

**Methods to Explore Families of Quasi-Periodic Orbits with
Applications in Astrodynamics**

by

David A. Lujan

B.S., Applied Mathematics, Missouri University of Science and
Technology, 2017

M.S., Applied Mathematics, University of Colorado Boulder, 2021

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Aerospace Engineering Sciences
2023

Committee Members:

Daniel J. Scheeres, Chair

James Meiss

Stephen Becker

Natasha Bosanac

Jay McMahon

Lujan, David A. (Ph.D., Aerospace Engineering Sciences)

Methods to Explore Families of Quasi-Periodic Orbits with Applications in Astrodynamics

Thesis directed by Prof. Daniel J. Scheeres

Quasi-periodic orbits are of great interest for mission design due to their prevalence in astrodynamics. Incorporating quasi-periodic orbits in the design process expands the design space and have the potential to decrease station-keeping efforts. Quasi-periodic orbits are computationally more complex than periodic orbits, however the benefits of utilizing quasi-periodic orbits and their invariant manifolds can outweigh the computational burden. Methods and tools to handle the larger design space are needed to make the study of families of quasi-periodic orbits tractable to mission designers. In this thesis we leverage single-parameter continuation of n -dimensional quasi-periodic invariant tori to compute quasi-periodic orbits with specific orbit frequencies and with specific orbital characteristics. Additionally, we formulate and solve optimization problems such that the optimization variables are the frequencies of the quasi-periodic orbits. The solution process incorporates a novel parametric constraint, which constrains the direction of travel in frequency space in continuation methods. Moreover, we develop search strategies which successively use parametric constraints to explore families of quasi-periodic orbits. Lastly, we leverage number theoretic properties of quasi-periodic orbits to avoid the effect of resonances in the continuation process. We compute the family of the Earth-Moon L_2 quasi-halo orbits in the circular restricted three-body problem to serve as a test bed and solution space for the work in this thesis.

Dedication

To everyone I have ever met. You have made me who I am for better or for worse.

Acknowledgements

I could not have made it to this point without the support of my friends, family, lab, professors, committee, and advisor. Thank you for everything you've done for me.

My greatest appreciation goes to my friends Jesse Greaves, Damennick Henry, Chloe Long, Jordan Murphy, Omar Nawaz, Kate Schmidt, and Nakul Sridhar for providing years of lifelong memories and friendships. There are many more names to list, and I apologize for leaving anyone out, but these people have been there for me at my worst, have seen me at my best, and significantly influenced who I have become.

I am lucky to be part of the legacy of CSML and am grateful for the opportunities this lab has afforded me and will continue to afford me, and for that I am eternally indebted to my advisor. I would like to thank my advisor Dan Scheeres, committee members James Meiss, Natasha Bosanac, Stephen Becker, and Jay McMahon, and former lab member Nicola Baresi. The discussions with them helped me to progress in my research. Without them, I would not have accomplished as much as I did.

I would like to specifically thank my lab mates Damennick Henry and Luke Peterson for numerous discussions which deepened my understanding, for bouncing ideas off each other, for editing papers, and providing sanity checks. Without all of this my work would have suffered.

Contents

Chapter

1	Introduction	1
1.1	Introduction and Background	1
1.2	Contributions	5
1.3	Organization	7
1.4	Publications	9
2	Preliminaries	11
2.1	Introduction	11
2.2	The Torus	12
2.3	Dynamical Systems Theory	13
2.3.1	Linearization	15
2.3.2	Invariant Sets	17
2.3.3	Equilibrium Points	17
2.3.4	Periodic Orbits	20
2.3.5	Quasi-Periodic Orbits	23
2.3.6	Invariant Stable and Unstable Manifolds	30
2.3.7	Common Torus Terminology	32
2.4	The Circular Restricted Three-Body Problem	35
2.5	Smooth Manifolds	37

2.5.1	Implicit Function Theorem	38
2.5.2	Tangent Space of a Smooth Manifold	39
2.5.3	Retractions	40
2.5.4	Gradients on Smooth Manifolds	41
2.6	Computation of Quasi-Periodic Orbits	42
2.6.1	Continuation	42
2.6.2	Correction Procedure	45
2.7	Optimization Theory	48
2.7.1	Unconstrained Optimization	49
2.7.2	Constrained Optimization	50
3	Study of the Earth-Moon L_2 Quasi-Halo Orbits	52
3.1	Introduction	52
3.2	Computation of n -Dimensional Quasi-Periodic Invariant Tori	56
3.2.1	Quasi-Periodicity	60
3.2.2	Continuity Constraints	63
3.2.3	Phase Constraints	64
3.2.4	Consistency Constraints	65
3.2.5	Parametric Constraints	65
3.2.6	Family Tangents	66
3.2.7	Stability Analysis	68
3.3	Amplitudes of an Invariant Surface	68
3.4	The Halo Orbit Family	71
3.5	The 2-Parameter Family of Quasi-Halo Orbits	74
3.5.1	Exploring the Family	74
3.5.2	Family Surface	81
3.5.3	Characteristics	85

3.5.4	Geometry	88
3.5.5	Period Doubling Halo Orbit Family	93
3.5.6	Stability Analysis	96
3.5.7	Survey of Hyperbolic Manifolds	99
3.5.8	9:2 Resonant Synchronous Quasi-Halo Orbits	107
3.5.9	Conclusion	110
3.6	The Dynamics in the Vicinity of the Stable Halo Orbits	111
3.6.1	Frequencies	113
3.6.2	Amplitudes	114
3.6.3	Jacobi Constant	117
3.6.4	Stability Bifurcation	119
3.6.5	Region of Stability	122
3.6.6	Geometry	125
3.6.7	Relative Motion	130
3.6.8	Error Analysis	132
3.6.9	Conclusion	133
4	Resonance Avoidance Methods	136
4.1	Introduction	136
4.2	Irrational Numbers in Floating Point Arithmetic	138
4.3	Picking Irrational Frequencies	140
4.4	Hopping Beyond Resonances	142
4.5	Results of Implementation	145
4.6	Discussion	150
4.7	Conclusion	152
5	A Retraction on Families of Quasi-Periodic Orbits	154
5.1	Introduction	154

5.2	Solution Development	156
5.2.1	Manifolds	156
5.2.2	Continuation Equations	157
5.2.3	Initial Direction	161
5.2.4	Stopping Criterion	162
5.2.5	Step Size Computation	162
5.3	The Retraction	164
5.4	Validation	166
5.4.1	Targeting Orbit Frequencies in a Family of Quasi-Periodic Orbits	167
5.4.2	Targeting Distance on a Subset of a Family of Quasi-Periodic Orbits	169
5.5	Conclusion	172
6	Optimization Over Families of Quasi-Periodic Orbits	174
6.1	Introduction	174
6.2	Problem Statement	177
6.3	Derivative With Respect to the Parameters	179
6.4	Targeting Characteristics	181
6.4.1	Algorithm	183
6.4.2	Validation	185
6.4.3	Discussion	186
6.5	Unconstrained Optimization	190
6.5.1	Algorithm	191
6.5.2	Validation	193
6.6	Constrained Optimization	195
6.6.1	Algorithm	195
6.6.2	Validation	197
6.7	Conclusion	200

7	Concluding Remarks	202
7.1	Conclusion	202
7.2	Avenues for Future Work	206
	Bibliography	209
	Appendix	
A	Publicly Available Code	220

Tables

Table

3.1	Maximum and minimum values of the Jacobi constant, perilune, and apolune obtained by the quasi-halo family.	85
4.1	Performance of RAv and Plain variants.	148
6.1	Test cases for targeting solution characteristics	185
6.2	Test cases for unconstrained optimization	193
6.3	Test cases for constrained optimization	198

Figures

Figure

1.1	Depiction of a quasi-periodic torus (a) and a quasi-periodic orbit (b).	2
2.1	A 2-dimensional torus with constant θ_0 (red) and θ_1 (blue) circles.	13
2.2	A 2-dimensional quasi-periodic torus with a partial trajectory.	14
2.3	Depiction of the distinct eigenspaces surrounding an equilibrium point.	19
2.4	Eigenvalue pairs of the monodromy matrix in a Hamiltonian system.	23
2.5	Figure of linear invariant curves with different step sizes relative to a periodic orbit (a) and in configuration space (b).	25
2.6	Construction of a 2-dimensional linear invariant torus about a periodic orbit (a) and a complete linear invariant torus (b).	26
2.7	Unstable directions to initialize the unstable manifold on an invariant curve (a) and on a quasi-periodic orbit (b).	33
2.8	Coordinate system and equilibrium points of the circular restricted three-body problem.	36
2.9	Depiction of the tangent space of a manifold at the point \mathbf{z} .	40
2.10	Ambiguity with moving on a manifold. In (a) there is a region of points which are considered close to the approximated point $\tilde{\mathbf{z}}$. In (b) there are multiple curves which can be followed to get to a point on the manifold.	41
2.11	Schematic of a predictor-corrector algorithm.	43
2.12	Depiction of a single shooting arc (a) broken into m multiple-shooting arcs (b).	46

3.1	Partition of a flattened 2-dimensional torus with $N_1 = N_2 = 5$	57
3.2	Example of computing amplitudes of an invariant surface with $n = 2$, $N_1 = 6$, and $N_2 = 4$	70
3.3	A stable halo orbit with the linear center manifolds from its monodromy matrix. . .	71
3.4	The frequencies of the center manifolds of the halo orbit family.	72
3.5	Frequencies of the stable L_2 halo orbit family and the frequencies of their normal center subspaces.	73
3.6	Correspondence between the halo orbit family in frequency space (left) and in coordinate space (right).	74
3.7	Quasi-halo orbits computed using different parametric constraints: constant ω_0 (a), constant ω_1 (b), constant Jacobi energy (c), and constant slope (d).	76
3.8	Search strategies utilizing different parametric constraints: stair step (a), resonance bouncing (b), and roots (c).	79
3.9	Family surface of Earth-Moon L_2 quasi-halo orbits parameterized by the orbit frequencies.	81
3.10	Family surface emanating from CM-A overlaid with resonance lines.	83
3.11	Comparison of L_2 quasi-halo orbit family between our results in (a) to the results of Gómez and Mondelo in [75] in (b).	84
3.12	Jacobi constant across the family surface in non-dimensional units.	86
3.13	Perilune across the family surface in non-dimensional units.	87
3.14	Apolune across the family surface in non-dimensional units.	87
3.15	Three different branches of quasi-halo orbits extending from the same halo orbit. . .	89
3.16	Surface of box sizes of the L_2 quasi-halo orbits.	91
3.17	Example of how to interpret the surface of box sizes in Figure 3.16.	91
3.18	Locations of the minimum and maximum amplitude of the invariant curves across the 2-parameter family of L_2 quasi-halo orbits.	93

3.19	Period-doubling family of periodic orbits that bifurcates from the L_2 halo orbits broken into four segments.	94
3.20	Period-doubling quasi-halo orbit.	95
3.21	Lyapunov exponent across the family surface at a top-down view (top row) and a 3-d view (bottom row).	97
3.22	Lyapunov exponent across the family surface in the stable region of the halo orbits at a top-down view (top row) and a 3-d view (bottom row).	98
3.23	Family surface with stars indicating the sampled orbits used to investigate the behavior of the hyperbolic manifolds along with the regions identified as having distinct behavior.	100
3.24	W_-^u and W_+^u for three different quasi-halo orbits from the three different regions identified to have distinct behavior.	102
3.25	Lunar surface intersection points with W_-^u and W_+^u within 108.56 days for four different quasi-halo orbits.	104
3.26	A quasi-halo orbit overlaid with the locations on the orbit where W_-^u impacts the lunar surface (left) and the parameterization of the orbit with the W_-^u takeoff points (right).	105
3.27	The inclination of the instantaneous Keplerian orbit about the Moon at the point of lunar impact (top) and the Δv needed to land on the lunar surface with no velocity (bottom).	106
3.28	9:2 halo orbit and Moon with lunar intersections (left) with a closeup on the Moon (right).	108
3.29	Family of 9:2 synchronous quasi-halo orbits at different views with the 9:2 halo orbit show in plot (d).	109
3.30	Lunar intersections for a sample of the 9:2 synchronous quasi-halo family.	110
3.31	Frequencies of the four branches with different views.	113
3.32	Three constant ω_2 2-d branches showing the non-uniqueness of the orbit frequencies.	115

3.33	Distribution of amplitudes for the branches of 3-d quasi-halos.	116
3.34	Maximum amplitudes within each of the four computed branch types.	117
3.35	Histogram of location where maximum and minimum amplitude along each 2-d quasi-periodic orbit occurs (a) and a plot of the invariant curves in configuration space with maximum and minimum amplitude for each 2-d quasi-periodic orbit (b).	118
3.36	Plots displaying the Jacobi constant for the 2-d quasi-halo branches.	119
3.37	Plots displaying the Jacobi constant for the 3-d quasi-halo branches.	120
3.38	Stability transition in the 2-d quasi-halo branches.	120
3.39	Eigenvalues near the identified bifurcation points from the visual analysis for each constant ω_2 2-d branch.	121
3.40	Eigenvalues near the identified bifurcation points from the visual analysis for each constant ω_1 2-d branch.	122
3.41	Empirical estimates of the size of the stable region surrounding the stable halo orbits.	123
3.42	Comparison of the 97% line with the computed orbit amplitudes.	124
3.43	Surface (with shadows) made by the last invariant curve of each branch in the constant ω_2 2-d family (a) and in the constant ω_1 2-d family (b).	127
3.44	Examples of 3-dimensional quasi-halo orbits in configuration space.	128
3.45	Invariant surfaces within a constant ω_2 3-d branch (a-c) and within a constant ω_1 3-d branch (d-f).	129
3.46	Survey of invariant surfaces in the 3-dimensional quasi-halo family.	130
3.47	Representative example of the relative motion for the 3-d quasi-halo orbits.	131
3.48	The error over time between six computed invariant surfaces and the propagated surfaces.	134
4.1	Example of picking step size to minimize chance of stepping near a resonance. Plot (b) is a zoomed in copy of plot (a).	142

4.2	Comparing the solutions in frequency space between the resonance avoidance variant (a) to the plain variant (b).	146
4.3	Comparing the branches in frequency space between the resonance avoidance variant (a) to the plain variant (b).	147
4.4	Comparison of performance between the plain variant and the resonance avoidance variant (b). The x -axis is the branch number. The branch numbers are given in plot (a).	149
5.1	The parameter and state manifolds defined by Equations (5.1) and (5.5).	158
5.2	Depiction of the direction constraint in the constrained case. The constrained manifold $\bar{\mathcal{Q}}$ shown against the unconstrained manifold \mathcal{Q} in (a). The projection of $\delta\mathbf{q}_t$ onto the tangent space, resulting in $\delta\mathbf{q}_t^{\parallel}$, and the projection of the continuation path onto $\delta\mathbf{q}_t^{\parallel}$ in (b).	159
5.3	Targeting the frequencies of a quasi-periodic orbit.	169
5.4	Targeting the distance along a constant Jacobi energy line.	171
6.1	Equality and inequality constraints reduce the feasible set of solutions.	179
6.2	Example of two parametric constraints on a 2-dimensional manifold, leading to a single solution \mathbf{z}^*	182
6.3	Results of case 1.	187
6.4	Results of case 2.	187
6.5	Results of case 3.	188
6.6	Results of case 4.	188
6.7	Results of case 5.	189
6.8	Results of case 6.	189
6.9	Results of case 1. Plot (a) is a zoomed out version of plot (b).	194
6.10	Results of case 2. Plot (a) is a zoomed out version of plot (b).	194

6.11 Orbit frequency iterates to minimize the difference in the Jacobi constant with an ω_1 inequality constraint.	198
6.12 Orbit frequency iterates to minimize the difference in the Jacobi constant with an ω_0 equality constraint and ω_1 inequality constraint.	199
6.13 Orbit frequency iterates to minimize the unstable eigenvalue with Jacobi constant and ω_0 inequality constraints.	200

Chapter 1

Introduction

1.1 Introduction and Background

Dynamical systems theory offers an abundance of tools which has accelerated research in astrodynamics enabling the construction of complex mission designs which were previously unimaginable. Dynamical systems theory provides a systematic way of discovering dynamical structures which are known to organize the dynamics of the phase space. In astrodynamics these are, namely, equilibrium points, periodic orbits, quasi-periodic orbits (QPOs), and their stable and unstable manifolds. By leveraging the tools afforded by dynamical systems theory, researchers and mission designers are able to take advantage of these structures to accomplish science objectives and efficiently explore space.

The first mission to utilize dynamical systems theory in the trajectory design process was Genesis which injected into a halo orbit about the Sun-Earth L_1 Lagrange point [119]. The transfer and return trajectory from this orbit utilized the orbit's stable and unstable manifolds. Since that time, a number of missions have utilized dynamical systems theory to design trajectories to and from the vicinity of the libration points. One notable mission is ARTEMIS which first launched as THEMIS and later repurposed for a lunar flyby [40]. ARTEMIS became the first mission to achieve orbit around the Earth-Moon L_1 and L_2 Lagrange points, utilizing a low-energy transfer between them. Looking forward there are many more libration point missions planned [162, 132, 7, 6, 5] such as NASA's flagship Artemis program. The Lunar Gateway is set to be launched to an Earth-Moon near rectilinear halo orbit (NRHO) about the L_2 Lagrange point [167, 169].




Figure 1.1: Depiction of a quasi-periodic torus (a) and a quasi-periodic orbit (b).

Of the dynamical structures utilized by libration point missions, quasi-periodic orbits are quite challenging to compute, more so than periodic orbits. The challenge is due to their inherent property of never repeating themselves. To conceptualize a quasi-periodic orbit consider the 2-dimensional torus in plot (a) in Figure 1.1. A trajectory, in cyan, lies on the surface of this torus. The trajectory evolves in such a way that the trajectory will never repeat itself, covering the surface in infinite time. The torus and the trajectory undergo smooth deformations, consistent with the dynamics a spacecraft experiences, resulting in a quasi-periodic orbit in plot (b). Despite the complicated geometry of the quasi-periodic orbit, it can be easily visualized as motion on a torus. Due to the consistency between the trajectories on the torus and on the quasi-periodic orbit, we call the torus a quasi-periodic invariant torus. Quasi-periodic invariant tori are characterized by having a rationally independent set of frequencies which govern the trajectory on the surface of the torus and give insight into the motion of a spacecraft.

At the expense of a higher computational complexity, quasi-periodic orbits offer numerous benefits over equilibrium points and periodic orbits. To understand the benefits let us first examine a commonly studied multi-body dynamical system in astrodynamics, the circular restricted three-body problem. In the circular restricted three-body problem there are five equilibrium points, and by Lyapunov's center theorem [159, 87, 134], about each equilibrium point there are 1-parameter

families of periodic orbits. Under certain circumstances (mentioned in Chapter 2) there exists a 1-parameter family of 2-dimensional quasi-periodic orbits about a periodic orbit. Sweeping across the family of periodic orbits results in a 2-parameter family of quasi-periodic orbits. Therefore, in the circular restricted three-body problem quasi-periodic orbits are more plentiful than periodic orbits, highlighting the greater abundance of options for mission designers to utilize when considering quasi-periodic orbits in the design process.

The (un)stable manifold extending from an equilibrium point is a 1-dimensional curve, whereas the (un)stable manifold extending from a periodic orbit forms a 2-dimensional surface in phase space. Likewise, the (un)stable manifold of an n -dimensional quasi-periodic orbit forms a $n + 1$ -dimensional surface. Considering quasi-periodic orbits form higher-parameter families than periodic orbits then there are vastly more low-energy trajectories which more densely fill phase space, providing versatility to the mission designer in designing complex missions.

Moreover, periodic orbits generically become quasi-periodic when periodic or quasi-periodic perturbations are added to the system's dynamics [107]. Farquhar notes this behavior in [65] in which he studies the use of halo orbits for Lunar operations. He also adds that including these higher order effects presents a more dynamically accurate solution which lowers station-keeping costs by saying, "The accurate nominal trajectories are actually quasi-periodic and can be obtained by finding particular solutions to equations of motions that include the effects of nonlinearities, lunar eccentricity, and the sun's gravitational field. If a satellite is forced to follow an inaccurate nominal path, the stationkeeping fuel expenditure can become quite high."

It is evident that to leverage the dynamics of a system to design low-fuel missions it is critical to understand quasi-periodic orbits and their manifolds. This motivates the creation of tools to aid mission designers and motivates this thesis.

The numerical computation of quasi-periodic orbits is a fairly recent development and made great strides in the early 2000s with continued development now [71, 83]. We provide a brief background on the methods to compute quasi-periodic orbits, however the reader is referred to the recent book written by Àlex Haro et al. for an in depth treatment on the history of computing quasi-

periodic orbits and the different methods available [84]. Early methods relied on semi-analytical methods such as Lindstedt-Poincaré series expansions and reduction to the center manifold [63, 74, 103]. The benefit of these methods is that they can compute orbits with very high accuracy, however they suffer from having small radii of convergence as more terms are incorporated into the solution and thus struggle to compute orbits with rather large amplitudes. Moreover, the formulation of the solutions have to be modified any time the type of orbit changes or the dynamical system changes.

From the year 2000 and onward the developed methods primarily focus on fully numerical implementations. A popular collection of methods, known as parameterization methods, has the advantage of being fast, backed by rich theory, and simultaneously computing the stability information for reducible tori [44, 45, 57, 82, 84, 83, 104, 115, 46, 172]. However, the downside of parameterization methods is they have to be curated to the type of dynamical system. Moreover, other available methods have a simpler implementation than the parameterization methods, making the parameterization method less favorable from a practitioner’s point of view.

Another popular method is a flow map method first conceived by Castellà and Jorba [101, 47] and independently developed by Gómez and Modelo [75], and later improved upon by Olikara and Scheeres [140]. This method has been successfully used in a variety of contexts [22, 19, 139, 135, 121, 131, 20, 16, 166]. It has the advantage of being simple to implement, applicable to a range of dynamical systems, is indifferent to the stability type of the computed orbits, and computes the stability of the orbits as a byproduct. The downside is that the method is $\mathcal{O}(n^3)$ in the number of computations and $\mathcal{O}(n^2)$ in the required storage space. In comparison, the parameterization methods are $\mathcal{O}(n \log n)$ in the number of computations and $\mathcal{O}(n)$ in the storage requirement. Despite this drawback we opt to use this flow map method to compute quasi-periodic orbits here. However, the tools and methods developed in this thesis are amenable to the various quasi-periodic orbit computational schemes.

With the advent of powerful modern personal computers the barrier to studying quasi-periodic orbits has greatly diminished. This is likely a reason, alongside the development of efficient computational methods, that quasi-periodic orbits are gaining popularity in the astrodynamics re-

search community. The utilization of quasi-periodic orbits and their stable and unstable invariant manifolds in space missions will become more widespread as more research on the application of quasi-periodic orbits is published. This is a motivation for this thesis which aims to develop tools to explore families of quasi-periodic orbits. Many of the tools used and developed here can be used in a variety of dynamical systems, but this thesis will exclusively focus on astrodynamics, and in particular the circular restricted three-body problem. The presentation of the material here is kept general where possible and will be drawn to the specific case of the circular restricted three-body problem.

1.2 Contributions

Thesis Statement: Quasi-periodic orbits generally form multi-parameter families of orbits in astrodynamics, expanding the design space, presenting challenges for mission designers. Leveraging the natural parameterization of quasi-periodic orbits, their orbit frequencies, the design space becomes simpler to explore and find specific solutions.

The goals of this thesis is to make simple the inherently complex nature of quasi-periodic orbits in astrodynamics, and to present practical and efficient tools for their computation in mission design, thus lowering the barrier for their use in space missions. We accomplish these goals by presenting an existing quasi-periodic orbit computational algorithm in generality in great detail (Chapter 3), studying a 2- and 3-parameter family of quasi-periodic orbits (Chapter 3), improving the continuation of families of quasi-periodic orbits (Chapters 3 & 4), and developing tools to compute specific quasi-periodic orbits within single- and multi-parameter families (Chapters 5 & 6). A detailed description of the contributions are given below.

Detailed description of an n -dimensional invariant torus algorithm: Much research has used and described the quasi-periodic orbit computational method, commonly referred to as GMOS, however it has only been used to compute diffeomorphisms of 2- and 3-dimensional quasi-periodic invariant tori, and the body of literature usually leaves out critical information necessary for the numerical implementation of the algorithm. As research in astrodynamics progresses the interest

in studying more complex dynamical systems will increase, and thus the interest in computing invariant tori of dimension greater than 3 will increase. A literature review has revealed that the generalization of this algorithm has not been presented in such detail before. Therefore, in this thesis a very detailed generalization of the algorithm applicable to n -dimensional tori will be given allowing researchers to more easily implement it.

In-depth studies of relevant families of quasi-periodic orbits: NASA’s Artemis program has sparked interest in the Earth-Moon L_2 NRHO family. This 1-parameter family is surrounded by a 2-parameter family of quasi-halo orbits, and in the vicinity is also a 3-parameter family of stable quasi-halo orbits. Therefore, a study of the 2- and 3- parameter families of the Earth-Moon L_2 quasi-halo orbits is conducted to better understand the space in which the Lunar Gateway will be staged in. This study examines a novel parametric constraint to fix the continuation direction in frequency space, examines a previously unstudied portion of the 2-parameter family of quasi-halo orbits along with a sub-family of the 3-parameter family of quasi-halo orbits, computes several orbital characteristics, makes observations about the family, and speaks to applications to mission design. This work presents a paradigm in which to study and visually represent data for 2-parameter families efficiently.

Improved performance of the continuation of quasi-periodic invariant tori: One major challenge in the continuation of families of quasi-periodic orbits is the presence of resonances. Resonances inhibit the ability of algorithms to compute quasi-periodic invariant tori, thereby halting the continuation procedure. A method to determine the irrationality of floating point numbers is leveraged to construct two novel methods to avoid the issues of resonances. The first method uses a line search to pick a step size which predicts a torus with the “most” irrational set of frequencies, promoting better-conditioned computations which lead to faster matrix inversions and less Newton iterations for convergence. The second method determines when the continuation is near a resonance and decides when to increase the step size to jump beyond the resonance. Together these two methods improve the continuation of families of quasi-periodic orbits by mitigating the effects of resonances.

Enhanced mission design capabilities applied to families of orbits: An algorithm is developed which adapts standard continuation methods to be able to precisely compute specific orbits based on a target set of orbit frequencies. This has multiple applications and a couple of them are explored here. Both applications aim to enhance the capability of mission designers to find orbits from within families of orbits which satisfy constraints which can be derived from mission constraints. The methods used are not restricted to quasi-periodic orbits and astrodynamics; they are general enough for a class of problems satisfying the conditions of the implicit function theorem. The first application is a modified Newton's method to compute an orbit from a family of quasi-periodic orbits with desired orbital characteristics. The second application imports techniques from optimization to optimize over a family of quasi-periodic orbits with equality and inequality constraints.

1.3 Organization

This section summarizes the remaining chapters of the thesis.

Chapter 2 introduces the notation which is used throughout this thesis and outlines the background material which is fundamental to the research presented here. Section 2.2 introduces the notion of a torus, and defines different types of tori. Section 2.3 gives a discourse on dynamical systems theory which provides the framework in which quasi-periodic invariant tori are computed. Section 2.4 presents the dynamical system in which orbits are computed. Section 2.5 goes over the implicit function theorem and the manifolds which define families of orbits. Section 2.6 presents the ideas of continuation methods and shooting methods which are vital to the computation of families of quasi-periodic orbits. Lastly, Section 2.7 provides the elementary principles of optimization and known algorithms for constrained and unconstrained optimization.

