pairs of charts which do not intersect, these charts will "obviously" not intersect in the sense that their images in phase space will be very far apart. We exploit this using a fast algorithm for identifying many such pairs and in this case skip the slower Newton-based intersection attempt.

To be more precise, consider an arbitrary polynomial $P : [-1, 1]^2 \to \mathbb{R}$ defined by

$$P(s,t) = \sum_{m=0}^{M} \sum_{n=0}^{N} a_{m,n} s^n t^m \quad a_{m,n} \in \mathbb{R}.$$

We define the ℓ_1 *box* for P to be

$$B_P = [a_{0,0} - r, a_{0,0} + r] \qquad \text{where } r = \sum_{(m,n)\neq(0,0)} \left|a_{m,n}\right|.$$

The significance of B_P is that we have the bound

$$\left|P(s,t) - a_{0,0}\right| \leq r \qquad \text{for all } (s,t) \in [-1,1]^2$$

or equivalently, $P(s,t) \in B_P$ for all $(s,t) \in [-1,1]^2$. Analogously, we extend this to higher dimensions component-wise and apply this to geometrically rule out pairs of charts which can not intersect because their images are "well separated." Specifically, consider a pair of stable/unstable charts

$$\Gamma^s(s,t) = \sum_{m=0}^{M} \sum_{n=0}^{N} a_{m,n} s^n t^m \qquad \Gamma^u(s,t) = \sum_{m=0}^{M} \sum_{n=0}^{N} b_{m,n} s^n t^m.$$

which have ℓ_1 boxes described by rectangles in $\mathbb{R}^4$ and satisfying $\Gamma^s(s,t) \in B_{\Gamma^s}$, and $\Gamma^u(s,t) \in B_{\Gamma^u}$. Then, if the set distance, $d(B_{\Gamma^s}, B_{\Gamma^u})$ is large enough, we can conclude that Γ^s, Γ^u do not intersect.

Using ℓ_1 boxes has two advantages. The first is that computing and checking ℓ_1 boxes for pairwise intersections is much faster than our Newton-like intersection method. This is due to the fact that for each coordinate the box radius, r, is equivalently computed as

$$r = a_{0,0} + ||P||_{\ell_1} - \left|a_{0,0}\right|$$

which is extremely fast to compute using modern implementations. Determining whether two boxes intersect or not is also fast due to efficient interval arithmetic libraries such as the INTLAB library for MATLAB (Rump 1999) which was utilized in our implementation.

The second advantage is that an ℓ_1 box is typically a very coarse enclosure for the true values of P. This "problem" is often referred to as the data-dependence problem or the wrapping effect. In our situation, however, we consider the coarseness to be a feature since it makes our numerical estimates more conservative. Thus, we are able to rule out many pairs of charts which clearly do not intersect without eliminating false negatives.

In practice, a single pairwise ℓ_1 box intersection check is approximately 1,000 times more efficient than the Newton-based scheme and this method rules out around 90 percent of nonintersecting chart pairs. Moreover, the ℓ_1 box for each chart can be computed only once during the atlas construction and stored. This leaves the cost of a single box intersection check as the only significant computational operation.

Finally, we remark that once ℓ_1 boxes have been computed and stored for each chart in both atlases, one can make careful use of the triangle inequality to reduce the computation even further. This provides roughly an additional order of magnitude improvement in the efficiency of our algorithm which could be crucial to the feasibility of mining extremely large atlases. However, we took limited advantage of this fact in the present work.

5.2.2 Fundamental domains

The other main source of efficiency gain in our algorithm relies on using the dynamics explicitly. Recalling our notation in Sect. 4, assume $\mathcal{A}^s$ is the stable manifold atlas which we have computed to include the L_s^{th} generation and let $W_k^s(\mathbf{x}_0)$ denote the k^{th} generation local stable manifold. Then, $W_k^s(\mathbf{x}_0)$ is a fundamental domain for $W^s(\mathbf{x}_0)$. In other words, if $\mathbf{x}(t)$ is any orbit which satisfies $\lim_{t\to\infty} \mathbf{x}(t) = \mathbf{x}_0$ and if $\mathbf{x}(0) \neq \mathbf{x}_0$, then there exists $t_k \in \mathbb{R}$ such that $\mathbf{x}(t_k) \in W_k^s(\mathbf{x}_0)$. Of course, the same claim holds for the unstable manifold. Taken together, if we assume we have computed the unstable manifold, $\mathcal{A}^u$, up to the L_u^{th} generation, then we have the following observation.

Proposition 2 *Let $\mathbf{x}(t)$ be a transverse homoclinic to $\mathbf{x}_0$. Then $\mathbf{x}(t) \in W^s(\mathbf{x}_0) \cap W^u(\mathbf{x}_0)$ for all $t \in \mathbb{R}$. Let $W_0^{s,u}(\mathbf{x}_0), W_1^{s,u}(\mathbf{x}_0), \ldots, W_{L_{s,u}}^{s,u}(\mathbf{x}_0)$ denote the generation sequence of local stable/unstable manifolds. Then exactly one of the following is true.*

- *There exists k_s, k_u and $t_0 \in \mathbb{R}$, such that $\mathbf{x}(t_0) \in W_{k_s}^s(\mathbf{x}_0) \bigcap W_{k_u}^u(\mathbf{x}_0)$ and $k_s + k_u$ is constant for all pairs (k_s, k_u) which satisfy this property.*
- *There exists $t_0 \in \mathbb{R}$ such that for all $0 \leq k_s \leq L_s$, and $0 \leq k_u \leq L_u$, we have*

$$\mathbf{x}((-\infty, t_0)) \bigcap W_{k_s}^s(\mathbf{x}_0) = \emptyset \quad \text{and} \quad \mathbf{x}((t_0, \infty)) \bigcap W_{k_u}^u(\mathbf{x}_0) = \emptyset.$$

Proposition 2 says that any transverse homoclinic for $\mathbf{x}_0$ satisfying the second condition is a connection which does not intersect in the atlases which we have computed. Restricting to those that do, this proposition says that there is a "first" generation for both the stable and unstable atlases for which the connection will appear.

The significance of this situation is that we need only do pairwise comparisons between stable/unstable charts one generation at a time. Thus, the computational complexity for mining intersections between the two atlases has computational complexity of order $\mathcal{O}(K_s K_u (L_s + L_u))$ where K_s, K_u are the sizes of the largest stable/unstable generations, respectively. This is a dramatic improvement over the naive solution of checking every pair in both atlases which has complexity on order $\mathcal{O}(L_s L_u K_s K_u)$.

Remark 10 We note that often the atlases we compute in practice do not technically satisfy the fundamental domain property. This is due to the fact that sections of manifold boundary which pass near a primary are "cut out" as described in Sect. 4.2.3. Nevertheless, this has no impact on our mining algorithm. Specifically, each generation is still a fundamental domain for the subset of the global manifold which satisfies the speed constraint. Thus, mining for connections via "leapfrogging" through pairwise generations is still assured to find all connections which are present in the computed atlases, and therefore, all connections which satisfy the speed constraint.

Remark 11 The result in Proposition 2 gives rise to a natural mining algorithm. Namely, at each generation, all chart pairs are compared and transverse intersections are identified. It follows that once a transverse intersection is identified, then the next/previous generation must also contain an orbit segment corresponding to the same homoclinic. Hence, in addition to gaining a computational speedup, exploiting the fundamental domain property also ensures that all homoclinics identified are distinct. This follows from the existence of the minimum value for $k_s + k_u$ in Proposition 2.

Furthermore, this observation yields a method of resolving the ambiguous case in which the Newton intersection method finds a zero for G but condition 1 from Sect. 5.1 is not satisfied. Specifically, if $G(\hat{s}, \hat{t}, \hat{\sigma}) \approx 0$ and $\Gamma^u(\hat{s}, \hat{t}) \approx \Gamma^s(\hat{\sigma}, 0) \approx 0$, then we may follow the suspected intersection through earlier-/later-generation charts until the sign condition can be verified or refuted in appropriate predecessor/successor charts. Lastly, we mention that by storing "parent/child" information about the charts in the atlas, we can perform the search just described in post-processing.

5.3 The symmetric case: locating, refining, and classifying, connections

We now describe the homoclinic mining procedure in the case of the triple Copenhagen problem. Assume that we have computed stable/unstable atlases denoted by $\mathcal{A}^s$, $\mathcal{A}^u$, respectively. Each atlas is of the form described in Sect. 4, i.e., each atlas is a union of chart maps having the form, $\Gamma^{s,u} : D \to \mathbb{R}^4$ with $\Gamma^{s,u}(D) \subset W^{s,u}(\mathbf{x_0})$.

We begin with a lemma to motivate the choice to grow each atlas in the symmetric case and then do continuation as opposed to growing the atlas for nonsymmetric cases.

Lemma 1 *Assume f is the symmetric CRFBP vector field, i.e., $m_1 = m_2 = m_3 = \frac{1}{3}$ and define two linear maps, $\varphi^{\pm} : \mathbb{R}^4 \to \mathbb{R}^4$ by $\varphi^{\pm}(x, \dot{x}, y, \dot{y}) = \varphi^{\pm}(\mathbf{x}) = R^{\pm}\mathbf{x}$ where $R^{\pm}$ is the matrix given by*

$$R^{\pm} = \begin{pmatrix} \cos(\pm\theta) & 0 & -\sin(\pm\theta) & 0 \\ 0 & \cos(\pm\theta) & 0 & -\sin(\pm\theta) \\ \sin(\pm\theta) & 0 & \cos(\pm\theta) & 0 \\ 0 & \sin(\pm\theta) & 0 & \cos(\pm\theta) \end{pmatrix} \qquad \theta = \frac{2\pi}{3},$$

then $\varphi^{\pm}$ is a rotational conjugacy for f and $\varphi^{\pm} \circ f(\mathbf{x}) = f \circ \varphi^{\pm}(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^4$. In particular, if γ parameterizes a homoclinic orbit for $\mathcal{L}_0$, then $\varphi^{\pm} \circ \gamma$ are parameterizations for

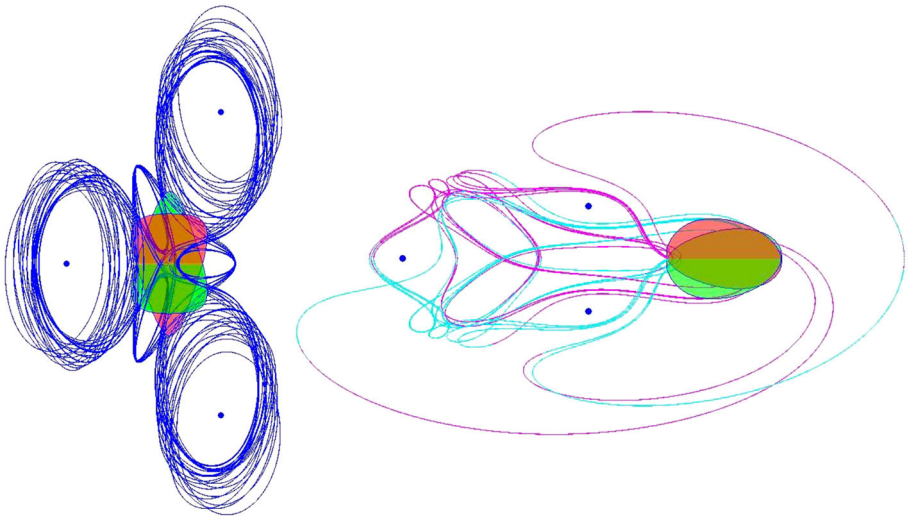

Fig. 9 Asymptotic orbits in configuration space—top-down view of the connections: red and green disks represent the parameterized local stable/unstable manifolds at the libration points. The blue/magenta curves represent the portion of the connecting orbit off the local invariant manifolds—that is, the part found by solving the projected boundary value problem as discussed in Sect. 2.3. Left: 42 shortest homoclinic connecting orbits at $\mathcal{L}_0$ (up to rotational symmetry). Another 84 connections are obtained by $\pm 120^\circ$ rotations. The ± 120 rotations are not plotted as they only thicken the blue-shaded region. Right: 23 shortest homoclinic connecting orbits at $\mathcal{L}_5$. Another 46 connections at $\mathcal{L}_4$ and $\mathcal{L}_6$ are obtained by $\pm 120^\circ$ rotations. The initial guess for the boundary value solver comes from atlases obtained by integrating the local unstable/stable manifolds for $\pm T = 5$ time units. The mining procedure is discussed in Sect. 5.2. All reference to color refers to the online version

two additional, distinct "symmetric" homoclinic orbits for $\mathcal{L}_0$. Moreover, if γ parameterizes a homoclinic orbit for $\mathcal{L}_5$, then $\varphi^+ \circ \gamma$ and $\varphi^- \circ \gamma$ parameterize symmetric homoclinics for $\mathcal{L}_4$ and $\mathcal{L}_6$.

