

University of Southampton Research Repository

Copyright © and Moral Rights for this thesis and, where applicable, any accompanying data are retained by the author and/or other copyright owners. A copy can be downloaded for personal non-commercial research or study, without prior permission or charge. This thesis and the accompanying data cannot be reproduced or quoted extensively from without first obtaining permission in writing from the copyright holder/s. The content of the thesis and accompanying research data (where applicable) must not be changed in any way or sold commercially in any format or medium without the formal permission of the copyright holder/s.

When referring to this thesis and any accompanying data, full bibliographic details must be given, e.g.

Thesis: Jack Tyler (2022) "Exploitation of Three-body Dynamics in Space Mission Design", University of Southampton, Faculty of Engineering and Physical Sciences, PhD Thesis.



University of Southampton

Faculty of Engineering and Physical Sciences

School of Engineering

A thesis for the degree of Doctor of Philosophy

Exploitation of Three-body Dynamics in Space Mission Design

Doctoral Thesis of:

Jack Tyler, BEng

ORCiD: [0000-0002-0979-7533](https://orcid.org/0000-0002-0979-7533)

8 January 2023

Discover the force of the skies O Men: once recognised it can be put to use.
– Johannes Kepler

Abstract

Exploitation of Three-body Dynamics in Space Mission Design

by Jack Tyler

With renewed interest in space exploration, the question of designing efficient transfers between celestial bodies is as relevant as ever. Combined with the rise of scientific computing in the latter half of the 20th century, significant attention has been given to using modern computing methods to design more efficient orbital transfers to reduce costs and improve overall mission lifetime.

Preliminary mission design is often performed in simplified, time-independent models of motion. In these, classical dynamical systems theory identifies dynamical structures which can be used to create low-energy transfers between points in space, or used to create orbits that exist as a delicate balance of gravity to achieve mission objectives. However, in more realistic, time-dependent models such structures are not guaranteed to exist. Attention has thus been given to techniques well-developed in fluid dynamics to identify similar structures in astrodynamics systems, but numerical and computational difficulties have frustrated these efforts.

This PhD is separated into two parts. The first studies the use of time-independent dynamical structures in space mission design in combination with high-performance computing techniques. An intensive optimisation procedure is used to construct transfers that use the invariant manifolds to retrieve asteroids into two of the equilibrium points of the Sun-Earth system. This PhD improves on the state of the art in this field by improving the methods used to find and construct the retrieval transfers. As a result, 27 more asteroids that are considered ‘easily retrievable’ are found, and the velocity required to compute the transfers is generally reduced for those already considered easily retrievable. Moreover, it is revealed that these transfers exist across a range of transfer times, allowing greater flexibility for mission designers.

The second part of this thesis uses techniques from fluid mechanics to find analogous structures to those used in the asteroid retrieval study directly in time-dependent models of motion, rather than needing to use simplified models. This thesis makes two improvements to the current body of research: the first is the presentation of an improved numerical method to compute Lagrangian Coherent Structures (LCS) in three-dimensional dynamical systems called DA-LCS. This numerical method uses a direct computer implementation of an algebra of polynomials to compute more accurate and less numerically noisy quantities that signal LCS in general dynamical systems, and greatly outperforms standard approaches in astrodynamics systems. Since the relevant quantities are computed as polynomial expansions, the method also allows the computation of all relevant quantities completely automatically.

This numerical method is then applied to a series of test cases from astrodynamics, where three-dimensional LCS is constructed and shown to perform the role of generalised unstable manifolds in astrodynamics systems by separating qualitatively different behaviour. The effect of orbit parameterisation and integration time is also elaborated in an effort to provide the space community with in-depth knowledge of how to use LCS in astrodynamics in future studies.

Table of Contents

List of Figures	xi
List of Tables	xix
Declaration of Authorship	xxiii
1 Introduction	3
2 The Three-body Problem	11
2.1 Introduction	11
2.2 Derivation of the Equations of Motion	13
2.2.1 Keplerian Motion	13
2.2.2 Derivation of the ER3BP Equations of Motion	15
2.2.3 The Circular-Restricted Three-body Problem	19
2.3 Transforming between inertial and rotating reference frames	21
2.4 Conclusion	23
3 Dynamical Features and Dynamical Indicators in Astrodynamics	27
3.1 Introduction	27
3.2 Dynamical features in autonomous systems	28
3.3 Dynamical features and flow heuristics in nonautonomous systems	32
3.4 Lagrangian Coherent Structures	35
3.5 Conclusion	42
4 Asteroid Retrieval Missions using the Invariant Manifolds of the Circular Problem	47
4.1 Background	47
4.2 The Methodology	52
4.2.1 Obtaining Near-Earth Object Data	52
4.2.2 Orbit and manifold generation	53
4.2.3 Filtering Unsuitable NEOs	57
4.2.4 Finding optimal transfers	62
4.3 Implementation ‘tips’	64
4.4 Results	65
4.5 Pareto Front Generation	69
4.5.1 Naïve MIDACO	69
4.5.2 Naïve scanning	71
4.5.3 Global scanning	72
4.5.4 Fully converged global scan	72
4.6 Conclusion	76

5 An improved numerical method for three-dimensional hyperbolic Lagrangian Coherent Structures	81
5.1 Background	81
5.2 A review of the mathematical background and notation	84
5.3 Differential algebra	85
5.3.1 Flow expansion using DA	88
5.3.2 Computing expansions of leading eigenvectors of $C_{t_0}^T$ to arbitrary order	89
5.4 Lagrangian Coherent Structures	92
5.5 Application to Arnold-Beltrami-Childress Flows	95
5.5.1 Steady Arnold-Beltrami-Childress flow	96
5.5.2 Periodic Arnold-Beltrami-Childress Flow	104
5.5.3 Chaotically-forced Arnold-Beltrami-Childress flow	110
5.6 Conclusion	112
6 Lagrangian Coherent Structures in the Elliptic-Restricted Three-body Problem	117
6.1 Introduction	117
6.2 Definition of Initial Conditions	118
6.2.1 Orbital elements	120
6.2.2 Spherical coordinates	121
6.2.3 Cartesian coordinates	122
6.3 Producing the three-dimensional surface	123
6.4 A proof-of-concept	124
6.4.1 Results	127
6.5 Pushing the envelope	130
6.5.1 The LCS Structure	131
6.5.2 Effect of orbit parameterisation and integration time	135
6.5.3 Implementation considerations	138
6.6 Conclusion	141
7 Conclusions	145
7.1 Limitations and future work	146
Appendix A Derivation of the Tisserand parameter from Keplerian elements	155
Appendix B Inverting a matrix of DA objects	161

List of Figures

1.1	Number of satellites deployed per year; recent years have seen a marked increase in space missions, in part due to the reduction in launch costs. Figure adapted from Our World in Data [1].	4
2.1	Graphical representation of a subset of the Keplerian orbital elements. Not pictured are the semimajor axis a and the eccentricity e .	14
2.2	The geometry of an ellipse, as used in the solution of the orbit equation for $0 < e < 1$. a is the semi-major axis of the orbit, and e the eccentricity. The periapsis is the closest point on the orbit to the central body, and the apoapsis the farthest.	15
2.3	The coordinate systems considered in the derivation of the equations of motion of the CR3BP. In grey is a fixed (inertial) system with its origin at the barycentre of m_1 and m_2 ; in black is the rotating (synodic) frame with m_1 and m_2 set fixed.	16
2.4	The location of the Lagrangian (equilibrium) points in the CR3BP. L_1 through L_3 are known as the collinear points, and L_4 and L_5 the equilateral solutions.	20
3.1	Example periodic orbits that exist about Sun-Earth L_1 and L_2 . Pictured are the vertical and planar Lyapunov family, and the quasiperiodic ‘Halo’ orbits. These families are used as destinations for asteroids in chapter 4.	29
3.2	Stable manifold associated to a Halo orbit about Sun-Earth L_1 , represented as its projection onto spatial coordinates.	31
3.3	Example LD field on the Sun-Mars ER3BP (left) and the ‘separatrices’ obtained after performing an edge detection algorithm with manually-determined behaviours overlaid in different colours (right). The integrand used is $ v ^{\frac{1}{2}}$, which highlights features well. However, the choice of integrand is generally not known <i>a priori</i> and must practically be selected through trial-and-error. Figure taken from Quinci, Merisio, and Toppotto [2].	33
3.4	Comparison between a surface that maximises repulsion and a surface that causes trajectories to ‘shear’ across its surface. The former is the desired behaviour of a repelling hyperbolic LCS, the second is the behaviour of a shearing surface, which can in some cases be erroneously highlighted by the FTLE field.	36
3.5	Elliptic LCS, which produce time-dependent generalisations of the KAM torus. Shown here is the elliptic LCS for the chaotically-forced Arnold-Beltrami Chil-dress flow, taken from Blazevski and Haller [3].	37
3.6	Geometry of a surface \mathcal{M} , including the definition of the normal N and tangent spaces T to a point x_0 on the surface used in definitions of LCS from an FTLE ridge.	38
3.7	Previous attempts at computing LCS in the Elliptic-Restricted Three-body Problem. Both authors cited numerical issues as preventing the construction of the full LCS.	41

4.1	The number of NEOs detected over time by the Centre for Near-Earth Object Studies (CNEOS), grouped by diameter. Figure adapted from NASA JPL [4]. . .	49
4.2	The families of periodic orbits studied in this investigation: planar and vertical Lyapunov orbits and Northern and Southern ‘Halo’ orbits. Only the orbits about L_2 are shown for brevity, and only orbits used in previous literature are used to ensure a like-for-like comparison. Highlighted in red are portions of the two retrieval trajectories for asteroids 2020 DW (planar Lyapunov) and 2006 RH120 (Northern Halo), and the nominal trajectory for 2006 RH120 (dashed blue). . . .	54
4.3	The $\pm \frac{\pi}{8}$ planes used as the reference fixed section in this paper: $+\frac{\pi}{8}$ is for manifolds belonging to orbits about L_2 and $-\frac{\pi}{8}$ is for orbits about L_1 . Outside of the ‘cone’ formed by the two planes, the dynamics can be assumed to be two-body only around the Sun; inside of the cone, full CR3BP dynamics are used.	56
4.4	The Hohmann transfers considered in the pre-filtering methodology; the blue orbit is the orbit of the NEO; the green orbit is the orbit defined by the equivalent Keplerian elements of the manifold target point at the chosen fixed section; and the red arc is the Hohmann transfer. Since the orientation of the semimajor axes of the orbits is ignored, four possible combinations of transfer exist as one of the inclination change combinations may be discarded as more expensive.	58
4.5	Schematic of the Hohmann transfer between two circular, coplanar orbits. The initial orbit has radius r_1 and the final orbit radius r_2	59
4.6	Geometry of the inclination change problem between two orbits with the same orbital speed. The inclination change is assumed to be a pure rotation about a common point on the two orbits.	60
4.7	Pre-filter cost (black), Jacobi energy (blue) and semi-major axis (red) for asteroid 3410533. Preliminary analysis shows the potential for pre-filter cost to vary with time as a result of changing orbital elements from close approaches to planets, for example. In this case the asteroid does not pass the pre-filter when run after a certain date. Even filters used in prior works and based on the Jacobi energy (blue) would vary with time.	61
4.8	Intersections of stable manifolds and the $\pi/8$ plane for selected periodic orbits from the planar Lyapunov family (green), Northern Halo family (blue), and vertical Lyapunov family (red). Design variables $n_{\text{mnfd}} \in [1, 360]$ and $K \in [1, 8000]$ are indicated.	62
4.9	Number of retrieval candidates (light blue) and EROs (dark blue) for a given pre-filter $\Delta\nu$ threshold. While 66% of EROs are found for 700 m/s (dashed line), a higher threshold of 3000 m/s ensures no EROs are excluded prematurely, albeit at significant computational cost in the pre-filter.	66
4.10	Distribution of retrievable NEOs over the range of optimal capture $\Delta\nu$. The dark blue bars correspond to EROs, but many more NEOs have very similar capture $\Delta\nu$. EROs are subject to more rigorous optimisation, causing the gap between dark blue and light blue bars.	68
4.11	Pareto front for 2011 BL45 computed by MIDACO’s in-built Pareto front generation and MIDACO’s PlotTool. Only the points that define the front are included, and intermediate information on the dominating solutions is lost.	69
4.12	Preliminary transfer time t_t vs. capture $\Delta\nu$ space for 2011 BL45, found by continuing the solution at each value of t_t from the optimal solution at the previous t_t . This produces an approximately periodic cost across the range of t_t , with large differences between the lowest and highest capture $\Delta\nu$ s.	70

4.13	Preliminary structure of the $\Delta\nu$ vs. transfer time space for asteroid 2009 BD, when only the optimal solution at the previous transfer time is used as an initial guess for the parameters in the optimisation. The orbit of the NEO is in black, in red is the equivalent orbit of the manifold target point, and in blue is the transfer orbit between those points.	71
4.14	Pareto front for 2011 BL45 when generated using multiple initial guesses for each value of t_t . There is a pronounced ‘floor’ in the optimisation but transfers with similar t_t can have significantly different $\Delta\nu$	73
4.15	Pareto front for NEO 2011 BL45. While not an ERO, with an optimal capture $\Delta\nu$ of approximately 625 m/s, it still displays the same behaviour. Non-EROs may, therefore, offer similar flexibility to the mission designer as EROs.	73
4.16	Pareto front for 2012 TF79, an ERO with all transfers below the ERO threshold, including a zero-time, single impulse transfer. The Pareto front is shown as a black line.	74
5.1	The approximation of divided differences to computing derivatives across a two-dimensional grid on which the values of the function values are known. In practice, the offsets $\delta x_0, \delta x_1$ are often determined through trial-and-error, but are a key parameter in the approximation.	82
5.2	Comparison between computing $1/(2+1)$ in the field of real numbers \mathbb{R} and the the field of floating-point approximations to \mathbb{R}, \mathbb{F} . Each operation in \mathbb{R} has an adjoint operation in \mathbb{F} . Figure taken from [5].	86
5.3	Comparison of the manipulation of k -times differentiable functions in C^r and DA using the example of $1/(x+1)$; operations in C^k have adjoint operations in the DA space \mathcal{P} . The final result of the manipulation in DA is a Taylor polynomial which locally approximates the function in C^k up to a given expansion order. Figure taken from [5].	87
5.4	Example of C++ code using DA via the DACE [6] (left) and the output (right) to show how operator overloading makes using DA in modern codes straightforward, and how the output is a high-order truncated Taylor polynomial of $\sin(1+x)$	87
5.5	Relative error across all polynomial orders in successive applications of $[C_{t_0}^T]$ to an initial guess containing only the floating-point dominant eigenvector at the expansion point as the constant part. Higher expansion orders (black) can all be seen converging at around the expected convergence rate λ_n/λ_{n-1} (dashed red) towards the floating-point floor.	91
5.6	Poincaré section (return map) for the steady ABC flow on the $z = 0$ plane; generated using a 15×15 grid of initial points with integration time $T = 1500$	96
5.7	Finite-time Lyapunov fields for the steady ABC flow from $t = 0$ to $T = 3$ using DA-LCS and divided differences. The relative error is below 3×10^{-5} , suggesting that computing the FTLE associated with C_0^3 using divided differences is not a major source of error in this example.	97
5.8	Helicity fields for the steady ABC flow from $t_0 = 0$ to $T = 3$ using DA-LCS and divided differences. Again both strongly agree, showing that DA-LCS is working. The DA-LCS structure is a little smoother along the main ridge on the right compared to divided differences, making the identification of seed points more robust.	98
5.9	The seed points (black) for the strainline ODE obtained for different values of helicity thresholds for the steady ABC flow on the $z = 0$ plane, and the reference solution from Blazevski and Haller [3].	99

5.10	Final, filtered strainlines for the steady ABC flow on the $z = 0$ plane computed using DA-LCS. The structure is formed of approximately 53 strainline segments.	100
5.11	The strainline structure found for the steady ABC flow overlaid on two plots highlighting the underlying dynamics.	101
5.12	Finite-time Lyapunov exponent fields for the periodic ABC flow from $t_0 = 0$ to $T = 4.0$, obtained using DA-LCS and divided differences. Again, the FTLE field agrees between divided differences and DA-LCS, suggesting divided differences on the correct auxiliary grid in this case accurately approximates the FTLE associated with C_0^4 .	103
5.13	Helicity fields for the periodic ABC flow computed using DA-LCS and divided differences from $t_0 = 0$ to $T = 4.0$. Here DA-LCS highlights in particular the main ridge on the right more clearly and smoothly than divided differences.	105
5.14	Final strainlines for the periodic ABC flow on the $z = 0$ plane computed using DA-LCS, after filtering. The strainline structure is composed of approximately 50 strainline segments.	106
5.15	The seed points (black) for the strainline ODE obtained for different values of helicity thresholds for periodic ABC flow on the $z = 0$ plane. The helicity thresholds in this paper are chosen to give good coverage of the first ‘continuous’ structure that emerges. In the bottom-right is the reference from Blazevski and Haller [3].	107
5.16	Finite-time Lyapunov exponent field for the chaotically-forced ABC flow and an integration time from $t_0 = 0$ to $T = 5.0$. Differences are beginning to become visible on the ‘ends’ of the main wishbone-like structure when using divided differences due to the spiky FTLE values.	108
5.17	Helicity fields for the chaotically-forced ABC flow from $t_0 = 0$ to $T = 5.0$. DA-LCS produces visibly better-defined ridges, helping to robustly identify seed points. However, the relative errors are very high due to the spiky nature of the ridges in both DA-LCS and divided differences.	109
5.18	The seed points (black) for the strainline ODE obtained for different values of helicity thresholds for the forced ABC flow on the $z = 0$ plane.	111
5.19	Final strainline structure on the $z = 0$ plane for the chaotically-forced ABC flow computed using DA-LCS. The structure is formed of 57 individual strainline segments.	112
6.1	Graphical depiction of the parameterisation of an initial condition using standard Keplerian elements. The light green plane represents an example hyperplane of fixed i .	119
6.2	Graphical depiction of the spherical parameterisation of an initial condition. The light green ‘cone’ represents an example hyperplane of fixed ϕ .	121
6.3	Graphical depiction of the Cartesian parameterisation of an initial condition. The light green planes represent example hyperplanes of fixed z .	122
6.4	Diagrams to show how the full LCS surface is reconstructed by interpolating between strainlines on neighbouring reference planes.	123
6.5	Finite-time Lyapunov exponent field for the Elliptic-Restricted Three-body Problem on the $\phi = 115^\circ$ plane from $t_0 = v_0 = 0$ to $T = v = 2\pi$. While the structure is qualitatively the same, the ridges in the FTLE field are more well-defined with DA-LCS.	125

6.6	$-\log H_{\zeta_n}$ for the Elliptic-Restricted Three-body Problem on the $\phi = 115^\circ$ plane from $t_0 = \nu_0 = 0$ to $T = \nu = 2\pi$. No defined ridges of low helicity are visible with divided differences, but DA-LCS can extract well-defined ridges on which the strainline integration can be performed without stopping prematurely.	126
6.7	Strainlines on the $\phi = 115^\circ$ plane for the Elliptic-Restricted Three-Body Problem computed using DA-LCS. Divided differences is unable to generate any strainlines, but with DA-LCS the structure of the LCS can be deduced readily and with only 8 strainlines.	128
6.8	Strainlines with the origins of sample trajectories in the rotating frame overlaid. Mars is highlighted in red in the centre of the Figure. The LCS is acting as expected in separating regions of qualitatively different behaviour, validating the strainlines found.	129
6.9	Rendering of the full 3D LCS for the ER3BP test case. The main structure is the two large curving ‘arms’, which previous literature has struggled to reproduce [7, 8].	130
6.10	The full structure of the LCS obtained, which is comprised of an inner section and two large ‘arms’ which emanate from the inner section. The region in red is the point-of-view of the section plot shown in more detail in Figure 6.12.	132
6.11	A section view of one of the hyperplanes of fixed inclination with points that form the strainlines shown in black to highlight more clearly how the surface is reconstructed.	133
6.12	Inner section of the LCS (highlighted in red in Figure 6.10) with the ‘arms’ removed and three sample trajectories overlaid to show how it is acting as a separatrix. The trajectory taken from inside the LCS (red) remains orbiting Mars while the two trajectories that begin nearby but outside the inner section (blue) both immediately escape.	133
6.13	Trajectories in the synodic frame of points taken from different parts of the LCS arms. Each trajectory is qualitatively different and the LCS is again acting as a separatrix.	134
6.14	The effect of integration time on the helicity field for the ER3BP. The initial true anomaly is $t_0 = 0$ and the final true anomaly is either $T = \pi/4$ (68 days, top row), or $T = 2\pi$ (one Martian year, bottom row). In general, sufficient time is required to remove the effect of the transformations on the final derivatives. Regions of low helicity (dark blue) imply the location of an LCS.	136
6.15	The helicity fields presented for each of the transformations using two unit lengths: Mars radius and Sun-Mars distance. One can aid the conditioning of relevant terms by scaling the unit length such that $C_{t_0}^T$ is well-conditioned in parameterisations with some bounded terms.	139

List of Tables

4.1	Unique identifiers K used to identify periodic orbits and the orbit types they represent.	56
4.2	The ranges of Jacobi energy for which orbits were generated, inclusive. 1000 orbits were generated per family with equi-spaced x-coordinates between the extrema given above. The energy ranges for Halo orbits extend to both the Northern (L1N and L2N for L_1 and L_2) and Southern types (L1S and L2S for L_1 and L_2).	57
4.3	Definitions of the five design variables used in the optimisation procedure. Taken together, n_{mnfd} and t_{end} uniquely specify the insertion point in the stable manifold of the K^{th} target orbit. Dates are given in the form YYYY/MM/DD.	64
4.4	List of $\Delta\nu$ -optimal solutions found for all of the EROs in this work; I find 17 new trajectories with retrieval costs under 100 m/s. Note that while these transfers may be optimal in terms of $\Delta\nu$, many more solutions exist for other transfer times with similar retrieval velocities. All dates are given in the form YYYY/MM/DD HH:MM:SS.	67
4.5	Table of the $\Delta\nu$ -optimal solutions for the EROs found, and the ratio of single-impulse (S.I.) $\Delta\nu$ to the optimal $\Delta\nu$; many single-impulse transfers exist at similar or less cost than two-impulse transfers. 24 EROs are retrieved into different periodic orbits for the single-impulse transfer.	75
5.1	Core time required to compute the LCS on one reference plane for the steady ABC flow using divided differences and DA-LCS on Intel Xeon E5-2670 processors. While DA-LCS is slower to determine the initial H_{ζ_3} field, it is quicker at the integration of a representative set of strainlines and can grow much longer strainlines with the same number of evaluations of Equation 5.13 as divided differences. Importantly, divided differences requires significant grid size tuning, which may make the time required to determine H_{ζ_3} slower overall when used practically.	101

Declaration of Authorship

I declare that this thesis and the work presented in it is my own and has been generated by me as the result of my own original research.

I confirm that:

1. This work was done wholly or mainly while in candidature for a research degree at this University;
2. Where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
3. Where I have consulted the published work of others, this is always clearly attributed;
4. Where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
5. I have acknowledged all main sources of help;
6. Where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
7. Parts of this work have been published as:
 - Tyler, Jack and Wittig, Alexander. On asteroid retrieval missions enabled by invariant manifold dynamics. *Acta Astronautica*, 183:43–51, 2021. ISSN 00945765. doi: [10.1016/j.actaastro.2021.03.002](https://doi.org/10.1016/j.actaastro.2021.03.002)
 - Tyler, Jack and Wittig, Alexander. An improved numerical method for three-dimensional hyperbolic Lagrangian Coherent Structures using Differential Algebra. Accepted for publication in the *Journal of Computational Science*. In press.
 - Tyler, Jack and Wittig, Alexander. Three-dimensional Lagrangian Coherent Structures in the Elliptic-Restricted Three-body Problem. Under review, submitted to: *Celestial Mechanics and Dynamical Astronomy*.
 - Tyler, Jack and Wittig, Alexander. Flexible and Single-Impulse Transfers for Asteroid Retrieval using the Invariant Manifolds of the Circular-Restricted Three-Body Problem. In: *31st AAS/AIAA Space Flight Mechanics Meeting, Charlotte, North Carolina*. American Institute of Aeronautics and Astronautics, 2021.
 - Three-dimensional Lagrangian Coherent Structures in the ER3BP using Differential Algebra. In: *32nd AAS/AIAA Space Flight Mechanics Meeting, Charlotte, North Carolina*. American Institute of Aeronautics and Astronautics, 10-14 August 2022.

- Using Differential Algebra to Compute Lagrangian Coherent Structures for Mission Design and Analysis. In: *73rd International Astronautical Congress, Paris, France*. International Astronautical Federation, 18-22 September 2022.

Signed:

Date: 8 January 2023.

1

Introduction

The economics of access to space has been revolutionised in the past decade, with the cost of launching to orbit dramatically decreasing as a result of rapid re-usability of launch vehicles and growing commercial competition. As early as the latter half of the 1960s, NASA realised that for access to space and the resources it offers to become commonplace, keeping costs low was to be paramount. Amid a backdrop of a dwindling budget, NASA instigated the first studies into the Space Transportation System ('Space Shuttle'), a partially-reusable craft that would be recycled between launches to avoid the costly need to construct one-time-use launch vehicles. NASA designers believed that this would revolutionise access to space by dramatically reducing the cost of launching a kilogram to low-Earth orbit to as low as \$3100 / kg at 2022 values of the dollar. The difficulty of designing for partial reusability, conflicting design requirements from NASA's various governmental partners, and a launch rate well below the 24 per year estimated during the initial design phase, meant that the cost per kilogram ended up being \$105 041.36 [9].

When the Space Shuttle's high costs became clear, President Ronald Reagan signed the Commercial Space Launch Act in 1985, designed to formally encourage private, commercial entities to enter the space launch sector. It was hoped that the introduction of free-market competition would drive down prices to launch to orbit. The (now) Lockheed Martin Atlas II was introduced in 1988, and the McDonnell (Boeing) Delta II rocket followed in 1989, which were workhorses for the US government, US military, and commercial partners. In Europe, Arianespace became the first commercial launch service provider in 1982. Later efforts from Orbital Sciences, who launched their air-launch rocket Pegasus in 1991, further expanded the US launch market. The Russian Soyuz and Proton rockets remained available for commercial launches in the 1980s and 1990s, respectively.

However, costs did not decrease below approximately \$18 000/kg to low-Earth orbit, and launch rates remained low (Figure 1.1). In 2008, SpaceX launched the first privately-funded liquid-fuelled rocket into orbit, with the long-term goal of partial reusability to reduce the cost of access to space. In 2015, SpaceX performed the first landing of a first stage of a rocket. By doing so,

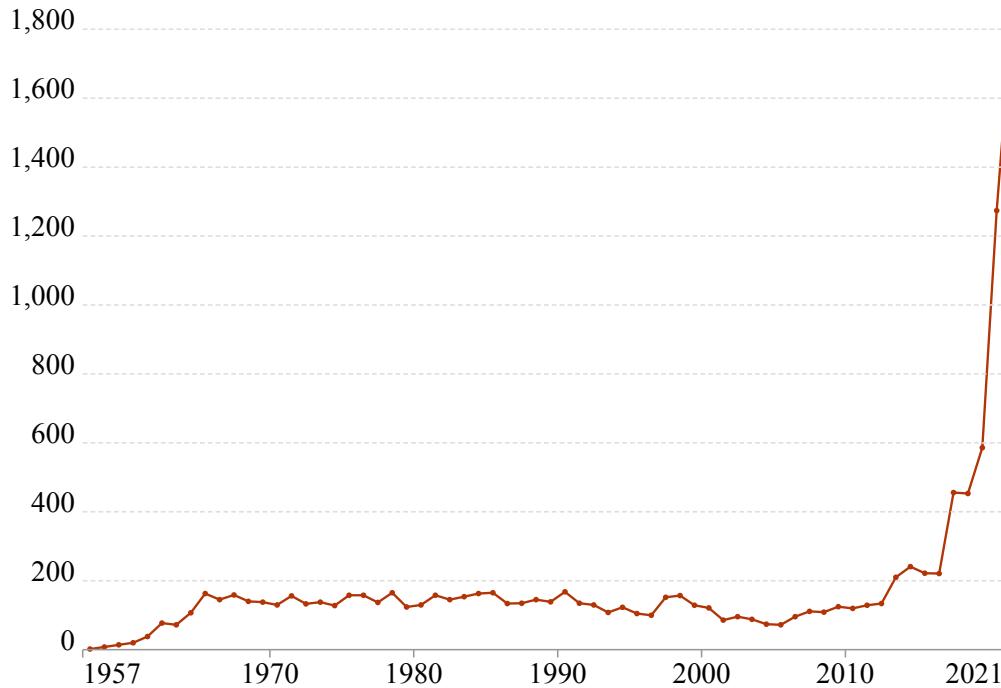


Figure 1.1: Number of satellites deployed per year; recent years have seen a marked increase in space missions, in part due to the reduction in launch costs. Figure adapted from Our World in Data [1].

the first stage can be reused with minor reconditioning, rather than requiring the full manufacture of a replacement. This has dramatically reduced the cost of access to space to below \$2000/ kg. Similar efforts at reusability are currently being emulated by United Launch Alliance, Arianespace, and Rocket Lab, to remain financially viable. The latter is a New Zealand-based space startup developing dedicated launch services for small, low-cost satellites. In parallel, companies such as Virgin Orbit and Blue Origin are also beginning to provide reusable launch vehicles designed for small, low-cost satellites and CubeSats, which now form a key part of the space launch market¹. As a result of this dramatic reduction in cost of access to space, and a growing numbers of operators conducting launches, the number of satellites being deployed per year has increased sharply (Figure 1.1).

However, while getting into orbit is a key part of a successful space mission and certainly worthy of significant research and development to reduce costs, it is only one part of the puzzle and in some cases only 30-40% of the cost [10]. Most spacecraft will have to use on-board propulsion at some point in their lifetime to transfer from the orbit the launch vehicle provided to its nominal operational orbit, and also to maintain the nominal orbit after orbital perturbations. Since space missions end when spacecraft are depleted of their fuel, fuel-efficient orbital transfers can increase spacecraft lifetime for a given amount of propellant, delivering further returns on the initial investment. Alternatively, the operator can reduce the amount of propellant carried on-board at launch, further reducing launch costs.

¹SpaceX's Transporter-1 mission in 2021 deployed 143 satellites from a single, 5-tonne launch.

Designing fuel-efficient trajectories thus plays a key role in modern trajectory design. In preliminary trajectory design, where the focus is on rapid exploration of the design space to find feasible trajectories, mission designers often work in highly-simplified models of motion. The simplest possible is the two-body problem, where the spacecraft moves under the influence of a single gravitational body at any one time. In this formulation, analytical transfers and analytical expressions for motion exist which can be ‘patched’ together and optimised for minimum total change in velocity (Δv) and thus the lowest fuel usage. Once the trajectory is known in the simplified model, it can be ‘corrected’ into higher-fidelity models.

However, in search of more efficient transfers and mission plans, mission designers have begun to use the dynamics of multiple planets and Moons to assist the spacecraft in its transfer, as careful use of the system’s dynamics can reduce overall velocity requirements. This can either be by exploiting the gravity of another body to assist in an orbital transfer or orbital insertion, or alternatively the dynamics of multiple bodies can also assist in designing bespoke orbits that only exist as a result of the delicate gravitational balance between them.

Since this relies on the gravitational balance of at least two planets, the simplest system for which these techniques are valid is the three-body problem, a case of Newton’s n -body problem for $n = 3$. This model of motion has a rich mathematical history, dating back to the works of Kepler, Lagrange and Poincaré. Further assumptions on the orbit of the two planets can yield either an autonomous dynamical system if the planets orbit each other in a circle (the ‘circular’ problem), or a non-autonomous system if they orbit each other in an ellipse (the ‘elliptic’ problem). However, in these systems the equations of motion do not have an analytic solution, so trajectories must be found numerically. Trajectories must also still be corrected into higher-fidelity, more realistic models for final, ‘real-world’ mission plans.

The circular problem has received significant attention in preliminary trajectory design. As an autonomous dynamical system, dynamical features such as equilibria, periodic orbits and invariant manifolds exist which perform a key role in modern space mission design. These can dramatically reduce overall propellant usage (chapter 3) by allowing propellant-free transfers and orbital insertions about the Solar System, using only natural dynamics. These periodic orbits and their characteristics can also provide significant advantages for some missions over more traditional orbits.

Separately, the field of computing – and particularly scientific computing² – has become increasingly prevalent in modern society. Whereas the most powerful computer in the 1950s, the IBM704, could perform 2.5×10^6 floating-point operations per second at maximum output, 2022 saw the first ‘exascale’ supercomputer become operational and capable of $\sim 10^{18}$ floating-point operations per second. This availability of computational resources means that brute-force searches for feasible trajectories can be performed, and significant databases of trajectories that exploit dynamical features can be created and manipulated. This rise in computing has proved

²‘Scientific computing’ derives from ‘scientific computers’, a class of primarily IBM supercomputers that were equipped with the optional extra of a floating-point arithmetic unit in the 1950s and 1960s.

seminal in many fields, and actually motivated many numerical studies of periodic orbits and transfers to-and-from them 1970s and 1980s. However, high-performance computing techniques are not yet widespread in astrodynamics. The first question this motivates is: can modern, high-performance techniques be used to generate additional insight in autonomous astrodynamics systems?

However, even if additional insight is found, there is no guarantee that these trajectories can be directly corrected into higher-fidelity models. Even when they can, the features being exploited are not native to the higher-fidelity model, but rather a continuation of the behaviour of another system. Working directly in the time-dependent models can be nontrivial and numerically challenging, and there is no guarantee on the existence of the invariant manifolds in systems with arbitrary time-dependence, which are dynamical structures heavily used in mission design in autonomous systems.

But not all is lost: finding time-dependent generalisations of the invariant manifolds is an active area of research in fluid dynamics. This motivates a second question: can some of these techniques be used in astrodynamics to find time-dependent generalisations of such structures directly in the time-dependent models? Being able to construct similar trajectories to those in the time-independent models directly in the time-dependent models would mean one could use the underlying dynamics of the system directly, rather than relying on being corrected from lower-fidelity models, respecting the intricacies of the system.

The structure of the thesis now follows. First, in chapter 2, I formally introduce and derive the equations of motion for the circular and elliptic problems, as well as introduce their respective reference frames. Next, in chapter 3, I review many of the dynamical phenomena used in modern mission design in the circular problem, highlighting the role that periodic orbits and stable and unstable invariant manifolds play in the mission design process. The chapter then goes on to examine techniques from the fluid dynamics community, who seek analogous structures to the invariant manifolds in time-aperiodic dynamical systems. As part of this, the Lagrangian Coherent Structure (LCS) is identified as a promising time-dependent generalisation of the stable and unstable manifolds, but which has thus far proved numerically troublesome to calculate in space mission design.

Three novel investigations follow that aim to answer the two key questions of this thesis.

In the first investigation in chapter 4, I use the classical invariant manifolds in combination with high-performance computing to intensively optimise trajectories which transport asteroids into periodic orbits in the circular problem. This chapter represents the most in-depth search for such retrieval trajectories known to the author as of the writing of this thesis. Content from this chapter has been published as Tyler and Wittig [11] and presented at a specialist astrodynamics conference as Tyler and Wittig [12].

Chapter 5 then uses Differential Algebra, a method for approximating functions directly in a computer using high-order polynomials, to improve the numerical performance of finding three-dimensional LCS. This method outperforms standard approaches for common fluid dynamics test cases in the literature while reducing user involvement by computing relevant quantities automatically. Content from this chapter has been published in the Journal of Computational Science as Tyler and Wittig [13] and presented partly at an astrodynamics conference as Tyler and Wittig [14].

I then apply this new algorithm to two examples from an astrodynamics system, where it continues to provide insight, in chapter 6. This allows mission designers to exploit the underlying dynamics of the system without first needing to go via a simplified autonomous problem. The aim of this chapter is to provide the mission design community with the tools to reproduce studies such as those found in chapter 4 in the future. I uncover the role that these structures play in space mission design, and investigate in-depth many of the free parameters available in their construction. This chapter also provides the community with mitigation strategies to avoid common numerical pitfalls. Content from this chapter has been submitted to Celestial Mechanics and Dynamical Astronomy as Tyler and Wittig [15] and presented at the International Astronautical Congress as Tyler and Wittig [16], along with some content presented as Tyler and Wittig [14].

2

The Three-body Problem

This chapter presents a review of the three-body problem, a subset of Newton's n -body problem for $n = 3$. A brief history of the problem is given, and the equations of motion for two formulations of the three-body problem are derived for use in later chapters.

2.1 Introduction

Newton's restricted n -body problem is one of the most general cases of purely gravitational attraction in the solar system. It models the effect of n bodies of mass m_i influencing the motion of a small spacecraft with position vector \mathbf{r} . If the bodies are at a position \mathbf{r}_i from the spacecraft, in an inertial frame of reference the problem can be modelled using the differential equation

$$\ddot{\mathbf{r}} = - \sum_{i=0}^n \frac{Gm_i}{r_i^3} \mathbf{r}_i. \quad (2.1)$$

For large n , in heavily perturbed environments, or when the spacecraft and the planet are very close together, studying this system directly can be numerically difficult, computationally expensive, and time-consuming.

Thus, in the practical design of space missions, trajectory designers often seek to perform preliminary analysis in simplified models of motion [17, 18, 19]. The simplest of these is the two-body problem ($n = 2$), the movement of a mass under the gravitational influence of only one other larger, more massive body that dominates the dynamics. In this system, spacecraft trace out Keplerian motion where orbits exist analytically as conic sections under the solution of simple equations, and analytic methods such as Lambert's problem exist to construct transfers. Mission designers can join segments of these trajectories under the influence of different bodies together – a technique known as patched conics – to achieve good approximations to mission plans in the full model [20].

However, it has been shown that lower- Δv – and thus lower-cost – transfers can be achieved by designing missions in systems for which $n \geq 3$ [20]. The three-body problem, the specialisation for $n = 3$, is a formulation of motion which has a rich theoretical background and the lowest number of bodies for which such interesting dynamical phenomena first start to emerge. Leonhard Euler was the first to propose the three-body problem, and he discovered the first class of periodic orbit in the general three-body problem in his 1767 work [21]: three bodies, placed on a straight line, would stay on this line for certain initial conditions while rotating about their centre of mass. Joseph-Louis Lagrange would later find another class of periodic motion in the three-body problem, where objects located at the vertices of an equilateral triangle may, depending on their initial conditions, retain their original positions relative to each other [22].

Both Euler and Lagrange then made significant contributions to the formulation of the Circular-Restricted Three-body Problem (CR3BP), a specialisation of the three-body problem. Two further assumptions are made: the first is that the mass of the third body is restricted, or considered negligible. This is an assumption with great relevance to spaceflight problems given the relative masses of spacecraft and the objects they orbit. The second assumption is that the remaining two bodies, known as primaries, orbit each other in circles. Euler was the first to use a synodic (rotating) co-ordinate system to describe the motion of the third mass where the distance between the two massive bodies are set fixed on the x -axis with a distance of unity between them, which greatly simplifies the equations of motion.

Lagrange showed that there are five equilibrium points in the CR3BP where the gravitational influence of two bodies ‘cancels’. In 1836, Carl Gustav Jacob Jacobi used Euler’s synodic co-ordinate system to formulate the Circular Problem’s only integral. The *Jacobi integral*, and the related *Jacobi constant* were then used by Hill [23] to introduce zero-velocity surfaces and forbidden regions, where the motion of an object – in Hill’s case an asteroid – is bounded as a result of its energy.

Henri Poincaré in his *Les Méthodes Nouvelles de la Mécanique Céleste*¹ identified and studied periodic orbits in the restricted three-body problem, and defined several new methods for studying motion in the system [24]. This work led to renewed interest in the problem, and was swiftly followed by several publications which examined the existence and stability of several new classes of periodic orbits in the CR3BP (on the existence: [25, 26, 27], on the stability: [28, 29]). More recent work used the advances in computing in the late 20th century to identify and analyse yet more types of periodic orbits, including many which are used in modern mission design (see chapter 4). New families are still being found to this day [30].

This thesis studies two specialisations of the restricted three-body problem, the Circular-Restricted Three-body Problem (CR3BP), which has a key role in modern mission design, and the Elliptic-Restricted Three-body Problem (ER3BP), a natural extension of the CR3BP that introduces a time-dependency in the equations of motion. An in-depth introduction to the dynamical features

¹In English: New Methods in Celestial Mechanics.

and dynamical indicators used to perform mission design and analyse the behaviour of these systems is presented in chapter 3; this chapter serves only to introduce their formulation, derive their equations of motion, and highlight common reference frames used in mission design that will be of use in the remainder of the thesis.

2.2 Derivation of the Equations of Motion

This section presents a derivation of the equations of motion for the Elliptic-Restricted Three-body Problem, a time-dependent case of the three-body problem in which the two large masses orbit each other on ellipses. The CR3BP is then introduced as a special case of the ER3BP.

The derivation begins by considering two-body (Keplerian) motion, since in the three-body problem the two large bodies orbit their common barycentre in two-body motion. Relevant quantities to describe the positions and characteristics of the bodies on their respective orbits, are stated, and this result is then used in defining the equations of motion for both the CR3BP and the ER3BP.

2.2.1 Keplerian Motion

If the only forces acting on an object as it proceeds around its orbit is the gravitation of a central body, then the orbit is a Keplerian orbit. The solution for the motion of the object then follows the equation of a conic section. In the following I give a brief review of Keplerian motion. Since many of the relevant relations in Keplerian motion are standard in the literature (see, for example, Curtis [17], Vallado [31] or Bate et al. [32]) and Keplerian motion does not form a key part of this thesis, only the main results that will be required later are stated.

A Keplerian orbit can be represented using six parameters, a , e , i , Ω , ω , ν (Figure 2.1). The semi-major axis a and eccentricity e determine the size and shape of the ellipse representing the orbit, respectively (Figure 2.2). The inclination of the orbit i is the angle measured between the reference plane and the plane of the orbit. This angle is measured at the ascending node, the intersection of the orbit with the reference plane when the object is travelling ‘upwards’. The angle Ω gives the location of the ascending node with respect to the reference direction, and the angle ω gives the location of the periapsis from the ascending node. The reference direction is typically the vernal equinox, often denoted by Υ , and is the position of the Sun at the vernal point. The true anomaly ν represents the angle of the spacecraft from periapsis. Other anomalies are often used, such as the mean anomaly M and elliptic anomaly E . These hold applications in wider space mission design; for more information on these, see Curtis [17]. Standard orbital element conversion routines exist in Curtis [17], Vallado [31], and Bate et al. [32] that can be used to convert between orbital elements and Cartesian position and velocity.

Only $e < 1$ is considered in this thesis, resulting in the orbit either transcribing an ellipse ($e \neq 0$) or a circle ($e = 0$). For completeness, the case of $e = 1$ transcribes a parabola, and $e > 1$ a hyperbola. In purely Keplerian motion, the first five elements given above are constant, with ν periodic

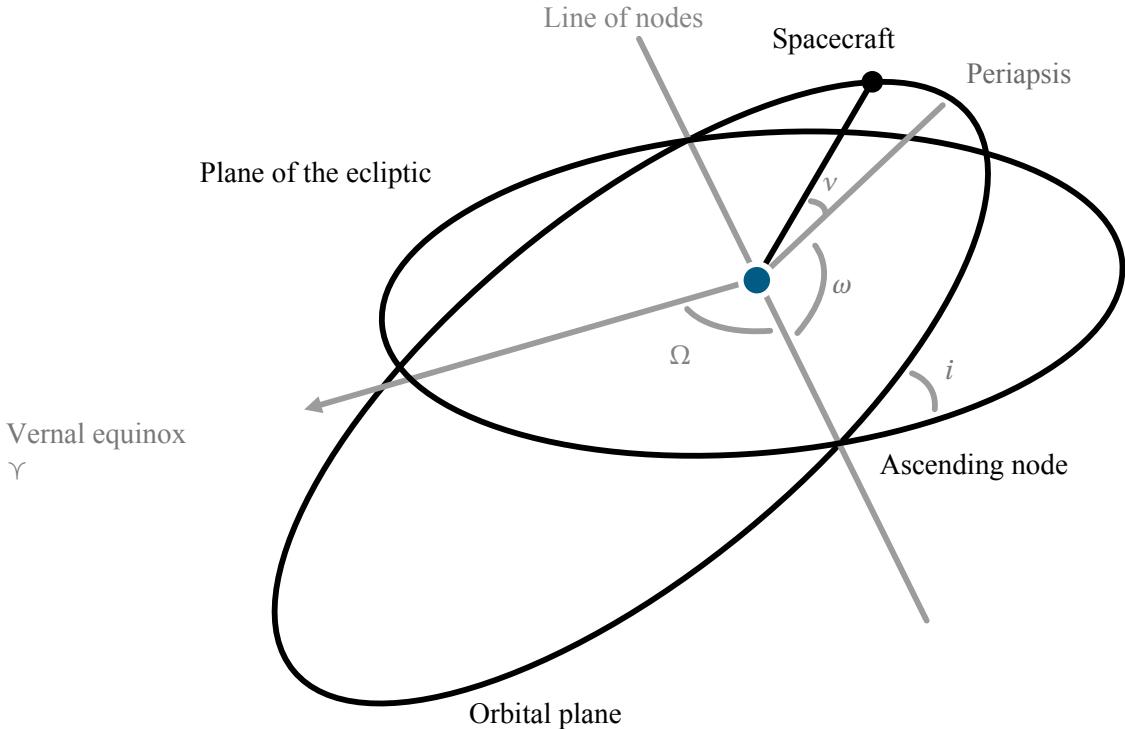


Figure 2.1: Graphical representation of a subset of the Keplerian orbital elements. Not pictured are the semimajor axis a and the eccentricity e .

such that $\nu \in [0, 2\pi]$. One can also obtain osculating Keplerian orbital elements, which are time-varying but correspond to the equivalent Keplerian orbit for a spacecraft with a given position and velocity at any instant. These are useful for describing the current motion of an object even in perturbed environments (see chapter 4).

Some key relations for Keplerian motion are now given. Consider an object of mass m orbiting about a central body with position vector \mathbf{r} and velocity vector $\dot{\mathbf{r}} = \mathbf{v}$. This object has specific angular momentum

$$\mathbf{h} = \mathbf{r} \times \mathbf{v}. \quad (2.2)$$

The direction of the angular momentum vector $\hat{\mathbf{h}}$ is by definition normal to the orbital plane. The magnitude of the angular momentum vector depends only on the component of the velocity perpendicular to the radius, v_{\perp} , and so the scalar angular momentum, h , is

$$h = r v_{\perp} = r^2 \frac{dv}{dt} \quad (2.3)$$

which will be required later in the derivation of the ER3BP equations of motion. The specific angular momentum is constant along an orbit, i.e. $\dot{h} = 0$.

Given that the solution for the motion of the object in Keplerian motion is a conic section, standard ellipse relations can be used to determine the radius of the object at a given true anomaly ν .

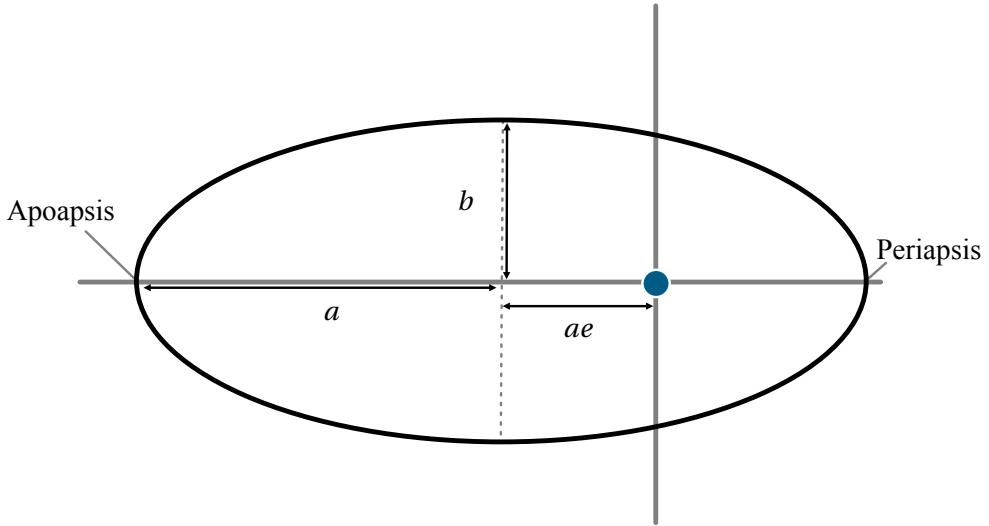


Figure 2.2: The geometry of an ellipse, as used in the solution of the orbit equation for $0 < e < 1$. a is the semi-major axis of the orbit, and e the eccentricity. The periapsis is the closest point on the orbit to the central body, and the apoapsis the farthest.

The orbit equation is defined as

$$r(\nu) = \frac{h^2}{GM} \frac{1}{1 + e \cos \nu}. \quad (2.4)$$

An alternative formulation of Equation 2.4 is

$$r = a \frac{1 - e^2}{1 + e \cos \nu} \quad (2.5)$$

where $p = h^2/GM = a(1 - e^2)$ is the semi-latus rectum of the orbit. Equation 2.4 has a minimum at $\nu = 0$ and maximum at $\nu = 180^\circ$. The minimum distance to the central body is known as periapsis, and the maximum distance is known as apoapsis.

2.2.2 Derivation of the ER3BP Equations of Motion

The orbit equation and relations concerning the conservation of angular momentum are used as part of the derivation of the equations of motion for both the Elliptic-Restricted Three-body Problem and the Circular-Restricted Three-body Problem. The derivation begins by considering two masses orbiting about their centre of mass under Keplerian motion, denoted m_1 and m_2 , about which the motion of a third mass, m_3 , is investigated. m_1 and m_2 are also known as the primaries of the problem, and it is enforced that $m_1 \geq m_2 \gg m_3$. The mass parameter $\mu = m_2 / (m_1 + m_2)$ parameterises the system, and m_2 would orbit m_1 on an ellipse of eccentricity e_p .

The equations of motion for the ER3BP are defined in a rotating coordinate system defined at the barycentre of m_1 and m_2 , where the $+x$ direction points from m_1 to m_2 and where the location of m_1 and m_2 are fixed, as in Figure 2.3. The z -direction points normal to the orbital plane of m_1 and m_2 , and the y -axis completes the right-handed orthonormal triad. The subscript R is used

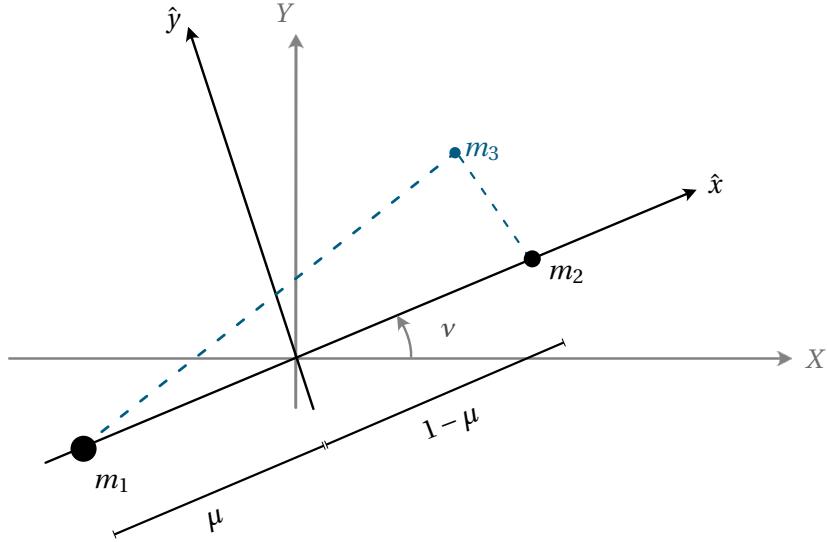


Figure 2.3: The coordinate systems considered in the derivation of the equations of motion of the CR3BP. In grey is a fixed (inertial) system with its origin at the barycentre of m_1 and m_2 ; in black is the rotating (synodic) frame with m_1 and m_2 set fixed.

to represent quantities in the rotating frame and the subscript I is used to represent quantities in the inertial frame. Given an angular velocity $\dot{\nu}$ of the rotating frame about the z -axis, such that $d\nu/dt = \dot{\nu}\hat{e}_z$ and $\hat{e}_z = (0, 0, 1)$ points along \hat{z} , the first derivative of the position vector \mathbf{r} can be transformed using

$$\frac{d\mathbf{r}}{dt}_I = \frac{d\mathbf{r}}{dt}_R + \frac{d\nu}{dt} \times \mathbf{r}. \quad (2.6)$$

Differentiating Equation 2.6 yields

$$\frac{d^2\mathbf{r}}{dt^2}_I = \frac{d}{dt} \left(\frac{d\mathbf{r}}{dt}_R \right) + \frac{d^2\nu}{dt^2} \times \mathbf{r} + \frac{d\nu}{dt} \times \frac{d\mathbf{r}}{dt}_I. \quad (2.7)$$

Combining Equations 2.6 and 2.7 with the equations of motion for the n -body problem about the barycentre of m_1 and m_2

$$\frac{d^2\mathbf{r}}{dt^2}_I = -\frac{Gm_1\mathbf{r}_1}{r_1^3} - \frac{Gm_2\mathbf{r}_2}{r_2^3} \quad (2.8)$$

gives

$$\frac{d^2\mathbf{r}}{dt^2}_R = -G \left(\frac{m_1}{r_1^3} \mathbf{r}_1 + \frac{m_2}{r_2^3} \mathbf{r}_2 \right) - 2 \frac{d\nu}{dt} \times \frac{d\mathbf{r}}{dt}_R - \frac{d\nu}{dt} \times \frac{d\nu}{dt} \times \mathbf{r} - \frac{d^2\nu}{dt^2} \times \mathbf{r}. \quad (2.9)$$

The final three terms on the right-hand side of Equation 2.9 are the Coriolis acceleration, centrifugal acceleration, and acceleration due to the angular acceleration due to the fact that the rotating frame rotates at a variable rate (Euler force). Since the remaining analysis is only manipulations of these terms, the subscripts are dropped for clarity.

It is numerically beneficial to normalise distance, time and mass in Equation 2.9. The normalising variables are given by

$$\bar{r}(\nu) = \frac{a(1 - e_p^2)}{1 + e_p \cos \nu} \quad (2.10)$$

$$\bar{t} = \left(\frac{dv}{dt} \right)^{-1} \quad (2.11)$$

$$\bar{m} = m_1 + m_2 \quad (2.12)$$

such that

$$\mathbf{r} = \bar{r} \mathbf{r}^* \quad (2.13)$$

$$t = \bar{t} \tau \quad (2.14)$$

$$m = \bar{m} m^*. \quad (2.15)$$

The variables a , e_p and ν are the semi-major axis, eccentricity and true anomaly of the orbit of m_2 about m_1 , since they orbit each other in two-body motion. Since the time and radius are not constant, the ER3BP is studied in a rotating coordinate system which is non-uniformly rotating and pulsating. I note that the normalised time τ progresses at the same rate as the true anomaly.

Given Equation 2.11 it is possible to relate the normalised time to the dimensional time via

$$\frac{d}{dt} = \frac{d\tau}{dt} \frac{d}{d\tau} = \frac{dv}{dt} \frac{d}{d\tau} \quad (2.16)$$

$$\frac{d^2}{dt^2} = \frac{d^2v}{dt^2} \frac{d}{d\tau} + \left(\frac{dv}{dt} \right)^2 \frac{d^2}{d\tau^2}. \quad (2.17)$$

Using Equations 2.13-2.15 and 2.16-2.17, each of the terms in Equation 2.9 can be rewritten:

$$\frac{d\mathbf{r}}{dt} = \frac{d}{dt} \left(\bar{r} \frac{d\mathbf{r}^*}{dt} + \frac{d\bar{r}}{dt} \mathbf{r}^* \right) = \bar{r} \frac{d^2\mathbf{r}^*}{dt^2} + 2 \frac{d\bar{r}}{dt} \frac{d\mathbf{r}^*}{dt} + \frac{d^2\bar{r}}{dt^2} \mathbf{r}^* \quad (2.18)$$

$$= \bar{r} \left(\frac{d^2v}{dt^2} \frac{d\mathbf{r}^*}{d\tau} + \left(\frac{dv}{dt} \right)^2 \frac{d^2\mathbf{r}^*}{d\tau^2} \right) + 2 \frac{d\bar{r}}{dt} \frac{dv}{dt} \frac{d\mathbf{r}^*}{d\tau} + \frac{d^2\bar{r}}{dt^2} \mathbf{r}^* \quad (2.19)$$

$$2 \frac{d\nu}{dt} \times \frac{d\mathbf{r}}{dt} = 2 \frac{dv}{dt} \hat{\mathbf{e}}_z \times \left(\bar{r} \frac{d\mathbf{r}^*}{dt} + \frac{d\bar{r}}{dt} \mathbf{r}^* \right) \quad (2.20)$$

$$= 2 \frac{dv}{dt} \hat{\mathbf{e}}_z \times \left(\bar{r} \frac{dv}{dt} \frac{d\mathbf{r}^*}{d\tau} + \frac{d\bar{r}}{dt} \mathbf{r}^* \right) \quad (2.21)$$

$$\frac{d\nu}{dt} \times \frac{d\nu}{dt} \times \mathbf{r} = \bar{r} \left(\frac{dv}{dt} \right)^2 \hat{\mathbf{e}}_z \times \hat{\mathbf{e}}_z \times \mathbf{r}^* \quad (2.22)$$

$$\frac{d^2\nu}{dt^2} \times \mathbf{r} = \bar{r} \frac{dv}{dt} \hat{\mathbf{e}}_z \times \mathbf{r}^* \quad (2.23)$$

$$G \left(\frac{m_1}{r_1^3} \mathbf{r}_1 + \frac{m_2}{r_2^3} \mathbf{r}_2 \right) = \frac{G(m_1 + m_2)}{\bar{r}^2} \left(\frac{1 - \mu}{r_1^{*3}} \mathbf{r}_1^* + \frac{\mu}{r_2^{*3}} \mathbf{r}_2^* \right). \quad (2.24)$$

Substituting the above into Equation 2.9 and collecting terms yields

$$\left(\bar{r} \frac{d^2 v}{dt^2} + 2 \frac{d\bar{r}}{dt} \frac{dv}{dt} \right) \frac{dr^*}{d\tau} + \left(\bar{r} \frac{d^2 v}{dt^2} + 2 \frac{d\bar{r}}{dt} \frac{dv}{dt} \right) \hat{e}_z \times r^* + \bar{r} \left(\frac{dv}{dt} \right)^2 \frac{d^2 r^*}{d\tau^2} + 2\bar{r} \left(\frac{dv}{dt} \right)^2 \hat{e}_z \times \frac{dr^*}{d\tau} \quad (2.25)$$

$$= -G \frac{m_1 + m_2}{\bar{r}^2} \left(\frac{1 - \mu}{r_1^{*3}} \bar{r}_1^* + \frac{\mu}{r_2^{*3}} r_2^* \right) - \frac{d^2 \bar{r}}{dt^2} r^* - \bar{r} \left(\frac{dv}{dt} \right)^2 \hat{e}_z \times \hat{e}_z \times r^*. \quad (2.26)$$

Using some of the relations derived in the two-body problem, and recalling that the mass of the spacecraft is restricted or negligible, the above can be simplified further. As per Szebehely and Jefferys [33], first recall the conservation of angular momentum:

$$\frac{dh}{dt} = \frac{d}{dt} \left(\bar{r}^2 \frac{dv}{dt} \right) = \bar{r} \frac{d^2 v}{dt^2} + 2 \frac{d\bar{r}}{dt} \frac{dv}{dt} = 0 \quad (2.27)$$

and then the semi-latus rectum p and angular momentum:

$$h^2 = \left(\bar{r}^2 \frac{dv}{dt} \right)^2 = p G (m_1 + m_2) \Rightarrow \left(\bar{r}^2 \frac{dv}{dt} \right)^2 = a (1 - e^2) G (m_1 + m_2). \quad (2.28)$$

and its equation of motion

$$\frac{d^2 \bar{r}}{dt^2} - \bar{r} \left(\frac{dv}{dt} \right)^2 = -\frac{G (m_1 + m_2)}{\bar{r}^2} = -\frac{1}{a (1 - e_p^2)} \bar{r}^2 \left(\frac{dv}{dt} \right)^2. \quad (2.29)$$

Substituting these three relations into Equation 2.25 and fixing the location of m_1 on the x -axis at $x_1 = x - \mu$ and m_2 at $x_2 = x + 1 - \mu$ by convention [34], one finds that the equations of motion for the ER3BP can be written

$$x'' - 2y' = \frac{1}{1 + e_p \cos v} \left(x - \frac{1 - \mu}{r_1^3} (x + \mu) - \frac{\mu}{r_2^3} (x - 1 + \mu) \right) \quad (2.30)$$

$$y'' + 2x' = \frac{1}{1 + e_p \cos v} \left(y - \frac{1 - \mu}{r_1^3} y - \frac{\mu}{r_2^3} y \right) \quad (2.31)$$

$$z'' + z = \frac{1}{1 + e_p \cos v} \left(-\frac{1 - \mu}{r_1^3} z - \frac{\mu}{r_2^3} z \right) \quad (2.32)$$

where \square' denotes derivatives with respect to true anomaly v . Through the introduction of a pseudopotential function Ω , the equations of motion can be rewritten in a more elegant form.

$$\Omega = \frac{1}{1 + e_p \cos v} \left[\frac{1}{2} (x^2 + y^2 - z^2 e_p \cos v) + \frac{1 - \mu}{r_1} + \frac{\mu}{r_2} \right] \quad (2.33)$$

such that

$$x'' - 2y' = \frac{\partial \Omega}{\partial x} \quad (2.34)$$

$$y'' + 2x' = \frac{\partial \Omega}{\partial y} \quad (2.35)$$

$$z'' = \frac{\partial \Omega}{\partial z}. \quad (2.36)$$

The ER3BP is a time-periodic dynamical system; its uses in astrodynamics and attempts to profile its dynamics are presented in chapter 3.