Chapter 3 provides an in depth study of the Earth-Moon L_2 quasi-halo orbits in the circular restricted three-body problem tailored toward a mission design perspective. Section 3.2 generalizes an existing quasi-periodic orbit computational scheme to compute n -dimensional quasi-periodic invariant tori. Section 3.3 describes a method to compute the amplitudes of an invariant surface. Section 3.4 overviews the halo orbit family and points out important aspects for the remaining

sections. Section 3.5 presents findings for the 2-parameter family of quasi-halo orbits, while Section 3.6 presents findings on the 2- and 3-parameter family of quasi-halo orbits in the vicinity of the stable halo orbits.

Chapter 4 defines two methods to avoid the difficulties of computing quasi-periodic orbits near resonances. Section 4.2 discusses the problem of identifying a resonance between orbit frequencies. Section 4.3 presents a line search method for picking the most irrational set of orbit frequencies within a range of continuation step sizes. Section 4.4 presents a heuristic method for determining when a continuation scheme should increase the step size to jump beyond a resonance. Section 4.5 analyzes the performance of the methods on a test example. Section 4.6 discusses the impact of resonances on the computations of quasi-periodic orbits and provide advice for dealing with resonances in computing branches of quasi-periodic orbits. Lastly, Section 4.7 concludes the chapter.

Chapter 5 presents a novel algorithm for computing quasi-periodic orbits within single- and multiple- parameter families. This algorithm enables the precise computation of orbits within a family and is equivalent to a retraction function. Section 5.2 goes over the development of the retraction. Section 5.3 presents the algorithm for the retraction. Section 5.4 validates the retraction with two examples. The first example computes quasi-halo orbits from a 2-parameter family which have specific sets of orbit frequencies. The second example computes an orbit which is a specific distance from the initial orbit along a 1-dimensional submanifold of constant energy. Section 5.5 ends the chapter with a discussion of the advantages of the retraction and presents various uses for it.

Chapter 6 presents applications for the use of the algorithm developed in Chapter 5. The examples highlight the power of the retraction function and its use for developing tools to enhance mission design in multi-parameter families of quasi-periodic orbits. Section 6.2 defines the optimization problem. Section 6.4 details a modified Newton's method devised for computing solutions with specified characteristics. Several examples in the mindset of mission design are presented. Section 6.5 details a modified gradient descent algorithm which recasts an equality constrained problem into an unconstrained optimization problem. Examples of optimization problems are presented. Section

6.6 uses the modified gradient descent algorithm in an augmented Lagrangian method (ALM) to perform constrained optimization with inequality constraints. And again, examples of optimization problems are presented. Lastly, Section 6.7 ends the chapter with concluding remarks.

Lastly, Chapter 7 concludes this thesis by summarizing the main contributions, presenting ongoing challenges, and defining avenues for future work.

1.4 Publications

The following list of publications resulted from the work on this thesis.

Journal Publications

- A) Lujan, D., and Scheeres, D. J., “Earth-moon L_2 quasi-halo orbit family: characteristics and manifold applications”, *Journal of Guidance, Control, and Dynamics*, Vol. 45, No. 11, pp. 2029-2045, 2022;
- B) Lujan, D., and Scheeres, D. J., “Dynamics in the vicinity of the stable halo orbits”, *Journal of the Astronautical Sciences*, Vol. 70, No. 20, 2023;
- C) Lujan, D., and Scheeres, D. J., “Methods to target quasi-periodic orbits”, [*in preparation*];
- D) Lujan, D., and Scheeres, D. J., “Optimization over families of quasi-periodic orbits”, [*in preparation*];

Conference Publications

- A) Lujan, D., and Scheeres, D. J., “Frequency structure of the nrho family in the earth-moon system”, 2019 AAS/AIAA Astrodynamics Specialist Conference, August 11-15, 2019, Portland, Maine;
- B) Lujan, D., and Scheeres, D. J., “Global L_2 quasi-halo family and their characteristics”, 31st AAS/AIAA Spaceflight Mechanics Meeting, February 1-3, 2021, Virtual;

- C) Lujan, D., and Scheeres, D. J., “The earth-moon L_2 quasi-halo orbit family: characteristics and manifold applications”, 32nd AIAA/AAS Space Flight Mechanics Meeting, January 3-7, 2022, San Diego, California;
- D) Lujan, D., and Scheeres, D. J., “Dynamics in the Vicinity of the Stable Halo Orbits”, 73rd International Astronautical Congress, September 18-22, 2022, Paris, France;
- E) Lujan, D., and Scheeres, D. J., “Method to target quasi-periodic orbit frequencies within multi-parameter families”, 33rd AAS/AIAA Space Flight Mechanics Meeting, January 15-19, 2023, Austin, Texas;

Chapter 2

Preliminaries

2.1 Introduction

This chapter is devoted to defining the fundamental concepts and notation necessary to understand the work in this thesis. The hope is that the reader can read through this material with limited background knowledge and be able to understand the research methods used here, and walk away with an understanding of the numerical computation of quasi-periodic invariant tori. At the roots of this subject are dynamical systems theory, Kolmogorov-Arnold-Moser (KAM) theory, manifold theory, and numerical analysis. Dynamical systems theory and manifold theory will be treated in fair detail, while KAM theory will only be introduced at a high level, and no attention is brought to numerical analysis. Those familiar with the presented material will only need to skim through to learn the notation.

This chapter is organized as follows: Section 2.2 introduces the notion of a torus, and defines different types of tori. Section 2.3 gives a discourse on dynamical systems theory which provides the framework in which quasi-periodic invariant tori are computed. Section 2.4 presents the dynamical system in which orbits are computed. Section 2.5 goes over the implicit function theorem and the manifolds which define families of orbits. Section 2.6 presents the ideas of continuation methods and shooting methods which are vital to the computation of families of quasi-periodic orbits. Lastly, Section 2.7 provides the elementary principles of optimization and known algorithms for constrained and unconstrained optimization.

The information presented in this chapter is covered in much more detail in various textbooks.

The reader is strongly encouraged to read [133] for a treatment on the theory of dynamical systems theory applied to general dynamical systems. This book is where much of Section 2.3 comes from. See [154] for a treatment on dynamics, especially Hamiltonian dynamics, and techniques which are vital to the aspiring astrodynamist. See [158] to dive deeper into the numerical tools which are used to numerically study dynamical systems. The ideas in this book form the basis for computing families of orbits in dynamical systems. See [111] for a rigorous treatment on invariant tori in dynamical systems. See [39] for families of quasi-periodic invariant tori in dynamical systems with a focus on KAM theory in different settings. See [116] for a comprehensive background on manifold theory and differential geometry. A complementary text is [33] which presents optimization on smooth manifolds. Additionally, the reader is referred to the works of [25, 35, 120, 138] for background on classical optimization theory, including constrained and unconstrained optimization, linear and nonlinear programming, and convex optimization. Lastly, see [56, 12] for a theoretical treatment on KAM theory.

2.2 The Torus

At the heart of a quasi-periodic invariant torus is the topological object known as a torus. It is the Cartesian product between circles. A circle is defined as the space $\mathbb{S}^1 = \{\theta \mid 0 \leq \theta \leq 2\pi\}$. Then, an n -dimensional torus is defined as $\mathbb{T}^n = \underbrace{\mathbb{S}^1 \times \dots \times \mathbb{S}^1}_n$, which is called an n -torus for short. An n -torus is parameterized over its surface with n independent coordinates $\boldsymbol{\theta} = [\theta_0, \dots, \theta_{n-1}]$ with $0 \leq \theta_i \leq 2\pi$ for $i = 0, \dots, n-1$. See Figure 2.1 for an example of a 2-dimensional torus with constant θ_0 and θ_1 circles.

With the definition of a torus given we may proceed to define a quasi-periodic torus. A quasi-periodic torus is a special case of a torus in which the parameters $\boldsymbol{\theta}$ are governed by the set of dynamics

$$\dot{\boldsymbol{\theta}} = \boldsymbol{\omega} \tag{2.1a}$$

$$\dot{\boldsymbol{\omega}} = \mathbf{0} \tag{2.1b}$$




Figure 2.1: A 2-dimensional torus with constant θ_0 (red) and θ_1 (blue) circles.

wherein each θ_i evolves at a constant rate in time. The rates are called the torus frequencies and are denoted by the vector $\boldsymbol{\omega}$. This vector is also called the frequency vector. Furthermore, the non-resonance condition

$$\langle \boldsymbol{\omega}, \mathbf{k} \rangle \neq 0, \quad \forall \mathbf{k} \in \mathbb{Z} \setminus \{0\} \quad (2.2)$$

must be met by the frequencies to ensure the trajectory following Equation (2.1) densely covers the surface of the torus. When the non-resonance condition is met we call the torus an n -dimensional quasi-periodic torus (Figure 2.2) and call the frequencies incommensurate. In the case the non-resonance condition is not met then the dimension of the torus decreases to the number of rationally independent frequencies.

More machinery is required to fully define a quasi-periodic invariant torus. This machinery lies within dynamical systems theory and will be defined in Section 2.3.2.

2.3 Dynamical Systems Theory

A dynamical system is a set of rules which govern the time evolution of points in a phase space. In this thesis we only consider the phase space to be \mathbb{R}^d even though phase space can be on a general manifold \mathcal{M} . Let $\mathbf{x} \in \mathbb{R}^d$ be a point in phase space, then a general autonomous dynamical system takes the form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}), \quad (2.3)$$

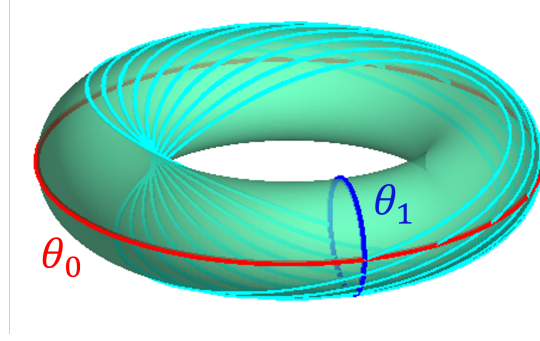


Figure 2.2: A 2-dimensional quasi-periodic torus with a partial trajectory.

where $t \in \mathbb{R}$ is time and $\mathbf{f} : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called the vector field. The solution flow (or simply flow) to Equation (2.3) is the differentiable mapping $\varphi : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ such that

$$(1) \quad \varphi_{t_0}(\mathbf{x}) = \mathbf{x}$$

$$(2) \quad \varphi_t(\varphi_s(\mathbf{x})) = \varphi_{t+s}(\mathbf{x}), \quad \forall t, s \in \mathbb{R},$$

where $\mathbf{x}(t) = \varphi_t(\mathbf{x}_0)$ gives the state of a particle at time t starting from the initial condition \mathbf{x}_0 at the initial time t_0 . Moreover, the vector field is recovered from

$$\mathbf{f}(\mathbf{x}) = \frac{d\varphi_t(\mathbf{x})}{dt}. \quad (2.4)$$

When the word “orbit” or “trajectory” is used we mean

$$\Gamma_x = \{\varphi_t(x) \mid t \in \mathbb{R}\}, \quad (2.5)$$

the collection of states in forward and backward time starting with \mathbf{x} at t_0 . When $\mathbf{x}(t)$ is referenced as the orbit it is implied to mean the orbit induced by $\mathbf{x}(t)$ in forward and backward time.

The dynamical systems of interest in this thesis are Hamiltonian systems. Before defining a Hamiltonian system, we give the definition of the broader category of dynamical systems called conservative systems. A conservative system is a class of dynamical systems in which there is a conserved quantity along the solution flow, meaning the following equation holds:

$$\nabla I \cdot \mathbf{f} = 0, \quad (2.6)$$

where $I = I(\mathbf{x})$ is constant along $\mathbf{x}(t)$.

A Hamiltonian system typically describes the motion of a particle under natural forces. Systems of this type have coordinates $\mathbf{q}, \mathbf{p} \in \mathbb{R}^d$ which define the position and momentum of a particle in phase space, respectively. In this case, phase space is always an even dimension equal to $2n$ and $\mathbf{x} = [\mathbf{q}^T \mathbf{p}^T]^T$. For a Hamiltonian system the conserved quantity is the Hamiltonian function

$$H(\mathbf{q}, \mathbf{p}, t) = \mathbf{p} \cdot \dot{\mathbf{q}} - \frac{1}{2}m\dot{\mathbf{q}} \cdot \dot{\mathbf{q}} + U(\mathbf{q}, t), \quad (2.7)$$

where U is the potential function of the forces acting on the particle. The dynamical system is generated from taking the partial derivative of the Hamiltonian with respect to \mathbf{q} and \mathbf{p} , resulting in

$$\dot{\mathbf{x}} = \begin{bmatrix} \mathbf{0}_n & \mathbb{I}_n \\ -\mathbb{I}_n & \mathbf{0}_n \end{bmatrix} \begin{bmatrix} \frac{\partial H}{\partial \mathbf{q}} \\ \frac{\partial H}{\partial \mathbf{p}} \end{bmatrix}, \quad (2.8)$$

where $\mathbf{0}_n$ is an $n \times n$ matrix of zeros and \mathbb{I}_n is the $n \times n$ identity matrix.

While the Hamiltonian formalism is not used in this thesis it is important to note when the term ‘‘Hamiltonian’’ is used to describe a system, force, or perturbation the implication is the object in question can be tied back to a Hamiltonian formalism and all the assumptions and results pertaining to Hamiltonian systems apply. Notably, the state transition matrix (STM) in a Hamiltonian system is a symplectic matrix. One important property of a symplectic matrix is the determinant always equals 1, meaning the matrix is always invertible. Another implication is its eigenvalues come in reciprocal pairs. Let λ be an eigenvalue, then $1/\lambda$ is also an eigenvalue. See [154] for more properties of symplectic matrices. The symplectic nature of the STM has important implications on the linear behavior surrounding a trajectory.

2.3.1 Linearization

A fundamental tool in the study of dynamical systems is the linearization of solutions, which provides insight into the local dynamics, and allows for the detection of nearby dynamical objects. Let $\mathbf{x}(t)$ be a solution to Equation (2.3). Suppose one is interested in how a small deviation $\delta\mathbf{x}_0 \in \mathbb{R}^d$

evolves in time then the linear dynamical system describing the motion of $\delta\mathbf{x}(t)$ is derived from the following:

$$\begin{aligned}\frac{d(\mathbf{x} + \delta\mathbf{x})}{dt} &= \mathbf{f}(\mathbf{x} + \delta\mathbf{x}) \\ \frac{d\mathbf{x}}{dt} + \frac{d\delta\mathbf{x}}{dt} &= \mathbf{f}(\mathbf{x} + \delta\mathbf{x}) \\ \mathbf{f}(\mathbf{x}) + \frac{d\delta\mathbf{x}}{dt} &= \mathbf{f}(\mathbf{x} + \delta\mathbf{x}).\end{aligned}\tag{2.9}$$

A Taylor series expansion of the right-hand side of Equation (2.9) reveals

$$\mathbf{f}(\mathbf{x} + \delta\mathbf{x}) = \mathbf{f}(\mathbf{x}) + \left. \frac{d\mathbf{f}(\mathbf{x})}{d\mathbf{x}} \right|_{\mathbf{x}(t)} \delta\mathbf{x} + \mathcal{O}(\|\delta\mathbf{x}\|^2).\tag{2.10}$$

Substituting Equation (2.10) into Equation (2.9), canceling like terms, and dropping higher-order terms leads to

$$\frac{d\delta\mathbf{x}}{dt} = \left. \frac{d\mathbf{f}(\mathbf{x})}{d\mathbf{x}} \right|_{\mathbf{x}(t)} \delta\mathbf{x}.\tag{2.11}$$

Letting $A(t) = \left. \frac{d\mathbf{f}(\mathbf{x}, t)}{d\mathbf{x}} \right|_{\mathbf{x}(t)}$, Equation (2.11) can be rewritten in the more familiar form

$$\delta\dot{\mathbf{x}} = A(t)\delta\mathbf{x},\tag{2.12}$$

where $A(t)$ is the Jacobian matrix of $\mathbf{f}(\mathbf{x})$ evaluated at $\mathbf{x}(t)$.

The general solution to Equation (2.12) is

$$\delta\mathbf{x}(t) = \Phi(t, t_0)\delta\mathbf{x}_0,\tag{2.13}$$

where $\Phi(t, t_0) : \mathbb{R}^d \rightarrow \mathbb{R}^d$ is called the fundamental matrix solution or the state transition matrix.

The STM is the solution to the following linear differential equation

$$\dot{\Phi} = A(t)\Phi, \quad \Phi(t_0, t_0) = \mathbb{I},\tag{2.14}$$

where $A(t)$ is as previously defined. The use of Φ in dynamical systems theory to help organize the phase space will be made clear in the next section when discussing the common invariant sets in astrodynamics.

2.3.2 Invariant Sets

The ability of dynamical systems theory to organize phase space by detecting and computing invariant sets is what makes dynamical systems theory so attractive to those studying dynamical systems. An invariant set is a set $\Lambda \subset \mathbb{R}^d$ such that $\varphi_t(\Lambda) = \Lambda$ for all t . Dynamical system theory provides systematic approaches to determining the invariant sets in a dynamical system.

We now have the machinery to introduce the notion of an invariant torus. An n -torus is said to be invariant of the vector field \mathbf{f} if there exists an invertible and differentiable function¹ $\boldsymbol{\tau} : \mathbb{T}^n \rightarrow \mathbb{R}^d$ such that the restriction of the vector field to $\mathcal{T} = \{\boldsymbol{\tau}(\boldsymbol{\theta}) \mid \boldsymbol{\theta} \in \mathbb{T}^n\}$ is mapped by the inverse of the differential map $D\boldsymbol{\tau}^{-1}$ into the constant vector field on the torus in Equation (2.1). The orbit in \mathbb{R}^d given by $\boldsymbol{\tau}$ is an invariant set under the flow φ , and the orbit in \mathbb{T}^n densely covers the surface of the torus with frequencies $\boldsymbol{\omega}$ satisfying Equations (2.1) and (2.2). Each point $\boldsymbol{\tau}(\boldsymbol{\theta}) = \mathbf{x}$ generates an orbit, and each orbit belongs to the same invariant set. The function $\boldsymbol{\tau}$ is commonly referred to as a torus function.

The value of n determines the type of invariant set the invariant torus is diffeomorphic to. For $n = 0$ the invariant set is an equilibrium point, for $n = 1$ the invariant set is a periodic orbit, and for $n \geq 2$ the invariant set is a quasi-periodic orbit. In the case of $n \geq 2$ we call the invariant torus a quasi-periodic invariant torus. Sections 2.3.3 to 2.3.6 are devoted to defining and describing the common invariant sets in astrodynamics, namely, equilibrium points, periodic orbits, quasi-periodic orbits, and the invariant stable and unstable manifolds emanating from these types of orbits.

2.3.3 Equilibrium Points

The simplest invariant set is the equilibrium point. An equilibrium point of Equation (2.3) is a point \mathbf{x}^* which satisfies

$$\mathbf{f}(\mathbf{x}^*) = \mathbf{0}. \quad (2.15)$$

¹ An invertible and differentiable function is called a diffeomorphism.

The invariant set defined by an equilibrium point is the singleton $\Lambda = \{\mathbf{x}^*\}$. An object at this point in phase space will never change its state unless perturbed by a force not considered in the dynamical model.

Studying the linearization of the dynamics about \mathbf{x}^* reveals information about the types of motion in the vicinity of \mathbf{x}^* . The resulting linear differential equation

$$\delta\dot{\mathbf{x}} = A\delta\mathbf{x} \quad (2.16)$$

is a linear time invariant system since the Jacobian matrix, or dynamics matrix, $A = \left. \frac{d\mathbf{f}(\mathbf{x})}{d\mathbf{x}} \right|_{\mathbf{x}^*}$ is constant. The space surrounding \mathbf{x}^* in the linear regime is composed of three distinct subspaces, also called eigenspaces, the stable subspace, E^s , the unstable subspace, E^u , and the center subspace, E^c . The direct product of these three subspaces compose \mathbb{R}^d with $\dim E^s + \dim E^u + \dim E^c = n$. The local dynamics is organized according to these three types of subspaces and the eigenvalues of A tell which subspaces are present around \mathbf{x}^* . Suppose $(\lambda, \boldsymbol{\psi})$ is an eigenpair of A , then the three subspaces are defined as follows:

$$E^s = \text{span}\{\text{Re}(\boldsymbol{\psi}), \text{Im}(\boldsymbol{\psi}) \mid \text{Re}(\lambda) < 0\}, \quad (2.17)$$

$$E^u = \text{span}\{\text{Re}(\boldsymbol{\psi}), \text{Im}(\boldsymbol{\psi}) \mid \text{Re}(\lambda) > 0\}, \quad (2.18)$$

$$E^c = \text{span}\{\text{Im}(\boldsymbol{\psi}) \mid \text{Re}(\lambda) = 0\}. \quad (2.19)$$

For any $\delta\mathbf{x}_0 \in E^s$ then $\lim_{t \rightarrow \infty} \|\delta\mathbf{x}(t)\| \rightarrow 0$, showing asymptotic stability. Conversely, for any $\delta\mathbf{x}_0 \in E^u$ then $\lim_{t \rightarrow \infty} \|\delta\mathbf{x}(t)\| \rightarrow \infty$, showing instability. For deviations which are in the center eigenspace the solution $\delta\mathbf{x}(t)$ remains bounded for all time if A is semisimple. An equilibrium point is deemed stable if *all* the eigenvalues are without positive real parts and A is semisimple. Otherwise, the equilibrium point is deemed unstable. The classification of stability becomes more involved when A is not semisimple. See Section 2.7 in [133] for more information on the classifications of linear stability. The symplectic nature of the STM has ramifications on the eigenvalues of the Jacobian matrix. The existence of an unstable eigenvalue λ_u implies the existence of a stable eigenvalue $\lambda_s = -\lambda_u$, enforcing the stable and unstable eigenspaces to appear

together. When a stable and unstable eigenvalue are paired together in the linearization the linear behavior is called a saddle. Moreover, the existence of a complex conjugate pair of eigenvalues λ and $\bar{\lambda}$ implies the existence of two additional complex conjugate pairs $-\lambda$ and $-\bar{\lambda}$. In the event that $\text{Re}(\lambda) = 0$ then the additional conjugate pair is the same as the original pair and do not appear in the spectrum of A . When $\text{Re}(\lambda) = 0$ the linear behavior is called a center. A depiction of the saddle and center behavior is in Figure 2.3. The black arrows depict the direction of motion a particle would experience in each subspace. An equilibrium point that is described with linear behavior *center* \times *center* \times *saddle* implies there are two independent center modes and one saddle mode of motion.




Figure 2.3: Depiction of the distinct eigenspaces surrounding an equilibrium point.

It is important to note that the eigenspaces are only invariant of the linear dynamics in Equation (2.16) and not invariant of the original nonlinear system. However, the vectors which span each subspace lie tangent to the nonlinear invariant manifolds ([133]) and provide good initial estimates of the nonlinear invariant manifolds in a small region of \mathbf{x}^* , after which the nonlinear dynamics warp these directions. The process to generate the nonlinear invariant manifolds from the linearization of other invariant objects will be discussed throughout the remainder of Section 2.3, however it is worth mentioning now the types of invariant structures present around an equilibrium point based on the eigenvalues of A when \mathbf{f} is Hamiltonian.

If there exists a λ such that $\text{Re}(\lambda) < 0$ or $\text{Re}(\lambda) > 0$ then both a stable and unstable

manifold exist. The combination of these two manifolds results in a hyperbolic manifold. If there exists a purely imaginary pair of eigenvalues then the invariant sets are periodic orbits as given by Lyapunov's center theorem [87, 134, 159]. This theorem states:

Let \mathbf{x}^ be an equilibrium point of $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t)$. Suppose $\lambda = i\omega$ is an eigenvalue of the Jacobian matrix. Then there exists a continuous one-parameter family of periodic orbits about \mathbf{x}^* . Moreover, the period of the periodic orbits approaches $2\pi/\omega$ as the periodic orbits approach \mathbf{x}^* .*

There is another condition on the remaining eigenvalues λ_i which says that ω/λ_i can not be an integer. Lyapunov's center theorem enables the detection of families of periodic orbits about equilibrium points.

2.3.4 Periodic Orbits

A point \mathbf{x} is said to be on a periodic orbit if there exists a $T > 0$ such that

$$\varphi_T(\mathbf{x}) = \mathbf{x}. \quad (2.20)$$

The variable T is called the period of the periodic orbit and is the minimum T necessary for Equation (2.20) to be satisfied, implying the orbit $\Gamma_{\mathbf{x}}$ consists of a closed curve in phase space. A point on a periodic orbit traces out this curve and repeats its state every T time units. A change of variables can be made so that the orbit is parameterized by an angular coordinate $\theta_0 \in [0, 2\pi]$ such that $\mathbf{x} = \mathbf{x}(\theta_0) = \mathbf{x}(\theta_0 + 2\pi)$. The torus function $\boldsymbol{\tau}(\theta_0)$ is given by $\boldsymbol{\tau}(\theta_0) = \mathbf{x}(T\frac{\theta_0}{2\pi})$.

Suppose \mathbf{x}^* is an equilibrium point which satisfies the conditions of Lyapunov's center theorem, then we know a one-parameter family of periodic orbits exists in the vicinity of \mathbf{x}^* . The theorem does not state how large this family of periodic orbits is or how to find the family. In general, one does not know how large a family of periodic orbits will be, however the linearization of $\mathbf{f}(\mathbf{x})$ about \mathbf{x}^* provides information on how to generate a guess of a nearby periodic orbit. The eigenvector $\boldsymbol{\psi}$ associated with the imaginary eigenvalue $i\omega$ of the dynamics matrix provides insight into the prediction of a point \mathbf{x} on a periodic orbit while $i\omega$ provides insight into the value of the period. Recall the linearization is only valid close to the equilibrium point, so only a small step size

Δs should be taken in the center subspace. The value of Δs varies from system to system, but a good value in the non-dimensional circular restricted three-body problem is between $1e-3$ and $1e-6$.

Equation (2.21) gives how to make a prediction of a periodic orbit about an equilibrium point with the imaginary eigenvalue $i\omega$ and corresponding imaginary eigenvector $\boldsymbol{\psi}$.

$$\tilde{\boldsymbol{x}}_0 = \boldsymbol{x}^* + \text{Im}(\boldsymbol{\psi})\Delta s \quad (2.21a)$$

$$\tilde{T} = \frac{2\pi}{\omega} \quad (2.21b)$$

In the above equations the $(\tilde{\cdot})$ above a variable indicates the quantity is an approximate solution to the nonlinear system. The process to correct the approximate solution and find invariant sets in the nonlinear system is not presented in this thesis, but it follows the ideas in Section 2.6.2. The interested reader is directed to [15] for an algorithm to find families of periodic orbits.

Suppose we have a periodic orbit $\boldsymbol{x}(t)$ with period T and want to study the linear behavior about the orbit. Following the linearization process in Section 2.3.1 gives the following periodic linear system.

$$\delta\dot{\boldsymbol{x}} = A(t)\delta\boldsymbol{x}, \quad \delta\boldsymbol{x}(t_0) = \delta\boldsymbol{x}_0, \quad A(t+T) = A(t) \quad (2.22)$$

The solution to Equation (2.22) is still the general solution to linear systems given in Equation (2.13), however the analysis to determine the eigenspaces is different from an equilibrium point. The theory to study periodic linear systems is known as Floquet theory [133]. A linear system is said to be Floquet or reducible if there exists a periodic change of coordinates $\boldsymbol{x} = P(t)\boldsymbol{y}$ with period T such that the linear system of Equation (2.22) reduces to a constant coefficient linear system. The eigenvalues of the dynamics matrix of the constant coefficient system describe the linear behavior about a periodic orbit.