The proof of Lemma 1 is included in "Appendix A." The significance of this symmetry is the fact that global stable/unstable atlases for the triple Copenhagen problem can be separated into three distinct equivalence classes where for $\mathbf{x}, \mathbf{y} \in W^*(\mathbf{x}_0)$, the equivalence relation $\mathbf{x} \sim \mathbf{y}$ is satisfied if and only if $\mathbf{x} \in \{\mathbf{y}, \varphi^+(\mathbf{y}), \varphi^-(\mathbf{y})\}$. Thus, each atlas is obtained by advection of only a single representative for each class. In other words, in the equal masses case, we only need to advect $\frac{1}{3}$ of each initial parameterization boundary to obtain the entire atlas. Specifically, we define

$$D' = \{\mathbf{z} \in D : 0 \leq \mathrm{Arg}(z_1) < \theta,\ z_2 = \mathrm{conj}(z_1)\}$$

and we globalize only $\partial D'$ to obtain a partial atlas, $\mathcal{A}'$. We can then access the full global atlas by applying φ^+, φ^- to each chart in $\mathcal{A}'$ and we set

$$\mathcal{A} = \mathcal{A}' \cup \varphi^+\left(\mathcal{A}'\right) \cup \varphi^-\left(\mathcal{A}'\right).$$

The advantage is a ninefold increase in computational efficiency for the atlas computation and a threefold improvement in efficiency for the atlas mining scheme. Applying the procedure for the triple Copenhagen problems results in the connecting orbits illustrated in Fig. 9. These results are further described and classified in the next section.

5.3.1 Quantitative/qualitative classifications of the homoclinic orbit set at $\mathcal{L}_{0,5}$

Suppose $\mathbf{x}_0 \in \mathbb{R}^4$ is an equilibrium solution and $W^{s,u}_{\text{loc}}(\mathbf{x}_0)$ a local stable/unstable manifold. Let γ be an orbit homoclinic to $\mathbf{x}_0$, and suppose that $T \in \mathbb{R}$ is the elapsed time from when γ passes through the boundary of the local unstable manifold to when γ passes through the boundary of the local stable manifold. Observe that if $W^s_{\text{loc}}(\mathbf{x}_0) \cap W^u_{\text{loc}}(\mathbf{x}_0) = \{\mathbf{x}_0\}$ and if the vector field is inflowing/outflowing on the boundaries of $W^{s,u}_{\text{loc}}(\mathbf{x}_0)$, respectively, then $T > 0$ is well defined.

When the local parameterizations intersect only at $\mathbf{x}_0$, it makes sense to talk about the "shortest" connection time," the "second shortest" connection time, and so on. This natural ordering on connection times provides a useful observable for classifying homoclinic connections relative to fixed local stable/unstable manifolds. Generically, we expect a one-to-one correspondence between connecting orbits and connection times, though this expectation will fail in the presence of symmetries as seen below.

In the CRFBP, when we "mine" the stable/unstable atlases for connecting orbits and order them by connection time we see something interesting. In each of the cases we studied, the shortest homoclinic orbits appear to organize the longer connections. Informally speaking, we find that a small number of short homoclinic orbits serve as a sort of alphabet of "letters," and the longer connections can be roughly identified as "words" in this alphabet.

For example, the first 42 homoclinic connecting orbits (up to symmetry) at $\mathcal{L}_0$ in the triple Copenhagen problem are classified in Table 5. These results are obtained by integrating initial local stable/unstable manifolds for ± 5 time units subject to the speed constraint, $\dot{x}^2 + \dot{y}^2 \leq 4$. Our method finds all of the connections satisfying these constraints. The classification is in terms of the connection time, the order of appearance, and a geometric description in terms of words and letters.

We give the names L_{0A} and L_{0B} to the shortest two connections at $\mathcal{L}_0$. These orbits are illustrated in Fig. 10 and have connection times approximately 1.717 and 2.331, respectively. Rotating either of these by $\pm 120°$ gives another connecting orbit with exactly the same shape and connection time. We refer to these rotations as $L_{0A^\pm}$ and $L_{0B^\pm}$. These six shortest connections—L_{0A}, L_{0B} and their symmetric counterparts—organize the rest of the homoclinic behavior seen at $\mathcal{L}_0$ as we now describe.

We associate the third shortest connection with the word $L_{0A^+} \cdot L_{0A}$ because the orbit moves off the unstable manifold appearing to follow L_{0A^+}, passes near the equilibrium at $\mathcal{L}_0$, and makes another excursion following L_{0A} before returning to the stable manifold. Similarly, we associate with the 5th longest connecting orbit the word L^2_{0B}, as this orbit moves off the unstable manifold and appears to follow L_{0B}, making two loops around the second primary before returning to the stable manifold. Heuristically speaking, L_{0A}, L_{0B} and their symmetric counterparts comprise a system of "homoclinic channels" or simple allowable motions and other homoclinic orbits seem to follow in their wake.

Table 6 records analogous information for the first 23 connections found at $\mathcal{L}_5$ in the triple Copenhagen problem. In this case there are six basic letters L_{5A}, L_{5B}, L_{5D}, L_{5D}, L_{5E}, L_{5F}. Words are formed for these letters just as discussed above. Applying $\pm 120°$ rotations produces connections with the same shapes and connection times at $\mathcal{L}_4$ and $\mathcal{L}_6$, respectively. We stress that this description of the connecting orbits in terms of words and letters, while intuitively appealing, is based on qualitative observations and is subordinate to the rigid quantitative classification of the orbits by connection time.

Remark 12 Several comments about the results reported in Tables 5 and 6 are in order.

Table 5 Classification of the connecting orbits for $\mathcal{L}_0$: advecting one-third of the boundary of the local stable/unstable manifolds for $T = \pm 5$ time units and imposing a speed threshold of 2 reveals the 42 homoclinic connections illustrated in the left frame of Fig. 9

Connection	Connection time	Letter or word
1st	1.717	L_{0A}
2nd	2.331	L_{0B}
3rd	4.198	$L_{0A+} \cdot L_{0A}$
4th	4.520	$L_{0A+} \cdot L_{0B}$
5th	4.715	L_{0B}^2
6th	5.643	L_{0A}^2
7th	6.132	$L_{0A} \cdot L_{0B}$
8th	6.132	$L_{0B-} \cdot L_{0A}$
9th	6.583	$L_{0A-} \cdot L_{0A}$
10th	6.627	$L_{0A+} \cdot L_{0B} \cdot L_{0A}$
11th	6.628	$L_{0B-} \cdot L_{0B}$
12th	6.684	$L_{0A-} \cdot L_{0A+} \cdot L_{0A}$
13th	6.760	$L_{0A+} \cdot L_{0B}^2$
14th	6.846	L_{0B}^3
15th	7.009	$L_{0B+} \cdot L_{0A+} \cdot L_{0A}$
16th	7.336	$L_{0B+} \cdot L_{0A+} \cdot L_{0B}$
17th	7.038	$L_{0B+} \cdot L_{0A}$
18th	7.038	$L_{0A-} \cdot L_{0B}$
19th	7.490	$L_{0B+} \cdot L_{0B}$
20th	8.119	$L_{0B} \cdot L_{0A}$
21st	8.125	$L_{0A+}^2 \cdot L_{0A}$
22nd	8.125	$L_{0A+} \cdot L_{0A}^2$
23rd	8.296	$L_{0A} \cdot L_{0B} \cdot L_{0A}$
24th	8.448	$L_{0B} \cdot L_{0A}^2$
25th	8.453	$L_{0A} \cdot L_{0B}^2$
26th	8.453	$L_{0B-}^2 \cdot L_{0A}$
27th	8.499	$L_{0B}^3 \cdot L_{0A}$
28th	8.499	$L_{0A+} \cdot L_{0B}^3$
29th	8.583	L_{0B}^2
30th	8.614	$L_{0A+} \cdot L_{0A} \cdot L_{0B}$
31st	8.732	$L_{0A+} \cdot L_{0B}^3$
32nd	8.794	$L_{0A} \cdot L_{0B-} \cdot L_{0B}$
33rd	8.937	$L_{0B} \cdot L_{0A} \cdot L_{0B}$
34th	8.953	$L_{0B-} \cdot L_{0B}^2$
35th	8.953	$L_{0B-}^2 \cdot L_{0B}$
36th	9.065	$L_{0A} \cdot L_{0A-} \cdot L_{0A}$
37th	9.340	$L_{0A-} \cdot L_{0B}^2$
38th	9.88	$L_{0B-} \cdot L_{0A-} \cdot L_{0A}$
39th	9.444	$A^- \cdot B^+ \cdot A^+ \cdot B$

Table 5 continued

Connection	Connection time	Letter or word
40th	9.495	$A \cdot A^- \cdot A^+ \cdot B$
41st	9.579	$B^+ \cdot A^+ \cdot B^2$
42nd	9.579	$(B^+)^2 \cdot A^+ \cdot B$

The rows of the table give data for the homoclinics, ordered by connection time. For each orbit described in the table, there are two additional orbits with exactly the same connection time, obtained by $\pm 120°$ rotations. Taken with their symmetric counterparts, the orbits given here are all of the connecting orbits satisfying the connection time and speed constraints. The three columns in the table report a connection's order of appearance, its connection time, and our qualitative description of the connection as a word built from the two shortest homoclinics—the letters shown in Fig. 10. The connections associated with longer words are separated and illustrated in Figs. 11, 12, 14, 15 and 16

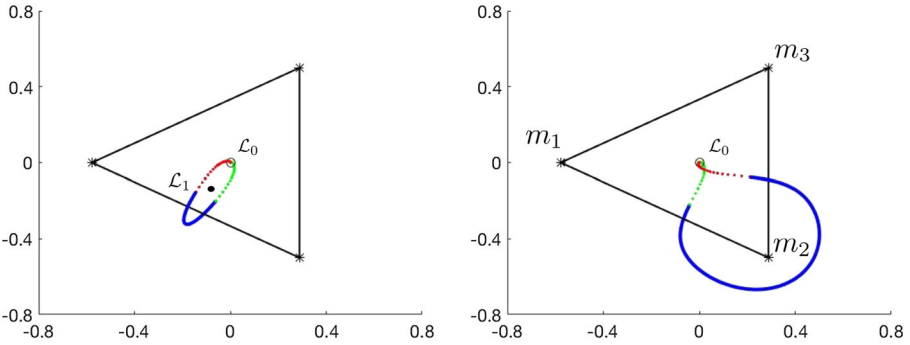

Fig. 10 Shortest connections at $\mathcal{L}_0$: we distinguish the six shortest connecting orbits at $\mathcal{L}_0$ in the triple Copenhagen problem. Two are shown above. The four others are rotations of these by ± 120. See also the left frames in Fig. 20. The blue portion of the orbit is the segment found by solving the boundary value problem. The dotted red/green lines are the asymptotic portions on the parameterized local unstable/stable manifolds, respectively. The asymptotic portions are recovered using the conjugacy provided by the parameterization method. We refer to the orbit on the left as L_{0A} and the orbit on the right as L_{0B}. The rotations by $\pm 120°$ we refer to as $L_{0A^\pm}$ and $L_{0B^\pm}$. Observe that L_{0A} has winding number 1 with respect to the libration point $\mathcal{L}_1$, while $L_{0A^\pm}$ wind once around $\mathcal{L}_2$ and $\mathcal{L}_3$, respectively. L_{0B} on the other hand winds once around m_2 while $L_{0B^\pm}$ wind once around m_3 and m_1, respectively. As the next five figures illustrate, the six basic connecting orbits organize all the connections we found at $\mathcal{L}_0$. All references to color refer to the online version

– **Additional symmetries** Some of the orbits, for example the 27th and 28th shortest orbits at $\mathcal{L}_0$ and the 21st and 22nd shortest orbits at $\mathcal{L}_5$, have reported the same connection times to three decimal places. In fact the connection times agree to within numerical errors. This is because the equal mass problem has reversible symmetries that we are not exploiting in our computations. Rather these serve as a check on the numerics.
– **Connection time versus ordering** While the connection times reported in these tables depend on the choice of local stable/unstable manifold, it should be remarked that, as long as the parameterization method is used to represent the local manifolds as discussed in Sect. 3, the ordering of the connections does not change. It is easy to check that the boundary of the parameterized manifolds are inflowing/outflowing and that the manifolds intersect only at the equilibrium solution. Moreover, since the eigenvalues are complex conjugate, the local parameterizations are unique up to the choice of a single eigenvector

Table 6 Classification of the connecting orbits for $\mathcal{L}_5$: the 23 homoclinic connections which appear on the right side of Fig. 9 satisfying the same connection time and speed constraints as in the $\mathcal{L}_0$ case

Connection	Connection time	Letter or word
1st and 2nd	4.802	L_{5A} and L_{5B}
3rd and 4th	4.943	L_{5C} and L_{5D}
5th and 6th	5.261	L_{5E} and L_{5F}
7th	6.028	$L_{5D} \cdot L_{5C}$
8th	8.204	$L_{5A} \cdot L_{5B}$
9th	8.331	$L_{5A} \cdot L_{5D}$
10th	8.456	$L_{5C} \cdot L_{5D}$
11th	8.917	$L_{5E} \cdot L_{5F}$
12th	8.934	$L_{5A} \cdot L_{5D}$
13th	9.156	$L_{5C} \cdot L_{5D}$
14th	9.324	$L_{5E} \cdot L_{4E} \cdot L_{5C}$
15th	9.363	$L_{5A} \cdot L_{5D} \cdot L_{5C}$
16th	9.387	$L_{5A} \cdot L_{6A} \cdot L_{5B}$
17th	9.429	L_{5C}^2
18th	9.429	L_{5D}^2
19th	9.487	$L_{5D} \cdot L_{5C} \cdot L_{5D}$
20th	9.487	$L_{5C} \cdot L_{5D} \cdot L_{5C}$
21st	9.554	$L_{5A} \cdot L_{6A} \cdot L_{5D}$
22nd	9.554	$L_{5C} \cdot L_{6A} \cdot L_{5B}$
23rd	9.629	$L_{5C} \cdot L_{6A} \cdot L_{5D}$

In this case, the 120° symmetry does not produce additional connections for $\mathcal{L}_5$ but rather, rotation of each connection produces symmetric homoclinics for both $\mathcal{L}_4$ and $\mathcal{L}_6$. The columns are similar to those in Table 5 and the longer words associated with $\mathcal{L}_5$ homoclinics illustrated in Figs. 19 and 20

scale factor. By choosing the unit disk as the domain of the parameterization, the scaling the only free parameter in the problem. Decreasing the scaling by a factor of $\tau > 0$ is equivalent to flowing the boundary by the same time τ. So: rescaling the eigenvectors changes all the times of flight by exactly the same amount, hence does not reorder them.