2.2.3 The Circular-Restricted Three-body Problem

If one assumes that the two primaries given in the derivation above orbit each other in a circle rather than an ellipse (i.e. $e_p = 0$), a further specialisation of the n -body problem – the Circular-Restricted Three-body Problem (CR3BP) – is obtained. The CR3BP is a well-studied system for preliminary mission design, since it is the simplest system for which dynamical phenomena begin to exist.

By setting $e_p = 0$ in the derivation above, the motion of m_2 about m_1 is no longer dependent on true anomaly and thus the coordinate system no longer pulsates. Moreover, the equations of motion become autonomous. The equations of motion have the form

$$\ddot{x} - 2\dot{y} = \frac{\partial \Omega_C}{\partial x} \quad (2.37)$$

$$\ddot{y} + 2\dot{x} = \frac{\partial \Omega_C}{\partial y} \quad (2.38)$$

$$\ddot{z} = \frac{\partial \Omega_C}{\partial z}. \quad (2.39)$$

where

$$\Omega_C = \left[\frac{1}{2} (x^2 + y^2) + \frac{1-\mu}{r_1} + \frac{\mu}{r_2} \right] \quad (2.40)$$

and \square signifies derivatives with respect to (non-dimensional) time. There is a wealth of literature available on the CR3BP, which forms a large portion of the preliminary mission design process for modern space missions. Its dynamical phenomena are also well-understood; I provide a review in chapter 3.

The expression for Jacobi's integral, the CR3BP's single integral of motion derived by Carl Gustav Jacob Jacobi, can be derived by multiplying the equation of motion by \dot{x} , \dot{y} and \dot{z} , respectively

$$\dot{x}\ddot{x} + \dot{y}\ddot{y} + \dot{z}\ddot{z} = \dot{x}\frac{\partial \Omega_C}{\partial x} + \dot{y}\frac{\partial \Omega_C}{\partial y} + \dot{z}\frac{\partial \Omega_C}{\partial z}. \quad (2.41)$$

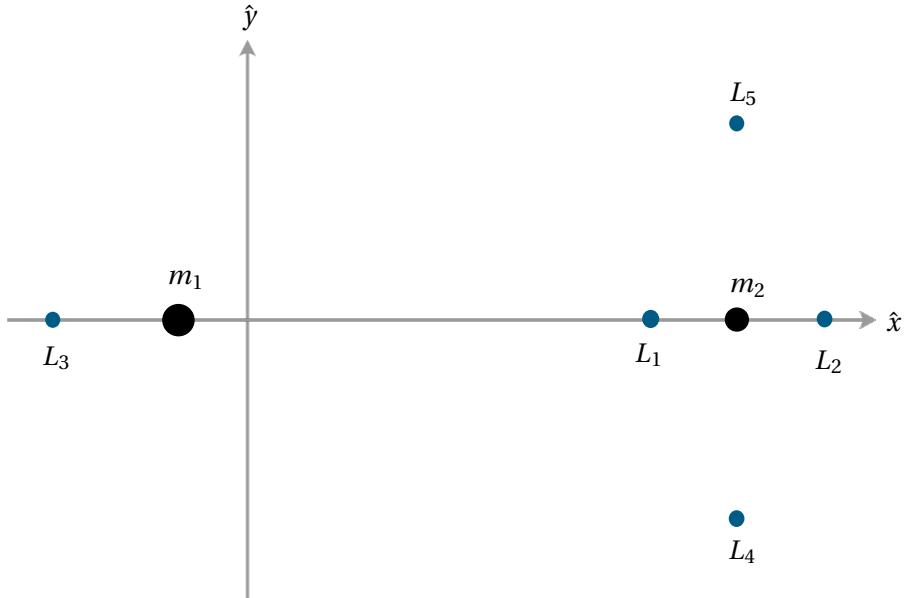


Figure 2.4: The location of the Lagrangian (equilibrium) points in the CR3BP. L_1 through L_3 are known as the collinear points, and L_4 and L_5 the equilateral solutions.

Integrating the equation above gives

$$\dot{x}^2 + \dot{y}^2 + \dot{z}^2 = 2\Omega_C - \mathcal{J} \quad (2.42)$$

where \mathcal{J} is Jacobi's constant, analogous to energy. This constant has a range of uses, such as defining regions where particle can and cannot enter (see, for example, Koon et al. [35]), and in reducing the degrees of freedom of the system to allow studying dynamics on a given section (Short, Howell, and Tricoche [36]). The Jacobi constant is used extensively in chapter 4 to identify asteroids which may be retrievable in the CR3BP.

If one searches for solutions to

$$\frac{\partial\Omega_C}{\partial x} = \frac{\partial\Omega_C}{\partial y} = \frac{\partial\Omega_C}{\partial z} = 0 \quad (2.43)$$

one will find five equilibria of the CR3BP, known as the Lagrange points (Figure 2.4). Numbered L_1 through L_5 , three of these lie on the line joining the primaries and two lie on the edges of an equilateral triangle formed with the m_1 - m_2 barycentre. The equilateral points are unstable for ratios of mass of the smaller body to the larger less than 24.96, and have generally been less well-studied in the literature and for practical uses. Overall, the Lagrangian points have been the destination of choice for no fewer than 17 space missions as of the writing of this thesis.

The dynamics about the collinear points are of type saddle \times saddle \times centre, and each of the equilibria admit periodic orbits about them. These orbits have long been classified into families and are of great importance in modern mission design (chapter 3). Invariant manifolds also emanate from many of these orbits, and these structures have been well-investigated in space

mission design and form a core part of chapter 4. I give a review of these techniques in chapter 3.

The CR3BP readily admits variational equations which can be used to obtain linearised derivatives of the final condition with respect to the initial condition. They can be found by solving a differential equation of the form

$$\dot{\Phi}(t_0, t) = A(t) \Phi(t_0, t) \quad (2.44)$$

where $\Phi(t_0, t)$ is the state transition matrix evaluated from an initial time t_0 to a final time t and $A(t)$ the Jacobian of the CR3BP equations of motion. From the above, two properties can be defined. Namely,

$$\Phi(t_0, t_f) = \Phi(t, t_f) \Phi(t_0, t) \quad (2.45)$$

$$\Phi(0, 0) = I_{6 \times 6} \quad (2.46)$$

where $I_{6 \times 6}$ is the 6×6 identity matrix. The Jacobian takes the form

$$A(t) = \left[\begin{array}{c|c} 0_{3 \times 3} & I_{3 \times 3} \\ \hline \nabla^2 \Omega_C & U \end{array} \right] \quad (2.47)$$

where $I_{3 \times 3}$ is a 3×3 identity submatrix, $0_{3 \times 3}$ a 3×3 submatrix of zeros, and $\nabla^2 \Omega$ and U are 3×3 submatrices of partial derivatives of the pseudo-potential and

$$U = \begin{bmatrix} 0 & 2 & 0 \\ -2 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}. \quad (2.48)$$

By augmenting the initial state vector in the numerical integrator with the additional differential equations that arise from the above ODE for $\dot{\Phi}$, the linear derivatives of a trajectory can be computed; for a system with a phase space dimension of n , an additional n^2 equations are required.

2.3 Transforming between inertial and rotating reference frames

As mentioned in their respective derivations, the equations of motion for the CR3BP and the ER3BP exist in a rotating (synodic) coordinate system. However, it can often be more convenient to investigate system behaviour in the inertial frame. Converting between the inertial frame and the rotating frame is possible using standard Euler rotation matrices. Only the conversion from inertial frames centred on m_2 is considered in the below, but the process can be adjusted to begin about an inertial frame centred on m_1 through replacement of the relevant angles with those of m_1 , and the replacement of the coefficient $1 - \mu$ with μ .

Considering the geometry of the reference frames presented in Figure 2.3, and beginning with some position $\mathbf{x}_I = (x, y, z)$ and velocity $\mathbf{v}_I = \dot{\mathbf{x}}_I = (\dot{x}, \dot{y}, \dot{z})$ in the inertial frame about m_2 , the transformation has three steps:

1. Translate the initial position \mathbf{x} from being about m_2 to being about the $m_1 - m_2$ barycentre (BC).
2. Rotate \mathbf{x} clockwise by an angle ν .
3. Normalise the dimensional position of \mathbf{x} to match the normalised units of the rotating frame.

The transformation is a composite translation, rotation and scaling:

$$\mathbf{x}_{ER3BP} = \frac{R_z(\nu)}{r(\nu)} (\mathbf{x} + \mathbf{x}_{m_2}) \quad (2.49)$$

where from inspection of Figure 2.3

$$\mathbf{x}_{m_2} = r(\nu) \begin{bmatrix} \cos \nu \\ \sin \nu \\ 0 \end{bmatrix} (1 - \mu) \quad (2.50)$$

where $(1 - \mu)$ is the portion of the distance between m_1 and m_2 that lies between the barycentre and m_2 . Thus,

$$\mathbf{x}_{ER3BP} = \underbrace{\frac{R_z(\nu)}{r(\nu)}}_{\text{rotation and scaling}} \left\{ \underbrace{\mathbf{x} + r(\nu) \begin{bmatrix} \cos \nu \\ \sin \nu \\ 0 \end{bmatrix} (1 - \mu)}_{\text{translation from } m_2 \text{ to BC}} \right\} \quad (2.51)$$

To simplify, I cancel the trivial $\frac{r}{r}$ term in the translation and perform the rotation of ν to the position vector of m_2 , such that

$$\mathbf{x}_{ER3BP} = \frac{R_z(\nu) \mathbf{x}}{r(\nu)} + \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix} (1 - \mu). \quad (2.52)$$

To transform the velocity \mathbf{v} into the rotating-pulsating frame, I take the derivative of Equation 2.52. Importantly, the velocity is defined with respect to time, but in the synodic frame the derivatives are with respect to true anomaly ν . Via the chain rule the derivatives of Equation 2.52 are

$$\mathbf{x}'_{ER3BP} = \frac{R'_z \mathbf{x}}{r} + \underbrace{\frac{R_z \mathbf{x}'}{r}}_{(a)} + \frac{R_z \mathbf{x}}{r'} \quad (2.53)$$

where the dependence on v has been dropped for r and R_z . Using the relationship between time and true anomaly, term (a) can be represented as

$$\frac{R_z \mathbf{x}'}{r} = \frac{R_z \dot{\mathbf{x}}}{r \dot{v}} = \frac{R_z \mathbf{v}}{r \dot{v}}. \quad (2.54)$$

and thus the transformation of the velocity is given by

$$\mathbf{x}'_{ER3BP} = \frac{R'_z \mathbf{x}}{r} + \frac{R_z \mathbf{v}}{\dot{v} r} + \frac{R_z \mathbf{x}}{r'}. \quad (2.55)$$

The rate of change of v with dimensional time is given with respect to the angular momentum of the system in combination with the orbit equation

$$\dot{v} = \frac{GM^{\frac{1}{2}} (1 + e)^{\frac{1}{2}}}{r^{\frac{3}{2}}} = \frac{GM^{\frac{1}{2}} (1 + e_p \cos v)^2}{a^{\frac{3}{2}} (1 - e_p^2)^{\frac{3}{2}}} \quad (2.56)$$

and also

$$\frac{d}{dv} \frac{1}{r(v)} = \frac{-e_p \sin v}{a(1 - e_p^2)}. \quad (2.57)$$

For $e_p = 0$ (i.e. the CR3BP), one obtains

$$\dot{v} = \sqrt{\frac{GM}{a^3}} \quad (2.58)$$

and

$$\frac{d}{dv} \frac{1}{r(v)} = 0 \quad (2.59)$$

which yields a transformation of

$$\mathbf{x}_{CR3BP} = \frac{R'_z \mathbf{x}}{r} + \frac{R_z \mathbf{v}}{r} \sqrt{\frac{a^3}{GM}}. \quad (2.60)$$

2.4 Conclusion

This chapter has introduced the circular- and elliptic-restricted three-body problems, special cases of the n -body problem for $n = 3$ and which have received particular interest in space mission design. The equations of motion have also been derived, and these systems are used in later chapters as the underlying dynamical systems for various investigations.

3

Dynamical Features and Dynamical Indicators in Astrodynamics

This chapter presents a two-part overview of analysis methods used in the literature for the Circular- and Elliptic-Restricted Three-body Problem. The first part introduces the structures predicted by classical dynamical systems theory that have been used in the analysis of the time-independent Circular problem and highlights their use in modern mission design in preparation for their use in chapter 4. The second part highlights techniques used by the fluid dynamics community to find similar structures to those used in the time-independent system in systems with arbitrary time-dependence, which forms the key focus of chapters 5 and 6.

3.1 Introduction

In the analysis of dynamical systems, particularly in the field of space mission design, one often wishes to answer several questions. Firstly, one often wishes to profile the dynamics by identifying where trajectories are similar and where trajectories have qualitatively different outcomes. If the regions where trajectories are all similar can be identified, one can begin to identify possible sets of orbits that satisfy mission constraints. Next, it may be beneficial to learn what drives the behaviour of the system and its trajectories in the short- and long-term.

In time-independent systems, the above role is performed by the stable and unstable manifolds of the system, which act to separate phase space and govern nearby dynamics. This chapter presents an overview of the use of the structures predicted by classical dynamical systems theory in time-independent systems in astrodynamics: the role that periodic orbits and stable and unstable manifolds play in modern mission design and practical, real-world missions is elaborated.

In hyperbolic¹ time-periodic systems these structures become more numerically difficult to find, as one now seeks surfaces invariant in the extended phase space of position and time. Often, the invariant manifolds are found by finding fixed points on a Poincaré map over the system period. In time-aperiodic systems, there is no guarantee of the existence of such structures as flow features often do not exist over the lifetime of the flow, or the long-term behaviour of the system is undefined. However, analysing time-dependent dynamical systems to examine what is governing, promoting or inhibiting transport is common in the study of both experimental and theoretical fluid dynamics. Some of these techniques have become widespread in the field of fluid dynamics, and some have begun to be applied to astrodynamics problems to attempt to find analogous structures to the manifolds used widely in space mission design directly in the time-dependent models of motion.

The aim of this chapter is twofold: first, it seeks to motivate the reader that phenomena predicted by dynamical systems theory hold relevance and provide advantages in modern space mission design. It then seeks to provide an overview of dynamical indicators from fluid dynamics that have been used to attempt to find generalised stable and unstable manifolds in astrodynamics systems. In both cases, I do not attempt to highlight all studies known to the author. Rather, I try to give an overview of the key studies that drove the development of such ideas forward, and sufficient information to provide a solid theoretical foundation for the coming chapters and to place this thesis in the context of the wider literature. More specific information required for each investigation is included at the beginning of each chapter that follows.

3.2 Dynamical features in autonomous systems

It was seen in the previous chapter that the Circular-Restricted Three-body Problem is an autonomous dynamical system which has five equilibria, the Lagrangian points. As predicted by dynamical systems theory, these equilibria admit periodic orbits and associated stable and unstable manifolds about them. These play a significant role in modern mission design, which is elaborated in the coming section.

I begin with a survey of the role of the periodic orbits. Initial mathematical analyses of the periodic orbits of the CR3BP were motivated by Poincaré's comprehensive study of the problem in the late 19th and early 20th Century. In 1897, Darwin [25] performed a comprehensive, manual search for periodic orbits and evaluated their stability. In 1912, Moulton [27] continued Darwin's work by determining periodic orbits in the case where $\mu \rightarrow 0$, with potential applications to triple star systems. Lyapunov [29] performed further analysis of the periodic orbits after predicting their existence².

¹Hyperbolic systems are characterised by having everywhere expanding or contracting directions. Given that Hamiltonian systems preserve area by Liouville's theorem, one can generally regard Hamiltonian systems – and thus the CR3BP and ER3BP – as hyperbolic systems.

²The vertical and planar Lyapunov orbits are a set of in- and out-of-plane orbits that are named after Aleksandr Lyapunov.

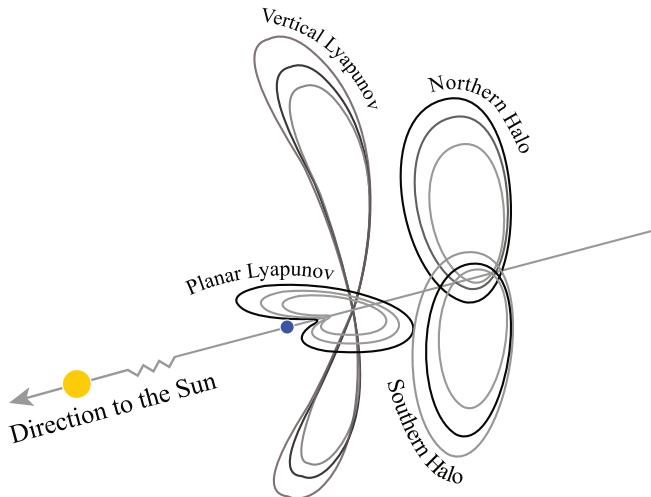


Figure 3.1: Example periodic orbits that exist about Sun-Earth L_1 and L_2 . Pictured are the vertical and planar Lyapunov family, and the quasiperiodic ‘Halo’ orbits. These families are used as destinations for asteroids in chapter 4.

However, it was not until the rise of computing in the latter half of the 20th century that periodic orbits began to be thoroughly investigated numerically. Hénon [37, 38, 39, 40] and Michalodimitrakis [41] performed in-depth numerical investigations into the construction and the stability of planar and vertical families of orbits about the collinear Lagrange points in the late 1960s and early 1970s; some example periodic orbits are provided in Figure 3.1.

In 1968, the doctoral thesis of Robert Farquhar studied the use of ‘Halo’ orbits about the Earth-Moon L_2 , to act as destinations for communication gateways for the Apollo missions when they went ‘behind’ the Moon where radio communication is not possible. The amplitude (the magnitude of the out-of-plane motion) for such an orbit is sufficiently large it can maintain direct line-of-sight with both the Apollo spacecraft and the Earth. In 1973, Farquhar computed a high-order approximation to quasi-periodic motion near the Earth-Moon L_1 and L_2 points, and went on to send the ISEE-3 spacecraft to a ‘Halo’ orbit about the Sun-Earth L_1 point in 1978 to study the relationship between the solar wind and Earth’s magnetosphere [42]³. By placing the probe at Sun-Earth L_1 , ISEE-3 was offered a view of the solar wind that would be permanently ‘in front’ of the Earth and was thus outside of the magnetosphere’s bow shock. As part of this mission, Richardson [43] developed a high-order approximation for periodic motion about all the collinear Lagrangian points which has since served to generate initial guesses for designing practical trajectories about these points. Howell [44] used this approximation to perform a thorough numerical investigation of the ‘Halo’ orbits, which feature heavily in modern trajectory plans. A more generic method to construct arbitrary quasiperiodic orbits was presented by Howell in 1987 [45].

³This spacecraft was later renamed ICE and sent to study the tail of comet Giacobini-Zinner. The spacecraft was designed to return to Earth in 2014, at which point a citizen science initiative (‘ISEE-3 reboot project’) attempted to regain contact with it to place it back into a Halo orbit. The author saved £50 from his weekly pocket money to back this project on Kickstarter in 2014.

Using orbits about the Lagrange points as destinations for operational spacecraft can offer significant advantages. For example, Sun-Earth L_1 points offer uninterrupted views of the Sun and the solar wind, exploited by missions such as SOHO (Solar and Heliospheric Observatory) [46], ACE (Advanced Composition Explorer) and WIND. The Sun-Earth L_2 point has the Sun permanently blocked by the Earth, which makes it an excellent choice for observatories (e.g. DSCOVR and the James Webb Space Telescope [47].) Moreover, some of the orbits about the collinear Lagrange points have excellent stability properties, requiring minimal amounts of propellant to keep the spacecraft in the desired orbit. This was one of the key drivers behind NASA selecting a Near-Rectilinear Halo Orbit, a highly-stable quasiperiodic orbit about Earth-Moon L_2 . The stability of these orbits means that SOHO has enough fuel on board to remain in service until at least 2025, ACE at least 2024, and WIND until 2070, providing excellent mission lifetimes as a result of leveraging structures predicted by dynamical systems theory.

One can also transfer between equilibrium points. Launched in 1994 to the Sun-Earth L_1 point, the WIND mission performed a transfer from its nominal home at L_1 to the vicinity of L_2 in 2003-2004 to further study variations in solar wind before being returned to L_1 . NASA's GENESIS mission performed a heteroclinic transfer – one that begins and ends at two different equilibrium points – in 2003 to return to Earth from L_1 via L_2 , a type of transfer first proposed by Canalias and Masdemont [48] and Davis et al. [49]. Typically, these transfers exist along the intersection of the stable and unstable manifolds of orbits which exist about each equilibrium point.

Stable and unstable manifolds are formed of the infinite set of trajectories which asymptotically approach or depart an orbit as $t \rightarrow \infty$. Found by studying the eigenvalues and eigenvectors of the state transition matrix evaluated from a point on a given periodic orbit for one full orbital period, they play a key role in governing nearby dynamics by acting as separatrices in phase space and allowing propellant-free transfers to-and-from periodic orbits [50]. Conley [51] was the first to study structures around the Lagrange points. He classified trajectories near the Lagrange points as ones which cross or do not cross the region, or which asymptotically approach the periodic orbit. He also constructed a practical mission plan to show how one could transfer a spacecraft from the Earth to the Moon using these orbits for lower cost than traditional methods of directly injecting into a trans-Lunar orbit. To the author's knowledge, this was the first 'low-energy transfer' mission plan, a term given to a broad class of trajectories which consume far less propellant than an otherwise direct transfer using traditional methods [52].

The first approach to explicitly use the invariant manifolds to compute a transfer from the Earth to a Halo orbit was Gómez et al. [53]. In their paper, Gómez et al. integrate the structure of a stable manifold associated with a Halo orbit about Sun-Earth L_1 backward in time until points on the manifold fulfil a criteria relating to its minimum distance to the Earth. These points then form targets in an optimisation procedure from an orbit about the Earth onto the stable manifold, and from there asymptotically about the Halo orbit using only the dynamics of the system.

This work was followed by Howell, Barden, and Lo [54] in 1997 who extended this notion to

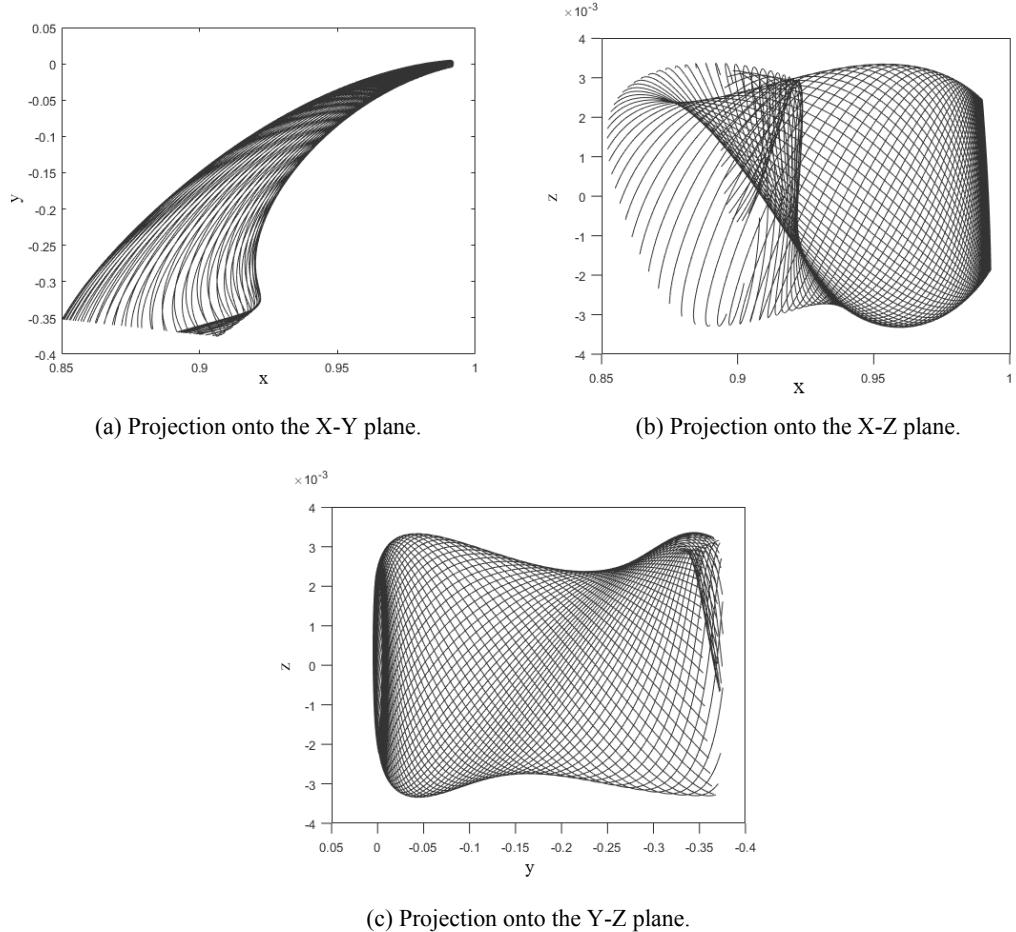


Figure 3.2: Stable manifold associated to a Halo orbit about Sun-Earth L_1 , represented as its projection onto spatial coordinates.

Earth-to-Halo and Halo-to-Earth transfers at Sun-Earth L_2 . This paper extended the previous work by also including a return trajectory along the unstable manifold of the Halo about L_2 , as well as by demonstrating its application to the preliminary trajectory design of Suess-Urey, which later became the GENESIS mission. The full trajectory for the GENESIS mission remained qualitatively similar, and in Howell, Barden, and Wilson [55] the designers remarked that the trajectory provided ‘unique’ opportunities for recovering the solar wind particles about Sun-Earth L_1 , while allowing a great deal of insight through the trajectory design process. Similar methods were employed for the ISEE-3 and NEAR missions in Farquhar, Muñonen, and Richardson [56] and Dunham, McAdams, and Farquhar [57].

Gómez et al. [50] used the invariant manifolds to determine a ‘Petit Grand Tour’ of the Jovian Moons. A series of three-body sub-problems in the Jupiter-Europa-Ganymede-spacecraft system were found to well approximate the overall dynamics and the resulting transfer consumed much less propellant than the equivalent Hohmann transfer, and with more control over the ‘three-dimensional details’ of the trajectory compared to standard approaches.

The ballistic capture mechanism, which emerged in 1990 after the orbiter ‘Hagoromo’ failed

following separation from a parent spacecraft, Hiten, has also been linked with the intersections between invariant manifolds [58, 59]. Hiten was sent on a transfer that used only system dynamics to insert itself into orbit around the Moon following an unsolicited proposal to the spacecraft's operators by Edward Belbruno. Similar transfers have been used in the trajectory design for the GRAIL [60], SMART-1 [61], CAPSTONE and KPLO missions. The use of the invariant manifolds to achieve ballistic or nearly propellant-free transfers has featured heavily in preliminary mission design studies to reduce overall capture Δv [18, 19, 35, 48, 62, 63, 64, 65, 66, 67, 68, 69]. It has performed a particularly key role in the investigation of asteroids and other large bodies, since their masses make traditional methods of orbital insertion difficult [70, 71, 72, 73, 74]; two asteroids, 2006 RH120 and 2020 CD3, have even been temporarily captured by the Earth. An investigation into the use of the manifolds in this role is performed in chapter 4, which also gives an overview into the practical construction of periodic orbits and their associated invariant manifolds.

3.3 Dynamical features and flow heuristics in nonautonomous systems

While the periodic orbits and invariant manifolds of the Circular-Restricted Three-body Problem have been well-studied in the literature, in some cases the assumption that the two primaries orbit each other in circles is inappropriate (for example, the Sun-Mercury system has $e_p = 0.2$). When moving to the Elliptic-Restricted Three-body Problem, which is a more realistic model for eccentric systems and relaxes the assumption that the primaries orbit each other in circles, the system becomes time-periodic. As discussed earlier, the invariant manifolds – found by studying trajectory behaviour as $t \rightarrow \pm\infty$ in the CR3BP – become more numerically difficult to find directly in the ER3BP due to this time-dependence. Often, mission designers correct trajectories that use the time-independent dynamical structures to their counterparts in the elliptic problem [75, 76, 77], but this relies on the dynamics of the simplified system rather than using the ER3BP dynamics directly.

In systems with arbitrary time-dependence, there is no guarantee that these invariant manifolds exist at all [78]. Locating similar structures to the invariant manifolds in systems with arbitrary time-dependence has been, and remains, a great challenge in the field of dynamical systems theory and fluid dynamics [79]. Recent literature has begun to use some of these fluid dynamics techniques in the context of studying time-dependent systems in space mission design. A review of the most common methods for finding these analogous structures in fluid dynamics, with an emphasis on those which have already seen use in spaceflight problems, follows.

In lieu of the invariant manifolds, the fluid dynamics literature seeks coherent structures, which are in their broadest sense structures in the flow which persist over long lifetimes and which exert significant influence on nearby trajectories [80, 81, 82, 83, 84]. The search for coherent structures begun with simple, statistical methods. Mezić and Wiggins [85] found approximations to the finite-time invariant manifolds based on the average velocities of a grid of initial conditions.

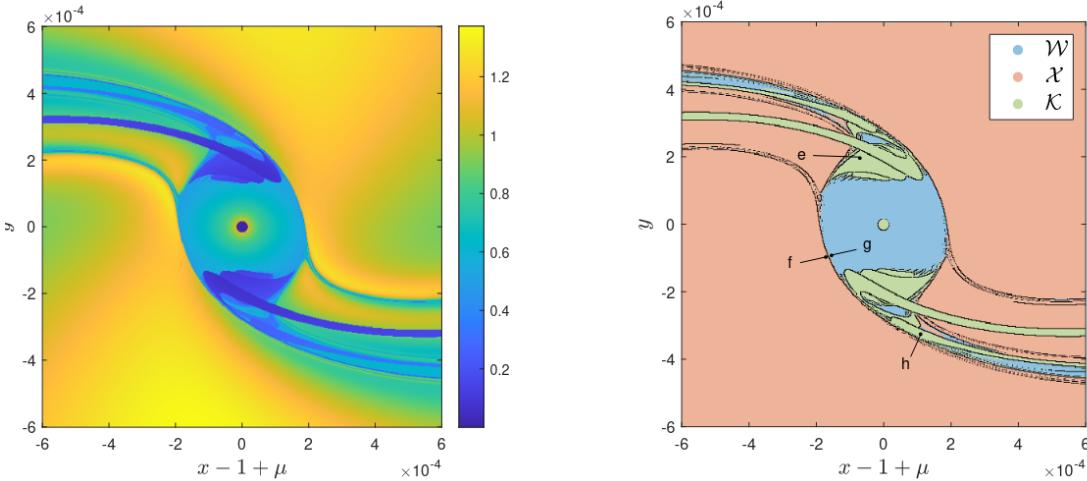


Figure 3.3: Example LD field on the Sun-Mars ER3BP (left) and the ‘separatrices’ obtained after performing an edge detection algorithm with manually-determined behaviours overlaid in different colours (right). The integrand used is $\|v\|^{\frac{1}{2}}$, which highlights features well. However, the choice of integrand is generally not known *a priori* and must practically be selected through trial-and-error. Figure taken from Quinci, Merisio, and Topputo [2].

Regions with similar average velocities are taken to have similar finite-time behaviour, and thus the boundaries of these regions could be said to approximate the invariant manifolds in globally separating the phase space and thus separating qualitatively different behaviours.

This approach was developed in the following years in parallel with more principled techniques related to flow properties. For example, in Mendoza and Mancho [86] and Mendoza, Mancho, and Rio [87] the authors used the arc-length of individual fluid trajectories to profile the dynamical features of the flow in discrete ocean datasets. Again, the premise is that trajectories that begin close but with significantly different arc lengths are likely separated by some form of dynamical structure. In this case, the authors found good agreement between this heuristic and known structures in the flow, including invariant manifolds and hyperbolic and non-hyperbolic flow regions.

The Lagrangian Descriptor (LD) followed in 2013 as the extension of this concept to other bound, positive quantities. In Mancho et al. [88] a variety of arclengths were introduced, such as the sum of velocities and accelerations of trajectories, and investigated for their potential to similarly highlight regions of qualitatively different behaviour. The authors show through several examples that depending on the integrand chosen, one can identify different features in the flow.

The LD has seen widespread use in recent years, including in modelling ocean flows, oil spills and the behaviour of Monsoons [86, 89, 90], cardiovascular flows [91] and wider dynamical systems and classical mechanics applications [92, 93, 94]. In space mission design, two MSc theses [95, 96] and related preprints⁴ [2, 97] used the LD to identify ballistic capture sets and bounded motion in binary asteroid systems. In the former, different types of orbit exhibit signifi-

⁴These papers are still under review at the time of writing, 17 July 2022.

cantly different values of LD, leading to discontinuities in their values. Standard edge detection algorithms can then be used to extract the separatrices from the LD field. This approach yields accurate separatrices between these orbits once the edge detection algorithm threshold has been adjusted. In the latter, the LD scalar field was used to identify different families of periodic orbits; the largest discontinuities in the field identified different families without *a priori* determination of individual orbits. However, the correct identification of integrand to use was crucial in both cases, with some integrands not producing usable insight.

Several issues arise with the LD. Firstly, the selection of the integrand is nontrivial and must be chosen without any *a priori* insight, often through trial and error. The LD is also non-objective, which means that its result changes depending on the reference frame chosen [98]. Given the range of reference frames used in mission design, this makes a reliable application of the LD difficult. Nonetheless, the method is very simple to compute and implement and requires little additional computational complexity compared to the integration of the reference trajectory. It may thus serve a use as a metric for a ‘first guess’ [79].

There are other non-objective flow analysis techniques that are widely used in fluid dynamics but not discussed in further detail here due to both their nonobjectivity, which has been identified as a potential issue, and also due to their lack of use in astrodynamics. These include the Eulerian Okubo-Weiss criterion [99], which can quantify the amount of deformation or rotation at a given point in the flow and thus delineate hyperbolic and elliptic regions, the Q-criterion [100] and $\lambda - 1$ criterion [101] for vortex detection, Patchiness [102] for analyses similar to the LD, and Mesochronic Analysis, which analyses the eigenstructure of the flow map of a given trajectory [103] to determine broad flow characteristics.

Shortly after the original work of Mezić in 1999, Bowman [104] suggested the use of the linear separation between trajectories that begin next to each other on a grid (a ‘finite strain map’) along with the edge detection step of the LD to determine the location of coherent structures. Define the finite strain map as

$$\frac{|\mathbf{x}(t_0, \mathbf{x}_0; T)_2 - \mathbf{x}(t_0, \mathbf{x}_0; T)_1|}{|\mathbf{x}_{0,2} - \mathbf{x}_{0,1}|} \quad (3.1)$$

for two trajectories $\mathbf{x}_{0,2}$ and $\mathbf{x}_{0,1}$ that begin infinitesimally close and evaluated from time t_0 to time T . For $T < t_0$, maxima of the strain map highlight the stable manifolds, and for $T > t_0$ maxima highlight the unstable manifolds. Lekien and Ross [105] praised this approach for being able to determine coherent structures for even incredibly noisy datasets with a significant number of features, as well as being numerically simple to evaluate and relatively simple to implement.

Improvement on this linear separation estimate however is relatively straightforward [106]. Two modern indicators are often used to quantify separation, the finite-time Lyapunov Exponent (FTLE) [81, 107] and the finite-size Lyapunov Exponent (FSLE) [108]. The former quantifies the amount of separation between two trajectories which begin infinitesimally close after a given time, and the latter measures how much time is required for the separation to grow to a fixed size. The FTLE (and closely-related Fast Lyapunov Indicator) has been used in a variety of

dynamical systems [78, 109, 110, 111] and astrodynamics roles [112, 113, 114, 115] for quantifying chaoticity. While the FTLE is an objective measure, the following chapters highlight how it can in some cases present false positives, and suggest an alternative structure defined by fluid dynamics as a structure with no false positives.

3.4 Lagrangian Coherent Structures

Motivated by use of an FTLE as an improved separation metric on the finite strain map, an influential series of papers by Haller [79, 81, 116, 117, 118] defined the Lagrangian Coherent Structure (LCS), a set of structures which exert maximal influence on nearby trajectories. The first LCS defined was the hyperbolic LCS, the locally most repelling or attracting surfaces in a flow. This generalises the concept of the unstable and stable manifolds respectively to systems with arbitrary time dependence and defined over a fixed time period of investigation. These surfaces were defined to be necessarily ridges (maxima, repelling LCS) or trenches (minima, attracting LCS) of the FTLE field.

Additional theoretical development soon followed. Shadden, Lekien, and Marsden [82] examined the flux across an LCS to quantify how effectively repelling LCS act as barriers to transport in flows. They determined that even systems that are time-aperiodic have magnitudes lower flux through the repelling LCS than surrounding surfaces, lending weight to the idea of ridges of FTLE being indicators of such hyperbolic LCS. Lekien, Shadden, and Marsden [119] extended this concept to n -dimensional systems based on the FTLE, finding codimension-1 surfaces that maximise the FTLE in the system. These findings that FTLE necessarily implied the location of such coherent structures were further reinforced by numerical [120, 121] and experimental [122] studies.

As a result, hyperbolic LCS derived from the FTLE field quickly found use in many different fields, particularly in datasets where position and velocities are only known at fixed points. Investigations were published in cardiovascular flows [123], tropical cyclone modelling and weather forecasting [90, 124, 125, 126], aerodynamics studies [127], and the analysis of turbulent flows and PIV data [128, 129, 130, 131]. However, in the early 2010s, publications began to assert that the FTLE alone is not sufficient to identify hyperbolic LCS. Since the FTLE is simply a measure of deformation, it cannot distinguish between stretch and strain (Figure 3.4a) and shear (Figure 3.4b). The latter of these is deformation along the surface rather than orthogonal to it, and is not associated with the behaviour of a hyperbolic LCS. Thus, the FTLE can yield multiple ‘false positives’ for hyperbolic LCS [98], but no false negatives. Haller and Sapsis [133] presented several counterexamples where the FTLE ridges do not identify hyperbolic LCS. A later paper, Haller and Beron-Vera [134], identified that any randomly chosen material surface admits zero material flux, a key finding in the original theoretical study of Shadden, Lekien, and Marsden [82]. Norgard and Bremer [135] identifies that second-derivative ridges of the FTLE field, of which Shadden’s definition of repelling LCS relied upon, do not exist in generic fluid

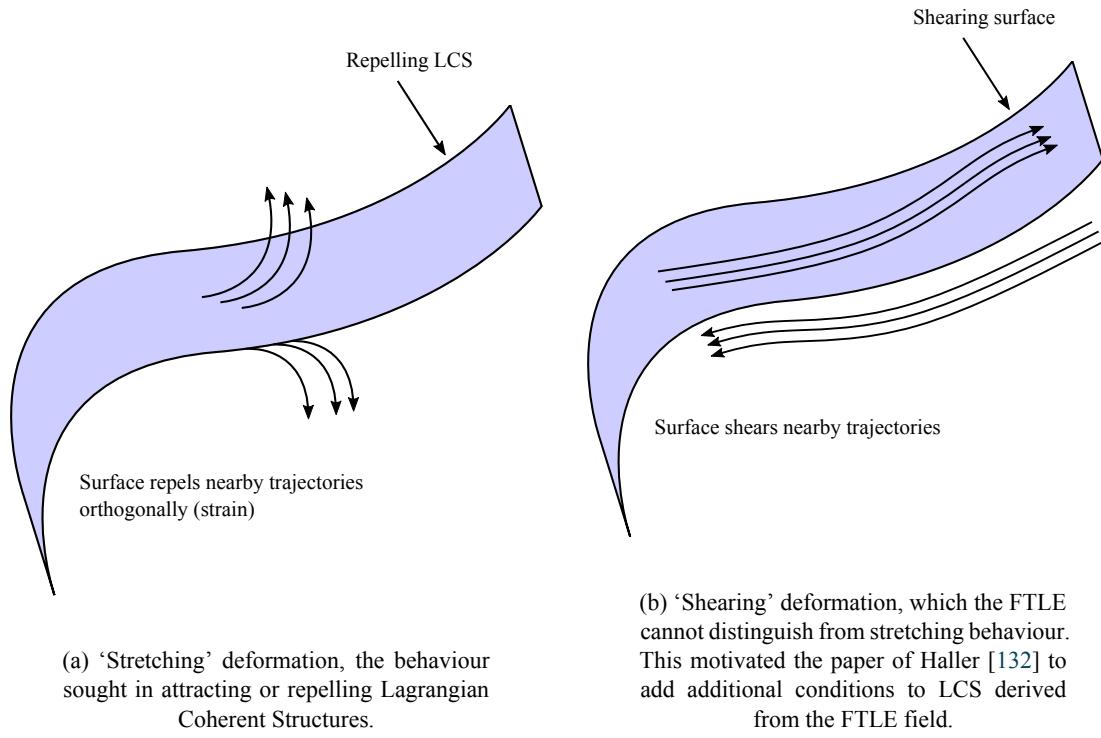


Figure 3.4: Comparison between a surface that maximises repulsion and a surface that causes trajectories to ‘shear’ across its surface. The former is the desired behaviour of a repelling hyperbolic LCS, the second is the behaviour of a shearing surface, which can in some cases be erroneously highlighted by the FTLE field.

flows. Despite Karrasch and Haller [136] finding that attracting hyperbolic LCS do, in fact, lie along trenches of the FTLE field, the authors also remark that it is incredibly difficult to attain the required numerical resolution in the FTLE field to be able to infer the result directly from the FTLE field. In astrodynamics research, Parkash [8], Gondelach, Armellin, and Wittig [112], and Pérez-Palau [137] have shown examples of the FTLE either being a false positive or being affected by things other than the underlying dynamics.

In response, Haller released several publications that strengthen the definition of a hyperbolic LCS, defining them directly from the variational theory of deformation. Haller [132] and Farazmand and Haller [138] presented the mathematical conditions for n -dimensional LCS, along with additional criteria that FTLE ridges must meet to signal a repelling or attracting LCS, taking into account the direction of maximal stretching of the flow. By considering the direction of deformation as well as the magnitude, one can distinguish between stretch and shear. Around the same time, two additional definitions of LCS were introduced. First introduced in Haller and Beron-Vera [134], the elliptic LCS generalises the Kolmogrov-Arnold-Moser (KAM) tori of time-independent dynamical systems (Figure 3.5). A discussion of the relevance of the KAM torus to astrodynamics is outside the scope of this work, but here it suffices to describe the KAM torus as a doughnut-shaped surface for which quasi-periodic orbits move upon even under the presence of perturbation. For a more rigorous introduction on the KAM tori, the reader is referred to one of many excellent overviews [139, 140, 141]. The parabolic LCS delineates

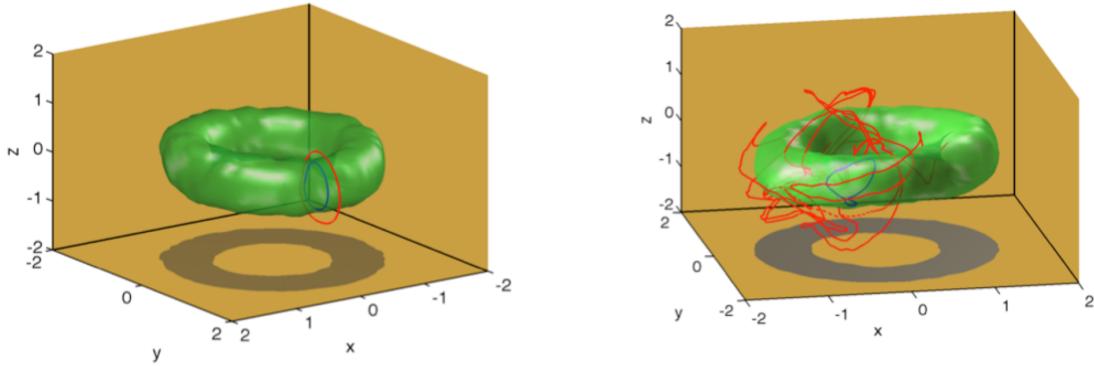


Figure 3.5: Elliptic LCS, which produce time-dependent generalisations of the KAM torus. Shown here is the elliptic LCS for the chaotically-forced Arnold-Beltrami-Childress flow, taken from Blazevski and Haller [3].

‘jet-type cores’, series of trajectories which neither stretch nor shear [142]. Such features are typically only present in geophysical flows and are thus less well-studied in the literature.

Given the importance of the role the hyperbolic invariant manifolds have been seen to play in space mission design in the preceding subsection, this thesis focuses only on the hyperbolic LCS. To continue the introduction of hyperbolic LCS, its formulation is now introduced in more detail, following the works of Haller [132] and Farazmand and Haller [138]. I first introduce a conceptual definition for hyperbolic LCS, then review the theoretical formulation for n -dimensional flows, since these are well-studied in fluid dynamics and provide the main theoretical underpinning of the variational theory. A specialisation of the method for $n = 3$ is then introduced based upon the work of Blazevski and Haller [3].

The Cauchy-Green Strain Tensor

Recalling that a hyperbolic LCS is the locally most attracting or repelling surface in a flow, then the unit normal to the surface at every point must necessarily be aligned with the direction of minimum or maximum deformation along a trajectory, respectively. To quantify this deformation, consider two curves $\mathbf{x}(s)$ and $\mathbf{X}(s)$ parameterised by some scalar parameter $s \in [0, 1]$. The curve $\mathbf{x}(s)$ is an undeformed curve defined at some initial time t_0 , and $\mathbf{X}(s)$ is the same curve at a later time T after deformation under a dynamical system f , such that

$$\dot{\mathbf{x}} = f(\mathbf{x}, t), \mathbf{x} \in D \subset \mathbb{R}^n, t \in [t_0, t_0 + T] \quad (3.2)$$

with flow map

$$F: \begin{cases} D \mapsto D \\ \mathbf{x}_0 \mapsto \mathbf{x}(t_0, \mathbf{x}_0; T) \end{cases} \quad (3.3)$$

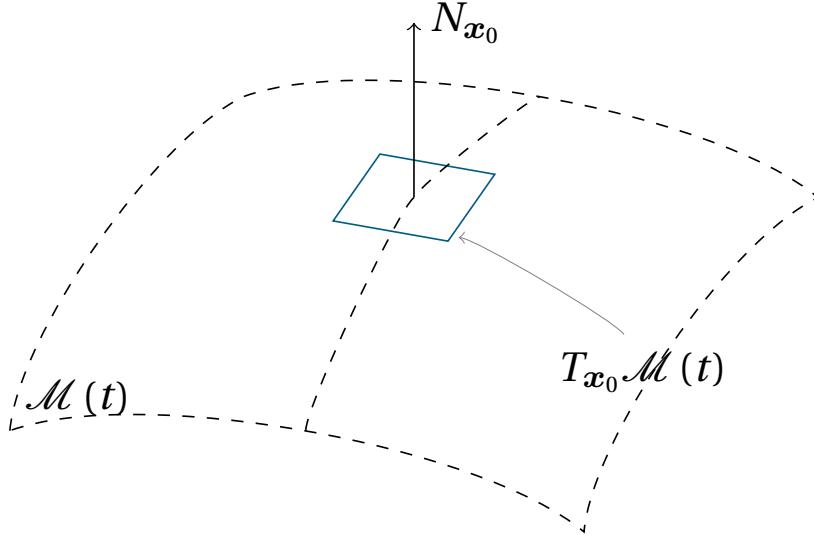


Figure 3.6: Geometry of a surface \mathcal{M} , including the definition of the normal N and tangent spaces T to a point x_0 on the surface used in definitions of LCS from an FTLE ridge.

where a trajectory of the system starting at position x_0 at time t_0 up to a time T is denoted as $\mathbf{x}(t_0, x_0; T)$. The length of the undeformed curve \mathbf{x} is

$$\int_0^1 \left| \frac{\partial \mathbf{x}}{\partial s} \right| ds = \int_0^1 \sqrt{\frac{\partial \mathbf{x}}{\partial s} \frac{\partial \mathbf{x}}{\partial s}} ds = \int_0^1 \sqrt{\frac{\partial \mathbf{x}}{\partial s} \cdot I \cdot \frac{\partial \mathbf{x}}{\partial s}} ds \quad (3.4)$$

where I is the identity matrix. On the other hand, the deformed curve \mathbf{X} has length

$$\int_0^1 \left| \frac{\partial \mathbf{X}}{\partial s} \right| ds = \int_0^1 \sqrt{\frac{\partial \mathbf{X}}{\partial s} \frac{\partial \mathbf{X}}{\partial s}} ds = \int_0^1 \sqrt{\left(\frac{\partial \mathbf{X}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial s} \right) \cdot \left(\frac{\partial \mathbf{X}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial s} \right)} ds \quad (3.5)$$

such that

$$\int_0^1 \left| \frac{\partial \mathbf{X}}{\partial s} \right| ds = \int_0^1 \sqrt{\frac{\partial \mathbf{x}^\top}{\partial s} \left[\left(\frac{\partial \mathbf{X}}{\partial \mathbf{x}} \right)^\top \cdot \left(\frac{\partial \mathbf{X}}{\partial \mathbf{x}} \right) \right] \frac{\partial \mathbf{x}}{\partial s}} ds = \int_0^1 \sqrt{\frac{\partial \mathbf{x}^\top}{\partial s} \cdot C_{t_0}^T \frac{\partial \mathbf{x}}{\partial s}} ds. \quad (3.6)$$

The quantity in Equation 3.6 that replaced the identity matrix in Equation 3.4 is the right Cauchy-Green Strain Tensor (CGST)

$$C_{t_0}^T := (\nabla \mathbf{F}_{t_0}^T)^\top (\nabla \mathbf{F}_{t_0}^T) \quad (3.7)$$

where the subscript t_0 and superscript T is used to quantify the times associated with the dynamical system that $\mathbf{x}(s)$ is subject to. The quantity $\nabla \mathbf{F}_{t_0}^T$ is the first derivative of the flow map, $\nabla \mathbf{F}_{t_0}^T = \frac{d\mathbf{X}}{d\mathbf{x}}$, which is also known as the state transition matrix in astrodynamics.

By inspection of the integrand, it is evident that $C_{t_0}^T$ quantifies the deformation of the length of the initial curve. By its definition, it is symmetric and positive definite and thus has all real eigenvalues of magnitude greater than zero. The remainder of this thesis uses the notation that its eigenvalues are $\lambda_1, \lambda_2, \dots, \lambda_n$ and the associated eigenvectors are $\zeta_1, \zeta_2, \dots, \zeta_n$ such that $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$. The CGST is invariant under changes in coordinates and frames [136].

The FTLE $\sigma_{t_0}^T$ introduced earlier is also derived from $C_{t_0}^T$ as

$$\sigma_{t_0}^T := \frac{1}{2(T - t_0)} \log(\lambda_n). \quad (3.8)$$

Since the attracting and repelling hyperbolic LCS are necessarily orthogonal to the minimum and maximum directions of deformation, it is conceptually clear that they are necessarily orthogonal to ζ_1 or ζ_n , respectively. Mathematically, Haller defines the repelling and attracting LCS as maximisers and minimisers of the repulsion ratio \mathcal{R} , respectively

$$\mathcal{R} = \frac{1}{\sqrt{\left\langle \hat{n}_0, [C_{t_0}^T]^{-1} \hat{n}_0 \right\rangle}} \quad (3.9)$$

for a surface with initial unit normal \hat{n}_0 .

For dimensions greater than two, it can be shown (Haller [132]) that for an eigenvector defined at a given point \mathbf{x}_0 to be a maximiser of the repulsion ratio it must necessarily satisfy four conditions on a surface $\mathcal{M}(t_0)$ at a given time t_0 over a time T :

1. $\lambda_n(\mathbf{x}_0) > 1, \lambda_n(\mathbf{x}_0) > \lambda_{n-1}(\mathbf{x}_0) > \dots > \lambda_1(\mathbf{x}_0)$
2. $\langle \zeta_n(\mathbf{x}_0), \nabla^2 \lambda_n(\mathbf{x}_0) \zeta_n(\mathbf{x}_0) \rangle < 0$
3. $\nabla \lambda_n(\mathbf{x}_0) \parallel \mathcal{M} \implies \nabla \lambda_n \in T_{\mathbf{x}_0} \mathcal{M}$
4. $\zeta_n(\mathbf{x}_0) \perp \mathcal{M} \implies \zeta_n(\mathbf{x}_0) \in N_{\mathbf{x}_0} \mathcal{M}.$

Condition (1) requires there to be a dominant eigenvalue of distinct magnitude to allow a unique ‘dominant direction’ of deformation to exist. Condition (2) is derived from the FTLE and enforces that the curve is on a ridge of the FTLE (maximal stretch or shear), and Condition (3) and Condition (4) ensure the correct orientation of the surface to the required directions to resolve between stretch and shear, where \parallel is the common shorthand for ‘parallel to’, i.e. $\nabla \lambda_n$ lies in the tangent space of \mathcal{M} (Figure 3.6). Condition (3) and (4) are the additional variational criteria that must be satisfied for a FTLE ridge to be a repelling hyperbolic LCS and enforce orthogonality to ζ_n ; other equivalent forms of the conditions given above can be derived but the LCS as derived from a FTLE ridge is used here to highlight the additional requirements for the FTLE that must be satisfied. This method has no known false positives or negatives, unlike the FTLE.

The author does not know of any study which has applied the conditions given above directly for $n > 2$. The work of Blazevski and Haller [3] defined a specific case of the formulation of the LCS for $n = 3$, seeking surfaces that are orthogonal to ζ_n directly rather than applying Conditions (1–4) given above. They find that, as per Palmerius, Cooper, and Ynnerman [143], points on surfaces which are everywhere orthogonal to a given vector field ξ necessarily satisfy a zero

helicity criterion H :

$$H_\xi = \langle \nabla \times \xi, \xi \rangle = 0. \quad (3.10)$$

This quantity is defined from Stoke's theorem as follows.

Proposition 1. Let w be a vector field defined in \mathbb{R}^3 , and S a surface orthogonal to w with unit normal at every point \hat{n} . For any $\xi \in S$, H_ξ must vanish on S

$$H_w = \langle \nabla \times w(\xi), w(\xi) \rangle. \quad (3.11)$$

Proof. Consider some open neighbourhood $D \subset S$ of ξ in S . By Stoke's theorem, we have:

$$\int_D (\nabla \times w) \cdot \hat{n} = \int_C w \cdot dr. \quad (3.12)$$

Since w is orthogonal to S , then $\int_C w \cdot dr = 0$. Consider also that $w = \langle w, \hat{n} \rangle \hat{n}$, which yields:

$$\int_D \frac{1}{\langle w, \hat{n} \rangle} H_w dA = 0. \quad (3.13)$$

Since D was defined arbitrarily, then H_w must be zero on S . \square

The helicity will be used extensively in chapters 5 and 6. In the case of repelling LCS, one seeks level sets of zero H_{ζ_n} ; for attracting LCS, one seeks level sets of zero H_{ζ_1} . It can be shown (Appendix C of Blazevski and Haller [3]) that points which satisfy the zero helicity criterion on reference hyperplanes placed in the domain are global maximisers or minimisers of repulsion⁵. By orienting different hyperplanes in the flow, one can interpolate between the intersections of the LCS with the reference plane to produce the full, three-dimensional structure. This three-dimensional formulation has been applied in the literature directly to simple test cases from fluid dynamics and shown to be effective.

While no LCS is known to the author to have been constructed for astrodynamics systems using the variational theory of LCS in three dimensions, significant contributions have been made by other authors that show the feasibility of LCS providing insight in time-dependent astrodynamics systems. Early applications of the LCS to astrodynamics systems used LCS defined with respect to the FTLE. The first known to the author was Gawlik et al. [144] who studied repelling forward-time LCS in the Earth-Moon and Sun-Mercury planar ER3BP using ridges of the FTLE to find LCS, motivated by early planning for the BepiColombo mission: since the eccentricity of Mercury about the Sun is 0.2, the CR3BP is not a suitable assumption for its motion around the Sun, and as a result similar structures to the invariant manifolds were sought in the ER3BP. Using the formula for flux across an LCS suggested by Shadden, Lekien, and Marsden [82], Gawlik

⁵The zero helicity points for the ‘middle’ eigenvector ζ_2 define elliptic LCS, which maximise shear and generalise the KAM torus or the parabolic jet.

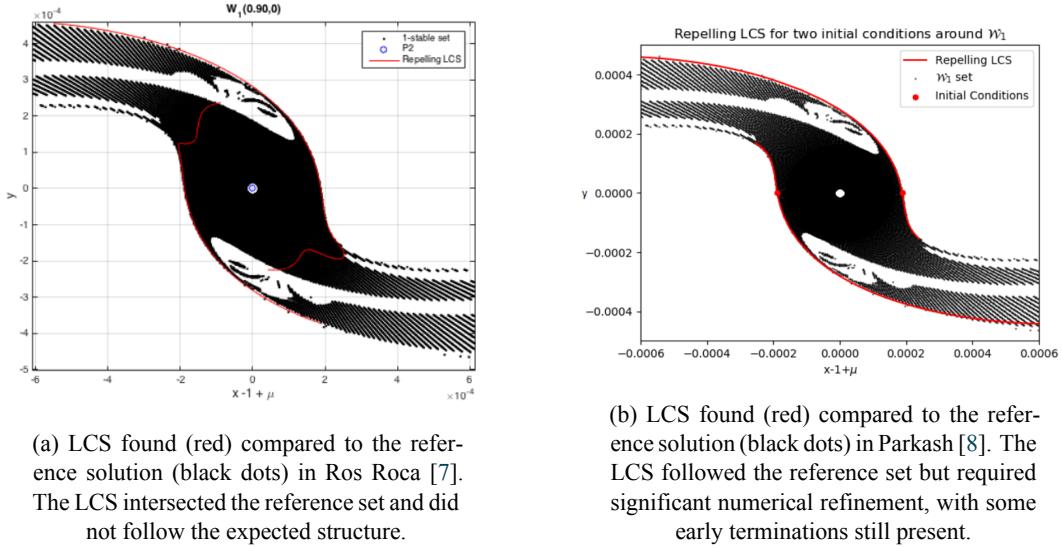


Figure 3.7: Previous attempts at computing LCS in the Elliptic-Restricted Three-body Problem. Both authors cited numerical issues as preventing the construction of the full LCS.

showed that repelling hyperbolic LCS do indeed act analogously to invariant manifolds in separating regions of qualitatively different behaviour. Despite the FTLE not *de facto* highlighting LCS, this study showed the potential effectiveness of the LCS in time-dependent astrodynamics systems.

Later studies followed from Cody Short. In a series of papers released around the same time the definition of the LCS was strengthened to include additional criteria from variational theory, Short, Howell, and Tricoche [36] began by investigating the FTLE field on Poincaré sections in combination with GPU parallelism. Again, it was shown that even on arbitrarily-oriented sections of position space the LCS still continues to separate regions of qualitatively different behaviour. A follow-up journal paper, Short and Howell [145], demonstrated that this also holds in mixed position-velocity phase space, and that additional information derived from nearby FTLE values can be used to aid the analysis of determining the stability of periodic orbits. A later paper by Short et al. [113] collaborated with Daniel Blazevski and George Haller in the construction of strainlines on a fixed section of phase space in four-dimensional mixed position-velocity space, with the initial conditions sampled on regions of high FTLE rather than from the variational theory. Nonetheless, this paper identified that the strainline approach holds relevance in astrodynamics, later reinforced by a similar study by Oshima and Yanao [146] who applied the same technique to the more complex planar bicircular-restricted four-body problem, and found that the strainlines still separate qualitatively different behaviour in even higher fidelity models.

Additional work in categorising the LCS from the FTLE field in astrodynamics has been performed since. Qingyu, Mingpei, and Ming [147] extracted the ridges of the FTLE field using a Particle Swarm optimisation algorithm in the planar and hyperbolic restricted three-body problem. Onozaki, Yoshimura, and Ross [148] used the FTLE to determine tube-like structures that mimic the invariant manifolds in the four-body problem. Pérez, Gómez, and Masdemont [149]

used hyperbolic LCS to detect analogues to invariant manifolds from FTLE ridges, again finding good agreement except in ‘spurious points’. Such points were later refined by applying some of the additional criteria for LCS from the FTLE field following Haller [132], which removed some of these spurious points, and another filter by comparing to the FTLE computed at a different time. Additional dynamical indicators similar to LCS, based largely on polynomial algebra (section 5.3.1, known as ‘Jet Transport’ in those works) were then introduced to mimic LCS structures in Pérez, Gangnet, and Blake [150], but the author is not aware of the use of any of these indicators beyond that investigation.

The only attempts known to the author to compute LCS from their variational theory directly in astrodynamics systems are two MSc theses. Ros Roca [7] was the first, attempting to use LCS to provide automatic insight on the ballistic capture phenomenon. He found that identifying seed points using divided differences was numerically very challenging, even when using variational equations and analytic Jacobians to obtain the derivatives of the flow itself. As a result, the LCS was very difficult to find, and even when it could be found to exist, it did not match expectations from the literature (Figure 3.7). A follow-up thesis in 2019, Parkash [8], continued the study by using an in-depth pointwise interpolation and refinement algorithm, and obtained better results, but ones which seemed to terminate early.