Determining the eigenspace around a periodic orbit relies on studying the eigenvalues of the STM evaluated over one period, $\Phi(T+t_0, t_0)$. The resulting matrix is called the monodromy matrix and is denoted by M . Suppose $(\lambda, \boldsymbol{\psi})$ is an eigenpair of M , then λ is called a Floquet multiplier of M . The successive application of M to $\boldsymbol{\psi}$ results in

$$M^k \boldsymbol{\psi} = \lambda^k \boldsymbol{\psi} = e^{k \ln \lambda} \boldsymbol{\psi}, \quad (2.23)$$

where $\ln \lambda$ is called the Floquet exponent and is a special case of the Lyapunov exponent. The Lyapunov exponent characterizes the long-term behavior of a system and quantifies the rate at which nearby orbits leave an area.

The classification of the eigenspace relies on the values of the Floquet multipliers and differs from the classification for an equilibrium point. Let $(\lambda, \boldsymbol{\psi})$ be an eigenpair of M , then the three subspaces are defined as follows:

$$E^s = \text{span}\{\text{Re}(\boldsymbol{\psi}), \text{Im}(\boldsymbol{\psi}) \mid |\lambda| < 1\}, \quad (2.24)$$

$$E^u = \text{span}\{\text{Re}(\boldsymbol{\psi}), \text{Im}(\boldsymbol{\psi}) \mid |\lambda| > 1\}, \quad (2.25)$$

$$E^c = \text{span}\{\text{Re}(\boldsymbol{\psi}), \text{Im}(\boldsymbol{\psi}) \mid |\lambda| = 1\}. \quad (2.26)$$

Any eigenvalues inside the unit circle are deemed stable modes of motion. Conversely, any eigenvalues outside the unit circle are deemed unstable modes of motion. Lastly, eigenvalues on the unit circle are deemed to be center modes. Similar to the case of equilibrium points, the eigenvectors lie tangent to the nonlinear invariant manifolds, and provide approximations to the nearby invariant objects.

In Hamiltonian systems M is a symplectic matrix. The result of which is the eigenvalues of M come in reciprocal pairs as noted before. Therefore, stable and unstable manifolds appear together and form a hyperbolic manifold. Additionally, one pair of eigenvalues is always a unity pair [154]. Figure 2.4 shows the distinct eigenvalue pairs which can possibly come from M in a Hamiltonian system. Stable eigenvalues are denoted by λ_s , unstable eigenvalues are denoted by λ_u , non-unity center eigenvalues are denoted by λ_c , and unity eigenvalues are denoted by λ .

The idea of Lyapunov's center theorem can be extended to make a statement about one-parameter families of quasi-periodic orbits:

Let \boldsymbol{x} be a periodic orbit with period T of the system $\dot{\boldsymbol{x}} = \boldsymbol{f}(\boldsymbol{x}, t)$. Suppose $\lambda = a + ib$ is an eigenvalue of the monodromy matrix such that $b \geq 0$ and $\rho = \arctan(b/a) \in \mathbb{R} \setminus \mathbb{Q}$. Then there exists a nearly continuous one-parameter family of quasi-periodic orbits with two frequencies $\boldsymbol{\omega} = (\omega_0, \omega_1)$. Moreover, the frequencies of the orbits approach $(2\pi/T, \rho/T)$ as the orbits approach \boldsymbol{x} .




Figure 2.4: Eigenvalue pairs of the monodromy matrix in a Hamiltonian system.

2.3.5 Quasi-Periodic Orbits

An orbit $\mathbf{x}(t)$ is said to be quasi-periodic if it can be represented as a Fourier series given as

$$\mathbf{x}(t) = \sum_{\mathbf{k}} \mathbf{a}_{\mathbf{k}} e^{i(\boldsymbol{\omega}, \mathbf{k})t}, \quad (2.27)$$

where $\mathbf{k} \in \mathbb{Z}^n$ is a multi-index ranging over all the integers, the coefficients $\mathbf{a}_{\mathbf{k}}$ decay exponentially with the growth of $|\mathbf{k}| = \sum_{i=0}^{n-1} |k_i|$, and $\boldsymbol{\omega} \in \mathbb{R}^n$ constant. The vector $\boldsymbol{\omega}$ is called the frequency vector and satisfies the non-resonance condition of Equation (2.2). A quasi-periodic orbit has no period because there does not exist a T such that $\mathbf{x}(t+T) = \mathbf{x}(t)$.

Similar to a periodic orbit, a change of variables can be made such that the orbit is parameterized by angular coordinates $\boldsymbol{\theta} \in \mathbb{T}^n$, where $\boldsymbol{\theta}(t) = \boldsymbol{\omega}t$. With these coordinates we can associate the torus function

$$\boldsymbol{\tau}(\boldsymbol{\theta}) = \sum_{\mathbf{k}} \tilde{\mathbf{a}}_{\mathbf{k}} e^{i(\boldsymbol{\theta}, \mathbf{k})} \quad (2.28)$$

which is 2π periodic in each angular coordinate. The term “orbit frequencies” is synonymous with “torus frequencies” and are also called the “internal frequencies” or “intrinsic frequencies”. The term “internal” refers to the fact that those frequencies have center modes which have already been excited while there may be other available “external” or “normal” center modes to excite. We have already encountered this idea with both equilibrium points and periodic orbits whose linearization

revealed a center mode of motion normal to the respective orbit.

Just as an equilibrium point with a center mode has closed curves surrounding them, so does the point on a periodic orbit (with a center mode) where the monodromy matrix is computed. In this case the closed curve is called an invariant curve rather than a periodic orbit. The invariant curve is a slice of a 2-dimensional invariant torus with the coordinate θ_0 fixed and parameterized by $\theta_1 \in [0, 2\pi]$. As the point on the periodic orbit flows forward in time the invariant curve flows forward too, sweeping out the rest of the quasi-periodic orbit. This idea extends to cases when a periodic orbit has $n - 1$ center modes in the linearization. The result is an invariant surface diffeomorphic to an $(n - 1)$ -dimensional invariant torus. For this reason the term “invariant surface” is used to describe the slice of an n -dimensional invariant torus when at most $n - 1$ coordinates are fixed.

Let $\mathbf{x}(t_0) = \mathbf{x}_0$ be a point on a periodic orbit with period T whose monodromy matrix M has $n - 1$ pairs of complex conjugate eigenvalues λ_i and $\bar{\lambda}_i$, $i = 1, \dots, n - 1$, with associated eigenvectors $\boldsymbol{\psi}_i$ and $\bar{\boldsymbol{\psi}}_i$, respectively. For each λ_i suppose $\text{Im}(\lambda_i) > 0$, $\rho_i = \arg(\lambda_i) \notin \mathbb{Q}$, and $\frac{\rho_i}{\rho_j} \notin \mathbb{Q}$ for $i \neq j$. The quantity $\boldsymbol{\rho} = [\rho_1, \dots, \rho_{n-1}]^T$ is called the rotation vector with each element called a rotation number. Then an $n - 1$ dimensional linear invariant surface is given in two flavors in Equation (2.29). Equation (2.29a) gives the linear invariant surface relative to the point \mathbf{x}_0 while Equation (2.29b) gives the full states of the surface in phase space (Figure 2.5).

$$\tilde{\boldsymbol{\chi}}_{\text{rel}}(\bar{\boldsymbol{\theta}}) = \sum_{i=1}^{n-1} \Delta s_i (\text{Re}(\boldsymbol{\psi}_i) \cos \theta_i - \text{Im}(\boldsymbol{\psi}_i) \sin \theta_i) \quad (2.29a)$$

$$\tilde{\boldsymbol{\chi}}(\bar{\boldsymbol{\theta}}) = \mathbf{x}_0 + \sum_{i=1}^{n-1} \Delta s_i (\text{Re}(\boldsymbol{\psi}_i) \cos \theta_i - \text{Im}(\boldsymbol{\psi}_i) \sin \theta_i) \quad (2.29b)$$

In the equations above, $\bar{\boldsymbol{\theta}} \in \mathbb{T}^{n-1}$, $\Delta s_i \in \mathbb{R} \setminus \{0\}$ are step sizes to control the amplitude in each center subspace, and the $\tilde{(\cdot)}$ above a variable means it is an approximate solution to the nonlinear system. Varying the Δs_i 's results in a continuous $(n - 1)$ -parametric family of invariant surfaces. It should be noted that the invariant surface in Equation (2.29a) is invariant of M and not the nonlinear dynamics. Equation (2.29b) serves as an initial guess of an invariant surface invariant of the nonlinear dynamics provided the Δs_i 's are small enough. Again, good values here

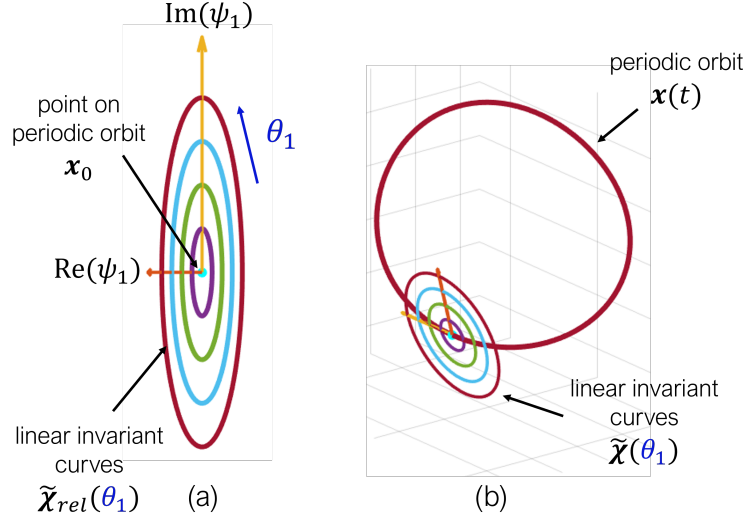


Figure 2.5: Figure of linear invariant curves with different step sizes relative to a periodic orbit (a) and in configuration space (b).

are between $1e-3$ and $1e-6$. Smaller step sizes may be needed if the correction procedure is having issues converging.

Each application of M to $\bar{\chi}_{\text{rel}}$ gives

$$M^k \bar{\chi}_{\text{rel}}(\bar{\theta}) = \bar{\chi}_{\text{rel}}(\bar{\theta} + k\rho), \quad (2.30)$$

indicating points on an invariant surface do not leave the surface, and the points are shifted on the torus by ρ each time M is applied. Applying M corresponds to x_0 completing a full revolution around the periodic orbit and the invariant surface returning to where it began.

Now, the full invariant torus (invariant of the linearized dynamics) is given by

$$\tilde{\tau}(\theta_0, \bar{\theta} + \rho) = x \left(T \frac{\theta_0}{2\pi} \right) + \Phi \left(T \frac{\theta_0}{2\pi}, 0 \right) \bar{\chi}_{\text{rel}}(\bar{\theta}), \quad (2.31)$$

where $\bar{\theta} = [\theta_1, \dots, \theta_{n-1}]^T$ and Φ is the STM of the periodic orbit.

The invariant torus has frequencies $\omega = [\frac{2\pi}{T}, \frac{\rho_1}{T}, \dots, \frac{\rho_{n-1}}{T}]^T$. The period T of the periodic orbit is referred to as the stroboscopic time of the quasi-periodic orbit. It is the time taken for an invariant surface to get back to its starting position. The stroboscopic time is given by the following equation.

$$T = \frac{2\pi}{\omega_0} \quad (2.32)$$

Figure 2.6 depicts the construction of a 2-dimensional linear invariant torus $\tilde{\tau}(\theta_0, \theta_1)$ with several constant θ_0 slices (a), and a complete torus with a few trajectories on the surface and the invariant curve $\tilde{\chi}(\theta_1)$ (b). As time evolves the trajectories wrap around the surface of the torus. Upon completing one stroboscopic period the trajectories return to the invariant curve in the stroboscopic map, however their locations are shifted by the rotation number ρ_1 .

The stroboscopic time is constant in the linear system, but when transitioned to the nonlinear system the stroboscopic time depends on the amplitudes of the invariant tori. Moreover, not all the linear invariant tori survive the effects of the nonlinearity, leaving behind a Cantor family of quasi-periodic invariant tori. These results follow from KAM theory [12, 110, 108, 107, 39, 80, 143].

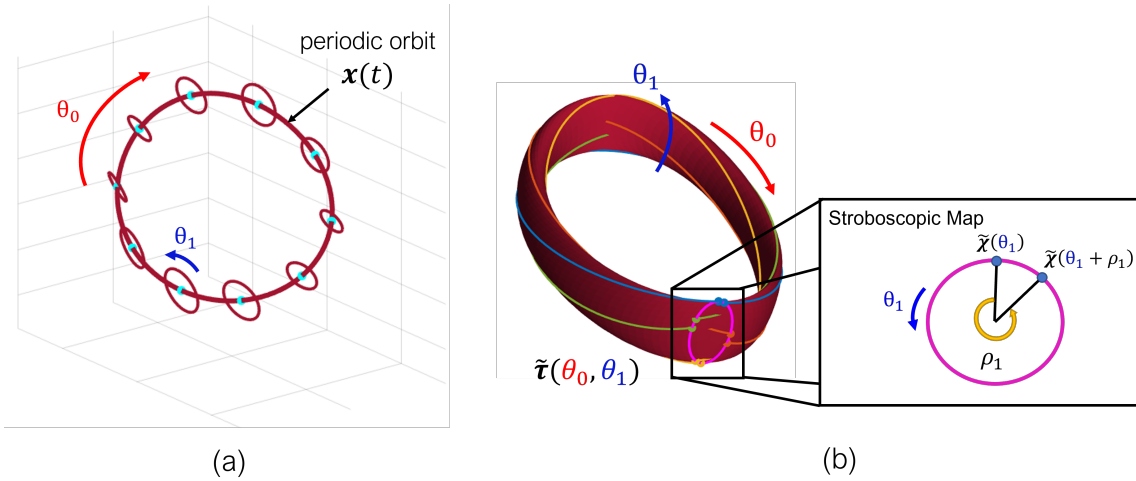


Figure 2.6: Construction of a 2-dimensional linear invariant torus about a periodic orbit (a) and a complete linear invariant torus (b).

A process to compute families of quasi-periodic invariant tori in the nonlinear system is explained in Section 2.6, however we present the fundamental equation which, when satisfied, determines when a torus is an invariant of Equation (2.3) has been found. First, let us introduce some notation. For any torus function $\tau(\theta)$ there infinitely many parameterizations of the torus such that $\tau(\hat{\theta} + \alpha) = \tau(\theta)$. Let $\mathcal{F}_n(\omega) = \{\tau(\theta) = x(t) \mid x(t) \text{ is a quasi-periodic orbit with frequencies } \omega\}$

be the space of torus functions $\tau : \mathbb{T}^n \rightarrow \mathbb{R}^d$ which associate quasi-periodic tori to quasi-periodic trajectories in Equation (2.3) with frequencies ω . Let $\mathcal{X}_{n-1}(T, \rho) = \{\chi(\bar{\theta}) = \tau(\theta_0, \bar{\theta}) \mid \tau \in \mathcal{F}_n(\omega)\}$ be the space of invariant surface functions with

$$\rho = [\omega_1 T, \dots, \omega_{n-1} T]^T, \quad (2.33)$$

where T is the stroboscopic time defined in Equation (2.32). Furthermore, let $R_{-\rho} : \mathcal{X}_{n-1}(T, \rho) \rightarrow \mathcal{X}_{n-1}(T, \rho)$ be the transformation $\chi(\bar{\theta}) \mapsto \chi(\bar{\theta} - \rho)$ which rotates the coordinates $\bar{\theta}$ backward by ρ . In the linear case, the action of the monodromy matrix on the linear invariant surface rotates the coordinates $\bar{\theta}$ to $\bar{\theta} + \rho$. In the nonlinear case the solution flow φ_T has the same effect on the nonlinear invariant surface given in Equation (2.34). The solution flow φ_T is a map called the stroboscopic map with stroboscopic time T .

$$\varphi_T(\chi(\bar{\theta})) = \chi(\bar{\theta} + \rho) \quad (2.34)$$

Finally, we introduce the quasi-periodicity constraint

$$R_{-\rho} \varphi_T(\chi(\bar{\theta})) - \chi(\bar{\theta}) = \mathbf{0}. \quad (2.35)$$

Any (χ, T, ρ) which satisfies Equation (2.35) produces an invariant surface. The full invariant torus τ is recovered by flowing χ forward in time from t_0 to $T + t_0$ (Equation (2.36)). The frequencies are recovered from Equations (2.32) and (2.33) which uses the stroboscopic time to transform the rotation vector into the frequency vector.

$$\tau(2\pi(t - t_0)/T, \bar{\theta} + \omega(t - t_0)) = \varphi_t(\chi(\bar{\theta})), \quad t \in [t_0, T + t_0] \quad (2.36)$$

Suppose we know an invariant surface $\chi(\bar{\theta})$ with a rotation vector ρ which satisfies Equation (2.35) and wish to study how deviations $\psi(\bar{\theta})$ behave linearly under the stroboscopic map φ_T . To study the linear behavior about an invariant surface we Taylor expand $\varphi_T(\chi(\bar{\theta}) + \psi(\bar{\theta}))$ about $\chi(\bar{\theta})$.

$$\varphi_T(\chi(\bar{\theta}) + \psi(\bar{\theta})) = \varphi_T(\chi(\bar{\theta})) + D_{\chi} \varphi_T(\chi(\bar{\theta})) \psi(\bar{\theta}) + \mathcal{O}(\|\psi(\bar{\theta})\|^2)$$

Substituting $\varphi_T(\chi(\bar{\theta})) = \chi(\bar{\theta} + \rho)$ from Equation (2.34) and dropping higher order terms gives

$$\varphi_T(\chi(\bar{\theta}) + \psi(\bar{\theta})) = \chi(\bar{\theta} + \rho) + D_{\chi}\varphi_T(\chi(\bar{\theta}))\psi(\bar{\theta}).$$

Here, we let $\varphi_T(\chi(\bar{\theta}) + \psi(\bar{\theta})) = \chi(\bar{\theta} + \rho) + \tilde{\psi}(\bar{\theta} + \rho)$. Note that $\tilde{\psi}(\bar{\theta} + \rho)$ is not the same as $\psi(\bar{\theta} + \rho)$ because it includes the effects of any rotations, stretching, and shrinking induced by the flow. After canceling like terms the resulting system is

$$\tilde{\psi}(\bar{\theta} + \rho) = D_{\chi}\varphi_T(\chi(\bar{\theta}))\psi(\bar{\theta}) \quad (2.37)$$

$$= \Phi(T, 0; \bar{\theta})\psi(\bar{\theta}) \quad (2.38)$$

$$= A(\bar{\theta})\psi(\bar{\theta}), \quad (2.39)$$

where $\Phi(T, 0; \bar{\theta})$ is the STM associated with a point on the invariant surface at a fixed value of $\bar{\theta}$.

Examining Equation (2.39) it is clear this is the linearization of the stroboscopic map and governs how deviations behave around the invariant surface. The system is a discrete quasi-periodic linear system and needs special treatment to determine the linear behavior around χ . Equation (2.39) can be rewritten as the linear quasi-periodic skew-product

$$\hat{\psi} = A(\theta)\psi \quad (2.40)$$

$$\hat{\theta} = \theta + \rho, \quad (2.41)$$

where the $\hat{(\cdot)}$ above a quantity represents an updated quantity after the discrete map has been applied. The $\bar{(\cdot)}$ notation has been dropped as it is assumed in the linear system we are dealing with invariant surfaces and discrete maps instead of invariant tori and flows. In [102] Equation (2.40) is called reducible if and only if there exists a quasi-periodic change of variables $\psi = C(\theta)\mathbf{y}$ such that Equation (2.40) becomes

$$\hat{\mathbf{y}} = B\mathbf{y} \quad (2.42)$$

$$\hat{\theta} = \theta + \rho \quad (2.43)$$

where $B = C^{-1}(\theta + \rho)A(\theta)C(\theta)$ is constant. We call B the Floquet matrix of the torus and its eigenvalues the Floquet multipliers. The dynamics of Equation (2.40) are described by the

dynamics of Equation (2.42). The issue is determining $C(\boldsymbol{\theta})$ to reduce the system to constant coefficient. However, Jorba in [102] proves the eigenvalues and eigenvectors can be found by solving a generalized eigenvalue problem.

Revisiting Equation (2.39), the term $\tilde{\boldsymbol{\psi}}(\bar{\boldsymbol{\theta}} + \boldsymbol{\rho})$ is rewritten as $R_{\boldsymbol{\rho}}\tilde{\boldsymbol{\psi}}(\bar{\boldsymbol{\theta}})$. Substituting this into Equation (2.39) and dropping the $(\bar{\cdot})$ notation gives

$$R_{\boldsymbol{\rho}}\tilde{\boldsymbol{\psi}}(\boldsymbol{\theta}) = A(\boldsymbol{\theta})\boldsymbol{\psi}(\boldsymbol{\theta}). \quad (2.44)$$

Suppose $\boldsymbol{\psi}(\boldsymbol{\theta})$ is chosen carefully such that it is an eigenfunction with eigenvalue η of the generalized eigenvalue problem

$$\eta R_{\boldsymbol{\rho}}\boldsymbol{\psi}(\boldsymbol{\theta}) = A(\boldsymbol{\theta})\boldsymbol{\psi}(\bar{\boldsymbol{\theta}}). \quad (2.45)$$

It is shown in [102] that if $(\eta, \boldsymbol{\psi}(\boldsymbol{\theta}))$ is an eigenpair of Equation (2.45) then so is $(\eta e^{i\langle \mathbf{k}, \boldsymbol{\rho} \rangle}, \boldsymbol{\psi}(\boldsymbol{\theta}) e^{i\langle \mathbf{k}, \boldsymbol{\rho} \rangle})$ for $\mathbf{k} \in \mathbb{Z}^{n-1}$, showing the eigenvalues form concentric rings in the complex plane centered on the origin. Moreover, if η is an eigenvalue of B then it is also an eigenvalue of Equation (2.45). Therefore, the eigenvalues λ of B which characterize the linear stability of $\boldsymbol{\chi}$ are contained within the spectrum of Equation (2.45). Furthermore, the eigenfunctions $\boldsymbol{\psi}(\boldsymbol{\theta})$ form the columns of $C(\boldsymbol{\theta})$, and are the functions which determine the eigenspaces of $\boldsymbol{\chi}$. Jorba presents a method to sort $\eta e^{i\langle \mathbf{k}, \boldsymbol{\rho} \rangle}$ to find the appropriate η such that $\eta = \lambda$ is in the spectrum of B . See [102] for the details and more on the linear stability of invariant surfaces. The presentation of the material is for invariant curves, but the results hold for higher dimensions.

The classification of the linear behavior about a quasi-periodic invariant torus relies on the eigenvalues of B . Let λ be an eigenvalue of B with associated eigenfunction $\boldsymbol{\psi}(\boldsymbol{\theta})$, then the three subspaces are defined as follows:

$$E^s(\boldsymbol{\theta}) = \text{span}\{\text{Re}(\boldsymbol{\psi}(\boldsymbol{\theta})), \text{Im}(\boldsymbol{\psi}(\boldsymbol{\theta})) \mid |\lambda| < 1\}, \quad (2.46)$$

$$E^u(\boldsymbol{\theta}) = \text{span}\{\text{Re}(\boldsymbol{\psi}(\boldsymbol{\theta})), \text{Im}(\boldsymbol{\psi}(\boldsymbol{\theta})) \mid |\lambda| > 1\}, \quad (2.47)$$

$$E^c(\boldsymbol{\theta}) = \text{span}\{\text{Re}(\boldsymbol{\psi}(\boldsymbol{\theta})), \text{Im}(\boldsymbol{\psi}(\boldsymbol{\theta})) \mid |\lambda| = 1\}. \quad (2.48)$$

Now, each subspace varies continuously on $\boldsymbol{\theta}$ (i.e. the location on the invariant surface). The

radius of the rings give insight into the stability of χ . Rings of radius greater than 1 represent unstable modes of motion. Rings of radius less than 1 represent stable modes of motion. Lastly, rings of radius 1 represent center modes of motion. The magnitudes of the rings quantify how strongly stable or unstable a particular mode is. In a Hamiltonian system the eigenvalues come in reciprocal pairs. So an unstable ring of radius r is necessarily paired with a stable ring of radius $1/r$. Furthermore, for an n -dimensional quasi-periodic invariant torus there are $n - c$ pairs of unity eigenvalues, where c is the number of forcing frequencies in the dynamical system.

Let $\chi(\bar{\theta})$ be an $(n - 1)$ -dimensional invariant surface of an n -dimensional quasi-periodic invariant torus with internal frequencies ω . Suppose Equation (2.40) is reducible to Equation (2.42). Further, suppose there exists r complex eigenvalues λ_i of B on the unit circle with eigenfunctions $\psi(\bar{\theta})$, and each λ_i has an associated rotation number $\rho_i = \arg(\lambda_i) \notin \mathbb{Q}$ and none of the rotation numbers are commensurate. The frequencies $\omega_i = \frac{\rho_i}{T}$ associated to λ_i are called the normal frequencies of the quasi-periodic invariant torus. Then an approximation of an $n + r - 1$ -dimensional invariant surface is given in a similar fashion to Equation (2.29b). However, r additional coordinates θ_i , $i = n, n + 1, \dots, n + r - 1$ need to be considered to construct the new invariant surface diffeomorphism. Let $\tilde{\theta} \in \mathbb{T}^{n+r-1}$ where $\tilde{\theta} = [\bar{\theta}^T, \theta_n, \dots, \theta_{n+r-1}]^T$. An approximation for a new invariant surface is given by

$$\tilde{\chi}(\tilde{\theta}) = \chi(\bar{\theta}) + \sum_{i=1}^r \Delta s_i (\operatorname{Re}(\psi_i(\bar{\theta})) \cos(\theta_{n-1+i}) - \operatorname{Im}(\psi_i(\bar{\theta})) \sin(\theta_{n-1+i})) \quad (2.49)$$

with rotation vector $\tilde{\rho} = [\rho^T, \rho_1, \dots, \rho_r]^T$. Upon correcting the approximate invariant surface to satisfy the nonlinear dynamics one can easily recover the full invariant torus. The frequencies of this torus are recovered from Equation (2.33), all of which are now considered internal frequencies.

2.3.6 Invariant Stable and Unstable Manifolds

In Sections 2.3.3 to 2.3.5, we present the linear subspaces about different types of orbits. These subspaces are only valid locally to each orbit, but their directions are tangent to the nonlinear invariant manifolds [133]. In each of those sections we show how to generate initial guesses for the

nonlinear objects residing in the center manifolds. This section presents the method for computing the nonlinear stable and unstable invariant manifolds emanating from quasi-periodic orbits. From here on out we refer to these simply as the stable and unstable manifolds. The method can easily be modified to compute stable and unstable manifolds emanating from periodic orbits and equilibrium points.

Before going into the method we first introduce the definition of the stable and unstable manifolds W^s and W^u , respectively. To measure the distance between a point \mathbf{x} in phase space and an invariant set Λ we define $\rho_\Lambda : \mathbb{R}^d \rightarrow \mathbb{R}$ to be

$$\rho_\Lambda(\mathbf{x}) = \min(\|\mathbf{x} - \mathbf{y}\| \mid \mathbf{y} \in \Lambda). \quad (2.50)$$

Then the stable and unstable manifolds of a smooth vector field are defined as

$$W^s(\Lambda) = \{\mathbf{x} \notin \Lambda : \lim_{t \rightarrow \infty} \rho_\Lambda(\varphi_t(\mathbf{x})) = 0\} \quad (2.51)$$

$$W^u(\Lambda) = \{\mathbf{x} \notin \Lambda : \lim_{t \rightarrow -\infty} \rho_\Lambda(\varphi_t(\mathbf{x})) = 0\} \quad (2.52)$$

A stable manifold is the collection of orbits that are forward (in time) asymptotic to an invariant set. Conversely, an unstable manifold is the collection of orbits that are backward asymptotic to an invariant set. These manifolds are themselves invariant sets according to the definition provided in Section 2.3.2.