– **Qualitative classification** The decomposition of the connecting orbits into words is performed "by eye" in the present work. That is, we simply inspect the connections and describe what we see. We now sketch a method which could be used to formalize our qualitative description and note that the idea is computationally feasible. Recall that if γ is a simple closed rectifiable curve in the plane (with counter clockwise orientation), and $z_0 = x_0 + iy_0$ is a point not on γ, then the number of times that γ winds around z_0 is counted by the integral

$$\mathrm{Ind}_\gamma(z_0) = \frac{1}{2\pi i} \int_\gamma \frac{1}{z - z_0} \, \mathrm{d}z.$$

So, observe for example that the curve L_{0A} winds once around the libration point $\mathcal{L}_1$, while L_{0A^+} and L_{0A^-} each wind once, respectively, around $\mathcal{L}_2$ and $\mathcal{L}_3$. Similarly, the curve L_{0B} winds once around the second primary, while L_{0B^+} and L_{0B^-} wind once each around the third and first primaries. Then for a homoclinic orbit γ at $\mathcal{L}_0$ define the integer vector $(\alpha_1, \alpha_2, \alpha_3, \beta_1, \beta_2, \beta_3) \in \mathbb{Z}^6$, where $\alpha_j = \mathrm{Ind}_\gamma(\mathcal{L}_j)$ for and $\beta_j = \mathrm{Ind}_\gamma(P_j)$ both

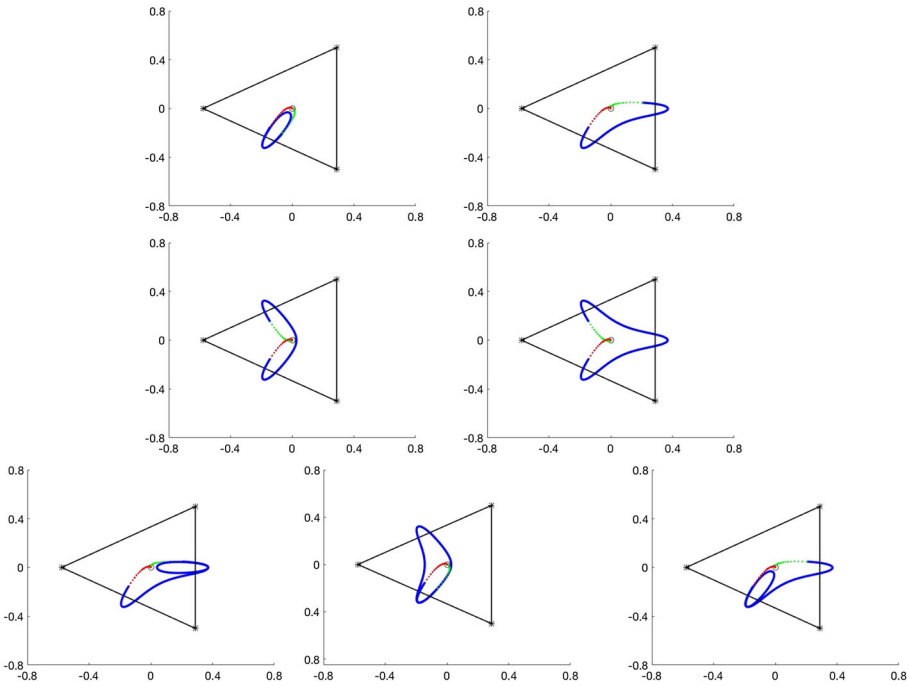

Fig. 11 Homoclinic orbits at $\mathcal{L}_0$ associated with L_{0A}: the figure illustrates seven homoclinic connecting orbits all built from L_{0A} and its rotations $L_{0A\pm}$ in the following sense: each of these orbits has nontrivial winding number with respect to one or more of the inner libration points, but none of these orbits wind around any of the primaries. That is, these orbits are "built from pieces of L_{0A} and its rotations, but have no contribution from L_{0B} and its rotations. (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) The orbits give the following symbol sequences: L_{0A}^2—6th connection (top left), $L_{0A^+} \cdot L_{0A}$—3rd (top right), $L_{0A^-} \cdot L_{0A}$—9th (center left), $L_{0A^-} \cdot L_{0A^+} \cdot L_{0A}$–12th connection (center right), $L_{0A^+}^2 \cdot L_{0A}$—21st (bottom left), $L_{0A} \cdot L_{0A^-} \cdot L_{0A}$—36th (bottom center), $L_{0A^+} \cdot L_{0A}^2$—22nd (bottom right). In each case rotation of each by $\pm 120^\circ$ gives a distinct connection (not shown). All references to color refer to the online version

for $j = 1, 2, 3$. (Here P_j are the coordinates of the jth primary). Then α_1 counts the contribution of L_{0A} to γ, while β_1 counts the contribution of L_{0B} and so on. This description amounts to an Abelianization of the previous notion of words/letters, as the winding vector looses track of the order of the letters in the word. (It is often the case that mechanical calculation of topological data is facilitated by passing to an Abelianization). This notion is extended to the homoclinic orbits at $\mathcal{L}_5$ in a similar way, see Fig. 18.

– **Blue skies** The main theorem of Henrard (1973), already mentioned in the introduction, gives that there is a family of periodic orbits accumulating to each of the homoclinic orbits found by our procedure. In some cases we can venture a guess as which families of periodic orbits accumulate to which homoclinic. For example, Fig. 17 illustrates the orbit L_{0A} and $L_{0A\pm}$ along with the planar Lyapunov families attached to the inner libration points $\mathcal{L}_{1,2,3}$. The results suggest that the planar families may accumulate at to these homoclinics. Similarly, comparing the orbits L_{5E} and L_{5F} in the bottom left and right frames of Fig. 18 with the planar Lyapunov families at $\mathcal{L}_{7,9}$ illustrated in the left frame of Fig. 3 suggests that these may be the families of periodic orbits attached to these homoclinics. The orbits L_{0B}, and $L_{0B^\pm}$, as well as the orbits L_{5A} and L_{5B} must be

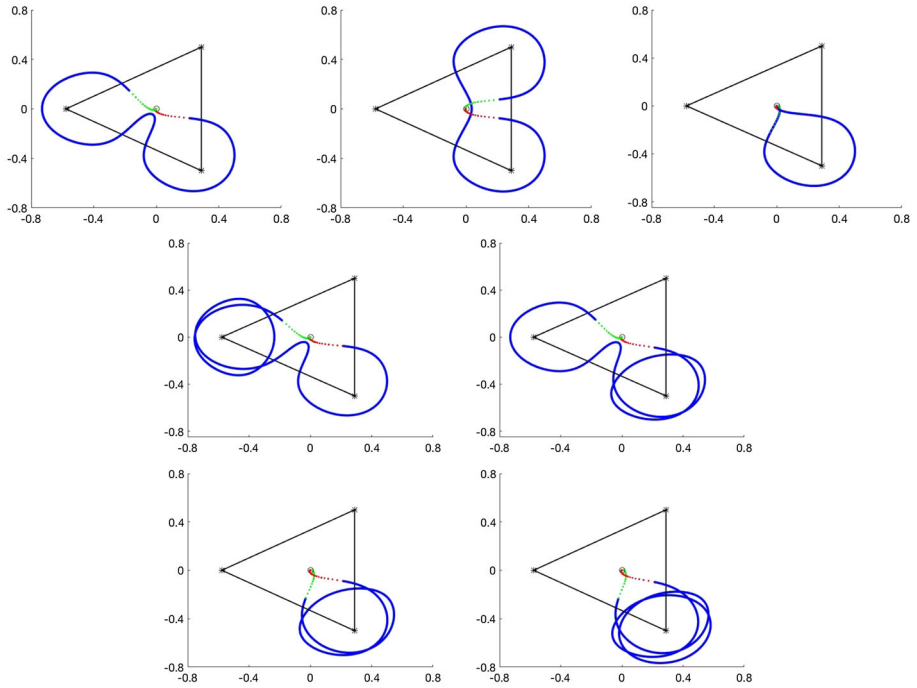

Fig. 12 Homoclinic orbits at $\mathcal{L}_0$ associated with L_{0B}: the figure illustrates seven connections built from L_{0B} and its rotations. So, each of these orbits has nonzero winding about one or more primary—that is L_{0B} behavior—and none of them have any winding about the inner libration points—that is no L_{0A} behavior. (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) The orbits give the following symbol sequences: $L_{0B^-} \cdot L_{0B}$—11th connection (top left), $L_{0B^+} \cdot L_{0B}$—19th (top center), L_{0B}^2—29th (top right), $L_{0B^-}^2 \cdot L_{0B}$—35th (center left), $L_{0B^-} \cdot L_{0B}^2$—34th connection (center right), L_{0B}^2—5th connection (bottom left), L_{0B}^3—14th (bottom right). In each case rotation of each by $\pm 120^\circ$ gives a distinct connection (not shown). Observe that the symbol sequences need not be unique. For example, the top right and bottom left orbits have the same winding, but different times of flight. (See also Fig. 13 and Table 5). All references to color refer to the online version

the limits of families of periodic orbits winding around the primary bodies. Making a systematic study of the periodic families associated with the homoclinic orbits discussed here would make a nice topic for a future study.

Remark 13 (Some related work on asymptotic orbits) There are interesting similarities between some of the orbits discussed above, and some asymptotic orbits already discovered in Baltagiannis and Papadakis (2011b). The interested reader might, for example, compare the homoclinic orbit on the bottom right frame of Fig. 20 with the heteroclinic termination orbit illustrated in Figure 5 of Baltagiannis and Papadakis (2011b). (To make such a comparison one has to "flip" Figure 5 of Baltagiannis and Papadakis (2011b) 180° about the y-axis as the two papers use different normalizations of the four-body problem. Also, their L_3 is our $\mathcal{L}_5$). In that study the heteroclinic is discovered by numerical continuation of the author's f_{10} family of periodic orbits: a family of orbits with winding number one about all three of the primary masses. We note that our homoclinic of Fig. 20 is similar, but that the $\pm 120^\circ$ rotational symmetry broken. We conjecture that there are three families of periodic orbits

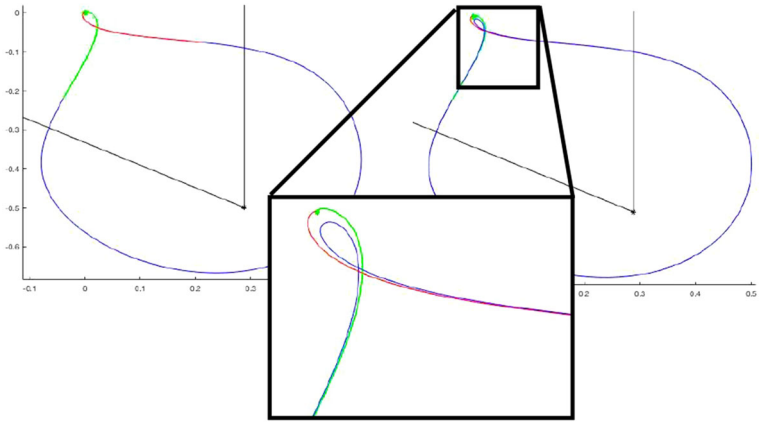

Fig. 13 A close return to $\mathcal{L}_0$: while the top right frame in Fig. 12—illustrating the L_{0B}^2 orbit—looks like a repeat of L_{0B}, further examination reveals that the orbits are distinct. The L_{0B} and L_{0B}^2 orbits are here shown side-by-side. The longer L_{0B}^2 orbit on the right follows L_{0B}, makes are "flyby" of the libration point, then a second excursion, finally returning to the stable manifold. A closeup of a neighborhood of $\mathcal{L}_0$—shown in the inlay—illustrates the flyby. These orbits illustrate the fact that two very similar looking orbits can have very different connection times. Indeed, very simple looking orbits can have long connection times if they spend a long time in the neighborhood of a libration point where dynamics move slow

bifurcating from the f_{10} family after a symmetry breaking, and that these families terminate on the homoclinic of Fig. 20 (bottom right) and its rotation by $\pm 120°$ counterparts.