In both cases, the authors cited numerical issues in preventing the construction of the LCS. Addressing these numerical issues may thus provide insight into LCS in space mission design, allowing missions to be constructed directly in time-dependent models of motion, as well as having potential benefits for those outside of astrodynamics whom also experience numerical issues in their construction [151, 152]. This apparent gap in the literature will be addressed in this thesis in chapters 5 and 6.

3.5 Conclusion

This chapter has introduced the dynamical structures and dynamical indicators that have been widely used in astrodynamics thus far. In time-independent systems, the equilibria of the system, periodic orbits and their associated stable and unstable invariant manifolds have been widely used to profile dynamical behaviour and to achieve propellant-free ‘ballistic’ captures in modern mission design, as well as unique trajectories and orbits that help satisfy mission objectives.

It was then revealed that in time-dependent systems, there is no guarantee on the existence of these structures, and analogous indicators originally developed for fluid mechanics have begun to be adapted into spaceflight problems. An overview of many that have been used in astrodynamics has been given, before identifying the Lagrangian Coherent Structure as a time-dependent generalisation of the concept of the invariant manifolds. However, the use of these in astrodynamics has been made difficult by numerical issues in previous investigations, even in two-dimensions. Remedyng these numerical difficulties could allow for these structures – well-established in fluid dynamics – to be applied directly to time-dependent models of motion.

The coming chapter will use the invariant manifolds of the time-independent CR3BP in the context of asteroid retrieval missions. Given that these structures are well-established in space mission design, and the search space for asteroid retrieval trajectories can quickly become large, these structures form a foundation on which to base the investigation as to whether high-performance computing can provide additional insight in astrodynamics.

4

Asteroid Retrieval Missions using the Invariant Manifolds of the Circular Problem

This chapter presents an investigation into the use of dynamical structures in the CR3BP reviewed in the previous chapter to retrieve asteroids and comets that pass close to the Earth. This builds upon a body of research in the literature that examines retrieval with current propulsion technologies. The methodology used to find retrievable asteroids is modified to improve their identification, and high-performance computing is used to improve the effectiveness of the search for retrieval trajectories. The chapter goes on to analyse search strategies for these transfers and identifies new characteristics which improve the robustness of asteroid retrieval trajectories using such dynamical structures. Content from this chapter has been published as Tyler and Wittig [11] and presented at the 31st AAS/AIAA Space Flight Mechanics Meeting in 2021 as Tyler and Wittig [12].

4.1 Background

In recent years, there has been significant interest in the investigation and exploration of Near-Earth Objects (NEOs), asteroids and comets which pass within 1.3 Astronomical Units (AU) of the Sun. Much of this is motivated by their composition, which mirrors the structure of the early Solar System and thus offers unique, invaluable insight into the history of our stellar neighbourhood [153].

This interest is such that many major space agencies have either conducted or planned missions to investigate asteroids, including multiple missions to return samples of asteroids back to the Earth for further analysis. The National Aeronautics and Space Administration (NASA) sent the first mission to investigate a NEO in 1985, with the International Cometary Explorer (ICE)

spacecraft flying through the plasma tail of Giacobini-Zinner. The 1996 mission Near-Earth Asteroid Rendezvous (NEAR), also designed by NASA, was the first spacecraft to orbit and land on a NEO in 1997 and 2001, respectively. Five years later, the Stardust mission was the first to return samples from a NEO. The Japanese Aerospace Exploration Agency (JAXA) later followed these missions with Suisei in 1986, a fly-by, and the Hayabusa-1 and Hayabusa-2 sample return missions [154]. The European Space Agency launched the infamous Rosetta mission in 2004, which performed the first soft-landing on a comet in 2014 [155]. NASA's ongoing OSIRIS-REx mission will return approximately 60 g of samples to the Earth in 2023 [156].

NEOs have long been classified into families, typically grouped by their orbital properties. Atens are NEOs which cross the orbit of the Earth, with semi-major axes less than 1 AU (Astronomical Units, 1AU = 149597870.7 km), and an aphelion¹ distance greater than 0.983 AU. The Apollo family also crosses the Earth, but their semi-major axes are greater than 1 AU, and their perihelion distance is less than 1.017 AU. Amors have orbits completely outside of the Earth, such that their semi-major axes are greater than 1 AU, and their perihelion distance is between 1.017 and 1.3 AU.

These groups are relatively well-established, with their prototype asteroids discovered in 1976 for Aten-type asteroids, and 1932 for both Apollo- and Amor-type asteroids. A newer family of NEOs was introduced in 2003 after the discovery of the first NEO of its kind. Known as Atiras, the orbits of these NEOs are wholly within the Earth's orbit, such that their semi-major axes are less than 1 AU, and their aphelion distance is less than 0.983 AU.

While these families are the most commonly accepted, others have been proposed. Arjuna-type objects, of which there are several hundred expected to exist, have semi-major axes between 0.985 and 1.013 AU, orbital eccentricities below 0.1, and inclinations less than 8.46 degrees [157]. They are so defined to be similar to the orbit of the Earth, and are expected to be excellent candidates for asteroid retrieval missions. However, few of these types of asteroids have been discovered due to their small size [158].

A similar family of Small Earth-Approachers was suggested by Brasser and Wiegert [159]. The definition is similar to the Arjuna-type asteroids, with $0.95 \leq a \leq 1.05$ AU, $0 \leq e \leq 0.1$ and $0^\circ \leq i \leq 10^\circ$. An additional requirement that asteroids have diameters less than 50 m is placed on these objects.

Other authors have suggested classifying NEOs by the hazard they pose to the Earth. Potentially Hazardous Objects (PHAs) have their distance between the NEO and the Earth at any given time less than 0.05 AU. Their absolute magnitude must also be equal to or brighter than 22 [160]. As of June 2022, approximately 2000 PHAs are known to exist. They have demanded extra attention in the years following the 2013 Chelyabinsk superbolide, an atmospheric air burst of a

¹The suffix of 'apoapsis' and 'periapsis' are often changed to denote the host object. The -helion suffix implies the Sun is the host object.

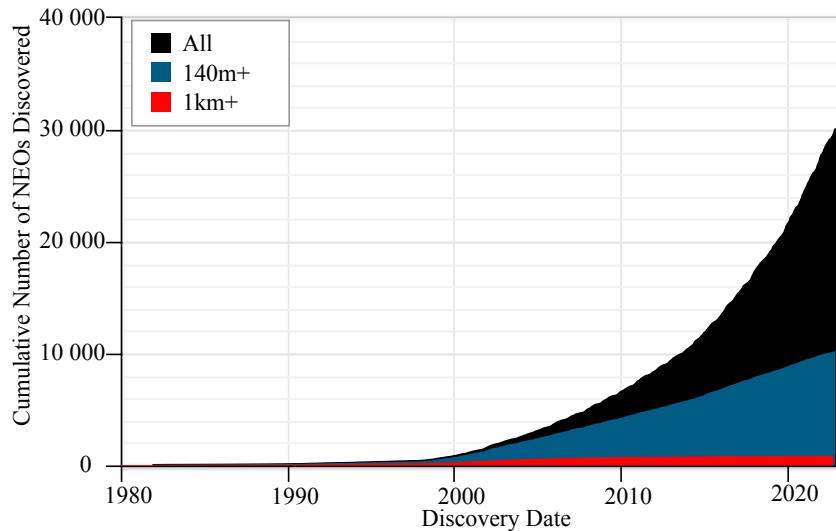


Figure 4.1: The number of NEOs detected over time by the Centre for Near-Earth Object Studies (CNEOS), grouped by diameter. Figure adapted from NASA JPL [4].

meteor in Russia caused by an asteroid just 20 metres wide and releasing energy equivalent to 500 kilotons of TNT [161].

Many organisations – both private and public – have thus become interested in the mitigation of the dangers PHAs and generic NEOs pose in the event of a collision with the Earth, a field known as planetary protection. NASA has a dedicated telescope, NEOWISE, to detect and catalogue NEOs, of which approximately 30000 have been found (Figure 4.1). The development of vehicles for deflecting asteroids on a collision course with the Earth is currently ongoing for such an event (for example, the B612 foundation²), and asteroid exploration provides both technology demonstrations and crucial information as to material properties of potential threats. These deflection methods are also of use in asteroid retrieval missions.

Several methods of deflection have been proposed for retrieval and planetary protection. The ‘simplest’ of these techniques, the use of traditional propulsion systems, provides a system for deflection which requires relatively little technological development [162]. However, the total impulses required for such objects are often too high for chemical propulsion, and constant low-thrust systems may be more suitable [163]. The trajectory and attitude control requirements for retrieval methods based on direct contact for low-thrust have been studied extensively [164]. Other methods that involve direct contact with the NEO, such as the ‘lasso’ method intended to be used in the now-cancelled NASA Asteroid Redirect Mission [165, 166, 167], chained series of nuclear explosions [168], and kinetic impacts [169], have been studied in the literature.

Methods that do not require direct contact have also been proposed, as such methods may not be suitable for asteroids with unstable or weak surface composition. For example, the gravity tractor, devised in 2005 by former Astronaut Ed Lu, uses the mutual gravitational attraction between a ‘shepherding’ spacecraft and a NEO [170]. The spacecraft, ideally of relatively high

²<https://b612foundation.org/>

mass, maintains a close position to the asteroid using a set of electric propulsion systems to allow station-keeping. The mutual gravitation between the NEO and the spacecraft then induces small deflections in the trajectory of the NEO, the magnitude of which can be changed using the mass of the spacecraft and the distance between the spacecraft and the NEO.

Some authors have argued several drawbacks for this idea. Firstly, the mass of the spacecraft itself may be need to be impractically high, depending on both the deflection magnitude, and the time over which the deflection can be performed [171]. The station-keeping can also be difficult to control. A separate concept known as the Ion Beam Shepherd (IBS) has been proposed to mitigate these drawbacks of the gravity tractor [172]. It uses a pair of identical ion propulsion units: one is aimed directly at the NEO, and the other is aimed directly opposite to provide stationkeeping capabilities. The force of the ion beam impinging on the NEO then provides the necessary deflection. The magnitude of the deflection depends only on the power and propellant available to the propulsion unit(s), which can also be used as part of the orbital transfer to the NEO. This concept has also been studied for active debris removal [173].

However, the impulses required for retrieval can still be large. The retrieval must, therefore, be into orbits which can be reached with minimal transfer velocity: these are the ‘low-energy transfers’ mentioned earlier. In 2011, Sanchez and McInnes [174] investigated the possibility of using periodic orbits about the Sun-Earth L_1 and L_2 and their associated stable invariant manifolds to retrieve NEOs. The study investigated the Jacobi energy of a set of pre-computed invariant manifolds in the Sun-Earth system in combination with a Keplerian Kick Map, a function for rapidly approximating the change in the orbital elements of a Keplerian orbit under (in this case) a third-body perturbation [175]. The study concluded that approximately 0.1% of the then-known NEO population lies on regions near portions of invariant manifolds of the Sun-Earth system. At this point, no transfers were constructed, and the NEO population was estimated using a theoretical distribution model provided in Bottke et al. [176].

In 2013, García Yáñez, Sanchez, and McInnes [71] performed the first quantitative search for asteroids retrievable using the invariant manifolds of the Sun-Earth CR3BP, using the then-known NEO database. They proposed a new class of NEO, the ‘Easily Retrievable Object’ (ERO), NEOs which can be retrieved using the invariant manifolds of the Sun-Earth CR3BP for a total cost of less than 500 m/s. An optimisation procedure between NEOs and pre-computed invariant manifolds is employed to find EROs, with the transfers modelled as impulsive two-impulse Lambert arcs. The NEO database is also pre-filtered using a Hohmann transfer approximation to remove NEOs unlikely to become EROs. Despite the relatively small NEO database size, 12 EROs were found.

This was followed by the use of low-thrust propulsion to assist the retrieval of EROs in Mingotti, Sánchez, and McInnes [177]. An optimal control problem is formulated to use constant low-thrust to transfer the NEO onto a point in a target manifold. The transfers that resulted were typically of lower overall $\Delta\nu$ than those seen in García Yáñez, Sanchez, and McInnes [71],

and allowed for trajectories not possible with either a standard two-body transfer, or impulsive transfers leveraging the invariant manifolds alone. Retrieval was only studied for the 12 EROs identified in the 2013 study.

A study in 2016 examined the control implications for low-thrust retrieval trajectories, and its implications for mission design [178]. Using a stochastic Monte Carlo method, the effect of perturbations to the initial state of the asteroid, as well its estimated mass, were analysed. The study found that some retrieval trajectories could not be controlled with low-thrust, and suggests that the controllability of EROs also needs to be investigated when choosing targets for retrieval missions.

A later study in 2016 from one of the authors of the controllability study, Sánchez and García Yáñez [179], established the state-of-the-art in the field before the publication of Tyler and Wittig [11]. The analysis in 2013 was repeated with a more up-to-date NEO database, and the trajectories were again continued into low-thrust. The controllability of the transfers was not considered. The low-thrust problem was converted from a time-continuous optimal control problem to a non-linear programming problem. Its key advances on previous literature were the consideration of NEO masses, based on their apparent magnitudes, in combination with increasing the number of EROs to 17 from the approximately 16000 NEOs known at the time.

While these studies considered in detail the retrievability and controllability of these trajectories, there has not yet been known to the author a study on the robustness of these trajectories: in the event of deviations from a nominal mission plan, do these NEOs still remain easily retrievable for different transfer times, for example? Moreover, some simplifying assumptions were made in the 2016 study to limit computation time, which will be elaborated in the coming sections. If high-performance computing techniques are used to lift these assumptions, can greater improvements be made?

This chapter outlines extensions to the previous literature in this field based upon the work of two research outputs completed as part of this PhD [11, 12]. To summarise this chapter's findings, it is revealed that the transfer cost for the majority of EROs is approximately constant, and this extends to zero-time transfers, where many EROs can be captured with a single impulsive manoeuvre at only marginally higher cost than the optimal two-impulse transfer. This greatly increases flexibility for mission designers. In the process of this extension, 27 more EROs are found, and as a result of leveraging high-performance computing, 17 new capture trajectories with a retrieval cost of less than 100 m/s are identified.

The structure of this chapter is as follows: firstly, the methods used to construct periodic orbits and their invariant manifolds in the Circular-Restricted Three-body Problem, which serve as destinations for the retrieved NEOs, are reviewed. Then, the chapter outlines how the transfers are constructed, including the assumptions made, and the optimisation procedure. Analysis of the transfers then follows.

4.2 The Methodology

The methodology for finding and analysing EROs is elaborated in this section. Conceptually, the methodology is a three-step ‘pipeline’: firstly, information on the orbital characteristics of NEOs is required (section 4.2.1), and sets of periodic orbits and their stable invariant manifolds must be generated to serve as retrieval targets (section 4.2.2). Next, to reduce the computational expense a fast approximator (pre-filter) must be developed and implemented to remove NEOs unlikely to become EROs to ensure the computationally-expensive optimisation procedure is not performed unnecessarily (section 4.2.3). NEOs which pass this filter are known as ‘retrieval candidates’. Lastly, an optimisation procedure must be performed to identify the subset of EROs from the NEO database (section 4.2.4). Each of these procedures is now discussed in more detail.

4.2.1 Obtaining Near-Earth Object Data

Two primary sources are commonly used for NEO orbital data: the International Astronomical Union’s Minor Planets Centre (MPC) database³, and the NASA Jet Propulsion Laboratory HORIZONS system.

The MPC is responsible for the identification, designation and orbit computation for all minor planets, comets and outer irregular natural satellites of the major planets. A ‘snapshot’ of the orbital elements for each NEO are available as comma-separated value files, which are easily downloaded, manipulated and interfaced with analysis codes in a variety of languages. However, these elements are ‘frozen’ at the given time, and do not reflect changes in NEO orbits in the future as a result of orbital perturbations from close approaches, for example.

As an alternative, NASA’s Jet Propulsion Laboratory provides ephemeris for most of these objects, made available via the JPL HORIZONS system⁴. Ephemeris files contain high-precision positions of the NEO for, in this case, up to 100 years into the future. The HORIZONS system is accessible via a Telnet interface, and NEO data can only be downloaded one object at a time. Moreover, there is significant user interaction required, as options must be selected by the user to determine the NEO desired, the time range of the ephemeris, and the coordinate system used. An FTP transfer is then required to retrieve the ephemeris from a different server.

There were 22,406 NEOs in the Minor Bodies Database as of 24th February 2020. Manually downloading the ephemeris for each of these objects from the HORIZONS system is not feasible, and so previous studies on the topic used either only the ‘frozen’ orbital elements from the MPC [71], or the ‘frozen’ orbital elements for the prefilter (section 4.4) on the entire Minor Bodies Database, and downloaded the ephemeris only for the $\sim 10^2$ retrieval candidates [179].

However, using the frozen orbital elements from the MPC assumes that the orbital elements at the epoch of the study do not change throughout the investigation time-scale, or between

³<https://www.minorplanetcenter.net/iau/mpc.html>

⁴<http://horizons.jpl.nasa.gov>

the time the elements are downloaded and the time the prefilter is performed. It is not clear *a priori* whether this assumption is valid (this is explored in section 4.4), or whether the change is significant enough to dramatically change the retrievability of a NEO. Thus, a script was written in Expect, an extension of the Tcl/k scripting language, to automatically download the high-precision ephemerides from the HORIZONS system. Expect allows the automation of interactive applications, such as FTP and Telnet. The script opens an interface with the NASA HORIZONS system, requests information on each NEO in the MBD in turn, and then opens an FTP connection to the required server to download the ephemeris.

Approximately one ephemeris file can be generated and downloaded every two seconds with the Expect script, which was extended to be executed by a multi-threaded computer to increase throughput. However, this places significant stress on the public JPL HORIZONS system; the time between transfers is extended to comply with access limits for JPL HORIZONS. Approximately two days is required to download all the ephemeris.

The ephemeris were requested between calendar dates Jan 1, 2025 and Jan 1, 2100, with the default time-step. The coordinate system selected was the inertial ecliptic J2000 frame, centred on the Sun. The upper calendar time of Jan 1, 2100 is selected to be consistent with previous literature to provide a valid comparison. The lower bound in the literature is Jan 1, 2020, but this was raised as it precedes the date of this study: the new bound was selected arbitrarily but with consideration to the time between the start of this study and the publication of the previous paper on this topic [179]. The NASA-internal Spacecraft and Planet Kernel (SPK) ID is used to uniquely identify NEOs.

4.2.2 Orbit and manifold generation

Families of planar and vertical Lyapunov, and Northern and Southern Halo orbits about L_1 and L_2 and their associated stable invariant manifolds are generated to serve as retrieval targets for asteroids (Figure 4.2). The techniques used to do this are common in the literature, but are briefly reviewed here for completeness [19, 35, 44, 180]. While other sets of orbit families exist, such as the Near-Rectilinear Halo orbits planned to be used for the NASA Lunar Gateway [181], these orbit families are chosen to match the literature exactly and provide a valid comparison.

One typically obtains a periodic orbit in two stages. First, an approximation to motion is used to generate an initial guess for an initial condition on the orbit. When this initial guess is placed into the full dynamics it is unlikely that the orbit will remain fully periodic, and so a correction algorithm such as Newton's method is used to enforce periodicity in the full dynamics.

The most common approximation to motion about the collinear Lagrangian points⁵ is due to Richardson [43], who constructed an approximation when studying the reference orbit for the

⁵The problem of generating approximations to motion about the Lagrange points is still generating research interest; Servadio, Arnas, and Linares [182] recently used the Koopman operator to design a new approximation as recently as 2022.

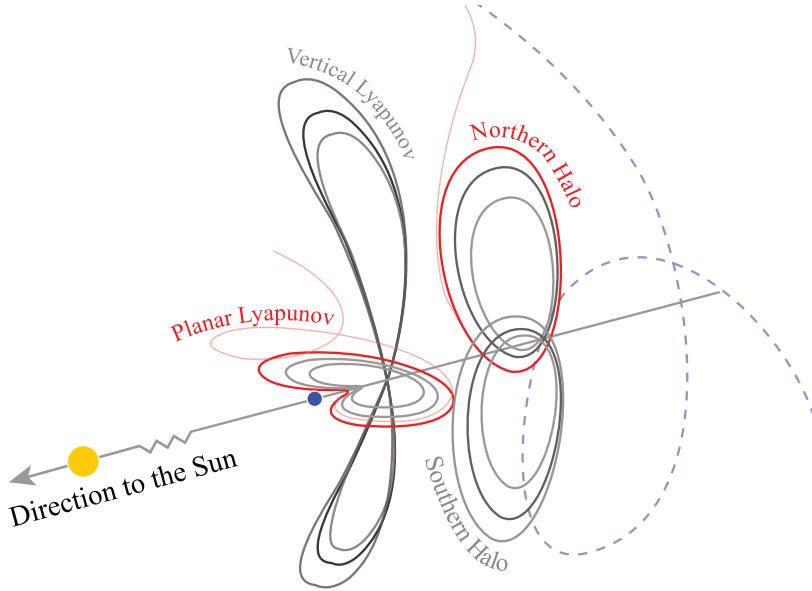


Figure 4.2: The families of periodic orbits studied in this investigation: planar and vertical Lyapunov orbits and Northern and Southern ‘Halo’ orbits. Only the orbits about L_2 are shown for brevity, and only orbits used in previous literature are used to ensure a like-for-like comparison. Highlighted in red are portions of the two retrieval trajectories for asteroids 2020 DW (planar Lyapunov) and 2006 RH120 (Northern Halo), and the nominal trajectory for 2006 RH120 (dashed blue).

ISEE-3 mission [183]. He provides relatively simple, variable-order approximations to motion that practically depend on the user giving two parameters that represent the in- and out-of-plane amplitudes of the orbit. These parameters here were either taken from reference software⁶ or obtained through trial-and-error to find an appropriate initial guess.

Planar Lyapunov orbits are characterised by having no out-of-plane amplitude, existing entirely in the $x - y$ plane, and so to generate these one provides no out-of-plane amplitude. As per Kim and Hall [184], Richardson’s linear approximation is typically sufficient for generating planar Lyapunov orbits about these points. Halo orbits are characterised by equal out-of-plane and in-plane frequencies, and vertical Lyapunov orbits are characterised by unequal out-of-plane and in-plane frequencies. These orbits require a third-order approximation, also made available by Richardson, to generate a suitable initial guess, avoiding secular terms in the approximation. I note that one can also construct Halos and vertical Lyapunov orbits by progressively increasing the amplitude of the planar Lyapunov until a critical amplitude is reached and the orbits bifurcate, but this approach is not used in this work.

When placed into the full dynamics, the approximation often needs to be corrected to ensure it forms a fully closed orbit that returns (at least very close to) its initial condition after one orbital period. As per Bernelli Zazzera, Topputto, and Massari [183], given that the CR3BP equations

⁶Made available as part additional material for the book ‘Dynamical Systems, the Three-Body Problem and Space Mission Design’ [35].

of motion present the following symmetry

$$(x, y, z, \dot{x}, \dot{y}, \dot{z}, t) \mapsto (x, -y, z, -\dot{x}, \dot{y}, -\dot{z}, -t) \quad (4.1)$$

then a periodic orbit that crosses the $x - z$ plane twice and perpendicular to the plane is periodic. This means the desired state vector at the $x - z$ plane is

$$\mathbf{x} = (x, 0, z, 0, v, 0) \quad (4.2)$$

for an arbitrary x and z position and velocity v . If, when integrating to the $x - z$ plane, there are additional non-zero terms of velocity in the x -direction \dot{x}^* and z -direction \dot{z}^* , i.e. the state vector is of the form

$$\mathbf{x}^* = (x, 0, z, \dot{x}^*, \dot{y}^*, \dot{z}^*) \quad (4.3)$$

then by considering a first-order expansion of this state, and recalling from section 2.2 the state transition matrix gives the linear relationship between final and initial conditions, one can derive the following relationship between the correction to the initial condition $\delta \mathbf{x}_0 = (\delta x, 0, \delta z, 0, \delta v, 0)$ required to remove a perturbation $\delta \mathbf{x}_d = (0, 0, 0, \dot{x}^*, 0, \dot{z}^*)$

$$\delta \mathbf{x}_0 = \Phi(0, \tau/2)^{-1} \delta \mathbf{x}_d \quad (4.4)$$

where $\Phi(0, \tau/2)$ is the state transition matrix evaluated over the half-orbital period τ that it takes to achieve the $x - z$ crossing. If the time $\tau/2$ is not known, the events functionality of modern numerical integrators (e.g. SciPy or Matlab) can be used to terminate integration on reaching the $x - z$ plane.

The above equation has three initial states available to target two final states. One of the correction terms is thus held to zero to produce an iterative technique. In this case, δx is set to zero and δz and δv are varied to produce the desired periodicity at the plane crossing.

Since the CR3BP is a Hamiltonian system no periodic orbit exists in isolation. Once one orbit is known, the x -coordinate held fixed in the correction procedure can be adjusted to find a nearby periodic orbit, and so on until an entire family of orbits has been found. The orbits that exist within set bounds of Jacobi constant given in Table 4.2 are used as retrieval targets in this chapter; while these bounds are arbitrary, they match previous literature on the topic to provide a like-for-like comparison. The increment δx is selected to give 1000 orbits between the bounds of Jacobi constant for each family. Individual orbits are identified using a unique index K across all 8 families, such that $1 \leq K \leq 8000$ (Table 4.1). This number of orbits was selected to correspond approximately to those found in the previous literature to enable a fair comparison.

By evaluating the state transition matrix from any point on a periodic orbit for one full orbital period, the monodromy matrix is obtained, which can be used to quantify local stability. In six-dimensional motion, the eigenvalues of M for the periodic orbits studied here exist as two complex conjugate pairs, two unity eigenvalues and two non-unity reciprocal eigenvalues. The

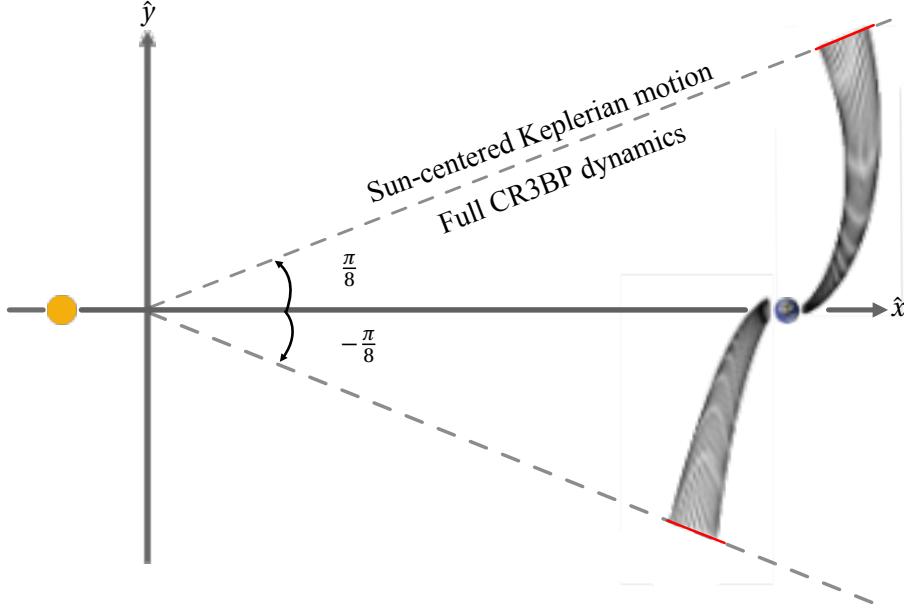


Figure 4.3: The $\pm \frac{\pi}{8}$ planes used as the reference fixed section in this paper: $+\frac{\pi}{8}$ is for manifolds belonging to orbits about L_2 and $-\frac{\pi}{8}$ is for orbits about L_1 . Outside of the ‘cone’ formed by the two planes, the dynamics can be assumed to be two-body only around the Sun; inside of the cone, full CR3BP dynamics are used.

Range of K	Orbit type
[1, 1000]	Planar Lyapunov about L_1
[1001, 2000]	Planar Lyapunov about L_2
[2001, 3000]	Northern Halo about L_1
[3001, 4000]	Southern Halo about L_1
[4001, 5000]	Northern Halo about L_2
[5001, 6000]	Southern Halo about L_2
[6001, 7000]	Vertical Lyapunov about L_1
[7001, 8000]	Vertical Lyapunov about L_2

Table 4.1: Unique identifiers K used to identify periodic orbits and the orbit types they represent.

largest real eigenvalue corresponds to the unstable direction, and the smallest the stable direction. By perturbing a position and velocity on the orbit in the direction of the corresponding stable or unstable eigenvector \mathbf{v}_s or \mathbf{v}_u at that point, one obtains an initial condition for a trajectory that lies on the stable \mathbf{x}_0^s or unstable manifold \mathbf{x}_0^u , respectively

$$\mathbf{x}_0^s = \mathbf{x}_0 \pm \varepsilon \mathbf{v}_s \quad (4.5)$$

$$\mathbf{x}_0^u = \mathbf{x}_0 \pm \varepsilon \mathbf{v}_u \quad (4.6)$$

where $\varepsilon = 10^{-6}$ is used to limit the magnitude of the perturbation to remain in the local space; 10^{-6} is used by convention [18, 20, 179, 185]. The above procedure is performed for each of the 8000 periodic orbits found previously. 360 points per periodic orbit are used to determine initial conditions on the stable manifold. The seed points are located at regular time intervals

Orbit Family	Jacobi Energy Ranges	Shorthand
L_1 Planar	[3.0003, 3.00087]	L1P
L_2 Planar	[2.99985, 3.00087]	L2P
L_1 Halo	[3.00042, 3.00082]	L1N/L1S
L_2 Halo	[3.00025, 3.00082]	L2N/L2S
L_1 Vertical Lyapunov	[3.0002, 3.00087]	L1V
L_2 Vertical Lyapunov	[2.99935, 3.00087]	L2V

Table 4.2: The ranges of Jacobi energy for which orbits were generated, inclusive. 1000 orbits were generated per family with equi-spaced x-coordinates between the extrema given above. The energy ranges for Halo orbits extend to both the Northern (L1N and L2N for L_1 and L_2) and Southern types (L1S and L2S for L_1 and L_2).

around the periodic orbit. These points are assigned an index n_{mnfd} for each periodic orbit, with $1 \leq n_{\text{mnfd}} \leq 360$. This gives 2888000 points across all periodic orbits.

The points that form initial conditions for the stable manifold are integrated backwards in time to a fixed $\pi/8$ section (Figure 4.3) about the Sun-Earth line using the 7th/8th order Dormand-Prince numerical integrator made available as part of the Boost ODEINT C++ package with absolute and relative tolerances of 10^{-13} . The $+\pi/8$ plane is used for manifolds belonging to orbits about L_2 and the $-\pi/8$ plane is used for manifolds belonging to orbits about L_1 . The rationale for this is as follows: in later stages of the optimisation, it is beneficial to be able to use simpler, analytic approximations to motion. The $\pm\pi/8$ plane was chosen in previous literature to bound conditions on the manifold away from regions where three-body dynamics dominates. At the $\pm\pi/8$ section and beyond (away from Earth), the dynamics can be assumed to be solely two body with no significant loss in accuracy [71]. It will be seen later the role this plane plays in simplifying the optimisation procedure by being able to assume analytic relationships between the initial and final states.

The intersection points are transformed from the synodic frame of the CR3BP into a Sun-centred inertial frame using the Sun-Earth (ecliptic) plane as the reference $x - y$ plane to simplify further analysis and allow easier interfacing with the data from the HORIZONS system.

4.2.3 Filtering Unsuitable NEOs

It is unnecessary and computationally infeasible to perform a full optimisation procedure on each of the 22406 NEOs included in the Minor Bodies Database. Therefore, the MBD is filtered using a prefilter to the optimal transfer cost, which is designed to be cheap to evaluate and which reasonably approximates the optimal transfer cost to rapidly reduce a list of 10^5 NEOs to a list of 10^2 retrieval candidates.

The earliest investigation into EROs used a prefilter based on the Jacobi constant of the NEO. Through the Tisserand parameter [186], one can use the orbital elements of the NEO to obtain

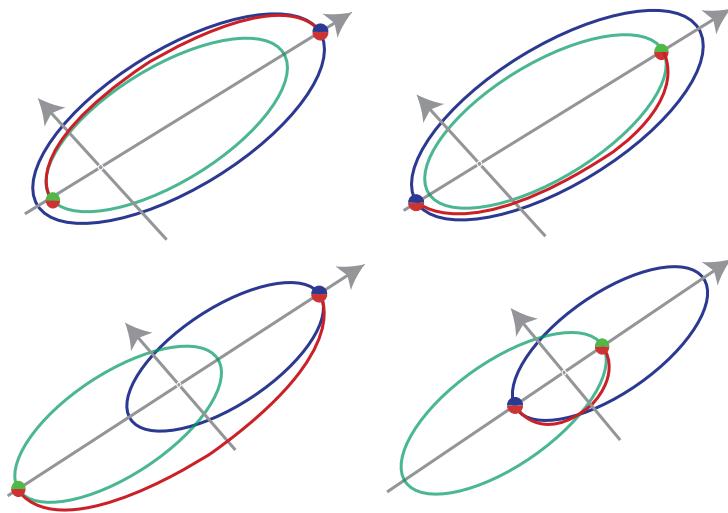


Figure 4.4: The Hohmann transfers considered in the pre-filtering methodology; the blue orbit is the orbit of the NEO; the green orbit is the orbit defined by the equivalent Keplerian elements of the manifold target point at the chosen fixed section; and the red arc is the Hohmann transfer. Since the orientation of the semimajor axes of the orbits is ignored, four possible combinations of transfer exist as one of the inclination change combinations may be discarded as more expensive.

an approximation to its Jacobi constant

$$\mathcal{J} \approx \frac{1}{a} + 2\sqrt{a(1-e^2)} \cos i \quad (4.7)$$

with the semi-major axis expressed in AU⁷. Any NEO whose Jacobi constant lies within the ranges of \mathcal{J} for which orbits were constructed previously is deemed to be a retrieval candidate.

However, this method does not consider the geometry of the NEO orbits, only their Jacobi constant. Orbit with very similar Jacobi constant may have very different orbital properties, making a transfer expensive. Thus, in Sánchez and García Yáñez [179], the prefilter was modified to use an idealised Hohmann transfer between the orbit of the NEO and each of the points in the target manifold data set. The equivalent position and velocity of the manifolds in the inertial frame is converted into its equivalent Keplerian elements at that point and used to calculate the Hohmann transfer cost. This implicitly assumes that the transfer happens at either aphelion or perihelion of both the NEO's and target point's orbits, ignoring any phasing between the two.

The Hohmann transfer is a method for transferring between two circular, co-planar orbits of radius r_1 and r_2 about a common central body in two-body dynamics. The transfer comprises two impulses, transferring on to and off of an intermediate transfer arc (Figure 4.5). The Δv required to perform the Hohmann transfer may be derived from the conservation of energy in a two-body system [17]. Using an energy balance equation (the *vis-viva* equation) on the initial and final orbits, and substituting in expressions for the semi-major axis of the initial a_0 , intermediate

⁷This formulation is originally due to Kresák [187] and is a special case for which $\mu \ll 1$. A derivation is presented in Appendix A.

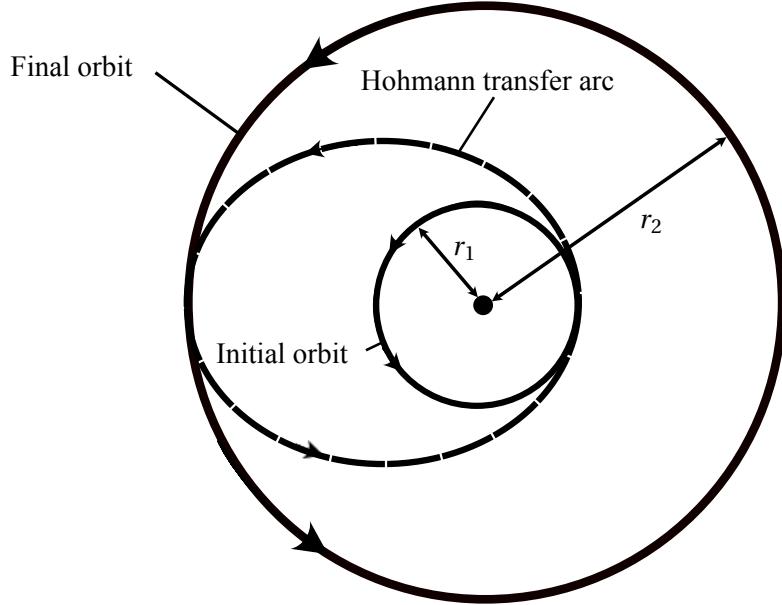


Figure 4.5: Schematic of the Hohmann transfer between two circular, coplanar orbits. The initial orbit has radius r_1 and the final orbit radius r_2 .

a_{int} , and final orbits a_f , leads to the relation for the two impulses

$$\Delta v_1 = \sqrt{GM\left(\frac{2}{r_1} - \frac{1}{a_{\text{int}}}\right)} - \sqrt{GM\left(\frac{2}{r_1} - \frac{1}{a_0}\right)} \quad (4.8)$$

$$\Delta v_2 = \sqrt{GM\left(\frac{2}{r_2} - \frac{1}{a_f}\right)} - \sqrt{GM\left(\frac{2}{r_2} - \frac{1}{a_{\text{int}}}\right)} \quad (4.9)$$

$$a_{\text{int}} = \frac{r_1 + r_2}{2} \quad (4.10)$$

where GM is the standard gravitational parameter of the host body, in this case the Sun. Since these orbits are not necessarily co-planar, differences in inclination between the two orbits are removed using an impulse which performs an instantaneous, pure rotation of the velocity vector of the orbit at the aphelion or perihelion of the transfer orbit (the orbital speed before and after the manoeuvre is unchanged). The pure rotation is assumed to take place only once – i.e. the inclination change manoeuvre is not split between aphelion or perihelion – and is performed such that the total capture Δv is minimised. Considering the geometry of the initial and final velocities of NEO at the relevant point of inclination change (Figure 4.6), the angle Δi can be inferred to be

$$\sin \frac{\Delta i}{2} = \frac{\Delta v_i/2}{v} \quad (4.11)$$

and knowing that the velocity change occurs at either apoapsis or periapsis only, one can rearrange to find

$$\Delta v_i = 2\sqrt{\frac{GM}{a}} r^* \sin(\Delta i/2)$$

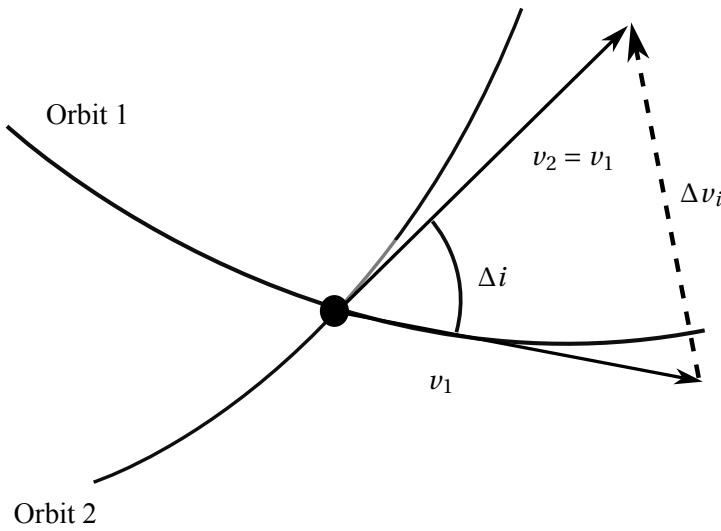


Figure 4.6: Geometry of the inclination change problem between two orbits with the same orbital speed. The inclination change is assumed to be a pure rotation about a common point on the two orbits.

with a the semi-major axis of the orbit at the point the transfer is made, and i_0 and i_f the initial and final inclinations, respectively. The variable r^* is

$$r^* = \begin{cases} r_a/r_p & \text{if the inclination change is at perihelion} \\ r_p/r_a & \text{if the inclination change is at aphelion} \end{cases} \quad (4.12)$$

with r_a and r_p the aphelion and perihelion distance for the given orbit, respectively.

The total cost of the transfer is then computed using

$$\Delta v_t = \sqrt{\Delta v_{a,1}^2 + \Delta v_{i,1}^2} + \sqrt{\Delta v_{a,2}^2 + \Delta v_{i,2}^2}.$$

The inclination change is only performed once, such that one of $\Delta v_{i,1}$, $\Delta v_{i,2}$ is zero. In total, there are 8 possible combinations of transfer: 4 combinations of orienting the semi-major axes of the NEO and the target manifold point, and 2 combinations per orientation corresponding to whether the inclination change occurs at the beginning or end of the intermediate transfer arc. However, one of the locations of the inclination changes can always be identified as being more expensive based on the radial distance to the Sun. This leaves only four transfers per NEO and target manifold to be considered, as shown in Figure 4.4. In total, this means that for each of the 2888000 target points generated previously, 4 sets of transfers must be calculated for a total of 11552000 trial computations.

At this point, the literature is extended. Since the Hohmann transfer approximation depends on the radius, semimajor axis and inclination of both the osculating NEO orbit and the target manifold point, there is the possibility that the pre-filter output may vary with time as the NEOs

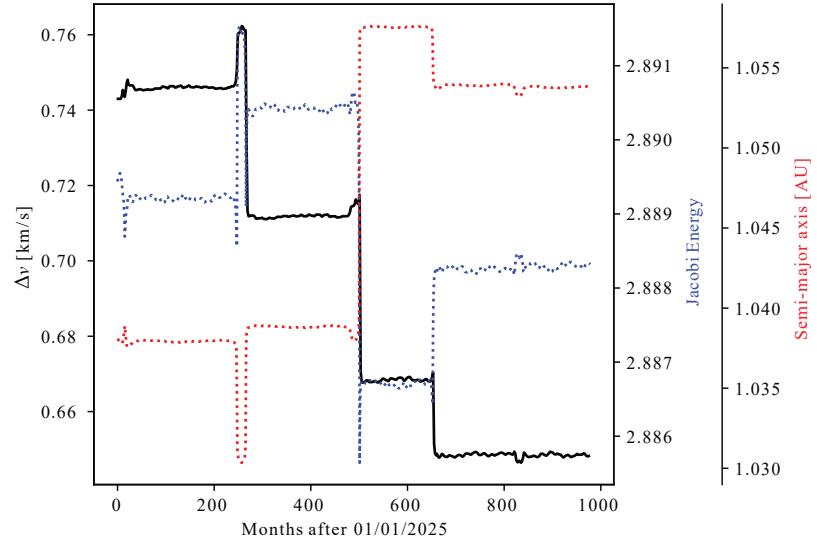


Figure 4.7: Pre-filter cost (black), Jacobi energy (blue) and semi-major axis (red) for asteroid 3410533. Preliminary analysis shows the potential for pre-filter cost to vary with time as a result of changing orbital elements from close approaches to planets, for example. In this case the asteroid does not pass the pre-filter when run after a certain date. Even filters used in prior works and based on the Jacobi energy (blue) would vary with time.

do not move on purely Keplerian orbits⁸. A preliminary analysis that shows the occurrence of this time-dependency for asteroid 3410533 is presented in Figure 4.7, where the semi-major axis of the orbit of the asteroid changes as a result of close approaches with the Earth with time. This changes the estimated pre-filter cost when using the Hohmann transfer approximation. The Jacobi constant is also changed, which means this phenomenon also occurs for the previous pre-filters based on Jacobi constant. The pre-filter is thus repeated once every 28 days in the time-span January 1, 2025 and January 1, 2100, with 28 days selected arbitrarily. The time ranges were chosen to match those used in the optimisation procedure (section 4.2.4). This requires 978 evaluations of the pre-filter for a total of 11297856 000 transfers to be constructed. While this greatly increases computation time, it is necessary to ensure no NEOs are excluded prematurely.

The pre-filter is implemented in Fortran and interfaced with the SPICE library [189] to access the high-precision ephemerides. To reduce computation time – particularly for when repeating the pre-filter at regular intervals – parallelisation is implemented via the MPI and OpenMP libraries⁹. The program is written to support vectorisation of many of quantities automatically on relevant processors. Following the vectorisation efforts, approximately one week of CPU time is required to run the asteroid pre-filter at an interval of 28 days on 2.0 GHz Intel Xeon E5-2670 processors. By use of 120 cores, the pre-filter can be computed in 1.5 days of wall time.

⁸While not considered in their original paper, this time dependency is also acknowledged by one of the original authors of Sánchez and García Yáñez [179] in a later publication [188].

⁹The small time required to evaluate each transfer and the need to evaluate the solution in double-precision meant that GPU parallelism did not provide a major advantage over CPU parallelism.

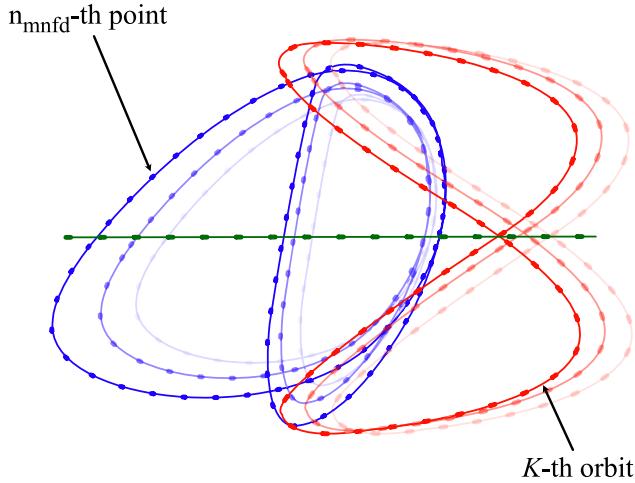


Figure 4.8: Intersections of stable manifolds and the $\pi/8$ plane for selected periodic orbits from the planar Lyapunov family (green), Northern Halo family (blue), and vertical Lyapunov family (red). Design variables $n_{\text{mnfd}} \in [1, 360]$ and $K \in [1, 8000]$ are indicated.

Previous literature used 700 m/s as the threshold for the idealised Hohmann transfer to designate a NEO as a retrieval candidate. However, this chapter uses a higher threshold of 3 km/s to classify a NEO as a retrieval candidate to ensure that no EROs are excluded prematurely. This chapter also uses the ephemeris files rather than the MBD. An analysis into the behaviour of the prefilter for various transfer thresholds is presented along with the results of applying the prefilter in section 4.4.

4.2.4 Finding optimal transfers

Once the list of retrieval candidates is known, an optimisation procedure is performed to determine whether any transfers exist between the retrieval candidate and any point in any target manifold below the ERO threshold.

By leveraging that the dynamics outside of the $\pm\pi/8$ plane can be assumed to be Keplerian, a Lambert arc is used to model the transfer between a point on the NEO's orbit and a point in the target manifold. A Lambert arc constructs an equivalent orbit that passes through two known positions with a known time to move from the first position to the next. The method essentially solves a boundary value problem for the differential equation that describes Keplerian motion, and uses a two impulses to perform the transfer: one impulse at the beginning of the arc, Δv_{start} , and one at the end of the Lambert arc Δv_{end} . For more information on the derivation and uses of Lambert's problem, the reader is referred to Curtis [17]. To ensure correctness in the assumption of Keplerian dynamics, the transfers are constructed only to-and-from points which lie outside of the $\pm\pi/8$ plane discussed previously: ensuring that simpler and quicker to evaluate two-body transfers can be used is the motivation behind the use of this plane.

To parameterise the transfer for use in an optimisation routine, five variables are used (Table 4.3). The epoch of the first impulsive manoeuvre t_0 determines the position of the NEO from

its ephemeris, and the transfer duration t_t determines the length of the Lambert arc. The three remaining variables t_{end} , n_{mnfd} and K determine the target (final) point: t_{end} is the time in non-dimensional, CR3BP units to arrive at the $\pm\pi/8$ plane from the target point. The variable n_{mnfd} and K carry the same meaning as in section 4.2.2, and identify a point in a target manifold on a target periodic orbit (Figure 4.8).

To improve the possible set of final orbits, provided the underlying dataset is dense enough one can interpolate between points on adjacent manifolds and across adjacent periodic orbits to greatly increase the number of possible target states [185]. For this purpose, a cubic B-spline is used whenever the optimiser selects non-integer values of n_{mnfd} or K to determine a local approximation to the manifold at that point on the relevant $\pm\pi/8$ plane. To avoid errors in the numerical approximation, the interpolant is not allowed to span two different types of orbital families. Non-zero values of t_{end} are numerically integrated backwards from the pre-computed points on the $\pm\pi/8$ plane ‘on-demand’ to further increase the set of possible points. Since there is no guarantee that each orbital transfer now targets an exact discretisation of the manifold in the dataset, all transfers are manually validated in a separate Python script.

The constraints on these variables are given in Table 4.3. The constraint on t_0 was used so it corresponds to the same bounds used in the previous literature, although the lower bound here is raised from Jan 1, 2020, for practical reasons. The constraint on t_t was chosen to be any valid transfer length up to the longest transfer found in other literature, at 1500 days. Similarly, t_{end} was constrained to be approximately equal, in non-dimensional units, to the longest ‘coast’ phase after inserting onto the manifold found in other literature. An additional constraint was placed on t_0 in the previous literature to reduce the search space and computational cost by limiting t_0 to the first astronomical synodic period of the NEO and the Earth. This constraint is lifted here to promote exploration of the full parameter space, at the cost of increased computation time.

Two objectives are included in the optimisation problem: the first is the total Δv cost of the two-impulse Lambert transfer, which was considered in the literature and is required to designate a NEO as an ERO. The transfer time t_t is also considered, which is an extension to the literature and aims to answer questions regarding the robustness of the retrieval trajectory with changing mission plans. Including the transfer time allows a trade-off between transfer time *versus* cost, as shown in the Pareto fronts in section 4.5.

The optimisation problem is solved using the off-the-shelf software MIDACO (Mixed-Integer Distributed Ant Colony Optimiser) [190]. MIDACO can perform both global and local optimisation by altering the update parameters for the solver. MIDACO was selected after testing its performance against other off-the-shelf optimisers. The Boender-Rinnooy-Stougie-Timmer (BRST) optimisation algorithm uses a stochastic search in combination with sampling, clustering, and local searches, and is particularly effective for objective functions where the gradient is unavailable or otherwise difficult to compute [191]. A range of optimisers made available in the C++ PaGMO package [192] were also tested, including another ant colony optimisation

Design variable	Definition	Constraint
t_0	The epoch of the first manoeuvre of the transfer.	$2025/01/01 \leq t_0 \leq 2100/01/01$
t_t	Duration of the Lambert arc, in days.	$0 \leq t_t \leq 1500$ days.
t_{end}	Time in non-dimensional units to arrive at the $\pm\pi/8$ plane.	$-25 \leq t_{\text{end}} \leq 0$
n_{mnfd}	Identifies a trajectory in the stable manifold of the periodic orbit.	$1 \leq n_{\text{mnfd}} \leq 360$
K	Identifies a periodic orbit.	$1 \leq K \leq 8000$

Table 4.3: Definitions of the five design variables used in the optimisation procedure. Taken together, n_{mnfd} and t_{end} uniquely specify the insertion point in the stable manifold of the K^{th} target orbit. Dates are given in the form YYYY/MM/DD.

implementation and Particle Swarm Optimisation (PSO) algorithm. MIDACO was found to outperform all other algorithms tested when applied to this problem, although the PSO algorithm from the PaGMO package was used to ensure that the solutions found by MIDACO were at least local minima.

The optimisation is very computationally intensive, and requires dedicated computing facilities to complete in reasonable time. OpenMP and MPI are again used to reduce computation times. The cubic B-spline interpolants are evaluated using the Fortran B-spline library [193], and the Lambert arcs are computed using the Fortran Astrodynamics Toolkit [194]. Numerical integration is performed with the 7th/8th-order Dormand-Prince numerical integration scheme provided by the ODEINT C++ library used previously, again interfaced to Fortran. The relative and absolute tolerances in the integration are set to 10^{-13} . Approximately 2.5 days of CPU time is required to run the initial global search on a 2.0 GHz Intel Xeon E5-2670 processor for all the retrieval candidates found by the pre-filter. This corresponds to approximately 10 minutes of global optimisation per candidate, which was found to be sufficient for MIDACO to identify a suitably close solution to the global minimum.

4.3 Implementation ‘tips’

This section covers some additional implementation ‘tips’ for many of the concepts introduced previously.

The software developed for this investigation is written in mixed Fortran and C. The latter language is used for the MIDACO C optimiser and to interface with the Boost ODEINT library for numerical integration. NASA’s SPICE library is used to interface with the high-precision ephemerides.

To reduce computation time, the code is written with support for the OpenMP and MPI parallel programming paradigms, along with additional support for vectorisation on Intel CPUs which

is specialised by the compiler for the processor architecture being used. I found the pre-filter to be best parallelised by assigning each worker (core) a single NEO to estimate the transfer cost of at any time. This ensures that the database of target points, which is very large in memory, is traversed in an order amenable to keeping some data in CPU caches. It also makes the problem trivially parallelisable. The pre-filter is terminated at the first instance of an object passing below the 3km/s threshold set to avoid unnecessary computations. Since this can lead to unbalanced workloads, it is best to allow cores to ‘steal work’ from others by keeping a central pool of NEOs that are still to be evaluated which can be sampled from at any time. This minimises the number of workers idling waiting for others to finish unnecessarily.

The optimisation procedure is numerically more intensive and requires more care in its implementation. Several options for parallelism are available at this point. MIDACO includes parallelism inside the optimiser, automatically evaluating many different calls to the objective function in parallel. However, since the objective function is relatively cheap to compute, the overheads associated with the data sharing to do this dominate compute time. Instead, for maximum performance one should spawn a MIDACO process on each worker, and evaluate the optimisation in serial on each worker.

4.4 Results

The pre-filter identifies 792 retrieval candidates when applied to the NEO catalogue. The number of candidates that are obtained for a given $\Delta\nu$ threshold can be seen in Figure 4.9: while the number of candidates increases with pre-filter threshold, there are diminishing returns in terms of the number of EROs found. No additional EROs are found between thresholds of 3 km/s and 5 km/s, which motivated the use of 3 km/s as the limit for this work.

Only 46 of these candidates (5.8%) are retained when lowering the threshold to 0.7 km/s as in [179], an increase of 13 relative to their result. This increase is to be expected, since the NEO database has grown in size by approximately 40% since that study. Discarding all candidates discovered after the publication date yields 36 candidates, an increase of only 3.

Five candidates that pass the pre-filter in the previous literature – 3332535, 3339082, 3520667, 3634612, and 3719226 – are not recovered here, even when running the pre-filter specifically at the date of the previous study. These missing candidates do not become EROs in the previous literature, so this does not affect the resulting list of EROs.

The time-dependency described previously yields an additional two candidates. This shows that while a more comprehensive list of candidates can be achieved by considering the changing NEO orbits over time, the significant increase in computing time may not justify the result: a reasonable approximation to the full list of candidates can be found by running the pre-filter just once.

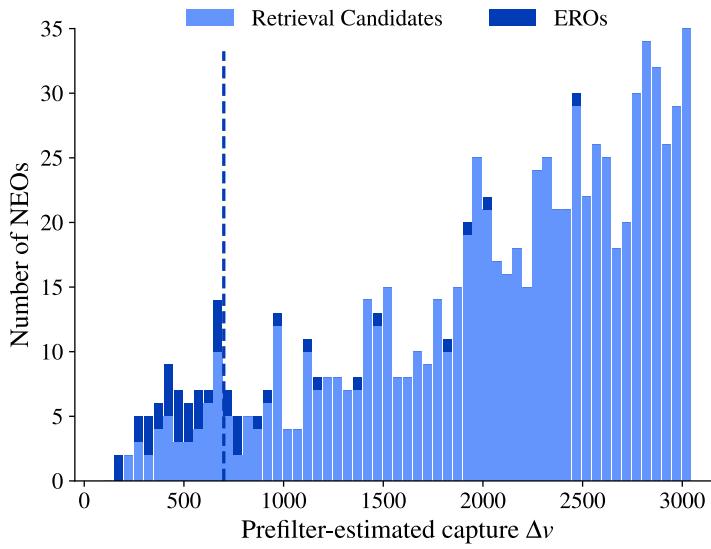


Figure 4.9: Number of retrieval candidates (light blue) and EROs (dark blue) for a given pre-filter Δv threshold. While 66% of EROs are found for 700 m/s (dashed line), a higher threshold of 3000 m/s ensures no EROs are excluded prematurely, albeit at significant computational cost in the pre-filter.

Of the 792 retrieval candidates, 44 go on to become EROs after the full optimisation procedure. The details of the Δv -optimal transfers for each ERO identified are given in Table 4.4. The lowest Δv found is 13.97 m/s, and the highest 445.35 m/s.

With the exception of a single ERO, 2011 BL45, the full list of EROs in Sánchez and García Yáñez [179] are reproduced, with an average retrieval Δv improvement of 176.7 m/s ($\sigma = 147.63$ m/s). This investigation also finds 17 new capture trajectories with a Δv of less than 100 m/s. Such objects may be interesting for future study, as they represent significant improvements over the previously-known capture Δv .

Why 2011 BL45 could not be recovered as an ERO (lowest $\Delta v = 636$ m/s) is not immediately clear, but could be explained by revisions to the orbital data of the NEO since the previous study.

Only 31 of these EROs are recoverable with a pre-filter threshold of 0.7 km/s, which suggests that the threshold used in previous literature limits the final number of EROs by 34%. Note that this still represents an improvement over previous literature of 14 EROs.

The additional constraint was placed on t_0 in the previous literature to reduce the search space and computational cost by limiting t_0 to the first astronomical synodic period of the NEO and the Earth was found to increase the optimal Δv by approximately 50 m/s, as the behaviour of the NEO is not perfectly periodic between astronomical synodic periods.

SPK ID	$\Delta v[m/s]$	Transfer epoch	Transfer end	Arrival at $\pm \frac{\pi}{8}$ plane	Orbit
3405338	13.97	2083/08/31 16:46:35	2086/11/12 16:46:35	2089/06/27 14:38:03	L1S
3843865	15.25	2093/06/22 11:37:34	2096/06/08 11:37:34	2096/08/01 04:44:51	L2S
3733264	23.68	2037/10/07 12:46:09	2040/08/23 12:46:09	2044/02/08 23:50:25	L1V
3698849	23.89	2091/05/27 20:52:28	2095/07/05 20:52:28	2096/02/06 05:44:02	L1S
3054374	27.39	2029/02/04 02:42:50	2030/10/05 02:42:50	2032/08/04 16:14:46	L1N
3315004	27.55	2074/09/30 11:54:17	2075/03/22 11:54:17	2078/05/17 22:15:21	L1S
3843512	33.06	2087/07/16 05:05:26	2090/04/14 05:05:26	2092/06/05 09:31:46	L1V
3076774	41.60	2091/09/26 03:41:05	2093/06/06 03:41:05	2095/07/07 12:39:19	L1V
3435539	42.23	2029/06/16 05:53:13	2030/07/20 05:53:13	2030/08/04 15:16:37	L1V
3405189	50.94	2087/10/22 15:41:02	2091/06/26 15:41:02	2093/01/24 07:46:57	L1N
3836857	51.52	2049/01/24 00:51:25	2051/10/17 00:51:25	2052/11/10 00:13:14	L2V
3702319	56.93	2080/12/13 23:13:48	2084/01/05 23:13:48	2085/03/28 03:06:42	L1V
3403148	63.83	2025/03/19 08:05:03	2027/03/11 08:05:03	2027/03/13 20:43:40	L2N
3696401	69.84	2073/11/14 22:32:44	2076/03/14 22:32:44	2079/12/15 16:03:31	L1N
3840784	81.52	2083/02/05 15:22:51	2086/04/24 15:22:51	2089/10/10 05:13:44	L1V
3759358	89.76	2048/06/09 21:15:26	2051/07/02 21:15:26	2052/04/18 09:36:45	L2N
3610175	90.82	2097/06/19 02:55:06	2100/12/02 02:55:06	2103/11/20 14:56:24	L2V
3893945	94.96	2099/10/21 08:33:59	2103/02/09 08:33:59	2106/12/15 20:49:44	L1N
3726012	108.25	2047/04/27 02:52:56	2051/01/14 02:52:56	2053/06/13 02:13:48	L1N
3582145	111.98	2039/11/11 04:56:22	2040/07/18 04:56:22	2040/07/18 08:34:39	L1V
3991650	112.27	2073/12/06 05:09:00	2076/11/03 05:09:00	2079/12/29 17:58:33	L2P
3697803	115.13	2038/07/11 09:56:29	2040/03/31 09:56:29	2041/03/09 15:04:05	L1V
3989308	116.72	2082/08/03 00:48:08	2086/07/09 00:48:08	2088/11/26 02:02:44	L2V
3735181	123.71	2064/12/07 21:54:01	2068/09/14 21:54:01	2071/12/06 08:15:26	L1P
3879383	128.16	2051/07/27 05:54:09	2053/12/07 05:54:09	2057/07/19 05:53:24	L1N
3444297	136.52	2074/07/15 01:23:25	2076/10/29 01:23:25	2080/04/03 12:48:42	L1N
3618493	140.38	2085/10/31 13:57:51	2089/08/16 13:57:51	2090/11/08 03:07:06	L1S
3519518	140.95	2057/09/13 00:43:54	2060/12/01 00:43:54	2061/09/02 11:08:22	L1P
3549639	146.66	2092/11/17 04:32:51	2095/12/09 04:32:51	2098/08/30 16:03:43	L2V
3005816	162.60	2084/12/21 20:46:39	2085/07/06 20:46:39	2086/02/01 10:31:22	L1V
3390109	166.09	2090/01/13 14:48:27	2091/02/27 14:48:27	2092/02/19 22:39:48	L2V
3556127	169.01	2073/01/15 19:47:19	2073/10/30 19:47:44	2073/11/24 15:04:24	L2V
3825158	170.93	2095/05/25 09:19:32	2099/04/20 09:19:32	2102/08/26 09:36:23	L1P
3551168	180.83	2035/02/16 23:00:39	2036/08/25 23:00:39	2037/05/30 06:05:40	L2V
3550232	195.40	2059/10/03 05:07:28	2062/03/25 05:07:28	2065/01/24 03:42:44	L1N
3625129	203.01	2090/12/27 06:25:33	2092/02/26 06:25:33	2092/05/11 15:40:21	L1S
3781983	245.83	2099/06/01 04:29:49	2102/07/12 15:22:23	2105/07/20 03:40:37	L1P
3826850	254.37	2069/08/12 23:03:50	2071/04/19 23:03:50	2074/05/24 00:54:36	L2V
3745994	301.71	2030/01/29 13:28:25	2033/10/17 13:28:25	2035/07/21 02:59:40	L2S
3781977	305.79	2072/12/16 21:24:12	2073/07/23 21:24:12	2075/01/11 19:26:12	L2V
3748467	332.02	2045/07/11 09:31:04	2047/07/20 09:31:04	2049/07/17 15:15:49	L1P
3648046	355.29	2041/11/26 12:08:11	2045/09/23 12:08:11	2049/04/05 01:23:13	L2V
3568303	418.53	2063/03/19 15:57:36	2065/12/08 15:57:36	2067/12/16 17:26:46	L2V
3719859	445.35	2036/03/14 17:07:21	2039/03/02 17:07:21	2040/11/08 13:58:20	L2N

Table 4.4: List of Δv -optimal solutions found for all of the EROs in this work; I find 17 new trajectories with retrieval costs under 100 m/s. Note that while these transfers may be optimal in terms of Δv , many more solutions exist for other transfer times with similar retrieval velocities.