Equations (2.51) and (2.52) provide insight as how to generate the stable and unstable manifolds for a quasi-periodic orbit. In order to generate stable and unstable manifolds, one needs to find appropriate points near the initial orbit and numerically integrate backwards and forwards in time, respectively. The appropriate points are generated from the subspaces $E^s(\bar{\boldsymbol{\theta}})$ and $E^u(\bar{\boldsymbol{\theta}})$ in Equations (2.46) and (2.47). Approximate points on the manifolds are generated from functions $\mathbf{v}^{(\cdot)}(\bar{\boldsymbol{\theta}}) \in E^{(\cdot)}(\bar{\boldsymbol{\theta}})$, where (\cdot) is a placeholder for s and u . Note that \mathbf{v} only initializes directions off of the invariant surface and must be mapped across the entire invariant torus.

We restrict our attention to the case that $E^s(\bar{\boldsymbol{\theta}})$ and $E^u(\bar{\boldsymbol{\theta}})$ are one-dimensional vector spaces for a fixed value of $\bar{\boldsymbol{\theta}}$. Doing so simplifies the discussion and this is the case encountered in this

thesis. Let $\lambda^{(\cdot)} \in \mathbb{R}$ be an eigenvalue of Equation (2.45) with the associated eigenfunction $\psi^{(\cdot)}(\bar{\theta})$. Then $\psi^{(\cdot)}(\bar{\theta})$ necessarily takes on real values. The eigenfunction $\psi^{(\cdot)}(\bar{\theta})$ must be mapped across the surface of the torus in order to generate initial points covering the manifold. The mapping is accomplished by left multiplying $\psi^{(\cdot)}(\bar{\theta})$ by the STM $\Phi(t, t_0; \bar{\theta})$. The stretching or shrinking effects due to $\lambda^{(\cdot)}$ must be taken into account, so the initial points on the manifold have consistent displacement sizes. After the stretching or shrinking has been taken care of then these directions are multiplied by a small step size and added to the torus function $\tau(\theta)$.

Note that since $\psi^{(\cdot)}(\bar{\theta})$ is a basis vector for a fixed value of $\bar{\theta}$ then $-\psi^{(\cdot)}(\bar{\theta})$ is as well, so we need to add and subtract $\psi^{(\cdot)}(\bar{\theta})$ to generate the full manifold. We write $W_+^{(\cdot)}$ when the portion of the manifold is generated by adding $\psi^{(\cdot)}(\bar{\theta})$ and write $W_-^{(\cdot)}$ when the portion of the manifold is generated by subtracting $\psi^{(\cdot)}(\bar{\theta})$. We call $W_+^{(\cdot)}$ and $W_-^{(\cdot)}$ half-manifolds. The full manifold is the union of each half-manifold, $W^{(\cdot)} = W_+^{(\cdot)} \cup W_-^{(\cdot)}$.

Let $\mathbf{w}^{(\cdot)} : \mathbb{T}^n \times \mathbb{R} \rightarrow \mathbb{R}^d$ be a diffeomorphism which maps from a cylinder to phase space. This function is often referred to as a cylinder function. Approximate points on the half-manifolds are given by

$$\mathbf{w}_{\pm}^{(\cdot)}(\theta, t_0) = \tau(\theta) \pm \Delta s \left(\lambda^{(\cdot)} \right)^{\frac{-\theta_0}{2\pi}} \Phi \left(T \frac{\theta_0}{2\pi}; \bar{\theta} \right) \psi^{(\cdot)}(\bar{\theta}). \quad (2.53)$$

Good values of Δs depend on the system, but are likely between 1e-4 and 1e-8. The value of Δs determines the amount of truncation error due to the linear approximation of the nonlinear invariant manifold. To generate the manifold at time t one simply needs to flow $\mathbf{w}_{\pm}^{(\cdot)}(\bar{\theta}, t_0)$ from time t_0 to t .

Figure 2.7 depicts the perturbing directions of an unstable eigenfunction across the invariant curve in plot (a) and across the entire quasi-periodic orbit in plot (b) for a discretized quasi-halo orbit in the Earth-Moon system of the circular restricted three-body problem.

2.3.7 Common Torus Terminology

In literature there are many types of studied tori. This section is devoted to defining commonly encountered tori which are used in this thesis.




Figure 2.7: Unstable directions to initialize the unstable manifold on an invariant curve (a) and on a quasi-periodic orbit (b).

2.3.7.1 Diophantine Tori

An important type of torus encountered in literature dealing with the computation of quasi-periodic invariant tori is a Diophantine torus. It is closely related to a quasi-periodic torus, however stricter non-resonance conditions are placed on the frequency vector, called the Diophantine condition (Equation (2.54)). The Diophantine condition is one of the two main assumptions on tori which KAM theory can deal with [12, 56, 80]. In short, KAM theory establishes the existence and persistence of quasi-periodic invariant tori in dynamical systems.

An n -torus is called a Diophantine n -torus if

$$|\langle \boldsymbol{\omega}, \mathbf{k} \rangle| \geq c |\mathbf{k}|^{-\gamma}, \quad \gamma > n - 1 \quad (2.54)$$

for all $\mathbf{k} \in \mathbb{Z}^n \setminus \{\mathbf{0}\}$, where $c > 0$, $\gamma > 0$, and $|\mathbf{k}| = \sum_{i=0}^n |k_i|$. Clearly, Diophantine tori are a subset of quasi-periodic tori, but the frequency vector is constrained to be a certain distance away from resonances between the frequencies.

The set of tori satisfying the Diophantine condition can be shown to be a Cantor manifold. A Cantor manifold is not smooth by the standard definition presented in calculus, but it is smooth in the sense of Whitney [168, 143, 39], meaning a function defined on the set of Diophantine tori can be extended to a smooth function. The idea of a Cantor manifold will be mentioned throughout the thesis as this directly relates to families of quasi-periodic orbits.

2.3.7.2 KAM Tori

A KAM torus is the type of torus which is studied in KAM theory. KAM tori satisfy, in addition to the Diophantine condition of Equation (2.54), a non-degeneracy condition [39, 12, 80, 56]. The non-degeneracy condition states an invariant torus's frequencies change with the amplitude of the torus and the invariant tori are uniquely parameterized by their frequencies. These tori, when subject to sufficiently small and smooth perturbations, deform and persist in the new dynamical system. KAM tori have been widely studied in literature and are of great importance in astrodynamics [107, 108, 81, 46, 102, 66, 111].

2.3.7.3 Reducible Tori

A reducible, or Floquet, torus is an invariant torus whose linearization can be reduced to the linear system of constant coefficients (Equation (2.42)) by a quasi-periodic change of variables [39].

2.3.7.4 Lagrangian Tori

A Lagrangian torus is a maximal-dimensional invariant torus [38]. For example, In the circular restricted three-body problem the maximum dimension of tori is 3. Moreover, In autonomous Hamiltonian systems with n degrees of freedom and r non-resonant perturbing frequencies the maximum dimension of tori is $n + r$ [104].

2.3.7.5 Elliptic Tori

An elliptic torus is an invariant torus which only has center modes of motion in its linearization, corresponding to the eigenvalues of B residing only on the unit circle [107].

2.3.7.6 Whiskered Tori

A whiskered torus is an invariant torus which has stable or unstable modes of motion in its linearization, resulting in the torus having stable or unstable manifolds, respectively, in the nonlinear system. The manifolds appear as “whiskers” emanating from the torus.

2.3.7.7 Partially Elliptic and Partially Hyperbolic Tori

A partially elliptic or partially hyperbolic torus is an invariant torus which is not elliptic, and will necessarily have whiskers, making them synonymous with whiskered tori.

2.4 The Circular Restricted Three-Body Problem

It is common practice to study simplified dynamical models to gain an understanding of the dynamics of higher fidelity models which represent more realistic environments that spacecraft will be exposed to. This thesis considers the circular restricted three-body problem as a test bed for the development, testing, and analysis of tools which can be applied in a variety of dynamical systems. The circular restricted three-body problem has been studied since the 1700s by famous mathematicians such as Euler, Lagrange, and Poincaré. Many books have been published on the circular restricted three-body problem [24, 129, 153, 161, 150, 165], and many papers describe the circular restricted three-body problem [94, 96, 37, 29, 75]. For this reason we omit a from-first-principles development of the equations of motion and provide only a brief description of the circular restricted three-body problem.

The circular restricted three-body problem is a Hamiltonian dynamical system which describes the motion of a massless particle due to gravitational forces from two massive bodies, P_1 and P_2 , called the primaries. They have masses m_1 and m_2 such that $m_1 \geq m_2$. The primaries move in circular orbits about their common center of mass according to Keplerian dynamics. The reference frame is a uniformly rotating frame with origin at the center of mass. The positive x -axis lies along the vector pointing from P_1 to P_2 . The positive z -axis is aligned with the orbital angular momentum vector of the primaries, and the positive y -axis completes the right-handed coordinate

system. The equations of motion in dimensionless form are

$$\begin{aligned}\ddot{x} &= x + 2\dot{y} - \frac{(1-\mu)(x+\mu)}{r_1^3} - \frac{\mu(x+\mu-1)}{r_2^3} \\ \ddot{y} &= y - 2\dot{x} - \frac{(1-\mu)y}{r_1^3} - \frac{\mu y}{r_2^3} \\ \ddot{z} &= -\frac{(1-\mu)z}{r_1^3} - \frac{\mu z}{r_2^3}\end{aligned}\tag{2.55}$$

where $\mu = \frac{m_2}{m_1+m_2}$ is the dimensionless mass distribution of the system, $r_1 = \sqrt{(x+\mu)^2 + y^2 + z^2}$ is the distance between the massless particle and P_1 , and $r_2 = \sqrt{(x+\mu-1)^2 + y^2 + z^2}$ is the distance between the massless particle and P_2 . In the non-dimensional form the dynamics admit one dimensionless constant of motion called the Jacobi constant [150]. The equation for the Jacobi constant used here is

$$J = 2\left(\frac{1-\mu}{r_1} + \frac{\mu}{r_2}\right) + x^2 + y^2 - (\dot{x}^2 + \dot{y}^2 + \dot{z}^2)\tag{2.56}$$




Figure 2.8: Coordinate system and equilibrium points of the circular restricted three-body problem.

The circular restricted three-body problem has five equilibrium points called the Lagrange points. Three of them lie along the x -axis and are called the co-linear points labeled L_1 , L_2 , and L_3 . Each of these equilibrium points have *center* \times *center* \times *saddle* behavior, meaning there are two center subspaces (E^c) and one hyperbolic subspace ($E^s \oplus E^u$) in the linearization. The other

two points are called the equilateral points as each point forms an equilateral triangle with the two primaries. These points are labeled as L_4 and L_5 and each have normal behavior of type $center \times center \times center$ for values of $\mu < 0.0385208965$, after which the normal behavior is of type $center \times saddle \times saddle$ [161, 150]. The Lagrange points are depicted in Figure 2.8 along with the primaries and reference frame of the circular restricted three-body problem.

The circular restricted three-body problem equations of motion exhibit symmetry. This symmetry is seen from the change of variables $(x, y, z, t) \mapsto (x, -y, z, -t)$ which preserves the equations of motion. Moreover, this symmetry is the same reason that the halo orbits have a northern and southern family, two families that are symmetric about the z -axis [37]. Additionally, for symmetric orbits centered on the x -axis only the stable or unstable invariant manifold needs to be computed; the other can be determined from the symmetry.

2.5 Smooth Manifolds

Chapter 6 focuses on the use of gradients to enable optimization over families of quasi-periodic orbits. Gradients are taken of functions whose arguments are quasi-periodic orbits. To perform the optimization we need

- (1) Information about the smoothness of the families of quasi-periodic orbits, and
- (2) To know how to define these gradients.

We have already seen in Section 2.3.7.1 that a family of quasi-periodic orbits form a Cantor manifold and are smooth in the sense of Whitney, allowing us to treat families of quasi-periodic orbits as smooth manifolds.

The objective of Section 2.5 is to develop the mathematical framework which enables us to define gradients of functions whose domain are families of quasi-periodic orbits. Moreover, Sections 2.5.1 and 2.5.2 help with the mathematical development in Section 2.6. We do not present a self-contained background as much of the material covered in this section comes from Boumal's easy-to-read book on optimization on smooth manifolds [33]. We do, however, provide the basics

which are necessary for this thesis. We encourage the reader to fill in the blanks by reading Chapter 3 and Section 7.7 in [33].

2.5.1 Implicit Function Theorem

Let $\mathbf{x} \in \mathbb{R}^D$ be state variables and $\mathbf{q} \in \mathbb{R}^p$ be a set of parameters². Let U be an open set in $\mathbb{R}^D \times \mathbb{R}^p$ and define the function

$$\mathbf{F} : U \rightarrow \mathbb{R}^D. \quad (2.57)$$

Further, define the set \mathcal{M} to be

$$\mathcal{M} = \{(\mathbf{x}, \mathbf{q}) \in U \mid \mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\}. \quad (2.58)$$

Under the conditions of the implicit function theorem (IFT), the state variables are uniquely parameterized locally by the parameters and we say that \mathcal{M} is an implicitly defined smooth manifold. As we will see in Section 2.6 the system of equations which define families of quasi-periodic orbits is represented by functions of this type, and we may treat the Cantor manifold with the same calculus with which we treat smooth manifolds \mathcal{M} .

Theorem 2.5.1 (Implicit Function Theorem). *Let U be an open set in $\mathbb{R}^D \times \mathbb{R}^p$ with coordinates (\mathbf{x}, \mathbf{q}) . Let \mathbf{F} in Equation (2.57) be continuously differentiable. Let $(\mathbf{a}, \mathbf{b}) \in U$ such that $\mathbf{F}(\mathbf{a}, \mathbf{b}) = \mathbf{0}$. If $D_{\mathbf{x}}\mathbf{F}(\mathbf{a}, \mathbf{b})$ is invertible, then there exist an open set $X \subset \mathbb{R}^D$ containing \mathbf{a} and an open set $Q \subset \mathbb{R}^p$ containing \mathbf{b} satisfying the following.*

(1) *There is a unique function $\mathbf{g} : Q \rightarrow \mathbb{R}^D$ such that $\mathbf{g}(\mathbf{b}) = \mathbf{a}$ and $\mathbf{F}(\mathbf{g}(\mathbf{q}), \mathbf{q}) = \mathbf{0}$ for all $\mathbf{q} \in Q$.*

(2) *The function \mathbf{g} is continuously differentiable and for all $\mathbf{q} \in Q$*

$$D_{\mathbf{q}}\mathbf{g}(\mathbf{q}) = -D_{\mathbf{x}}^{-1}\mathbf{F}(\mathbf{g}(\mathbf{q}), \mathbf{q})D_{\mathbf{q}}\mathbf{F}(\mathbf{g}(\mathbf{q}), \mathbf{q}). \quad (2.59)$$

² The parameters are often called control variables and design variables. These terms are used interchangeably.

The conclusions of the implicit function theorem say the state variables \mathbf{x} are uniquely parameterized locally by the parameters³, and the set of solutions \mathcal{M} define a p -dimensional smooth embedded submanifold of Euclidean space [85, 33]. Manifolds of this type are referred to here simply as smooth manifolds. The invertibility of $D_{\mathbf{x}}\mathbf{F}$ can easily be checked by computing the rank of the matrix. If the rank of $D_{\mathbf{x}}\mathbf{F}$ is equal to \mathcal{D} then it is full rank and invertible. Equation (2.59) presents how to compute the derivatives of the states with respect to the parameters. In practice, we do not use this equation to do this. Instead, we retrieve the derivative from the tangent space of \mathcal{M} .

2.5.2 Tangent Space of a Smooth Manifold

For the remainder of Section 2.5 let us forget about \mathbf{x} and \mathbf{q} and let $\mathbf{z} \in \mathbb{R}^{m+k}$, so that $\mathbf{F} : \mathbb{R}^{m+k} \rightarrow \mathbb{R}^k$. We modify Equation (2.58) to become

$$\mathcal{M} = \{\mathbf{z} \in \mathbb{R}^{m+k} \mid \mathbf{F}(\mathbf{z}) = \mathbf{0}\}. \quad (2.60)$$

Define $T_{\mathbf{z}}\mathcal{M}$ to be the tangent space of \mathcal{M} at the point \mathbf{z} . To determine $T_{\mathbf{z}}\mathcal{M}$, let us perform a Taylor series expansion of \mathbf{F} about the point $\mathbf{z} + \boldsymbol{\alpha}$ with $\mathbf{z} \in \mathcal{M}$.

$$\mathbf{F}(\mathbf{z} + \boldsymbol{\alpha}) = \mathbf{F}(\mathbf{z}) + D\mathbf{F}(\mathbf{z})\boldsymbol{\alpha} + \boldsymbol{\alpha}^T D^2\mathbf{F}(\mathbf{z})\boldsymbol{\alpha} + \mathcal{O}(\|\boldsymbol{\alpha}\|^3) \quad (2.61)$$

$$= D\mathbf{F}(\mathbf{z})\boldsymbol{\alpha} + \boldsymbol{\alpha}^T D^2\mathbf{F}(\mathbf{z})\boldsymbol{\alpha} + \mathcal{O}(\|\boldsymbol{\alpha}\|^3) \quad (2.62)$$

To find a first-order approximation of $\mathbf{F}(\mathbf{z} + \boldsymbol{\alpha})$ such that $\mathbf{F}(\mathbf{z} + \boldsymbol{\alpha}) \approx \mathbf{0}$ it is required that $D\mathbf{F}(\mathbf{z})\boldsymbol{\alpha} = \mathbf{0}$. It follows that $T_{\mathbf{z}}\mathcal{M}$ is

$$T_{\mathbf{z}}\mathcal{M} = \{\mathbf{z}' \mid D\mathbf{F}(\mathbf{z})\mathbf{z}' = \mathbf{0}\} \quad (2.63)$$

the null space of the Jacobian matrix of \mathbf{F} evaluated at \mathbf{z} . We write $(\cdot)'$ to indicate that a quantity is in a tangent space.

³ This condition is quite similar to the non-degeneracy condition in KAM theory where the parameters are the torus frequencies. See Section 2.3.7.2.




Figure 2.9: Depiction of the tangent space of a manifold at the point z .

The tangent space is an m -dimensional Euclidean subspace of \mathbb{R}^{m+k} with linearly independent basis vectors $z'_j \in \mathbb{R}^{m-k}$, $j = 1, \dots, m$ (Figure 2.9). The tangent space provides linear information about \mathcal{M} in a small neighborhood of z . Then for any $z' \in T_{z_0}\mathcal{M}$ the quantity

$$\tilde{z} = z_0 + z' \quad (2.64)$$

is a first-order approximation of another point $z \in \mathcal{M}$.

2.5.3 Retractions

Given the approximation of a point \tilde{z} from Equation (2.64) we need a way to move back to the manifold. There are many points z on \mathcal{M} which are near \tilde{z} , however we are interested in a map $R_{z_0} : T_{z_0}\mathcal{M} \rightarrow \mathcal{M}$ such that for each $z' \in T_{z_0}\mathcal{M}$ the point $R_{z_0}(z') = z$ is unique ((a) in Figure 2.10). Moreover, we require the path on \mathcal{M} which R_{z_0} follows be a smooth curve $c(t)$ such that $c(0) = z_0$ and $c'(0) = z'$ ((b) in Figure 2.10). According to Boumal in [33], the function which has these properties is called a retraction.

In this thesis a retraction is used to computationally move through a family of orbits to search for orbits with specific characteristics or search for optimal solutions. The property to move to a unique point on \mathcal{M} enables the precise computation of an orbit from within a family of orbits.




Figure 2.10: Ambiguity with moving on a manifold. In (a) there is a region of points which are considered close to the approximated point \tilde{z} . In (b) there are multiple curves which can be followed to get to a point on the manifold.

The property to move away from z_0 in a given direction on a manifold is “the basic operation of a gradient descent algorithm, and of essentially all optimization algorithms on manifolds” [33].

2.5.4 Gradients on Smooth Manifolds

Let \mathcal{M} be defined as in Equation (2.60) with $\mathbf{F} = [F_1, \dots, F_k]^T$ where each $F_i : \mathbb{R}^{m+k} \rightarrow \mathbb{R}$ is defined on a Euclidean space of dimension strictly greater than k . Let $f : \mathbb{R}^{m+k} \rightarrow \mathbb{R}$ be a continuously differentiable function defined on a Euclidean space and $\bar{f} = f|_{\mathcal{M}}$ be the restriction of f to \mathcal{M} . The gradient of \bar{f} , i.e. the gradient on the manifold, is the Euclidean gradient of f projected onto $T_z\mathcal{M}$. The projection onto $T_z\mathcal{M}$ depends on the Jacobian matrix of \mathbf{F} . The projection operator $\text{Proj}_z : \mathbb{R}^{m+k} \rightarrow T_z\mathcal{M}$ is defined as

$$\text{Proj}_z(\mathbf{v}) = \mathbf{v} - [D\mathbf{F}(z)]^T \left[[D\mathbf{F}(z)]^T \right]^\dagger \mathbf{v}, \quad (2.65)$$

where the \dagger is the Moore-Penrose inverse. Then the gradient of $f(z)$ is given as

$$\nabla \bar{f}(z) = \text{Proj}_z(\nabla f(z)). \quad (2.66)$$

In this thesis we do not use Equation (2.65). Recall from Equation (2.63) that $T_z\mathcal{M}$ is the null space of \mathbf{F} . Then the gradient of $\bar{f}(z)$ is simply the portion of the Euclidean gradient of $f(z)$ which

lies in $T_z\mathcal{M}$, equating to solving an underdetermined system linear system. The solution to the underdetermined system is

$$\nabla\bar{f}(z) = V \left[V^\dagger \nabla f(z) \right], \quad (2.67)$$

where the columns of V are composed of any linearly independent basis vectors spanning $T_z\mathcal{M}$.

2.6 Computation of Quasi-Periodic Orbits

This section is devoted to walking through the building blocks to be able to compute families of quasi-periodic orbits. These techniques are well studied and are used widely throughout many fields of engineering, mathematics, and science. The solution techniques presented in this chapter can be found in the books by Ascher et al. [13] and Seydel [158]. Both books cover the theory of ordinary differential equations (ODEs). Ascher et al. focus on solution methods for two point boundary value problems (TPBVPs) and provides a review of numerical analysis, whereas Seydel focuses on the continuation of solutions and bifurcations in dynamical systems. Together the information within these books form the foundation for most modern algorithms to compute families of solutions to ODEs.

Computing families of quasi-periodic orbits boils down to solving a TPBVP for ODEs in succession. We proceed in an outside-in approach to describe the components that come together to compute families of solutions in ODEs. Starting with the ideas of a continuation method, and then moving into the correction procedure.

2.6.1 Continuation

Most continuation methods compute solutions to form a 1-parameter subset of a family. We call a 1-parameter family a branch of solutions, as it resembles a branch from a tree. A multi-parameter continuation method was developed by Henderson [85] and used by Henry and Scheeres [89] to explore a 3-parameter family of heteroclinic connections between families of quasi-periodic orbits in the circular restricted three-body problem. Recently, Baresi et al. developed a continuation method to compute a Taylor polynomial of continuous solutions [17].

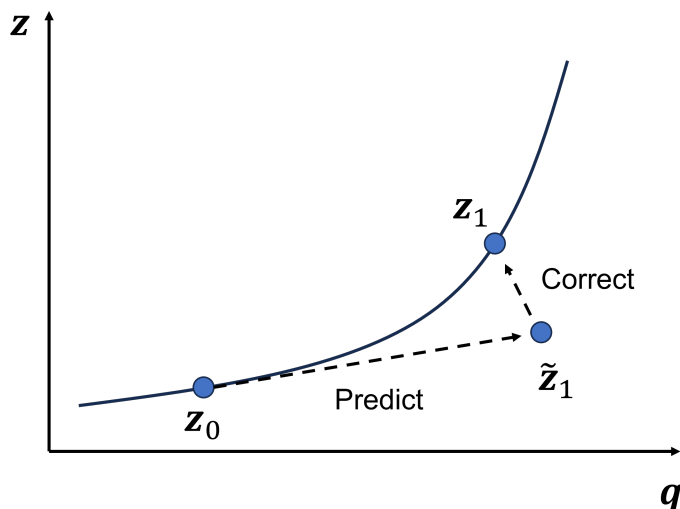


Figure 2.11: Schematic of a predictor-corrector algorithm.

A continuation method aims to compute a family of solutions from within a 1- or multi-parameter family. It predicts an initial guess \tilde{z}_k of a family member, which is then corrected and solved for by a correction procedure, and continues to make predictions (see Figure 2.11). The method continues until either the convergence to a true solution cannot be accomplished or a specified number of family members have been computed.

The main components of a continuation method are

- (1) Prediction of next solution
- (2) Parameterization strategy of the family
- (3) Correction to a true solution
- (4) Control of step size between family members

Let $\mathbf{F} : \mathbb{R}^s \rightarrow \mathbb{R}^{s-1}$, then \mathcal{M} as defined in Equation (2.60) has dimension 1, meaning \mathbf{F} defines a 1-parameter family. Let z_0 be a true or approximate family member of \mathcal{M} . This solution could come from the linear subspaces, such as in Equations (2.21) and (2.29b), or by some other means. The first step in the continuation method is to predict an initial guess of the next family

member \tilde{z}_1 according to

$$\tilde{z}_{k+1} = z_k + \Delta s_k z'_k, \quad (2.68)$$

where $z'_k \in T_{z_k} \mathcal{M}$ (Equation (2.63)), and $\Delta s_k \in \mathbb{R}$ is a step size.

At the moment, if \tilde{z}_1 is corrected then there are an infinite number of solutions z_1 can be. We must add another equation to \mathbf{F} so that the correction procedure has a unique solution to hone in on. We introduce the pseudo-arclength equation

$$s(z_{k+1}) = \langle z_{k+1} - z_k, z'_k \rangle - \Delta s_k = 0, \quad (2.69)$$

so that for each Δs_k there is a unique $z_{k+1} = z_{k+1}(\Delta s_k)$. We note that $\langle \cdot, \cdot \rangle$ is an appropriate inner product. The pseudo-arclength constraint measures the distance traveled between solutions and compares it to the step size used to predict the next solution. Moreover, Equation (2.69) allows the branch of solutions to be ordered on a line by the map $\zeta(\gamma_k)$ with $\gamma_k = \sum_{k=0} \Delta s_k$ such that $\zeta(\gamma_k) = z_k$. This is what is called a parameterization strategy for the branch of solutions.

We save a discussion on the correction strategy for the next section, and for now assume that we can solve for z_{k+1} . After z_{k+1} has been found we would like to adjust Δs_{k+1} so that we may predict solutions which are close to solutions on the branch. We use the following step size controller presented by Seydel in [158],

$$\sigma = N_{\text{opt}}/N_j, \quad (2.70a)$$

$$\bar{\sigma} = \begin{cases} 0.5, & \sigma < 0.5 \\ \sigma, & 0.5 \leq \sigma \leq 2 \\ 2, & \sigma > 2 \end{cases}, \quad (2.70b)$$

$$\Delta s_{k+1} = \bar{\sigma} \Delta s_k, \quad (2.70c)$$

where N_{opt} is a user defined value which determines how many iterations the corrector should take to converge, and N_j is the number of iterations the corrector actually took to converge. This step size controller is performance based, and has the advantage of decreasing the step size when the

topology of \mathbf{F} is constantly changing. On the other hand, it can increase the step size when the topology of \mathbf{F} is flat so that computational time is not wasted.