Similarly, the heteroclinic orbit illustrated in Figure 4 of Baltagiannis and Papadakis (2011b)—which is the termination of the author's f_5 family—is related the pair of homoclinic orbits $L_{E,F}$ illustrated in Fig. 18. To see this, imagine an orbit obtained by combining our L_{5F} with the orbit L_{4E}, that is our L_{5E} rotated by $-120°$ so that it is based at $\mathcal{L}_4$. The resulting union of curves has the same shape as the heteroclinic illustrated in Figure 4 of Baltagiannis and Papadakis (2011b). This suggests that the families of periodic orbits which terminate at our $L_{E,F}$ could emerge from the planar Lyapunov families after symmetry breaking.

In general we note that the homoclinic orbits tend to have less symmetry than the heteroclinic, so that studying the periodic orbits terminating at the homoclinics is a good way to obtain asymmetric periodic orbits—even in the symmetric versions of the problem. We also note that changing the mass parameters will tend to destroy heteroclinic connections, as the libration points will move into distinct energy levels. Homoclinic orbits on the other hand persist under generic Hamiltonian perturbations of the vector field. In particular, they persist after a small change in the mass ratios, facilitating numerical continuation as discussed below.

5.4 Numerical continuation of ensembles of connections

The fact that the homoclinic connecting orbits are formulated as solutions of boundary value problems makes parameter continuation natural. We give only an outline of our continuation algorithm, as numerical continuation of homoclinic orbits for Hamiltonian systems is described in great detail in the literature. References are discussed in the introduction.

Begin with an ensemble of connecting orbits for a libration point $\mathcal{L}$ at the mass parameters m_1, m_2, m_3 (initially we have $m_1 = m_2 = m_3 = 1/3$).

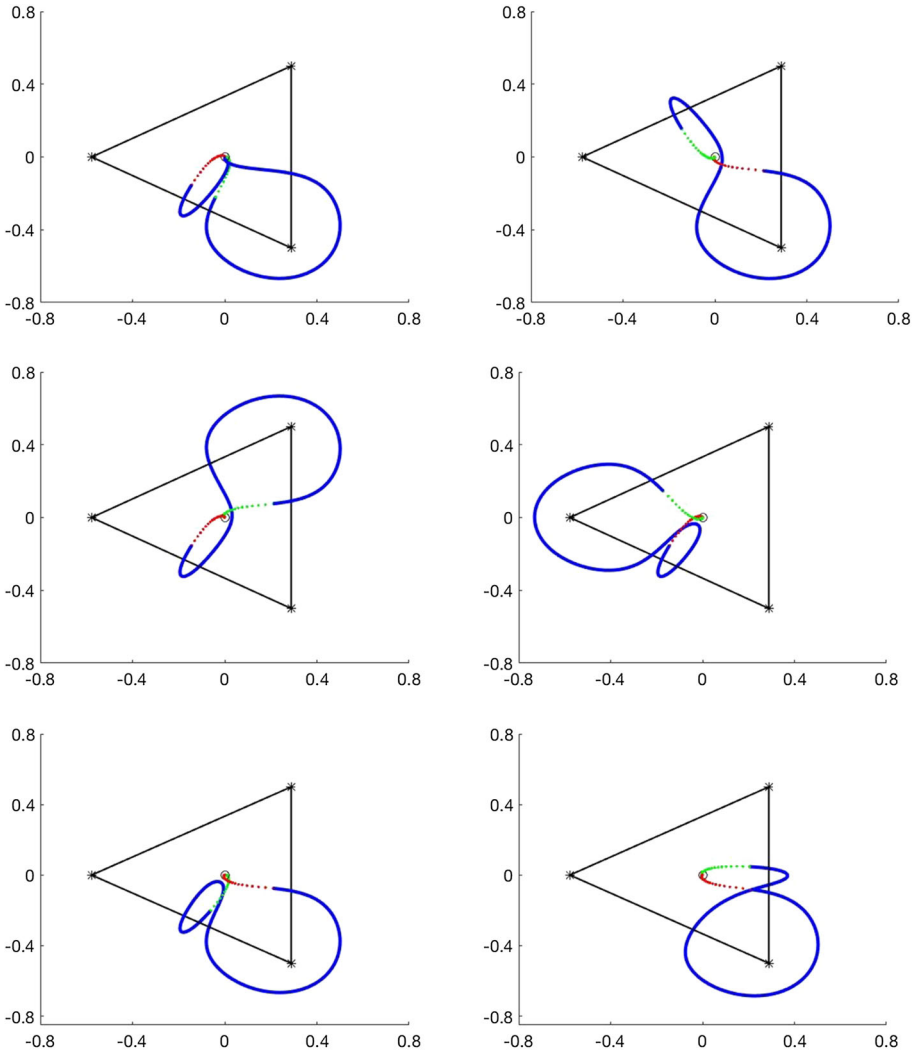

Fig. 14 Homoclinic orbits at $\mathcal{L}_0$ with AB words: six connections built from exactly one occurrence of L_{0A} (or a symmetric counterpart) and one occurrence of L_{0B} (or one of its symmetric counter parts). So, each of these orbits has winding number one about exactly one of the primaries—giving one instance of L_{0B} behavior—and winding number one about exactly one of the inner libration points—giving one instance of L_{0A} behavior. We refer to these as two-letter words. (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) We see the following symbol sequences: $L_{0B} \cdot L_{0A}$—20th connection (top left), $L_{0A^-} \cdot L_{0B}$—18th (top right), $L_{0B^+} \cdot L_{0A}$—17th (center left), $L_{0B^-} \cdot L_{0A}$—8th (center right), $L_{0A} \cdot L_{0B}$—7th (bottom left), $L_{0A^+} \cdot L_{0B}$—4th (bottom right). In each case rotation of each by $\pm 120^\circ$ gives a distinct connection (not shown). All references to color refer to the online version

– We choose a new parameter set $\bar{m}_1 = m_1 + \delta_1$, $\bar{m}_3 = m_3 + \delta_3$. Then we compute $\bar{m}_2 = 1 - \bar{m}_1 - \bar{m}_3$ and apply a first-order predictor corrector to find the libration point at the new parameter values. We numerically compute the eigenvalues and eigenvectors of the new libration point, and if it remains a saddle-focus (i.e., if there has been no bifurcation) we proceed.

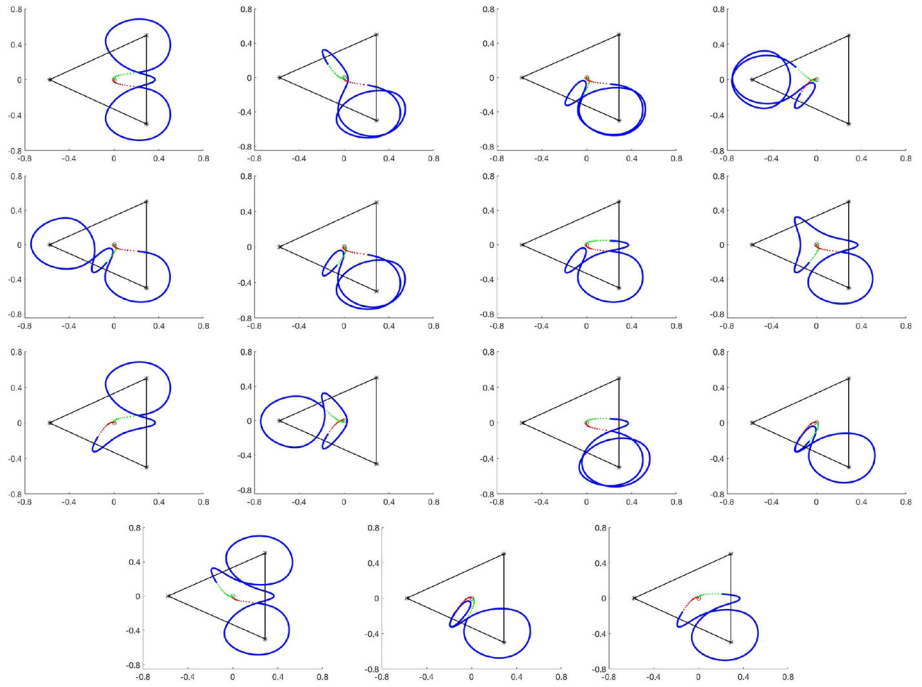

Fig. 15 Fifteen connections at $\mathcal{L}_0$ whose words have three or more symbols: orbits with multiple winding about one or more primaries as well as one or more inner libration point, and hence multiple L_{0A} and L_{0B} behaviors. We refer to these as three or four-letter words. (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) We see the following symbol sequences: $L_{0B^+} \cdot L_{0A^+} \cdot L_{0B}$—16th connection (first row, first column), $L_{0A^-} \cdot L_{0B}^2$—37th (first row, second column), $L_{0B} \cdot L_{0A} \cdot L_{0B}$—33rd (first row, third column), $L_{0B^-}^2 \cdot L_{0A}$—26th (first row, fourth column), $L_{0A} \cdot L_{0B^-} \cdot L_{0B}$—32nd (second row, first column), $L_{0A} \cdot L_{0B}^2$—25th (second row, second column), $L_{0A^+} \cdot L_{0A} \cdot L_{0B}$—30th (second row, third column), $L_{0A} \cdot L_{0A^-} \cdot L_{0A^+} \cdot L_{0B}$—40th (second row, fourth column), $L_{0B^+} \cdot L_{0A^+} \cdot L_{0A}$—15th (third row, first column), $L_{0B^-} \cdot L_{0A^-} \cdot L_{0A}$—38th (third row, second column), $L_{0A^+} \cdot L_{0B}^2$—13th (third row, third column), $L_{0A} \cdot L_{0B} \cdot L_{0A} \cdot$—23rd (third row, fourth column), $L_{0A^-} \cdot L_{0B}^+ \cdot L_{0A}^+ \cdot L_{0B}$—39th (fourth row, first column), $L_{0B} \cdot L_{0A}^2$—24th (fourth row, second column), $L_{0A}^+ \cdot L_{0B} \cdot L_{0A}$—10th (fourth row, third column), In each case rotation of each by $\pm 120°$ gives a distinct connection (not shown)

– We recompute the local invariant manifolds at the new parameter set. A good strategy is to compute the coefficients to order N_0 by recursively solving the homological equations. Initially we take the eigenvector scaling from the previous step and rescale if needed. For the higher-order coefficients we use the coefficients from the previous step. This gives as an initial guess for the Newton or pseudo-Newton method which usually converges very fast.
– The new local parameterizations provide the boundary conditions for the multiple shooting scheme for the homoclinic orbits. We take the connecting orbits from the previous step as the initial guesses for the Newton method at the current mass parameters. If necessary we can apply a first-order predictor corrector, but this is often unnecessary, due to the fact that the boundary value problem formulated with the high-order parameterizations of the local manifold is very well conditioned. Note that in a given continuation step, the same local parameterizations serve as the boundary conditions for the entire ensemble

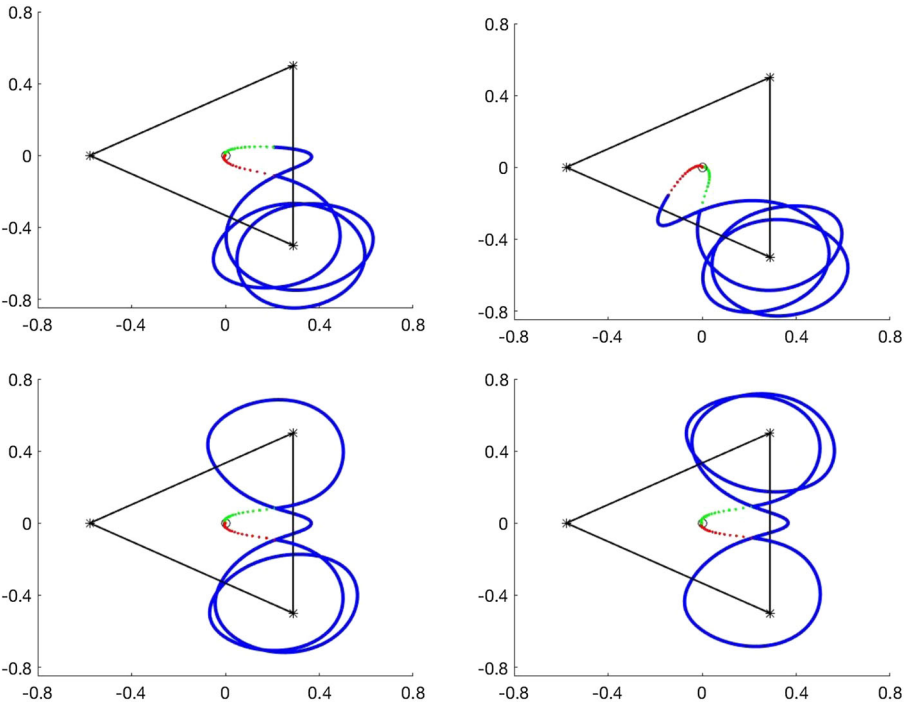

Fig. 16 Longest words the $\mathcal{L}_0$ atlas data: the four most complicated words in our $\mathcal{L}_0$ search—orbits whose words contain four letters. Each has triple winding about the primaries—either winding three about one primary mass or winding two about one followed by winding one about another. Each also has winding one about an inner libration point. It is interesting to note that these are not necessarily the longest orbits, in the sense of connection time. While the bottom left and right are the two longest orbits, the top left and right are only 27th and 28th, respectively. (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) We see the following symbol sequences: $L_{0A}^{+} \cdot L_{0B}^{3}$—28th connection (top left), $L_{0B}^{3} \cdot L_{0A}$—27th (top right), $L_{0B}^{+} \cdot L_{0A^{+}} \cdot L_{0B}^{2}$—41st (bottom left), $(L_{0B}^{+})^{2} \cdot L_{0A} \cdot L_{0B}$ 42nd (bottom right). In each case rotation of each by $\pm 120^{\circ}$ gives a distinct connection (not shown)

of connecting orbits. This justifies the cost of computing high-order representations of the manifolds.