All dates are given in the form YYYY/MM/DD HH:MM:SS.

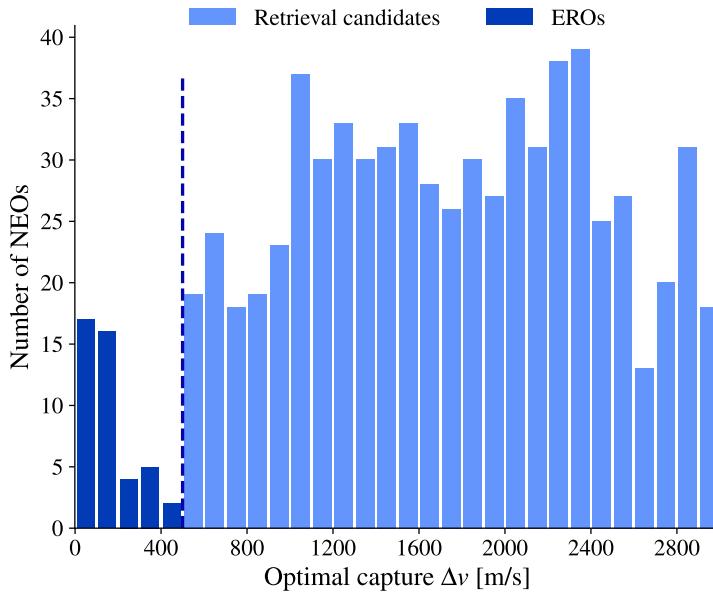


Figure 4.10: Distribution of retrievable NEOs over the range of optimal capture Δv . The dark blue bars correspond to EROs, but many more NEOs have very similar capture Δv . EROs are subject to more rigorous optimisation, causing the gap between dark blue and light blue bars.

Attention must be paid to the 500 m/s threshold used for defining a candidate, which was originally selected in [71]. There is no obvious engineering constraint to set the ERO threshold at that value. Figure 4.10 shows the number of NEOs in discrete bins of optimal capture Δv . By increasing the threshold by as little as 20%, the number of EROs increases by 50%. Note that in Figure 4.10, the dark blue region corresponds to EROs subjected to the more rigorous optimisation procedure to generate the Pareto fronts. This explains why these solutions exhibit much lower Δv , and the apparent gap in the distribution. It is expected that the solutions in the light-blue regions would similarly improve when subjected to the same, more rigorous optimisation procedure.

Instead of using an arbitrary fixed cutoff, a mission-by-mission approach could be used. The maximum retrieval Δv for a given mission profile may be used in combination with Figures 4.9–4.10 to identify the population of EROs for that mission and available computational resources. Such an approach would allow a more flexible and practical definition of an ERO: transfer costs of 499 m/s and 501 m/s are functionally the same, but the latter is deemed unretrievable by using a fixed cutoff.

All identified transfers are validated in the full CR3BP model through another local optimisation, including the influence of the dynamics of the third body. There is little difference ($\mathcal{O}(m/s)$) between the Keplerian approximation used in the optimisation and the full CR3BP transfers, further validating the choice of the $\pm\pi/8$ planes.

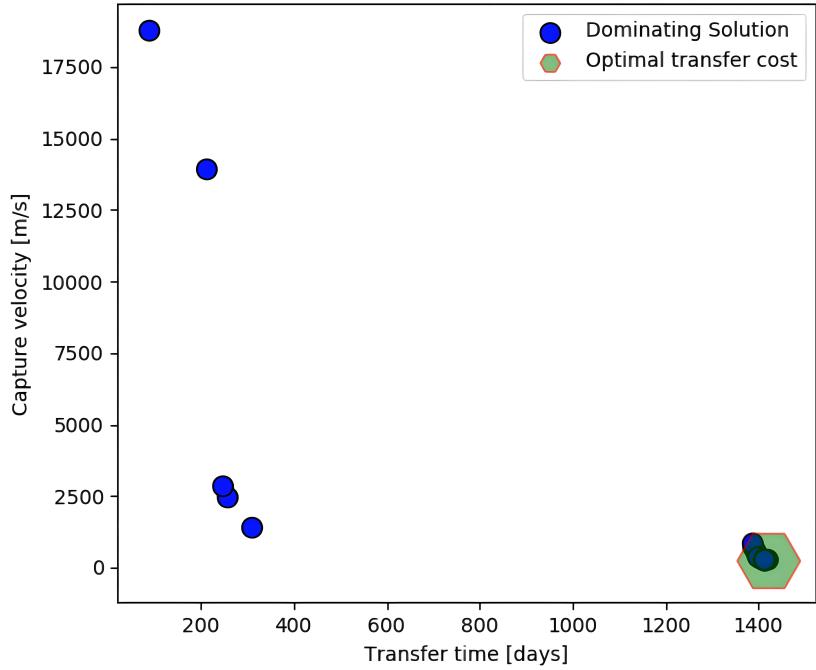


Figure 4.11: Pareto front for 2011 BL45 computed by MIDACO’s in-built Pareto front generation and MIDACO’s PlotTool. Only the points that define the front are included, and intermediate information on the dominating solutions is lost.

4.5 Pareto Front Generation

The previous sections of this chapter outlined the identification of single, $\Delta\nu$ -optimal retrieval trajectories. However, such trajectories are not of great practical use in space mission design. Firstly, there may be mission design requirements that constrain the maximum time between trajectories, such as solar incidence, the availability of on-board propellant over time, or even mission urgency in the case of planetary protection. Secondly, a $\Delta\nu$ -optimal may not be as robust to perturbations in either the manoeuvres or general control of the asteroid.

Recall that the case of asteroid retrieval trajectory control feasibility and sensitivity has been assessed in detail in the literature [73, 74, 164, 188]. What has not been considered, however, is the flexibility of such transfers. The remaining content of this chapter investigates how flexible these EROs are for practical mission design purposes by studying the Pareto frontier of the retrieval trajectories, trading off transfer time t_t and transfer cost. A series of strategies for searching for low- $\Delta\nu$ transfers in the space of t_t were analysed and are discussed in the following. The results from the final strategy are also analysed in detail.

4.5.1 Naïve MIDACO

The MIDACO optimiser used to identify optimal- $\Delta\nu$ transfers earlier in the chapter stores all the dominant solutions it finds during the optimisation process, and this can be extracted after the optimisation is complete. The number of points collected, and how points are collected, can be

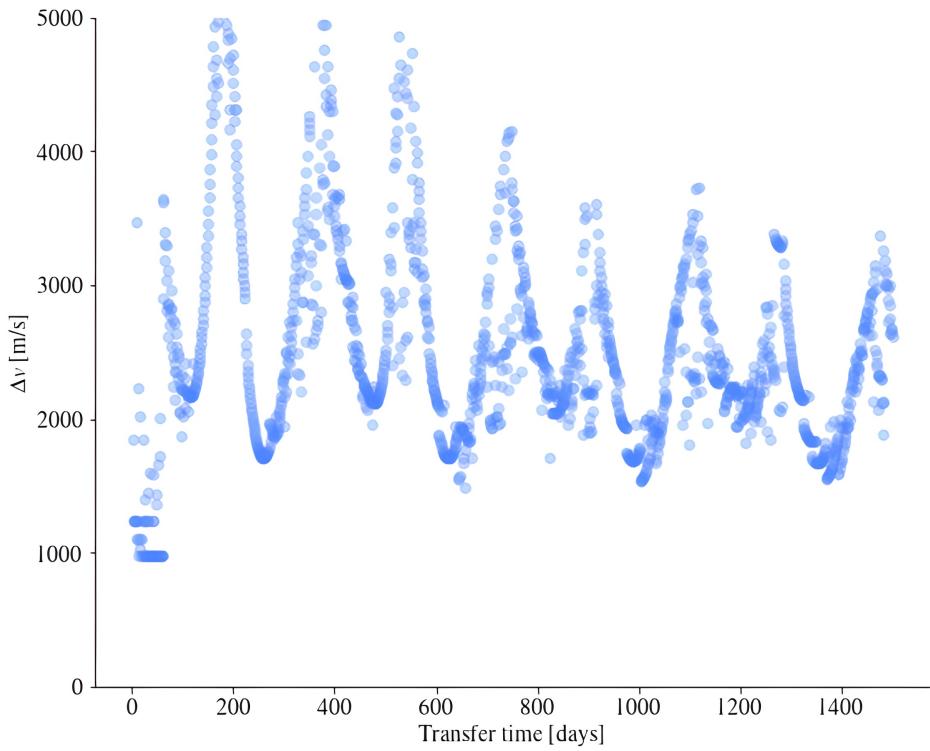


Figure 4.12: Preliminary transfer time t_t vs. capture Δv space for 2011 BL45, found by continuing the solution at each value of t_t from the optimal solution at the previous t_t . This produces an approximately periodic cost across the range of t_t , with large differences between the lowest and highest capture Δv s.

adjusted within the solver. MIDACO also provides PlotTool, a tool for plotting the Pareto fronts of the dominating solutions that MIDACO finds [195]. An example Pareto front for asteroid 2011 BL45 (SPK ID 3556125) is given in Figure 4.11. However, the dominating solutions returned are clustered around the lower and upper end of the t_t range. While this is similar to the Pareto front for transfer times larger than about 200 days as the points between the solutions in Figure 4.11 are dominated, the transfer cost for shorter transfers is vastly overestimated in these results, as will be shown later. The Figure also provides no insight as to why the Pareto front has the resulting structure.

To combat this, denser Pareto fronts are generated using a series of successive optimisations with the transfer time t_t fixed to an integer number of days between 0 and 1500, with all the other parameters introduced previously (t_0 , t_{end} , n_{mnfd} , K) left free. Since all other parameters are left free, transfers for subsequent values of t_t may not be just continuations of each other, but instead spread widely around the parameter space. Careful implementation of the optimisation method is required to find the solutions across the full range of parameters.

Several techniques are used to alter the structure of the Pareto fronts, and are discussed here with the example of 2011 BL45, which is not actually an ERO but exhibits the same behaviour. This is discussed in more detail later.

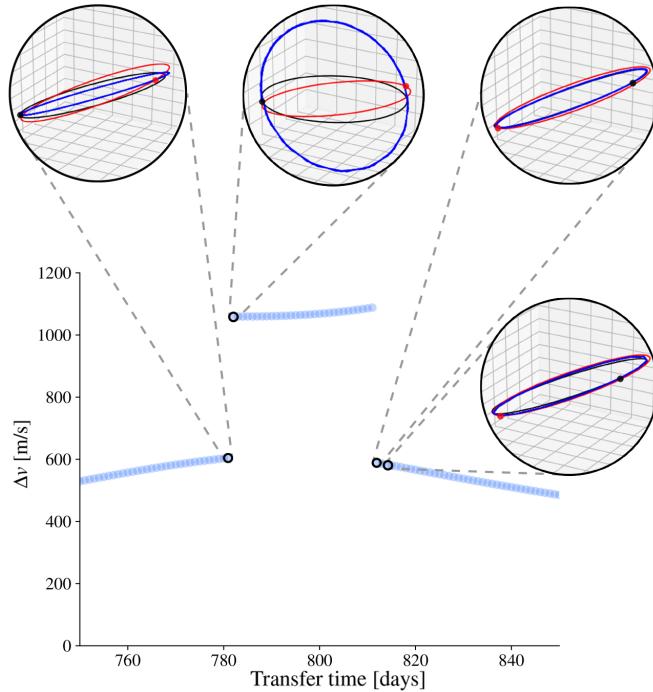


Figure 4.13: Preliminary structure of the Δv vs. transfer time space for asteroid 2009 BD, when only the optimal solution at the previous transfer time is used as an initial guess for the parameters in the optimisation. The orbit of the NEO is in black, in red is the equivalent orbit of the manifold target point, and in blue is the transfer orbit between those points.

4.5.2 Naïve scanning

The first attempt at Pareto front generation simply continued the search at each successive value of t_t from the optimal solution at the previous value of t_t . On the first iteration, the optimal solution found in the initial global search was used. The Pareto fronts are then ‘grown’ in two directions – increasing and decreasing transfer time – from this optimal solution. MIDACO was also configured such that its update rules are set to prioritise the currently best-known solution.

The results from this initial method are given in Figure 4.12, which displays a range of solutions between approximately 1000 m/s and 5000 m/s. The solution repeats almost periodically approximately every 200 days, or half an orbital period of the NEO. However, the transfer costs between different values of t_t are *a priori* expected to be similar, as the dense pre-computed grid of manifolds should contain a range of nearby points for which transfers exist with comparable Δv .

This is an issue with the use of a two-impulse Lambert arc in the optimiser cost function. As the transfer angle of the Lambert arc approaches ± 180 degrees, the Lambert arc approaches a singularity, where small changes in the start and end point of the arc yield very large changes in its inclination, potentially altering the cost of the transfer dramatically. This is shown graphically in Figure 4.13, which shows a preliminary Pareto front for asteroid 2009 BD (SPK ID 3444297) with the transfers themselves shown in blue in insets.

For 2011 BL45, this instability causes the optimiser to discover a new family of Lambert arcs at approximately $t_t = 100$ days, which are themselves relatively unstable. These solutions are then continued down to a transfer time of zero days. This instability can be avoided by seeding the optimiser with multiple initial guesses in order to examine regions away from the instability, and potentially through use of a three-impulse transfer.

4.5.3 Global scanning

To prevent premature convergence to local minima, and to avoid instabilities caused by the Lambert transfer, multiple initial guesses for each t_t are now included to promote exploration of the parameter space by the optimiser. The guesses and subsequent sets of design variables are evaluated in parallel, such that the total number of guesses and design variable sets is equal to the total number of cores.

One of these guesses is set to the previous optimal solution at the previous time-step (or the optimal solution from the initial global search if this is unavailable). The remaining guesses are initialised randomly throughout the parameter space. Subsequent sets of design variables are also evaluated simultaneously. The parameters for MIDACO are adjusted to force the solver to be less ‘greedy’, and place less emphasis on the currently best-known solution.

Figure 4.14 shows the resulting Pareto front, again for asteroid 2011 BL45, using this modified strategy of including multiple initial guesses. Solutions are included with transfer costs between 625 m/s and 2650 m/s. With this strategy, there is a pronounced ‘floor’ in the Pareto front around the optimal capture Δv of approximately 625 m/s. However, there is also significant variance in the solutions identified, and solutions close to each other in terms of transfer time are often very different in both transfer cost and the values of the design variables.

This is a result of the optimiser no longer converging at every time step due to the larger search space. In this figure, each value of t_t is generated using 60 seconds of wall-time on 48 cores, or approximately 70×10^3 function evaluations. Further analysis shows that approximately 500×10^3 function evaluations gives convergence for the majority of candidates. The strategy was modified to run for 10 minutes per value of t_t . While MIDACO does provide the ability to terminate based upon an estimation of convergence, it was found to be simpler to interface the codes with the IRIDIS HPC scheduling system when fixed times were used.

4.5.4 Fully converged global scan

The final Pareto fronts for asteroids 2011 BL45 and 2012 TF79 (SPK ID 3610175) are given as representative examples in Figures 4.15 and 4.16, respectively. 2012 TF79 admits an optimal capture Δv of 90.82 m/s, but is actually retrievable for any transfer time between 0 and 1500 days, including a zero-time, single-impulse transfer. In fact, 34 of the 44 EROs identified remain EROs for a zero-time, single impulse transfer, and these remain generally low-cost solutions.

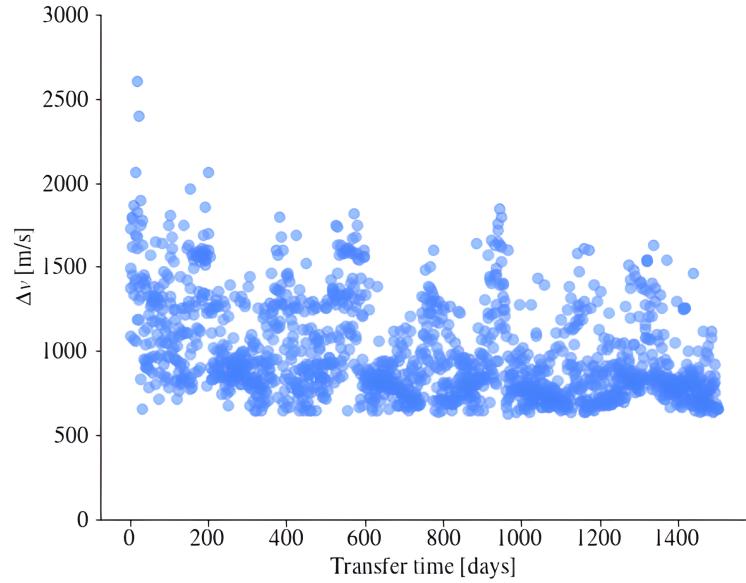


Figure 4.14: Pareto front for 2011 BL45 when generated using multiple initial guesses for each value of t_t . There is a pronounced ‘floor’ in the optimisation but transfers with similar t_t can have significantly different Δv .

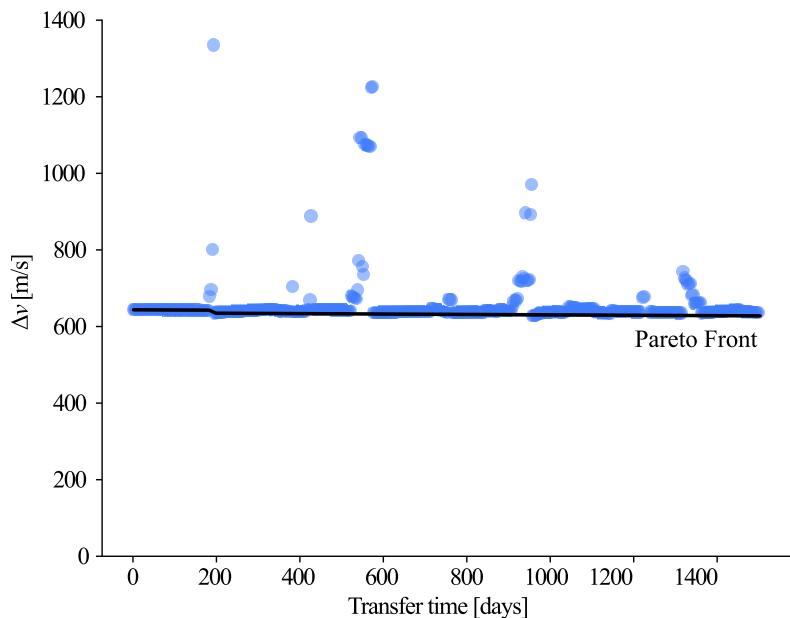


Figure 4.15: Pareto front for NEO 2011 BL45. While not an ERO, with an optimal capture Δv of approximately 625 m/s, it still displays the same behaviour. Non-EROs may, therefore, offer similar flexibility to the mission designer as EROs.

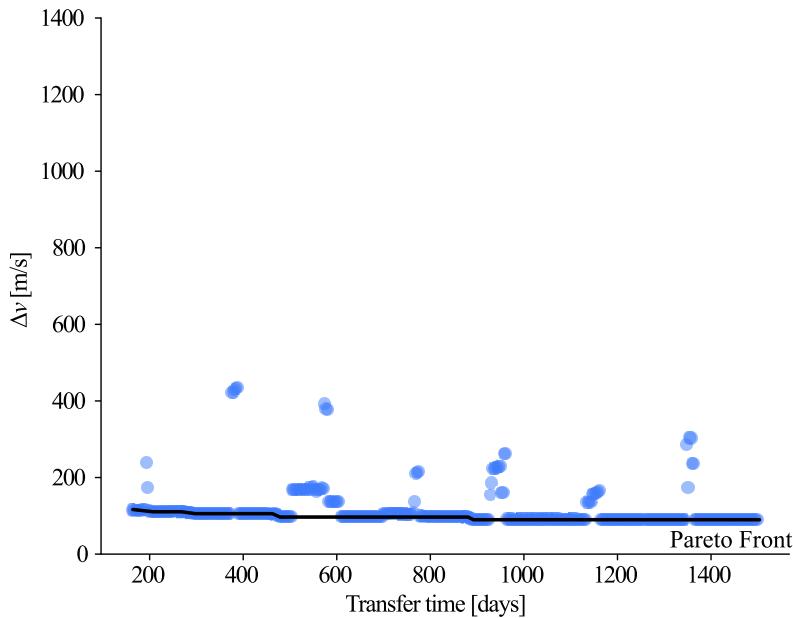


Figure 4.16: Pareto front for 2012 TF79, an ERO with all transfers below the ERO threshold, including a zero-time, single impulse transfer. The Pareto front is shown as a black line.

The $\Delta\nu$ -optimal solutions, and the ratio between the single-impulse $\Delta\nu$ and the optimal two-impulse $\Delta\nu$, is given in Table 4.5 for all the EROs identified. Note that the table of $\Delta\nu$ -optimal solutions in Table 4.4 was updated if the Pareto front search yielded a lower-cost solution than the initial global search.

One object, SPK ID 3648046, has a single-impulse transfer far higher than the optimal $\Delta\nu$. This was found to be due to an unstable Lambert arc – as discussed previously – affecting convergence at the previous value of $t_t = 1$ day. Given that there are only relatively few points in the manifold suitable for capture with a single impulse, the optimiser is dependent on continuing the solution from a transfer time of 1 day to 0 days. The optimiser was then unable to converge for the zero-impulse transfer. This solution could be improved significantly by ‘manual’ optimisation and thus represents an outlier.

In general, the optimal capture $\Delta\nu$ for a given ERO is not unique and is approximately constant through the t_t vs. $\Delta\nu$ space. This offers significant flexibility to the mission designer. These NEOs may also be more promising as retrieval targets, since they are more resilient to changes or to disruptions affecting mission timelines. These objects may also allow for material samples to be returned over a longer span of time without the need for full capture, a strategy which would see samples returned at regular intervals via the invariant manifolds of the Sun-Earth system, and return to a Lagrange point of the system using only natural dynamics¹⁰. The single-impulse

¹⁰This was later investigated in Ionescu, McInnes, and Ceriotti [74] and found to be a feasible method of sample return.

SPK ID	Optimal $\Delta\nu$ [m/s]	Optimal $\Delta\nu$ orbit	Single-impulse $\Delta\nu$ Optimal $\Delta\nu$	Single-impulse orbit
3405338	13.97	L1S	25.00	L2V
3843865	15.25	L2S	30.62	L2S
3733264	23.68	L1V	2.40	L1V
3698849	23.89	L1S	18.27	L2P
3054374	27.39	L1N	2.27	L1P
3315004	27.55	L1S	3.18	L1S
3843512	33.06	L1V	3.13	L1V
3076774	41.60	L1V	15.48	L2N
3435539	42.23	L1V	8.76	L2P
3405189	50.94	L1N	4.48	L1P
3836857	51.52	L2V	1.66	L2V
3702319	56.93	L1V	3.24	L1N
3403148	63.83	L2N	1.83	L2N
3696401	69.84	L1N	1.74	L1P
3840784	81.52	L1V	7.80	L1P
3759358	89.76	L2N	3.82	L2V
3610175	90.82	L2V	1.27	L2V
3893945	94.96	L1N	3.64	L1P
3726012	108.25	L1N	3.95	L2P
3582145	111.98	L1V	1.74	L1V
3991650	112.27	L2V	1.81	L2S
3697803	115.13	L1V	1.35	L1V
3989308	116.72	L2V	2.54	L1P
3735181	123.71	L1P	2.03	L1P
3879383	128.16	L1N	2.02	L1S
3444297	136.52	L1N	2.98	L2P
3618493	140.38	L1S	2.17	L1P
3519518	140.95	L1P	2.36	L1P
3549639	146.66	L2V	2.42	L1P
3005816	162.60	L1V	3.86	L1P
3390109	166.09	L2V	2.48	L2P
3556127	169.01	L2V	1.54	L2P
3825158	170.93	L1P	3.20	L1P
3551168	180.83	L2V	1.30	L2V
3550232	195.40	L1N	1.71	L1N
3625129	203.01	L1S	1.60	L1N
3781983	245.83	L1P	1.65	L1P
3826850	254.37	L2V	2.20	L1P
3745994	301.71	L2S	1.76	L1P
3781977	305.79	L2V	1.09	L2V
3748467	332.02	L1P	2.02	L1P
3648046	355.29	L2V	2032	L2V
3568303	418.53	L2V	1.23	L2V
3719859	445.35	L2N	1.35	L2N

Table 4.5: Table of the $\Delta\nu$ -optimal solutions for the EROs found, and the ratio of single-impulse (S.I.) $\Delta\nu$ to the optimal $\Delta\nu$; many single-impulse transfers exist at similar or less cost than two-impulse transfers. 24 EROs are retrieved into different periodic orbits for the single-impulse transfer.

transfers also often have different transfer properties, including the final periodic orbit, than the optimal two-impulse transfer.

This behaviour is to be expected from the use of the invariant manifolds to facilitate NEO capture. In backwards time, the structure of the stable manifold ‘stretches’ and spans large regions of space in the Sun-Earth system. Thus, for any possible target point, there exists other points in the set of pre-computed manifolds that are retrievable for similar Δv . A dense underlying grid of manifolds and target points that can be interpolated between increases the likelihood of discovering low-cost transfers for all possible transfer times, provided the parameter space is not too large for the optimiser.

The near-constant transfer cost also extends to objects that do not meet the traditional definition of an ERO. 2011 BL45, the NEO introduced earlier to discuss the preliminary tuning of optimiser behaviour is not an ERO, but exhibits the same behaviour as the EROs identified. In fact, the zero-time transfer and optimal capture Δv for this object differ by only 2.5%.

The orbit family used for capturing the NEO for minimum cost and the single-impulse transfer differ for 24 of the 44 EROs found in Table 4.5. While this could be enforced in the optimisation procedure if a specific final periodic orbit is desired, the optimiser was left free to find the minimum capture Δv . The transfer epoch and insertion point onto the manifold can also be significantly different from the Δv -optimal solutions.

Note also the ‘flat’ Pareto front as a result of the optimisation strategy and optimiser parameter tuning, when contrasted to the first approach to the Pareto front generation in Figure 4.12. Based on this work, with further computation time the lowest transfer Δv for some NEOs could be improved even further, which would also help ‘flatten’ the structure of their Pareto fronts. This includes the case of the single-impulse transfer for 3648046, introduced previously.

4.6 Conclusion

This chapter has shown how structures predicted by dynamical systems theory in time-independent astrodynamics systems can be used to enable the retrieval of asteroids in the CR3BP. In combination with high-performance computing, many of the limitations placed on the analysis pipeline to limit computational expense have been lifted, significantly raising the number of asteroids thought retrievable using this method.

To summarise, 27 more EROs, including 17 new capture trajectories below 100 m/s, have been identified, and previous solutions are improved by an average of 176.7 m/s. The Pareto fronts in the Δv *versus* transfer time space have been studied for the first time, and I identify that the optimal transfer cost for these EROs is not unique. A key result is the existence of single-impulse transfers for similar cost to their two-impulse counterparts, which could provide significant flexibility to mission designers.

However, it was noted that the invariant manifolds that underpin this work are not guaranteed to exist in systems with arbitrary time-dependence. The remaining chapters will aim to make the use of Lagrangian Coherent Structures more straightforward in astrodynamics systems to perform investigations similar to that performed in this chapter directly in time-dependent models of motion.

5

An improved numerical method for three-dimensional hyperbolic Lagrangian Coherent Structures

This chapter presents a novel use of Differential Algebra (DA), a method for storing, manipulating and evaluating polynomials in a computer directly, to improve the numerical performance of the algorithm for determining hyperbolic Lagrangian Coherent Structures in dynamical systems. This replaces the need to alter the parameters used for approximating relevant quantities with an automatic approach that requires no parameter tuning but which determines these quantities to higher accuracy. The chapter shows that the method reproduces and outperforms the results of standard approaches on a series of test cases from the literature. Content from this chapter has been accepted for publication in the Journal of Computational Science as Tyler and Wittig [13], and presented at the 2022 Space Flight Mechanics Meeting in Charlotte, North Carolina [14].

5.1 Background

The previous chapter has shown how the invariant manifolds of the CR3BP can be used to design practical space missions, in this case dramatically reducing the transfer cost of performing asteroid retrieval into periodic orbits about the Sun-Earth L_1 and L_2 points. The flexibility these structures offered was also elaborated. However, recall from chapter 3 that for time-independent systems the invariant manifolds are well-defined and partition the phase space. In systems with time-periodic behaviour, such structures become far more difficult to find. In time-aperiodic systems, there is no guarantee of their existence at all.

It was noted in chapter 3 that the fluid dynamics community has approached the challenge of finding similar structures as the invariant manifolds in flows with arbitrary time-dependence through a range of flow heuristics and definitions. One such type, the hyperbolic Lagrangian Coherent

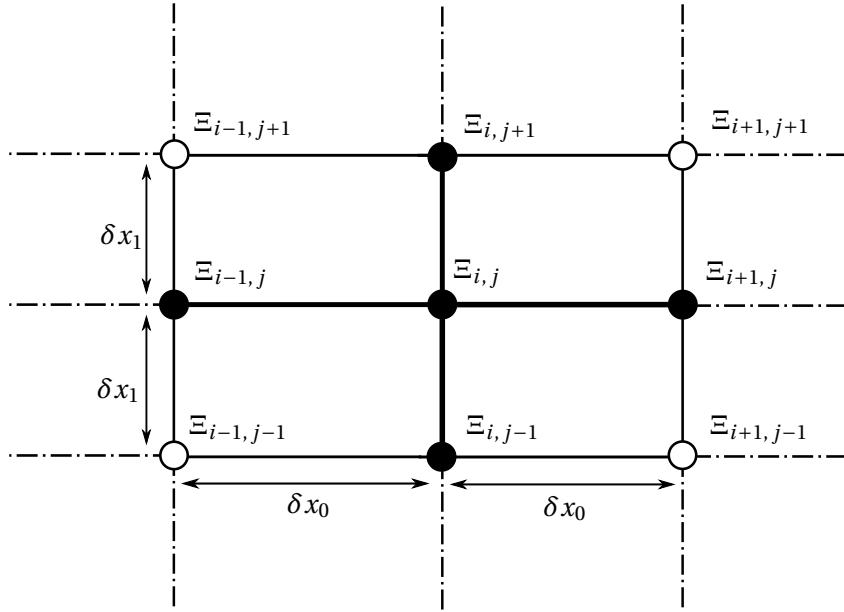


Figure 5.1: The approximation of divided differences to computing derivatives across a two-dimensional grid on which the values of the function values are known. In practice, the offsets $\delta x_0, \delta x_1$ are often determined through trial-and-error, but are a key parameter in the approximation.

Structure (LCS) (henceforth just LCS) was shown to be an effective theoretical grounding for such structures as yielding no false positives whilst generalising the behaviour of the hyperbolic invariant manifolds to systems with arbitrary time-dependence. Being able to use such structures would mean that missions such as those introduced in the previous chapter would be able to be constructed directly in higher-fidelity, time-dependent models rather than in simpler time-independent approximations to motion and then successively continued into higher-fidelity models.

Despite formulating a solid theoretical foundation for a global, objective approach to the computation of LCS in three dimensions in Blazevski and Haller [3], there are several computational complexities associated with computing LCS using this approach [138], such as the need to account for degenerate points (such as where $\lambda_3 = \lambda_2$) and orientation discontinuities in the eigenvector field of $C_{t_0}^T$, since the eigenvector is only defined up to a sign. In wider dynamical systems, Sánchez-Martín, Masdemont, and Romero-Gómez [152] who used two-dimensional LCS in a galactic bar model, remarked that accurately approximating derivatives for the required quantities was ‘crucial’, but required great care. Allshouse [196] who studied coherent structures in ocean flows found that the numerical accuracy required was such that estimating the quantities from experimental data was prohibitive, and only high-density velocity information from analytical equations of motion or regions with detailed radar measurements could be used. Even when Ghosh et al. [197] had these high-density measurements, only the FTLE could be used to signal LCS due to the loss in accuracy associated to recasting unstructured data to a structured grid for use in taking derivatives.

This numerical difficulty comes from needing to approximate high-order derivatives of flows. The additional criteria that must be satisfied in arbitrary n -dimensional flows involve the derivatives of the dominant eigenvalues and eigenvectors of the Cauchy-Green Strain Tensor, which is in effect a second derivative of the flow since one derivative is already required to compute the Cauchy-Green Strain Tensor. This quantity is very numerically difficult to compute, and is often approximated using finite-differences [98, 132, 152, 198, 199, 200]. That is, for a given quantity $\Xi(\mathbf{x})$ defined at some position grid position $\mathbf{x}_{i,j,\dots,z}$, if Ξ is known at neighbouring grid points with some position offset $\Delta = (\delta x_0, \delta x_1, \dots, \delta x_n)$, then $\nabla \Xi$ can be approximated as

$$\nabla \Xi(\mathbf{x}) \approx \begin{bmatrix} \frac{\Xi_{i+1,j,\dots,z} - \Xi_{i-1,j,\dots,z}}{2\delta x_0} & \frac{\Xi_{i+1,j,\dots,z} - \Xi_{i-1,j,\dots,z}}{2\delta x_1} & \dots & \frac{\Xi_{i+1,j,\dots,z} - \Xi_{i+1,j,\dots,z}}{2\delta x_n} \\ \frac{\Xi_{i,j+1,\dots,z} - \Xi_{i,j-1,\dots,z}}{2\delta x_0} & \frac{\Xi_{i,j+1,\dots,z} - \Xi_{i,j-1,\dots,z}}{2\delta x_1} & \dots & \frac{\Xi_{i,j+1,\dots,z} - \Xi_{i,j-1,\dots,z}}{2\delta x_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\Xi_{i,j,\dots,z+1} - \Xi_{i,j,\dots,z-1}}{2\delta x_0} & \dots & \dots & \frac{\Xi_{i,j,\dots,z+1} - \Xi_{i,j,\dots,z-1}}{2\delta x_n} \end{bmatrix} \quad (5.1)$$

which is presented graphically for a 2D grid in x and y in Figure 5.1.

The approximation of divided differences can be troublesome. Firstly, one needs to determine the offset Δ carefully, such that no information is lost by choosing a step-size that is too large, but ensuring that enough information is captured by not choosing a step-size that is too small. In practice, this offset is chosen through trial-and-error, which requires multiple trial computations. This consumes both computation time and user time and is generally difficult to determine *a priori*. It may also be necessary to use different grid offsets for computing different derivatives if derivatives of multiple parameters are required, or if behaviour exists at different length scales at different portions in the flow.

Moreover, near the separatrices the derivatives of the flow will necessarily become large. This can lead to the Cauchy-Green Strain Tensor becoming ill-conditioned, and exacerbate any numerical inaccuracies in the method. However, this is the exact location at which one would expect a hyperbolic LCS, which means accurate evaluation of the relevant quantities is essential for accurate numerical determination of the LCS. While other methods for approximating derivatives exist, such as the use of variational equations which yields derivatives as accurate as the propagation along the reference trajectory, this approach requires the derivation, implementation and integration of n^2 additional equations for the first derivatives of an n dimensional flow, and another $n^2(n+1)/2$ equations for the second flow derivatives.

Needing to integrate the trajectories required for LCS can increase computation time significantly. As a result, various techniques have been used to attempt to reduce the computational complexity. Lipinski and Mohseni [201] used a ridge-tracking algorithm that prevented the computation of trajectories away from the LCS, instead constructing each portion of the structure iteratively. Many different approaches to adaptive mesh refinement have also been used to decrease run-times and improve accuracy near complex features in the flow [202]; some of these

methods may reduce computing times by another order of magnitude [203]. All of these methods work in tandem with standard approaches to program parallelism on CPUs. Recently, Lin, Xu, and Fu [204] published work that details the parallelism of LCS computation on Graphics Processing Units (GPUs) using NVIDIA’s CUDA programming paradigm. However, this work used a simple geometrical approach to extract LCS directly from FTLE fields, an approach that parallelises trivially but with the same pitfalls that FTLE fields alone do not necessarily indicate LCSs.

Motivated by authors of previous studies that attempted to apply the variational theory of LCS to astrodynamics systems citing numerical difficulties [7, 8], this chapter reports on an investigation to determine whether modern numerical methods well-established in astrodynamics could be used to improve the numerical performance of algorithms designed to find LCS in dynamical systems with arbitrary time-dependence. It focuses on improving the numerical performance of the three-dimensional LCS formulation in Blazevski and Haller [3] rather than the n -dimensional formulation. This choice had two main motivators: firstly, parts of this method have been successfully applied to spaceflight problems in Short et al. [113]. Moreover, sampling a six-dimensional phase space with only three variables has shown to be possible in other investigations into ballistic capture [205, 206, 207, 208]. Nonetheless, the techniques outlined in the following maintain applicability to n -dimensional flows in the future.

The outcome of this chapter is a new numerical method for computing LCS using Differential Algebra, called DA-LCS, and the structure of this chapter now follows. First, section 5.3.1 briefly recapitulates how polynomial expansions of arbitrary flows of an ordinary differential equation (ODE) can be calculated, with applications to obtaining flow derivatives of arbitrary systems to machine precision. Next, section 5.3.2 introduces a novel use of DA to construct algebraic expansions of the leading eigenvector of a matrix of polynomials. Both of these techniques are then combined to form the DA-LCS algorithm for computing LCS in three-dimensional flows. In section 5.5, I demonstrate that this method works well in reproducing and improving upon results for commonly-used ‘toy’ problems from the literature.

5.2 A review of the mathematical background and notation

The mathematical background required to study LCS in dynamical systems is now reviewed to aid clarity after being first introduced in chapter 3. To summarise, the method studies the behaviour of a dynamical system

$$\dot{\mathbf{x}} = f(\mathbf{x}, t), \mathbf{x} \in D \subset \mathbb{R}^n, t \in [t_0, t_0 + T] \quad (5.2)$$

where f is a smooth vector field considered over some time T starting at time t_0 . Denoting a trajectory of the dynamical system starting at position \mathbf{x}_0 at time t_0 as $\mathbf{x}(t_0, \mathbf{x}_0; T)$, the flow map

of Equation 5.2 is given by

$$\mathbf{F}_{t_0}^T : \begin{cases} D \rightarrow D \\ \mathbf{x}_0 \mapsto \mathbf{x}(t_0, \mathbf{x}_0; T) \end{cases} \quad (5.3)$$

which is assumed to be at least twice continuously differentiable. The Jacobian of this flow map, $\nabla \mathbf{F}_{t_0}^T$, defines the right Cauchy-Green Strain Tensor $C_{t_0}^T$, which describes the local deformation of the flow at the end of a given trajectory.

$$C_{t_0}^T = (\nabla \mathbf{F}_{t_0}^T)^\top (\nabla \mathbf{F}_{t_0}^T) \quad (5.4)$$

with $^\top$ denoting the matrix transpose. $C_{t_0}^T$ is positive-definite and symmetric, with real eigenvalues $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_n$ and associated real eigenvectors $\zeta_1, \zeta_2, \dots, \zeta_n$.

The dominant eigenvalue λ_n can be used to calculate the finite-time Lyapunov exponent (FTLE), a measure of maximum separation of two particles advected forward under Equation 5.2 that start out infinitesimally close to each other:

$$\sigma_{t_0}^T = \frac{1}{2} \frac{\log \lambda_n}{T}. \quad (5.5)$$

Many previous studies have leveraged the FTLE field as a heuristic indication of high regions of separation in the flow. While the FTLE has been shown to be insufficient to indicate LCS alone [132], the FTLE is a commonly-used metric and is thus used to preliminarily highlight system behaviour.

5.3 Differential algebra

Given that approximating the derivatives required for the LCS has been shown in the literature to be numerically difficult and time-consuming in both user time and computation time when using standard approaches of divided differences, I now seek a way of taking more accurate derivatives of functions automatically¹.

Three options present themselves from the literature. The first is symbolic mathematics, where the floating-point numbers are replaced with a symbolic algebra and the computer directly manipulates algebraic expressions and symbols. Packages such as Python’s SymPy or Matlab’s symbolic math toolbox can make the implementation of these straightforward. This approach can yield high-order derivatives readily, but for long, complex functions, needing to manipulate these algebraic expressions directly can be computationally very expensive, and the LCS itself has already been shown to be computationally intensive itself.

Secondly, one could use a standard numerical automatic differentiation package. By using a custom datatype, one can record all the operations performed on a given input (stored on a ‘gradient tape’) and evaluate its derivative using a chain rule by traversing along the recorded operations

¹This is a relatively new field known as Differentiable Programming.

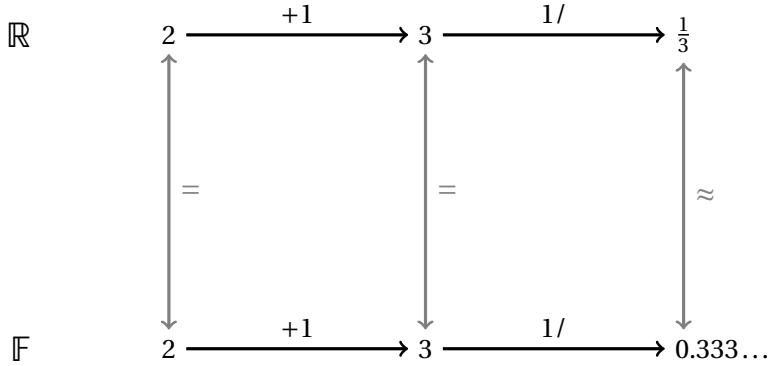


Figure 5.2: Comparison between computing $1/(2+1)$ in the field of real numbers \mathbb{R} and the the field of floating-point approximations to \mathbb{R} , \mathbb{F} . Each operation in \mathbb{R} has an adjoint operation in \mathbb{F} . Figure taken from [5].

when the derivative is requested. This is common in machine learning packages, such as TensorFlow or PyTorch, for determining Jacobians to use in gradient descent methods. While this is quicker than the approach of symbolic math since it relies on numerical manipulation of relevant quantities rather than an algebraic one, the majority of automatic differentiation packages that implement the above method only function to first order. This is because the number of arithmetic rules and data to store is quadratic in the desired order of the derivative.

As a middle ground between the speed of numeric automatic differentiation and the capability of symbolic mathematics, Differential Algebra (DA) is a technique that is becoming more widespread in the field of astrodynamics for the study and propagation of uncertainties [5, 209, 210, 211, 212, 213, 214] and as a tool for automatic differentiation using truncated Taylor polynomials. Originally designed to bridge the gap between numerical integration-based method and approximations using standard relations in modelling the optics of particle beams [215], DA is a technique for approximating functions directly in a computer environment, and is built on the founding principle that it is possible to extract more information from a function than just its mere values. Where the floating-point numbers \mathbb{F} approximate the field of real numbers \mathbb{R} , DA allows the approximation of the k -differentiable functions using high-order Taylor polynomials represented natively in a computer environment.

Consider two real numbers a and $b \in \mathbb{R}$. The approximation to a and b in a computational environment is their floating-point representation $\bar{a}, \bar{b} \in \mathbb{F}$, which essentially stores a set number of digits of its binary expansion. Any operation defined in \mathbb{R} , \square , has a corresponding operation in \mathbb{F} , \boxtimes , defined such that the result is another floating-point approximation of the operation on the real numbers a and b , i.e. $\bar{a} \times \bar{b}$ commutes with the floating-point representation of $a \times b$, $\overline{a \times b}$.

Similarly, now consider two functions, c and d , which are sufficiently smooth, k -differentiable functions of n variables: $c, d : \mathbb{R}^n \rightarrow \mathbb{R}$. In the DA framework, a computer operates on the multivariate Taylor expansion of c and d , $[c]$ and $[d]$, with corresponding operations to those

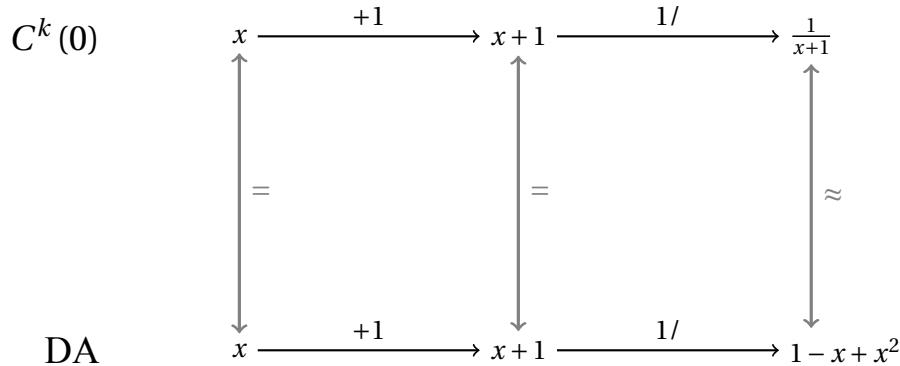


Figure 5.3: Comparison of the manipulation of k -times differentiable functions in C^r and DA using the example of $1/(x+1)$; operations in C^k have adjoint operations in the DA space \mathcal{P} . The final result of the manipulation in DA is a Taylor polynomial which locally approximates the function in C^k up to a given expansion order. Figure taken from [5].

	I	COEFFICIENT	ORDER	EXPONENTS
#include<iostream>	1	8.4147098480789650e-01	0	0
#include<dace/dace_s.h>	2	5.4030230586813977e-01	1	1
auto main() -> int {	3	-4.2073549240394825e-01	2	2
// 10th order in one variable	4	-9.0050384311356632e-02	3	3
DACE::DA::init(10, 1);	5	3.5061291033662352e-02	4	4
DACE::DA x = DACE::DA(1);	6	4.5025192155678318e-03	5	5
std::cout << sin(1 + x) << std::endl;	7	-1.1687097011220783e-03	6	6
return 0;	8	-1.0720283846590076e-04	7	7
}	9	2.0869816091465686e-05	8	8
	10	1.4889283120263995e-06	9	9
	11	-2.3188684546072984e-07	10	10

Figure 5.4: Example of C++ code using DA via the DACE [6] (left) and the output (right) to show how operator overloading makes using DA in modern codes straightforward, and how the output is a high-order truncated Taylor polynomial of $\sin(1 + x)$.

defined in the real function space, such that the operation of $[c] \cdot [d]$ commutes with the DA representation of the product $[c \cdot d]$.

An example to demonstrate how real numbers are approximated in a computer environment is provided in Figure 5.2 for the evaluation of the expression $1/(x+1)$ for $x = 2$ in \mathbb{F} and \mathbb{R} . In Figure 5.2, one begins with $x = 2$, performs the operation $+1$ to obtain three, and then perform the operation $1/$ to compute the final expression. In \mathbb{R} , the user obtains the solution $1/3$, and in \mathbb{F} the user obtains the solution $0.333\dots$ up to the limit of precision of the type. The final result of the evaluation in floating-point arithmetic is an approximation of the real computation.

Analogously, in Figure 5.3 the expression $1/(1+x)$ is evaluated in the space $C^k(0)$ of real functions, and a DA representation of expansion order 2. Beginning with the function $c(x) = x$, the operations $+1$ and $1/$ are performed, yielding $1/(x+1)$ in the real function space, and $1-x+x^2$ in the DA arithmetic. The result of the DA arithmetic is the Taylor expansion of $1/(x+1)$ which represents the function exactly at $x = 0$, and approximates the function locally near $x = 0$ with an error of $\mathcal{O}(x^3)$. The coefficients of the expansion are computed automatically without any further input from the user (Figure 5.4).

Differential Algebra contains all the necessary operations to store, manipulate and evaluate polynomials in a computer environment [216]. This includes operations for common intrinsic functions such as division, square roots, trigonometric functions, and exponentials, as well as operations for differentiation and integration [5]. Some classically floating-point algorithms have been reformulated to make use of the DA framework [217]. Since the algebra includes a derivative operator – hence Differential Algebra – and the coefficients of the expansion are computed automatically, DA can be considered a form of automatic forward differentiation.

5.3.1 Flow expansion using DA

A key advantage of using DA is that the derivatives of flows with respect to the initial conditions can be obtained automatically and without any further effort from the user, beyond implementing the system's governing equations and the numerical integration scheme in DA arithmetic. This is a technique that has been widely used in the literature [137, 214, 218, 219, 220, 221, 222], and is known as flow expansion.

To illustrate this concept of flow expansion, a demonstration is now performed using the simplest of the Runge-Kutta family of numerical integrators, the forward Euler scheme. Suppose one wants to solve the following initial value problem (IVP) numerically:

$$\begin{cases} \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t) \\ \mathbf{x}(t_i) = \mathbf{x}_i. \end{cases} \quad (5.6)$$

A single step in the forward Euler scheme at the i -th point on the trajectory is given explicitly by

$$\mathbf{x}_i = \mathbf{x}_{i-1} + \Delta t \mathbf{f}(\mathbf{x}_{i-1}) \quad (5.7)$$

which can be expressed as a function of the initial condition \mathbf{x}_0 ,

$$\mathbf{x}_f = \mathbf{x}_0 + \sum_{i=0}^n \Delta t \mathbf{f}(\mathbf{x}_0 + i \cdot \Delta t) \quad (5.8)$$

i.e. the initial condition is simply a sequence of operations on the initial condition. The above is straightforward to extend to any other explicit integration scheme, where the solution at the next time-step is a simple function of the solution at the current time. It is in theory possible to also extend the above to implicit methods where one must solve an equation at every time-step, but care must be taken to ensure the equation is solved in all orders of expansion.

If the initial condition \mathbf{x}_0 is set to be a DA representation of the initial condition by substituting the initial value with the DA identity, $[\mathbf{x}(t_0)] = \mathbf{x}(t_0) + \delta\mathbf{x}$, then \mathbf{x}_f becomes a DA representation of the final condition as a function of the initial condition, $[\mathbf{x}_f]$. Applying the partial derivative operator made available as part of the DA arithmetic, ∂_j , means the polynomial representing the final condition as a function of the initial condition can be differentiated completely automatically. Since the coefficients of the polynomial are stored to machine precision, the derivatives

will also be available to machine precision without the need for any further effort from the user, such as deriving and implementing the variational equations for the flow.

To implement the above practically, the numerical integrator must support DA operations and be supplied with a DA initial condition. One can either write their own numerical integrator designed exclusively for DA, or exploit modern templated C++ libraries such as Boost C++. The latter's 7th/8th order Dormand-Prince method is used throughout this thesis, and requires only one modification to support DA compared to native double-precision types¹. When computing the estimated error in a given step, the L_2 norm is usually evaluated. Evaluating the usual L_2 norm of a vector $|\mathbf{x}| = \sqrt{\sum_{i=0}^n x_i^2}$ in DA yields another DA object representing a polynomial. As there is no ordering on the space of polynomials, this cannot be directly compared to some tolerance. Instead, an additional definition of the norm of a DA object is made which maps it into the non-negative real numbers. In this thesis, the norm of a DA object is taken to be the largest absolute value of any coefficient of the expansion in any order. Considering all orders in the norm allows the usual step-size control algorithms of embedded Runge-Kutta methods to control the error in all orders of the expansion, rather than just the constant part.

DA has also been applied to numerous other astrodynamics problems, such as the Two-Point Boundary Value Problem; spacecraft guidance and state estimation; trajectory optimisation; and orbital conjunctions [220, 223, 224, 225, 226].

5.3.2 Computing expansions of leading eigenvectors of $C_{t_0}^T$ to arbitrary order

Since derivatives of polynomials are straight forward to compute, one can apply the partial derivative operator to differentiate the j -th variable of an expansion, ∂_j , making it particularly easy to assemble an expansion of $C_{t_0}^T$. This means the Jacobian can be evaluated directly as

$$[\nabla \mathbf{F}_{t_0}^T]_{ij} = \partial_j [\mathbf{x}]_{t_0, i}^T \quad (5.9)$$

from which a polynomial expansion of $C_{t_0}^T$ can be assembled

$$[C_{t_0}^T] = [\nabla \mathbf{F}_{t_0}^T]^\top [\nabla \mathbf{F}_{t_0}^T]. \quad (5.10)$$

Note that the constant part of $[C_{t_0}^T]$ is the CGST at the expansion point accurate to machine precision, that is it is the same as would be approximated with via divided differences or variational equations. The remaining higher order terms represent an expansion of the CGST in the neighbourhood around the expansion point².

To compute the LCS, the derivatives of the leading eigenvector of $C_{t_0}^T$ with respect to position are required. While divided differences can in principle again be used to obtain these derivatives,

¹This was motivated by the recommendation made by Løken [151], who performed an excellent review of numerical integration schemes for propagating strainlines.

²Investigations as part of this thesis has shown the convergence radius to generally only be accurate with a radius of \mathcal{O} (metres) and so it is not used directly here.

the method is susceptible to numerical noise and it is difficult to determine the most appropriate grid sizes to use. Moreover, eigenvectors are only defined up to a sign, and thus care must be taken when taking the derivatives that nearby eigenvectors have ‘smooth’ changes in orientation.

It is possible, in combination with a backward propagation pass, to compute derivatives of eigenvectors and eigenvalues directly with an automatic differentiation framework using gradient tapes (see, for example, Xie, Liu, and Wang [227]). However, this requires manual differentiation of some of the relations *a priori* to ensure that only the derivatives of the desired eigenvalues/eigenvectors are obtained and not the derivatives of the full spectrum. Moreover, I seek derivatives of the leading eigenvector, and using automatic differentiation will necessarily also include information on the eigendecomposition process.

The use of analytic relations to compute eigenvalues and eigenvectors is also possible in theory for positive symmetric-definite matrices [228]. One can recast the characteristic polynomial of $\det(A - I\lambda)$ into a form where it can be solved using Cardano’s formula, a general solution for roots of cubic equations. Once the eigenvalues are known, it is straightforward to find the eigenvectors. However, deriving the relations required is non-trivial, and the derivatives of the eigenvector require further manipulation and derivations. Importantly, one of the coefficients in the solution can become very large as the relative magnitudes of the eigenvalues in the spectrum increase, which leads to significant error in the calculation [229]. This is worsened when using this result to compute eigenvectors. Thus, while mathematically accurate, its numerical implementation can lead to significant errors. Kopp [228] attempted to tackle this to use the analytical approach more widely by suggesting an eigendecomposition routine that first attempts to use Cardano’s formula, falling back to a QR decomposition if the estimated error exceeds a certain bound³. However, the QR decomposition would need to be differentiated using the gradient tape approach given above, at which point the advantages of using the analytical approach are lost. Kopp actually recommends against the use of Cardano’s formula in practical applications due to accuracy concerns.

Instead, a novel application of DA is used to obtain an expansion of the leading eigenvector of a matrix of DAs, which then can once again be differentiated directly in DA. Given that only the derivatives of the dominant eigenvector of $C_{t_0}^T$ are required, a modified power (von Mises) iteration [230] is performed, which is a well-established algorithm in standard floating-point operations [231] for obtaining leading eigenvectors. This avoids the need for any *a priori* knowledge and avoids derivatives associated with the decomposition being present in the final result, negating the downside of standard automatic differentiation approaches. It also obtains the eigenvectors to machine precision, retaining much of the accuracy of the analytic formulation.

Power iteration performs the repeated evaluation of an arbitrary starting vector b_0 through a matrix A to obtain an approximation to its dominant unit eigenvector b through the recurrence

³The error bound is $\|\zeta_n\|^2 \leq 2^8 \varepsilon \Lambda^2$, where ε is the machine precision, $\Lambda^2 = \max(\lambda_n^2, \lambda_n)$, and ζ_n is the unnormalized eigenvector.

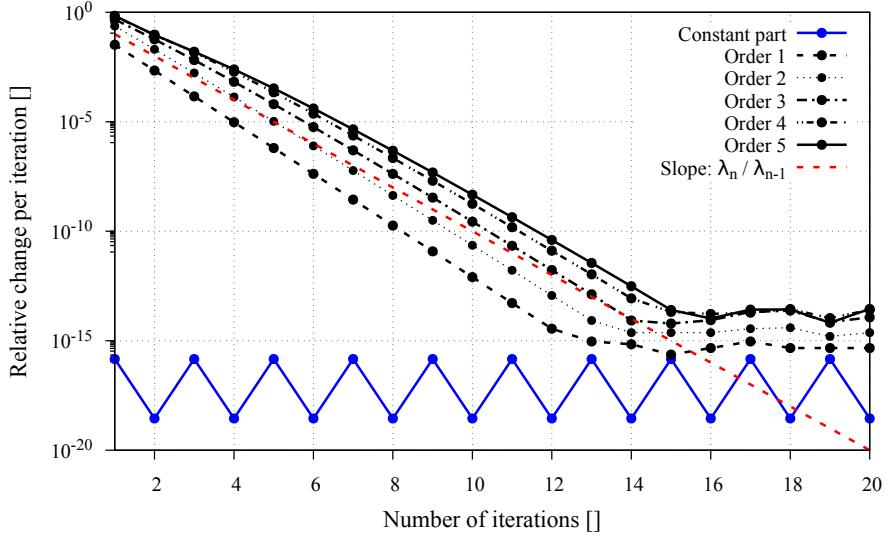


Figure 5.5: Relative error across all polynomial orders in successive applications of $[C_{t_0}^T]$ to an initial guess containing only the floating-point dominant eigenvector at the expansion point as the constant part. Higher expansion orders (black) can all be seen converging at around the expected convergence rate $\lambda_n / \lambda_{n-1}$ (dashed red) towards the floating-point floor.

relation

$$\mathbf{b}_{m+1} = \frac{A\mathbf{b}_m}{\|A\mathbf{b}_m\|} \quad (5.11)$$

where $\|\cdot\|$ represents a vector norm, here taken to be the L_2 norm, the vector \mathbf{b}_0 is an arbitrary initial vector, and m is the number of iterations. The vector \mathbf{b} will converge provided that the starting vector \mathbf{b}_0 has a nonzero component in the direction of the dominant eigenvector, and A has a unique largest eigenvalue by absolute value. The theoretical convergence rate of the method between successive iterations is the ratio of the dominant eigenvalue to the second dominant eigenvalue. Practically, the recurrence relation is iterated until the stopping condition $\|\mathbf{b}_{m+1} - \mathbf{b}_m\| \leq \varepsilon$ is valid, where $\varepsilon > 0$ is a pre-set tolerance and the norm is again taken to be an L_2 norm.

To convert this algorithm to DA, let A now be a DA matrix with DA objects in each entry, $[A]$. Iterating it on a DA vector $[\mathbf{b}_0]$ will yield a DA vector $[\mathbf{b}]$ corresponding to the dominant eigenvector of $[A]$ with a polynomial expansion in each entry, that is it is the recurrence relation

$$[\mathbf{b}]_{m+1} = \frac{[A][\mathbf{b}]_m}{\|[A][\mathbf{b}]_m\|}. \quad (5.12)$$

Note that here the norm in the denominator is simply a DA evaluation of the L_2 (Euclidean) norm $\|\mathbf{x}\| = \sqrt{\sum_{i=0}^n [x]_i^2}$. The stopping condition is generalised from floating-point computation such that iteration is performed until there is no change in any order in any entry of $([\mathbf{b}]_{m+1} - [\mathbf{b}]_m)$ above a pre-set tolerance $\varepsilon > 0$. ε is set to be 10^{-12} in this thesis.

To speed up convergence, and because eigenvector solvers for floating-point computations are readily available and highly efficient, the initial guess for $[\mathbf{b}_0]$ is set to have a constant part equal to the dominant eigenvector of the constant part of $[A]$, since by construction this will be the

constant part of $[\mathbf{b}]$. This is not required for the method to converge, but is included since the eigendecomposition of the constant part is performed anyway to find the FTLE in this chapter.