2.6.2 Correction Procedure

We now turn attention to the correction procedure so that one may actually solve \mathbf{F} for a true solution. The correction procedure is dependent on the form of \mathbf{F} . Solving for solutions in a dynamical system amounts to solving a TPBVP. A TPBVP has the general form

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, t; \mathbf{q}), \quad t_0 < t < t_f \quad (2.71a)$$

$$\mathbf{g}(\mathbf{x}(t_0), \mathbf{x}(t_f), \mathbf{q}) = \mathbf{0}, \quad (2.71b)$$

where $\mathbf{q} \in \mathbb{R}^p$ are constants the trajectory $\mathbf{x} \in \mathbb{R}^d$ depends on. We are interested in finding a solution $\mathbf{z}(t) = [(\mathbf{x}(t))^T, \mathbf{q}^T]^T$ such that $\mathbf{x}(t)$ obeys the vector field \mathbf{f} , and \mathbf{z} satisfies the boundary constraints in Equation (2.71b). TPBVPs type are commonly solved using shooting methods, however collocation methods are commonly used when the problem at hand requires more care. This thesis uses shooting methods, so we will describe these methods.

2.6.2.1 Single-Shooting

A single-shooting method takes an initial guess of the solution $\mathbf{z}(t_0)$ and integrates it to time t_f . Then we have

$$\mathbf{x}(t_f) = \varphi_{t_f}(\mathbf{x}(t_0)), \quad (2.72)$$

and can remove the dependency on $\mathbf{x}(t_f)$. Single-shooting gets its name from the fact that $\mathbf{x}(t_0)$ is integrated in a “single shot” to the final time.

After the integration an updated guess is determined via a Newton’s method. Let \mathbf{z}^j be the j th guess to the solution $\mathbf{z}(t_0)$ which satisfies the boundary conditions. Then the update \mathbf{z}^{j+1} is given by

$$\mathbf{z}^{j+1} = \mathbf{z}^j - [D\mathbf{g}(\mathbf{z}^j)]^{-1}\mathbf{g}(\mathbf{z}^j). \quad (2.73)$$

The shooting process is repeated until $\|\mathbf{g}(\mathbf{z}^{j+1})\| \leq \varepsilon$. A good value of the tolerance level is $0 < \varepsilon \leq 1\text{e-}10$. A single-shooting method relies on a good initial guess and does not perform well when the dynamics are highly unstable. A way to reduce the effect of the instability and improve the convergence is to use a multiple-shooting method.

2.6.2.2 Multiple-Shooting

A multiple-shooting method breaks the trajectory $\mathbf{x}(t)$ into m independent arcs at the times $t_i, i = 0, \dots, m-1$. Each arc is defined by the initial point \mathbf{x}_i , and each arc is flowed from time t_i to t_{i+1} with the final points of each arc denoted by $\varphi_{\Delta t_{i+1}}(\mathbf{x}_i)$. Figure 2.12 contains an image of a single-shooting arc (a) broken into multiple-shooting arcs (b).




Figure 2.12: Depiction of a single shooting arc (a) broken into m multiple-shooting arcs (b).

Now the boundary constraint takes the form $\mathbf{g}(\mathbf{x}_0, \mathbf{x}_{m-1}, \mathbf{q}) = \mathbf{0}$. In addition to the boundary

conditions needing to be satisfied, it is also required that continuity constraints are satisfied. Let

$$\mathbf{z} = \begin{bmatrix} \mathbf{x}_0 \\ \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_{m-1} \\ \mathbf{q} \end{bmatrix},$$

then the continuity constraints $\mathbf{c} : \mathbb{R}^{md+p} \rightarrow \mathbb{R}^{(m-1)d}$ are given by

$$\mathbf{c}(\mathbf{z}) = \begin{bmatrix} \varphi_{\Delta t_1}(\mathbf{x}_0) - \mathbf{x}_1 \\ \vdots \\ \varphi_{\Delta t_{m-1}}(\mathbf{x}_{m-2}) - \mathbf{x}_{m-1} \end{bmatrix} = \mathbf{0} \quad (2.74)$$

Equation (2.74) ensures that the entire trajectory is continuous so that it obeys the vector field in Equation (2.71a). The vector \mathbf{z} is corrected until both Equations (2.71b) and (2.74) are satisfied. To find the updated solution with Newton's method it is necessary to solve the augmented system of nonlinear equations

$$\mathbf{G}(\mathbf{z}) = \begin{bmatrix} \mathbf{g}(\mathbf{x}_0, \mathbf{x}_{m-1}, \mathbf{q}) \\ \mathbf{c}(\mathbf{z}) \end{bmatrix} = \mathbf{0}. \quad (2.75)$$

Therefore, it is necessary to compute the derivative of Equation (2.74). The $(m-1)d \times (md+p)$ Jacobian matrix of the continuity constraints has the form

$$D\mathbf{c}(\mathbf{z}) = \begin{bmatrix} \Phi(t_1, t_0; \mathbf{x}_0) & -\mathbb{I}_s & \mathbf{0} & \dots & \dots & \mathbf{0} & \frac{\partial \varphi_{\Delta t_1}(\mathbf{x}_0)}{\partial \mathbf{q}} \\ \mathbf{0} & \Phi(t_2, t_1; \mathbf{x}_1) & -\mathbb{I}_s & \mathbf{0} & \dots & \vdots & \frac{\partial \varphi_{\Delta t_2}(\mathbf{x}_1)}{\partial \mathbf{q}} \\ \vdots & \mathbf{0} & \ddots & \ddots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \ddots & \Phi(t_{m-2}, t_{m-3}; \mathbf{x}_{m-3}) & -\mathbb{I}_s & \mathbf{0} & \frac{\partial \varphi_{\Delta t_{m-2}}(\mathbf{x}_{m-3})}{\partial \mathbf{q}} \\ \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} & \Phi(t_{m-1}, t_{m-2}; \mathbf{x}_{m-2}) & -\mathbb{I}_s & \frac{\partial \varphi_{\Delta t_{m-1}}(\mathbf{x}_{m-2})}{\partial \mathbf{q}} \end{bmatrix}, \quad (2.76)$$

where $\Phi(t_{i+1}, t_i; \mathbf{x}_i)$ is the STM corresponding to \mathbf{x}_i and evaluated from time t_i to t_{i+1} . Equation (2.73) is carried out until $\|\mathbf{G}(\mathbf{z}^{j+1})\| < \varepsilon$.

2.7 Optimization Theory

Optimization is widely used throughout the world, and most people reading this thesis have likely encountered it in some form in their coursework or otherwise in their life. The theory of optimization defines what a well-posed optimization problem is, and allows one to answer the question, “what is the best solution to my problem?”

Optimization problems can be grouped into many classes, but here we only consider the dichotomy of unconstrained and constrained problems. Unconstrained optimization considers the entire solution space, while constrained optimization presents restrictions on the solution space. Because of this, the solutions to an unconstrained and constrained problem can be vastly different even when the objective function is the same.

An optimization problem has the form

$$\begin{aligned} \min_{\mathbf{x} \in \mathbb{R}^d} \quad & f(\mathbf{x}) \\ \text{subject to} \quad & g_i(\mathbf{x}) = a_i, \quad i = 1, \dots, r \\ & h_j(\mathbf{x}) \leq b_j, \quad j = 1, \dots, s. \end{aligned} \tag{2.77}$$

We call the vector \mathbf{x} the optimization variables, the function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ the cost or objective function, the functions $g_i : \mathbb{R}^d \rightarrow \mathbb{R}$ the equality constraint functions, and the functions $h_i : \mathbb{R}^d \rightarrow \mathbb{R}$ the inequality constraint functions. When both r and s are equal to zero then we say that Problem (2.77) is an unconstrained problem. Otherwise, we call it a constrained problem.

We define the set of feasible solutions S to be

$$S = \begin{cases} \mathbb{R}^d, & r = s = 0 \\ \{\mathbf{x} \mid g_i(\mathbf{x}) = a_i, \quad h_j(\mathbf{x}) \leq b_j\}, & \text{otherwise} \end{cases}. \tag{2.78}$$

Any $\mathbf{x} \in S$ is called a feasible solution. Let $\mathbf{x}^* \in S$ be a feasible solution to Problem (2.77), and $U \subset S$ be an open set containing \mathbf{x}^* . We call \mathbf{x}^* a local minimum if $f(\mathbf{x}^*) \leq f(\mathbf{x})$ for any

$\mathbf{x} \in U$. Likewise, we call \mathbf{x}^* a global minimum if $f(\mathbf{x}^*) \leq f(\mathbf{x})$ for any $\mathbf{x} \in S$. In this thesis we are only concerned with cost functions and constraint functions that are continuously differentiable on U , leading to the first order optimality conditions: if \mathbf{x}^* is a local optimum point then $\nabla f(\mathbf{x}^*) = \mathbf{0}$. We are interested in finding local minima in the set of feasible solutions. We say that \mathbf{x}^* is a solution to Problem (2.77) if $\mathbf{x}^* \in S$ and \mathbf{x}^* is a local minimum.

There are many methods to solve Problem (2.77), and we present only one method to solve each the unconstrained and the constrained problem. We use the ideas of these methods to solve the formulated optimization problems in Chapter 6.

2.7.1 Unconstrained Optimization

When $r = s = 0$ Problem (2.77) reduces to

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}). \quad (2.79)$$

A classic method to solve Problem (2.79) is gradient descent. Its objective is to search along the direction of steepest descent to arrive at a point such that $\nabla f(\mathbf{x}) = \mathbf{0}$. It is not guaranteed to find a global optima, only a local optima. In its basic form gradient descent assumes that f is continuously differentiable on \mathbb{R}^d , and it ignores equality and inequality constraints. There are ways to improve the performance of gradient descent, such as using projected gradient descent or acceleration methods.

Let \mathbf{x}_k be a guess of the solution to Problem (2.79). The direction $\mathbf{d} = -\nabla f(\mathbf{x}_k)$ is the direction of steepest descent and is called the search direction. One could let the next iterate $\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{d}$, however performance can be improved when a line search is done in the search direction. That is, we compute the set

$$L_\alpha = \{f(\mathbf{y}(\gamma)) \mid \mathbf{y}(\gamma) = \mathbf{x}_k + \gamma\mathbf{d}, \gamma \in [0, \alpha]\} \quad (2.80)$$

over a specified range of values, and let \mathbf{x}_{k+1} be the point which attains the minimum value of f in the set L_α . In many situations there is no way to obtain an analytical expression for \mathbf{x}_{k+1} ,

and we rely on numerical procedures to find the next iterate. We call the procedure to find \mathbf{x}_{k+1} LINESEARCH, and leave the details for the procedure until Chapter 6.

An algorithm for gradient descent is given in Algorithm 1. We call the algorithm a Euclidean gradient descent (EGD) because the cost function depends only on the optimization variables. In contrast, Chapter 6 develops an optimization problem which depends on state variables \mathbf{x} and parameters \mathbf{q} , but only the parameters are used as the optimization variables.

Algorithm 1 Euclidean Gradient Descent

```

1: procedure EGD( $\mathbf{x}_0, f, \varepsilon$ )
2:    $k = 0$ 
3:   while  $\|\nabla f(\mathbf{x}_k)\| \geq \varepsilon$  do
4:      $\mathbf{d} = -\nabla f(\mathbf{x}_k)$ 
5:      $\mathbf{x}_{k+1} = \text{LINESEARCH}(f, \mathbf{d})$ 
6:      $k \leftarrow k + 1$ 
7:   end while
8:   return  $\mathbf{x}_k$ 
9: end procedure

```

We note that Newton's method can be used to solve unconstrained optimization problems. We have already seen Newton's method in Section 2.6.2.1 presented as a root-finding method. To solve Problem 2.79, Newton's method computes the gradient and Hessian of the cost function. For the optimization problems in Chapter 6 computing the Hessian is not desirable. However, as a root-finding method, Newton's method can be thought of as solving for the roots of the gradient of some unknown cost function. In Equation (2.73) the function \mathbf{g} can be associated to this gradient.

2.7.2 Constrained Optimization

If either r or s is not equal to 0 then we have a constrained problem. The constrained problem is fundamentally different from the unconstrained problem since \mathbf{x} is not allowed to roam around \mathbb{R}^d . Gradient descent on its own is not suitable for these problems since it may find a solution which satisfies $\nabla f(\mathbf{x}) = \mathbf{0}$, but is not in the feasible set. We need a way to ensure the solution lies inside the feasible set.

Popular methods to enforce the constraints are barrier methods and penalty methods. Both

methods solve a sequence of unconstrained optimization problems, called subproblems, with a modified cost function. Barrier methods, also called interior-point methods, start with an initial guess $\boldsymbol{x}_0 \in S$ and enforce that the solution to each subproblem lies inside S . This is done by imposing a large cost on feasible points that lie ever closer to the boundary of S . On the other hand, penalty methods can start with an initial guess outside the feasible set. A penalty is applied to points which lie outside the feasible set. The solution to each subproblem determines how the penalty changes. The penalty continues to change until the solution resides inside the feasible set.

Chapter 3

Study of the Earth-Moon L_2 Quasi-Halo Orbits

3.1 Introduction

Szebehely eloquently states in his comprehensive exposition on the restricted problem of three bodies, “the aim of dynamics is to characterize dynamical systems by describing the totality of motions and discussing their properties; to fulfill this purpose numerical integration is one of the powerful tools” [161]. He could not have had more foresight in his assertion 57 years ago when numerical explorations were just beginning. Each family of quasi-periodic orbits is unique and as such each have their own advantages and disadvantages for spacecraft missions. Computing a family in its entirety provides a global view of that family and allows for quick preliminary evaluation in the early stages of trajectory design. Only when families are computed as fully as possible can we gain a complete picture of the dynamics.

To put the context of this chapter into perspective a literature review of computing families of orbits and high-dimensional tori in astrodynamics is provided. First, a non-comprehensive review on the computation of 1-parameter families is presented, followed by a review on the computation of multi-parameter families. Lastly, a review on the work of 3+ dimensional invariant tori is provided.

Szebehely in [161] covered the history and the works of many researchers between the years 1897 and 1965 devoted to the computation of periodic orbits in the circular restricted three-body problem. In 1966, Farquhar computed halo orbits [64]. In 1973, Farquhar and Kamel used a Lindstedt-Poincaré series to compute a 1-parameter family of quasi-halo orbits about the L_2 Lagrange point in the Earth-Moon system with the addition of the lunar orbital eccentricity and the

Sun's gravitational field [63]. In the late 70s and early 80s Breakwell, Brown, and Howell further studied the 1-parameter families of halo orbits by employing numerical methods [37, 94, 96]. In 1987, Howell and Pernicka developed a numerical method to compute Lissajous orbits [97]. In 1997, Hénon presented a comprehensive list of periodic orbit families in the restricted three-body problem [86]. In 2000, Castellà and Jorba introduced a new method of computing quasi-periodic orbits and used it to explore 1-parameter families of Lissajous orbits about the dynamical equivalents of the Earth-Moon L_4 and L_5 points in the Bicircular Problem. Their method computed quasi-periodic orbits based on the Fourier coefficients of an invariant curve [101, 47]. In 2016, Baresi and Scheeres computed a 1-parameter family of 3-dimensional quasi-periodic orbits in the vicinity of the complex rotator 4179 Toutatis [19]. In 2021, McCarthy and Howell computed in the Earth-Moon circular restricted three-body problem a 1-parameter family of 3-dimensional quasi-halos with constant ω_0 and Jacobi energy fixed to the generating periodic orbit [130]. In 2021, Rosales et al. computed the dynamical equivalents to the 1-parameter family halo orbits in the Earth-Moon Bicircular Problem [148].

In 2001, Gómez and Mondelo used the method in [47] to compute 2-parameter families of Lissajous and quasi-halo orbits in the Earth-Moon circular restricted three-body problem by means of single-parameter continuation [75]. A similar analysis by the same authors was conducted in 2005 which looked at the same families but in Hill's problem [73]. In this same vein, in 2021, Haro and Mondelo in [83] computed the L_1 Lissajous family in the Earth-Moon system using a different method, called the parameterization method, to compute the family and compare it to the results in [75]. In this work they evaluated dynamical observables which indicate when tori break down at the end of the family. In 2018, Bosanac used a method similar to that of Gómez and Mondelo to compute 2-parameter families of 2-dimensional quasi-periodic orbits about resonant orbits in the Earth-Moon and Sun-Earth circular restricted three-body problem [29]. In 2022, Ming et al. computed a 2-parameter subset of the 3-parameter family of 3-dimensional quasi-distance retrograde orbits by fixing ω_0 to match the frequency of the generating periodic orbit, leaving two additional parametric constraints to fully constrain their problem [135]. They constructed

two pseudo-arclength constraints and devised a 2-parameter continuation method to explore the two additional parameters with single-parameter continuation. Their method performs the same function as the constant slope parametric constraint developed by Lujan and Scheeres in [124] by using the two pseudo-arclength equations to enforce a constant search direction in the continuation. Similarly, Baresi and Scheeres employed two pseudo-arclength constraints in [19], but they did not constrain the search direction. The direction was allowed to change throughout the continuation as two pseudo-arclength step sizes were adaptively updated.

In 1998, Jorba and Villanueva used normal form methods to compute 3-dimensional quasi-periodic orbits about a linearly stable periodic orbit of L_5 in the circular restricted three-body problem [109]. In 2005, Gabern and Jorba used normal form methods to compute 3-, 4-, and 5-dimensional quasi-periodic orbits about the dynamical equivalent of the Sun-Jupiter L_5 point subject to perturbations by Saturn and Uranus [69]. Likewise, in 2009, Jorba and Olmedo used a parameterization method to compute 3- and 4-dimensional quasi-periodic orbits of the dynamical equivalent of the Sun-Jupiter L_5 point subject to perturbations by Saturn, Uranus, Neptune, and Earth [104]. In these papers the authors showed off the ability to compute higher-dimensional invariant tori with their developed methods, however they did not study the 3-parameter families of invariant tori which densely fill the space, nor the 1- and 2-parameter families that are known to foliate the center manifold [107].

While many papers have published methods to compute quasi-periodic orbits there are various works which use these methods to study families of quasi-periodic orbits and how to efficiently utilize them for mission design. Some applications include trajectory design [19, 23, 68, 79], transfer design [28, 89, 90, 130, 131, 139], and formation flying [22, 20, 21, 88, 91, 92].

It is evident the most commonly studied families are 1-parameter families even though it is known that a higher density of orbits exist in 2- and 3-parameter families in astrodynamics [38, 107, 108]. One of the reasons is the increase in computational complexity which arises when moving from a 1-parameter family to a multi-parameter family. The increase in the dimension of the family means there are more dimensions to explore, and vastly more family members need to

be computed to fully explore the family. Additionally, the method to explore higher-parameter families becomes nontrivial since the search direction is no longer forward or backward along a line. Rather, there are infinite directions, resulting in infinite choices in the direction to explore. To get a clear picture of the dynamics it is necessary to understand ways in which to compute and study multi-parameter families.

Moreover, most studies which compute quasi-periodic orbits focus on 2-dimensional quasi-periodic orbits even when 3+ dimensional quasi-periodic orbits exist in their respective systems. One of the reasons is the increase in the number of points needed to represent the invariant surface which increases the number of integrations and storage space. However, access to machines with several cores should eliminate the barrier caused by the time required for integration since each point can be integrated independently of one another. In the near future there will be increased interest in computing invariant tori of dimension greater than 2 as a result of studying higher-fidelity multi-body systems. Therefore, it is necessary to develop tools now to compute higher-dimensional invariant tori. Doing so will assuredly lead to the discovery of new families of orbits, and subsequently more options for mission designers.

The literature review highlights the prevalence of research of 1-parameter families of orbits and reveals the lack of research on higher-parameter families of higher-dimensional tori. This chapter studies the use of various parametric constraints on their effectiveness to explore multi-parameter families. Due to the Lunar Gateway, the area near the Earth-Moon L_2 9:2 NRHO is of particular interest in the community [76, 30, 31, 54, 55, 163, 174, 32, 173, 141]. The halo orbit family is a 1-parameter family of periodic orbits about which 2- and 3-parameter families of quasi-halo orbits exist. Therefore, this chapter focuses on the 2- and 3-parameter Earth-Moon L_2 quasi-halo families in the circular restricted three-body problem to further understand the dynamical environment surrounding the Gateway. Gómez and Mondelo in [75] partially studied the 2-parameter family in this same system, however their study stopped short of the whole family. They found for certain energy values the Lissajous and quasi-halos are bounded by the planar Lyapunov orbit with the same energy level, but did not elaborate further on the family and their

characteristics. In the spirit of “describing the totality of motions and discussing their properties” this study paints a more complete picture of the family and examines them from a mission design perspective.

This chapter is organized as follows: Section 3.2 generalizes an existing quasi-periodic orbit computational scheme to compute n -dimensional quasi-periodic invariant tori. Section 3.3 describes a method to compute the amplitudes of an invariant surface. Section 3.4 overviews the halo orbit family and points out important aspects for the remaining sections. Section 3.5 presents findings for the 2-parameter family of quasi-halo orbits, while Section 3.6 presents findings on the 2- and 3-parameter family of quasi-halo orbits in the vicinity of the stable halo orbits.

3.2 Computation of n -Dimensional Quasi-Periodic Invariant Tori

The presentation of quasi-periodic orbits in Chapter 2 was given in a functional form, however in practice we cannot compute invariant tori this way. The function must be sampled at discrete points with a fine enough partition so that the approximate invariant torus is accurately represented. We use a flow map method to compute n -dimensional quasi-periodic invariant tori which reduces the problem to working with invariant surfaces of dimension $n - 1$. We present the method by Olikara and Scheeres [140], who originally presented the method for n -dimensional quasi-periodic invariant tori, in great detail, and in an algorithmic way. In this section we describe in detail the equations needed and how to compute them at a level which should be easy to implement on a computer. We present the equations as they would be used in a multiple-shooting. The single-shooting version is obtained as a special case of the multiple-shooting version and is not difficult to derive.

For ease of notation we present the equations assuming the dimension of the invariant surface is n , meaning the dimension of the invariant torus is $n + 1$. Let $\tau \in \mathcal{T}_{n+1}(\boldsymbol{\omega})$ be an $(n + 1)$ -dimensional invariant torus with frequencies $\boldsymbol{\omega}$, then $\tau : \mathbb{R}^{n+1} \rightarrow \mathbb{R}^d$. Let $\chi \in \mathcal{X}_{n+1}(\boldsymbol{\rho})$ be an invariant surface with the rotation vector $\boldsymbol{\rho}$ such that $\chi(\boldsymbol{\theta}) = \tau(0, \boldsymbol{\theta})$. Let $\mathbf{N} = (N_1, \dots, N_n) \in \mathbb{N}^n$ be a set of integers defining the number of points to be sampled in each dimension of \mathbb{T}^n , and

$\mathcal{N} = \prod_{i=1}^n N_i$. We define $\mathbf{j} = (j_1, \dots, j_n)$ to be a multi-index where each $j_i \in \{0, 1, \dots, N_i - 1\}$. Then we define $\boldsymbol{\theta}_{\mathbf{j}} = (\theta_{1,j_1}, \theta_{2,j_2}, \dots, \theta_{n,j_n}) = 2\pi(j_1/N_1, j_2/N_2, \dots, j_n/N_n) \in \mathbb{T}^n$, containing the points which partition the n -torus. Figure 3.1 shows an example of a partition on a flattened 2-dimensional torus with $N_1 = N_2 = 5$.




Figure 3.1: Partition of a flattened 2-dimensional torus with $N_1 = N_2 = 5$.

We stack all the points $\boldsymbol{\chi}(\boldsymbol{\theta}_{\mathbf{j}})$ into the vector $\mathbf{X} \in \mathbb{R}^{d\mathcal{N}}$ according to Algorithm 2. The order in which the points are stacked is important. If the points are stacked differently than what is shown here, then the resulting Fourier computations will not work, and modifications will be needed. In the algorithm the functions $\text{HCAT}(A, B) = [A|B]$ and $\text{VCAT}(A, B) = \begin{bmatrix} A \\ B \end{bmatrix}$ perform the function of horizontally and vertically concatenating the arrays A and B , respectively. The function $\text{LENGTH}(A)$ returns the number of elements in an array. The function $\text{ONES}(a, b; c)$ returns an $a \times b$ array filled with the value c . Lastly, the notation $a : b : c$ produces the column vector $[a, a + b, a + 2b, \dots, c]^T$, the notation $A = []$ creates an empty array, and $A[a, :]$ returns the a th row of A .

Algorithm 2 Stack Points of an Invariant Surface

```

1: procedure STACK( $\mathbf{N}, \chi(\boldsymbol{\theta})$ )
2:    $n \leftarrow \text{LENGTH}(\mathbf{N})$ 
3:    $J \leftarrow (1 : 1 : N_n)$ 
4:   for  $k$  in  $(n - 1 : -1 : 1)$  do
5:      $A \leftarrow []$ 
6:     for  $j$  in  $(1 : 1 : N_k)$  do
7:        $p \leftarrow \prod_{i=k+1}^n N_i$ 
8:        $B \leftarrow \text{HCAT}(\text{ONES}(p, 1; j), J)$ 
9:        $A \leftarrow \text{VCAT}(A, B)$ 
10:    end for
11:     $J \leftarrow A$ 
12:  end for
13:   $\mathbf{X} \leftarrow []$ 
14:   $p \leftarrow \prod_{i=1}^n N_i$ 
15:  for  $k$  in  $(1 : 1 : p)$  do
16:     $j \leftarrow J[1, :]$ 
17:     $\mathbf{X} \leftarrow \text{HCAT}(\mathbf{X}, \chi(\boldsymbol{\theta}_j))$ 
18:  end for
19:  return  $\mathbf{X}$ 
20: end procedure

```

As an example, suppose $\mathbf{N} = (2, 2, 3)$, then we have

$$\mathbf{X} = \begin{bmatrix} \chi(\theta_{1,1}, \theta_{2,1}, \theta_{3,1}) \\ \chi(\theta_{1,1}, \theta_{2,1}, \theta_{3,2}) \\ \chi(\theta_{1,1}, \theta_{2,1}, \theta_{3,3}) \\ \chi(\theta_{1,1}, \theta_{2,2}, \theta_{3,1}) \\ \chi(\theta_{1,1}, \theta_{2,2}, \theta_{3,2}) \\ \chi(\theta_{1,1}, \theta_{2,2}, \theta_{3,3}) \\ \chi(\theta_{1,2}, \theta_{2,1}, \theta_{3,1}) \\ \chi(\theta_{1,2}, \theta_{2,1}, \theta_{3,2}) \\ \chi(\theta_{1,2}, \theta_{2,1}, \theta_{3,3}) \\ \chi(\theta_{1,2}, \theta_{2,2}, \theta_{3,1}) \\ \chi(\theta_{1,2}, \theta_{2,2}, \theta_{3,2}) \\ \chi(\theta_{1,2}, \theta_{2,2}, \theta_{3,3}) \end{bmatrix} = \begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x}_3 \\ \mathbf{x}_4 \\ \mathbf{x}_5 \\ \mathbf{x}_6 \\ \mathbf{x}_7 \\ \mathbf{x}_8 \\ \mathbf{x}_9 \\ \mathbf{x}_{10} \\ \mathbf{x}_{11} \\ \mathbf{x}_{12} \end{bmatrix}.$$

For m multiple-shooting segments we need to sample $\boldsymbol{\tau}$ at m values of θ_0 . The sampling does not need to be uniform. Recall from Equation (2.31) we let $\mathbf{x}(t) = \mathbf{x}(T\frac{\theta_0}{2\pi})$, showing the time associated with each segment is $t_j = T\frac{\theta_{0,j}}{2\pi}$, and the time each segment needs to be propagated is $\Delta_{j+1} = t_{j+1} - t_j$. With this information we construct the invariant surfaces $\chi_j(\boldsymbol{\theta}) = \boldsymbol{\tau}(\theta_{0,j}, \boldsymbol{\theta})$. Then using Algorithm 2 we construct the stacked vectors $\bar{\mathbf{X}}_j$ and let

$$\mathbf{X} = \begin{bmatrix} \bar{\mathbf{X}}_0 \\ \vdots \\ \bar{\mathbf{X}}_{m-1} \end{bmatrix}_{[dm\mathcal{N} \times 1]} \quad (3.1)$$

We proceed by giving the equations and describing their purpose in the torus computations.