– Once we have applied Newton to all the connections in the ensemble, we are ready to take a new step. If Newton fails to converge for any of the connecting orbits, we have to decide if we throw the orbit away, or if we recompute with smaller $\delta_1, \delta_2, \delta_3$.

We also remark that the atlas is not recomputed at the new mass parameter set. That is, we continue only the connecting orbits—the intersections of the stable unstable manifolds—not the manifolds themselves. Continuation of ensembles of connections is much cheaper than recomputing the atlas each time we change parameters.

Results of several numerical continuations are illustrated in Figs. 21 and 22. As we change the masses we break the rotational symmetry of the Triple Copenhagen problem and the symmetric counterparts resolve into distinct connection, no longer obtainable by rotations of a single representative. During the numerical continuation we sometimes encounter bifurcations of the connecting orbits themselves, which involve no bifurcation of the underlying equilibrium. Figure 23 illustrates a common scenario where a family of homoclinic orbits undergoes a doubling bifurcation. These bifurcations seem very common and we have not

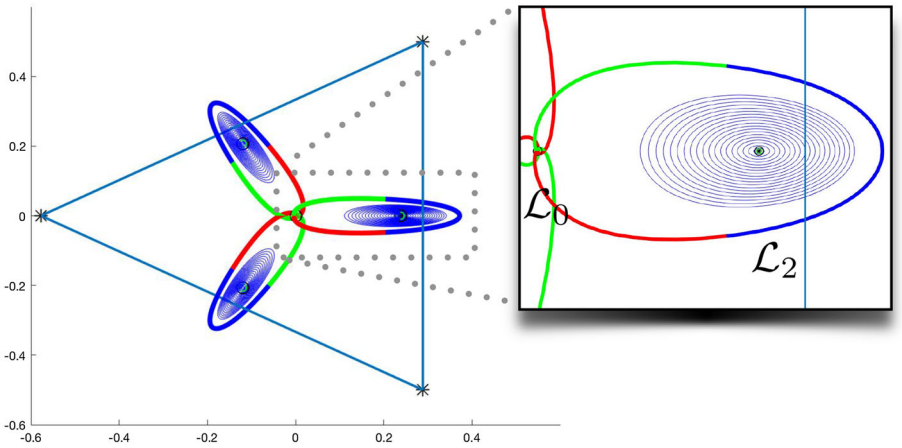

Fig. 17 Blue sky catastrophes at $\mathcal{L}_0$: the thick blue lines correspond to the portion of the homoclinic orbit computed using the BVP approach. The green and red lines correspond to the asymptotic portion of the orbit on the parameterized local unstable and stable manifolds, respectively. Left: the three shortest homoclinic orbits at $\mathcal{L}_0$. Note that the orbits are rotations by 120° of one another. The figure also includes the planar Lyapunov families of orbits about the inner libration points $\mathcal{L}_{1,2,3}$. These periodic orbits were computed using a center manifold reduction, and we have not applied numerical continuation to the boundary. Nevertheless the images suggest that the planar Lyapunov families may accumulate that the L_{0A} homoclinic orbits. Right: closeup on a neighborhood of $\mathcal{L}_0$ and $\mathcal{L}_2$. The homoclinic orbit L_{0A^+} clearly has winding number one about $\mathcal{L}_2$. All reference to color refers to the online version

made a systematic effort to track them. This would make an interesting topic for a future study.

6 Conclusions

In this work we implemented a numerical method for computing an atlas for the stable/unstable manifold attached to a libration point in the CRFBP. We consider saddle-focus equilibria, as in this case topological theorems give rich dynamical structure near a transverse homoclinic. We then developed algorithms for searching or "mining" the atlas for approximate connections. After an approximate connection is found we sharpen it using a Newton scheme for an appropriate boundary value problem. The procedure is completely automated and locates all the homoclinic orbits up to a specified integration time. (To make the calculations less stiff we ignore orbits passing too close to the primaries). The resulting collection of homoclinic orbits is ordered by connection time. A different choice of local parameterization may yield different connection times, but the ordering of these connections is universal. This last comment requires that the local stable/unstable manifolds are computed using the parameterization method.

Our method locates dozens of distinct homoclinic connections and we consider their qualitative properties in addition to their connection times. The geometry of the homoclinic orbit set is organized by the several shortest connections, in the sense that they form a system of channels that other connections appear to follow. We decompose the motions of the longer connections into words built from these simple letters and discuss briefly how this

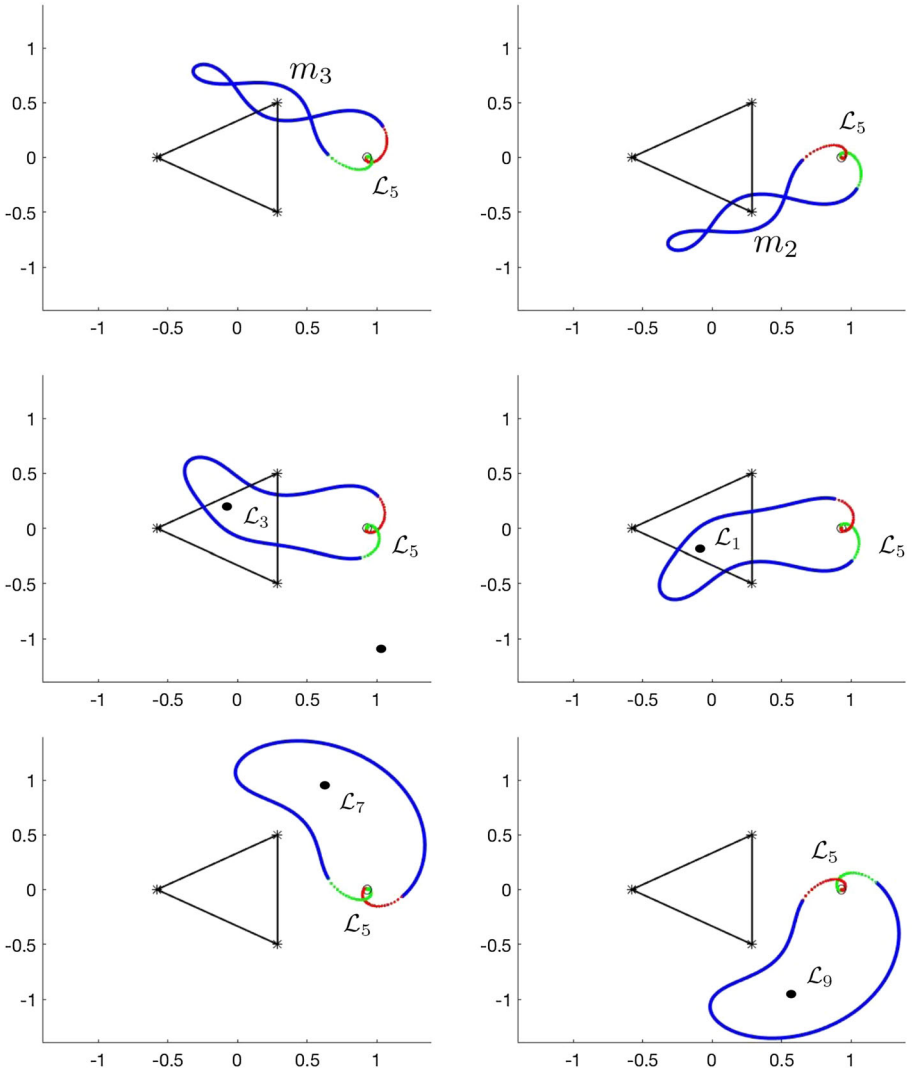

Fig. 18 Fundamental connections at $\mathcal{L}_5$: the six shortest connecting orbits at $\mathcal{L}_5$ for the triple Copenhagen problem. Shortest connections at $L_{4,6}$ are obtained by rotation by ± 120. These six orbits and their rotations organize all the connections we find at $\mathcal{L}_5$, as illustrated in the next two figures. (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) We refer to these orbits as L_{5A}, L_{5B} (top left and right), L_{5C}, L_{5D} (middle left and right), and L_{5E}, L_{5F} (bottom left and right). Moreover, by rotating the pictures by $\pm 120^\circ$ we obtain orbits which we refer to as $L_{4ABCDEF}$ and $L_{6ABCDEF}$. Observe that L_{5A} and L_{5B} have winding number 1 with respect to the second and third primaries. Moreover, they are the only orbits of the six basic words which wind around any primary. Similarly, L_{5C} and L_{5D} wind once around the inner libration points $\mathcal{L}_3$ and $\mathcal{L}_1$ once, respectively. They are the only basic words at $\mathcal{L}_5$ with this property. Finally, L_{5E} and L_{5F} are distinguished by the fact that they have winding number one with respect to the outer libration points $\mathcal{L}_7$ and $\mathcal{L}_9$, respectively. All references to color refer to the online version

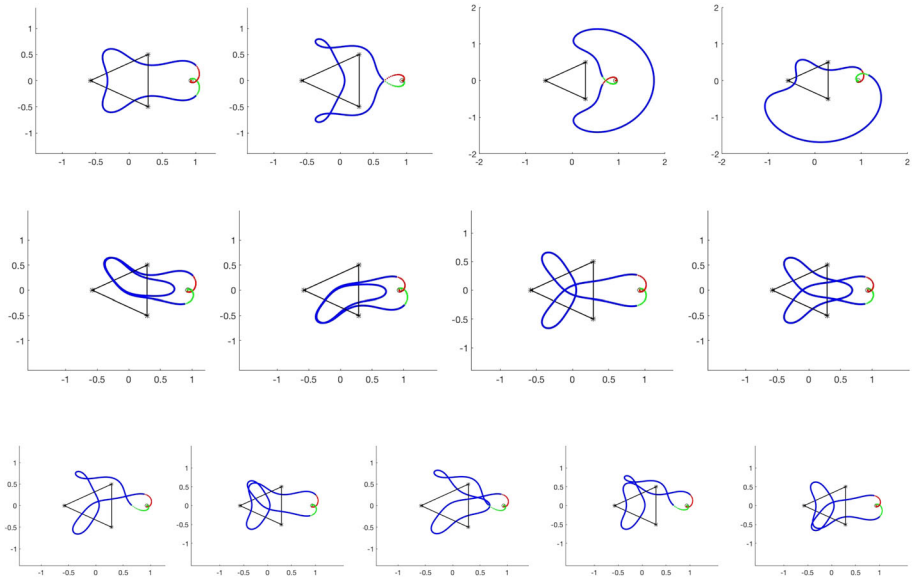

Fig. 19 Homoclinic orbits with two- or three-letter words at $\mathcal{L}_5$: each of the orbits shadows two or three of the six basic shapes shown in Fig. 18 (or one of their rotations). Rotation of each by $\pm 120°$ give a connection at $L_{4,6}$. The symbol sequences for these orbits are $L_{5D} \cdot L_{5C}$—7th connection (first row, first column), $L_{5A} \cdot L_{5B}$—8th (first row, second column), $L_{5E} \cdot L_{5F}$—11th (first row, third column), $L_{5F} \cdot L_{4E} \cdot L_{5C}$—14th (first row, fourth column), L_{5C}^2—17th (second row, first column), L_{5D}^2—18th (second row, second column), $L_{5C} \cdot L_{5D}$—10th (second row, third column), $L_{5C} \cdot L_{5D}$—13th (second row, fourth column), $L_{5A} \cdot L_{5D}$—9th (third row, first column), $L_{5C} \cdot L_{5D} \cdot L_{5C}$—20th (third row, second column), $L_{5A} \cdot L_{5D}$—12th (third row, third column), $L_{5A} \cdot L_{5D} \cdot L_{5C}$—15th (third row, fourth column), $L_{5D} \cdot L_{5C} \cdot L_{5D}$—19th (third row, fifth column). (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) Again we see that a symbol sequence can appear more than once. For example, the frames in the second row third and fourth column both have $L_{5C} \cdot L_{5D}$. However, the orbits are distinguished by their connection time (see Table 6). All references to color refer to the online version

decomposition could be calculated in an automated way using the formulation of the winding number as a complex line integral.