An example of the convergence of this method is illustrated in Figure 5.5, which shows the maximum relative change of coefficients in $[\mathbf{b}]$ separated by their expansion order over repeated application of $[C_{t_0}^T]$ to the initial guess of the dominant eigenvector of a trajectory in the periodic ABC flow (section 5.5.2). The theoretically expected rate of convergence λ_n/λ_{n-1} can be seen in the plot as a dashed red line. All orders converge at approximately the expected rate and the floating-point portion of the expression converges instantly as it was already set to the double-precision representation of the leading eigenvector.

Once the eigenvector $[\zeta_n]$ is expanded to at least first order, the curl $\nabla \times \zeta_n$ of the eigenvector field, which is used in the LCS construction (section 5.4), can be computed by simply applying the DA partial derivative operator ∂_j again. To obtain the value of $\nabla \times \zeta_n$ at the expansion point, the flow map $\mathbf{F}_{t_0}^T$ must be computed at least to order 2. This is because one derivative is taken in the construction of $C_{t_0}^T$ (section 5.3.1), and another is then taken in $\nabla \times \zeta_n$, both of which reduce the order of the expansion by one.

5.4 Lagrangian Coherent Structures

In this section I give the method for computing hyperbolic LCS in three-dimensional systems. I begin by formally reviewing their mathematical construction following [3] and highlight how this is implemented algorithmically and practically following the literature [3, 132, 138]. I then highlight how DA is used to enhance the algorithm in DA-LCS.

The full, three-dimensional hyperbolic LCS, which is the locally maximally repelling or attracting surface over a given time interval $[t_0, t_0 + T]$, is constructed from its intersections with a family of hyperplanes \mathcal{S} . These intersections are called *reduced strainlines* or *reduced stretchlines*. Interpolating between these intersections yields the full, three-dimensional structure of the LCS. In the following, I give the mathematical formulation and perform all the analysis for repelling LCS, whose structure is derived from the dominant eigenvector ζ_3 and whose intersections with \mathcal{S} are the reduced strainlines. The selection of the repelling LCS compared to attracting LCS was motivated by two factors. First, the repelling LCS has already been shown to highlight regions of qualitatively different behaviour in preliminary applications to spaceflight problems [7, 8, 36, 113, 145]. Secondly, the original literature for LCS in three dimensions – Blazevski and Haller [3] – does not provide any results for the attracting LCS, only repelling and elliptic. Thus, there is no known data to compare the implementation against. Nonetheless, a similar procedure applies to ζ_1 (reduced stretchlines) to obtain attracting LCS by following the procedure below after inverting $C_{t_0}^T$ using the algorithm made available in Appendix B.

Mathematically, at any point s on a hyperplane, one can define the reduced strainline that passes through that point from the definition of the repelling LCS as being necessarily orthogonal to ζ_3

and of course lying within the hyperplane. This is true for any point on the strainline, allowing their parameterisation to be described by the ODE

$$\mathbf{s}' = \hat{\mathbf{n}}_{\mathcal{S}} \times \zeta_3 \quad (5.13)$$

where $\hat{\mathbf{n}}_{\mathcal{S}}$ is the unit normal to the surface at \mathbf{s} . Points with zero *helicity* H_{ζ_3} lie on surfaces which are maximisers of repulsion

$$H_{\zeta_3} = \langle \nabla \times \zeta_3, \zeta_3 \rangle, \quad (5.14)$$

where $\langle \cdot, \cdot \rangle$ is the inner product. Reduced strainlines which begin from points with zero helicity are thus the intersection of the LCS with that hyperplane. The strainlines forming part of the LCS on each hyperplane are then interpolated to produce the full 3D structure of the LCS.

Practically, Algorithm 1 is used to solve for the two quantities above following the literature [3, 132, 138]. First, points \mathbf{s} are sampled on each hyperplane in \mathcal{S} on a uniformly-spaced grid, and the derivative of the flow map the derivative of the flow map $\nabla \mathbf{F}_{t_0}^T$ and thus $C_{t_0}^T$ are computed at each point (Lines 3 and 4). The dominant eigenvector of $C_{t_0}^T$, ζ_3 , and its curl, $\nabla \times \zeta_3$, is obtained and then used to compute the helicity (Lines 5 and 6).

Numerically, the selection of zero-helicity initial conditions for the strainline ODE in Equation 5.13 is relaxed by allowing them to begin from points where the helicity H_{ζ_3} is below some tolerance $\alpha > 0$ (Line 7). The ODE in Equation 5.13 is then rewritten in discretised form and solved numerically as

$$\mathbf{s}'_i = \text{sign}(\zeta_{i,3} \cdot \zeta_{i-1,3}) \hat{\mathbf{n}}_{\mathcal{S}} \times \zeta_{i,3} \quad (5.15)$$

where \mathbf{s}_i is the i -th point on the strainline L and the term $\text{sign}(\zeta_{i,3} \cdot \zeta_{i-1,3})$ is introduced to enforce continuity in the vector field by selecting the direction most closely aligned with the previous tangent vector. Since the eigenvector is only defined up to the sign, the strainline is integrated in both directions corresponding to $\pm \zeta_3$ to capture the entire strainline structure. The numerical integration of the ODE along the strainline continues until the sum of the helicity at each \mathbf{s}_i divided by the number of steps performed (i) rises above α (Line 10).

The trajectories obtained through the strainline ODE are segments of strainlines forming the LCS. However, since different initial points can belong to the same strainline, the trajectories often overlap. They must, therefore, be filtered to provide a single, continuous curve. Given a suitable metric d_F of how close two strainline segments are, the shorter of the two strainlines is discarded whenever d_F is below some threshold $\delta > 0$ (Line 15).

I now highlight the additions to the algorithm in DA-LCS. In previous literature, divided differences was used to numerically approximate $\nabla \mathbf{F}_{t_0}^T$ in Line 3, and the quantity $\nabla \times \zeta_3$ in Line 5 [98, 113, 147, 232]. This can lead to significant numerical error, particularly when computing the second derivative required for $\nabla \times \zeta_3$. Divided differences can either be applied on the

Algorithm 1 High-level algorithm for computing three-dimensional LCS

Require: $\mathcal{S}, \alpha, \delta, t_0, T$

- 1: **for** hyperplane S in \mathcal{S} **do**
- 2: **for** point s on hyperplane S **do**
- 3: Compute $\nabla \mathbf{F}_{t_0}^T$ at s
- 4: $C_{t_0}^T \leftarrow \left(\nabla \mathbf{F}_{t_0}^T \right)^\top \left(\nabla \mathbf{F}_{t_0}^T \right)$
- 5: $\zeta_3 \leftarrow$ dominant eigenvector of $C_{t_0}^T$
- 6: $|H_{\zeta_3}| \leftarrow \langle \nabla \times \zeta_3, \zeta_3 \rangle$
- 7: **if** $H_{\zeta_3} \leq \alpha$ **then**
- 8: $L \leftarrow$ new strainline starting at s
- 9: **while** average of H_{ζ_3} along $L \leq \alpha$ **do**
- 10: Extend L via ODE 5.13
- 11: **end while**
- 12: Add L to set of strainlines on S , \mathcal{L}_S
- 13: **end if**
- 14: **end for**
- 15: Filter \mathcal{L}_S for duplicate strainlines L using distance metric and threshold δ
- 16: **end for**
- 17: Interpolate strainlines in \mathcal{L}_S and $\mathcal{L}_{S'}$ for all adjacent hyperplanes S and S' in \mathcal{S}

Algorithm 2 Computation of $\left[C_{t_0}^T \right]$

Require: x_0, t_0, T, f

- 1: $[x_0] \leftarrow x_0 + \delta x$ ▷ Supplement initial condition with DA identity
- 2: $[x] \leftarrow$ Integration of $[x_0]$ from t_0 to $t_0 + T$ under f
- 3: **for** element $[x]$ in $[x]$ **do**
- 4: **for** DA variable j **do**
- 5: $\left[\nabla \mathbf{F}_{t_0}^T \right]_{ij} \leftarrow \partial_j [x]_i$
- 6: **end for**
- 7: **end for**
- 8: $\left[C_{t_0}^T \right] \leftarrow \left(\left[\nabla \mathbf{F}_{t_0}^T \right] \right)^\top \left(\left[\nabla \mathbf{F}_{t_0}^T \right] \right)$ ▷ $C_{t_0}^T$ is the constant part of $\left[C_{t_0}^T \right]$

same grid on which points are sampled, or on a finer grid used solely for the purpose of approximating the derivatives. In DA-LCS, the technique of flow expansion described in section 5.3.1 is used to compute $\nabla \mathbf{F}_{t_0}^T$ as an expansion at each grid point and around each grid point, as given in Algorithm 2. This provides an automatic representation of this quantity accurate to machine precision, and without the need to alter grid sizes through trial-and-error. It thus also provides $C_{t_0}^T$ and ζ_3 in Lines 4 and 5 to high accuracy for use in the strainline ODE. To improve the numerically-challenging computation of $\nabla \times \zeta_3$, Line 6 is computed to machine precision using the power iteration described in section 5.3.2, and again without the need to alter grid sizes through trial-and-error. A detailed algorithm for the computation of $\nabla \times \zeta_3$ is given in Algorithm 3.

In [3], the distance metric used was the Hausdorff distance, a measure of similarity between two curves. Despite the Fréchet distance being recognised as a better measure of similarity than the

Algorithm 3 Computation of $\nabla \times \zeta_3$

Require: $\left[C_{t_0}^T \right], \varepsilon$

- 1: $C_{t_0}^T \leftarrow$ constant part of $\left[C_{t_0}^T \right]$
- 2: $\zeta_3 \leftarrow$ dominant eigenvector of $C_{t_0}^T$ ▷ Solved using standard floating-point techniques
- 3: $[b_0] \leftarrow \zeta_3$ in constant part ▷ Initial guess
- 4: $n \leftarrow 0$ ▷ Number of iterations
- 5: $\mathcal{E} \leftarrow \infty$ ▷ Error
- 6: **while** $\mathcal{E} \geq \varepsilon$ **do**
- 7: $[b_{n+1}] = \left[C_{t_0}^T \right] [b_n] / \left\| \left[C_{t_0}^T \right] [b_n] \right\|$ ▷ $\left\| \cdot \right\|$ is the L_2 norm evaluated in DA
- 8: $\mathcal{E} \leftarrow$ max relative change of coefficient in $[b_{n+1}] - [b_n]$ w.r.t. $[b_{n+1}]$
- 9: $[b_n] \leftarrow [b_{n+1}]$
- 10: $n \leftarrow n + 1$
- 11: **end while**
- 12: $[\zeta_3] \leftarrow [b_n]$
- 13: **for** i -th element in ζ_3 **do**
- 14: **for** j -th DA variable **do**
- 15: Append $\partial_i [\zeta_3]_j - \partial_j [\zeta_3]_i$ to \mathcal{C} ▷ Stores $\nabla \times \zeta_3$
- 16: **end for**
- 17: **end for**
- 18: **return** \mathcal{C}

Hausdorff distance in trajectory clustering problems [233, 234], we find we obtain qualitatively better strainlines through use of the Hausdorff distance. It is defined as follows: given two curves A and B , the Hausdorff filtering distance d_F between A and B is then defined such that

$$d_F = \max \left\{ \max_{x \in A} \left(\min_{y \in B} d_E(x, y) \right), \max_{x \in B} \left(\min_{y \in A} d_E(x, y) \right) \right\} \quad (5.16)$$

where d_E is the Euclidean distance. As per other investigations into computing LCS [138], we also enforce a minimum strainline length of δ ; the rationale behind this is given in section 5.5.1.

5.5 Application to Arnold-Beltrami-Childress Flows

To show the numerical performance of DA-LCS and ensure it reproduces the results from the literature, I now apply the standard approach of divided differences and the DA-LCS method to several variations of the Arnold-Beltrami-Childress (ABC) flow, a relatively straightforward system studied in Blazevski and Haller [3] as test cases. For each example, I present the equations of motion, the FTLE field using DA-LCS and divided differences, and the relative error of the application of the approximation of divided differences to the FTLE field, using DA-LCS as the baseline. I also present the helicity fields for each method and the resulting strainlines. The results obtained using divided differences each use the manually determined optimal grid size that produces the qualitatively ‘best’ results, to allow for a fair comparison. Grid sizes between 0.1 and 5 times the nominal grid size were analysed. No such adjustments are needed when using DA-LCS.

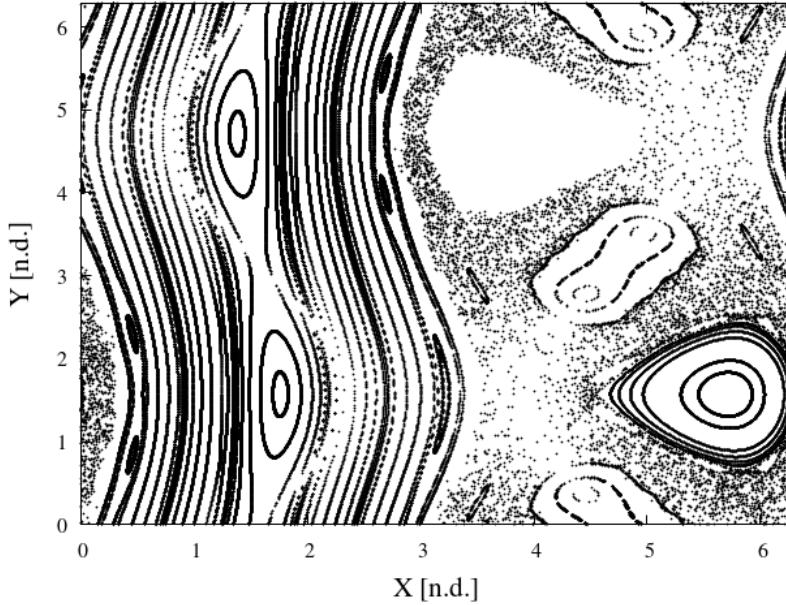


Figure 5.6: Poincaré section (return map) for the steady ABC flow on the $z = 0$ plane; generated using a 15×15 grid of initial points with integration time $T = 1500$.

5.5.1 Steady Arnold-Beltrami-Childress flow

The first system considered is the steady Arnold-Beltrami-Childress flow, as presented in Blazevski and Haller [3]. The ABC flow is an exact solution to Euler's equation, and its equations of motion in Cartesian coordinates are

$$\dot{x} = A \sin z + C \cos y \quad (5.17)$$

$$\dot{y} = B \sin x + A \cos z \quad (5.18)$$

$$\dot{z} = C \sin y + B \cos x \quad (5.19)$$

with parameter values $A = \sqrt{3}$, $B = \sqrt{2}$, $C = 1.0$ as in Blazevski and Haller [3]. To illustrate the behaviour of this system, the Poincaré section in the x - y plane is shown in Figure 5.6, computed from a regular 15×15 grid of initial points and an integration time of $T = 1500$.

For the LCS computation, matching previous literature the set of reference planes are taken to be

$$\mathcal{S} = \{(x, y, z) \in [0, 2\pi]^3 : z \in \{0, 0.005, 0.01, \dots, 0.1\}\},$$

that is the x - y plane evenly spaced along the z axis. However, within each plane I alter the grid size used. Blazevski and Haller [3] use a 500×500 grid on which to compute the underlying helicity field, and then sample seed points for the ODE in Equation 5.13 on a reduced grid of 600×10 . While sampling every point on a dense grid is numerically inefficient, to simplify analysis, ensure all of the flow's behaviour is captured, and to work off of the assumption of no *a priori* knowledge all stages of the analysis are performed on a 1000×1000 grid defined for each hyperplane in \mathcal{S} . In practice, additional information about the system may be available to

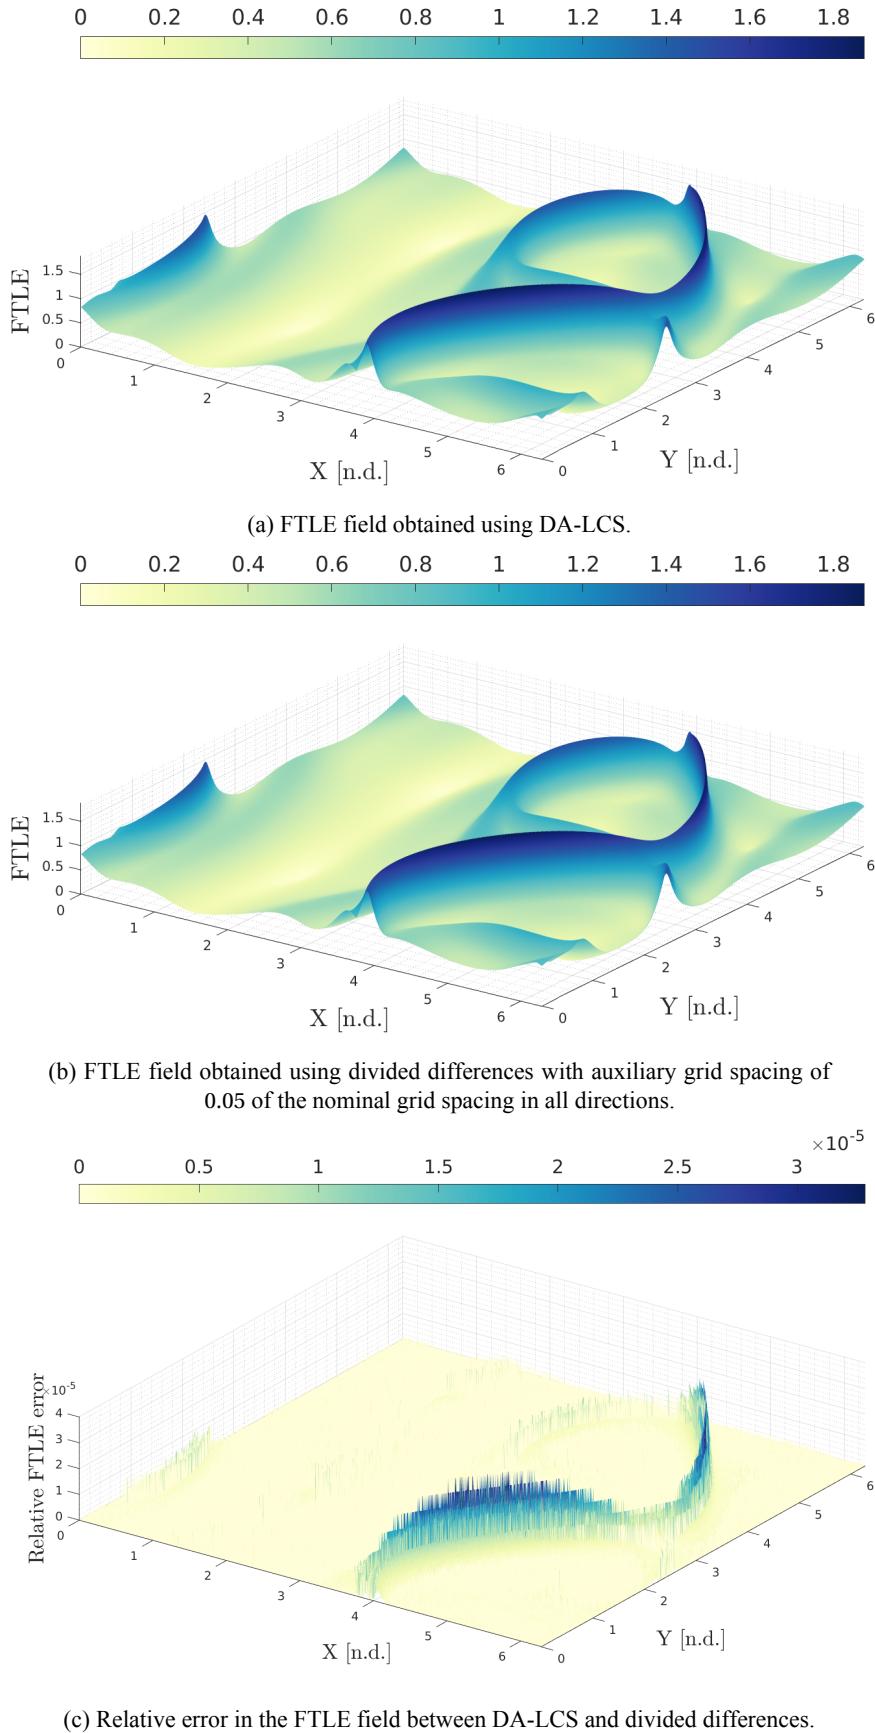


Figure 5.7: Finite-time Lyapunov fields for the steady ABC flow from $t = 0$ to $T = 3$ using DA-LCS and divided differences. The relative error is below 3×10^{-5} , suggesting that computing the FTLE associated with C_0^3 using divided differences is not a major source of error in this example.

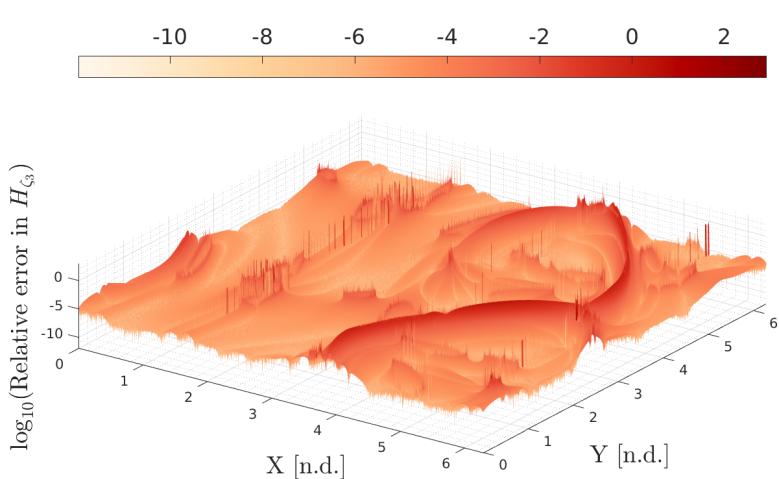
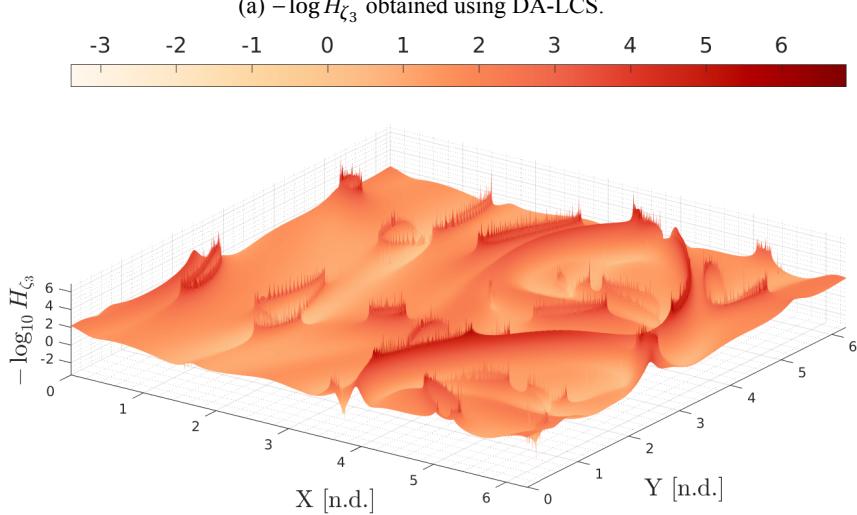
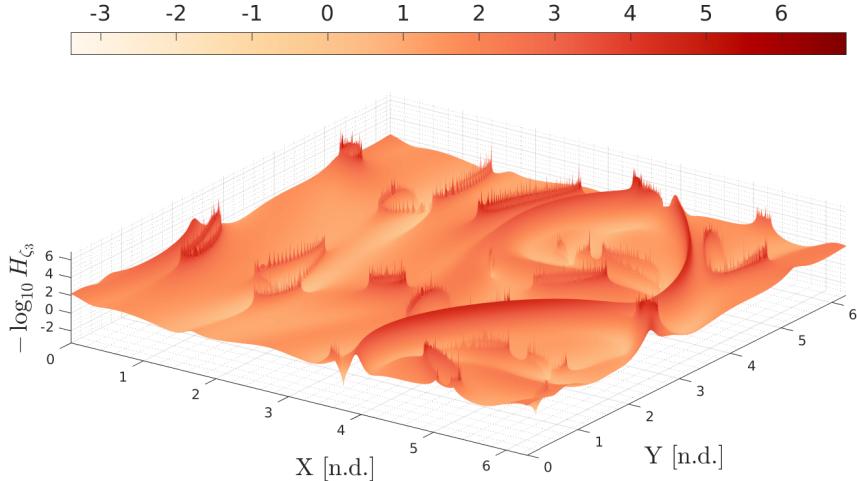


Figure 5.8: Helicity fields for the steady ABC flow from $t_0 = 0$ to $T = 3$ using DA-LCS and divided differences. Again both strongly agree, showing that DA-LCS is working. The DA-LCS structure is a little smoother along the main ridge on the right compared to divided differences, making the identification of seed points more robust.

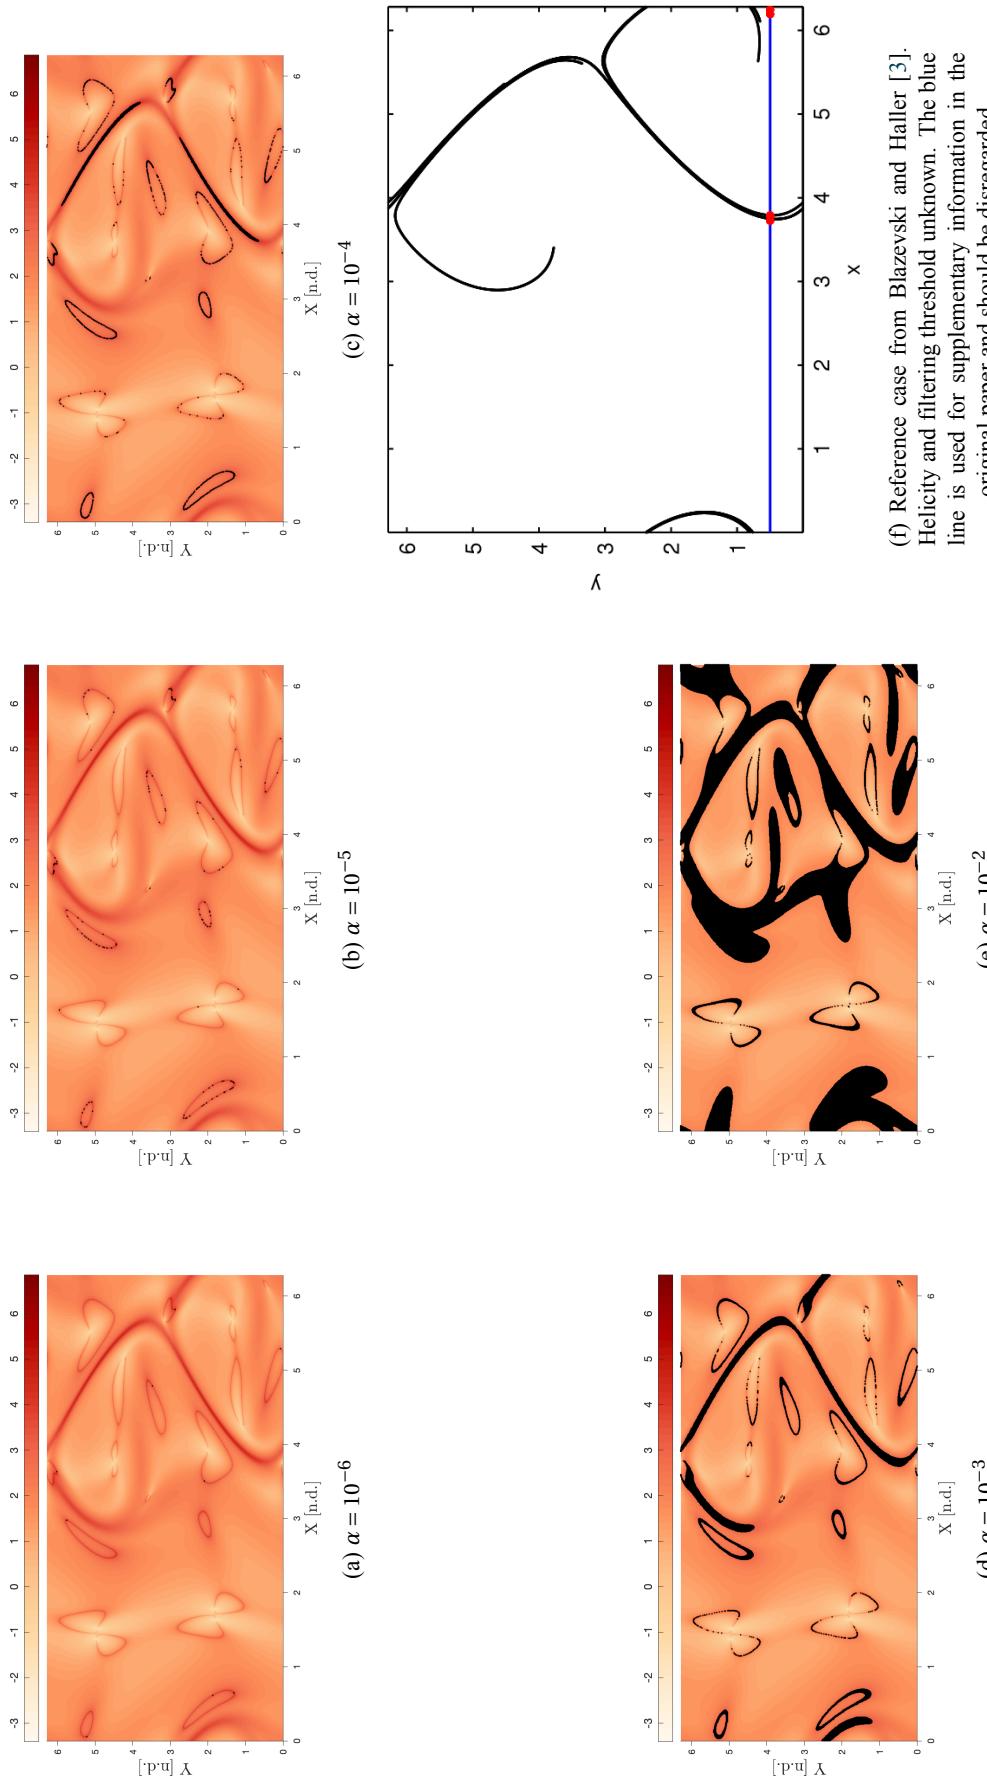


Figure 5.9: The seed points (black) for the streamline ODE obtained for different values of helicity thresholds for the steady ABC flow on the $z=0$ plane, and the reference solution from Blazevski and Haller [3].

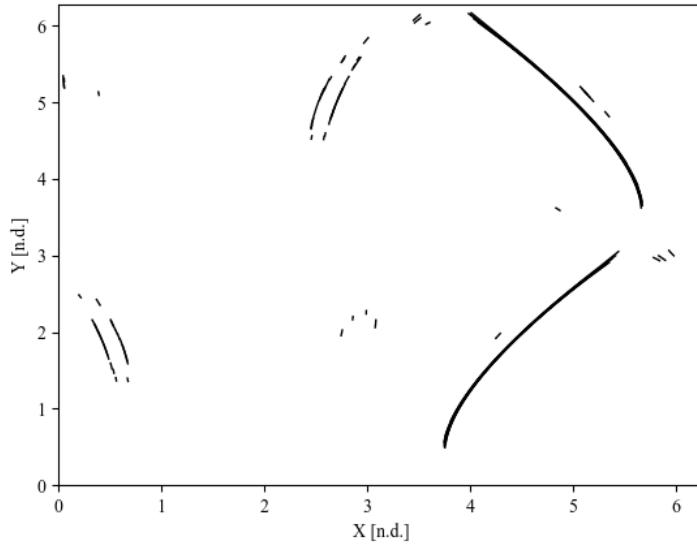


Figure 5.10: Final, filtered strainlines for the steady ABC flow on the $z = 0$ plane computed using DA-LCS. The structure is formed of approximately 53 strainline segments.

search more efficiently for LCS seed points, such as searching on a fixed line or only in a certain region of flow.

The system defined by Equations 5.17-5.19 is integrated forward for 3 non-dimensional time units using the DA-compatible numerical integrator introduced previously, with an integration tolerance of 10^{-13} . A helicity tolerance of $\alpha = 10^{-4}$ is applied to determine seed points and terminate the numerical integration. A minimum distance of $d_F = 0.04$ is used in the strainline segment filtering. Both of these parameters are chosen from visual examination of the helicity field and resulting strainline structure for all of the examples in this chapter.

The FTLE fields on the $z = 0$ plane for this flow, computed using DA-LCS and divided differences, are shown in Figure 5.7a and 5.7b, respectively. I show the relative error between the use of DA-LCS and the use of divided differences on the manually-determined ‘optimal’ grid-size in Figure 5.7c. The relative error is very low throughout the field, which suggests that the computation of C_0^3 and its dominant eigenvalue agrees across the two methods.

In the DA-LCS and divided difference helicity fields on the $z = 0$ plane, shown in Figures 5.8a and 5.8b respectively, some first differences can be seen. The quantitative differences are highlighted in a plot of the relative error between them in Figure 5.8c. The two methods qualitatively agree on the structure of the field, showing that DA-LCS is working. However, the DA-LCS method produces smoother peaks and ridges in the field for the primary features in the flow. This is particularly visible on the main ridge in the bottom right corner around $X = 4$ and $Y = 1$. This makes the identification of seed points in the flow more straightforward and more robust when using DA-LCS.

Method		Time to compute H_{ζ_3} field [s]	Time to compute 100 strainlines [s]	Average function evaluations per unit length
Divided differences	differ-	224.902	4684.689	7498.171
DA-LCS		611.360	1108.571	78.392

Table 5.1: Core time required to compute the LCS on one reference plane for the steady ABC flow using divided differences and DA-LCS on Intel Xeon E5-2670 processors. While DA-LCS is slower to determine the initial H_{ζ_3} field, it is quicker at the integration of a representative set of strainlines and can grow much longer strainlines with the same number of evaluations of Equation 5.13 as divided differences. Importantly, divided differences requires significant grid size tuning, which may make the time required to determine H_{ζ_3} slower overall when used practically.

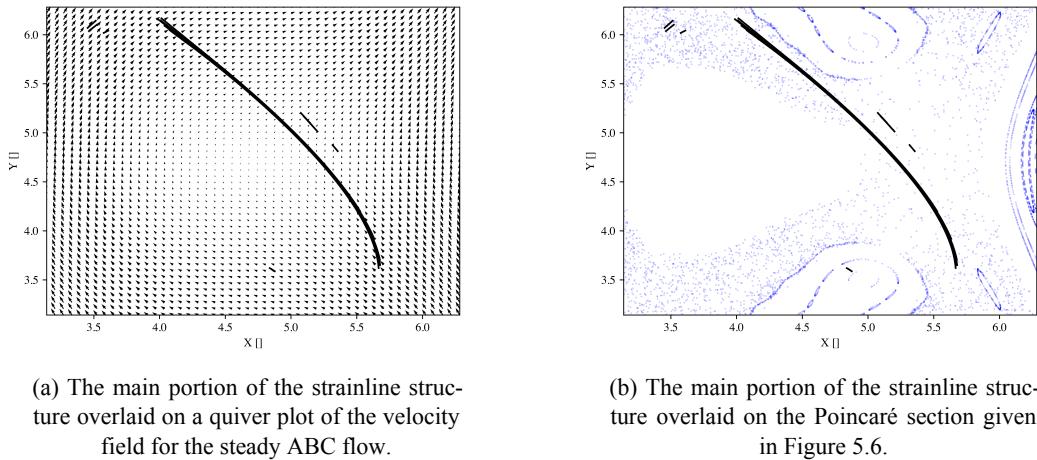


Figure 5.11: The strainline structure found for the steady ABC flow overlaid on two plots highlighting the underlying dynamics.

To highlight the differences between the results of DA-LCS and the literature, the effect of changing the helicity threshold α for the steady ABC flow and the reference set of strainlines from Blazevski and Haller [3] is given in Figure 5.9 with the example of the $z = 0$ plane. The resulting strainlines on the $z = 0$ plane for this flow are shown in Figure 5.10, and follow the expected structure from the helicity field presented in Figure 5.8a using 53 strainline segments. The main structure of the strainlines are also given in Figure 5.11 overlaid on the velocity field for the steady ABC flow and the Poincaré section computed previously. Interestingly, there are no immediately visible features in either of these plots that would suggest the existence of an LCS at these points, though I note that the Poincaré section is computed over a significantly longer time period than the strainlines.

While the strainlines match in identifying parts of the ‘wishbone’-like structure on the right-hand side of the plot, there are differences to reference set in the literature. I now provide several suggestions on where these differences may arise.

Firstly, there is not as much of the main wishbone structure contained in the strainlines in Figure

5.10 as in the reference. However it can be seen that adjusting α in Figure 5.9, which is a key parameter in determining the LCS, can signal more of the wishbone structure as the LCS on the $z = 0$ plane and match the literature more closely. Further detail on how I choose the threshold is given in the following test case. I note that the lower peaks of helicity when using divided differences, with reference to the relative error in the helicity field in Figure 5.8c, means that a given threshold also encompasses more of this structure automatically when using divided differences.

Next, I note the existence of several ‘loops’ in the helicity field computed using DA-LCS, particularly in the left-hand side of the field, that are not present in the literature. The strainline segments at these points often grow transverse to the ridges at certain points and do not track along the ridge as would be expected. This behaviour is also present when computing the helicity with divided differences and so can be assumed to be the behaviour of the underlying system rather than DA-LCS. Given this, and the existence of similar structure in the helicity field in Palmerius, Cooper, and Ynnerman [143], I do not investigate this issue further.

I believe these small strainline segments are not present in Blazevski and Haller [3] due to a combination of three reasons. First, much of this structure will have been missed by the largely reduced 600×10 grid resolution used there. Secondly, in many of Haller’s investigations [132, 138, 199, 235] a minimum length on the LCS is enforced to keep only those surfaces which exert a maximal influence on flow behaviour. This would remove all of these small loops while retaining only the main wishbone-like structure. While this is not explicitly mentioned in Blazevski and Haller [3], a combination of this approach with adjustment of the helicity threshold could yield the structure from the literature exactly. Lastly, many of the loops arise from seed points that are at the top of ‘spikes’ in the helicity field. It can be the case that if the numerical integrator used to integrate the strainline ODE takes too large of a step due to an improper step size or too high an error tolerance, then it immediately leaves the low-helicity region in both directions of $\pm \zeta_n$, and the integration is immediately terminated without taking a single valid step below α . This also explains why some seed points do not yield strainlines.

The ‘spiky’ nature of these loops, found partially as a result of the dense sampling of initial conditions, also means that those strainline segments that do grow are of exceedingly small length rather than being a continuous, low-helicity structure, as numerical integration of the strainline ODE is terminated immediately upon leaving the ‘tip’ of the spike. Ensuring that only strainlines which form part of a larger structure which exerts maximal influence on nearby flow are captured, as per the definition of an LCS, motivates the use of a minimum strainline length equal to the filtering distance. In practice, one may wish to only filter the subset of longest strainlines, or only those with minimal average helicity. The use of a reduced grid, or a more sophisticated search strategy for low-helicity points, would also avoid the inclusion of such spurious points.

The computational and numerical performance of DA-LCS is now discussed, using the steady ABC flow as an example. The total time to compute the full LCS on 48 2.0GHz Intel Xeon

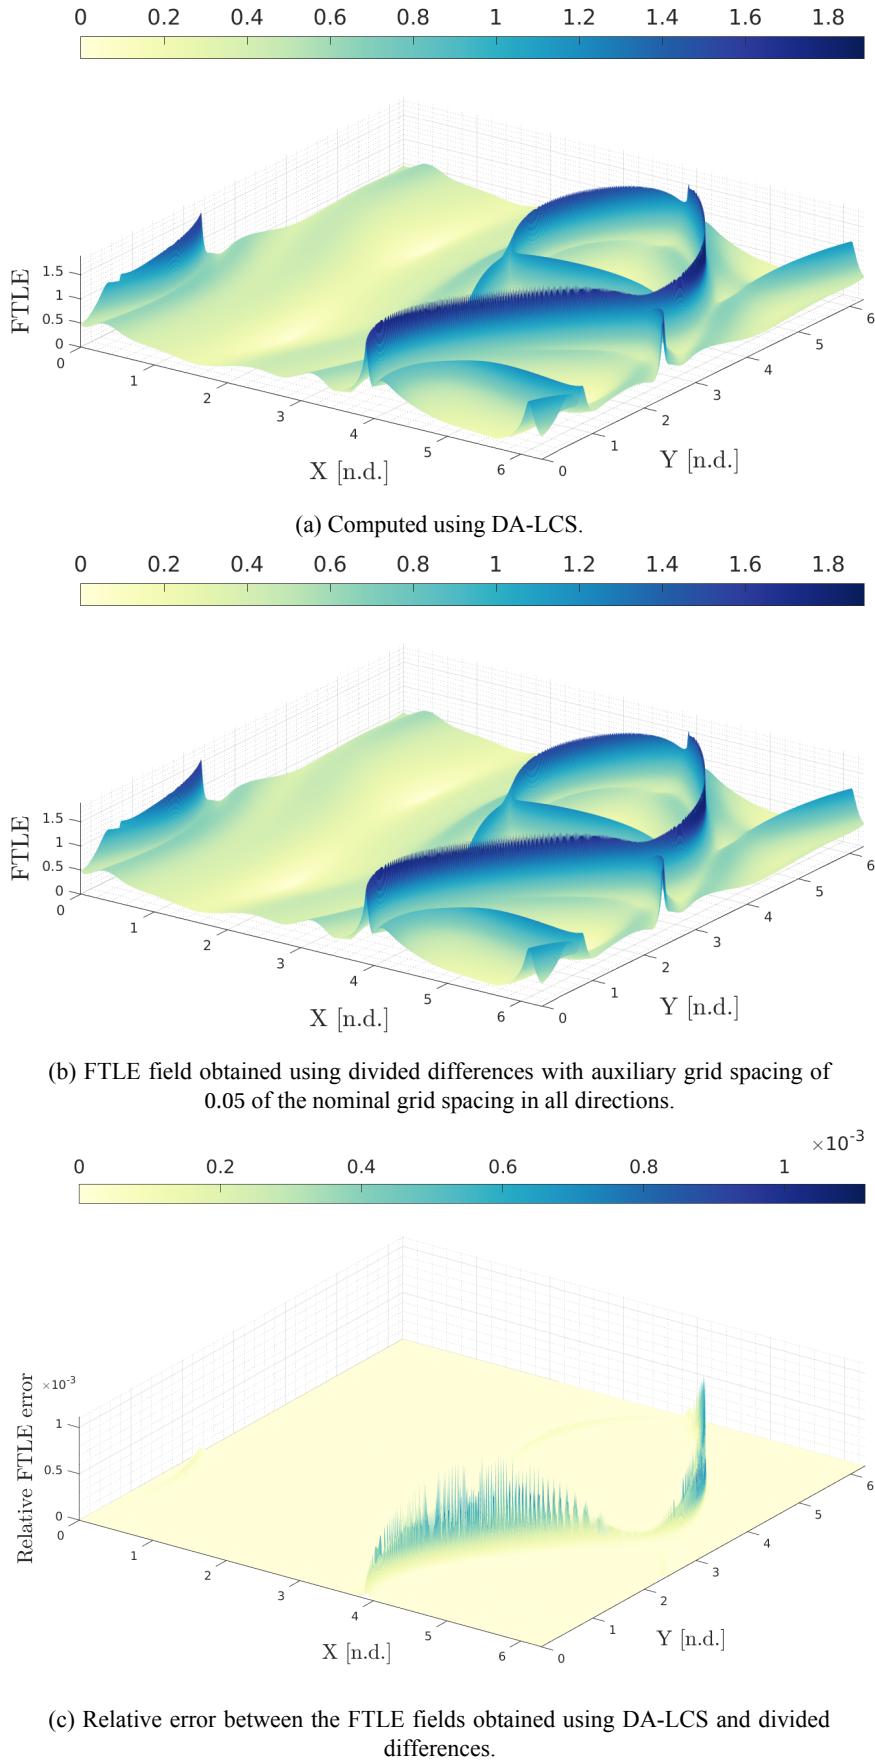


Figure 5.12: Finite-time Lyapunov exponent fields for the periodic ABC flow from $t_0 = 0$ to $T = 4.0$, obtained using DA-LCS and divided differences. Again, the FTLE field agrees between divided differences and DA-LCS, suggesting divided differences on the correct auxiliary grid in this case accurately approximates the FTLE associated with C_0^4 .

E5-2670 processors is given in Table 5.1, broken down by the time required to obtain the initial H_{ζ_3} field and then a representative set of 100 strainlines. The set of 100 strainlines is chosen to be the 100 points with lowest H_{ζ_3} , integrated until the running average of helicity rises above 10 times the initial value. Visual inspection of the initial conditions confirms that the seed points are sufficiently ‘close’ in both divided differences and DA-LCS that they are assumed to represent the same behaviour.

I find that DA-LCS is slower than divided differences for computing the initial helicity field since two orders are computed, requiring more CPU instructions per operation, and because fewer optimisations can be made by the compiler compared to native double-precision types. However, since DA-LCS requires no tuning of grid size, this computational deficit is eliminated as soon as more than two trial computations of the LCS using divided differences has to be performed to obtain the ‘optimal’ grid size in every dimension. Moreover, owing to better numerical performance, the strainline integration is approximately four times faster using DA-LCS than using divided differences, since the integrator can take larger steps than with divided differences while still controlling the error in the integration of Equation 5.13. I also find that the strainlines obtained with DA-LCS are on average 10 times longer than when using divided differences for the representative set here; this may mean that more sophisticated search methods for identifying seed points, such as the method of searching on a fixed line mentioned earlier, would be more feasible in DA-LCS. Both improvements in strainline integration are due to the elimination of the numerical noise introduced by divided differences, which is not present in DA-LCS.

5.5.2 Periodic Arnold-Beltrami-Childress Flow

A time-periodic version of the Arnold-Beltrami-Childress flow is now considered, with equations of motion

$$\dot{x} = (A + 0.1 \sin t) \sin z + C \cos y \quad (5.20)$$

$$\dot{y} = B \sin x + (A + 0.1 \sin t) \cos z \quad (5.21)$$

$$\dot{z} = C \sin y + B \cos x. \quad (5.22)$$

The hyperplanes \mathcal{S} and grids are the same as in the case of the steady ABC flow, but now with integration times $t_0 = 0$ and $T = 4$.

Mirroring the analysis in the steady case, the FTLE fields for both DA-LCS and divided differences are shown in Figures 5.12a and 5.12b, respectively, with the relative error presented in Figure 5.12c. Again, there is little qualitative difference between the two fields. The differences in smoothness in the helicity fields are, however, more pronounced between Figures 5.13a and 5.13b. The main wishbone-like structure is particularly ‘spiky’ when using divided differences, which conceptually requires the strainline integration to track a numerically noisy ridge. With DA-LCS, there is a smooth, well-defined ridge of consistently low helicity for the algorithm to track with much lower numerical noise; in fact, the helicity threshold here is approximately two

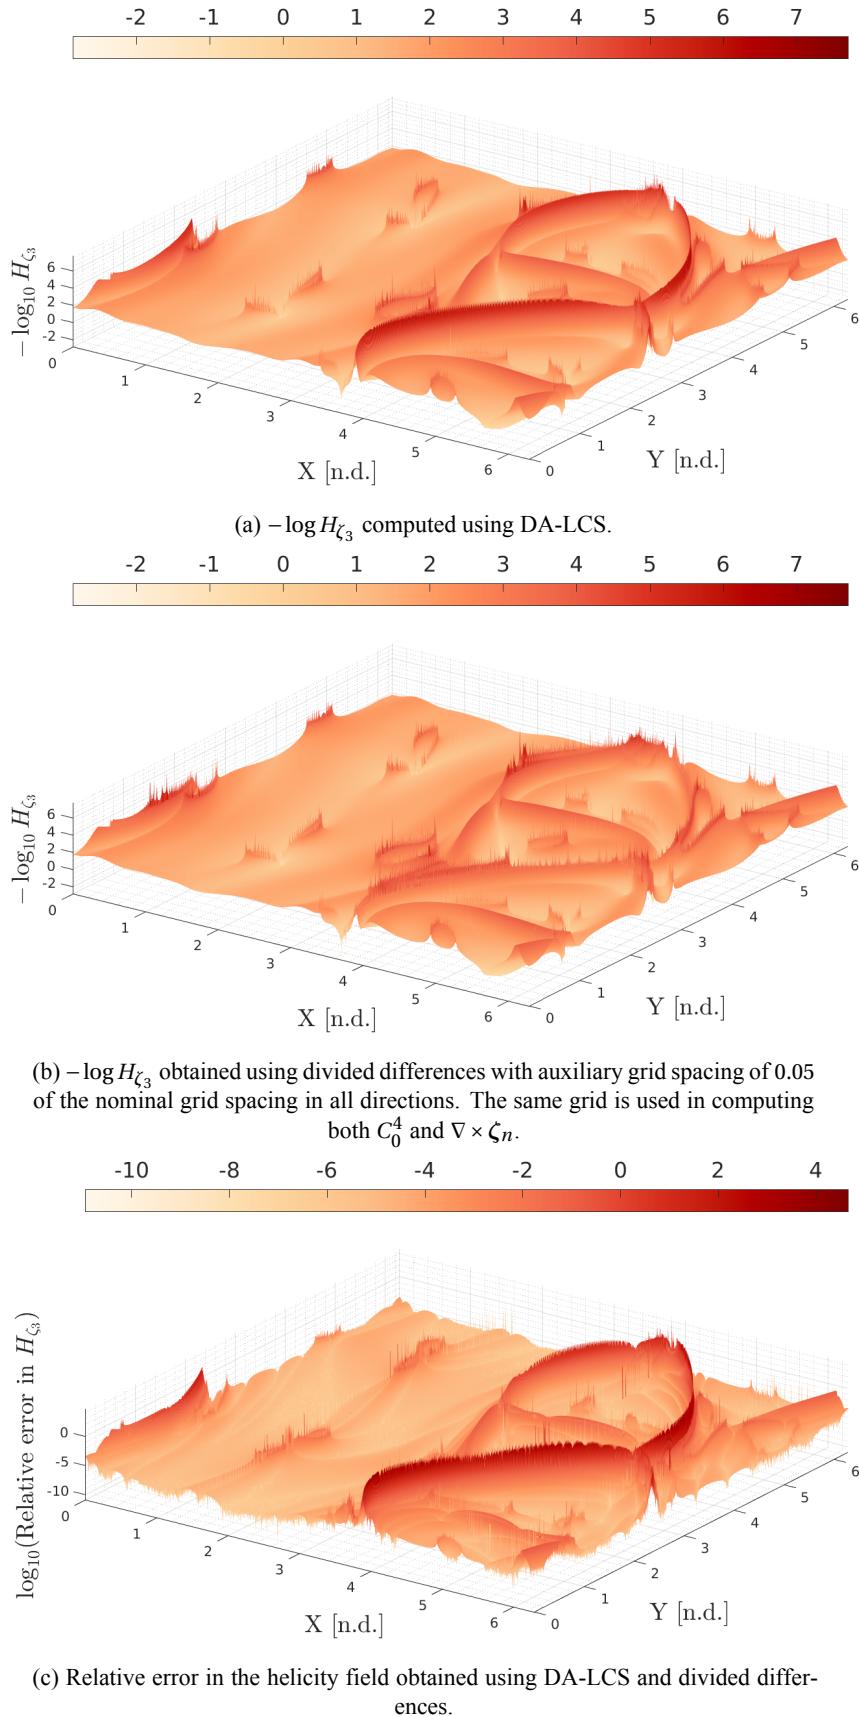


Figure 5.13: Helicity fields for the periodic ABC flow computed using DA-LCS and divided differences from $t_0 = 0$ to $T = 4.0$. Here DA-LCS highlights in particular the main ridge on the right more clearly and smoothly than divided differences.

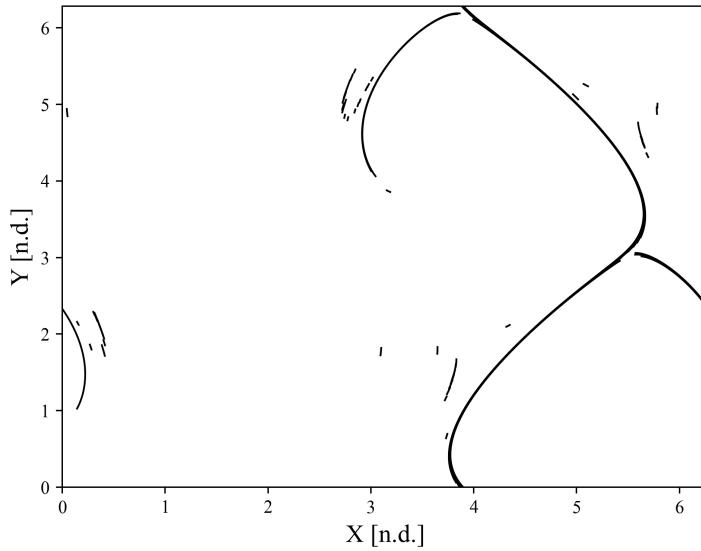


Figure 5.14: Final strainlines for the periodic ABC flow on the $z = 0$ plane computed using DA-LCS, after filtering. The strainline structure is composed of approximately 50 strainline segments.

orders of magnitude lower than used in literature but recovers qualitatively similar structures, and divided differences reports a relative error to DA-LCS of up to 10^4 on the main ridge. This means that the identification of seed points and the tracking of low-helicity ridges is again more robust in DA-LCS, which comes without the need to adjust grid sizes or derive and implement variational equations.

The effect of changing α for the periodic ABC flow is given in Figure 5.15 with the reference set of strainlines from the literature. A helicity tolerance of $\alpha = 5 \times 10^{-5}$ is used, with a distance threshold $d_F = 0.04$. In general, I select the helicity threshold to uncover the first emergence of a continuous structure, in this case the wishbone-like structure on the right-hand side of the Figure. While the threshold could have been raised to encompass more of this structure, the repelling LCS is necessarily the most repelling structure and thus lowest-helicity structure in a flow. I therefore choose not to include this structure in the LCS on the $z = 0$ reference plane.

Finally, the strainlines on the $z = 0$ plane for this system computed using DA-LCS are shown in Figure 5.14. Approximately 50 strainline segments determine the full strainline structure on the $z = 0$ plane for this example.

There are again differences in the strainline structure compared to the literature. As before, adjustment of α , the use of divided differences which signals more of the ridge as being at the same helicity than DA-LCS (Figures 5.13b and 5.13c), or the enforcement of a minimum length for strainlines would yield structures closer to the literature.

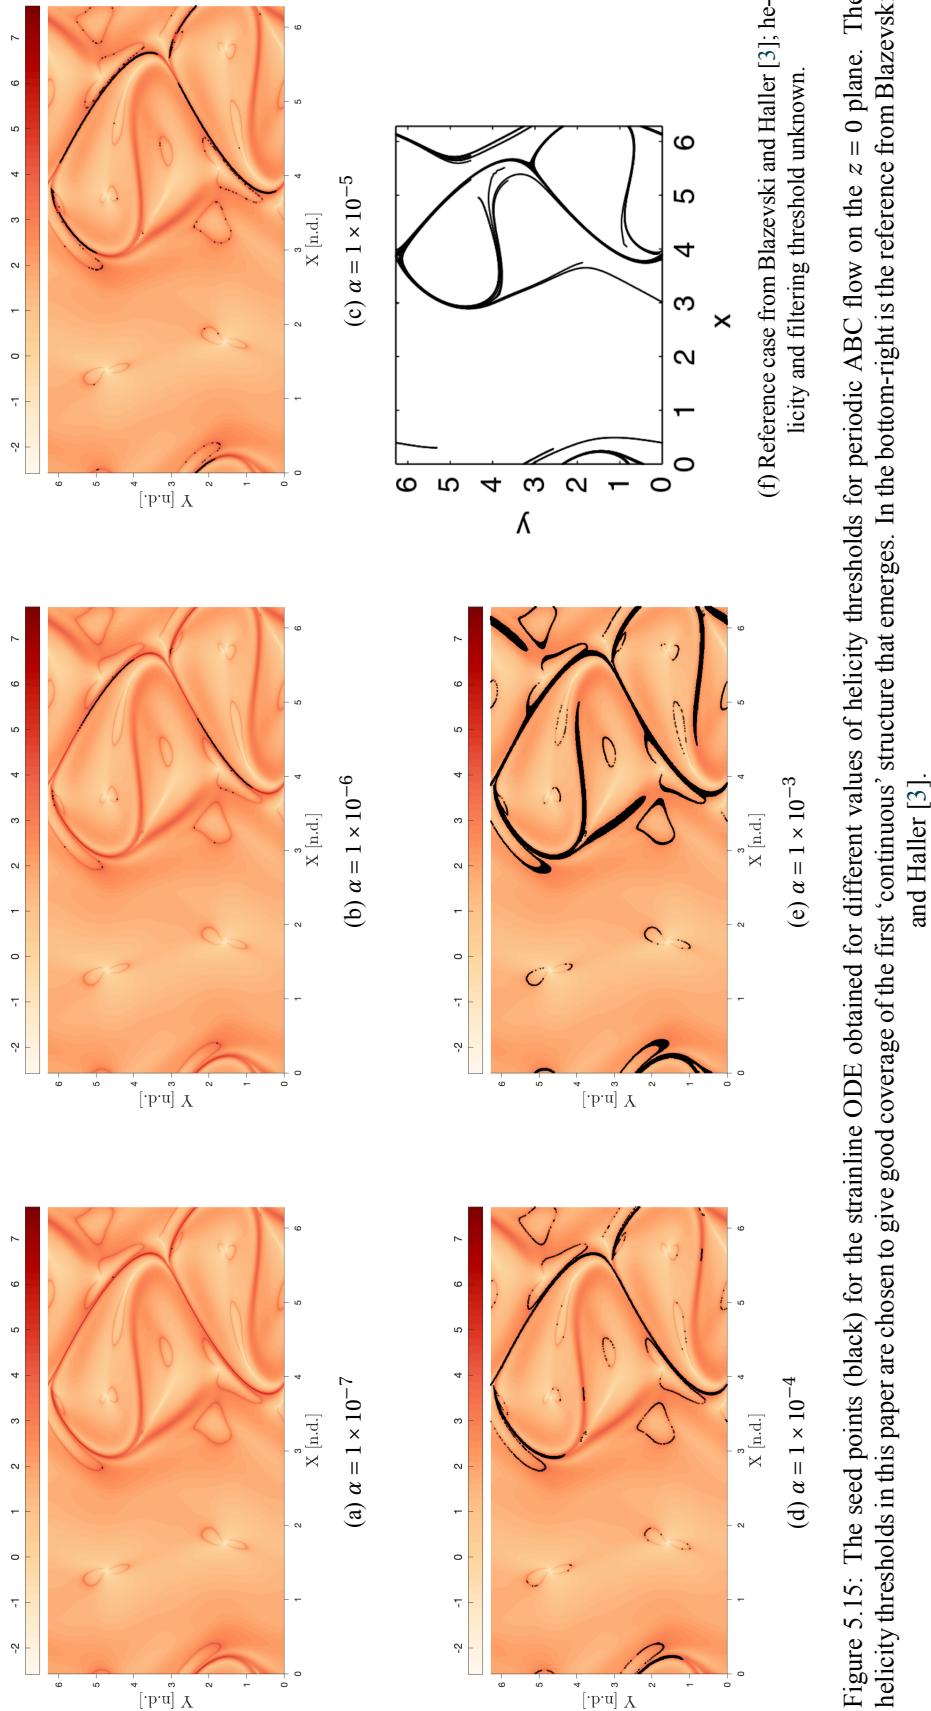


Figure 5.15: The seed points (black) for the strainline ODE obtained for different values of helicity thresholds for periodic ABC flow on the $z=0$ plane. The helicity thresholds in this paper are chosen to give good coverage of the first 'continuous' structure that emerges. In the bottom-right is the reference from Blazevski and Haller [3].