Let

$$\mathbf{z} = \begin{bmatrix} \mathbf{X} \\ T \\ \boldsymbol{\rho} \\ \boldsymbol{\omega} \end{bmatrix} \quad (3.2)$$

be the vector containing all the points on the invariant surface, the stroboscopic time, the rotation vector, and the frequency vector. The vector \mathbf{z} contains all the information necessary to fully define the invariant torus and construct the diffeomorphism between torus space and phase space. We note that including both (T, ρ) and $\boldsymbol{\omega}$ is redundant, however doing so is convenient. One can omit either set if desired and make the corresponding changes, but it is necessary to include $\boldsymbol{\omega}$ for the developments in Chapters 5 and 6.

3.2.1 Quasi-Periodicity

We have already seen the quasi-periodicity constraint in functional form in Equation (2.35). We now explain this equation from the perspective of a computer implementation. Equation (2.35) becomes

$$\mathbf{G}(\mathbf{z}) = R_{-\rho} \varphi_{\Delta t_m}(\bar{\mathbf{X}}_{m-1}) - \bar{\mathbf{X}}_0 = \mathbf{0}, \quad (3.3)$$

where $\varphi_{\Delta t_m}(\bar{\mathbf{X}}_{m-1})$ is a $d\mathcal{N} \times 1$ stacked vector where each of the \mathcal{N} points in the invariant surface have been propagated from time t_{m-1} to time T .

To understand the rotation operator $R_{-\rho}$ we first begin by describing a discrete Fourier transform (DFT) and its inverse. The reader is referred to the books of Amidror [4] and Rao et al. [145] to learn more about the DFT, its properties, and its uses. The DFT computes the Fourier coefficients from a discretely sampled function or signal. The inverse transform has the opposite effect as the DFT so that the sampled points are recovered in their original domain. There are several ways to compute the DFT, however we present one possible way here. Consider a single Fourier coefficient c_k of an invariant surface \mathbf{X} for $d = n = 1$ uniformly sampled at N_1 points so that $\theta_{1,j_1} = 2\pi \frac{j_1}{N_1}$, for $j_1 = 0, \dots, N_1 - 1$,

$$c_k = \sum_{j_1=0}^{N_1} \chi(\theta_{1,j_1}) e^{-2\pi i k \frac{j_1}{N_1}}. \quad (3.4)$$

We define $\mathbf{c} = [c_{-\frac{N_1-1}{2}}, c_{-\frac{N_1-1}{2}+1}, \dots, -1, 0, 1, \dots, c_{\frac{N_1-1}{2}-1}, c_{\frac{N_1-1}{2}}]^T$, then

$$\mathbf{c} = W_{N_1} \mathbf{X}, \quad (3.5)$$

where \mathbf{X} is the stacked vector of the points sampled from the invariant surface, and

$$W_{N_1} = \begin{bmatrix} w^0 & w^{-\frac{N_1-1}{2}} & w^{-2\frac{N_1-1}{2}} & \dots & w^{-(N_1-1)\frac{N_1-1}{2}} \\ w^0 & w^{-\frac{N_1-1}{2}+1} & w^{-2(\frac{N_1-1}{2}+1)} & \dots & w^{-(N_1-1)(\frac{N_1-1}{2}+1)} \\ \vdots & \vdots & \vdots & & \vdots \\ w^0 & w^0 & w^0 & \dots & w^0 \\ \vdots & \vdots & \vdots & & \vdots \\ w^0 & w^{\frac{N_1-1}{2}-1} & w^{2(\frac{N_1-1}{2}-1)} & \dots & w^{(N_1-1)(\frac{N_1-1}{2}-1)} \\ w^0 & w^{\frac{N_1-1}{2}} & w^{2\frac{N_1-1}{2}} & \dots & w^{(N_1-1)\frac{N_1-1}{2}} \end{bmatrix}, \quad (3.6)$$

with $w = e^{-i\frac{2\pi}{N_1}}$. The matrix W_{N_1} is the $N_1 \times N_1$ DFT to convert a signal to its Fourier coefficients.

The inverse DFT is given by

$$W^{-1} = \frac{1}{N_1} W_{N_1}^*, \quad (3.7)$$

where $(\cdot)^*$ is the complex conjugate transpose operator. The invariant surface in phase space is recovered by

$$\mathbf{X} = W_{N_1}^{-1} \mathbf{c}. \quad (3.8)$$

Now we derive the discrete form of the rotation operator $R_{-\rho}$. Recall the property of the rotation operator from Section 2.3.5

$$R_{-\rho_1} \chi(\theta_{1,j_1}) = \chi(\theta_{1,j_1} - \rho_1). \quad (3.9)$$

Expanding $\chi(\theta_{1,j_1} - \rho_1)$ as a Fourier series gives

$$\chi(\theta_{1,j_1} - \rho_1) = \frac{1}{N_1} \sum_{k=-\frac{N_1-1}{2}}^{\frac{N_1-1}{2}} c_k e^{ik(\frac{2\pi j_1}{N_1} - \rho_1)} \quad (3.10)$$

$$= \frac{1}{N_1} \sum_{k=-\frac{N_1-1}{2}}^{\frac{N_1-1}{2}} c_k e^{2\pi i k \frac{j_1}{N_1}} e^{-ik\rho_1} \quad (3.11)$$

Stacking all the $\chi(\theta_{1,j_1} - \rho_1)$ into a vector and writing the result in matrix form gives

$$R_{-\rho_1} \mathbf{X} = W_{N_1}^{-1} \text{diag}(e^{-i\mathbf{k}\rho_1}) \mathbf{c}, \quad (3.12)$$

where $\mathbf{k} = [-\frac{N_1-1}{2}, -\frac{N_1-1}{2} + 1, \dots, -1, 0, 1, \dots, \frac{N_1-1}{2} - 1, \frac{N_1-1}{2}]$ and $\text{diag}(\cdot)$ is a diagonal matrix.

We denote

$$Q_{-\rho_1} = \text{diag}(e^{-i\mathbf{k}\rho_1}) \quad (3.13)$$

the rotation matrix which rotates the Fourier coefficients by an angle $-\rho_1$. Using Equation (3.5) to replace \mathbf{c} in Equation (3.12) gives

$$R_{-\rho_1} \mathbf{X} = W_{N_1}^{-1} Q_{-\rho_1} W_{N_1} \mathbf{X}. \quad (3.14)$$

Therefore, the rotation operator is given by

$$R_{-\rho_1} = W_{N_1}^{-1} Q_{-\rho_1} W_{N_1}. \quad (3.15)$$

Extending to the case where $d > 1$ and $n > 1$ the DFT in Equation (3.6) becomes

$$W = W_{N_1} \otimes W_{N_2} \otimes \dots \otimes W_{N_n} \otimes \mathbb{I}_d, \quad (3.16)$$

where \otimes is the Kronecker product. The inverse DFT in Equation (3.7) becomes

$$W^{-1} = \frac{1}{\mathcal{N}} W^*, \quad (3.17)$$

the rotation matrix in Equation (3.13) becomes

$$Q_{-\rho} = \text{diag}(e^{-i\mathbf{k}\rho_1}) \otimes \text{diag}(e^{-i\mathbf{k}\rho_2}) \otimes \dots \otimes \text{diag}(e^{-i\mathbf{k}\rho_n}) \otimes \mathbb{I}_d. \quad (3.18)$$

Finally, the n -dimensional rotation operator for a d -dimensional phase space is given as

$$R_{-\rho} = W^{-1} Q_{-\rho} W. \quad (3.19)$$

For the correction process we need the derivative of $\mathbf{G}(\mathbf{z})$ in Equation (3.3) with respect to each variable in \mathbf{z} . The partial derivatives of \mathbf{G} are

$$\frac{\partial \mathbf{G}}{\partial \mathbf{X}} = \left[-\mathbb{I}_{dN} \quad \underbrace{\mathbf{0}_{dN \times dN} \dots \mathbf{0}_{dN \times dN}}_{m-2} \quad R_{-\rho} \Phi(T, t_{m-1}; \bar{\mathbf{X}}_{m-1}) \right] \quad (3.20)$$

$$\frac{\partial \mathbf{G}}{\partial T} = R_{-\rho} \mathbf{f}(\bar{\mathbf{X}}_{m-1}) \quad (3.21)$$

$$\frac{\partial \mathbf{G}}{\partial \rho} = \left[-\mathcal{D}_{\rho_1} R_{-\rho} \bar{\mathbf{X}}_{m-1} \quad -\mathcal{D}_{\rho_2} R_{-\rho} \bar{\mathbf{X}}_{m-1} \quad \dots \quad -\mathcal{D}_{\rho_n} R_{-\rho} \bar{\mathbf{X}}_{m-1} \right] \quad (3.22)$$

$$\frac{\partial \mathbf{G}}{\partial \omega} = \mathbf{0}_{dN \times n}, \quad (3.23)$$

where

$$\mathcal{D}_{\rho_1} = W^{-1}(\text{idia}g([1, \dots, N_1 - 1]) \otimes \mathbb{I}_{N_2} \otimes \dots \otimes \mathbb{I}_{N_n})W \quad (3.24a)$$

$$\mathcal{D}_{\rho_2} = W^{-1}(\text{i}\mathbb{I}_{N_1} \otimes \text{diag}([1, \dots, N_2 - 1]) \otimes \mathbb{I}_{N_3} \otimes \dots \otimes \mathbb{I}_{N_n})W \quad (3.24b)$$

$$\mathcal{D}_{\rho_i} = W^{-1}(\text{i}\mathbb{I}_{N_1} \otimes \mathbb{I}_{N_2} \otimes \dots \otimes \text{diag}([1, \dots, N_i - 1]) \otimes \mathbb{I}_{N_{i+1}} \otimes \dots \otimes \mathbb{I}_{N_n})W, \quad i = 3, \dots, n - 2 \quad (3.24c)$$

$$\mathcal{D}_{\rho_{n-1}} = W^{-1}(\text{i}\mathbb{I}_{N_1} \otimes \dots \otimes \mathbb{I}_{N_{n-3}} \otimes \text{diag}([1, \dots, N_{n-1} - 1]) \otimes \mathbb{I}_{N_n})W \quad (3.24d)$$

$$\mathcal{D}_{\rho_n} = W^{-1}(\text{i}\mathbb{I}_{N_1} \otimes \dots \otimes \mathbb{I}_{N_{n-1}} \otimes \text{diag}([1, \dots, N_n - 1]))W. \quad (3.24e)$$

3.2.2 Continuity Constraints

The continuity constraints arise from using a multiple-shooting method. These equations are given in Equation (2.74). The derivatives of the continuity constraints are in Equation (2.76). Here we elaborate on the partial derivatives in the last column of Equation (2.76) which are the partial derivatives of the flow with respect to the parameters in the vector \mathbf{z} . The parameters are T , $\boldsymbol{\rho}$, and $\boldsymbol{\omega}$. Without knowing the form of the dynamical system most of these derivatives must be left general, however the stroboscopic time will play a role in every dynamical system, so we give the partial derivative of the flow with respect to T in Equation (3.27).

Each Δt_i makes up some portion of the stroboscopic time such that

$$T = \sum_{i=1}^m \Delta t_i = T \sum_{i=1}^m a_i. \quad (3.25)$$

It is clear that the a_i 's must sum to 1 for the equation to hold. Moreover, it is clear that

$$\Delta t_i = a_i T. \quad (3.26)$$

Therefore, taking the partial derivative of the flow with respect to T gives

$$\frac{\partial \varphi_{\Delta t_i}(\bar{\mathbf{X}}_{i-1})}{\partial T} = \frac{1}{a_i} \mathbf{f}(\bar{\mathbf{X}}_{i-1}), \quad i = 1, \dots, m - 1. \quad (3.27)$$

In Equation (3.27) the term $\mathbf{f}(\bar{\mathbf{X}})$ is a $d\mathcal{N} \times 1$ stacked vector which evaluates the vector field at each point in the invariant curves and for each multiple-shooting segment. It is assumed that each point is evaluated at its respective time. For a uniform sampling of m segments we have $a_i = \frac{1}{m}$.

3.2.3 Phase Constraints

In Section 2.3.5 we state any invariant surface with stroboscopic time T and rotation vector $\boldsymbol{\rho}$ will satisfy the quasi-periodicity constraint, i.e. if $\boldsymbol{\tau} \in \mathcal{T}_{n+1}(\boldsymbol{\omega})$ is a torus function with incommensurate frequencies $\boldsymbol{\omega}$ then any $\boldsymbol{\chi} \in \mathcal{X}_n(T, \boldsymbol{\rho})$ satisfies Equation (2.35). The constraint is indifferent to the choice of parameterization on the torus. Therefore, there are an infinite number of solutions to the quasi-periodicity constraint. To resolve this issue phase constraints are added to the system of equations. Phase constraints fix the parameterization on the torus. Suppose $(\tilde{\mathbf{X}}, \tilde{T}, \tilde{\boldsymbol{\rho}})$ is a solution which satisfies Equation (3.3). This solution is typically either generated from a linear approximation according to Equation (2.29b) or is the previous solution in a continuation method. The vector $\tilde{\mathbf{X}}$ necessarily defines a parameterization over the torus. Then the phase constraints ensure that the next solution $(\mathbf{X}, T, \boldsymbol{\rho})$ to be computed has the same parameterization by finding the solution which is closest to $\tilde{\mathbf{X}}$.

The phase constraints are given as

$$\mathbf{p}(\mathbf{z}) = \begin{bmatrix} p_0(\mathbf{z}) \\ \vdots \\ p_n(\mathbf{z}) \end{bmatrix} = \begin{bmatrix} \left\langle \mathbf{X} - \tilde{\mathbf{X}}, \frac{\partial \tilde{\mathbf{X}}}{\partial \theta_0} \right\rangle \\ \vdots \\ \left\langle \mathbf{X} - \tilde{\mathbf{X}}, \frac{\partial \tilde{\mathbf{X}}}{\partial \theta_n} \right\rangle \end{bmatrix} = \mathbf{0}, \quad (3.28)$$

where

$$\frac{\partial \tilde{\mathbf{X}}}{\partial \theta_i} = \mathcal{D}_{\rho_i} \tilde{\mathbf{X}}, \quad i = 1, \dots, n \quad (3.29a)$$

$$\frac{\partial \tilde{\mathbf{X}}}{\partial \theta_0} = \frac{1}{\omega_0} \left(\mathbf{f}(\tilde{\mathbf{X}}) - \sum_{i=1}^n \omega_i \frac{\partial \tilde{\mathbf{X}}}{\partial \theta_i} \right). \quad (3.29b)$$

The number of phase constraints and which phase constraints are included in the system of equations to solve in the TPBVP depends on the dynamical system in which invariant tori are being computed. For systems which explicitly depend on ω_i or θ_i the corresponding phase constraint is not included. In a system of this type the dependence on the angle or frequency fixes that coordinate. In an autonomous Hamiltonian system all n phase constraints are needed. Suppose the periodic forcing of Jupiter is added to the system. Then that forcing frequency becomes a frequency of any

invariant torus in the system, and the phase constraint corresponding to that coordinate is removed from the system of equations.

The partial derivative of each p_i with respect to \mathbf{X} is given by

$$\frac{\partial p_i}{\partial \mathbf{X}} = \frac{\partial \tilde{\mathbf{X}}}{\partial \theta_i}. \quad (3.30)$$

The remaining partial derivatives are zero.

3.2.4 Consistency Constraints

In Section 3.2 we mentioned the redundancy between including both $(T, \boldsymbol{\rho})$ and $\boldsymbol{\omega}$ as independent variables, since they are correlated through Equations (2.32) and (2.33). To address this issue we add consistency constraints to ensure their correspondence between the stroboscopic time, rotation vector, and frequency vector.

The consistency constraints are

$$\mathbf{k}(\mathbf{z}) = \begin{bmatrix} T\omega_0 - 2\pi \\ T\omega_1 - \rho_1 \\ \vdots \\ T\omega_n - \rho_n \end{bmatrix} = \mathbf{0} \quad (3.31)$$

The derivatives of the consistency constraints are straightforward to compute.

3.2.5 Parametric Constraints

The quasi-periodicity constraint, the phase constraints, and the consistency constraints are sufficient to compute an orbit within a p -parameter family of quasi-periodic orbits, however to enable single-parameter continuation we must include p additional constraints. We call these constraints parametric constraints as they allow us to identify a single member from within a family, thus parameterizing the family. For a Hamiltonian system with d degrees of freedom the maximum size a family can be is a d -parameter family. In autonomous Hamiltonian systems, such as the circular restricted three-body problem, p -dimensional invariant tori come in p -parameter families.

Consider, again, that the periodic forcing of Jupiter is added to the system. Then $p+1$ -dimensional quasi-periodic invariant tori form p -parameter families.

The first and most important parametric constraint in continuation is the pseudo-arclength constraint. This constraint was encountered in Equation (2.69), however we amend it to be

$$s(\mathbf{z}) = \frac{1}{m\mathcal{N}} \langle \mathbf{X} - \tilde{\mathbf{X}}, \tilde{\mathbf{X}}' \rangle + (T - \tilde{T})\tilde{T}' + \langle \boldsymbol{\rho} - \tilde{\boldsymbol{\rho}}, \tilde{\boldsymbol{\rho}}' \rangle + \langle \boldsymbol{\omega} - \tilde{\boldsymbol{\omega}}, \tilde{\boldsymbol{\omega}}' \rangle - \Delta s = 0. \quad (3.32)$$

Recall that $(\cdot)'$ is a tangent space vector. We often call this vector a family tangent vector when using continuation because it predicts how to approximate the next family member in a branch. The method to generate a family tangent vector is given in Section 3.2.6.

The remaining parametric constraints $\mathbf{g}(\mathbf{z}) = [g_1(\mathbf{z}), \dots, g_{p-1}(\mathbf{z})]$ can take on any form as long as they construct a full rank Jacobian matrix in the resulting system of equations during the entire continuation procedure. These equations add flexibility in the types of solutions obtained in the continuation procedure.

3.2.6 Family Tangents

Putting all the equations together forms the system of equations $\mathbf{F} : \mathbb{R}^{dm\mathcal{N}+2(n+1)} \rightarrow \mathbb{R}^{dm\mathcal{N}+2(n+1)+p}$,

$$\mathbf{F}(\mathbf{z}) = \begin{bmatrix} \mathbf{G}(\mathbf{z}) \\ \mathbf{c}(\mathbf{z}) \\ \mathbf{p}(\mathbf{z}) \\ \mathbf{k}(\mathbf{k}) \\ s(\mathbf{z}) \\ \mathbf{g}(\mathbf{z}) \end{bmatrix}. \quad (3.33)$$

We stress that the exact form of \mathbf{F} depends on the dynamical system, the dimension of the tori being computed, and the dimension of the orbit family. Equation (3.33) gives the full system of equations which defines a single quasi-periodic orbit from within a family of quasi-periodic orbits. However, for the continuation we need the tangent space of the 1-parameter family in order to compute

branches of solutions. We recover the systems of equations defining the branch of solutions by removing the pseudo-arclength constraint, giving

$$\mathbf{F}_1(\mathbf{z}) = \begin{bmatrix} \mathbf{G}(\mathbf{z}) \\ \mathbf{c}(\mathbf{z}) \\ \mathbf{p}(\mathbf{z}) \\ \mathbf{k}(\mathbf{k}) \\ \mathbf{g}(\mathbf{z}) \end{bmatrix}. \quad (3.34)$$

Equation (3.34) forms an implicitly defined smooth embedded submanifold (in the sense of Whitney) of dimension 1. For other reasons we are interested in the tangent space of the full family of quasi-periodic orbits. Removing the remaining parametric constraints gives the system of equations

$$\mathbf{F}_p(\mathbf{z}) = \begin{bmatrix} \mathbf{G}(\mathbf{z}) \\ \mathbf{c}(\mathbf{z}) \\ \mathbf{p}(\mathbf{z}) \\ \mathbf{k}(\mathbf{k}) \end{bmatrix}. \quad (3.35)$$

Equation (3.35) forms an implicitly defined smooth embedded submanifold (in the sense of Whitney) of dimension p .

The initialization of the continuation method requires an initial or approximate solution \mathbf{z} and a family tangent vector \mathbf{z}' to provide an approximation of the next family member. If the solution is an approximate solution derived from the linearization about a periodic orbit, then the family tangent vector is generated from Equation (2.29a) as a stacked vector according to Algorithm 2 with $T' = 0$, $\boldsymbol{\rho}' = \mathbf{0}$, and $\boldsymbol{\omega}' = \mathbf{0}$. In this case we choose the step size to be quite small since the family tangent vector is likely not pointing in the direction to satisfy the parametric constraints \mathbf{g} . If the solution is a true solution satisfying \mathbf{F}_1 , then a family tangent vector is the basis vector of the 1-dimensional null space of the Jacobian matrix.

Once the continuation procedure is underway and has computed at least the first true solution \mathbf{F} then the family tangent vector can be computed as the difference between two consecutive solutions (Equation (3.36)) or, again, come from the null space of the Jacobian matrix of \mathbf{F}_1 .

$$\tilde{z}' = z - \tilde{z} \quad (3.36)$$

3.2.7 Stability Analysis

Once a quasi-periodic orbit has been computed we can easily solve the discretized version of the eigenvalue problem in Equation (2.45). Given the discretized invariant surface $\mathbf{X} \in \mathbb{R}^{dm\mathcal{N}}$ of an invariant torus with frequencies $\boldsymbol{\omega}$ we compute the stability matrix

$$G_{\boldsymbol{\omega}} = R_{-\rho}\Phi, \quad (3.37)$$

where Φ is the $d\mathcal{N} \times d\mathcal{N}$ block diagonal matrix containing the STM of each point in $\bar{\mathbf{X}}_0$ evaluated over one stroboscopic time.

To find the Floquet multipliers and the corresponding eigenvectors we compute the eigenvalues and eigenvectors of $G_{\boldsymbol{\omega}}$. Just as in the continuous case, the eigenvalues in the discretized case form rings in the complex plane, and the same rules apply for determining the type of motion each ring corresponds to. Using the method of Jorba in [102] we filter and sort the eigenvalues to find the Floquet multipliers λ . The stable and unstable manifolds are initialized in the same manner as in Equation (2.53). It is important to note the eigenvectors corresponding to λ are $d\mathcal{N}$ -element vectors with each d set of values corresponding to the tangent direction of the invariant manifold at the corresponding point in $\bar{\mathbf{X}}_0$. We call each d set of values a fiber of the invariant manifold, whereas the collection of all the fibers is called a bundle. When the invariant manifold is a (un)stable manifold we call the bundle the (un)stable bundle, and for a center manifold we call the bundle the center bundle.

3.3 Amplitudes of an Invariant Surface

The size of a quasi-periodic orbit is not a well-defined quantity. There are different ways to quantify the size of a quasi-periodic orbit. For 2-d quasi-halo orbits one might be interested in the average distance from the surface of the orbit to the halo orbit. However, even this is not well-

defined. The distance could be measured from the quasi-halo orbit to the closest point on the halo orbit, or the distance could be measured to the point on the halo orbit which corresponds to the same θ_0 value. Moreover, this metric is only defined when the stroboscopic time of the quasi-halo shares the same value with the period of the halo orbit. In other dynamical systems there may not be any periodic orbits to compare the quasi-periodic orbits with. Therefore, we present a method to measure the amplitudes of each dimension of an n -dimensional quasi-periodic orbit, irrespective of any external point of reference, so that we may quantify the size of a quasi-periodic orbit.

An algorithm is given to compute the amplitudes of an n -dimensional discretized invariant surface \mathbf{X} . First, given a set of integers $(N_1, \dots, N_k) \in \mathbb{N}^k$ that form an evenly spaced grid over \mathbb{T}^k define the multi-index $\mathbf{j} = (j_1, \dots, j_k)$ that belongs to the set $J_k \equiv \{\mathbf{j} \in \mathbb{Z}^k | 0 \leq j_i < N_i \text{ for } i = 1, \dots, k\}$. Then define the vector of angles $\boldsymbol{\theta}_{\mathbf{j}} = (\theta_{j_1}, \dots, \theta_{j_k}) = 2\pi(j_1/N_1, \dots, j_k/N_k) \in \mathbb{T}^k$. Then $\mathbf{X}(\boldsymbol{\theta}_{\mathbf{j}})$ is a single point on the invariant surface, and the collection of points representing a k -dimensional invariant surface is $\mathbf{X} = \{\mathbf{X}(\boldsymbol{\theta}_{\mathbf{j}}) | \mathbf{j} \in J_k\}$. The notation $\mathbf{Y}_{\mathbf{j}}(\boldsymbol{\theta}_k) = \{\mathbf{X}(\boldsymbol{\theta}_{\mathbf{j}}, \theta_{k,l}) | l = 0, \dots, N_k - 1\}$ represents a curve along the k^{th} dimension of the invariant surface with fixed angles $\boldsymbol{\theta}_{\mathbf{j}}$ where $\mathbf{j} \in J_{k-1}$ and $\mathbf{Y}_{\mathbf{j}}(\theta_{k,i})$ is a single point on the curve at $\theta_{k,i} = 2\pi i/N_k$.

Algorithm 3 Amplitudes of an Invariant Surface

Require: $\mathbf{X}, N_1, \dots, N_n$

- 1: **for** k in $(n : -1 : 2)$ **do**
 - 2: **for all** $\mathbf{j} \in J_{k-1}$ **do**
 - 3: $\mathbf{c}_{k,\mathbf{j}} \leftarrow \frac{1}{N_k} \sum_{i=0}^{N_k-1} \mathbf{Y}_{\mathbf{j}}(\theta_{k,i})$
 - 4: $r_{k,\mathbf{j}} \leftarrow \frac{1}{N_k} \sum_{i=0}^{N_k-1} \|\mathbf{Y}_{\mathbf{j}}(\theta_{k,i}) - \mathbf{c}_{k,\mathbf{j}}\|$
 - 5: **end for**
 - 6: $A_k \leftarrow \frac{1}{\prod_{i=1}^{k-1} N_i} \sum_{\mathbf{j} \in J_{k-1}} r_{k,\mathbf{j}}$
 - 7: $\mathbf{X} \leftarrow \{\mathbf{c}_{k,\mathbf{j}} | \mathbf{j} \in J_{k-1}\}$
 - 8: **end for**
 - 9: $\mathbf{c}_1 \leftarrow \frac{1}{N_1} \sum_{i=0}^{N_1-1} \mathbf{X}(\theta_{1,i})$
 - 10: $r_1 \leftarrow \frac{1}{N_1} \sum_{i=0}^{N_1-1} \|\mathbf{X}(\theta_{1,i}) - \mathbf{c}_1\|$
 - 11: $A_1 \leftarrow r_1$
 - 12: **return** A_1, \dots, A_n
-

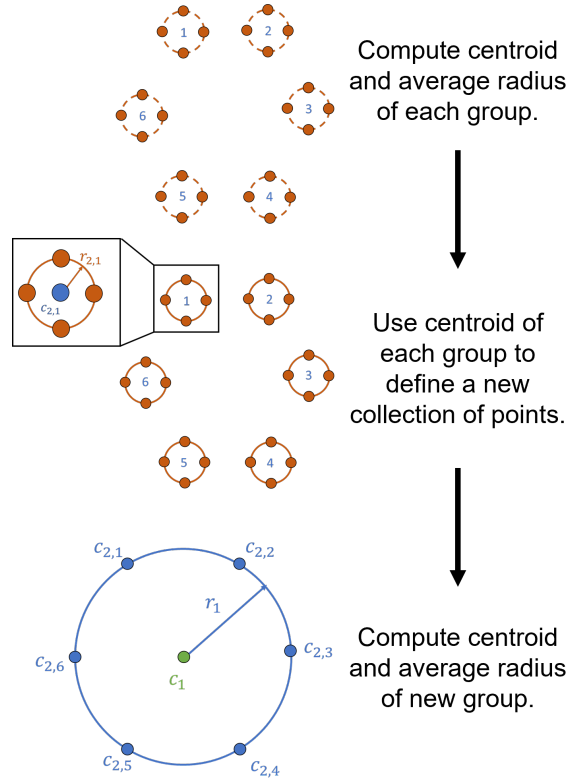


Figure 3.2: Example of computing amplitudes of an invariant surface with $n = 2$, $N_1 = 6$, and $N_2 = 4$.