Finally, we continue some of the orbits found in the equal mass case to other nonsymmetric mass values using a predictor corrector scheme for the boundary value problem. Rather than recomputing the entire atlas the continuation scheme only recomputes the initial parameterization at the new masses, and this can be done via a Newton scheme.

We remark that it would be easy to intersect that atlas data computed here with any desired surface of section. We have not used sections in the present work because (a) we wanted to find all the connections up to a certain integration time and a given section may find some orbits and not others and (b) projecting to a section may not preserve the ordering of the homoclinics by connection time. Moreover, while the present project focuses on the CRFBP—a two-degree-of-freedom Hamiltonian system—the scheme described here could be extended to higher-dimensional systems were sections are four or more dimensional and hence less useful for visualization purposes. In such a situation, for example the spatial CRFBP, it is desirable to have an automated procedure.

Other interesting topics of future research would be to combine our methods with more sophisticated continuation and branch following algorithms. It would also be nice to return to the ideas of Strömgren and examine the "tubes" of periodic orbits attached to each of

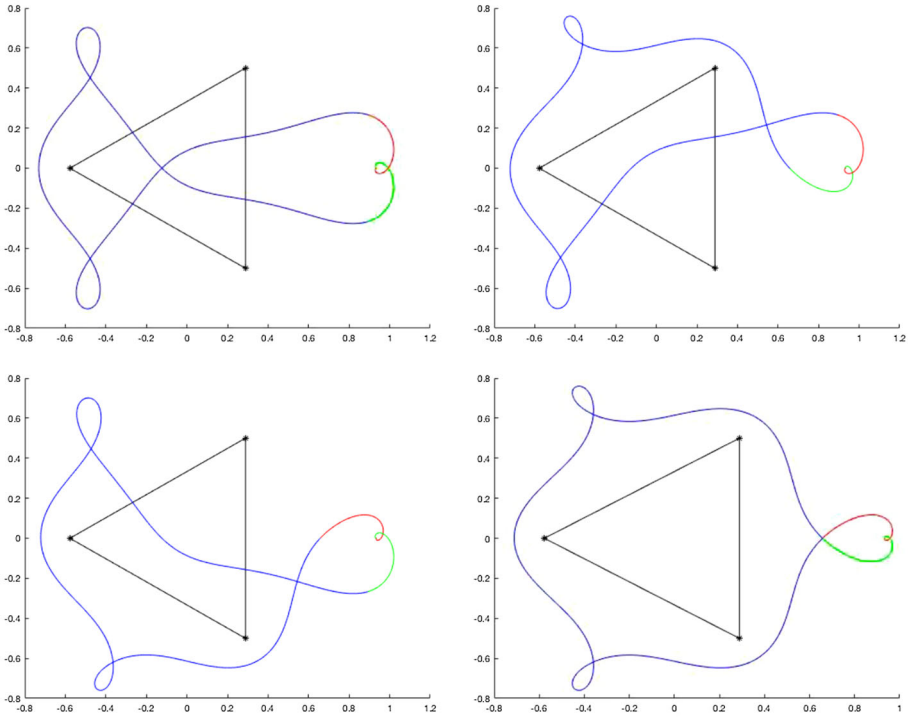

Fig. 20 Orbits at $\mathcal{L}_5$ which wind around the first primary: The orbit on the top left has nontrivial winding about the largest primary and no winding about the smaller primaries. The orbits in the top right and bottom left frames wind around the largest primary and one or the other smaller primaries. Finally the orbit in the bottom right frame winds around all three primaries. The symbol sequences for these orbits are $L_{5C} \cdot L_{6A} \cdot L_{5D}$—23rd and longest connection (top left), $L_{5A} \cdot L_{6A} \cdot L_{5D}$—21st (top right), $L_{5C} \cdot L_{6A} \cdot L_{5B}$—22nd (bottom right) $L_{5A} \cdot L_{6A} \cdot L_{5B}$—16th connection (bottom right). (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) Note that while orbit $L_{5A} \cdot L_{6A} \cdot L_{5B}$ is one of the most geometrically complicated in our search, it has only the 16th longest connection time. All references to color refer to the online version.

our homoclinic connections. These periodic families would themselves undergo bifurcations which one could try to follow numerically.

Another improvement to our method would be to remove the speed constraints on our manifold computations. This could be done by regularizing binary collisions. The idea would be that whenever a chart gets too close to a primary, then instead of subdividing we would change to the regularized coordinates where computations are less stiff. This idea of using such regularizations to improve numerics goes back at least to the work of Thiele. This would also provide a natural way for computing collision orbits between $\mathcal{L}_{0,5}$ and each of the primaries. A topic we have avoided via our imposed speed constraints. A modern implementation combined with our approach to computing atlases would be valuable, and is the subject of ongoing work.

If such advancements let us compute larger and more complete atlases, a very interesting question is to see if other "fundamental" connecting orbits appear. For example, at $\mathcal{L}_0$ all the connections we find shadow two basic orbits L_{0A}, L_{0B} and their symmetric counterparts. Is this true of all the connections? Or is this simply an artifact of the fact that we only consider

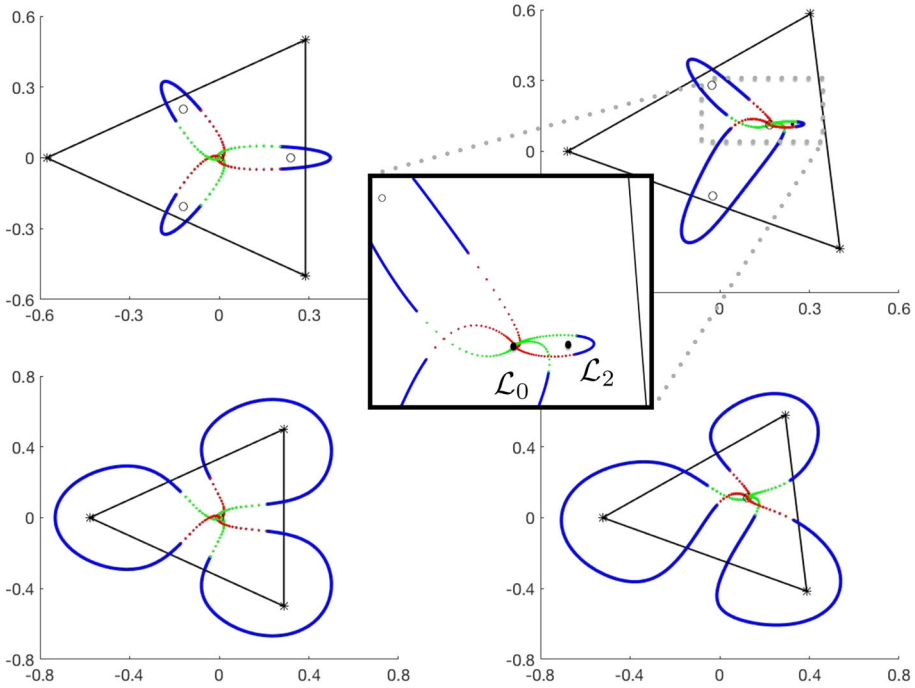

Fig. 21 Continuation at $\mathcal{L}_0$: of all six fundamental connecting orbits at $\mathcal{L}_0$ and their symmetric counterparts. Open circles in the top right frame depict the locations of the inner libration points $\mathcal{L}_{1,2,3}$. (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) Top/bottom left: fundamental connections in the triple Copenhagen problem—equal masses. Top right: final result of numerically continuing L_{0A} and its symmetric counterparts along the line in parameter space beginning at $m_1 = m_2 = m_3 = 1/3$ and ending at $m_1 = 0.415, m_2 = 0.3425, m_3 = 0.2425$—close to where $\mathcal{L}_0$ loses saddle-focus stability near $m_1 \approx 0.42$. During the continuation L_{0A+} shrank substantially. This is due to the fact that L_{0A+} winds around $\mathcal{L}_2$, which collides with $\mathcal{L}_0$ when $m_1 \approx 4.2$. A closeup of the situation is illustrated in the center frame. We observe that the connections L_{0A} and L_{0A-} are deformed much less dramatically. Bottom right: the result of numerically continuing L_{0B} and its symmetric counterparts along the line in parameter space beginning at $m_1 = m_2 = m_3 = 1/3$ and ending at $m_1 = 0.4, m_2 = 0.35, m_3 = 0.25$. These orbits are also deformed less dramatically, though the loops do seem to decrease in size according to the loss of mass in the respective primary, with the largest loop around m_1 and the smallest around m_3. All references to color refer to the online version

connections whose velocity is never too large? Will performing longer searches yield more fundamental letters for the alphabets at $\mathcal{L}_{0,5}$?

Of course with more computing power, one could perform the atlas computations at more values of the mass parameters, say for a mesh of ten or twenty different points in the simplex $m_1 + m_2 + m_3 = 1$. This would provide a more complete picture of the global orbit structure. Such a project would greatly benefit form a cluster computing implementation exploiting the data independence of the computations at different parameter sets, and indeed the independence of different portions of the atlas at a given parameter set. Numerical continuation could then be applied to "fill in the gaps" between the mesh points.

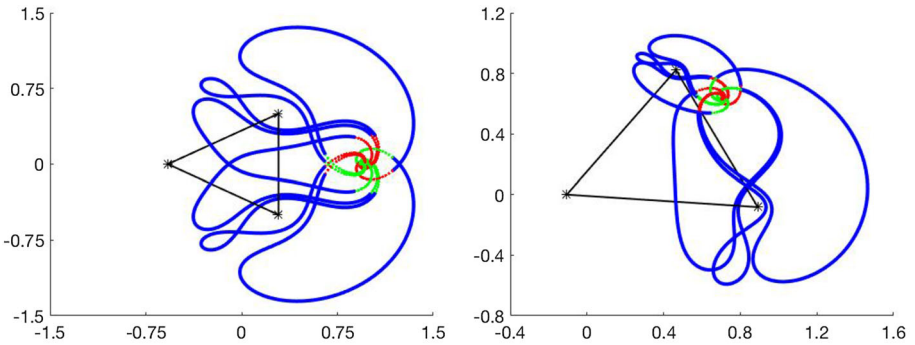

Fig. 22 Continuation at $\mathcal{L}_5$: continuation of the six basic homoclinic motions at $\mathcal{L}_5$, along the line in parameter space beginning at $m_1 = m_2 = m_3 = 1/3$ and ending at $m_1 = 0.89$, $m_2 = 0.1$, $m_3 = 0.01$. (Blue, green and red solid/dotted lines have the same meaning described in Fig. 10.) So, starting from the triple Copenhagen problem, we deform until almost ninety percent of the mass is in the first primary body—near the loss of saddle-focus stability of $\mathcal{L}_5$. Observe that the libration point $\mathcal{L}_5$ moves closer to the smallest primary m_3, and that the loops contract around the smallest primary, a similar situation to that discussed in the caption of Fig. 21. All references to color refer to the online version

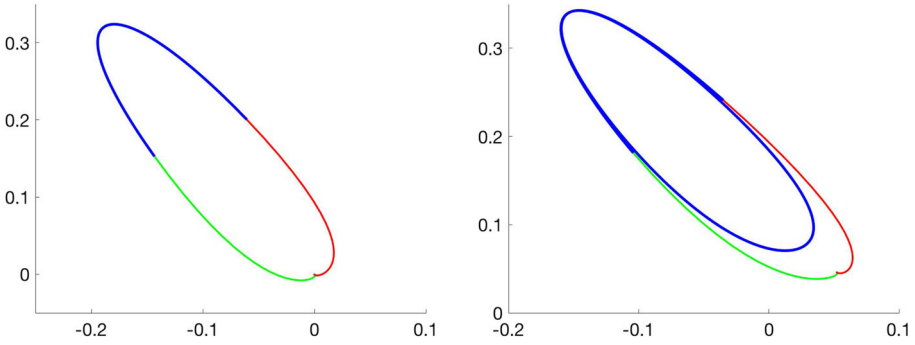

Fig. 23 Illustration of a typical bifurcation: on the left is a homoclinic orbit for $\mathcal{L}_0$ in the $m_1 = m_2 = m_3 = 1/3$ case. On the right is a connecting orbit for the parameter values $m_1 = 0.3617$, $m_2 = 0.34042$ and $m_3 = 0.298$, just after a homoclinic doubling bifurcation of the orbit on the left. Both orbits persist after the bifurcation, that is it seems to be a pitchfork. The new homoclinic has a close "flyby" of $\mathcal{L}_0$ before making a second excursion and finally landing on the stable manifold. In both frames the solid blue lines represent the portion of the orbit represented by the boundary value problem, while the green and red are portions on the unstable and stable manifolds, respectively. All references to color refer to the online version

Acknowledgements The authors would like to sincerely thank two anonymous referees who carefully read the submitted version of the manuscript. Their suggestions greatly improved the final version. The second author was partially supported by NSF Grant DMS-1813501. Both authors were partially supported by NSF Grant DMS-1700154 and by the Alfred P. Sloan Foundation Grant G-2016-7320.