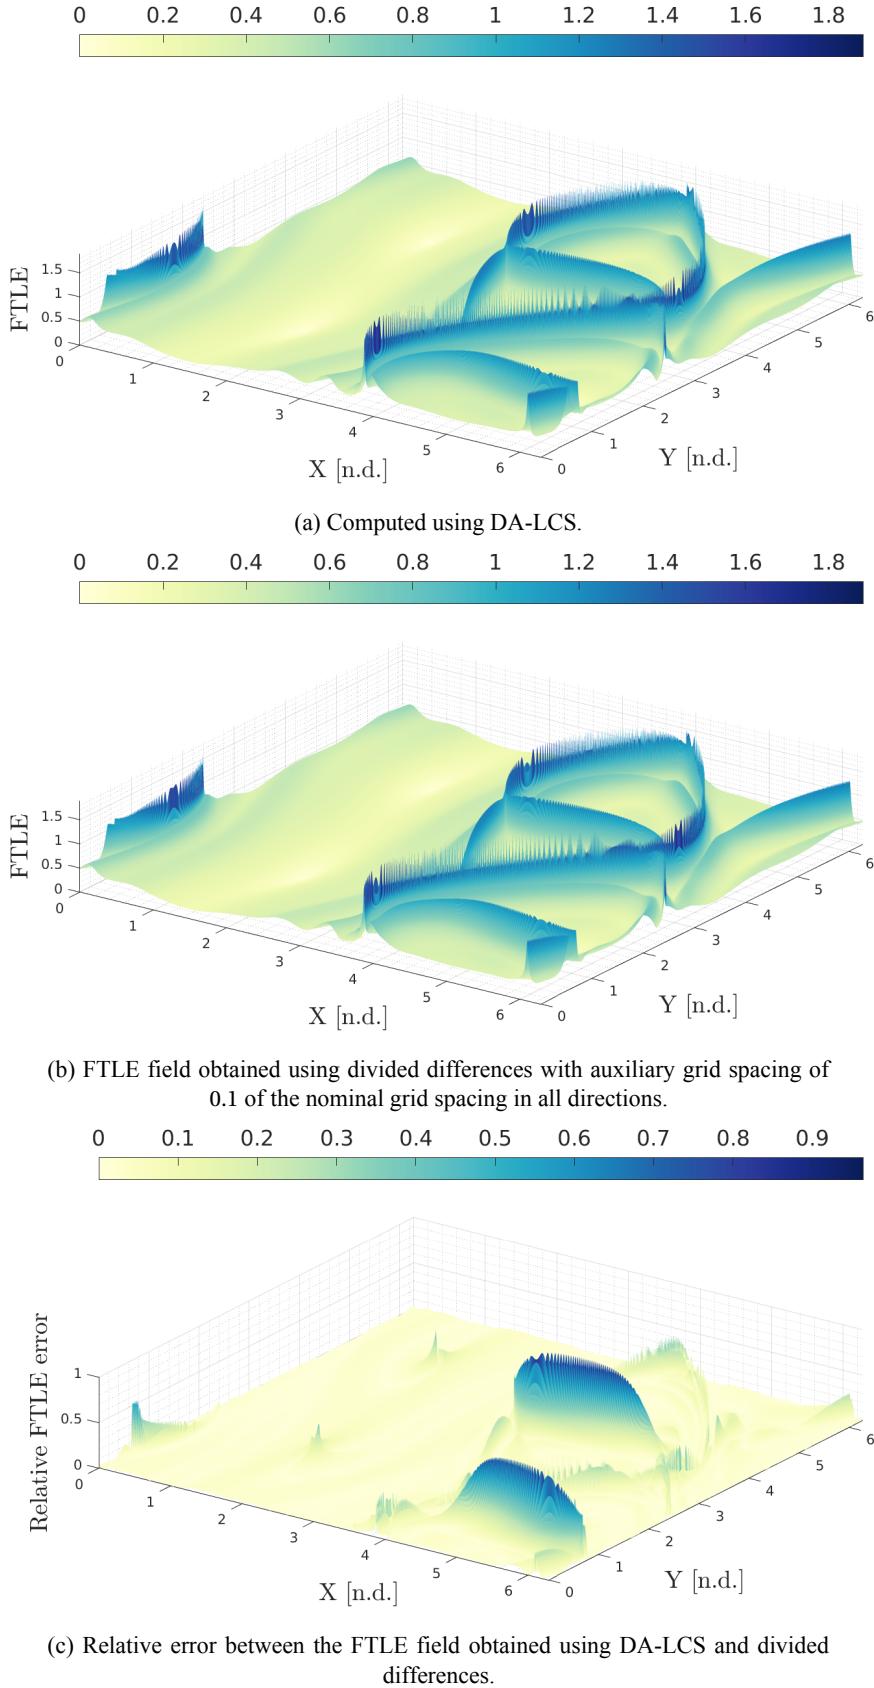


Figure 5.16: Finite-time Lyapunov exponent field for the chaotically-forced ABC flow and an integration time from $t_0 = 0$ to $T = 5.0$. Differences are beginning to become visible on the ‘ends’ of the main wishbone-like structure when using divided differences due to the spiky FTLE values.

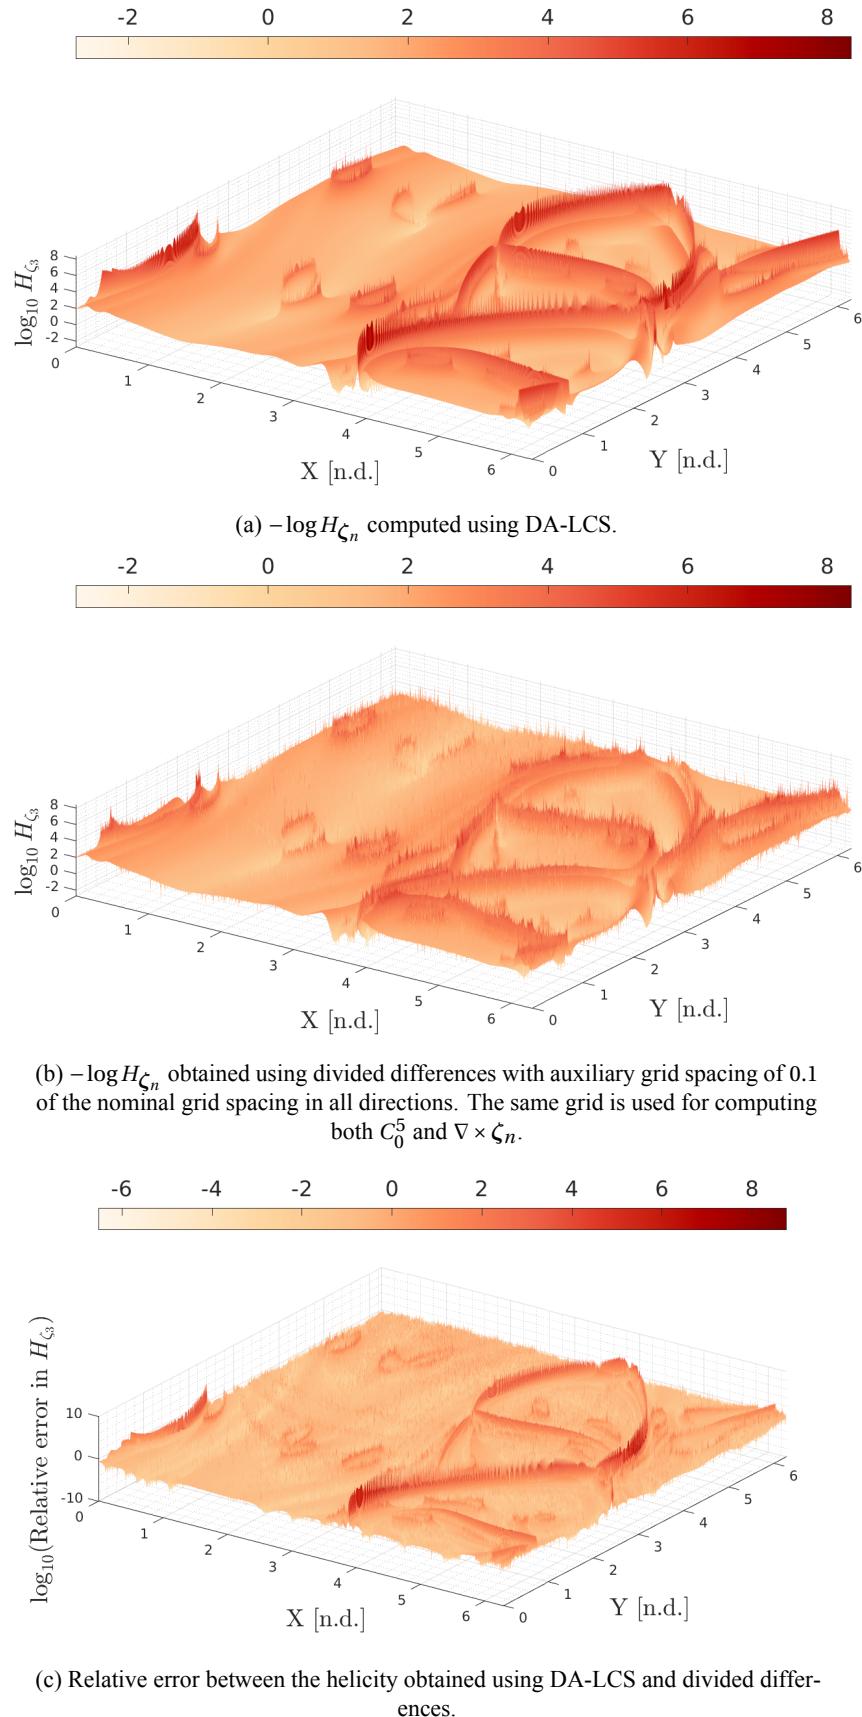


Figure 5.17: Helicity fields for the chaotically-forced ABC flow from $t_0 = 0$ to $T = 5.0$. DA-LCS produces visibly better-defined ridges, helping to robustly identify seed points. However, the relative errors are very high due to the spiky nature of the ridges in both DA-LCS and divided differences.

5.5.3 Chaotically-forced Arnold-Beltrami-Childress flow

Following Blazevski and Haller [3], I now demonstrate that DA-LCS is robust under perturbations from a chaotic forcing function $g(t)$. The motion is forced by a chaotic Duffing oscillator, with equations of motion given by

$$\dot{x} = (A + 0.1 \sin t) \sin z + C \cos y \quad (5.23)$$

$$\dot{y} = B \sin x + (A + 0.1 g(t)) \cos z \quad (5.24)$$

$$\dot{z} = C \sin y + B \cos x \quad (5.25)$$

where $g(t)$ is the x -coordinate of the solution to the Duffing equation

$$\ddot{x} = -\delta \dot{x} - \beta x - \alpha x^3 + \gamma \cos(\omega t). \quad (5.26)$$

with parameters $\alpha = 1$, $\beta = -1$, $\gamma = 0.3$, $\delta = 0.2$, $\omega = 1$.

The computational grid is again the same as for the previous test cases involving the ABC flow, including the hyperplanes $\mathcal{S} = \{(x, y, z) \in [0, 2\pi]^3 : z = s_1\}$, $s_1 = 0.0, 0.005, 0.01, \dots, 0.1$, but a longer integration time of $T = 5$ is used to match the literature. Again, a helicity tolerance of $\alpha = 5 \times 10^{-5}$ is used with a filtering distance of $d_F = 0.05$; the effect of changing the helicity threshold is presented in Figure 5.18.

The FTLE fields computed using DA-LCS and divided differences are again shown in Figure 5.16a and Figure 5.16b, respectively, with the relative error shown in Figure 5.16c. Some differences in the FTLE are beginning to emerge in portions of the main wishbone-like structure. The helicity fields are shown in Figures 5.17a and 5.17b, respectively, and the relative error is shown in Figure 5.17c. The helicity field in particular now exhibits a significant difference compared to the two previous cases. Using DA-LCS, I am able to resolve a relatively smooth ridge, whereas the use of divided differences leads to noticeable numerical noise throughout the field as well as an overall much higher helicity. Note that the relative error in the helicity field is high in part due to the spiky nature of the ridge, rather than just the relative performance of DA-LCS compared to divided differences. Thus, for this example no claims are made as to the robustness or performance of DA-LCS, other than that it is also able to resolve the LCS under a chaotic forcing function.

The strainlines for this system on the $z = 0$ plane computed using DA-LCS are presented in Figure 5.19. A total of 57 strainline segments give the full structure on the $z = 0$ plane.

I note significant differences in the strainlines for this example compared to the reference set, which do not disappear even when increasing α to the levels shown in this thesis. However, if the LCS is to be robust under a forcing function it should match closely the result of the periodic ABC flow, which the strainlines recovered using DA-LCS do. This discrepancy is thus not investigated further.

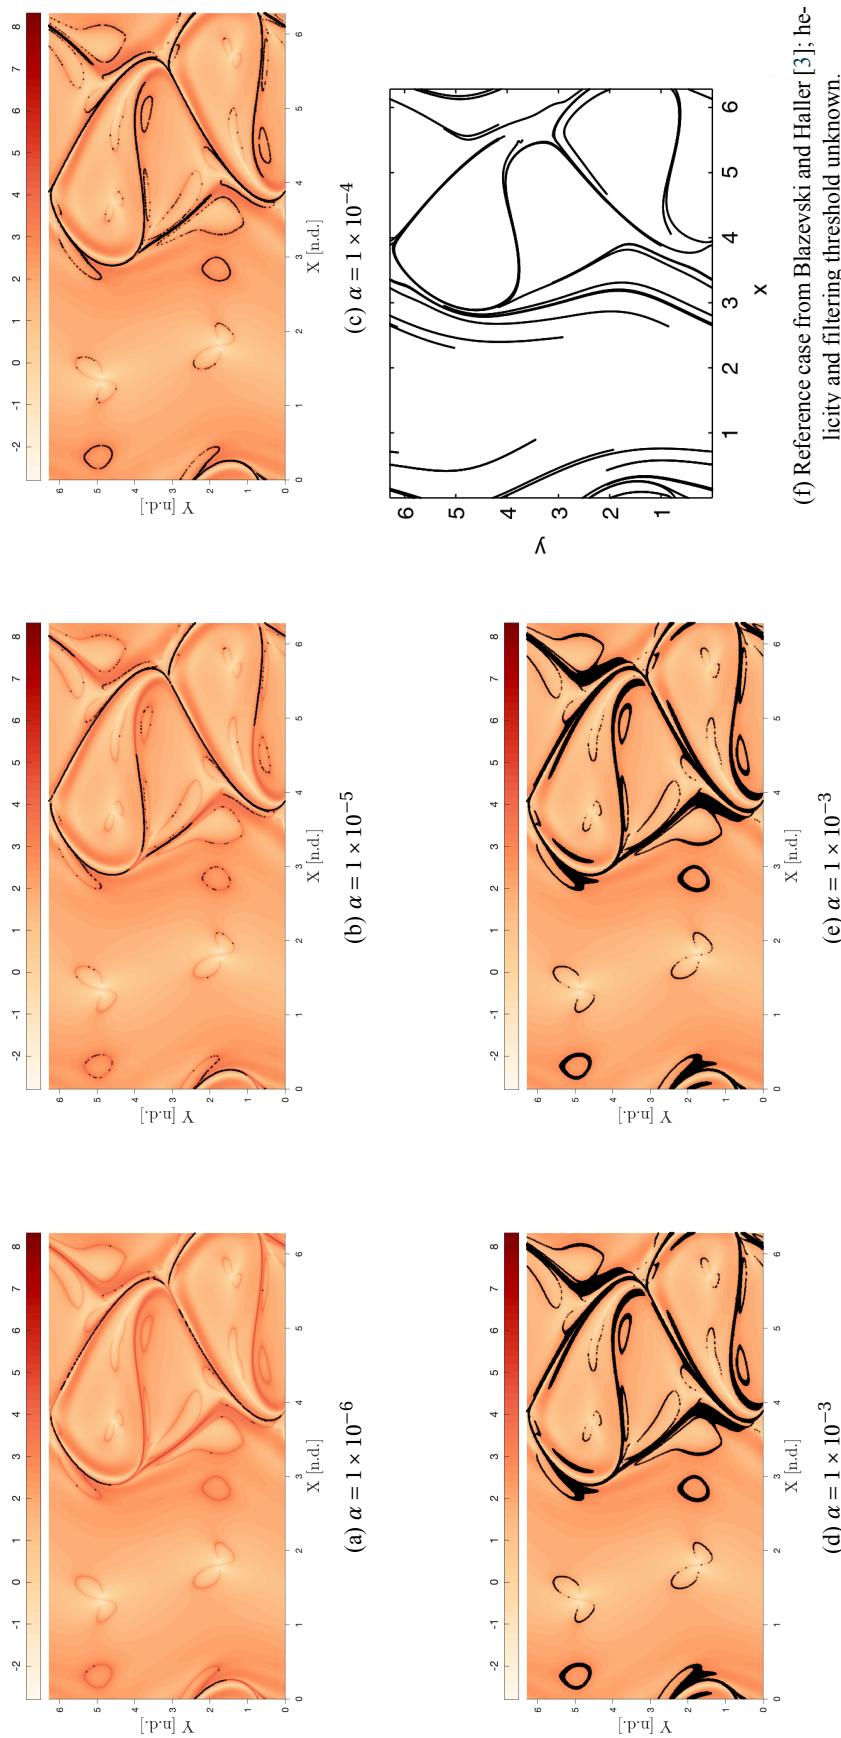


Figure 5.18: The seed points (black) for the strainline ODE obtained for different values of helicity thresholds for the forced ABC flow on the $z=0$ plane.

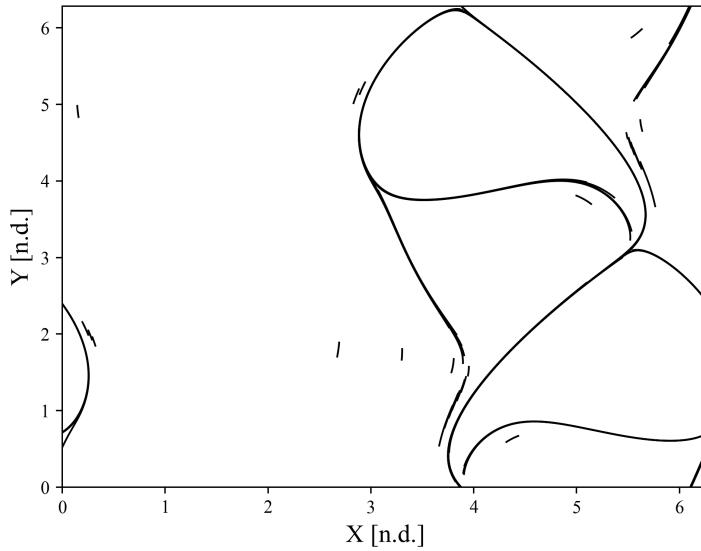


Figure 5.19: Final strainline structure on the $z = 0$ plane for the chaotically-forced ABC flow computed using DA-LCS. The structure is formed of 57 individual strainline segments.

5.6 Conclusion

This chapter has introduced DA-LCS, an improved numerical method for determining hyperbolic Lagrangian Coherent Structures in time-dependent dynamical systems. I introduced Differential Algebra, which uses the algebra of polynomials to approximate functions directly in a computer environment, and showed how it can be used to directly construct high-order Taylor expansions of the flow and its derivatives. I have shown that in combination with a modified power law one can construct a highly-accurate LCS based solely on the underlying dynamics of the system. The effectiveness of the method has been demonstrated through applications to common variations of the Arnold-Beltrami-Childress flow from the literature, highlighting more robust low-helicity regions and more accurately determining FTLE fields in the more complex examples. DA-LCS also constructs the LCS automatically and without any *a priori* information, requiring no additional implementation beyond the dynamics of the system.

In the following chapter, I apply DA-LCS to astrodynamics problems to assess its performance and whether it can reveal similar insight in the numerically challenging astrodynamics examples found in the literature.

6

Lagrangian Coherent Structures in the Elliptic-Restricted Three-body Problem

This chapter presents a deeper investigation into the use of Lagrangian Coherent Structures in the Elliptic-Restricted Three-body Problem. Following the last chapter, where it was shown that DA-LCS can accurately determine three-dimensional LCS on ‘toy’ problems from the literature, the method is applied to the ER3BP to examine whether it can provide insight in astrodynamics systems. As part of this investigation, the effect of integration time and orbit parameterisation is discussed, and mitigating strategies to avoid common numerical pitfalls as a result of limits of floating-point arithmetic are given. Parts of this chapter have been accepted for publication as Tyler and Wittig [13], submitted for publication as Tyler and Wittig [15] and presented at the 2022 Astrodynamics Specialist Conference in Charlotte, North Carolina in August 2022 [14] and the 73rd International Astronautical Congress in Paris, France in September 2022 [16].

6.1 Introduction

The new DA-LCS numerical method outlined in the previous chapter has been shown to accurately reproduce test cases from the literature without any *a priori* knowledge. More importantly, it has improved the ability to identify seed points and regions of consistently low-helicity on which the strainlines that form part of the LCS track, improving robustness, and greatly reduced the presence of numerical noise in the FTLE field in the most complex of examples. The equivalent quantities to the ‘helicity’ in the two-dimensional formulation of LCS were the terms that were most troublesome when attempted in the literature [7, 8]. Thus, this chapter now seeks to apply DA-LCS to three-dimensional astrodynamics systems to assess its performance in providing insight similar to the invariant manifolds used to perform the investigation in chapter 4.

This chapter is formed of two applications of DA-LCS to problems from astrodynamics. Both applications study the three-dimensional Sun-Mars Elliptic-Restricted Three-body Problem (section 2.2.2), mimicking the system investigated in Ros Roca [7] and Parkash [8]. The first investigation is designed to be a proof-of-concept to ascertain the suitability of DA-LCS in astrodynamics, making several simplifications about the initial velocity of the orbits to reduce the complexity of the problem while still providing an appropriate benchmark for the method. The second investigation builds on the first by aiming to determine the ‘operating envelope’ of the method, greatly increasing the numerical complexity by a more challenging and more physical choice of initial velocity. It then goes on to investigate practically what the effect of the repelling hyperbolic LCS is on neighbouring trajectories and how the definition of initial position changes the result found.

The structure of this chapter now follows. I first elucidate how the three-dimensional LCS algorithm can be used to sample the six-dimensional phase space of the ER3BP and introduce different ways of parameterising the initial condition in section 6.2. How the full, three-dimensional LCS is constructed from the strainlines is introduced in section 6.3, before a single parameterisation is used in a proof-of-concept study of the LCS in astrodynamics in section 6.4 to assess whether DA-LCS can provide insight and its relative performance to standard approaches. The more numerically complex example is shown in section 6.5, and a detailed analysis of the behaviour of the LCS in three dimensions, the effect of orbit parameterisation, and the integration time, is performed. Lastly, the common numerical pitfalls are presented in section 6.5.3.

6.2 Definition of Initial Conditions

The ER3BP lives in a phase space defined in \mathbb{R}^6 . However, the DA-LCS method is explicitly designed for a CGST that is 3×3 in dimension and hence represents the behaviour of a three-dimensional system. To simplify visualisation and enable analysis of the ER3BP, I seek a way of sampling the six-dimensional phase space with only three free variables. To achieve this, a three-dimensional sub-manifold is embedded in the six-dimensional phase space on which the LCS is computed. Using a sub-manifold to reduce the dimensionality of the problem is a technique used in both the MSc Theses and investigations into ballistic capture [205, 206, 236].

I arrive at this embedding by first parameterising a region of interest in three dimensions in a non-rotating frame about Mars, and then defining a transformation $\Psi : \mathbb{R}^3 \mapsto \mathbb{R}^6$ to complete the full phase space and transform the initial condition into the ER3BP rotating-pulsating frame. After propagation under the ER3BP equations of motion, the inverse transformation $\Pi : \mathbb{R}^6 \mapsto \mathbb{R}^3$ takes the final condition back into the original parameterisation. This process is shown below for an arbitrary starting parameterisation $(\kappa, \beta, \zeta) \in \mathbb{R}^3$:

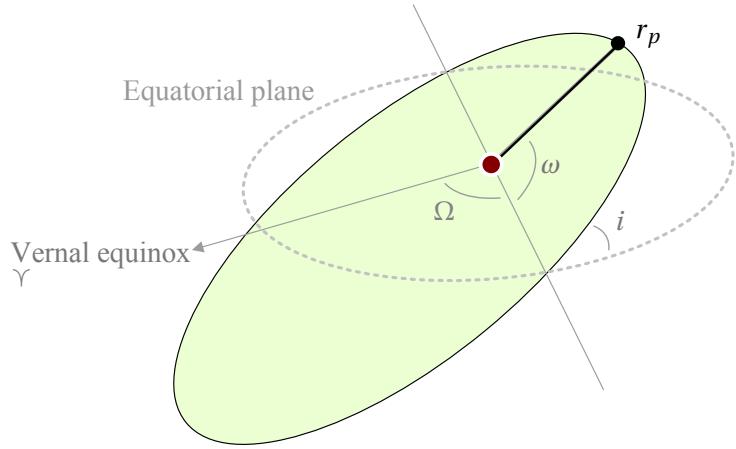


Figure 6.1: Graphical depiction of the parameterisation of an initial condition using standard Keplerian elements. The light green plane represents an example hyperplane of fixed i .

$$\begin{array}{ccc}
 (x, y, z, x', y', z') & \xrightarrow{f} & (x, y, z, x', y', z')_F \\
 \uparrow \Psi & & \downarrow \Pi \\
 (\kappa, \beta, \varsigma) & \xrightarrow{\Psi \circ f \circ \Pi} & (\kappa, \beta, \varsigma)_F
 \end{array}$$

where the subscript F implies final states. In the non-rotating frame, the distance unit is chosen such that the Sun-Mars distance is unity; the rationale behind this is given in section 6.5.3. I refer to $(\kappa, \beta, \varsigma)$ as the parameterisation of the manifold and the map $\Psi \circ f \circ \Pi$ as the transformation.

It must be noted that each choice of Π and Ψ introduced above induces a new dynamical system, $\Psi \circ f \circ \Pi$. While after sufficient time the time-dependent derivatives of f will dominate the time-independent derivatives Π and Ψ , it is not mathematically correct to compare the parameterisations to study f alone. Thus, comparisons given in the following sections are made with respect to the ability for each $\Pi - \Psi$ pair to highlight features of f useful for LCS computation in astrodynamics.

Three pairings of Ψ and Π are explored in this chapter and are elaborated in the following. A key advantage of DA-LCS is that for all Ψ and Π given below, provided the operations are coded using DA arithmetic the derivatives are computed fully automatically and so there is no need to manually derive or implement derivatives. This is particularly useful when considering that these transformations are in some cases non-trivial and otherwise very difficult to differentiate manually.

6.2.1 Orbital elements

In this transformation, the initial position and velocity are given by standard Keplerian orbital elements in the non-rotating frame, following their use in the literature to define initial conditions for ballistic capture in three dimensions [205].

The sub-manifold is parameterised by the radius of periapsis r_p , inclination i and argument of periapsis ω of a purely Keplerian orbit around Mars. The hyperplanes \mathcal{S} are given by fixed values of i (Figure 6.1). The full element set is obtained by fixing the remaining orbital elements.

The full orbital element set at each point is converted to equivalent Cartesian position and velocity in the non-rotating frame (subscript I), which is then transformed into the ER3BP rotating-pulsating frame (subscript R) and propagated under the full equations of motion (Equation 6.1).

$$\Psi_{OE} : \begin{Bmatrix} r_p \\ i \\ \omega \end{Bmatrix}_I \rightarrow \begin{Bmatrix} r_p \\ e \\ i \\ \Omega \\ \omega \\ M \end{Bmatrix}_I \rightarrow \begin{Bmatrix} x \\ y \\ z \\ \dot{x} \\ \dot{y} \\ \dot{z} \end{Bmatrix}_I \rightarrow \begin{Bmatrix} x \\ y \\ z \\ x' \\ y' \\ z' \end{Bmatrix}_R \quad (6.1)$$

$$\Pi_{OE} : \begin{Bmatrix} x \\ y \\ z \\ x' \\ y' \\ z' \end{Bmatrix}_R \rightarrow \begin{Bmatrix} x \\ y \\ z \\ \dot{x} \\ \dot{y} \\ \dot{z} \end{Bmatrix}_I \rightarrow \begin{Bmatrix} r_p \\ e \\ i \\ \Omega \\ \omega \\ M \end{Bmatrix}_I \rightarrow \begin{Bmatrix} r_p \\ i \\ \omega \end{Bmatrix}_I . \quad (6.2)$$

Π_{OE} is as in Equation 6.2, and takes the final condition in the rotating-pulsating frame of the ER3BP, transforms it back into the non-rotating frame, and converts this position and velocity into its equivalent instantaneous orbital elements. The radius of periapsis, inclination and argument of periapsis of the final condition are then isolated from the full set for computing $C_{t_0}^T$. The conversion to-and-from orbital elements and Cartesian position is performed using standard conversion algorithms presented in Curtis [17] modified to support DA arithmetic.

I note that this parameterisation is not valid at $i = 0$ or $i = 180^\circ$ due to singularities in the orbital elements. At these values, the motion is planar and ω is undefined, often arbitrarily set to zero. This means that $C_{t_0}^T$ does not have three distinct, real eigenvalues and thus the helicity is no longer a valid measure of orthogonality. The reference hyperplanes must thus be selected to avoid these values of inclination.

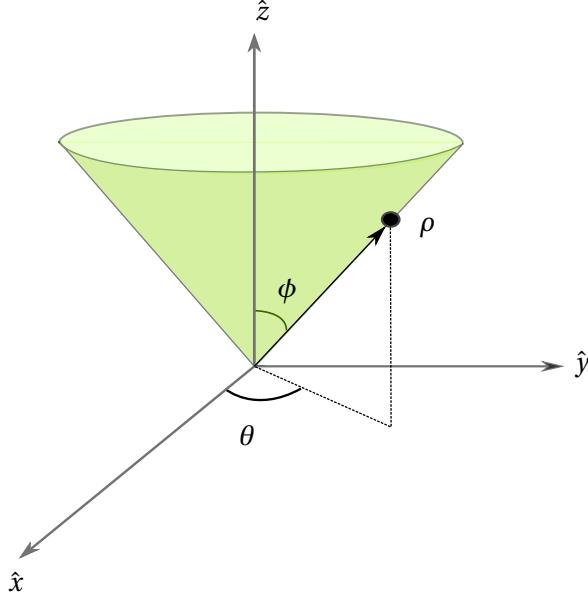


Figure 6.2: Graphical depiction of the spherical parameterisation of an initial condition. The light green ‘cone’ represents an example hyperplane of fixed ϕ .

6.2.2 Spherical coordinates

In this transformation, the sub-manifold is parameterised using spherical coordinates in the non-rotating frame $\psi = (\rho, \theta, \phi)$ where hyperplanes are given by fixed ϕ . I include a graphical depiction of this parameterisation in Figure 6.2 to aid clarity. This is a natural extension of the use of polar coordinates to compute ballistic capture sets in two dimensions in the literature [236]. The phase space is completed by uniquely associating a velocity with each point in physical space.

The velocity $v = (\dot{x}, \dot{y}, \dot{z})$ used in this transformation should match the velocity attached at each equivalent Cartesian position in the orbital element transformation, to provide a valid comparison. To ensure a fair test, this velocity should be defined independently of an knowledge of any of the orbital elements. This velocity will be selected differently in each application, and the specific formulation for each is presented later.

Once the velocity is known and the full phase space is complete, it is converted into the ER3BP rotating-pulsating frame and propagated under the equations of motion

$$\Psi_S: \begin{Bmatrix} \rho \\ \theta \\ \phi \end{Bmatrix}_I \rightarrow \begin{Bmatrix} x \\ y \\ z \end{Bmatrix}_I \rightarrow \begin{Bmatrix} x \\ y \\ z \\ \dot{x} \\ \dot{y} \\ \dot{z} \end{Bmatrix}_I \rightarrow \begin{Bmatrix} x \\ y \\ z \\ x' \\ y' \\ z' \end{Bmatrix}_R. \quad (6.3)$$

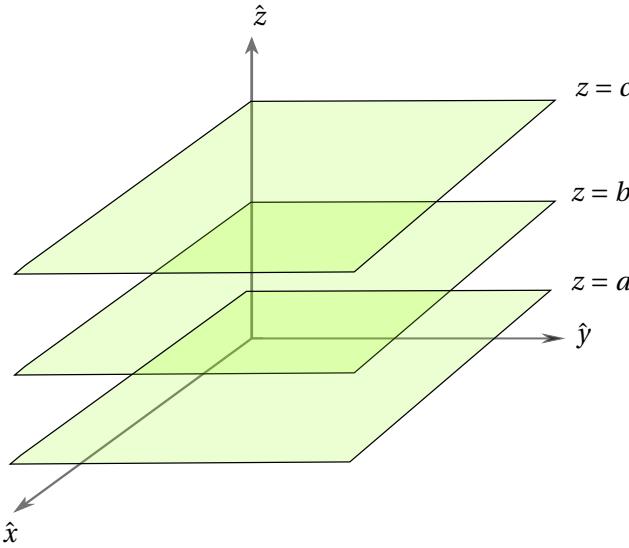


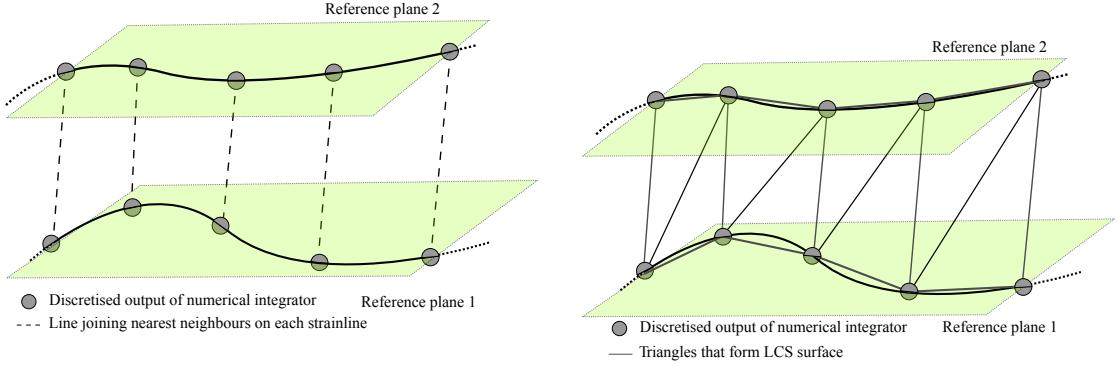
Figure 6.3: Graphical depiction of the Cartesian parameterisation of an initial condition. The light green planes represent example hyperplanes of fixed z .

The spatial dimensions are then converted back into spherical coordinates in the non-rotating frame to compute $C_{t_0}^T$

$$\Pi_S : \begin{Bmatrix} x \\ y \\ z \\ x' \\ y' \\ z' \end{Bmatrix}_R \rightarrow \begin{Bmatrix} x \\ y \\ z \\ \dot{x} \\ \dot{y} \\ \dot{z} \end{Bmatrix}_I \rightarrow \begin{Bmatrix} x \\ y \\ z \end{Bmatrix}_I \rightarrow \begin{Bmatrix} \rho \\ \theta \\ \phi \end{Bmatrix}_I . \quad (6.4)$$

6.2.3 Cartesian coordinates

The final transformation examined in this chapter defines an initial condition in the non-rotating frame using Cartesian coordinates $\mathbf{x} = (x, y, z)$ with hyperplanes given by fixed values of z (Figure 6.3). It is investigated here since it is conceptually simpler to visualise than the previous two approaches, but has not yet been investigated in the literature.



(a) The first step in reconstructing the full LCS structure is finding the nearest neighbours to points on adjacent reference planes by computing the Euclidean distance between them in parameter space.

(b) Triangles are formed between nearest neighbours in each reference plane, which can be written directly to a STL file and displayed and manipulated as a surface.

Figure 6.4: Diagrams to show how the full LCS surface is reconstructed by interpolating between strainlines on neighbouring reference planes.

The velocity to complete the full phase space is specified in the same manner as for spherical coordinates and elaborated later. Overall, this parameterisation is the $\Psi - \Pi$ pair

$$\Psi_C : \begin{Bmatrix} x \\ y \\ z \end{Bmatrix}_I \mapsto \begin{Bmatrix} x \\ y \\ z \\ \dot{x} \\ \dot{y} \\ \dot{z} \end{Bmatrix}_I \mapsto \begin{Bmatrix} x \\ y \\ z \\ x' \\ y' \\ z' \end{Bmatrix}_R \quad (6.5)$$

$$\Pi_C : \begin{Bmatrix} x \\ y \\ z \\ x' \\ y' \\ z' \end{Bmatrix}_R \mapsto \begin{Bmatrix} x \\ y \\ z \\ \dot{x} \\ \dot{y} \\ \dot{z} \end{Bmatrix}_I \mapsto \begin{Bmatrix} x \\ y \\ z \end{Bmatrix}_I \quad (6.6)$$

6.3 Producing the three-dimensional surface

The process for constructing the full structure of the LCS from its intersections with the reference planes is now elucidated in preparation for the coming subsections. Once the strainlines for each of these planes are computed, the full structure of the LCS is produced as a STL file, which is a common file format in 3D modelling applications. The STL file format contains a sequence of sets of three vertices that form triangles, and the collection of faces of the triangles are then

displayed as a full structure by freely-available viewers, such as Blender¹ or Meshlab².

Since this file format is relatively straightforward to create, the STL file and thus the full structure is created using a triangulation algorithm. The full structure is first separated into a collection of strainlines on each reference plane in parameter space and, by leveraging that I expect there to be a nearby strainline in a nearby reference plane, the nearest neighbour to each point on a strainline in the neighbouring reference plane is found (Figure 6.4a). This is performed by considering the Euclidean distance between points in adjacent reference planes in the parameter space of the submanifold.

Now the two points that should be interpolated between are known, a triangle is formed using the original point, its nearest neighbour in the next reference plane, and a point next to the original point in the same strainline (Figure 6.4b). This procedure is then repeated for every point in every strainline in the reference plane, and then repeated across all reference planes. By doing so, a triangulation of the full structure of the LCS can be constructed for use in visualisation software.

To avoid false reconstructions and to allow the structure to ‘end’ where required, a distance cutoff is used: nearest neighbours with Euclidean distances above the cutoff are not interpolated between. The choice of which hyperplane to begin the interpolation from can also be adjusted to provide the best results. These free parameters were adjusted by trial-and-error to produce the structures presented later in this chapter.

In some cases, it may still be necessary to remove reconstruction artefacts that occur as a result of the interpolation. This refinement is performed in Blender prior to rendering.

6.4 A proof-of-concept

The first application in this chapter is via a simplified selection of the initial velocity. I first parameterise a region of interest in three dimensions in an inertial frame about Mars where the Sun-Mars distance is unity using spherical coordinates. For this test case, given the Cartesian position $\mathbf{x} = (x, y, z)^\top$ corresponding to ψ

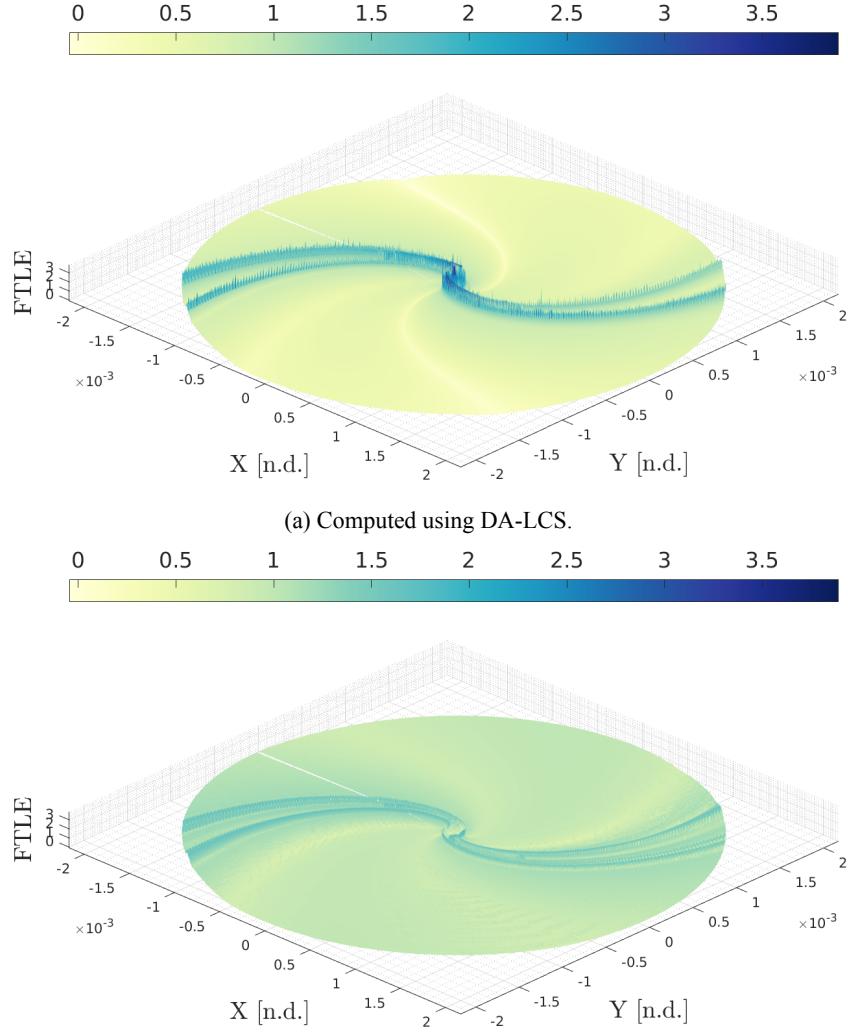
$$x = \rho \cos \theta \sin \phi \quad (6.7)$$

$$y = \rho \sin \theta \sin \phi \quad (6.8)$$

$$z = \rho \cos \phi \quad (6.9)$$

¹<https://blender.org>

²<https://meshlab.net>



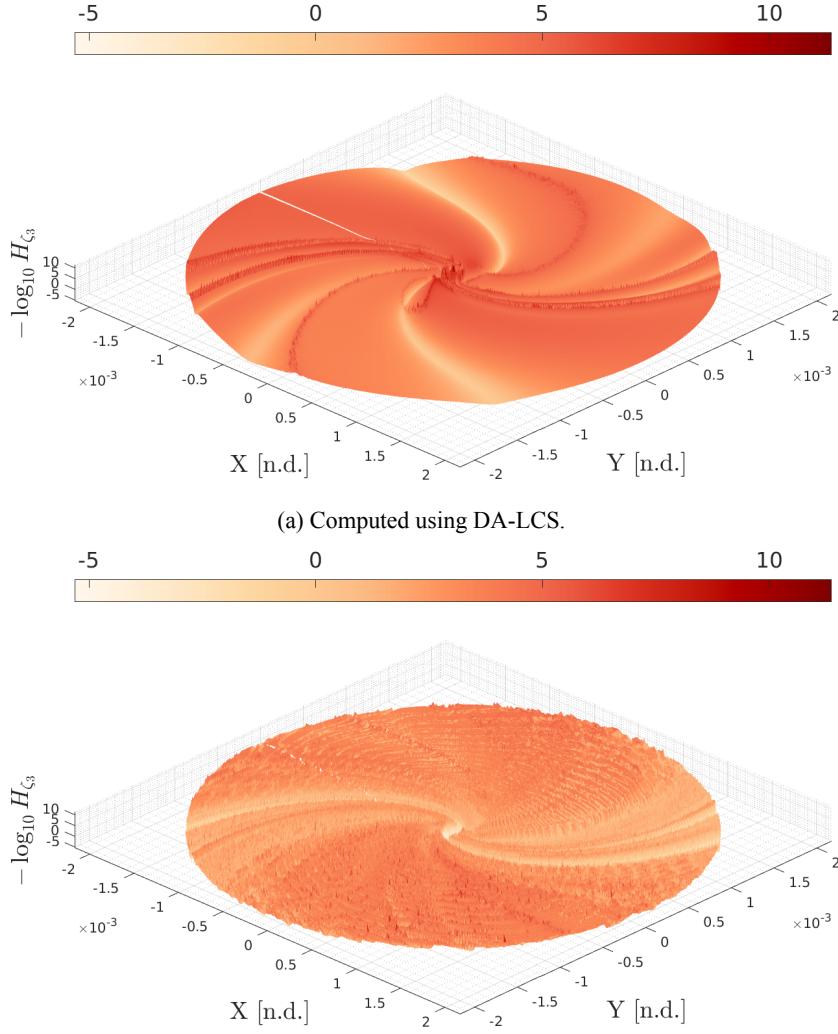
(b) FTLE field obtained using divided differences with auxiliary grid spacing of 0.05 of the nominal grid spacing in r and θ and the nominal grid spacing in ϕ .

Figure 6.5: Finite-time Lyapunov exponent field for the Elliptic-Restricted Three-body Problem on the $\phi = 115^\circ$ plane from $t_0 = v_0 = 0$ to $T = v = 2\pi$. While the structure is qualitatively the same, the ridges in the FTLE field are more well-defined with DA-LCS.

the velocity at this point $\mathbf{v}(\mathbf{x})$ is chosen to be

$$\mathbf{v}(\mathbf{x}) = \sqrt{GM \frac{(1+e)}{\rho^3}} \left[\begin{pmatrix} x \\ y \\ z \end{pmatrix} \times \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix} \right], \quad (6.10)$$

where the problem parameters $GM = 1.50499 \times 10^{-14}$ and $e = 0.9$ are the gravitational parameter of Mars and the eccentricity of orbits about Mars, respectively, defined in the non-rotating frame with Sun-Mars unity length. Conceptually, this fixes the velocity direction tangential to a cylinder around the z -axis, while the magnitude corresponds to the periapsis of a Keplerian orbit of eccentricity e . This choice of velocity vector reveals a ‘dynamically interesting’ region of behaviour about Mars where small changes in initial condition can lead to orbits escaping, or being



(b) Helicity field obtained using divided differences with auxiliary grid spacing of 0.05 of the nominal grid spacing in r and θ and the nominal grid spacing in ϕ . The same grid is used for computing both $C_0^{2\pi}$ and $\nabla \times \zeta_n$.

Figure 6.6: $-\log H_{\zeta_n}$ for the Elliptic-Restricted Three-body Problem on the $\phi = 115^\circ$ plane from $t_0 = \nu_0 = 0$ to $T = \nu = 2\pi$. No defined ridges of low helicity are visible with divided differences, but DA-LCS can extract well-defined ridges on which the strainline integration can be performed without stopping prematurely.

temporarily or permanently captured. At the same time, it generalises the choice of velocity in Ros Roca [7] and Parkash [8] by being the velocity of an object on the periapsis of an ellipse of eccentricity e , straightforwardly extended to three dimensions. It is also a simple-to-visualise velocity to aid understanding of LCS behaviour.

The above velocity attachment would not produce a valid LCS with the spherical parameterisation at $\phi = 90^\circ$, or the Cartesian parameterisation at $z = 0$ as the motion is purely planar and there would not be three distinct, real eigenvalues in $C_{t_0}^T$. The hyperplanes are thus selected to avoid these regions¹.

¹I note informally that interpolating between hyperplanes placed closely either side of these ‘troublesome’ hyperplanes yields a very good approximation to the LCS on the $\phi = 90^\circ$ and $z = 0$ hyperplanes.

The set of reference hyperplanes is defined as

$$\mathcal{S} = \{\psi \in [r, r_s] \times [0, 2\pi] \times [5^\circ, 15^\circ, \dots, 175^\circ]\}.$$

with hyperplanes given by fixed ϕ with a 600×600 equi-spaced grid in the remaining two variables per plane. The variables $r = 1.641 \times 10^{-5}$ and $r_s = 0.00513$ are the radius and the Hill sphere of Mars, respectively; the latter is the maximum distance from Mars at which it still dominates gravitational attraction. The initial integration time is set equal to $t_0 = v_0 = 0$ and the final time is $T = v = 2\pi$.

To reduce computational expense in computing the helicity field, I terminate integration of a given trajectory if the distance to Mars passes below 10% of the Mars radius, as the numerical integration of trajectories becomes exceedingly slow at these points. This is due to a combination of the integrator taking incredibly small time-steps as a result of needing to conserve errors in all orders of expansion, and the coefficients of the expansion becoming very large at distances close to Mars which is then carried through the rest of the integration, keeping these step-sizes small². This approach is also taken in the ballistic capture literature and the MSc theses, as trajectories which crash into the host planet are also generally unsuitable for space missions and can lead to numerical difficulties [7, 8, 205, 206, 236]. For correctness, while these points are typically of very low helicity, I do not consider these points valid seed points as they are not associated with the same integration time as the rest of the field. Strainlines are also terminated if they reach one of these points.

I use a helicity tolerance of $\alpha = 4 \times 10^{-6}$ on the $\phi = 115^\circ$ hyperplane. To reduce the computational expense, and motivated by the fact that DA-LCS produced longer strainlines for the steady ABC flow than divided differences, I use a reduced grid to sample seed points. This grid was selected with consideration to computing limits on IRIDIS and test runs of the strainline integration procedure and is varied per plane. The variable step size adjustment in the numerical integrator is also adjusted to increase the step size more slowly than the default settings for this test case by simply dividing the change in step size by 10 whenever it is increased, and multiplying the change in step size by 10 whenever it is decreased. Ideally, a very small fixed step size consistent with advancing the position by a fraction of the width of the low-helicity regions in ER3BP units would be used, but this is computationally infeasible if one wishes to sample every strainline. An idea for a more principled strategy to improve strainline integration is suggested in the conclusions of this thesis.

6.4.1 Results

The numerical performance of DA-LCS in an example astrodynamics system is now shown. The FTLE fields computed using DA-LCS and divided differences on the $\phi = 115^\circ$ plane is presented in Figures 6.5a and 6.5b, respectively. The structure found using DA-LCS agrees with

²In some cases, a single trajectory starting close to Mars has taken up to twelve hours to compute when allowed to go to 0.1% of the Mars radius.

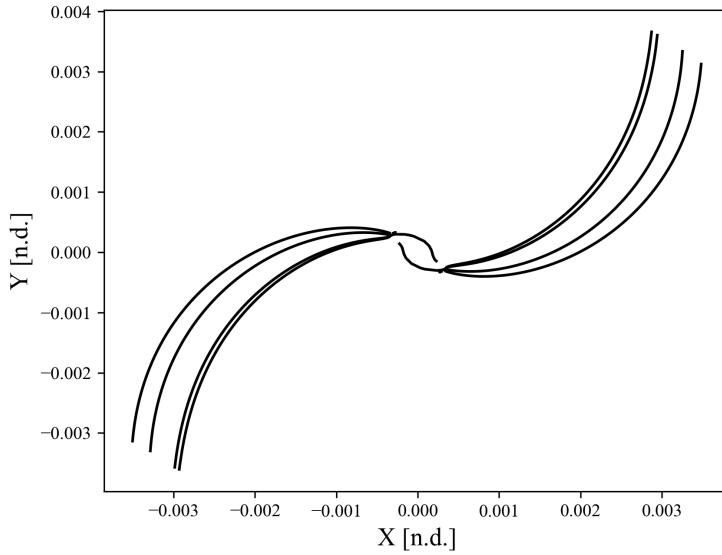


Figure 6.7: Strainlines on the $\phi = 115^\circ$ plane for the Elliptic-Restricted Three-Body Problem computed using DA-LCS. Divided differences is unable to generate any strainlines, but with DA-LCS the structure of the LCS can be deduced readily and with only 8 strainlines.

what would be expected from previous literature, with the structures in the two ‘arms’ being consistent with the transition between orbits that escape and are permanently or temporarily captured about m_2 [206]. Similar performance, albeit with poorer definition of the FTLE ridges, can be obtained using divided differences after tuning the grid sizes used to generate the derivatives. While the ER3BP does admit variational equations that can be integrated with the equations of motion which may improve the quality of the derivatives used to compute $C_0^{2\pi}$, these variational equations cannot be used to compute $\nabla \times \zeta_3$, which must still be approximated using divided differences and appear to produce the majority of the error for this test case. This is to be expected, as the estimation of second order derivatives using divided differences is more numerically difficult than first order.

Figure 6.6a presents the helicity field on the $\phi = 115^\circ$ plane for the ER3BP computed using DA-LCS. Like the FTLE field, the helicity highlights the ‘arms’ as being influential portions of flow. However, using divided differences to compute the helicity, given in Figure 6.6b, produces no meaningful insight into the helicity field even after tuning the grid sizes used; the numerical noise in the determination of the helicity reveals no distinct ridges along which the numerical integration can begin. This numerical improvement comes completely automatically, without the need to tune grid sizes.

The focus of this investigation is to determine the improvement in computing the helicity for astrodynamics systems. However, for completeness the final set of strainlines for this example computed using DA-LCS on the $\phi = 115^\circ$ plane are shown in Figure 6.7, and largely follow from the helicity field given earlier. *A posteriori* analysis of the strainlines show that they do indeed perform the role of LCS in separating regions of qualitatively different behaviour, validating the

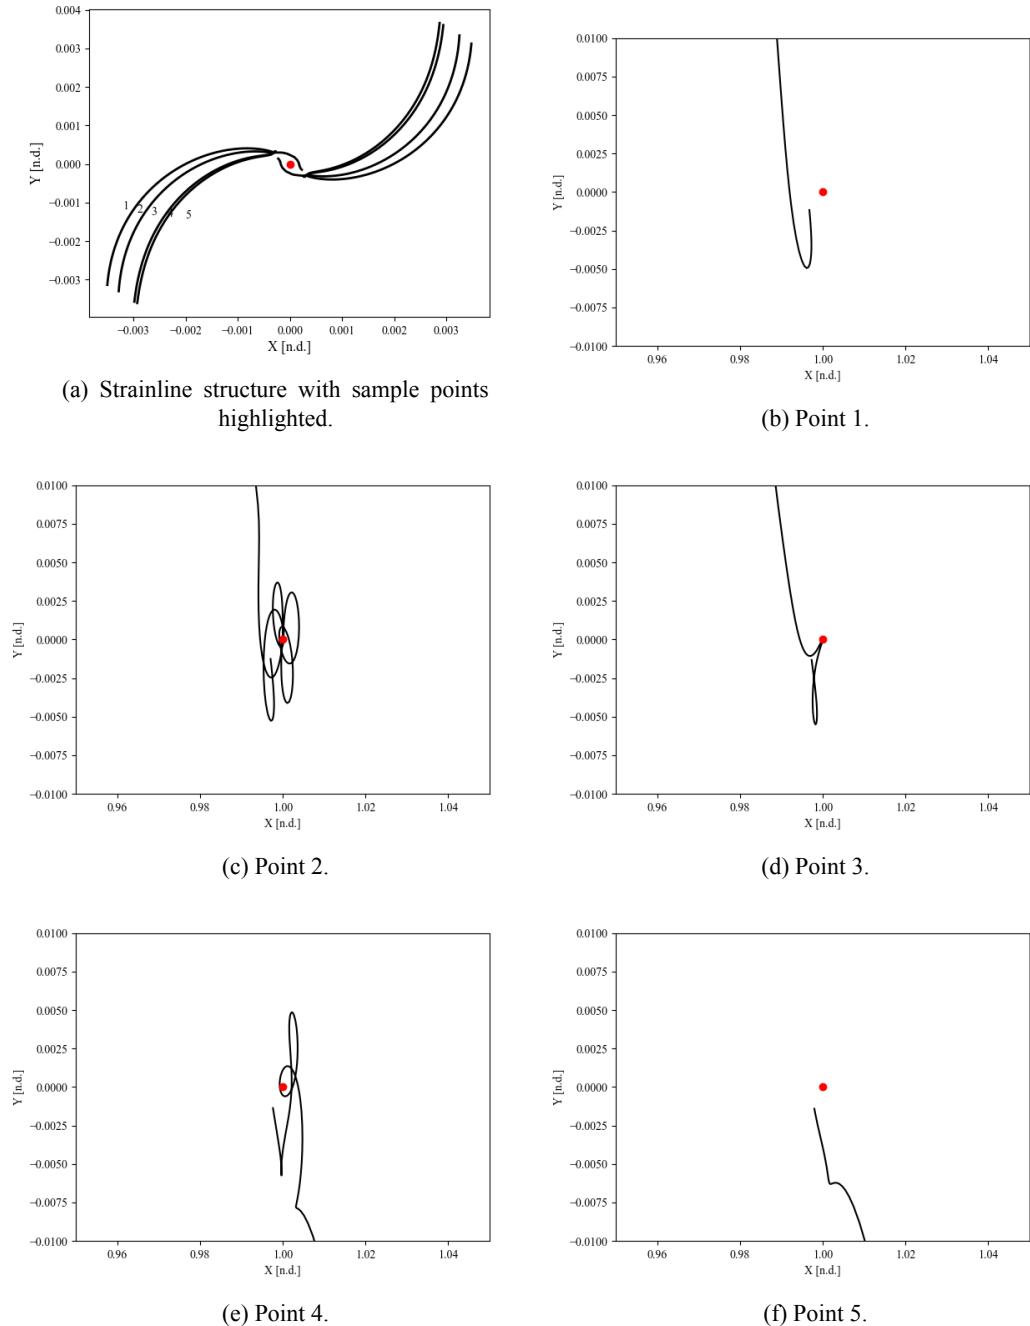


Figure 6.8: Strainlines with the origins of sample trajectories in the rotating frame overlaid. Mars is highlighted in red in the centre of the Figure. The LCS is acting as expected in separating regions of qualitatively different behaviour, validating the strainlines found.

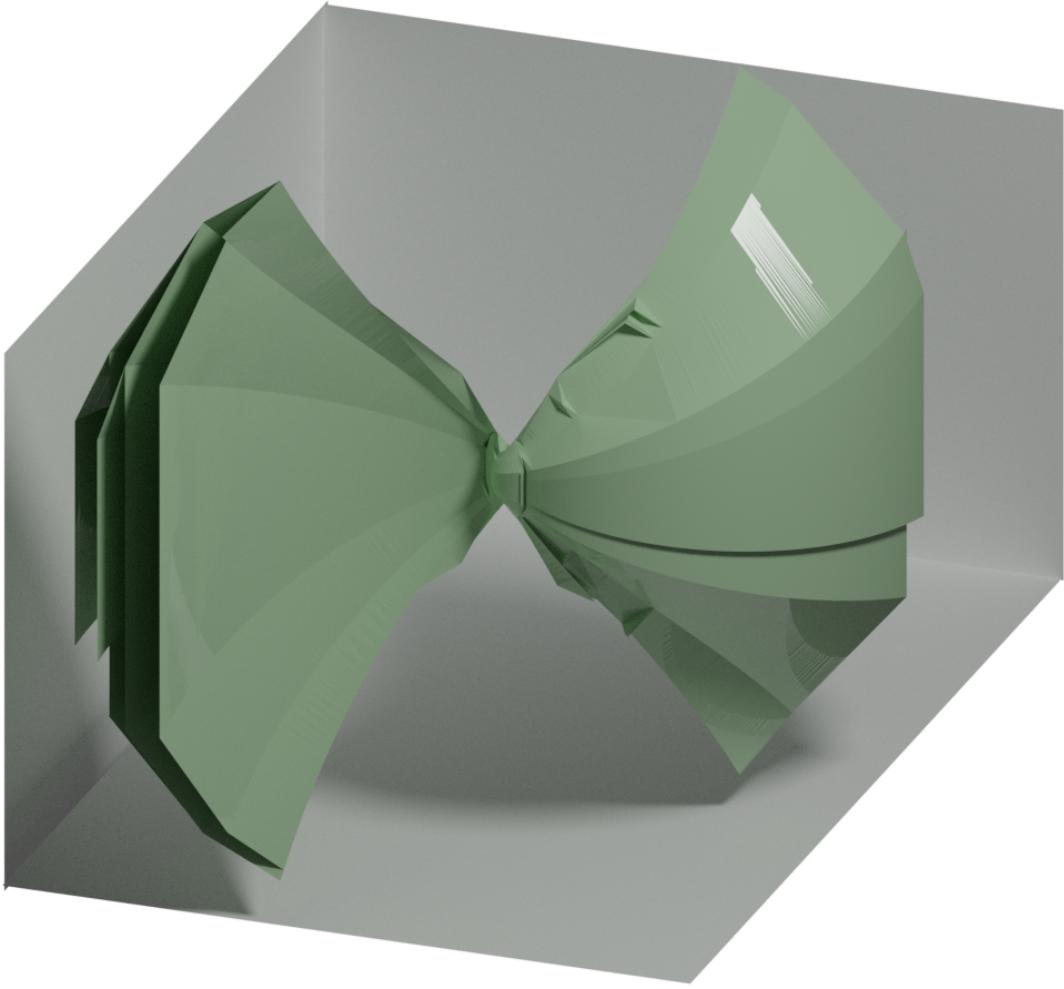


Figure 6.9: Rendering of the full 3D LCS for the ER3BP test case. The main structure is the two large curving ‘arms’, which previous literature has struggled to reproduce [7, 8].

choice of α (Figure 6.8) and their computation. A representative rendering of the full 3D LCS for this test case is shown in Figure 6.9.

6.5 Pushing the envelope

Now that it has been seen that DA-LCS can produce insight in example astrodynamics systems by improving the computation of H_{ζ_n} , I turn to a more principled example based on a more physical choice of velocity to examine its practical uses. I select the initial conditions to correspond to the orbital elements transformation Π_{OE} , Ψ_{OE} (Figure 6.1) for a more common definition of initial conditions, and increase the numerical complexity of the problem by increasing the rate at which orbits escape from Mars by selecting the eccentricity $e = 0.95$, $\Omega = 118^\circ$, and initial true anomaly of $\nu = 0$. This choice of Ω has been shown to maximise ‘interesting’ dynamical behaviour across all inclinations [205, 206], and the high eccentricity leads to orbits which escape into heliocentric

orbits rapidly. This presents a numerically very challenging test case to use to benchmark the LCS method against as the high rates of repulsion lead to large derivatives everywhere, not just at the separatrices. Being able to function even in such a challenging test case would highlight the numerical performance and robustness of DA-LCS.

I also wish to examine the effect of transformation on the LCS. To allow a valid comparison across transformations, I again require a way to associate a velocity with each point in the spherical and Cartesian transformations without information on the orbital elements. One can define this velocity in the non-rotating frame with respect to the initial Cartesian position $\mathbf{x} = (x, y, z)^\top$ as

$$\mathbf{v}(\mathbf{x}) = \sqrt{GM \frac{1+e}{\rho^3}} \begin{bmatrix} x \\ y \\ z \end{bmatrix} \times \hat{\mathbf{n}}(\mathbf{x}) \quad (6.11)$$

where GM is the standard gravitational parameter of Mars in the non-rotating frame, and e has the same meaning and values as before. The quantity $\hat{\mathbf{n}}(\mathbf{x})$ is the normal to the orbital plane at that point and can be computed by considering the vector that points along the line of nodes $\hat{\Omega} = (\cos \Omega, \sin \Omega, 0)^\top$

$$\hat{\mathbf{n}}(\mathbf{x}) = \text{sign}(z) \frac{\hat{\Omega} \times \mathbf{x}}{\|\hat{\Omega} \times \mathbf{x}\|}. \quad (6.12)$$

The value of Ω used is the same as in the orbital elements transformation, and the $\text{sign}(z)$ term is required to ensure the correct orientation of the unit normal. Equation 6.11 is well defined everywhere except on the line of nodes, where $\hat{\Omega} \times \mathbf{x} = \mathbf{0}$. Since there is no one unique inclination passing through these points, additional information from the orbital elements transformation would be required to attach an equivalent velocity at these points. The reference hyperplanes are thus selected or sampled in such a way to avoid these ill-defined points.