The process to compute the amplitudes of an n -dimensional invariant surface ($k = n$) computes the amplitudes in reverse order starting with A_n . By freezing the first $n - 1$ dimensions, a group of points is defined and forms a curve $\mathbf{Y}_j(\boldsymbol{\theta}_n)$. A centroid $c_{n,j}$ and average radius $r_{n,j}$ is computed from this curve and are recorded. The centroid and average radius are computed for each curve $\mathbf{Y}_j(\boldsymbol{\theta}_n)$ with $j \in J_{n-1}$. The amplitude A_n is the average of the radius of each curve. The set of points X is redefined to be the set of centroids $\{c_{n,j} | j \in J_{n-1}\}$. This process has effectively reduced the dimension of the invariant surface by one by averaging out the last dimension. The process continues until all amplitudes have been computed. A 3-d quasi-periodic invariant torus has an invariant surface that is a 2-dimensional, and as such, it has 2 amplitudes A_1 and A_2 . The amplitude A_0 of the 3-d quasi-halo orbits correspond approximately to the sizes of the halo orbits, so these amplitudes have been omitted. However, it is straightforward to compute A_0 ; one simply

needs to let \mathbf{X} be a discretized set of points covering the entire invariant torus. Figure 3.2 depicts the process of computing the amplitudes for a 2-d invariant surface.

3.4 The Halo Orbit Family

We are interested in the Earth-Moon system, so the value of μ used throughout this thesis is 0.012153599037880. The L_2 halo orbit family serve as the generating orbits for the L_2 quasi-halo orbit family, and their structure is important to the structure of the quasi-halo orbit family. Quasi-halo orbits inherit the stability properties of their generating halo orbit.

The halo orbits are a bifurcation from the planar Lyapunov orbits at the point where the frequency of the Lyapunov orbit matches the frequency of the center manifold emanating from the orbit. This center manifold produces out of plane oscillations which in turn gives rise to the 3-dimensional geometry of the halo orbits. From the bifurcation point the halo orbits grow larger in size becoming more vertical while moving closer to the Moon forming a continuous 1-parameter family.

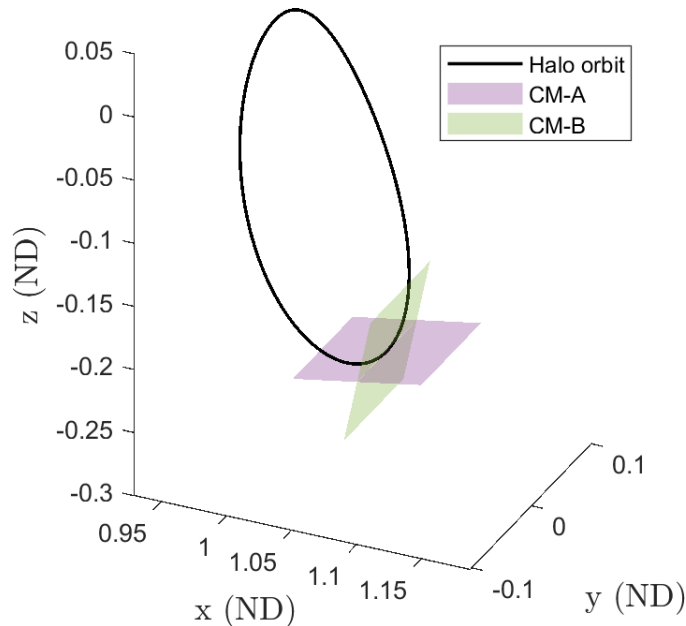


Figure 3.3: A stable halo orbit with the linear center manifolds from its monodromy matrix.




Figure 3.4: The frequencies of the center manifolds of the halo orbit family.

From the Lyapunov bifurcation the halo orbits start out as type *center* × *saddle*. Quasi-periodic motion is created by exciting the normal center manifold. Moving through the family there is a bifurcation and the halo orbits become type *saddle* × *saddle*. The region of orbits that have this unstable behavior is very small, and the family quickly reverts to type *center* × *saddle*. Continuing through the family there is another bifurcation and the orbits become stable with type *center* × *center*. The stable region was first discovered by Breakwell and Brown in [37], and the two normal center manifolds are visualized in Figure 3.3. In this region 2- and 3-dimensional quasi-halo orbits exist. Continuing through the family the stability is lost, reverting to type *center* × *saddle*, and remains this way until getting very close to the Moon; at which point the halo orbits become stable again. However, these stable orbits have a perilune (point of closest approach to the Moon) below the surface of the Moon, so they are not typically considered for operational applications.

The frequencies of the normal center manifolds are plotted against the frequencies of the halo orbits in Figures 3.4 and 3.5. The units of the frequencies are radians per non-dimensional time unit (rad/TU). The existence of two distinct center manifolds results in two distinct lines in the figure. The first center manifold is referred to as Center Manifold A (CM-A) while the second one

is referred to as Center Manifold B (CM-B). These lines are referred here as the halo lines since they represent the halo orbit family in frequency space. It is these halo lines that are used as the starting boundaries to present the frequencies of the quasi-halo family.

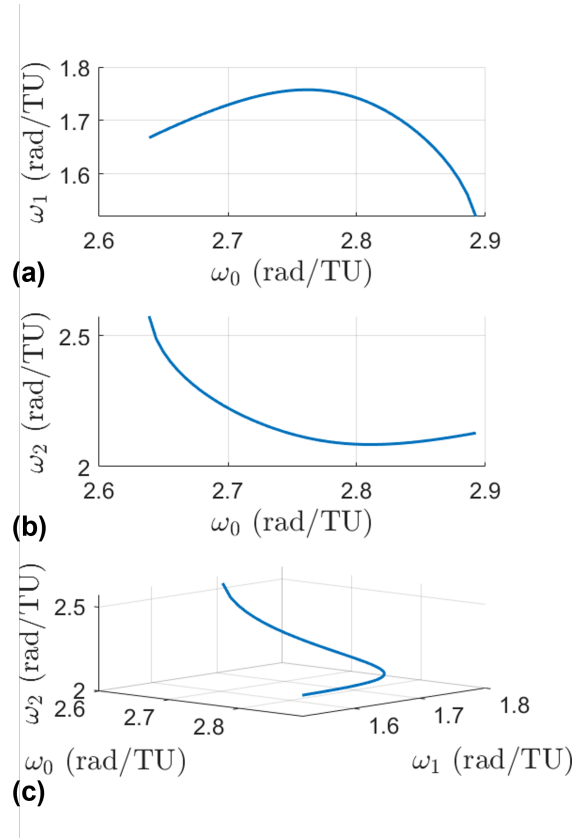


Figure 3.5: Frequencies of the stable L_2 halo orbit family and the frequencies of their normal center subspaces.

Halo orbits have a base frequency ω_0 (see Equation (2.32)). 2-dimensional quasi-halo orbits have an additional frequency ω_1 , while 3-dimensional quasi-halo orbits have two additional frequencies ω_1 and ω_2 . These frequencies arise from the normal center manifolds. The frequency ω_0 dictates the frequency of moving along the periodic orbit, while the frequencies ω_1 and ω_2 dictate the frequencies of oscillations about the halo orbit.

Figure 3.6 shows the halo lines in conjunction with the halo orbit family to put the lines into perspective. It is important to note that both plots represent the halo orbit family. The left plot

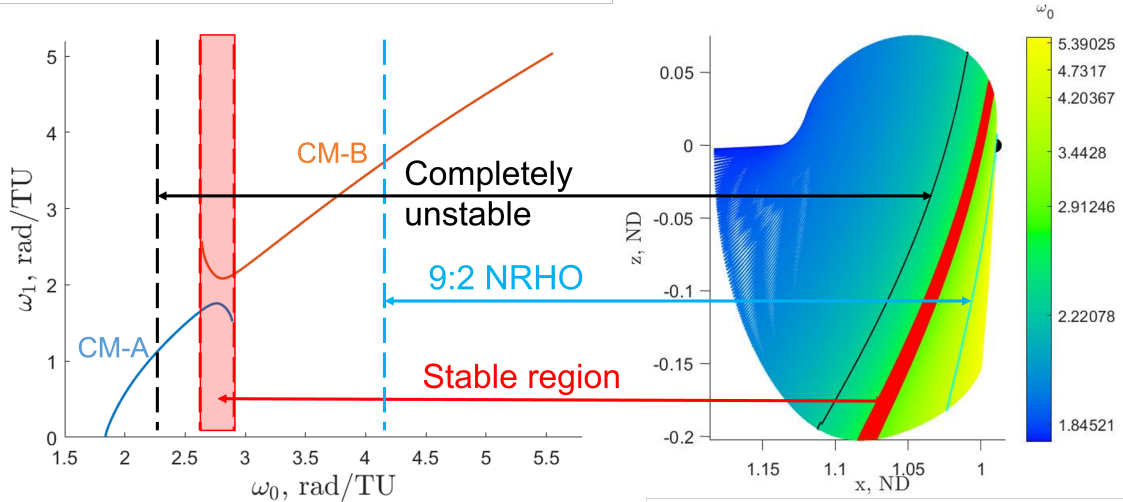


Figure 3.6: Correspondence between the halo orbit family in frequency space (left) and in coordinate space (right).

represents the family in frequency space while the right plot represents the family in coordinate space in an x - z view. Notable regions of the family are indicated on both plots and the orbits are colored according to their frequency ω_0 . One halo orbit to take note of is the 9:2 NRHO because this is where NASA plans to stage the Lunar Orbital Platform Gateway as a proving ground to test technologies that could enable future manned missions to Mars [169].

3.5 The 2-Parameter Family of Quasi-Halo Orbits

We have computed 500 halo orbits which span from the genesis of the family at the point of the planar Lyapunov stability bifurcation to the point where the orbits have perilune below the lunar surface. This provides numerous starting points from which to compute quasi-halo orbits. In the majority of the computation we used 41 points to discretize the invariant curves. A mix of single-shooting and multiple-shooting was used in the computation of quasi-halo orbits.

3.5.1 Exploring the Family

To understand efficient methods to explore families of quasi-periodic orbits we test the use of four different parametric constraints. The tested parametric constraints are given in Equations

(3.38) -(3.41). In the equations the symbol $(\cdot)^*$ represents a constant quantity. Equations (3.38)-(3.40) maintains a constant characteristic of the computed branch of orbits. These are the values of ω_0 , ω_1 , and the Jacobi energy, respectively. In Equation (3.40), $\bar{J}(\mathbf{X})$ averages the Jacobi constant among all points in \mathbf{X} . Equation (3.41) does not maintain a constant value of any orbital characteristic, but it ensures the branch of solutions follows straight line in frequency space at a specified slope m .

$$\omega_0 - \omega_0^* = 0 \quad (3.38)$$

$$\omega_1 - \omega_1^* = 0 \quad (3.39)$$

$$\bar{J}(\mathbf{X}) - J^* = 0 \quad (3.40)$$

$$\frac{\omega_1 - \omega_1^*}{\omega_0 - \omega_0^*} - m = 0 \quad (3.41)$$

It is easy to see that Equation (3.41) generalizes both Equations (3.38) and (3.39). Moreover, Equation (3.41) generalizes the constant rotation number constraint utilized by various researchers [29, 75, 83, 130]. To see this let $m = \frac{\omega_1^*}{\omega_0^*} = \frac{\rho_1^*}{2\pi}$. Then Equation (3.41) becomes

$$\begin{aligned} \frac{\omega_1 - \omega_1^*}{\omega_0 - \omega_0^*} &= \frac{\rho_1^*}{2\pi}, \\ \frac{2\pi\omega_1 - 2\pi\omega_1^*}{\omega_0 - \omega_0^*} &= \rho_1^*, \\ \frac{\rho_1\omega_0 - \rho_1^*\omega_0^*}{\omega_0 - \omega_0^*} &= \rho_1^*, \\ \rho_1\omega_0 - \rho_1^*\omega_1^* &= \rho_1^*\omega_0 - \rho_1^*\omega_0^*, \\ \rho_1\omega_0 &= \rho_1^*\omega_0 \\ \rho_1 &= \rho_1^*. \end{aligned}$$

In each of the parametric constraints we let the fixed value of the characteristic come from the generating halo orbit. The base frequency ω_0^* and the Jacobi constant J^* are computed directly from the halo orbit, while ρ_1^* and ω_1^* come from the normal center manifold. In the stable region there are two normal center manifolds, so we compute 2-dimensional quasi-halo orbits off of each center

manifold and use each center manifold's linear rotation number and frequency. For the constant slope constraint we chose $m = \frac{\omega_1^*}{\omega_0^*}$ so that constant rotation number branches were computed.




Figure 3.7: Quasi-halo orbits computed using different parametric constraints: constant ω_0 (a), constant ω_1 (b), constant Jacobi energy (c), and constant slope (d).

The results of using each parametric constraint are presented in Figure 3.7. The parametric constraint that explored the most area is the constant slope constraint for the case that each computed branch is keeping the rotation number constant. The reason for this is that each branch follows a straight line with slope $m = \frac{\omega_0}{\omega_1}$ to the origin in frequency space. In this space resonances are also straight lines intercepting the origin, however resonances only take on rational values for the slope. As long as a branch has a frequency ratio far enough away from a rational value that branch will be unaffected by the computational difficulties encountered near resonances. Gómez and Mondelo use a constant rotation number constraint and choose generating halo orbits with

particular rotation numbers to avoid computing branches near resonance lines. We did not compute halo orbits in this fashion so that we can study how well this constraint performs in more general cases.

The parametric constraint that explored the least area is the constant ω_1 constraint. This is because a constant ω_1 branch encounters resonances very quickly. Moreover, for the case of the L_2 quasi-halo orbits the orientation between a constant ω_1 branch and a resonance line is closer to parallel than it is to perpendicular, implying there is more “distance” to travel to get to the other side of a resonance line. Despite this the constant ω_1 constraint did well to fill out the bottom left portion of the family than did the constant slope constraint.

The constant ω_0 constraint performs well to explore the family, and even outperforms the constant slope constraint in the top right portion of the family, but falls short of when encountering resonances. This behavior is easily seen by the hard cutoffs and discontinuity between the branches, showing that it is easier to cross over resonance lines when the quasi-halo orbits are smaller. That is, until the resonances cross over the halo line. Moreover, the constant ω_0 constraint is unable to explore the portion of the family which is to the left of the left-most point on the halo line. The area of the family in this region does not have any generating halo orbits which have the same base frequency. If a quasi-halo orbit shares its ω_0 value with a halo orbit then they essentially have the same period and prevent secular drift between the two orbits, meaning constant ω_0 branches are excellent for studying formation flying when the chief spacecraft is assumed to be on a halo orbit.

The constant Jacobi energy constraint performs well by constructing long branches. It fills out the bottom left portion of the family near the halo line very well, but struggles to get past lower order resonances appearing at larger values of ω_1 . Constant Jacobi energy branches are particularly useful for studying families of orbits which can be reached given limitations on the propulsion of a spacecraft. Additionally, families of orbits at the same energy level provide the possibility for heteroclinic transfers between orbits.

It is not easy to quantify which parametric constraint is best because each constraint has its own pros and cons when it comes to studying families of quasi-periodic orbits. If the goal is

to explore the family as fully as possible, then the constant slope constraint is likely to achieve this goal. However, its exploration capability, just like the other parametric constraints, is limited to simple geometries of the parameter space of a family. To more fully explore families one needs to either use multi-parameter continuation, which has yet to be used to explore families of quasi-periodic orbits, develop search strategies which utilize multiple parametric constraints. In addition to either of these two options, one needs to find ways to reduce the effect of resonances which limit the efficacy of exploring families of quasi-periodic orbits. In this chapter we develop search strategies, and in Chapter 4 we address ways to reduce the effects of resonances.

The parametric constraints are swapped during the continuation to develop search strategies. The strategies developed and tested are named stair step, resonance bouncing, and roots. The stair step strategy swaps between Equation (3.39) and Equation (3.38) when the continuation method fails to find any more family members. The resonance bouncing strategy uses Equation (3.41) and changes the slope from m to $-m$ upon getting stopped by a resonance. The roots strategy is employed once a branch of solutions is stopped. The parametric constraint becomes Equation (3.41) and multiple branches are initialized with different slopes to explore in a variety of directions.

In each of these strategies it is advantageous to compute a new family tangent vector when the parametric constraint is swapped so that the next predicted family member is computed more easily. Let \tilde{z} be the current solution, then computing the null space of the Jacobian matrix of \mathbf{F}_p (Equation (3.35)) provides basis vectors for the tangent space of the 2-parameter family of quasi-halo orbits. We compute new basis vectors such that the last 2×2 block is the identity. We call these basis vectors the principle tangent basis, and denote them z'_{ω_0} and z'_{ω_1} . The principle tangent basis vector z'_{ω_i} provides a linear prediction of nearby solutions when solely ω_i is incremented. We leverage the principle tangent basis to construct a family tangent vector which points towards a solution satisfying the new parametric constraint.

Suppose we want to move in frequency space with slope m . We choose the family tangent according to Equation (3.42) by leveraging that the slope is simply the change in ω_1 divided by the

change in ω_0 .

$$\hat{z}' = \begin{cases} z'_{\omega_0}, & m = 0 \\ z'_{\omega_1}, & m = \pm\infty \\ z'_{\omega_0} + mz'_{\omega_1}, & \text{otherwise} \end{cases} \quad (3.42a)$$

$$z' = \frac{\hat{z}'}{\|\hat{z}'\|} \quad (3.42b)$$

The results of using each search strategy are presented in Figure 3.8. In plots (a) and (b) the black points are the orbits computed using the stair-step and resonance bouncing strategies, respectively. Plot (b) includes red lines indicating which resonances the solution branches are bouncing between. In plot (c) the black points show the solutions which only used a constant slope parametric constraint initialized from a the linearization about a halo orbit, while the red points show the extension of branches using the roots strategy.

The advantage of each strategy is that more area of the quasi-halo family is explored, and each has the potential to explore parameter spaces with more complex geometries by changing the search direction. The stair step strategy is seen to extend the branches to explore further than would have been done without swapping between the constant ω_0 and constant ω_1 constraints. A disadvantage of the stair step strategy is that it can get stuck at resonances.




Figure 3.8: Search strategies utilizing different parametric constraints: stair step (a), resonance bouncing (b), and roots (c).

The resonance bouncing strategy does well to stay within the bounds of two low-order resonances. The bounding resonances are the 3 : 1 and the 4 : 1 resonances and are labeled in plot (b) in Figure 3.8. By an $p : q$ resonance we mean that $\frac{\omega_0}{\omega_1} = \frac{p}{q}$. With the resonance bouncing strategy we can explore the family by filling out the regions between bounding resonances. It is unclear a priori which resonances will be strong enough to bound the continuation procedure, but with enough generating halo orbits this is not an issue.

It is seen in plot (b) that some branches escape the bounding resonances, showing the potential to move into other areas to explore. Moreover, it is seen that some branches are to the right side of the halo line. This happens when a branch of quasi-halo orbits approaches the halo orbit family. The amplitudes of the invariant curves approach zero and as a result become smaller than the error tolerance for the Newton's method. At this point any error in the quasi-periodicity constraint is below the error tolerance as well, becoming a "true" invariant curve. The branch continues as if nothing is wrong. Therefore, it is important to include a measurement on the size of invariant surfaces in the continuation method to ensure the invariant tori are not collapsing in a dimension.

Lastly, the roots strategy pushes the branches even further and fills out more of the family. The roots method is particularly advantageous to explore parameter spaces which are not geometrically convex since multiple branches are constructed to move in various directions. The strength of the roots method is not displayed here for two reasons. The first is that the constant rotation number constraint already filled out a large portion of the family, so there was not much space left to explore. Secondly, the geometry of the frequency space is largely convex such that using any one of the parametric constraints of Equations (3.38)-(3.41) by themselves is enough to largely explore the family.

Using a combination of various parametric constraints and search strategies we can explore more regions of the family and be more certain no areas are missing. There are many more parametric constraints and many more search strategies which can be constructed, however, for the L_2 quasi-halo family, the methods used here are sufficient for the exploration.

3.5.2 Family Surface

The family surface can be presented in various ways based on the parameterization of the family, but the natural way for families of quasi-periodic orbits is to show the surface in the space spanned by the quasi-periodic orbit orbit frequencies ω_0 and ω_1 . The entire range of the halo orbit family is used as a starting boundary to compute quasi-halo orbits. Family branches are computed using the four different parametric constraints of Equations (3.38)-(3.41) and the accumulation of the results are shown in Figure 3.9. Each data point in the figure represents a single quasi-halo with the corresponding frequency pair (ω_0, ω_1) . There are two separate surfaces because there are two distinct center manifolds that emanate from the halo orbits. The top plot shows the surfaces together whereas the bottom plots show each surface individually in more detail.

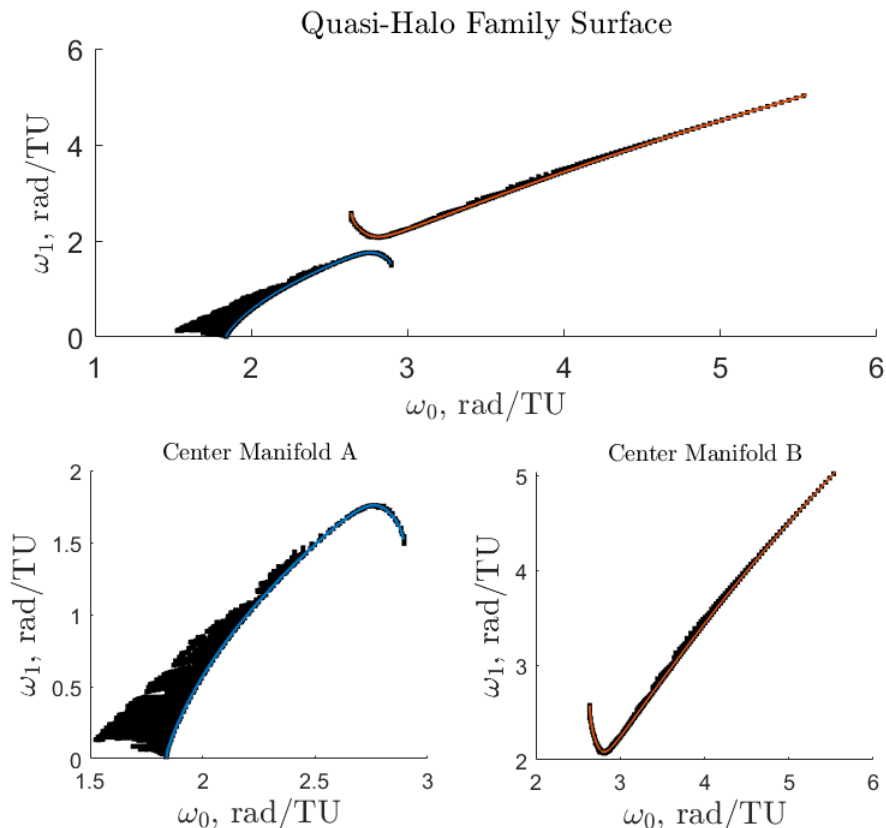


Figure 3.9: Family surface of Earth-Moon L_2 quasi-halo orbits parameterized by the orbit frequencies.

A large portion of the surface has ω_0 values less than the minimum ω_0 value of the halo orbit family. A quasi-halo with the same ω_0 value as a halo orbit have effectively the same period meaning spacecraft remain nearby for all time to a spacecraft on the halo orbit. This synchronicity is useful to be able to place satellites around the Lunar Gateway that can monitor the external health of the station or astronauts doing spacewalks. When the ω_0 values are not the same then the spacecraft on the quasi-halo will either move ahead or lag behind the chief spacecraft resulting in a secular type of drift. Eventually the two spacecraft will catch up to each other, but it will take a long time to do so. This means the areas of the surface directly above the halo line are suitable candidates for formation flying whereas the areas of the surface not above the halo line are not ideal for formation flying when the reference trajectory is a halo orbit. However, these quasi-halo orbits that are non-synchronous with halo orbits are synchronous with planar Lyapunov orbits. The synchronicity follows from the fact that the period of the planar Lyapunov orbits increases as the orbits grow larger from the libration point and continues to increase after the bifurcation of the halo orbit family.

The small displacement from the halo lines are not enough to say that these quasi-halos do not grow large, but it does show that there is not a lot of variability in the rates at which spacecraft move around on these quasi-halos. The sizes of the quasi-halo orbits need to be examined directly to make a conclusion about how large the quasi-halos grow in this region which is a study presented in Section 3.5.4.

The surface for CM-A (bottom left of Figure 3.9) is mostly triangular with some gaps. It is believed that most of these gaps are due to resonances between ω_0 and ω_1 . Recall that resonances between frequencies cause a degeneracy in quasi-periodic orbits, and for 2-dimensional quasi-periodic orbits this results in periodic orbits. There are resonances almost everywhere, but their effect is computationally noticeable at low-order resonances. Some resonance lines are plotted on top of the surface emanating from CM-A and shown in Figure 3.10. In this figure it can be seen that the resonance lines cause some branches to stop, yet different branches pass right through other resonances. The ability of resonances to stop the growth of branches seems to diminish near

the halo line. The exception to this is the 2:1 resonance that crosses at about the middle of the halo line which does not let any branch cross it. Near the crossing of the 2:1 resonance line and the halo line is the region of halo orbits that do not have center manifolds. Only in this area is it expected to not find quasi-halo orbits. The reason for the gaps and discontinuities is quasi-periodic orbit computation methods have problems computing quasi-periodic orbits near resonances. It is interesting to point out that on the resonance lines there exist families of periodic orbits with a larger period which can be found using high accuracy shooting methods.




Figure 3.10: Family surface emanating from CM-A overlaid with resonance lines.




Figure 3.11: Comparison of L_2 quasi-halo orbit family between our results in (a) to the results of Gómez and Mondelo in [75] in (b).

We compare our results to those in [75] in Figure 3.11. To do this we multiply the Jacobi constants by $-\frac{1}{2}$ to match their energy calculations. Looking at plot (b) in the figure there are four main regions where orbits do not appear. From top to bottom these correspond to the 2:1, 5:2, 3:1, and 4:1 resonances. These are low-order resonances that play a critical role in the organization of the dynamics, though they are not the only resonances that prevent quasi-halo orbits from being computed. For the range of halo orbits which are in common between our work and the work of Gómez and Mondelo they compute a larger portion of the quasi-halo family. This is because they used up to a maximum of 100 points to discretize their invariant surfaces and used 4 multiple-shooting segments. In this work we restrict to using 41 points and typically use single-shooting. This gives them an advantage in the numerical accuracy and allows them to compute quasi-periodic orbits closer to the point at which they break down. Gómez and Mondelo state that they stop computing halo orbits at the turning point in the Jacobi constant in the halo orbit family. Computing quasi-halo orbits from a wider range of halo orbits allows us to study the stable region within the halo orbit family, and allows us to study the quasi-halo orbits which share the same ω_0 value as the 9:2 NRHO.

Table 3.1: Maximum and minimum values of the Jacobi constant, perilune, and apolune obtained by the quasi-halo family.

Extremal Value	J (ND)	r_p (ND)	r_a (ND)
max	3.151818691968968	0.132410593278286	0.297495381129119
min	3.001239064876652	1.660669385605371e-04	0.154580747306807

When looking at the quasi-halo orbit family it is not completely understood why the boundary of the family exist in the way that it does. It is known that hyperbolic quasi-halo orbits are bounded by the planar Lyapunov orbits corresponding to the same energy level [75, 103]. Henon states that a family of periodic orbits terminates naturally due to either the family closing on itself or any of the size, energy, or period growing unbounded [86]. These same principles apply to families of quasi-periodic orbits. Moreover, it has been shown that quasi-periodic orbits break down when any of the bundles collide with each other, destroying the bundles and the orbit itself [46, 66, 81].