Compliance with ethical standards

Conflict of interest The authors of this manuscript certify that they have no affiliations with or involvement in any organization or entity with any financial interest (such as honoraria, educational grants, participation in speaker's bureaus, membership, employment, consultancies, stock ownership, or other equity interest, and expert testimony or patent-licensing arrangements), or nonfinancial interest (such as personal or professional relationships, affiliations, knowledge or beliefs) in the subject matter or materials discussed in this manuscript.

A Rotational symmetry for the equal mass case

Let $m_1 = m_2 = m_3 = 1/3$ and $\theta = \frac{2\pi}{3}$. Define the linear map, $\varphi : \mathbb{R}^4 \to \mathbb{R}^4$, by

$$\varphi(x, \dot{x}, y, \dot{y}) = \begin{pmatrix} \cos(\theta) & 0 & -\sin(\theta) & 0 \\ 0 & \cos(\theta) & 0 & -\sin(\theta) \\ \sin(\theta) & 0 & \cos(\theta) & 0 \\ 0 & \sin(\theta) & 0 & \cos(\theta) \end{pmatrix} \begin{pmatrix} x \\ \dot{x} \\ y \\ \dot{y} \end{pmatrix} = (\varphi_1, \varphi_2, \varphi_3, \varphi_4)^{\mathrm{T}}.$$

Note that φ acts as a rotation by θ in the (x, y) and $(\dot{x}, \dot{y})$ coordinate planes independently. Now, suppose that $\mathbf{x} : \mathbb{R} \to \mathbb{R}^4$ is a trajectory for f, then $\tilde{\mathbf{x}} = \varphi \circ \mathbf{x}$ is also a trajectory for f. Moreover, if $\mathbf{x} \subset W^{s,u}(\mathcal{L}_i)$ for $i \in \{0, 4, 5, 6\}$, then $\tilde{\mathbf{x}} \subset W^{s,u}(L_{\sigma(i)})$, where σ is the permutation given by $\sigma = (0)(4, 5, 6)$.

Proof Let $\hat{x} = (x, \dot{x}, y, \dot{y}) \in \mathbb{R}^4$ and suppose $\mathbf{x}$ is the trajectory through $\hat{x}$ satisfying $\mathbf{x}(0) = \hat{x}$. By definition, $\tilde{\mathbf{x}}(0) = \varphi(\mathbf{x}(0)) = \varphi(\hat{x})$, and we note that $\tilde{\mathbf{x}}$ will parameterize a trajectory for f if and only if $\tilde{\mathbf{x}}(t)$ is tangent to $f(\tilde{\mathbf{x}}(t))$ for all $t \in \mathbb{R}$. Thus, it clearly suffices to prove that $f \circ \varphi = \varphi \circ f$ holds for any $\hat{x}$ on $\mathbb{R}^4$.

With this in mind, define the planar rotation $\eta : \mathbb{R}^2 \to \mathbb{R}^2$ by

$$\eta(x, y) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} \eta_1(x, y) \\ \eta_2(x, y) \end{pmatrix},$$

Recall that for the symmetric mass case, we have equal masses given by $m_1 = m_2 = m_3 = \frac{1}{3}$. Set $m = \frac{1}{3}$, then the primaries are located at P_1, P_2, P_3 given by

$$P_1 = \left(-\frac{\sqrt{3}}{3}, 0\right) \quad P_2 = \left(\frac{\sqrt{3}}{6}, -\frac{1}{2}\right) \quad P_3 = \left(\frac{\sqrt{3}}{6}, \frac{1}{2}\right)$$

and note that $||P_1|| = ||P_2|| = ||P_3|| = \frac{1}{\sqrt{3}}$. Moreover, P_1, P_2, P_3 are vertices of an equilateral triangle and a direct computation shows that η acts as a cyclic permutation on the primary bodies in configuration space given by the cycle $\pi = (1, 2, 3)$. Recalling that $r_i(x, y) = \sqrt{(x - x_i)^2 + (y - y_i)^2} = ||(x, y) - P_i||$, it follows from this symmetry that for $i \in \{1, 2, 3\}$ we have

$$r_i \circ \eta(x, y) = ||\eta(x, y) - P_i|| = \left|\left|(x, y) - P_{\pi^{-1}(i)}\right|\right| = r_{\pi^{-1}(i)}. \tag{17}$$

Now, we recall that in the symmetric case, the CRFBP vector field is given by

$$f(x, \dot{x}, y, \dot{y}) = \begin{pmatrix} \dot{x} \\ 2\dot{y} + x - \frac{1}{3}\sum_{i=1}^{3} \frac{x - x_i}{r_i} \\ \dot{y} \\ -2\dot{x} + y - \frac{1}{3}\sum_{i=1}^{3} \frac{y - y_i}{r_i} \end{pmatrix},$$

which we write in scalar coordinates as $f = (f_1, f_2, f_3, f_4)$. Similarly, write $\varphi = (\varphi_1, \varphi_2, \varphi_3, \varphi_4)$ and we note that $(\varphi_1(\hat{x}), \varphi_3(\hat{x})) = \eta(x, y)$. Now, we check that $f_i \circ \varphi = \varphi_i \circ f$ holds for each $i \in \{1, 2, 3, 4\}$. For $i = 1$, we have the direction computation

$$\varphi_1 \circ f\left(\hat{x}\right) = \dot{x}\cos(\theta) - \dot{y}\sin(\theta) = f_1 \circ \varphi\left(\hat{x}\right).$$

Now, for $i = 2$ we first compute each expression

$$\varphi_2 \circ f\left(\hat{x}\right) = \left(2\dot{y} + x - \frac{1}{3}\sum_{i=1}^{3} \frac{x - x_i}{r_i(x, y)}\right)\cos(\theta) - \left(-2\dot{x} + y - \frac{1}{3}\sum_{i=1}^{3} \frac{y - y_i}{r_i(x, y)}\right)\sin(\theta)$$

$$f_2 \circ \varphi\left(\hat{x}\right) = 2(\dot{x}\sin(\theta) + \dot{y}\cos(\theta)) + x\cos(\theta) - y\sin(\theta) - \frac{1}{3}\sum_{i=1}^{3}\frac{\eta_1(x,y) - x_i}{r_i \circ \eta(x,y)}.$$

After canceling like terms in each expression, we are left to prove the following equality

$$\sum_{i=1}^{3}\frac{\eta_1(x,y) - x_i}{r_i \circ \eta(x,y)} = \cos(\theta)\sum_{i=1}^{3}\frac{x - x_i}{r_i(x,y)} - \sin(\theta)\sum_{i=1}^{3}\frac{y - y_i}{r_i(x,y)}. \tag{18}$$

Applying the result from (17) to the left side, we have

$$\sum_{i=1}^{3}\frac{\eta_1(x,y) - x_i}{r_{\pi^{-1}(i)}} = \frac{\eta_1(x,y) - x_1}{r_3(x,y)} + \frac{\eta_1(x,y) - x_2}{r_1(x,y)} + \frac{\eta_1(x,y) - x_3}{r_2(x,y)},$$

so that for each $i \in \{1,2,3\}$, the numerator for r_i is given by $\eta_1(x,y) - x_{\pi(i)}$. Now, we compute the numerators for $r_i(x,y)$ on the right-hand side as

$$\cos(\theta)(x - x_i) - \sin(\theta)(y - y_i) = \eta_1(x,y) - \eta_1(x_i,y_i) = \eta_1(x,y) - x_{\pi(i)}.$$

We conclude that the numerators for each r_i are equal, and therefore, the equality in (18) holds which proves that $\varphi_2 \circ f = f_2 \circ \varphi$. The proofs for the $i = 3,4$ cases are computationally similar to the corresponding proofs for $i = 1,2$ which concludes the proof that $f \circ \varphi = \varphi \circ f$, or equivalently, $\tilde{\mathbf{x}}$ is a trajectory for f.

To prove the second claim, fix $i \in \{0,4,5,6\}$ and suppose $\mathbf{x}(t) \to L_i$ as $t \to \infty$ implying that $\mathbf{x} \subset W^s(L_i)$. Let $\tilde{\mathbf{x}} = \varphi(\mathbf{x})$, and note that L_i is an equilibrium solution for f implying that $\mathbf{x}_{2,4}(t) \to 0$. Noting that η is a unitary operator, it follows that $\tilde{\mathbf{x}}_{2,4}(t) \to 0$ as well. Moreover, φ is a dynamical conjugacy implying that in configuration space we have

$$\lim_{t\to\infty}\tilde{\mathbf{x}}_{1,3}(t) = \lim_{t\to\infty}\eta\left(x(t),y(t)\right) = \eta(L_i).$$

Taken together it follows that $\eta(L_i)$ is again an equilibrium solution for f. Thus, η acts as a permutation on equilibria. A direct computation shows that $\eta(L_i) = L_{\sigma(i)}$ where σ is the permutation given by $\sigma = (0)(4,5,6)$. The preceding argument applies equally well to the unstable manifold of each equilibrium with $t \to -\infty$ which completes the proof of the second claim. □

B Power series manipulation, automatic differentiation, and the radial gradient

Our local invariant manifold computations are based on formal power series manipulations. The main technical challenge is to compute $f \circ P$ with P an arbitrary power series and f the vector field for the CRFBP. As usual in gravitational N body problems, the nonlinearity contains terms raised to the minus three halves power.

Consider two formal power series $P, Q\colon \mathbb{C}^2 \to \mathbb{C}$ given by

$$P(z_1,z_2) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty}a_{m,n}z_1^m z_2^n, \qquad \text{and} \qquad Q(z_1,z_2) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty}b_{m,n}z_1^m z_2^n,$$

where $a_{m,n}, b_{m,n} \in \mathbb{C}$ for all $(m,n) \in \mathbb{N}^2$. The collection of all formal power series forms a complex vector space, so that for any $\alpha, \beta \in \mathbb{C}$ we have that

$$(\alpha P + \beta Q)(z_1, z_2) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} \left(\alpha a_{m,n} + \beta b_{m,n}\right) z_1^m z_2^n.$$

The collection becomes an algebra when endowed with the Cauchy product

$$(P \cdot Q)(z_1, z_2) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} \left(\sum_{j=0}^{m}\sum_{k=0}^{n} a_{m-j,n-k} b_{jk} \right) z_1^m z_2^n. \tag{19}$$

We evaluate elementary functions of formal power series using a technique called *automatic differentiation* by many authors. Suppose, for example, we are given a formal series

$$P(z_1, z_2) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} p_{m,n} z_1^m z_2^n,$$

with $p_{0,0} \neq 0$. We seek the formal series coefficients $q_{m,n}$ of the function

$$Q(z_1, z_2) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} q_{m,n} z_1^m z_2^n = P(z_1, z_2)^{\alpha}, \qquad \alpha \in \mathbb{R}.$$

Our approach follows the discussion given by Haro et al. (2016). Consider the first-order partial differential operator

$$\nabla_{\text{rad}} P(z_1, z_2) = \nabla P(z_1, z_2) \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = z_1 \frac{\partial}{\partial z_1} P(z_1, z_2) + z_2 \frac{\partial}{\partial z_2} P(z_1, z_2),$$

which is referred to as *the radial gradient* of P. Evaluating on the level of formal power series leads to

$$\nabla_{\text{rad}} P(z_1, z_2) = \sum_{m=0}^{\infty}\sum_{n=0}^{\infty} (m+n) p_{m,n} z_1^m z_2^n.$$

Observe that

$$\begin{aligned} \nabla_{\text{rad}} Q(z_1, z_2) &= \nabla Q(z_1, z_2) \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \\ &= \nabla P(z_1, z_2)^{\alpha} \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} \\ &= \alpha P(z_1, z_2)^{\alpha - 1} \nabla P(z_1, z_2) \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}. \end{aligned}$$

Multiplying both sides of the equation by P, we obtain

$$P(z_1, z_2) \nabla Q(z_1, z_2) \begin{pmatrix} z_1 \\ z_2 \end{pmatrix} = \alpha Q(z_1, z_2) \nabla P(z_1, z_2) \begin{pmatrix} z_1 \\ z_2 \end{pmatrix}. \tag{20}$$

Here the fractional power is replaced by operations involving only differentiation and multiplication. This is the virtue of the radial gradient in automatic differentiation schemes. Plugging the power series expansions into Eq. (20) leads to

$$\left(\sum_{m=0}^{\infty}\sum_{n=0}^{\infty} p_{m,n} z_1^m z_2^n\right)\left(\sum_{m=0}^{\infty}\sum_{n=0}^{\infty}(m+n) q_{m,n} z_1^m z_2^n\right)$$
$$= \left(\sum_{m=0}^{\infty}\sum_{n=0}^{\infty}\alpha q_{m,n} z_1^m z_2^n\right)\left(\sum_{m=0}^{\infty}\sum_{n=0}^{\infty}(m+n) p_{m,n} z_1^m z_2^n\right),$$

and taking Cauchy products gives

$$\sum_{m=0}^{\infty}\sum_{n=0}^{\infty}\sum_{j=0}^{m}\sum_{k=0}^{n}(j+k) p_{m-j,n-k} q_{j,k} z_1^m z_2^n$$
$$= \sum_{m=0}^{\infty}\sum_{n=0}^{\infty}\sum_{j=0}^{m}\sum_{k=0}^{n}\alpha(j+k) q_{m-j,n-k} p_{j,k} z_1^m z_2^n.$$