6.5.1 The LCS Structure

For this example, I analyse the Sun-Mars ER3BP from an initial integration time of $t_0 = v_0 = 0$ to a final integration time of $T = v = 2\pi$ (1 full Martian year). Ψ and Π are the orbital elements transformation Ψ_{OE} and Π_{OE} , selected to highlight the ability for DA-LCS to use highly complex mappings that would be impractical to derive manually and still obtain insight automatically. The set of reference planes \mathcal{S} are defined such that

$$\mathcal{S}_{OE} = \{r_p \in [r, r_s], \omega \in [0, 2\pi], i \in [-85^\circ, -75^\circ, \dots, 85^\circ]\}.$$

The variables $r = 1.641 \times 10^{-5}$ and $r_s = 0.00513$ have the same meaning as before.

A 100×2500 grid in r_p and ω is used on each plane in \mathcal{S} on which to compute the helicity fields. As will be seen later, the width of the low-helicity region is approximately 40 km: accurately sampling transverse to this region motivated the use of this dense grid in ω . On the $i = 35^\circ$ hyperplane, a helicity threshold of 10^{-8} is used for the numerical integration, but to limit the

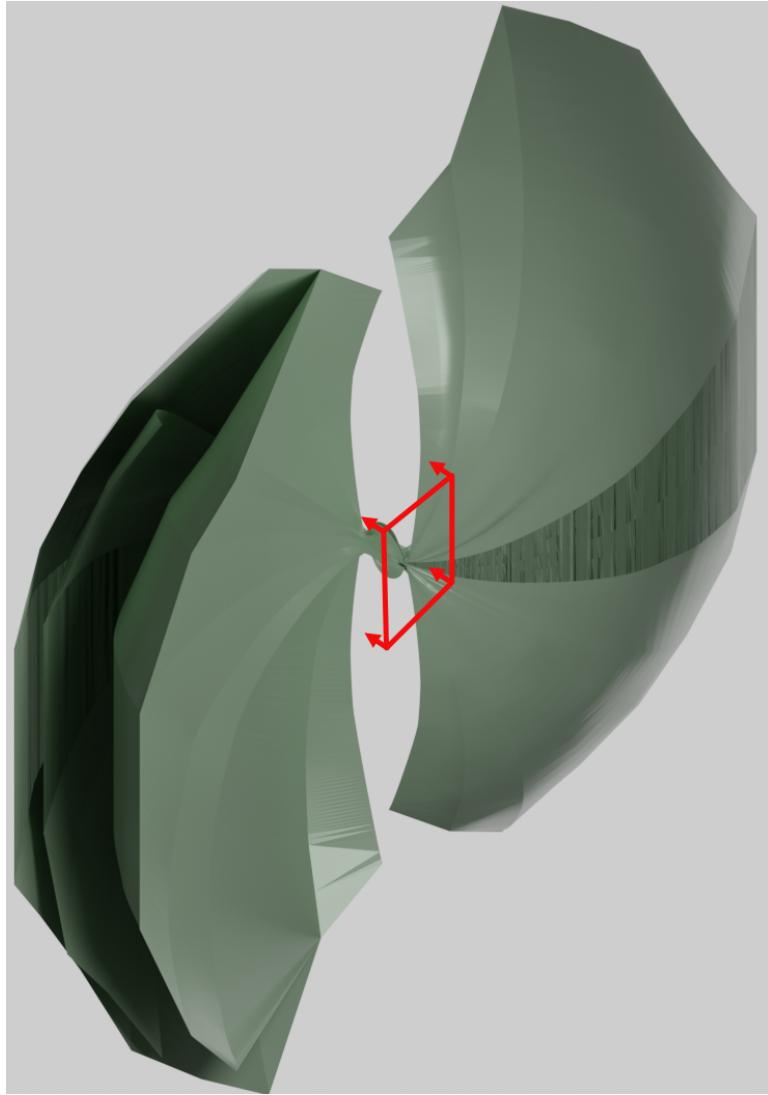


Figure 6.10: The full structure of the LCS obtained, which is comprised of an inner section and two large ‘arms’ which emanate from the inner section. The region in red is the point-of-view of the section plot shown in more detail in Figure 6.12.

computational expense the integration begins from points with helicity less than or equal to 10^{-9} . These thresholds are varied per plane based on visual examination of the helicity field.

A representative rendering of the full repelling LCS structure is given in Figure 6.10. To show more clearly how the LCS is constructed from the strainlines on the reference hyperplanes, an example set of strainlines is shown on a section view of the full LCS in Figure 6.11. The LCS is formed of two distinct sections: the ‘inner’ portion that encloses Mars (Figure 6.12) and the two extending arms that emanate from the inner structure (Figure 6.13); these two regions are now investigated in more detail.

In Figure 6.12 the inner structure is presented with the extending arms removed and example trajectories corresponding to some initial conditions inside and outside of the LCS are shown. In red a generic trajectory that begins inside the inner section is shown, which continually orbits

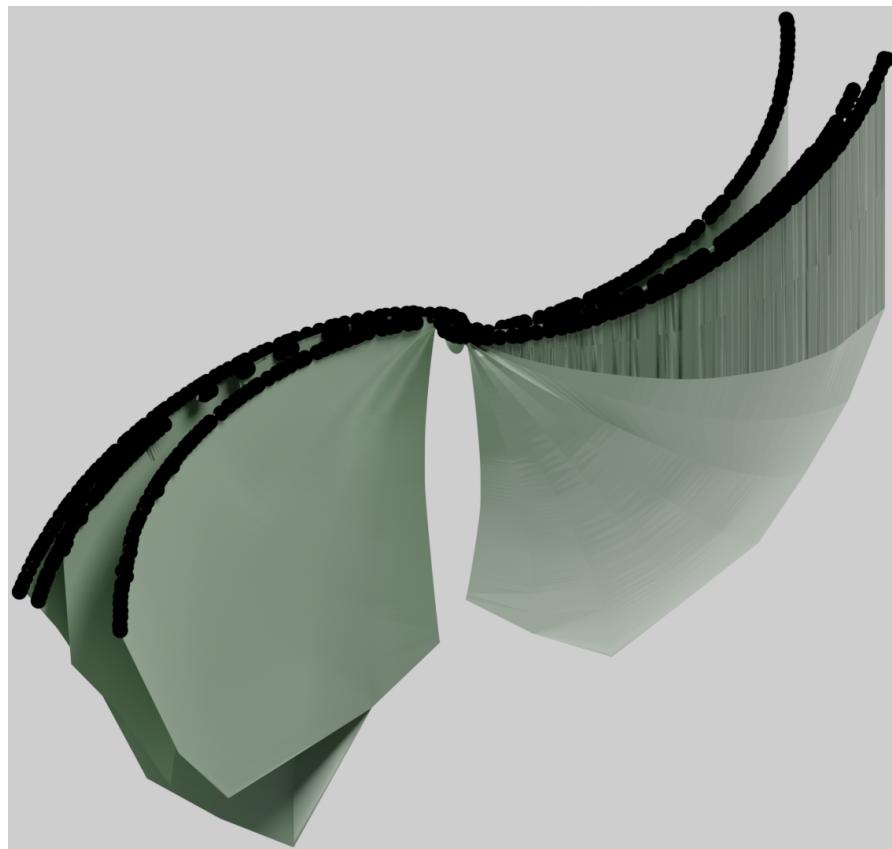


Figure 6.11: A section view of one of the hyperplanes of fixed inclination with points that form the strainlines shown in black to highlight more clearly how the surface is reconstructed.

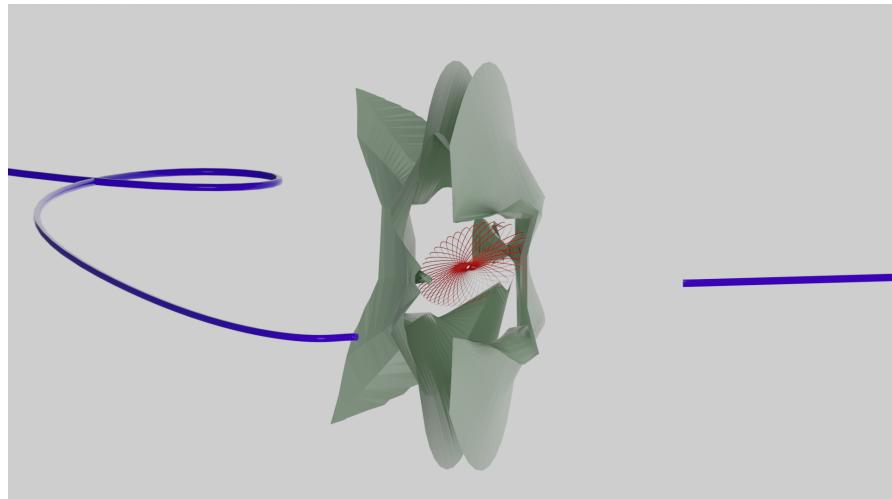


Figure 6.12: Inner section of the LCS (highlighted in red in Figure 6.10) with the ‘arms’ removed and three sample trajectories overlaid to show how it is acting as a separatrix. The trajectory taken from inside the LCS (red) remains orbiting Mars while the two trajectories that begin nearby but outside the inner section (blue) both immediately escape.

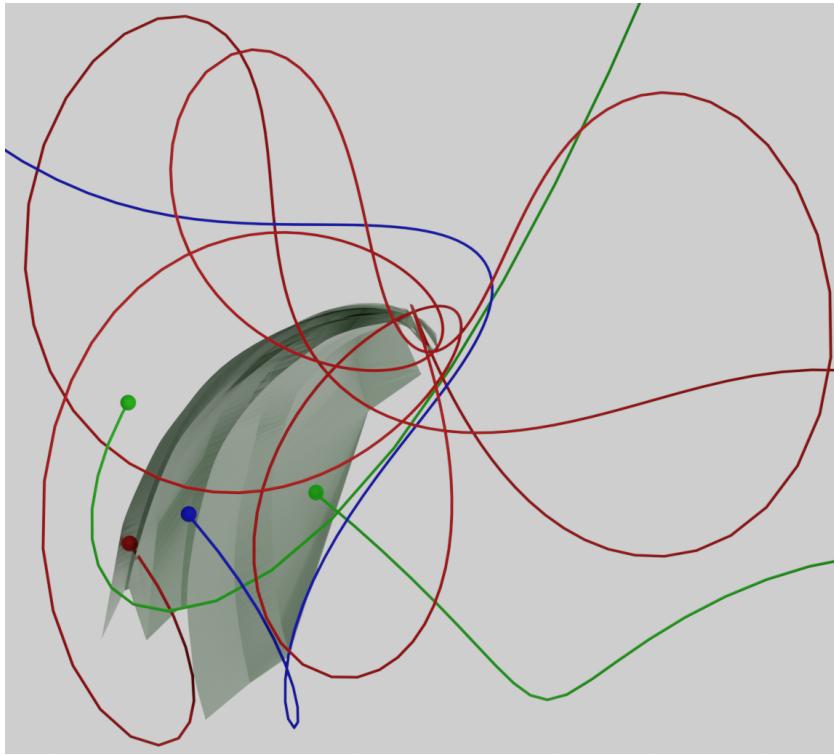


Figure 6.13: Trajectories in the synodic frame of points taken from different parts of the LCS arms. Each trajectory is qualitatively different and the LCS is again acting as a separatrix.

Mars over the integration period considered. Mars lies at the centre of the inner section. In blue are two trajectories that begin nearby but outside of the inner section, which both immediately escape from Mars' sphere of influence. The LCS here is acting as separator between trajectories which have qualitatively very different behaviour. This matches the expected structure from investigations into ballistic capture where the desired behaviour is defined beforehand [206], and again validates the helicity threshold used.

A more detailed presentation of one of the two ‘arms’ is given in Figure 6.13, again with example trajectories highlighted and the respective initial conditions indicated by spheres. In green are two trajectories which begin outside of the LCS structure and are qualitatively similar in immediately escaping the sphere of influence of Mars, as was the case in Figure 6.12. More significantly, the remaining blue and red trajectories that are taken from different parts of the arms are also qualitatively different to both themselves and the two green trajectories outside of the LCS, highlighting how effectively the LCS is partitioning dynamical behaviour. In general, these ‘arms’ separate orbits which complete a different number of revolutions from each other, including those which immediately escape from Mars. This again matches expectations from the literature [206], which identifies these as ‘stable sets’.

There are other advantages to determining separators of qualitatively different behaviour than simply finding orbits that differ in forward time. For example, a mission designer could make a mission more robust by targeting the ‘middle’ of one of the regions of the LCS. Since all the trajectories in the region are now known to have qualitatively the same behaviour in forward time,

any nearby points may be likely to continue to satisfy the mission plan in the event of deviations from a nominal trajectory. Moreover, the structure of the repelling LCS in backward time could be used in a similar manner to construct a full, and potentially robust, arrival procedure to the host planet.

I note differences from the LCS structure presented in the previous section and the LCS structure presented here. These are predominantly due to the different selection of velocity attachment, which in the previous section corresponds to selecting $\omega = 90^\circ$ and allowing Ω to vary such that $\Omega \in [0, 2\pi]$. The velocities associated with each Cartesian position are only the same between the two LCS structures at zero inclination, and $\omega = 90^\circ$ and $\omega = 270^\circ$; everywhere else, there is a non-zero velocity in z . At higher inclinations, the effect of non-zero z -velocities becomes more clear and the difference between the two structures becomes starker. Since the choice of orbital elements for the initial conditions, as per Luo and Toppotto [206], maximises ballistic capture across all inclinations, far more structure is seen at progressively higher inclinations.

6.5.2 Effect of orbit parameterisation and integration time

In the previous subsection, I showed the result of DA-LCS to the Sun-Mars ER3BP with a single choice of transformation Ψ_{OE} and Π_{OE} and a single prescribed time period of investigation. However, there are many free choices available to the mission designer when using LCS, such as the initial and final times to consider, and the choices of Ψ and Π . This subsection presents the effect of the choice of orbit transformation and integration time on the helicity field found. The structure of the LCS can then be inferred by regions of consistently low helicity, which are coloured dark blue in these plots to improve visualisation. An advantage of DA-LCS is that it makes the implementation of other transformations straightforward via automatic differentiation.

I compare each $\Psi - \Pi$ pair introduced in section 6.2 for two integration times. The integration starts at $t_0 = \nu_0 = 0$ and ends at the true anomaly corresponding to $T = \nu = \pi/4$ (68 days) and $T = 2\pi$ (1 Martian year) on the same set of initial conditions, to enable a fair comparison. The grid of initial conditions is chosen to be those which lie on the $i = 35^\circ$ hyperplane in the orbital elements transformation, again sampled on a 100×2500 grid with the same parameters as used in the construction of the full LCS. For each point on the hyperplane in the orbital elements transformation, I also compute the equivalent Cartesian position. This Cartesian position forms the initial condition for the Cartesian transformation, and converting this Cartesian position to spherical coordinates gives the equivalent initial condition in the spherical transformation. In this way, one can ensure that each transformation samples the same points to ensure a valid comparison. However, by choosing to replicate the points on the $i = 35^\circ$ plane exactly the points will necessarily pass through the line of nodes, at which point the velocity attachment for spherical and Cartesian coordinates is not valid (section 6.2.2). To avoid ill-defined results, points within 2° of the line of nodes are not evaluated. By use of the nominal hyperplanes in ϕ or z as discussed in sections 6.2.2 and 6.2.3 this problem does not arise.

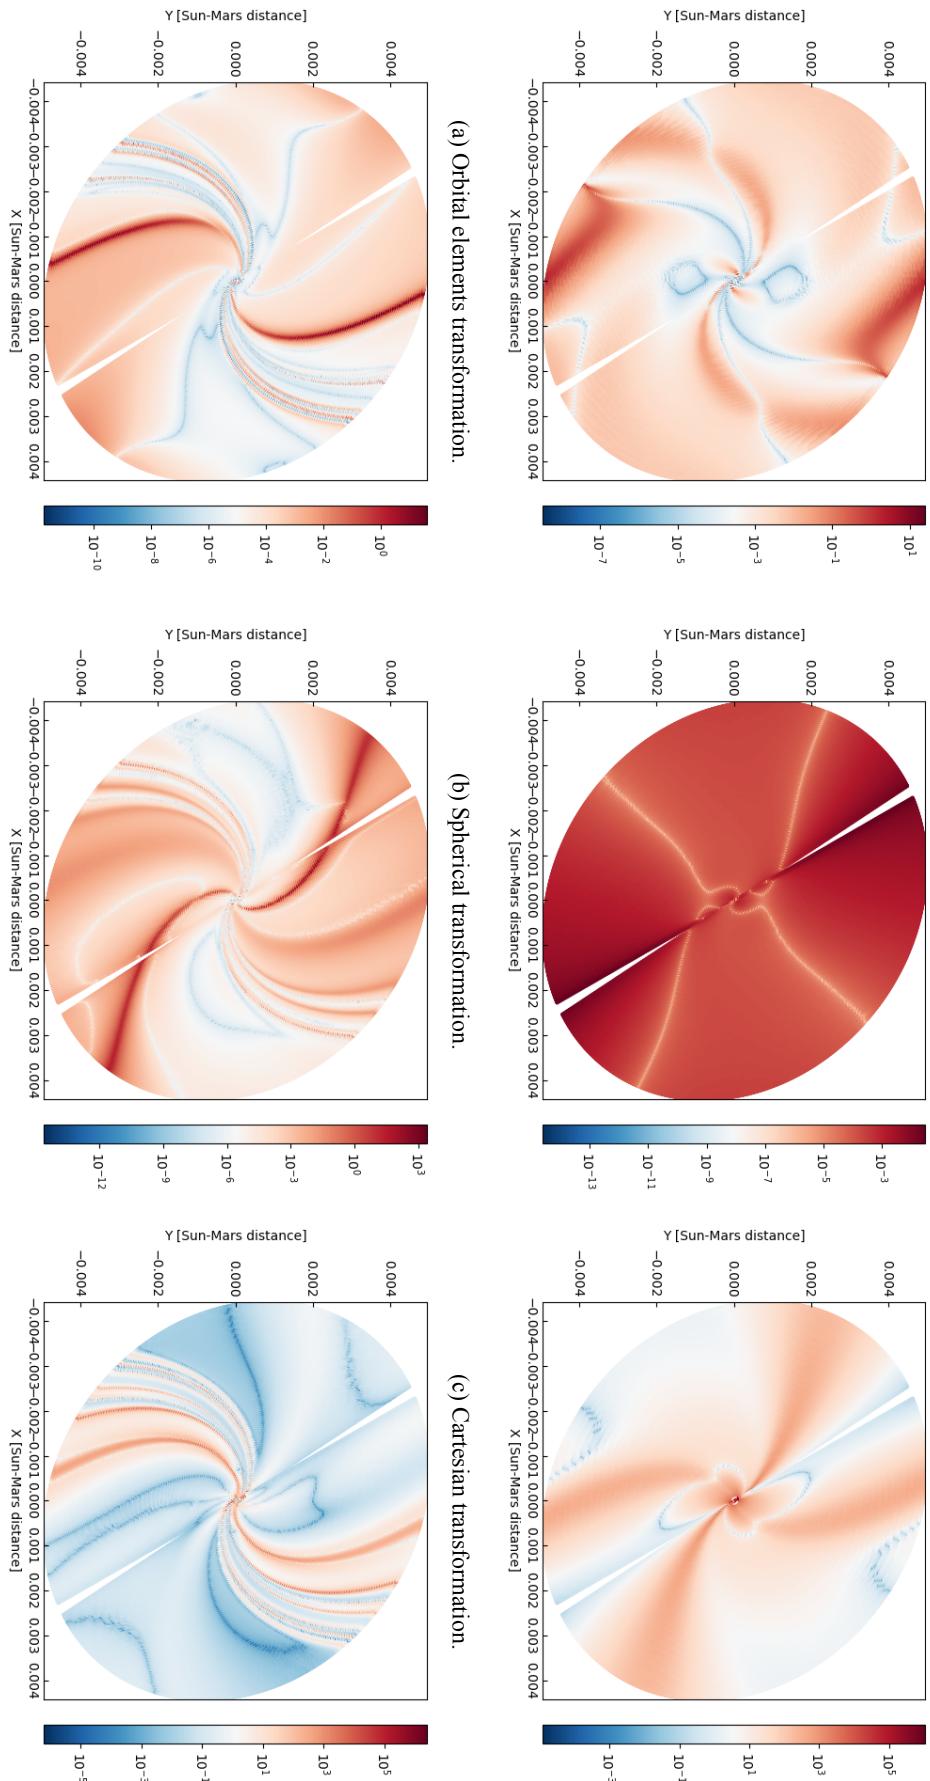


Figure 6.14: The effect of integration time on the helicity field for the ER3BP. The initial true anomaly is $t_0 = 0$ and the final true anomaly is either $T = \pi/4$ (68 days, top row), or $T = 2\pi$ (one Martian year, bottom row). In general, sufficient time is required to remove the effect of the transformations on the final derivatives. Regions of low helicity (dark blue) imply the location of an LCS.

Figure 6.14 presents the helicity fields for each of the transformations for both integration times. I find that the integration time T must be chosen to be sufficiently large to allow the time-dependent derivatives associated with the dynamical system in $\nabla \mathbf{F}_{t_0}^T$ to dominate the time-independent derivatives associated with the transformations Ψ and Π and obtain a system that closely resembles f . For $T = \pi/4$, sufficient time has not passed for the derivatives associated with the dynamics to dominate, and the result is inconsistent between the different transformations. For $T = 2\pi$, however, the results are consistently representing the behaviour of the dynamical system. For this example, $T = \pi$ is sufficient to ensure domination of the dynamics for all points in the domain for the examples shown here, though I note that many of the points farthest from Mars have orbital periods of thousands of days and much lower integration times are required for points closer to Mars.

The helicity fields signal qualitatively the same points as LCS, identifying the arms and the inner section as among the lowest helicity points for each transformation, despite differences in the actual value of helicity. The orbital elements and spherical transformations also quantitatively agree in both signalling the ‘arms’ and inner section as the strictly lowest regions of helicity in each case, as is to be expected. However, there is additional structure in the spherical transformation at this low helicity which does not agree with that predicted by the literature or the orbital elements case. With the Cartesian transformation this additional structure is signalled before the structure of the arms is completed. This structure is also discontinuous, despite the helicity being a continuous expression. By choosing a numerically challenging test case to benchmark DA-LCS on, a numerical error has been exposed that arises as a limit of floating-point arithmetic and manifests as ‘false positives’ in the helicity field. This error and mitigating strategies are discussed in section 6.5.3.

There is also other structure at higher helicities in the orbital elements case, which arises as a result of the transformation also including some information on the final velocity in the three variables that parameterise the submanifold. Since obtaining i and ω depends on the velocity, additional information on direction of escape and direction of final velocity is present compared to Π_S and Π_C . In those, all final velocity information is discarded (Equations 6.4 and 6.6). This additional structure extends that found in ballistic capture [205, 206], but these points do not feature in the full LCS: to be consistent with the original definition of the LCS, the helicity threshold is chosen to find only the most repelling structures. This yields the inner section and the ‘arms’ before any additional structure.

I now discuss another advantage of choosing a higher threshold to terminate the numerical integration than used to select seed points. Because the width of low-helicity regions is only 40 km, it means that strainlines exceed the helicity threshold and terminate quickly if the integrator obtains the next step with an error of more than approximately 1.2×10^{-7} in ER3BP units in position. This is already numerically sensitive, made worse by the fact that the integrator is adjusting its step size based on the error in evaluating an ODE involving only ζ_n , and is ‘unaware’ that it should instead be attempting to remain in the low-helicity regions defined by $\nabla \times \zeta_n$. With

reference to Figure 6.14d, the lowest-helicity regions are surrounded by points with helicities approximately six orders of magnitude higher. By increasing this threshold by less than six orders of magnitude (in the case of this plane, only one), the integrator is allowed to make small, recoverable errors to aid the integration procedure by slightly widening this ridge, but still terminate the integration immediately when it departs the low-helicity region that signals the LCS. The adjustment to the step size control mechanism used in the previous test case is also used here.

6.5.3 Implementation considerations

A potential numerical pitfall exists when numerically constructing LCS in astrodynamics problems and is elaborated in this subsection. I also provide guiding considerations on how to avoid these errors practically. I stress that since DA gives derivatives accurate to floating-point precision, the numerical difficulties given below are not as a result of DA-LCS, any particular software package or the mathematical definition of an LCS. Rather, they are a result of the finite precision of floating-point numbers and the large derivatives that arise in numerically challenging test cases.

Since computers approximate the field of real numbers using floating-point numbers, and these floating-point numbers have finite storage associated with them, there is a limit on the precision of numbers they can represent. Double-precision numbers have 64 bits of storage, which provides a maximum of 16 significant digits that can be represented accurately [237, 238]. As a result, in floating-point arithmetic, $1 + 1 \times 10^{-16} = 1$. Any error arising from this finite-precision representation compared to the evaluation in real numbers is known as round-off error [239].

Round-off error can occur in the evaluation of the helicity and construction of the LCS when $\nabla F_{t_0}^T$, $C_{t_0}^T$ or $\nabla \zeta_n$ become sufficiently ill-conditioned that the ratio of maximum to minimum absolute values of entries in those quantities approach 10^{16} in any relevant orders of the expansion. At this point, when the user requests a manipulation of these quantities, information on the smaller numbers is lost as a result of the finite precision. This error presents as false positives (false low values) in the helicity field.

I find that there are two mitigating strategies to prevent the ill-conditioning of these terms and prevent false positives in the helicity field in challenging test cases. The first is to ensure the proper conditioning of relevant quantities by adjusting the magnitude of the terms of the orbit parameterisation. In Figure 6.15, I present the helicity fields computed for each of the transformations with two different unit lengths: Mars radius and Sun-Mars distance at Mars periapsis to deliberately induce this round-off error. When using Mars radius as unit length for the orbital elements transformation (Figure 6.15a) and the spherical coordinates transformation (Figure 6.15b) which are both formed of a radius term and two angles, the helicity field is ‘inverted’, with low-helicity values in trajectories which all have qualitatively similar behaviour in immediately escaping, and higher values in the arms and inner section. A trajectory that immediately escapes from Mars will have a large final radius and associated derivatives, but with the angle terms still

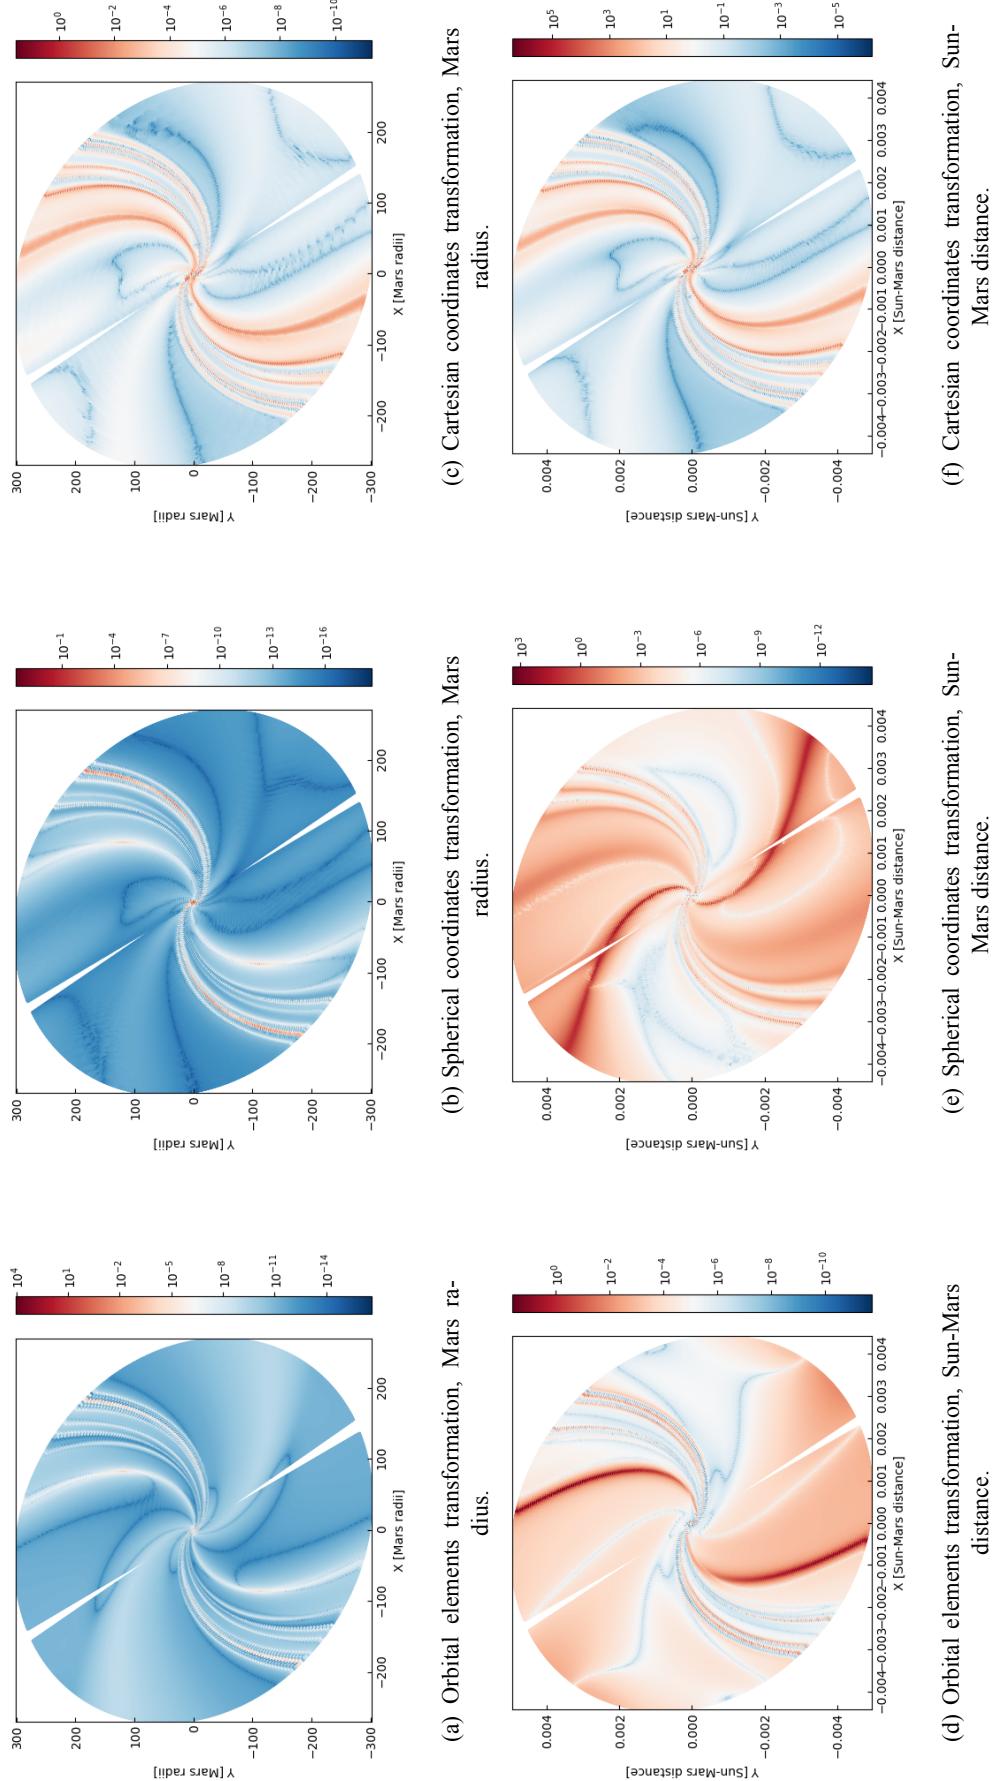


Figure 6.15: The helicity fields presented for each of the transformations using two unit lengths: Mars radius and Sun-Mars distance. One can aid the conditioning of relevant terms by scaling the unit length such that $C_{t_0}^T$ is well-conditioned in parameterisations with some bounded terms.

bounded in $[0, 2\pi)$. The terms in $C_{t_0}^T$ associated with the radius will thus dominate those associated with derivatives of the angles. When the helicity at this point is then evaluated, it becomes small not because of system dynamics, but because of a loss of precision of the arithmetic. Small numbers that would have otherwise contributed to the helicity have been lost to round-off error. Condition numbers \mathcal{C} for $\nabla \zeta_n$ of up to 10^{40} have been observed when using Mars radius as the unit length with the orbital elements transformation.

However, this only occurs in regions of large final radius, i.e. points which immediately escape. Where the final radius is lower, such as in the arms where orbits complete several revolutions before escaping, the relevant terms are sufficiently well-conditioned and thus helicity values are sufficiently free of round-off errors that the ‘correct’ order of magnitude of helicity is obtained at this point. Thus, lower helicities are present in regions of trajectories which all have qualitatively similar behaviour (immediately escaping) rather than in regions that separate the behaviour of trajectories, giving rise to the apparent inversion.

I find I can aid the conditioning of relevant terms in transformations where some of the variables that parameterise the submanifold remain bounded, as is the case for angles in the orbital elements and spherical coordinates transformations. For these, the distance units should be chosen such that the entries of $C_{t_0}^T$ are all of similar magnitude. This was achieved in this chapter by using the Sun-Mars distance as unit length. Changing only the length unit in the orbital elements transformation from Mars radius in Figure 6.15a to the Sun-Mars distance in Figure 6.15d reduces the condition number to approximately $\mathcal{C} = 10^{14}$ such that the helicity can be accurately represented and yields the expected result. I stress that these numerical effects are not specific to DA-LCS, any particular software package, or the mathematical definition of LCS. Rather, it is inherent to floating-point arithmetic and has been highlighted by the choice of a numerically challenging test case. A potential automatic mitigation strategy is to sample a greatly reduced grid of initial conditions and examine the mean or median condition number. If that condition number is above 10^{16} , the unit distance should be increased. This would allow a relatively rapid and ‘hands-off’ method for identifying a suitable unit length.

Scaling units to ensure the Jacobian is well-conditioned is more difficult in systems where all variables that parameterise the submanifold are unbounded, such as the Cartesian transformation. Here, where trajectories typically escape in the direction of x and y but stay close to the plane in z , terms in z can be dominated. Scaling the length unit uniformly will affect all terms equally and not change the condition number, which is above 10^{16} after even 68 days. This causes the poor numerical resolution in the Cartesian transformation observed earlier: scaling all terms of the parameterisation does not affect round-off error in regions of escape, causing the additional structure in trajectories that immediately escape to persist.

The second mitigating strategy is to use a transformation that is naturally more robust to ill-conditioning. Even when using Sun-Mars distance as the unit length in spherical coordinates (Figure 6.15e) there remains sufficient ill-conditioning in the trajectories which escape ($\mathcal{C} \sim$

10^{17}) that false positives are present at the same helicity as the arms and inner section. However, for the orbital elements transformation the helicity can be evaluated accurately at all points in the field. The numerical stability of orbital elements has been noted by other authors in similar situations (for example, Martin and Schaub [240]).

I therefore recommend against the use of the Cartesian transformation, except in isotropic systems or where one scales each term in the parameterisation non-uniformly using different unit lengths. In challenging test cases, I recommend the use of orbital elements to define initial conditions, which have produced usable insight even for the challenging problem parameters selected here. They also provide a more natural expression of defining initial conditions in the context of mission design. Nonetheless, other transformations may still yield better results depending on the system being studied.

I again stress that since DA provides derivatives accurate to machine precision, the above numerical difficulties are a result of double precision floating-point arithmetic. The use of quadruple precision (allowing up to 36 significant digits) would prevent round-off errors until approaching condition numbers of 10^{36} . However, quadruple precision is not currently implemented natively on most CPUs, and software solutions are generally computationally very intensive and do not necessarily guarantee reproducibility between platforms [237, 238].

6.6 Conclusion

This chapter has presented the the results of the application of DA-LCS, a new numerical method suggested earlier in this thesis for the computation of Lagrangian Coherent Structures in three dimensional astrodynamics systems. Two three-dimensional LCS in the Sun-Mars ER3BP have been constructed, corresponding to a proof-of-concept test case and a numerically challenging test case designed to benchmark the method against. Detailed analysis of the hyperbolic repelling LCS highlighted how it in this case acts as a separator between regions of qualitatively different behaviour, generalising the concept of the unstable manifold to systems with arbitrary time-dependence. This enables similar investigations to those in chapter 4 directly in time-dependent models of motion.

It was discovered that sufficient integration time is required to ensure the LCS found reflects the system dynamics, and that care must be taken in defining initial conditions to avoid regions of particularly poor numerical numerical performance due to limitations in floating-point arithmetic. While the LCS was determined here for ER3BP, there is no obstacle to computing it in other systems provided the arithmetic can be readily replaced by DA operations; high-fidelity ephemeris models have been tested with DA-LCS and shown to yield similar insight as for the ER3BP case studied here.

7

Conclusions

The initial goals of this thesis were to answer the following questions:

1. Can modern, high-performance computing methods give greater insight into dynamical structures in time-independent approximations to motion?
2. Is there a way in which we can perform this analysis directly in the time-dependent models of motion?

After reviewing the literature in chapter 3, the importance of the invariant manifolds in modern space mission design was highlighted. Chapter 4 used these invariant manifolds to assist with the retrieval of asteroids into periodic orbits about the Sun-Earth L_1 point. A gap in the literature was identified, in that significant assumptions about the properties of the transfer were made, potentially as a result of limited computational resources.

By using high-performance computing techniques – and thus providing an answer for the first goal of this thesis – many of the constraints these assumptions placed on the retrieval trajectories were lifted, expanding the search space for which valid retrieval missions could be constructed. Techniques for handling the increased search space were also introduced. As a result of this investigation, the number of asteroids retrievable using this method was increased by 147%. Moreover, additional properties about their transfer characteristics were revealed, including that many of the asteroids are retrievable across a range of transfer times, including a single impulse. The analysis approach presented in this thesis also holds wider applications, as disposal or deflection missions using the unstable manifolds can be constructed using the sample three-step ‘pipeline’.

Since this study relies on dynamical phenomena that is more numerically difficult to find in time-periodic systems – and which is not guaranteed to even exist in time-aperiodic systems – two investigations were performed to assess whether techniques that fluid dynamicists have

designed to tackle this problem could reasonably be applied to space flight problems. An improved numerical method that is usable for sufficiently-differentiable dynamical systems was introduced, which was found to reproduce and outperform sample cases from fluid dynamics that use divided differences, before greatly outperforming standard approaches on an example test case from astrodynamics.

Lastly, an in-depth investigation was performed to confirm that the hyperbolic repelling LCS does indeed perform the desired role in astrodynamics test cases in separating regions of qualitatively different behaviour, generalising the role of the hyperbolic manifolds. Implementation considerations and guiding principles for how to choose the parameterisations to use with LCS were outlined for use in future studies that may wish to work directly in time-dependent models of motion.

7.1 Limitations and future work

Suggestions for follow-up work and the limitations of the investigations that perform this thesis are now discussed.

Asteroid Retrieval

The asteroid retrieval work has provided an updated optimisation methodology for finding Easily-Retrievable Objects using high-performance computing techniques. However, there are several limitations in the method. The first is that the ranges of Jacobi constant used for generating periodic orbits and their associated manifolds were based on previous literature. However, it is possible to extend the lower bound of the Jacobi constant further, corresponding to orbits that move further away from the equilibrium points. This would yield more target periodic orbits and invariant manifolds, and with it more possible targets for retrieval. Additional families of periodic orbits other than the planar and vertical Lyapunov and ‘Halo’ orbits could be considered, as many other orbits have been investigated in mission design.

The stability of these orbits could also be investigated, as it may be favourable to place an asteroid on an unstable orbit such that ‘post-mission disposal’ also occurs naturally using the system dynamics. This would aid concerns raised around planetary protection, and the dangers associated with moving large and massive asteroids closer to the Earth. It may be possible to bias the optimisation procedure toward trajectories which have more desirable post-capture characteristics through the definition of a heuristic measure which considers post-capture dynamics. This heuristic could be based on the minimum distance to the Earth after completing a set number of revolutions in the orbit at the equilibrium point.

The 500 m/s threshold for labelling an object as an ERO and the threshold for labelling an object a ‘retrieval candidate’ were also set arbitrarily by García Yáñez, Sanchez, and McInnes [71]. Preliminary analysis as part of this thesis showed that the latter was set too high and restricted some retrieval candidates from becoming full EROs in the optimisation procedure, suggesting

instead 3 km/s. However, there has been no reconsideration of the 500 m/s retrieval limit. A reassessment of this threshold based on current technology readiness levels may be appropriate to confirm that 500 m/s is still an appropriate threshold, or whether it can be raised. The possibility of providing an impulse transfer of 500 m/s to objects which could weigh hundreds of tonnes may also not be feasible. Repeating this analysis using some of the retrieval technologies present in the literature, such as the gravity tractor or ion beam shepherd, would be a worthwhile pursuit. Moreover, one could return only small samples of the asteroid, rather than the entire object; this was investigated in Ionescu, McInnes, and Ceriotti [74].

The use of the impulsive transfer approximation is noted as a potential limitation, as the high mass of many of the EROs found makes impulsive transfers physically impossible. However, given that many of the transfers presented here have transfer times on the order of hundreds of days, constant low-thrust propulsion may be able to replace the two impulses for similar Δv , as performed in Sánchez and García Yáñez [179]. Given that this thesis has identified that transfers exist for similar cost across a range of transfer times, there should now be more transfers for which this is possible for each ERO.

Lastly, the asteroid investigation performed here found significant improvements on many of the transfers found in the literature when performing the full scan along transfer time. It is expected that similar improvements would be possible for objects that are classified as a retrieval candidate but do not become full EROs. This may mean the population of EROs could increase if such objects are subjected to the full optimisation procedure.

Lagrangian Coherent Structures

The study of LCS is more theoretical and less mature than the optimisation of asteroid retrieval trajectories, and because this is the first time known to the author that LCS have been constructed in three-dimensional astrodynamics systems, there are many avenues to take the work.

The most immediate investigation that should be performed is the extension of the concepts introduced as part of DA-LCS to n -dimensional LCS, i.e. the construction of LCS in the full phase space. Where this previously relied on quantities and derivatives that were numerically very difficult (such as $\nabla^2 \lambda_n$), the methods in DA-LCS make the computational of these straightforward. A preliminary extension of DA-LCS to six-dimensional flows in the ER3BP was performed with little implementation overhead to show that the concepts in DA-LCS extend to the required expressions in n -dimensional LCS. However, there is significant difficulty that comes with storing, computing and visualising six-dimensional structures. While storage could take place as a series of n -dimensional points, it may also be possible to use a DA representation of the surface to store a more exact version by directly representing the surface as a series of ‘patches’ described by polynomials rather than by interpolating between known points. The computation of these structures will not likely require much more computational effort as the integration of a n -dimensional system is still performed in this thesis and $C_{t_0}^T$ is still ‘small’. Nonetheless, the visualisation of n -dimensional surfaces, particularly for $n \geq 4$, will prove challenging.

In chapter 6, in the second test case I sampled the helicity on a regular 100×2500 grid to find seed points for the strainlines. However, the major advantage of an LCS is that strainlines can be propagated forward once a seed point is known, avoiding the computation of trajectories that do not otherwise correspond to the structure of the LCS. It would be a worthwhile pursuit to examine whether more sophisticated grid search strategies would reduce computation time, potentially exploiting a DA expansion of the helicity in combination with a gradient descent method to identify minima. Moreover, the extension of LCS to n -dimensions would avoid the need to define the mappings introduced in this thesis, and is numerically more simple to undertake than standard approaches if one follows the same methods used in DA-LCS. However, the visualisation of n -dimensional structures remains a challenge, and so significant effort would be required to store, manipulate and visualise these structures.

The above gradient descent method could also be used to improve the robustness of the strainline integration. It was mentioned in chapter 6 that the numerical integration of the strainlines can be difficult. This is because the numerical integration is being judged on its ability to track a ridge defined with respect to $\nabla \times \zeta_n$, when the integrator is attempting to control the error using an ODE defined with respect to ζ_n . Moreover, when evaluating the strainline ODE using a standard numerical integrator with variable step size, an estimate of the error is made by comparing the error in the obtained solution at two different orders of stepper (the truncation error), which is then used to refine the step size according to an update rule internal to the integrator. For the strainline ODE, outside of the low-helicity regions the eigenvectors are very ‘slowly’ and smoothly varying, which may lead to low truncation error and the step size being increased accordingly. Thus, once one step exceeds the low-helicity regions and enters the slowly-varying region, the integrator is likely to make even larger steps and further move away from the desired low-helicity points, terminating the integration. This can lead to very difficult and user-intensive integrations of the relevant strainlines as the update rule may need to be adjusted, or a fixed step-size used.

Instead, given that one knows the next point on the strainline should be on the local minimum of helicity, it would be interesting to see if the error between the solution predicted by the integrator and the local helicity minimum could be used to estimate the integrator error and perform step size adjustments¹, exploiting the gradient descent described above. This could also be extended to the relevant quantities for the n -dimensional formulation. This avenue of research may improve how ‘automatic’ the LCS determination method is.

The next avenue that could be taken is the application of DA-LCS to the attracting hyperbolic LCS in three dimensions. Preliminary investigations as part of the course of this PhD have shown that attracting LCS can be determined in the ER3BP, but it is currently unclear exactly what the surfaces are indicating. Moreover, the condition numbers associated with the relevant terms are orders of magnitude higher than required for the repelling LCS, which makes their

¹By way of analogy, this is akin to how Physics-informed Neural Networks penalise their error function with a known quantity about the underlying dynamics.

accurate computation more difficult. However, libraries such as CADNA [241] can monitor and compensate for the presence of round-off in common scientific computing languages, and the use of quadruple-precision rather than double-precision numbers would aid the avoidance of round-off errors until condition numbers as high as 10^{36} . The use of such techniques in LCS would enable its application to systems where the required terms are currently too ill-conditioned to provide accurate results; these techniques may also hold relevance to other analysis methods, since it is largely the derivatives of the underlying system that cause numerical difficulties rather than the analysis method itself.

Another topic of interest should be the extension of the methods made available in DA-LCS to other types of LCS. The elliptic LCS and parabolic LCS may hold some relevance to spaceflight problems, and exist in three dimensions as the zero-helicity points of the middle eigenvector of $C_{t_0}^T$.

Lastly, while this thesis has outlined the requisite tools to make the application of LCS to three-dimensional astrodynamics systems possible, I must acknowledge the method's limitations. These are generally grouped into two areas: the computational expense, and the requirement for the user to choose a threshold to denote "low-helicity" points.

Firstly, the LCS can be incredibly computationally expensive to compute in full for astrodynamics systems. While the helicity field is relatively quick to evaluate for the ER3BP shown here – on the order of hours for a single plane on a modern, high-performance desktop for the examples shown here – the computation of strainlines accurately is computationally very demanding as the LCS is defined on very narrow ridges of helicity, and the grid of seed points had to be limited in this thesis even with access to dedicated computing clusters. The use of the FTLE or LD, while not guaranteed to highlight only system dynamics, is computationally much quicker to evaluate as the former only requires first-order expansions, and the latter requires evaluation in the constant part only. If variational equations are known – which is often the case if one knows the equations of motion – or the user is willing to tune grid sizes to use divided differences, the use of DA can be removed, and massively parallel techniques on Graphics Processing Units (GPUs) can be used to greatly improve computational performance.

Secondly, the need to determine a helicity threshold is a major limitation of the method. While the helicity itself is objective and now free from errors due to tuning grid sizes, once the helicity threshold is chosen it is generally then the opinion of the user as to what constitutes the LCS. *In extremis*, I could select none or all of the helicity field to classify as an LCS, and either introduce false negatives by removing these points, or introduce false positives by erroneously including them. Indeed, in a paper released in 2020, George Haller himself recognised that incorporating tuneable parameters and thresholds is frustrating the efforts of coherent structure detection [235]². In the case of this thesis, *a posteriori* analyses of the behaviour of the LCS and examinations of

²He goes on to derive LCS as now also blocking diffusive transport in fluids, which is objective in terms of the mathematics without need for tuneable parameters. Unfortunately, this method relies on a vector field that follows the Cauchy momentum equation, which is not valid for astrodynamics systems.

the literature were used to show that the strainlines obtained were indeed an LCS by separating regions of different behaviour, but the endeavour would be made much more difficult if this information was unavailable.

However, there is in many fields a trade-off between accuracy and computational complexity, and the LCS, the FTLE and the LD have still emerged as effective tools in both the fluid dynamics and astrodynamics communities. Overall, the author sees the LCS as providing an analysis technique that can be trusted to produce usable insight based on intrinsic, mathematical properties of the underlying system at the cost of computational complexity. Of course, only the mission designer can decide how to determine the run-time *versus* accuracy trade-off based on the problem at hand, but this thesis makes available all the foundational methods, required prior information and implementation considerations required to use LCS in astrodynamics systems.

A

Appendix A: Derivation of the Tisserand parameter from Keplerian elements

This Appendix derives the Jacobi energy as a function of the osculating (instantaneous) orbital elements. This was used extensively in previous asteroid retrieval research to filter unsuitable NEOs from further optimisation procedures.

I begin from the expression for the Jacobi constant in an inertial frame due to Szebehely and Jefferys [33]. For a body with Cartesian position (x, y, z) in the Circular-Restricted Three Body Problem with mass parameter μ , this reads

$$\dot{x}^2 + \dot{y}^2 + \dot{z}^2 - 2(x\dot{y} - y\dot{x}) = \frac{2}{r_1} - \frac{2\mu}{r_1} + \frac{2\mu}{r_2} - C \quad (\text{A.1})$$

where r_1 and r_2 are, as before, the distance between the particle and the masses m_1 and m_2 . Under the unit normalisations in the CR3BP, one also has

$$G(m_1 + m_2) = 1. \quad (\text{A.2})$$

If one makes the assumption that $\mu \ll 1$, then in the inertial reference frame it can be assumed that the centre of the system can be taken to be the centre of mass of m_1 , and thus from the *vis-viva* equation

$$\dot{x}^2 + \dot{y}^2 + \dot{z}^2 = \frac{2}{r} - \frac{1}{a} \quad (\text{A.3})$$

since $r_1 \sim r$ if $\mu \sim 0$. Moreover, at every instant of time the angular momentum of the particle is given by its instantaneous orbital elements about m_1 as

$$h^2 = a(1 - e^2) \quad (\text{A.4})$$

and the quantity $x\dot{y} - y\dot{x}$ relates to the orbital inclination by

$$x\dot{y} - y\dot{x} = h \cos i. \quad (\text{A.5})$$

Substituting these into Equation A.1 yields

$$\frac{2}{r} - \frac{1}{a} - 2\sqrt{a(1-e^2)} \cos i = \frac{2}{r_1} - \frac{2\mu}{r_2} - C \quad (\text{A.6})$$

and since $r \approx r_1$ and any term that has μ as a coefficient can be deemed approximately zero. Thus,

$$C' = \frac{1}{2}a + \sqrt{a(1-e^2)} \cos i \approx J \quad (\text{A.7})$$

as was used in the literature.

B

Appendix B: Inverting a matrix of DA objects

It was noted in chapter 5 that the attracting LCS can be constructed by a similar procedure to the repelling LCS. By inverting the expansion of the Cauchy-Green Strain Tensor $[C_{t_0}^T]$ to obtain $[C_{t_0}^T]^{-1}$, the dominant eigenvector now signals the attracting LCS.

Standard double-precision inversion algorithms or algorithms designed to invert the constant part of $[C_{t_0}^T]$ do not respect that the expansion has higher-order terms. Instead, one needs to invert the constant portion of $[C_{t_0}^T]$ and then successively correct the higher-order terms to yield the full expansion of the inverse $[C_{t_0}^T]^{-1}$.

To derive this iterative technique, consider that $[C_{t_0}^T]$ is comprised of a constant part $C_{t_0}^T$, and higher-order terms $\delta C_{t_0}^T$. Its inverse is therefore $[C_{t_0}^T]^{-1} = (C_{t_0}^T)^{-1} + \delta(C_{t_0}^T)^{-1}$, such that

$$[C_{t_0}^T]^{-1} [C_{t_0}^T] = I \quad (B.1)$$

$$= ((C_{t_0}^T)^{-1} + \delta(C_{t_0}^T)^{-1})(C_{t_0}^T + \delta C_{t_0}^T) \quad (B.2)$$

$$= I + (C_{t_0}^T)^{-1} \delta C_{t_0}^T + \delta(C_{t_0}^T)^{-1} C_{t_0}^T + \delta(C_{t_0}^T)^{-1} \delta C_{t_0}^T \quad (B.3)$$

where I is the identity matrix. Cancelling the identity matrix on both sides and rearranging some terms gives

$$\delta(C_{t_0}^T)^{-1} C_{t_0}^T = (C_{t_0}^T)^{-1} \delta C_{t_0}^T + \delta(C_{t_0}^T)^{-1} \delta C_{t_0}^T. \quad (B.4)$$

By multiplying by the constant part of the inverse of the Cauchy-Green Strain Tensor, one obtains the relation

$$\delta(C_{t_0}^T)^{-1} = (C_{t_0}^T)^{-1} \delta C_{t_0}^T (C_{t_0}^T)^{-1} + \delta(C_{t_0}^T)^{-1} \delta C_{t_0}^T (C_{t_0}^T)^{-1}. \quad (B.5)$$

The first term in Equation B.5 is constant, and the second is contracting with an infinitely-small contraction factor $\delta C_{t_0}^T$. This suggests the iterative scheme

$$\delta (C_{t_0}^T)^{-1} = (C_{t_0}^T)^{-1} \delta C_{t_0}^T (C_{t_0}^T)^{-1} + \delta (C_{t_0}^T)^{-1} \delta C_{t_0}^T (C_{t_0}^T)^{-1}_{i-1}. \quad (\text{B.6})$$

The above scheme is superconvergent, requiring k iterations to correct k orders of $\delta (C_{t_0}^T)^{-1}$.

References

- [1] Our World in Data. *Annual number of objects launched into space*. 2022. url: <https://ourworldindata.org/grapher/early-number-of-objects-launched-into-outer-space> (visited on 09/03/2022).
- [2] Alessio Quinci, Gianmario Merisio, and Francesco Topputo. “Qualitative Study of Ballistic Capture at Mars via Lagrangian Descriptors”. Submitted, under review. 2022.
- [3] Daniel Blazevski and George Haller. “Hyperbolic and elliptic transport barriers in three-dimensional unsteady flows”. In: *Physica D: Nonlinear Phenomena* 273-274 (2014), pp. 46–62. issn: 01672789. doi: [10.1016/j.physd.2014.01.007](https://doi.org/10.1016/j.physd.2014.01.007).
- [4] NASA JPL. *Orbits of Potentially Hazardous Asteroids (PHAs)*. 2013. url: <https://www.jpl.nasa.gov/images/pia17041-orbits-of-potentially-hazardous-asteroids-phas> (visited on 10/18/2022).
- [5] Alexander Wittig et al. “Propagation of large uncertainty sets in orbital dynamics by automatic domain splitting”. In: *Celestial Mechanics and Dynamical Astronomy* 122.3 (July 2015), pp. 239–261. issn: 0923-2958. doi: [10.1007/s10569-015-9618-3](https://doi.org/10.1007/s10569-015-9618-3).
- [6] Mirco Rasotto et al. “Differential Algebra Space Toolbox for Nonlinear Uncertainty Propagation in Space Dynamics”. In: *International Conference on Astrodynamics Tools and Techniques* 6 (2016).
- [7] Xavier Ros Roca. “Computation of Lagrangian Coherent Structures with Application to Weak Stability Boundaries”. PhD thesis. Politecnico di Milano, 2015, p. 46.
- [8] Aaalok Parkash. “Application of Lagrangian Coherent Structures to the computation and understanding of ballistic capture trajectories”. MA thesis. Delft University of Technology, 2019, p. 110.
- [9] Pat Duggins. *Final Countdown: NASA and the End of the Space Shuttle Program*. University Press of Florida, 2009. isbn: 0813033845.
- [10] Javier Roa Vicens. “Regularization in Astrodynamics : applications to relative motion, low-thrust missions, and orbit propagation”. PhD thesis. Universidad Politecnica de Madrid, 2016, p. 254.

[11] Jack Tyler and Alexander Wittig. “On asteroid retrieval missions enabled by invariant manifold dynamics”. In: *Acta Astronautica* 183 (2021), pp. 43–51. issn: 00945765. doi: [10.1016/j.actaastro.2021.03.002](https://doi.org/10.1016/j.actaastro.2021.03.002).

[12] Jack Tyler and Alexander Wittig. “Flexible and Single-Impulse Transfers for Asteroid Retrieval using the Invariant Manifolds of the Circular-Restricted Three-Body Problem”. In: *31st AAS/AIAA Space Flight Mechanics Meeting, Charlotte, North Carolina*. American Institute of Aeronautics and Astronautics, 2021.

[13] Jack Tyler and Alexander Wittig. “An Improved Numerical Method for Three-dimensional Hyperbolic Lagrangian Coherent Structures using Differential Algebra”. In: *Journal of Computational Science* (2022).

[14] Jack Tyler and Alexander Wittig. “Three-dimensional Lagrangian Coherent Structures in Astrodynamics Applications using Differential Algebra”. In: *2022 AAS/AIAA Astrodynamics Specialist Conference, Charlotte, North Carolina*. American Institute of Aeronautics and Astronautics, 2022.

[15] Jack Tyler and Alexander Wittig. “Three-dimensional Lagrangian Coherent Structures in the Elliptic-Restricted Three-body Problem”. Submitted to *Celestial Mechanics and Dynamical Astronomy*, under review. 2022.

[16] Jack Tyler and Alexander Wittig. “Using Lagrangian Coherent Structures for Mission Design and Analysis”. In: *73rd International Astronautical Congress, Paris, France*. International Astronautical Federation, 2022.

[17] Howard D Curtis. *Orbital mechanics for engineering students*. 2nd Edition. Oxford: Elsevier, 2010, xiii, 722 p. isbn: 9780123747785 (hbk.)\r0123747783 (hbk.)\r9781856179546 (hbk. with on line testing)\r1856179540 (hbk. with on line testing).

[18] Ya Min Wang, Dong Qiao, and Ping Yuan Cui. “Design of low-energy transfer from lunar orbit to asteroid in the Sun-Earth-Moon system”. In: *Acta Mechanica Sinica/Lixue Xuebao* (2014). issn: 05677718. doi: [10.1007/s10409-014-0071-4](https://doi.org/10.1007/s10409-014-0071-4).

[19] Wang Sang Koon et al. “The Genesis Trajectory and Heteroclinic Connections”. In: (1999).

[20] Francesco Topputo. “Low-Thrust Non-Keplerian Orbits: Analysis, Design, and Control”. PhD Thesis. Milan: Politecnico di Milano, Mar. 2007.

[21] Leonard Euler. “De motu rectilineo trium corporum se mutuo attrahentium”. In: 11 (1767), pp. 144–151.

[22] Joseph-Louis Lagrange. *Essai sur le problème des trois corps*. Tech. rep. 1772.

[23] George William Hill. “On the part of the motion of the lunar perigee which is a function of the mean motions of the sun and moon”. In: *Acta Mathematica* 8.1 (1900), pp. 1–36. issn: 00015962. doi: [10.1007/BF02417081](https://doi.org/10.1007/BF02417081).

[24] Henri Poincaré. *Les Méthodes Nouvelles de la Mécanique Céleste*. Gauthier-Villars, 1899.

[25] George Howard Darwin. “Periodic Orbits”. In: *Acta Mathematica* 21.1 (1897), pp. 99–242. issn: 00015962. doi: [10.1007/BF02417978](https://doi.org/10.1007/BF02417978).

[26] Elis Strömgren. “Connaissance actuelle des orbites dans le problème des trois corps”. In: *Bull. Astron.* 9 (Jan. 1933), pp. 87–130.

[27] Forest Ray Moulton. “A Class of Periodic Orbits of Superior Planets”. In: *Transactions of the American Mathematical Society* 13.1 (1912), p. 96. issn: 00029947. doi: [10.2307/1988617](https://doi.org/10.2307/1988617).

[28] Tullio Levi-Civita. “Sopra alcuni criteri di instabilità”. In: *Annali di Matematica Pura ed Applicata* 5.1 (1901), pp. 221–305. issn: 16181891. doi: [10.1007/BF02419205](https://doi.org/10.1007/BF02419205).

[29] Aleksandr Lyapunov. “Problème général de la stabilité du mouvement”. In: *Annales de la Faculté des Sciences de Toulouse*, vol 9 (1907).

[30] Milovan Šuvakov and V. Dmitrašinović. “Three classes of newtonian three-body planar periodic orbits”. In: *Physical Review Letters* 110.11 (Mar. 2013). issn: 00319007. doi: [10.1103/PhysRevLett.110.114301](https://doi.org/10.1103/PhysRevLett.110.114301).

[31] David A Vallado. *Fundamentals of astrodynamics and applications*. Vol. 12. Springer Science & Business Media, 2001.

[32] Roger R Bate et al. *Fundamentals of astrodynamics: 2nd edition*. Courier Dover Publications, 2020.

[33] Victor Szebehely and William H. Jefferys. “Theory of Orbits: The Restricted Problem of Three Bodies”. In: *American Journal of Physics* 36.4 (1968), pp. 375–375. issn: 0002-9505. doi: [10.1119/1.1974535](https://doi.org/10.1119/1.1974535).

[34] Victor G Szebehely. *Theory of Orbits - The Restricted Problem of Three Bodies*. Academic Press, 1967.

[35] Wang Sang Koon et al. *Dynamical Systems, the Three-body Problem and Space Mission Design*. Marsden Books, 2008. isbn: 978-0-615-24095-4.

[36] Cody R. Short, Kathleen Howell, and Xavier M. Tricoche. “Lagrangian coherent structures in the restricted three-body problem”. In: *Advances in the Astronautical Sciences* 140 (2011), pp. 2111–2126. issn: 00653438.

[37] Michel Hénon. “Numerical Exploration of the Restricted Problem. V. Hill’s Case: Periodic Orbits and Their Stability”. In: *Astronomy and Astrophysics* (1969), pp. 223–238.

[38] Michel Hénon. “Numerical exploration of the restricted problem. VI. Hill’s case: Non-periodic orbits.” In: *Astronomy and Astrophysics* 9 (1970), pp. 24–36. issn: 0004-6361. doi: [10.1017/CBO9781107415324.004](https://doi.org/10.1017/CBO9781107415324.004).

[39] Michel Hénon. “Numerical Exploration of the Restricted Problem. V. Hill’s Case: Periodic Orbits and Their Stability”. In: *Astronomy and Astrophysics* (1969), pp. 223–238.

[40] T. Bray and C. Goudas. “Doubly Symmetric Orbits about the Collinear Lagrangian Points”. In: *The Astronomical Journal* 72.2 (1967).

[41] Manolis Michalodimitrakis. “Hill’s problem: Families of three-dimensional periodic orbits (part I)”. en. In: *Astrophysics and Space Science* 68.1 (Mar. 1980), pp. 253–268. issn: 0004-640X, 1572-946X. doi: [10.1007/BF00641660](https://doi.org/10.1007/BF00641660).

[42] Tycho T. von Rosenvinge, John C. Brandt, and Robert W. Farquhar. “The International Cometary Explorer Mission to Comet Giacobini-Zinner”. In: *Science* 232.4748 (Apr. 1986). Publisher: American Association for the Advancement of Science, pp. 353–356. doi: [10.1126/science.232.4748.353](https://doi.org/10.1126/science.232.4748.353).

[43] David L. Richardson. “Analytic construction of periodic orbits about the collinear points”. In: *Celestial Mechanics* 22.3 (1980), pp. 241–253. issn: 00088714. doi: [10.1007/BF01229511](https://doi.org/10.1007/BF01229511).

[44] Kathleen Howell. “Three-dimensional, periodic, ‘halo’ orbits”. In: *Celestial mechanics* 32.1 (Jan. 1984), pp. 53–71. issn: 1572-9478. doi: [10.1007/BF01358403](https://doi.org/10.1007/BF01358403).

[45] Kathleen Howell and H. Pernicka. “Numerical determination of Lissajous trajectories in the restricted three-body problem”. In: *Celestial Mechanics* 41 (1987), pp. 107–124. doi: [10.1007/BF01238756](https://doi.org/10.1007/BF01238756).

[46] Asif Azam Siddiqi. *Beyond Earth: A Chronicle of Deep Space Exploration, 1958-2016*. NASA SP. National Aeronautics and Space Administration, Office of Communications, NASA History Division, 2018. isbn: 978-1-62683-043-1.

[47] Wayne H. Yu and Karen Richon. “Launch Window Trade Analysis for the James Webb Space Telescope”. In: NTRS Author Affiliations: NASA Goddard Space Flight Center NTRS Report/Patent Number: GSFC-E-DAA-TN14126 NTRS Document ID: 20160001318 NTRS Research Center: Goddard Space Flight Center (GSFC). Laurel, MD, May 2014.

[48] Elisabet Canalias and Josep Masdemont. “Homoclinic and heteroclinic transfer trajectories between planar Lyapunov orbits in the Sun-Earth and Earth-Moon systems”. In: *Discrete and Continuous Dynamical Systems* 14.2 (2006), pp. 261–279. issn: 10780947. doi: [10.3934/dcds.2006.14.261](https://doi.org/10.3934/dcds.2006.14.261).

[49] Kathryn E. Davis et al. “The use of invariant manifolds for transfers between unstable periodic orbits of different energies”. In: *Celestial Mechanics and Dynamical Astronomy* 107.4 (2010), pp. 471–485. issn: 09232958. doi: [10.1007/s10569-010-9285-3](https://doi.org/10.1007/s10569-010-9285-3).

[50] Gerard. Gómez et al. “Invariant Manifolds, the Spatial Three-Body Problem and Space Mission Design”. en. In: ed. by David Bradley Spencer. Conference Name: AAS/AIAA Astrodynamics Specialists Conference Issue: 109 Meeting Name: AAS/AIAA Astrodynamics Specialists Conference Number: 109. San Diego: American Astronautical Society, 2001, pp. 3–22. isbn: 978-0-87703-488-9.

[51] C. C. Conley. “Low Energy Transit Orbits in the Restricted Three-Body Problem”. In: *SIAM Journal on Applied Mathematics* 16.4 (1968). Publisher: Society for Industrial and Applied Mathematics, pp. 732–746. issn: 0036-1399.

[52] Francesco Topputo, Massimiliano Vasile, and Franco Bernelli-Zazzera. “Low energy interplanetary transfers exploiting invariant manifolds of the restricted three-body problem”. In: *Journal of the Astronautical Sciences* 53.4 (2005), pp. 353–372. issn: 00219142. doi: [10.1007/bf03546358](https://doi.org/10.1007/bf03546358).

[53] Gerard Gómez et al. “Study of the transfer from the Earth to a halo orbit around the equilibrium point L1”. en. In: *Celestial Mechanics and Dynamical Astronomy* 56.4 (Aug. 1993), pp. 541–562. issn: 1572-9478. doi: [10.1007/BF00696185](https://doi.org/10.1007/BF00696185).

[54] Kathleen Howell, Brian Barden, and Martin Lo. “Application of Dynamical Systems Theory to Trajectory Design for a Libration Point Mission”. en. In: *The Journal of the Astronautical Sciences* 45.2 (June 1997), pp. 161–178. issn: 2195-0571. doi: [10.1007/BF03546374](https://doi.org/10.1007/BF03546374).

[55] Kathleen Howell, Brian Todd Barden, and R S Wilson. “Trajectory design using a dynamical systems approach with applications to GENESIS”. In: *AIAA Special Astrodynamics Conference*. Sun Valley, Idaho: AAS Publications Office, 1997, p. 33.

[56] Robert Farquhar, David Muñonen, and Daniel Richardson. “Mission design for a halo orbiter of the Earth”. In: *Journal of Spacecraft and Rockets* 14 (1977), pp. 170–177.

[57] David Dunham, James McAdams, and Robert Farquhar. “NEAR mission design”. In: *Johns Hopkins APL Technical Digest* 23 (Jan. 2002), pp. 18–33.

[58] Koon, Wang Sang et al. “Shoot the Moon”. In: *AAS/AIAA Astrodynamics Specialist Conference, Clearwater*. 2000, pp. 100–167.

[59] Edward Belbruno, Marian Gidea, and Francesco Topputo. “Weak stability boundary and invariant manifolds”. In: *SIAM Journal on Applied Dynamical Systems* 9.3 (2010), pp. 1061–1089. issn: 15360040. doi: [10.1137/090780638](https://doi.org/10.1137/090780638).

[60] Min Kun J. Chung et al. “Trans-lunar cruise trajectory design of GRAIL (Gravity Recovery and Interior Laboratory) mission”. In: *AIAA/AAS Astrodynamics Specialist Conference 2010*. 2010. isbn: 9781624101502. doi: [10.2514/6.2010-8384](https://doi.org/10.2514/6.2010-8384).

[61] G. D. Racca, G. P. Whitcomb, and B.H. Foing. “The ESA SMART-1 mission to the Moon with solar electric propulsion”. In: *Advances in Space Research* 23.11 (Jan. 1999), pp. 1865–1870. issn: 02731177. doi: [10.1016/S0273-1177\(99\)00544-X](https://doi.org/10.1016/S0273-1177(99)00544-X).

[62] Edward Belbruno and Brian G Marsden. “Resonance Hopping in Comets”. In: *The Astronomical Journal* 113.4 (1997), pp. 1433–1444. issn: 00046256. doi: [10.1086/118359](https://doi.org/10.1086/118359).

[63] Elisa Maria Alessi, Gerard Gomez, and Josep J Masdemont. “Leaving the Moon by means of invariant manifolds of libration point orbits”. In: *Communications in Nonlinear Science and Numerical Simulations* 14 (2009), pp. 4153–4167.

[64] Stefan Jerg, Oliver Junge, and Shane D. Ross. “Optimal capture trajectories using multiple gravity assists”. en. In: *Communications in Nonlinear Science and Numerical Simulation* 14.12 (Dec. 2009), pp. 4168–4175. issn: 1007-5704. doi: [10.1016/j.cnsns.2008.12.009](https://doi.org/10.1016/j.cnsns.2008.12.009).

[65] Francesco Topputo, Massimiliano Vasile, and Franco Bernelli-Zazzera. “Earth-to-moon low energy transfers targeting L1 hyperbolic transit orbits”. In: *Annals of the New York Academy of Sciences* 1065 (2005), pp. 55–76. issn: 00778923. doi: [10.1196/annals.1370.025](https://doi.org/10.1196/annals.1370.025).

[66] Rodney Anderson and Martin Lo. “Role of Invariant Manifolds in Low-Thrust Trajectory Design”. In: *Journal of Guidance, Control, and Dynamics* (2009). issn: 0731-5090. doi: [10.2514/1.37516](https://doi.org/10.2514/1.37516).

[67] Jeffrey S. Parker and Rodney L. Anderson. *Low-Energy Lunar Trajectory Design*. Tech. rep. July. 2014, pp. 1–396.

[68] Yi Qi and Shijie Xu. “Optimal earth-moon transfers using lunar gravity assist in the restricted four-body problem”. In: *Acta Astronautica* (2017). issn: 00945765. doi: [10.1016/j.actaastro.2017.02.002](https://doi.org/10.1016/j.actaastro.2017.02.002).

[69] Chiara Finocchietti, Pierpaolo Pergola, and Mariano Andrenucci. “Venus transfer design by combining invariant manifolds and low-thrust arcs”. In: *Acta Astronautica* 94 (2014), pp. 351–362.

[70] Daniel García Yáñez, Joan Pau Sanchez, and Colin R. McInnes. “Opportunities for asteroid retrieval missions”. In: *Asteroids: Prospective Energy and Material Resources* 9783642391 (2013), pp. 479–505. doi: [10.1007/978-3-642-39244-3__21](https://doi.org/10.1007/978-3-642-39244-3__21).

[71] Daniel García Yáñez, Juan Pao Sanchez, and Colin McInnes. “Easily retrievable objects among the NEO population”. In: *Celestial Mechanics and Dynamical Astronomy* (2013). issn: 09232958. doi: [10.1007/s10569-013-9495-6](https://doi.org/10.1007/s10569-013-9495-6).

[72] Daniel García Yáñez. “Exploiting Astrodynamics for the Manipulation and Exploration of Asteroids”. PhD thesis. University of Strathclyde, 2015, p. 204.

[73] Minghu Tan et al. “Direct and indirect capture of near-Earth asteroids in the Earth–Moon system”. In: *Celestial Mechanics and Dynamical Astronomy* 129 (2017), pp. 57–88. doi: [10.1007/s10569-017-9764-x](https://doi.org/10.1007/s10569-017-9764-x).

[74] Livia Ionescu, Colin R. McInnes, and Matteo Ceriotti. “A multiple-vehicle strategy for near-Earth asteroid capture”. en. In: *Acta Astronautica* (July 2022). issn: 0094-5765. doi: [10.1016/j.actaastro.2022.07.004](https://doi.org/10.1016/j.actaastro.2022.07.004).

[75] Stefano Campagnola, Martin Lot, and Paul Newton. “Subregions of motion and elliptic halo orbits in the elliptic restricted three-body problem”. In: *Advances in the Astronautical Sciences* 130 PART 2 (2008), pp. 1541–1555. issn: 00653438.

[76] Antonio Giorgilli et al. “Effective stability for a Hamiltonian system near an elliptic equilibrium point, with an application to the restricted three body problem”. In: *Journal of Differential Equations* 77.1 (1989), pp. 167–198. issn: 10902732. doi: [10.1016/0022-0396\(89\)90161-7](https://doi.org/10.1016/0022-0396(89)90161-7).

[77] I. A. Robin. “Bifurcations of plane with three-dimensional periodic orbits in the elliptic restricted problem”. In: *Celestial Mechanics* 23.1 (1981), pp. 97–106. issn: 00088714. doi: [10.1002/pamm.200700330](https://doi.org/10.1002/pamm.200700330).

[78] Francois Lekien and Shane D. Ross. “The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds”. In: *Chaos* 20.1 (2010). issn: 10541500. doi: [10.1063/1.3278516](https://doi.org/10.1063/1.3278516).

[79] George Haller. “Finding finite-time invariant manifolds in two-dimensional velocity fields”. en. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 10.1 (Mar. 2000), pp. 99–108. issn: 1054-1500, 1089-7682. doi: [10.1063/1.166479](https://doi.org/10.1063/1.166479).

[80] George Haller and Andrew Poje. “Finite time transport in aperiodic flows”. In: *Physica D: Nonlinear Phenomena* 119.3-4 (1998), pp. 352–380. issn: 01672789. doi: [10.1016/S0167-2789\(98\)00091-8](https://doi.org/10.1016/S0167-2789(98)00091-8).

[81] George Haller. “Distinguished material surfaces and coherent structures in three-dimensional fluid flows”. In: *Physica D: Nonlinear Phenomena* 149.4 (2001), pp. 248–277. issn: 01672789. doi: [10.1016/S0167-2789\(00\)00199-8](https://doi.org/10.1016/S0167-2789(00)00199-8).

[82] Shawn C. Shadden, Francois Lekien, and Jerrold E. Marsden. “Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows”. In: *Physica D: Nonlinear Phenomena* 212.3-4 (2005), pp. 271–304. issn: 01672789. doi: [10.1016/j.physd.2005.10.007](https://doi.org/10.1016/j.physd.2005.10.007).

[83] G. C. Yuan, L.J. Pratt, and C.K.R.T. Jones. “Barrier destruction and Lagrangian predictability at depth in a meandering jet”. en. In: *Dynamics of Atmospheres and Oceans* 35.1 (Feb. 2002), pp. 41–61. issn: 03770265. doi: [10.1016/S0377-0265\(01\)00082-3](https://doi.org/10.1016/S0377-0265(01)00082-3).

[84] Audrey Rogerson et al. “Lagrangian Motion and Fluid Exchange in a Barotropic Meandering Jet*”. en. In: *Journal of Physical Oceanography* 29.10 (Oct. 1999), pp. 2635–2655. issn: 0022-3670, 1520-0485. doi: [10.1175/1520-0485\(1999\)029<2635:LMAFEI>2.0.CO;2](https://doi.org/10.1175/1520-0485(1999)029<2635:LMAFEI>2.0.CO;2).

[85] Igor Mezić and Stephen Wiggins. “A method for visualization of invariant sets of dynamical systems based on the ergodic partition”. In: *Chaos* 9.1 (1999), pp. 213–218. issn: 10541500. doi: [10.1063/1.166399](https://doi.org/10.1063/1.166399).

[86] Carolina Mendoza and Ana M. Mancho. “Hidden Geometry of Ocean Flows”. In: *Phys. Rev. Lett.* 105 (3 July 2010), p. 038501. doi: [10.1103/PhysRevLett.105.038501](https://doi.org/10.1103/PhysRevLett.105.038501).

[87] Caroline Mendoza, Ana Mancho, and Maria-Helené Rio. “The turnstile mechanism across the Kuroshio current: analysis of dynamics in altimeter velocity fields”. In: *Nonlinear Processes in Geophysics* 17.2 (Mar. 2010), pp. 103–111. issn: 1607-7946. doi: [10.5194/npg-17-103-2010](https://doi.org/10.5194/npg-17-103-2010).

[88] Ana Mancho et al. “Lagrangian descriptors: A method for revealing phase space structures of general time dependent dynamical systems”. In: *Communications in Nonlinear Science and Numerical Simulation* 18 (12 2013). doi: [10.1016/j.cnsns.2013.05.002](https://doi.org/10.1016/j.cnsns.2013.05.002).

[89] Guillermo García-Sánchez et al. “Structured pathways in the turbulence organizing recent oil spill events in the Eastern Mediterranean”. In: *Scientific Reports* 12.1 (2022), p. 3662. issn: 2045-2322. doi: [10.1038/s41598-022-07350-w](https://doi.org/10.1038/s41598-022-07350-w).

[90] Coumba Niang et al. “Transport pathways across the West African Monsoon as revealed by Lagrangian Coherent Structures”. In: *Scientific Reports* 10.1 (2020), p. 12543. issn: 2045-2322. doi: [10.1038/s41598-020-69159-9](https://doi.org/10.1038/s41598-020-69159-9).

[91] Ahmed Darwish et al. *Extracting Lagrangian coherent structures in cardiovascular flows using Lagrangian descriptors*. 2021. doi: [10.1063/5.0064023](https://doi.org/10.1063/5.0064023).

[92] Víctor J García-Garrido and Stephen Wiggins. “Lagrangian descriptors and the action integral of classical mechanics”. In: *Physica D: Nonlinear Phenomena* 434 (2022), p. 133206. issn: 0167-2789. doi: <https://doi.org/10.1016/j.physd.2022.133206>.

[93] Matthaios Katsanikas et al. “Bifurcation study on a degenerate double van der Waals cirque potential energy surface using Lagrangian descriptors”. In: *Communications in Nonlinear Science and Numerical Simulation* 105 (2022), p. 106089. issn: 1007-5704. doi: <https://doi.org/10.1016/j.cnsns.2021.106089>.

[94] Priyanka Pandey, Shibabrat Naik, and Srihari Keshavamurthy. “Influence of low frequency modes on dynamical concertedness in double proton transfer dynamics”. In: *Communications in Nonlinear Science and Numerical Simulation* 109 (2022), p. 106326. issn: 1007-5704. doi: <https://doi.org/10.1016/j.cnsns.2022.106326>.

[95] Alessio Quinci. “Qualitative study of ballistic capture at Mars via Lagrangian descriptors”. MA thesis. Politecnico di Milano, 2021. doi: <http://hdl.handle.net/10589/186116>.

[96] Sebastiano Raffa. “Finding regions of bounded motion in binary asteroid environments using Lagrangian descriptors”. MA thesis. Politecnico di Milano, 2021. doi: <http://hdl.handle.net/10589/186129>.

[97] Sebastiano Raffa, Gianmario Merisio, and Francesco Topputo. “Finding Regions of Bounded Motion in Binary Asteroid Environment using Lagrangian Descriptors”. Submitted, under review. 2022.

[98] George Haller. “Lagrangian Coherent Structures”. In: *Annual Review of Fluid Mechanics* 47.1 (2015), pp. 137–162. doi: [10.1146/annurev-fluid-010313-141322](https://doi.org/10.1146/annurev-fluid-010313-141322).

[99] Akira Okubo. “Horizontal dispersion of floatable particles in the vicinity of velocity singularities such as convergences”. In: *Deep Sea Research and Oceanographic Abstracts* 17.3 (June 1970), pp. 445–454. issn: 0011-7471. doi: [10.1016/0011-7471\(70\)90059-8](https://doi.org/10.1016/0011-7471(70)90059-8).

[100] Julian Charles Roland Hunt, Alan A. Wray, and Parviz Moin. *Eddies, streams, and convergence zones in turbulent flows*. Tech. rep. NTRS Author Affiliations: Cambridge Univ. (England)., NASA Ames Research Center, Stanford Univ. NTRS Document ID: 19890015184 NTRS Research Center: Legacy CDMS (CDMS). NASA Ames Research Center, Dec. 1988.

[101] Jinhee Jeong and Fazle Hussain. “On the identification of a vortex”. en. In: *Journal of Fluid Mechanics* 285 (Feb. 1995). Publisher: Cambridge University Press, pp. 69–94. issn: 1469-7645, 0022-1120. doi: [10.1017/S0022112095000462](https://doi.org/10.1017/S0022112095000462).

[102] Naresh Malhotra, Igor Mezić, and Stephen Wiggins. “Patchiness: A new diagnostic for Lagrangian trajectory analysis in time-dependent fluid flows”. In: *International Journal of Bifurcation and Chaos in Applied Sciences and Engineering* 8.6 (1998), pp. 1053–1093. issn: 02181274. doi: [10.1142/S0218127498000875](https://doi.org/10.1142/S0218127498000875).

[103] Igor Mezić et al. “A New Mixing Diagnostic and Gulf Oil Spill Movement”. In: *Science* 330.6003 (2010), pp. 486–489. doi: [10.1126/science.1194607](https://doi.org/10.1126/science.1194607).

[104] Kenneth P Bowman. “Manifold Geometry and Mixing in Observed Atmospheric Flows”. en. In: *Journal of the Atmospheric Sciences* (1999), p. 29.

[105] Francois Lekien and Shane D. Ross. “The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds”. en. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 20.1 (Mar. 2010), p. 017505. issn: 1054-1500, 1089-7682. doi: [10.1063/1.3278516](https://doi.org/10.1063/1.3278516).

[106] Christopher Jones and Sean Winkler. “Invariant Manifolds and Lagrangian Dynamics in the Ocean and Atmosphere”. en. In: *Handbook of Dynamical Systems*. Vol. 2. Elsevier, 2002, pp. 55–92. isbn: 978-0-444-50168-4. doi: [10.1016/S1874-575X\(02\)80023-6](https://doi.org/10.1016/S1874-575X(02)80023-6).

[107] George Haller. “Lagrangian structures and the rate of strain in a partition of two-dimensional turbulence”. In: *Physics of Fluids* 13.11 (2001), pp. 3365–3385. issn: 10706631. doi: [10.1063/1.1403336](https://doi.org/10.1063/1.1403336).

[108] Tieh-Yong Koh and Bernard Legras. “Hyperbolic lines and the stratospheric polar vortex”. en. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 12.2 (June 2002), pp. 382–394. issn: 1054-1500, 1089-7682. doi: [10.1063/1.1480442](https://doi.org/10.1063/1.1480442).

[109] Guillaume Lapeyre. “Characterization of finite-time Lyapunov exponents and vectors in two-dimensional turbulence”. en. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 12.3 (Sept. 2002), pp. 688–698. issn: 1054-1500, 1089-7682. doi: [10.1063/1.1499395](https://doi.org/10.1063/1.1499395).

[110] Steven L. Brunton and Clarence W. Rowley. “Fast computation of finite-time Lyapunov exponent fields for unsteady flows”. In: *Chaos* 20.1 (2010). issn: 10541500. doi: [10.1063/1.3270044](https://doi.org/10.1063/1.3270044).

[111] Guoqiao You and Shingyu Leung. “An Improved Eulerian Approach for the Finite Time Lyapunov Exponent”. In: *Journal of Scientific Computing* 76.3 (2018), pp. 1407–1435. issn: 08857474. doi: [10.1007/s10915-018-0669-y](https://doi.org/10.1007/s10915-018-0669-y).

[112] David J. Gondelach, Roberto Armellin, and Alexander Wittig. “On the predictability and robustness of Galileo disposal orbits”. In: *Celestial Mechanics and Dynamical Astronomy* 131.12 (2019), pp. 1–30. issn: 15729478. doi: [10.1007/s10569-019-9938-9](https://doi.org/10.1007/s10569-019-9938-9).

[113] Cody R. Short et al. “Stretching in phase space and applications in general nonautonomous multi-body problems”. In: *Celestial Mechanics and Dynamical Astronomy* 122.3 (2015), pp. 213–238. issn: 15729478. doi: [10.1007/s10569-015-9617-4](https://doi.org/10.1007/s10569-015-9617-4).

[114] David Canales Garcia, Kathleen Howell, and Elena Fantino. “Leveraging Finite Time Lyapunov Exponent Maps to Design Tours Incorporating Three Moons”. In: *ASCEND2021*. American Institute of Aeronautics and Astronautics, 2021. doi: [10.2514/6.2021-4154](https://doi.org/10.2514/6.2021-4154).

[115] Rodney Anderson, Martin Lo, and George Born. *Application of local Lyapunov exponents to maneuver design and navigation in the three-body problem*. Tech. rep. NASA Jet Propulsion Laboratory, 2003.

[116] George Haller and G. Yuan. “Lagrangian coherent structures and mixing in two-dimensional turbulence”. In: *Physica D: Nonlinear Phenomena* 147.3-4 (2000), pp. 352–370. issn: 01672789. doi: [10.1016/S0167-2789\(00\)00142-1](https://doi.org/10.1016/S0167-2789(00)00142-1).

[117] George Haller. “Lagrangian coherent structures from approximate velocity data”. In: *Physics of Fluids* 14.6 (2002), pp. 1851–1861. issn: 10706631. doi: [10.1063/1.1477449](https://doi.org/10.1063/1.1477449).

[118] George Haller. “Exact theory of unsteady separation for two-dimensional flows”. In: *Journal of Fluid Mechanics* 512 (2004), pp. 257–311. issn: 00221120. doi: [10.1017/S0022112004009929](https://doi.org/10.1017/S0022112004009929).

[119] Francois Lekien, Shawn C. Shadden, and Jerrold E. Marsden. “Lagrangian coherent structures in n-dimensional systems”. In: *Journal of Mathematical Physics* 48.6 (2007). issn: 00222488. doi: [10.1063/1.2740025](https://doi.org/10.1063/1.2740025).

[120] Melissa Green, Clarence Rowley, and George Haller. “Detection of Lagrangian coherent structures in three-dimensional turbulence”. en. In: *Journal of Fluid Mechanics* 572 (Feb. 2007). Publisher: Cambridge University Press, pp. 111–120. issn: 1469-7645, 0022-1120. doi: [10.1017/S0022112006003648](https://doi.org/10.1017/S0022112006003648).

[121] Shawn C. Shadden, John O. Dabiri, and Jerrold E. Marsden. “Lagrangian analysis of fluid transport in empirical vortex ring flows”. en. In: *Physics of Fluids* 18.4 (Apr. 2006), p. 047105. issn: 1070-6631, 1089-7666. doi: [10.1063/1.2189885](https://doi.org/10.1063/1.2189885).

[122] Greg A. Voth, George Haller, and Jerry P. Gollub. “Experimental Measurements of Stretching Fields in Fluid Mixing”. In: *Physical Review Letters* 88.25 (June 2002). Publisher: American Physical Society, p. 254501. doi: [10.1103/PhysRevLett.88.254501](https://doi.org/10.1103/PhysRevLett.88.254501).

[123] Shawn C. Shadden and Charles A. Taylor. “Characterization of coherent structures in the cardiovascular system”. In: *Annals of Biomedical Engineering* 36.7 (2008), pp. 1152–1162. issn: 00906964. doi: [10.1007/s10439-008-9502-3](https://doi.org/10.1007/s10439-008-9502-3).

[124] Blake Rutherford, Gerhard Dangelmayr, and Michael T. Montgomery. “Lagrangian coherent structures in tropical cyclone intensification”. In: *Atmospheric Chemistry and Physics* 12.12 (2012), pp. 5483–5507. issn: 16807316. doi: [10.5194/acp-12-5483-2012](https://doi.org/10.5194/acp-12-5483-2012).

[125] Amir E. BozorgMagham and Shane D. Ross. “Atmospheric Lagrangian coherent structures considering unresolved turbulence and forecast uncertainty”. In: *Communications in Nonlinear Science and Numerical Simulation* 22.1-3 (May 2015), pp. 964–979. issn: 10075704. doi: [10.1016/j.cnsns.2014.07.011](https://doi.org/10.1016/j.cnsns.2014.07.011).

[126] L. A. Fiorentino et al. “Using Lagrangian Coherent Structures to understand coastal water quality”. In: *Continental Shelf Research* 47 (2012), pp. 145–149. issn: 02784343. doi: [10.1016/j.csr.2012.07.009](https://doi.org/10.1016/j.csr.2012.07.009).

[127] Han Tu et al. “Investigation of accelerating non-slender delta-wing planforms at high angle of attack using lagrangian coherent structures”. English (US). In: AIAA Scitech 2019 Forum (2019). Funding Information: This work was supported by the Office of Naval Research under ONR Award No. N00014-16-1-2732. Publisher Copyright: © 2019 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved.; AIAA Scitech Forum, 2019 ; Conference date: 07-01-2019 Through 11-01-2019. doi: [10.2514/6.2019-2165](https://doi.org/10.2514/6.2019-2165).

[128] Jeff D. Eldredge and Kwitae Chong. “Fluid transport and coherent structures of translating and flapping wings”. In: *Chaos* 20.1 (2010). issn: 10541500. doi: [10.1063/1.3273036](https://doi.org/10.1063/1.3273036).

[129] Gary Froyland, Simon Lloyd, and Naratip Santitissadeekorn. “Coherent sets for nonautonomous dynamical systems”. In: *Physica D: Nonlinear Phenomena* 239.16 (2010), pp. 1527–1541. issn: 01672789. doi: [10.1016/j.physd.2010.03.009](https://doi.org/10.1016/j.physd.2010.03.009).

[130] Manikandan Mathur et al. “Uncovering the Lagrangian skeleton of turbulence”. In: *Physical Review Letters* 98.14 (2007), pp. 1–13. issn: 00319007. doi: [10.1103/PhysRevLett.98.144502](https://doi.org/10.1103/PhysRevLett.98.144502).

[131] Roger Samelson. “Lagrangian Motion, Coherent Structures, and Lines of Persistent Material Strain”. In: *Annual Review of Marine Science* 5.1 (2013), pp. 137–163. issn: 1941-1405. doi: [10.1146/annurev-marine-120710-100819](https://doi.org/10.1146/annurev-marine-120710-100819).

[132] George Haller. “A variational theory of hyperbolic Lagrangian Coherent Structures”. In: *Physica D: Nonlinear Phenomena* 240.7 (Mar. 2011), pp. 574–598. issn: 01672789. doi: [10.1016/j.physd.2010.11.010](https://doi.org/10.1016/j.physd.2010.11.010).

[133] George Haller and Themistoklis Sapsis. “Lagrangian coherent structures and the smallest finite-time Lyapunov exponent”. In: *Chaos* 21.2 (2011), pp. 1–7. issn: 10541500. doi: [10.1063/1.3579597](https://doi.org/10.1063/1.3579597).

[134] George Haller and Francisco J. Beron-Vera. “Geodesic theory of transport barriers in two-dimensional flows”. In: *Physica D: Nonlinear Phenomena* 241.20 (2012), pp. 1680–1702. issn: 01672789. doi: [10.1016/j.physd.2012.06.012](https://doi.org/10.1016/j.physd.2012.06.012).

[135] Greg Norgard and Peer Timo Bremer. “Second derivative ridges are straight lines and the implications for computing Lagrangian Coherent Structures”. In: *Physica D: Nonlinear Phenomena* 241.18 (2012), pp. 1475–1476. issn: 01672789. doi: [10.1016/j.physd.2012.05.006](https://doi.org/10.1016/j.physd.2012.05.006).

[136] Daniel Karrasch and George Haller. “Do finite-size lyapunov exponents detect coherent structures?” In: *Chaos* 23.4 (2013). issn: 10541500. doi: [10.1063/1.4837075](https://doi.org/10.1063/1.4837075).

[137] Daniel Pérez-Palau. “Dynamical transport mechanisms in celestial mechanics and astrodynamics problems”. PhD thesis. Universitat de Barcelona, 2015.

[138] Mohammad Farazmand and George Haller. “Computing Lagrangian coherent structures from their variational theory”. In: *Chaos* 22.1 (2012). issn: 10541500. doi: [10.1063/1.3690153](https://doi.org/10.1063/1.3690153).

[139] Jacques Féjoz. “13. Introduction to KAM Theory with a View to Celestial Mechanics”. In: *13. Introduction to KAM Theory with a View to Celestial Mechanics*. De Gruyter, Jan. 2017, pp. 387–433. isbn: 978-3-11-043039-4. doi: [10.1515/9783110430394-013](https://doi.org/10.1515/9783110430394-013).

[140] Nicola Baresi, Zubin Olikara, and Daniel Scheeres. “Fully Numerical Methods for Continuing Families of Quasi-Periodic Invariant Tori in Astrodynamics”. In: *Journal of the Astronautical Sciences* 65 (2018). doi: [10.1007/s40295-017-0124-6](https://doi.org/10.1007/s40295-017-0124-6).

[141] Claude Froeschlé and Elena Lega. “On the Measure of the Structure around the Last Kam Torus before and after Its Break-Up”. In: *Celestial Mechanics and Dynamical Astronomy* 64.1 (Mar. 1996), pp. 21–31. issn: 1572-9478. doi: [10.1007/BF00051602](https://doi.org/10.1007/BF00051602).

[142] Francisco J. Beron-Vera et al. “Invariant-Tori-like Lagrangian Coherent Structures in Geophysical Flows”. In: *Chaos: An Interdisciplinary Journal of Nonlinear Science* 20.1 (Mar. 2010), p. 017514. issn: 1054-1500. doi: [10.1063/1.3271342](https://doi.org/10.1063/1.3271342).

[143] Karljohan Lundin Palmerius, Matthew Cooper, and Anders Ynnerman. “Flow field visualization using vector field perpendicular surfaces”. In: *Proceedings - SC'09: 25th Spring Conference on Computer Graphics* (2009), pp. 27–34. doi: [10.1145/1980462.1980471](https://doi.org/10.1145/1980462.1980471).

[144] Evan S. Gawlik et al. “Lagrangian coherent structures in the planar elliptic restricted three-body problem”. In: *Celestial Mechanics and Dynamical Astronomy* 103.3 (2009), pp. 227–249. issn: 09232958. doi: [10.1007/s10569-008-9180-3](https://doi.org/10.1007/s10569-008-9180-3).

[145] Cody R. Short and Kathleen Howell. “Lagrangian coherent structures in various map representations for application to multi-body gravitational regimes”. In: *Acta Astronautica* 94.2 (2014), pp. 592–607. issn: 00945765. doi: [10.1016/j.actaastro.2013.08.020](https://doi.org/10.1016/j.actaastro.2013.08.020).

[146] Kenta Oshima and Tomohiro Yanao. “Applications of gravity assists in the bicircular and bielliptic restricted four-body problem”. In: *Advances in the Astronautical Sciences*. Vol. 152. 2014, pp. 503–522.

[147] Qu Qingyu, Lin Mingpei, and Xu Ming. “Lagrangian Coherent Structures in the Planar Parabolic/Hyperbolic Restricted Three-Body Problem”. In: *Monthly Notices of the Royal Astronomical Society* (2020). issn: 0035-8711. doi: [10.1093/mnras/staa199](https://doi.org/10.1093/mnras/staa199).

[148] Kaori Onozaki, Hiroaki Yoshimura, and Shane Ross. “Tube Dynamics and Low Energy Earth-Moon Transfers in Teh 4-Body System”. In: *Advances in Space Research* 60 (2017), pp. 2117–2132. doi: <http://dx.doi.org/10.1016/j.asr.2017.07.046>.

[149] Daniel Pérez, Gerard Gómez, and Josep J. Masdemont. “Detecting invariant manifolds using Hyperbolic Lagrangian Coherent Structures”. In: *Advances in the Astronautical Sciences* 145 (2012), pp. 867–879. issn: 00653438.

[150] Patrick Pérez, Michael Gangnet, and Andrew Blake. “Poisson Image Editing”. In: *Image Processing On Line* 5 (2016), pp. 300–325. issn: 2105-1232. doi: [10.5201/ipol.2016.163](https://doi.org/10.5201/ipol.2016.163).

[151] Arne Magnus Tveita Løken. “Computing Hyperbolic Lagrangian Coherent Structures in Three- Dimensional Flow Systems”. MA thesis. 2018.

[152] P. Sánchez-Martín, J. J. Masdemont, and M. Romero-Gómez. “From manifolds to Lagrangian coherent structures in galactic bar models”. In: *Astronomy & Astrophysics* 618 (2018), A72. issn: 0004-6361. doi: [10.1051/0004-6361/201833451](https://doi.org/10.1051/0004-6361/201833451).

[153] Aline K. Zimmer. “Investigation of vehicle reusability for human exploration of Near-Earth Asteroids using Sun-Earth Libration point orbits”. In: *Acta Astronautica* (2013). issn: 00945765. doi: [10.1016/j.actaastro.2012.10.003](https://doi.org/10.1016/j.actaastro.2012.10.003).

[154] Yuichi Tsuda et al. “Hayabusa2–Sample return and kinetic impact mission to near-earth asteroid Ryugu”. In: *Acta Astronautica* 156 (Mar. 2019), pp. 387–393. issn: 00945765. doi: [10.1016/j.actaastro.2018.01.030](https://doi.org/10.1016/j.actaastro.2018.01.030).

[155] Andrea Accomazzo et al. “Rosetta operations at the comet”. In: *Acta Astronautica* (2015). issn: 00945765. doi: [10.1016/j.actaastro.2015.06.009](https://doi.org/10.1016/j.actaastro.2015.06.009).

[156] Dante S Lauretta. “OSIRIS-REx Asteroid Sample-Return Mission”. In: *Handbook of Cosmic Hazards and Planetary Defense*. Ed. by Joseph N Pelton and Firooz Allahdadi. Cham: Springer International Publishing, 2015, pp. 543–567. isbn: 978-3-319-03952-7. doi: [10.1007/978-3-319-03952-7__44](https://doi.org/10.1007/978-3-319-03952-7__44).

[157] Carlos de la Fuente Marcos and Raul de la Fuente Marcos. “Geometric characterization of the Arjuna orbital domain”. In: *Astronomische Nachrichten* 336.1 (2015), pp. 5–22. issn: 15213994. doi: [10.1002/asna.201412133](https://doi.org/10.1002/asna.201412133).

[158] Michael C.F. Bazzocchi and M. Reza Emami. “Study of arjuna-type asteroids for low-thrust orbital transfer”. In: *Journal of Spacecraft and Rockets* 55.1 (2018), pp. 37–48. issn: 15336794. doi: [10.2514/1.A33758](https://doi.org/10.2514/1.A33758).

[159] Ramon Brasser and Paul Wiegert. “Asteroids on Earth-like orbits and their origin”. In: *Monthly Notices of the Royal Astronomical Society* 386.4 (2008), pp. 2031–2038. issn: 00358711. doi: [10.1111/j.1365-2966.2008.13146.x](https://doi.org/10.1111/j.1365-2966.2008.13146.x).

[160] Task Force on Potentially Hazardous Near Earth Objects. *Report of the Task Force on Potentially Hazardous Near Earth Objects*. Tech. rep. British National Space Centre, 2000, p. 56.

[161] Peter Brown et al. “A 500-kiloton airburst over Chelyabinsk and an enhanced hazard from small impactors”. In: *Nature* 503 (2013). doi: [10.1038/nature12741](https://doi.org/10.1038/nature12741).

[162] Clark R. Chapman, Daniel D. Durda, and Robert E. Gold. *The comet/asteroid impact hazard: a systems approach*. Tech. rep. Space Engineering and Technology Branch: Johns Hopkins University Applied Physics Laboratory, 2001, pp. 1–18.

[163] Russell Schweickart et al. *Threat Mitigation: The Asteroid Tugboat*. 2006. doi: [10.48550/ARXIV.PHYSICS/0608156](https://doi.org/10.48550/ARXIV.PHYSICS/0608156).

[164] Matteo Ceriotti and Joan Pau Sanchez. “Control of asteroid retrieval trajectories to libration point orbits”. In: *Acta Astronautica* (2016). issn: 00945765. doi: [10.1016/j.actaastro.2016.03.037](https://doi.org/10.1016/j.actaastro.2016.03.037).

[165] Zaki Hasnain, Christopher A. Lamb, and Shane D. Ross. “Capturing near-Earth asteroids around Earth”. In: *Acta Astronautica* 81.2 (2012), pp. 523–531. issn: 00945765. doi: [10.1016/j.actaastro.2012.07.029](https://doi.org/10.1016/j.actaastro.2012.07.029).

[166] He Xi Baoyin, Yang Chen, and Jun Feng Li. “Capturing near earth objects”. In: *Research in Astronomy and Astrophysics* 10.6 (2010), pp. 587–598. issn: 16744527. doi: [10.1088/1674-4527/10/6/008](https://doi.org/10.1088/1674-4527/10/6/008).

[167] Neus Lladó et al. “Capturing small asteroids into a Sun-Earth Lagrangian point”. In: *Acta Astronautica* 95.1 (2014), pp. 176–188. issn: 00945765. doi: [10.1016/j.actaastro.2013.11.007](https://doi.org/10.1016/j.actaastro.2013.11.007).

[168] Thomas J. Ahrens and Alan W. Harris. “Deflection and fragmentation of near-Earth asteroids”. In: *Nature* 360.6403 (1992), pp. 429–433. issn: 00280836. doi: [10.1038/360429a0](https://doi.org/10.1038/360429a0).

[169] Juan L. Cano et al. “Mission analysis for the don quijote phase a study”. In: *International Astronautical Federation - 58th International Astronautical Congress 2007* 2 (2007), pp. 968–981.

[170] Edward T Lu and Stanley G Love. “Gravitational tractor for towing asteroids”. In: *Nature* 438.7065 (2005), pp. 177–178. issn: 14764687. doi: [10.1038/438177a](https://doi.org/10.1038/438177a).

[171] Claudio Bombardelli et al. “The ion beam shepherd: A new concept for asteroid deflection”. In: *Acta Astronautica* 90.1 (2013), pp. 98–102. issn: 00945765. doi: [10.1016/j.actaastro.2012.10.019](https://doi.org/10.1016/j.actaastro.2012.10.019).

[172] Claudio Bombardelli and Jesus Peláez. “Ion beam shepherd for asteroid deflection”. In: *Journal of Guidance, Control, and Dynamics* 34.4 (2011), pp. 1270–1272. issn: 15333884. doi: [10.2514/1.51640](https://doi.org/10.2514/1.51640).

[173] Claudio Bombardelli and Jesus Peláez. “Ion beam shepherd for contactless space debris removal”. In: *Journal of Guidance, Control, and Dynamics* 34.3 (2011), pp. 916–920. issn: 15333884. doi: [10.2514/1.51832](https://doi.org/10.2514/1.51832).

[174] Juan Pao Sanchez and Colin McInnes. “On the ballistic capture of asteroids for resource utilisation”. In: *62nd International Astronautical Congress 2011, IAC 2011* 6 (2011), pp. 4907–4921.

[175] Shane D. Ross and Daniel J. Scheeres. “Multiple gravity assists in the restricted three-body problem”. In: *Advances in the Astronautical Sciences* 127 PART 2.3 (2007), pp. 1881–1900. issn: 00653438.

[176] William F. Bottke et al. “Debiased orbital and absolute magnitude distribution of the near-Earth objects”. In: *Icarus* 156.2 (2002), pp. 399–433. issn: 00191035. doi: [10.1006/icar.2001.6788](https://doi.org/10.1006/icar.2001.6788).

[177] Giorgio Mingotti, Juan Pao Sánchez, and Colin McInnes. “Combined low-thrust propulsion and invariant manifold trajectories to capture NEOs in the Sun–Earth circular restricted three-body problem”. In: *Celestial Mechanics and Dynamical Astronomy* 120.3 (2014), pp. 309–336. issn: 15729478. doi: [10.1007/s10569-014-9589-9](https://doi.org/10.1007/s10569-014-9589-9).

[178] Matteo Ceriotti and Joan Pau Sanchez. “Control of asteroid retrieval trajectories to libration point orbits”. In: *Acta Astronautica* 126.March (Sept. 2016), pp. 342–353. issn: 00945765. doi: [10.1016/j.actaastro.2016.03.037](https://doi.org/10.1016/j.actaastro.2016.03.037).

[179] Joan Pau Sánchez and Daniel García Yáñez. “Asteroid retrieval missions enabled by invariant manifold dynamics”. In: *Acta Astronautica* 127 (2016), pp. 667–677. issn: 00945765. doi: [10.1016/j.actaastro.2016.05.034](https://doi.org/10.1016/j.actaastro.2016.05.034).

[180] Kathleen Howell, Brian Todd Barden, and Roby Wilson. “Trajectory design using a dynamical systems approach with applications to GENESIS”. In: *AIAA Special Astrodynamics Conference*. Sun Valley, Idaho: AAS Publications Office, 1997, p. 33.

[181] Michael E. Evans and Lee D. Graham. “A Flexible Lunar Architecture for Exploration (FLARE) supporting NASA’s Artemis Program”. In: *Acta Astronautica* 177 (Dec. 2020), pp. 351–372. issn: 0094-5765. doi: [10.1016/j.actaastro.2020.07.032](https://doi.org/10.1016/j.actaastro.2020.07.032).

[182] Simone Servadio, David Arnas, and Richard Linares. “Dynamics Near the Three-Body Libration Points via Koopman Operator Theory”. In: *Journal of Guidance, Control, and Dynamics* (July 2022), pp. 1–15. issn: 0731-5090. doi: [10.2514/1.G006519](https://doi.org/10.2514/1.G006519).

[183] Franco Bernelli Zazzera, Francesco Topputo, and Mauro Massari. *Assessment of Mission Design Including Utilization of Libration Points and Weak Stability Boundaries*. Tech. rep. 18147/04/NL/MV. Dipartimento di Ingegneria Aerospaziale, 2004.

[184] Mischa Kim and Christopher D Hall. “Lyapunov and Halo Orbits About L_2 ”. In: 2001, pp. 349–366. isbn: 0877034885.

[185] F Topputo. “Fast numerical approximation of invariant manifolds in the circular restricted three-body problem”. In: *Communications in Nonlinear Science and Numerical Simulation* 32 (2016), pp. 89–98. issn: 10075704. doi: [10.1016/j.cnsns.2015.08.004](https://doi.org/10.1016/j.cnsns.2015.08.004).

[186] A. Carusi, Lubor Kresák, and G. B. Valsecchi. “Conservation of the Tisserand Parameter at Close Encounters of Interplanetary Objects with Jupiter”. In: *Earth, Moon, and Planets* 68 (1995), pp. 71–94.

[187] Lubor Kresák. “Jacobian integral as a classificational and evolutionary parameter of interplanetary bodies”. In: *Bull. Astron. Inst. Czechosl* 23 (1972), pp. 1–33.

[188] Rita Neves and Joan Pau Sánchez. “Multifidelity design of low-thrust resonant captures for near-earth asteroids”. In: *Journal of Guidance, Control, and Dynamics* 42.2 (2019), pp. 335–346. issn: 07315090. doi: [10.2514/1.G003599](https://doi.org/10.2514/1.G003599).

[189] Charles H. Acton. “Ancillary data services of NASA’s navigation and Ancillary Information Facility”. In: *Planetary and Space Science* 44.1 SPEC. ISS. (1996), pp. 65–70. issn: 00320633. doi: [10.1016/0032-0633\(95\)00107-7](https://doi.org/10.1016/0032-0633(95)00107-7).

[190] Martin Schlueter et al. “MIDACO on MINLP space applications”. In: *Advances in Space Research* 51.7 (2013), pp. 1116–1131. issn: 02731177. doi: [10.1016/j.asr.2012.11.006](https://doi.org/10.1016/j.asr.2012.11.006).

[191] Guus Boender et al. “A stochastic method for global optimization”. In: *Mathematical Programming* 22.1 (Dec. 1982), pp. 125–140. issn: 00255610. doi: [10.1007/BF01581033](https://doi.org/10.1007/BF01581033).

[192] Francesco Biscani and Dario Izzo. “A parallel global multiobjective framework for optimization: pagmo”. In: *Journal of Open Source Software* 5.53 (2020), p. 2338. doi: [10.21105/joss.02338](https://doi.org/10.21105/joss.02338).

[193] Jacob Williams and pbregener. *jacobwilliams/bspline-fortran: 6.0.0*. Nov. 2019. doi: [10.5281/ZENODO.3546617](https://doi.org/10.5281/ZENODO.3546617).

[194] Jacob Williams. *jacobwilliams/Fortran-Astrodynamic Toolkit: 0.1*. Sept. 2020. doi: [10.5281/ZENODO.4047138](https://doi.org/10.5281/ZENODO.4047138).

[195] Martin Schleuter. *Multi-Objective Optimization with MIDACO*. 2009.

[196] Michael R Allshouse. “Detecting and analyzing coherent structures in two-dimensional dynamical systems”. PhD thesis. Massachusetts Institute of Technology, 2013.

[197] Anusmriti Ghosh et al. “Sensitivity and robustness of Lagrangian Coherent Structures in Coastal Water Systems”. In: *Environmental Fluid Mechanics* (2021), pp. 667–691. doi: [10.1007/s10652-021-09792-8](https://doi.org/10.1007/s10652-021-09792-8).

[198] Mohammad Farazmand and George Haller. “Computing Lagrangian coherent structures from their variational theory”. In: *Chaos* 22.1 (2012). issn: 10541500. doi: [10.1063/1.3690153](https://doi.org/10.1063/1.3690153).

[199] Mohammad Farazmand and George Haller. “Attracting and repelling Lagrangian coherent structures from a single computation”. In: *Chaos* 23.2 (2013). issn: 10541500. doi: [10.1063/1.4800210](https://doi.org/10.1063/1.4800210).

[200] Kristjan Onu, Florian Huhn, and George Haller. “LCS Tool: A computational platform for Lagrangian coherent structures”. In: *Journal of Computational Science* 7 (2015), pp. 26–36. issn: 18777503. doi: [10.1016/j.jocs.2014.12.002](https://doi.org/10.1016/j.jocs.2014.12.002).

[201] Doug Lipinski and Kamran Mohseni. “A ridge tracking algorithm and error estimate for efficient computation of Lagrangian coherent structures”. In: *Chaos* 20.1 (2010). issn: 10541500. doi: [10.1063/1.3270049](https://doi.org/10.1063/1.3270049).

[202] Philippe Miron et al. “Anisotropic mesh adaptation on Lagrangian Coherent Structures”. In: *Journal of Computational Physics* 231.19 (2012), pp. 6419–6437. issn: 0021-9991. doi: [10.1016/j.jcp.2012.06.015](https://doi.org/10.1016/j.jcp.2012.06.015).

[203] Christoph Garth et al. “Efficient computation and visualization of coherent structures in fluid flow applications”. In: *IEEE Transactions on Visualization and Computer Graphics* 13.6 (2007), pp. 1464–1471. issn: 10772626. doi: [10.1109/TVCG.2007.70551](https://doi.org/10.1109/TVCG.2007.70551).

[204] Mingpei Lin, Ming Xu, and Xiaoyu Fu. “GPU-accelerated computing for Lagrangian coherent structures of multi-body gravitational regimes”. In: *Astrophysics and Space Science* 362.4 (2017). issn: 1572946X. doi: [10.1007/s10509-017-3050-y](https://doi.org/10.1007/s10509-017-3050-y).

[205] Zong-Fu Luo et al. “Constructing ballistic capture orbits in the real Solar System model”. In: *Celestial Mechanics and Dynamical Astronomy* 120.4 (2014), pp. 433–450. issn: 15729478. doi: [10.1007/s10569-014-9580-5](https://doi.org/10.1007/s10569-014-9580-5).

[206] Zong-Fu Luo and Francesco Topputo. “Analysis of ballistic capture in Sun-planet models”. In: *Advances in Space Research* (2015). issn: 18791948. doi: [10.1016/j.asr.2015.05.042](https://doi.org/10.1016/j.asr.2015.05.042).

[207] Zong-Fu Luo and Francesco Topputo. “Capability of satellite-aided ballistic capture”. In: *Communications in Nonlinear Science and Numerical Simulation* 48 (2017), pp. 211–223. issn: 1007-5704. doi: <https://doi.org/10.1016/j.cnsns.2016.12.021>.

[208] Gianmario Merisio and Francesco Topputo. “Characterization of ballistic capture corridors aiming at autonomous ballistic capture at Mars”. In: Aug. 2021.

[209] Kyoko Makino and Martin Berz. “Remainder Differential Algebras and their Applications”. In: *Computational Differentiation: Techniques, Applications, and Tools* (1996), pp. 63–74.

[210] Kyoko Makino. “Rigorous Analysis of Nonlinear Motion in Particle Accelerators”. PhD thesis. Michigan State University, 1998.

[211] Giuseppe Di Mauro et al. “Nonlinear Control for Proximity Operations Based on Differential Algebra”. In: *Journal of Guidance, Control, and Dynamics* 38.11 (2015), pp. 2173–2187. issn: 0731-5090. doi: [10.2514/1.g000842](https://doi.org/10.2514/1.g000842).

[212] Alexander Wittig et al. “An automatic domain splitting technique to propagate uncertainties in highly nonlinear orbital dynamics”. In: *Advances in the Astronautical Sciences* 152 (2014), pp. 1923–1941. issn: 00653438.

[213] Mauro Massari, Pierluigi Di Lizia, and Mirco Rasotto. “Nonlinear Uncertainty Propagation in Astrodynamics Using Differential Algebra and Graphics Processing Units”. In: *Journal of Aerospace Information Systems* 14.9 (2017), pp. 493–503. doi: [10.2514/1.i010535](https://doi.org/10.2514/1.i010535).

[214] Roberto Armellin et al. “Asteroid close encounters characterization using differential algebra: The case of Apophis”. In: *Celestial Mechanics and Dynamical Astronomy* 107.4 (2010), pp. 451–470. issn: 09232958. doi: [10.1007/s10569-010-9283-5](https://doi.org/10.1007/s10569-010-9283-5).

[215] Martin Berz. “The method of power series tracking for the mathematical description of beam dynamics”. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 258.3 (Aug. 1987), pp. 431–436. issn: 0168-9002. doi: [10.1016/0168-9002\(87\)90927-2](https://doi.org/10.1016/0168-9002(87)90927-2).

[216] Alexander Wittig et al. “Rigorous and accurate enclosure of invariant manifolds on surfaces”. In: *Regular and Chaotic Dynamics* 15.2 (2010), pp. 107–126. issn: 15603547. doi: [10.1134/S1560354710020024](https://doi.org/10.1134/S1560354710020024).

[217] Francesco Cavenago et al. “On-board DA-based state estimation algorithm for spacecraft relative navigation”. In: *7th European conference for aeronautics and space sciences (EUCASS)* (2017), pp. 1–14. doi: [10.13009/EUCASS2017-607](https://doi.org/10.13009/EUCASS2017-607).

[218] Roberto Armellin et al. “Gravity assist space pruning based on differential algebra”. In: *Celestial Mechanics and Dynamical Astronomy* 106.1 (2009), pp. 1–24. issn: 09232958. doi: [10.1007/s10569-009-9235-0](https://doi.org/10.1007/s10569-009-9235-0).

[219] Daniel Pérez-Palau, Josep J. Masdemont, and Gerard Gómez. “Tools to detect structures in dynamical systems using Jet Transport”. In: *Celestial Mechanics and Dynamical Astronomy* 123.3 (2015), pp. 239–262. issn: 15729478. doi: [10.1007/s10569-015-9634-3](https://doi.org/10.1007/s10569-015-9634-3).

[220] Franco Bernelli Zazzera et al. *Trajectory Optimisation under Uncertainties*. Tech. rep. European Space Agency, 2012, p. 88.

[221] Paolo Lunghi, Michèle Lavagna, and Pierluigi Di Lizia. “Atmospheric entry guidance based on Differential Algebra for high elevation Mars landing”. In: *2018 Space Flight Mechanics Meeting*. January. 2018. doi: [10.2514/6.2018-1458](https://doi.org/10.2514/6.2018-1458).

[222] John Guckenheimer and Brian Meloon. “Computing Periodic Orbits and their Bifurcations with Automatic Differentiation”. In: *SIAM Journal on Scientific Computing* 22.3 (2003), pp. 951–985. issn: 1064-8275. doi: [10.1137/s1064827599359278](https://doi.org/10.1137/s1064827599359278).

[223] Pierluigi Di Lizia, Roberto Armellin, and Michele Lavagna. “Application of high order expansions of two-point boundary value problems to astrodynamics”. In: *Celestial Mechanics and Dynamical Astronomy* 102.4 (Dec. 2008), pp. 355–375. issn: 0923-2958. doi: [10.1007/s10569-008-9170-5](https://doi.org/10.1007/s10569-008-9170-5).

[224] Pierluigi Di Lizia et al. “High-order robust guidance of interplanetary trajectories based on differential algebra”. In: *Journal of Aerospace Engineering, Sciences and Applications* 1.1 (2012), pp. 43–57. issn: 2236577X. doi: [10.7446/jaes.a.0101.05](https://doi.org/10.7446/jaes.a.0101.05).

[225] Roberto Armellin et al. “An orbital conjunction algorithm based on Taylor models”. In: *Advances in the Astronautical Sciences*. Vol. 143. 2012, pp. 2127–2144. isbn: 9780877035817.

[226] Francesco Cavenago et al. “DA-based nonlinear filters for spacecraft relative state estimation”. In: *2018 Space Flight Mechanics Meeting*. February. 2018. doi: [10.2514/6.2018-1964](https://doi.org/10.2514/6.2018-1964).

[227] Hao Xie, Jin-Guo Liu, and Lei Wang. “Automatic differentiation of dominant eigen-solver and its applications in quantum physics”. In: *Physical Review B* 101.24 (June 2020). doi: [10.1103/physrevb.101.245139](https://doi.org/10.1103/physrevb.101.245139).

[228] Joachim Kopp. “Efficient numerical diagonalization of hermitian 3x3 matrices”. In: *International Journal of Modern Physics C* 19.03 (2008), pp. 523–548. doi: [10.1142/S0129183108012303](https://doi.org/10.1142/S0129183108012303).

[229] Karunia E. Lestari et al. “Generating roots of cubic polynomials by Cardano’s approach on correspondence analysis”. In: *Heliyon* 6.6 (2020), e03998. issn: 2405-8440. doi: <https://doi.org/10.1016/j.heliyon.2020.e03998>.

[230] Emelie Haynsworth and Alston S. Householder. “The Theory of Matrices in Numerical Analysis.” In: *The American Mathematical Monthly* 73.10 (1966). issn: 00029890. doi: [10.2307/2314680](https://doi.org/10.2307/2314680).

[231] Raymond H. Chan, Yuyang Qiu, and Guojian Yin. “Iterative Methods for Eigenvalues/Eigenvectors”. In: *arXiv preprint arXiv:1105.1185* (2018). doi: [10.1007/978-1-4939-7131-2__148](https://doi.org/10.1007/978-1-4939-7131-2__148).

[232] Mohammad Farazmand, Daniel Blazevski, and George Haller. “Shearless transport barriers in unsteady two-dimensional flows and maps”. In: *Physica D: Nonlinear Phenomena* 278-279 (2014), pp. 44–57. issn: 01672789. doi: [10.1016/j.physd.2014.03.008](https://doi.org/10.1016/j.physd.2014.03.008).

[233] Thomas Devogele et al. “Optimized Discrete Fréchet Distance between Trajectories”. In: *Proceedings of the 6th ACM SIGSPATIAL Workshop on Analytics for Big Geospatial Data*. BigSpatial’17. Redondo Beach, CA, USA: Association for Computing Machinery, 2017, pp. 11–19. isbn: 9781450354943. doi: [10.1145/3150919.3150924](https://doi.org/10.1145/3150919.3150924).

[234] Anne Driemel, Amer Krivosija, and Christian Sohler. “Clustering time series under the Fréchet distance”. In: *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms* 2 (2016), pp. 766–785. doi: [10.1137/1.9781611974331.ch55](https://doi.org/10.1137/1.9781611974331.ch55).

[235] George Haller et al. “Objective barriers to the transport of dynamically active vector fields”. In: *Journal of Fluid Mechanics* 905 (2020), A17. doi: [10.1017/jfm.2020.737](https://doi.org/10.1017/jfm.2020.737).

[236] Nicola Hyeraci and Francesco Topputo. “Method to design ballistic capture in the elliptic restricted three-body problem”. In: *Journal of Guidance, Control, and Dynamics* 33.6 (2010), pp. 1814–1823. issn: 15333884. doi: [10.2514/1.49263](https://doi.org/10.2514/1.49263).

[237] Jean-Michel Muller et al. *Handbook of Floating-Point Arithmetic*. Publication Title: Handbook of Floating-Point Arithmetic. Birkhäuser Boston, 2010. doi: [10.1007/978-0-8176-4705-6](https://doi.org/10.1007/978-0-8176-4705-6).

[238] Institute for Electronics and Electrical Engineers. *754-2019 - IEEE Standard for Floating-Point Arithmetic*. Tech. rep. ISBN: 9781504459242. IEEE, 2019.

[239] Ronald Bruce Lantz. “Quantitative Evaluation of Numerical Diffusion (Truncation Error)”. In: *Society of Petroleum Engineers Journal* 11.03 (Sept. 1971), pp. 315–320. issn: 0197-7520. doi: [10.2118/2811-PA](https://doi.org/10.2118/2811-PA).

[240] John R Martin and Hanspeter Schaub. *Periodic Orbit Discovery Enhanced by Physics-informed Neural Networks*. en. Conference preprint. 2022 Astrodynamics Specialise Conference, Charlotte, North Carolina, Aug. 2022.

[241] Fabienne Jézéquel and Jean-Marie Chesneaux. “CADNA: A Library for Estimating Round-off Error Propagation”. In: *Computer Physics Communications* 178.12 (June 2008), pp. 933–955. issn: 0010-4655. doi: [10.1016/j.cpc.2008.02.003](https://doi.org/10.1016/j.cpc.2008.02.003).