Match like powers to get

$$\sum_{j=0}^{m}\sum_{k=0}^{n}(j+k) p_{m-j,n-k} q_{j,k} = \sum_{j=0}^{m}\sum_{k=0}^{n}\alpha(j+k) q_{m-j,n-k} p_{j,k},$$

or

$$(m+n) p_{0,0} q_{m,n} + \sum_{j=0}^{m}\sum_{k=0}^{n}\hat{\delta}_{j,k}^{m,n}(j+k) p_{m-j,n-k} q_{j,k}$$
$$= \alpha(m+n) q_{0,0} p_{m,n} + \sum_{j=0}^{m}\sum_{k=0}^{n}\hat{\delta}_{j,k}^{m,n}\alpha(j+k) q_{m-j,n-k} p_{j,k},$$

for $m+n \geq 1$. Here

$$\hat{\delta}_{j,k}^{m,n} := \begin{cases} 0 & \text{if } \ j = m \text{ and } k = n \\ 0 & \text{if } \ j = 0 \text{ and } k = 0 \\ 1 & \text{otherwise} \end{cases}.$$

The $\hat{\delta}$ appears to remind us that terms of order (m, n) are extracted from the sum. Isolating $q_{m,n}$ gives

$$q_{m,n} = \alpha p_{0,0}^{\alpha-1} p_{m,n} + \frac{1}{(m+n) p_{0,0}} \sum_{j=0}^{m}\sum_{k=0}^{n}\hat{\delta}_{j,k}^{m,n}(j+k)\left(\alpha q_{m-j,n-k} p_{j,k} - p_{m-j,n-k} q_{j,k}\right), \tag{21}$$

for $m+n \geq 1$. Note that $q_{0,0} = p_{0,0}^{\alpha} \neq 0$ by hypothesis, so that the coefficients $q_{m,n}$ are formally well defined to all orders. Using the recursion given in Eq. (21) we can compute the formal series coefficients for Q for the cost of a Cauchy product. This allows us to compute power series representations for the nonlinear terms in $f(P)$ and $Df(P)$ in the CRFBP. Another approach which converts the CRFB field to a higher-dimensional polynomial field in discussed in Kepley and Mireles James (2018).

References

Abraham, R.H.: Chaostrophes, intermittency, and noise. In: Chaos, Fractals, and Dynamics (Guelph, Ont., 1981/1983), Volume 98 of Lecture Notes in Pure and Applied Mathematics, pp. 3–22. Dekker, New York (1985)

Álvarez-Ramírez, M., Vidal, C.: Dynamical aspects of an equilateral restricted four-body problem. Math. Probl. Eng. **23** (2009). Article ID 181360

Alvarez-Ramírez, M., Barrabés, E.: Transport orbits in an equilateral restricted four-body problem. Celest. Mech. Dyn. Astron. **121**(2), 191–210 (2015)

Alvarez-Ramírez, M., Delgado, J., Vidal, C.: Global regularization of a restricted four-body problem. Int. J. Bifurc. Chaos Appl. Sci. Eng. **24**(7), 1450092, 15 (2014)

Alvarez-Ramírez, M., García, A., Palacián, J.F., Yanguas, P.: Oscillatory motions in restricted n-body problems. J. Differ. Equ. **265**, 779–803 (2018)

Baltagiannis, A.N., Papadakis, K.E.: Equilibrium points and their stability in the restricted four-body problem. Int. J. Bifurc. Chaos Appl. Sci. Eng. **21**(8), 2179–2193 (2011a)

Baltagiannis, A.N., Papadakis, K.E.: Families of periodic orbits in the restricted four-body problem. Astrophys. Space Sci. **336**, 357–367 (2011b)

Barrabés, E., Mondelo, J.M., Ollé, M.: Numerical continuation of families of homoclinic connections of periodic orbits in the RTBP. Nonlinearity **22**(12), 2901–2918 (2009)

Barros, J.F., Leandro, E.S.G.: The set of degenerate central configurations in the planar restricted four-body problem. SIAM J. Math. Anal. **43**(2), 634–661 (2011)

Barros, J.F., Leandro, E.S.G.: Bifurcations and enumeration of classes of relative equilibria in the planar restricted four-body problem. SIAM J. Math. Anal. **46**(2), 1185–1203 (2014)

Burgos-García, J.: Families of periodic orbits in the planar Hill's four-body problem. Astrophys. Space Sci. **361**(11), 353, 21 (2016)

Burgos-García, J., Bengochea, A.: Horseshoe orbits in the restricted four-body problem. Astrophys. Space Sci. **362**(11), 212, 14 (2017)

Burgos-García, J., Delgado, J.: Periodic orbits in the restricted four-body problem with two equal masses. Astrophys. Space Sci. **345**(2), 247–263 (2013a)

Burgos-García, J., Delgado, J.: On the "blue sky catastrophe" termination in the restricted four-body problem. Celest. Mech. Dyn. Astron. **117**(2), 113–136 (2013b)

Burgos-García, J., Gidea, M.: Hill's approximation in a restricted four-body problem. Celest. Mech. Dyn. Astron. **122**(2), 117–141 (2015)

Burgos-García, J., Lessard, J.P., Mireles James, J.D.: Spatial periodic orbits in the equaliteral circular restricted four body problem: computer assisted proofs of existence. Celest. Mech. Dyn. Astron. **131**, 2 (2019)

Cabré, X., Fontich, E., de la Llave, R.: The parameterization method for invariant manifolds. I. Manifolds associated to non-resonant subspaces. Indiana Univ. Math. J. **52**(2), 283–328 (2003a)

Cabré, X., Fontich, E., de la Llave, R.: The parameterization method for invariant manifolds II. Regularity with respect to parameters. Indiana Univ. Math. J. **52**(2), 329–360 (2003b)

Cabré, X., Fontich, E., de la Llave, R.: The parameterization method for invariant manifolds. III. Overview and applications. J. Differ. Equ. **218**(2), 444–515 (2005)

Canalias, E., Masdemont, J.J.: Homoclinic and heteroclinic transfer trajectories between planar Lyapunov orbits in the sun–earth and earth–moon systems. Discrete Contin. Dyn. Syst. Ser. A **14**(2), 261–279 (2006)

Champneys, A.R., Kuznetsov, Y.A., Sandstede, B.: A numerical toolbox for homoclinic bifurcation analysis. Int. J. Bifurc. Chaos Appl. Sci. Eng. **6**(5), 867–887 (1996)

Cheng, X., She, Z.: Study on chaotic behavior of the restricted four-body problem with an equilateral triangle configuration. Int. J. Bifurc. Chaos Appl. Sci. Eng. **27**(2), 1750026, 12 (2017)

Darwin, G.H.: Periodic orbits. Acta Math. **21**(1), 99–242 (1897)

de la Llave, R., Mireles James, J.D.: Parameterization of invariant manifolds by reducibility for volume preserving and symplectic maps. Discrete Contin. Dyn. Syst. **32**(12), 4321–4360 (2012)

Devaney, R.L.: Homoclinic orbits in Hamiltonian systems. J. Differ. Equ. **21**(2), 431–438 (1976)

Devaney, R.L.: Blue sky catastrophes in reversible and Hamiltonian systems. Indiana Univ. Math. J. **26**(2), 247–263 (1977)

Doedel, E.J., Friedman, M.J.: Numerical computation of heteroclinic orbits. J. Comput. Appl. Math. **26**(1–2), 155–170 (1989). (Continuation techniques and bifurcation problems)

Doedel, E.J., Friedman, M.J., Kunin, B.I.: Successive continuation for locating connecting orbits. Numer. Algorithms **14**(1–3), 103–124 (1997). [Dynamical numerical analysis (Atlanta, GA, 1995)]

Gidea, M., Burgos, M.: Chaotic transfers in three- and four-body systems. Phys. A **328**(3–4), 360–366 (2003)

Haro, À., Canadell, M., Figueras, J.-L., Luque, A., Mondelo, J.-M.: The Parameterization Method for Invariant Manifolds, Volume 195 of Applied Mathematical Sciences. Springer, Cham (2016). (From rigorous results to effective computations)
Henrard, J.: Proof of a conjecture of E. Strömgren. Celest. Mech. **7**, 449–457 (1973)
Kalies, W., Kepley, S., Mireles James, J.D.: Analytic continuation of local (un)stable manifolds with rigorous computer assisted error bounds. SIAM J. Appl. Dyn. Syst. **17**(1), 157–202 (2018)
Kepley, S., Mireles James, J.D.: Chaotic motions in the restricted four body problem via Devaney's saddle-focus homoclinic tangle theorem. J. Differ. Equ. **226**(4), 1709–1755 (2019)
Koon, W.S., Lo, M.W., Marsden, J.E., Ross, S.D.: Heteroclinic connections between periodic orbits and resonance transitions in celestial mechanics. Chaos **10**(2), 427–469 (2000)
Leandro, E.S.G.: On the central configurations of the planar restricted four-body problem. J. Differ. Equ. **226**(1), 323–351 (2006)
Lerman, L.M.: Complex dynamics and bifurcations in a Hamiltonian system having a transversal homoclinic orbit to a saddle focus. Chaos **1**(2), 174–180 (1991)
Lerman, L.M.: Dynamical phenomena near a saddle-focus homoclinic connection in a Hamiltonian system. J. Stat. Phys. **101**(1/2), 357–372 (2000)
Lessard, J.-P., Mireles James, J.D., Ransford, J.: Automatic differentiation for Fourier series and the radii polynomial approach. Phys. D **334**(1), 174–186 (2016)
Moulton, F.R., Buchanan, D., Buck, T, Griffin, F.L., Longley, W.R., MacMillan, W.D.: Periodic orbits. Number Publication No. 161. Carnegie Institution of Washington (1920)
Papadakis, K.E.: Families of asymmetric periodic solutions in the restricted four-body problem. Astrophys. Space Sci. **361**(12), 377, 15 (2016a)
Papadakis, K.E.: Families of three-dimensional periodic solutions in the circular restricted four-body problem. Astrophys. Space Sci. **361**(4), 129, 14 (2016b)
Pedersen, P.: Librationspunkte im restringierten vierkörperproblem. Dan. Mat. Fys. Medd. **21**(6), 1–80 (1944)
Pedersen, P.: Stabilitätsuntersuchungen im restringierten vierkörperproblem. Dan. Mat. Fys. Medd. **26**(16), 1–38 (1952)
Rabe, E.: Determination and survey of periodic Trojan orbits in the restricted problem of three bodies. Astron. J. **66**, 500–513 (1961)
Rump, S.M.: INTLAB—INTerval LABoratory. In: Csendes, T. (ed.) Developments in Reliable Computing, pp. 77–104. Kluwer Academic Publishers, Dordrecht (1999). http://www.ti3.tu-harburg.de/rump/
She, Z., Cheng, X.: The existence of a Smale horseshoe in a planar circular restricted four-body problem. Celest. Mech. Dyn. Astron. **118**(2), 115–127 (2014)
She, Z., Cheng, X., Li, C.: The existence of transversal homoclinic orbits in a planar circular restricted four-body problem. Celest. Mech. Dyn. Astron. **115**(3), 299–309 (2013)
Shilńikov, L.P.: Existence of a countable set of periodic motions in a four-dimensional space in an extended neighborhood of a saddle-focus. Dokl. Akad. Nauk SSSR **172**, 54–57 (1967)
Shilńikov, L.P.: On the question of the structure of an extended neighborhood of a structurally stable state of equilibrium of saddle-focus type. Mat. Sb. (N.S.) **81**(123), 92–103 (1970a)
Shilńikov, L.P.: A contribution to the problem of the structure of an extended neighborhood of a rough equilibrium state of saddle-focus type. Math. USSR Sb. **10**(1), 91 (1970b)
Shilnikov, L.P., Shilnikov, A.L., Turaev, D.V.: Showcase of blue sky catastrophes. Int. J. Bifurc. Chaos Appl. Sci. Eng. **24**(8), 1440003, 10 (2014)
Simó, C.: Relative equilibrium solutions in the four-body problem. Celest. Mech. **18**(2), 165–184 (1978)
Strömgren, E.: Connaissance actuelle des orbites dans le probleme des trois corps. Bull. Astron. **9**, 87–130 (1934)
Szebehely, V.: Theory of Orbits. Academic Press Inc., Cambridge (1967)
Szebehely, V., Flandern, T.V.: A family of retegrade orbits around the triangular equilibrium points. Astron. J. **72**(3), 373–379 (1967)
Szebehely, V., Nacozy, P.: A class of E. Strömgren's direct orbits in the restricted problem. Astron. J. **77**(2), 184–190 (1967)
Van den Berg, J.B., Mireles James, J.D., Reinhardt, C.: Computing (un)stable manifolds with validated error bounds: non-resonant and resonant spectra. J. Nonlinear Sci. **26**, 1055–1095 (2016)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.