3.5.3 Characteristics

Recall that each point in Figure 3.9 corresponds to a unique orbit, and each of these orbits have measurable quantities that may be of interest for mission applications. The quantities computed here are the Jacobi constant (Figure 3.12), perilune distance (Figure 3.13), and apolune distance (Figure 3.14). In each of these figures the points are colored according to the values of each individual quantity in non-dimensional units (ND). The maximum and minimum values of the Jacobi constant, perilune, and apolune taken on by the computed family are listed in Table 3.1.

The behavior of the Jacobi constant for the family might look simple according to Figure 3.12, however it is nontrivial in the regions of the family surface that are very slim and therefore hard to see. Jacobi constant contours for CM-A outside the stable region make nearly straight diagonal lines leading downward away from the halo line. Considering constant ω_0 branches then the Jacobi constant decreases as the branches (and likewise the orbits) grow larger. For constant ω_0 branches up to $\omega_0 = 2.700793$ in the stable region of CM-A the Jacobi constant decreases as well, however for branches with ω_0 larger than this the trend in the Jacobi constant along branches switches and

begins to increase with some branches having a non-monotonic behavior in the transition zone. Jacobi constants along constant ω_0 branches in the stable region of CM-B only increase while branches outside the stable region transition to a decreasing trend.




Figure 3.12: Jacobi constant across the family surface in non-dimensional units.

The Jacobi constant is an energy-like quantity that dictates what areas of phase space are accessible, and as the value decreases more areas of space become accessible. Libration point orbits tend to have a decrease in the Jacobi constant as they grow larger from the libration point meaning more energy is needed to place spacecraft on those orbits and that their hyperbolic manifolds (if they exist) can extend to more areas of space. However, it is seen that there are branches of orbits where growing the orbits larger requires less energy to reach those orbits. Most of these branches are in or near the stable region, so these orbits could be targets for missions that require minimal station-keeping efforts.

The perilune and apolune are the closest and furthest approach distances respectively on an orbit with the center of the Moon. The behaviors of these plots are more physically intuitive. As the orbits grow larger from the halo orbits the perilunes decrease and the apolunes increase. As ω_0 increases the orbits are necessarily getting closer to the Moon since that is how the halo orbit family evolves. In Figure 3.13 the perilune decreases as the surface moves away from the halo lines

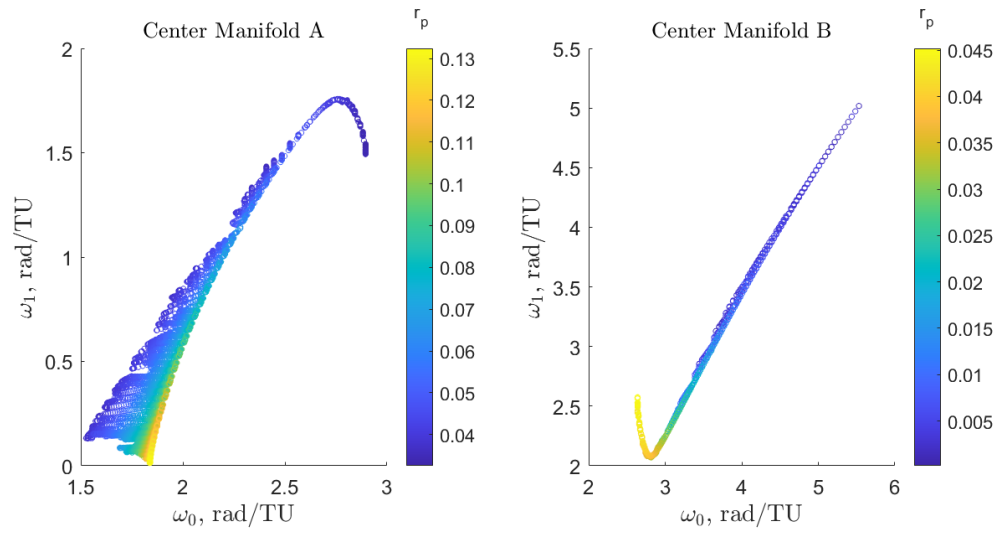


Figure 3.13: Perilune across the family surface in non-dimensional units.

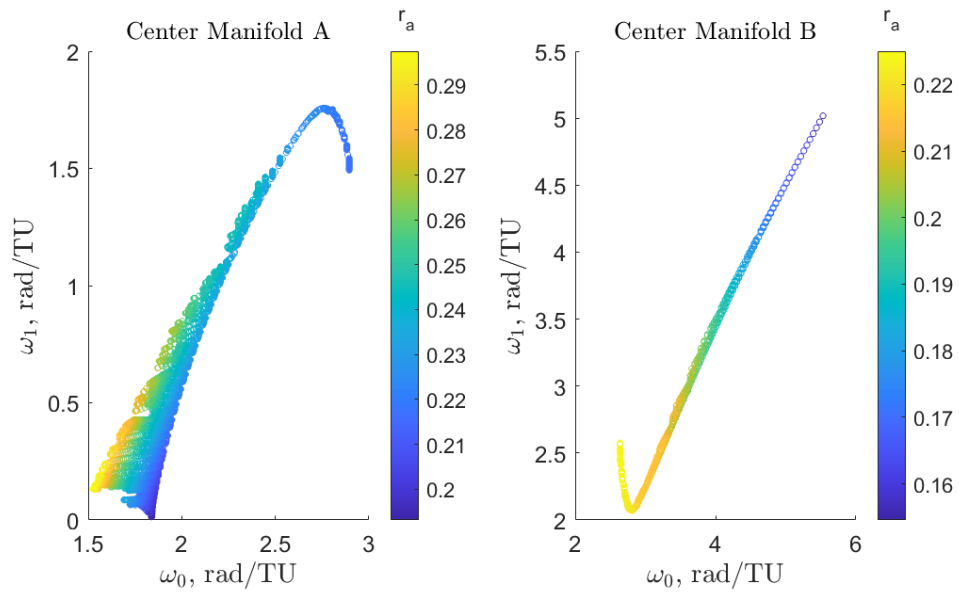


Figure 3.14: Apolune across the family surface in non-dimensional units.

(though less observable for CM-B) and decreases as ω_0 gets larger. In Figure 3.14 the apolune increases as the surface moves away from the halo lines, however the behavior as ω_0 increases is not as straight forward. Initially, the apolune increases as the ω_0 increases, but hits a maximum in the stable region of the halo orbits and begins to decrease as ω_0 increases further. This can be observed by looking at Figure 3.6. In non-dimensional units the radius of the Moon is about 0.00452, so an interesting observation is that there is an area of the family surface of CM-B composed with orbits that intersect the lunar surface. None of the quasi-halos emanating from CM-A have perilunes below this threshold.

Knowledge of these characteristics and undoubtedly others can be used to limit the search space of the family for orbits that meet preliminary mission design requirements. A Jacobi constant constraint could arise from an energy constraint due to the launch vehicle and spacecraft propulsion system or arise from a constraint on the hyperbolic manifold to reach a specific region in space. A perilune and apolune constraint could arise from viewing constraints to get images of the Moon at certain resolutions or arise from communication constraints.

3.5.4 Geometry

An example of three different family branches extending from the same halo orbit are in Figure 3.15 to show the variability in family members. The halo orbit has a frequency ω_0 of 1.868404, a Jacobi constant of 3.126576, and a center manifold frequency ω_1 of 0.166293. The three different branches are a constant ω_0 branch, a constant ω_1 branch, and a constant Jacobi energy branch. Each column represents a different family branch. The top row shows the branch on the family surface, the middle row shows the orbits in the branch colored according to the color gradient of the branch, and the bottom row shows the last family member of the branch. The family takes on a wide range of unique geometries which can not be achieved from the halo orbits. The quasi-halo orbits are self-intersecting near the x -axis crossings allowing for maneuvers to change phase. Near the x -axis crossings the halo orbit leaves and re-enters the branch of quasi-halos creating a potential for collisions and transfers onto other orbits. Additionally, the orbits in each branch intersect the

other orbits in the respective branches allowing for transfers to other orbits.

The constant ω_0 branch terminates at the 3:1 resonance and the effects of this resonance can be seen by the thicker bands in the last orbit giving a preview of what a 3:1 resonant periodic orbit looks like. This periodic orbit can be found by using a high accuracy predictor-corrector shooting method to propagate an initial condition on the dark band with a period of three times the stroboscopic time. The constant ω_1 branch terminates at a 10:1 resonance. The constant Jacobi energy branch terminates at the boundary of the family surface and not at a resonance. In this branch the value of ω_0 decreases below the halo orbit limit and the value of ω_1 tends to zero. This means the branch moves toward a planar Lyapunov orbit which can be seen by the thicker planar band in the last family orbit.

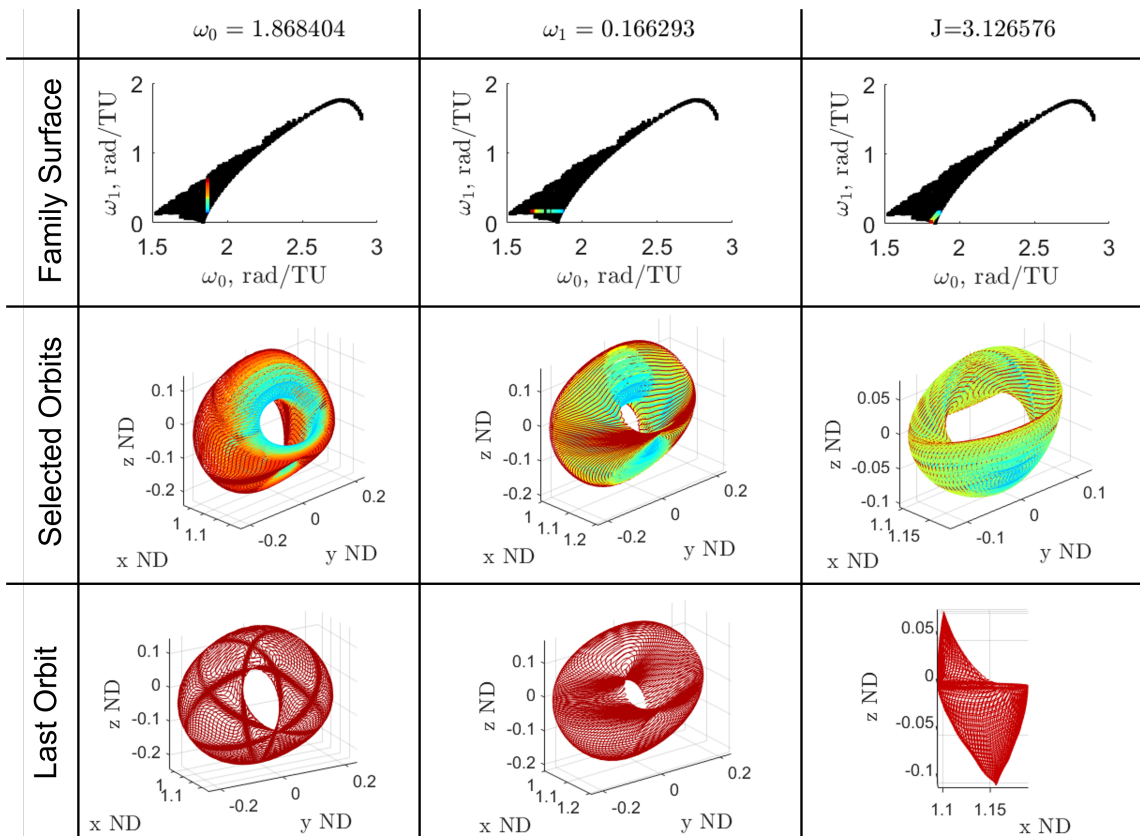


Figure 3.15: Three different branches of quasi-halo orbits extending from the same halo orbit.

Another way the geometry is analyzed is by computing the box size of each computed orbit.

The box size of an orbit is a 3-element set (l_x, l_y, l_z) that describes the lengths in the x , y , and z directions of the smallest box that encapsulates the orbit. Computing these values for each family member makes a set of values that can be plotted in a 3-dimensional space to create a 2-dimensional surface. This surface is in Figure 3.16 and each point on the surface is colored according to the Jacobi constant of the corresponding orbit. The surface does not give insight into the actual shapes of the quasi-halos but it does give a complete picture of the sizes of the quasi-halo orbits. The figure also contains three shadows that are projections of the surface on different planes. Additionally, the box sizes of the halo orbits are plotted on top of the shadows and colored according to the halo lines of Figure 3.4 to help the reader put this surface into correspondence with the family surface in the top plot of Figure 3.9. This type of surface can be generated for other 3-element data sets such as the maximum or minimum x - y - z coordinates of each orbit or even a 4-element set where the fourth variable is plotted as a color just like the Jacobi constant in Figure 3.16.

Following the surface along the boundary created by the halo lines the halo orbits grow in size while decreasing in the Jacobi constant. The largest halo orbit has a Jacobi constant of about 3.04. Continuing along the boundary the halo orbits get smaller while the Jacobi constant continues to decrease. At the point where the Jacobi constant achieves its minimum value the halos become stable. Past this point the Jacobi constant begins to increase in value while the halo orbits get smaller. After observing this behavior it is evident that the portion of the surface that belongs to the CM-B family surface is the top right hook-like portion of the surface. The rest of the surface then belongs to the CM-A family surface. An example of how to visualize the information from Figure 3.16 is in Figure 3.17. A box is constructed from the origin to a point on the surface and the corresponding quasi-halo orbit is plotted inside it and colored according to its Jacobi constant. It is then easy to see that the quasi-halo orbits belonging to CM-B are relatively small when compared to the orbits of CM-A. Similar to the Jacobi constant, perilune, and apolune the box sizes can be used in preliminary mission design to reduce the search space of quasi-periodic orbits to find orbits that meet a size requirement that stems from a communications requirement.

Another way to measure the size of a quasi-periodic orbit is by the maximum and minimum




Figure 3.16: Surface of box sizes of the L_2 quasi-halo orbits.




Figure 3.17: Example of how to interpret the surface of box sizes in Figure 3.16.

amplitudes of the invariant surfaces. For 2-dimensional invariant tori there is only one dimension to the invariant surfaces. Figure 3.18 shows the family of quasi-halo orbits in frequency space colored according to the location the invariant curve achieves its maximum and minimum values. The location is measured by θ_0 and is only specified between 0 and π . This is because the quasi-halo orbits exhibit symmetry about the x - z plane and $\theta_0 = 0$ occurs as the bottom of the orbit centered on the x -axis. At the beginning of the halo orbit family the minimum locations take on values near π and steadily decrease until the 2:1 resonance; at which point there is a sharp decrease to 0. At the beginning of the stable region and beyond the minimum locations are near 0. Interestingly, a similarly smooth behavior is not noticed with the maximum locations. There are three distinct regions which take on values of 0, $\frac{\pi}{2}$, and π . In the same figure we have provided examples of quasi-halo orbits from within each region and have plotted the invariant curves with the maximum and minimum sizes in each orbit.

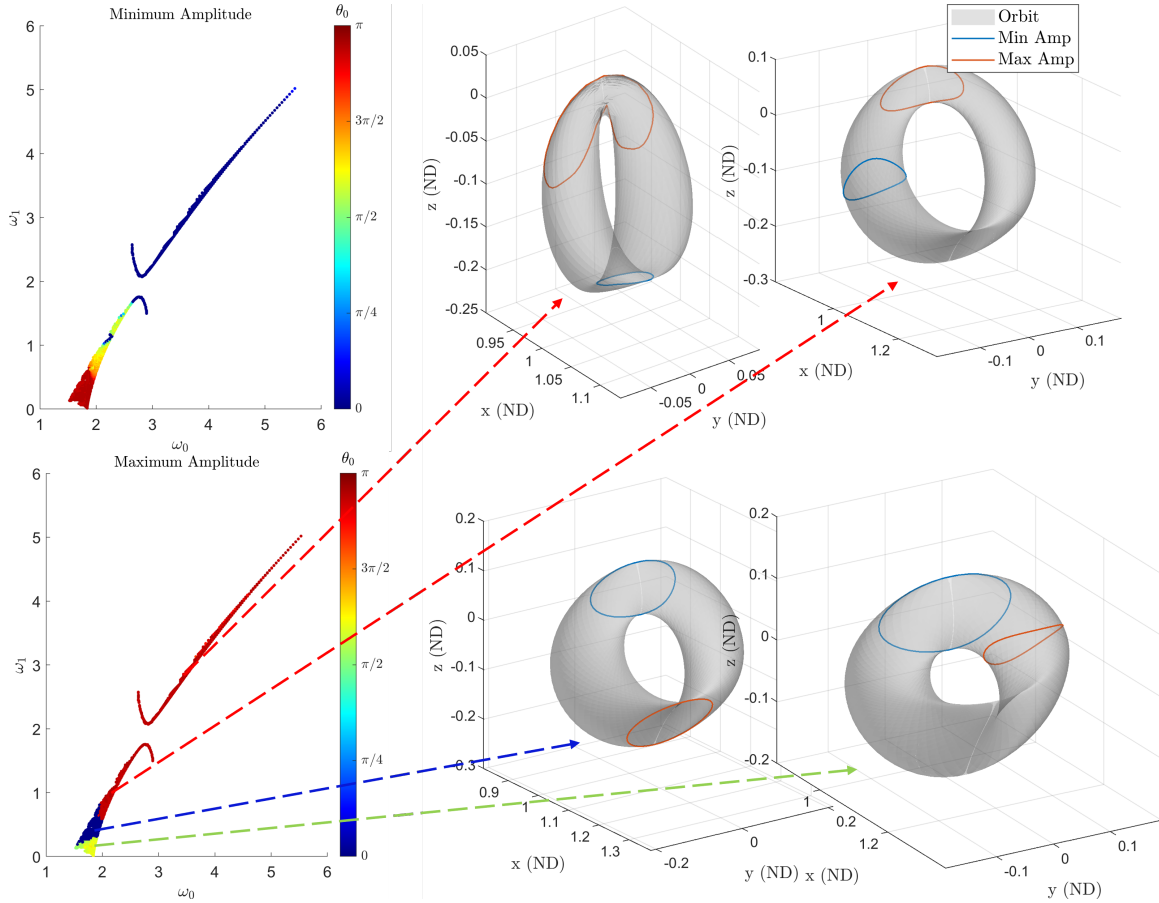


Figure 3.18: Locations of the minimum and maximum amplitude of the invariant curves across the 2-parameter family of L_2 quasi-halo orbits.

3.5.5 Period Doubling Halo Orbit Family

Periodic orbits with normal center manifolds with rational rotation numbers give rise to 1-parameter families of periodic orbits with larger periods. Common occurrences of these bifurcations are the period-doubling bifurcation, period-tripling bifurcation, and period-quadruplication bifurcation. Here we show the period-doubling halo orbit family and compute quasi-periodic orbits generated from this family. The purpose is not to show the entire family in a manner like the quasi-halo family, but to show these families exist and can generate families of quasi-periodic orbits different from the quasi-halo orbits.

The period-doubling family is very large and takes on a variety of shapes, so the family is




Figure 3.19: Period-doubling family of periodic orbits that bifurcates from the L_2 halo orbits broken into four segments.

broken into four segments based on the geometrical shape of the orbits and how the family grows (Figure 3.19). The family begins with the dark orbits in plot (a) and grow larger. The thick dark ring in the plot is the remnant of the halo orbit that the family bifurcated from. Moving clockwise to plot (b) the dark orbits continue the family where plot (a) left off with the light colored orbits. This pattern continues going from plot (b) to (c) and from plot (c) to (d). The orbits in plot (a) contain a center manifold, so these orbits give rise to a 2-parameter family of period-doubling quasi-halo orbits. An example of one of these quasi-periodic orbits is in Figure 3.20. The period-doubling halo and quasi-halo orbits are not a new discovery since they are presented by Gómez and Modelo in [75] among the period-tripling and period-quadrupling counterparts, however the period-doubling halo orbit family are displayed here.




Figure 3.20: Period-doubling quasi-halo orbit.

Generating an initial guess to a periodic orbit that has a $p : q$ resonance between the orbit frequency and its center manifold frequency requires the center eigenvector $\boldsymbol{\psi}$ of the monodromy matrix of a periodic orbit $\mathbf{y}(t)$ with period T . The initial guess of the state $\tilde{\mathbf{y}}(t_0)$ and period \tilde{T} is

generated from the following equations

$$\tilde{\mathbf{y}}(t_0) = \mathbf{y}(t_0) + \Delta s \text{Re}(\boldsymbol{\psi})$$

$$\tilde{T} = \frac{p}{q} T$$

This information is then put into the same predictor-corrector shooting algorithm used to generate the family of halo orbits.

3.5.6 Stability Analysis

The stability of each orbit is assessed from the eigenvalues of the stability matrix $G_{\boldsymbol{\omega}}$. As mentioned in Section 2.3.5 the eigenvalues come in rings, however only six eigenvalues are needed since the phase space is 6-dimensional. A method to sort the eigenvalues and pick the correct representative set is presented by Jorba in [102] and is the method used here. Once the set of representative eigenvalues is known then the orbit is classified as either hyperbolic, meaning it has eigenvalues off the unit circle indicating the existence of a hyperbolic manifold, or elliptic, meaning it has eigenvalues solely on the unit circle indicating stability and the existence of an additional center manifold where 3-d invariant tori reside.

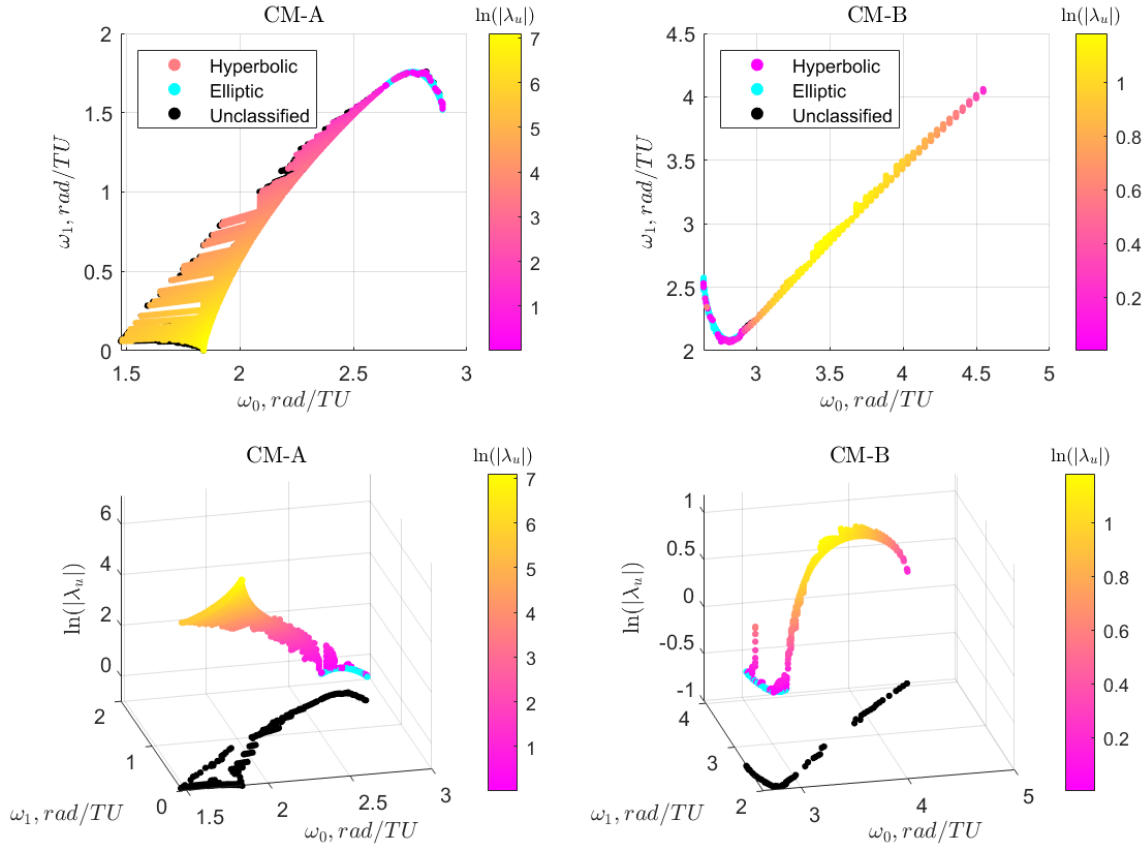


Figure 3.21: Lyapunov exponent across the family surface at a top-down view (top row) and a 3-d view (bottom row).

Using the method of Jorba on all orbits for which G_ω is computed each orbit is classified and the Lyapunov exponent is computed. The Lyapunov exponent is calculated as the natural logarithm of the magnitude of the unstable eigenvalue [102]. The results of this across the whole family surface are in Figures 3.21 while a zoomed in portion of just the stable region are in 3.22. In each figure the left column shows the stability of the family surface for CM-A, while the right column is for CM-B. The top row shows a top-down view while the bottom row shows a 3-dimensional view. Orbits classified as hyperbolic are colored and plotted on the z -axis according to the Lyapunov exponent while orbits classified as elliptic intrinsically have a Lyapunov exponent of 0 and have a fixed color to distinguish them from the hyperbolic points. Additionally, there are unclassified points that do not have a proper set of eigenvalues after using the method of Jorba. They have a

fixed color different from the hyperbolic and elliptic points and are assigned a value of -1 so that they are separated 3-dimensionally from the points that are classified. Note that -1 is not the value of the Lyapunov exponent.




Figure 3.22: Lyapunov exponent across the family surface in the stable region of the halo orbits at a top-down view (top row) and a 3-d view (bottom row).

An interesting feature of Figure 3.22 is that the stable quasi-halos are easily identified and that many orbits transition from stable to unstable in the region of the stable halo orbits. In this region we suspect 3-d invariant tori exist up to the stability bifurcation points. The information in Figure 3.22 can then be compared with the box sizes to infer the size of the 3-dimensional quasi-halo orbits. We compare the sizes of the orbits which have a stability bifurcation to the 3-d quasi-halo orbits in Section 3.6.4. Moreover, for constant ω_0 and constant ω_1 branches the orbits become more stable as the orbits grow larger. By comparing this figure to the constant Jacobi energy surface in

Figure 3.12 the conclusion is made that constant Jacobi energy branches become more unstable as the orbits grow larger.

Many points are unclassified near resonances, the surface boundaries, and in the stable region of the halo orbits showing a shortcoming of Jorba's eigenvalue sorting method. To validate the stability of the orbits in the stable region of the halo orbits numerical propagation tests are carried out by perturbing points on the orbits by about $1e-10$ and propagating those points for up to 27 years (1000 non-dimensional time units). If the orbit assumes a hyperbolic manifold then perturbations are made along the unstable eigenvector, and all other orbits have perturbations made in random directions. Analyzing data from these tests show that nearly every branch of quasi-halo orbits extending from a stable halo orbit has a stability bifurcation that transitions from stable to unstable.

3.5.7 Survey of Hyperbolic Manifolds

The previous sections of results deals with the quasi-halo orbits themselves and several quantities associated with each orbit to give a global perspective of the Earth-Moon L_2 quasi-halos. This section does not investigate the family as a whole, but instead looks at the hyperbolic manifold of a variety of orbits to generalize the behavior to the family. The hyperbolic manifold can be utilized to get to and depart from a specific orbit as well as make homoclinic connections - transfers between the same orbit - and heteroclinic connections - transfers between two orbits. This work looks at applications using the hyperbolic manifold to make orbit-to-Moon transfers which are trajectories on the hyperbolic manifold to move between a quasi-halo orbit and the Moon.




Figure 3.23: Family surface with stars indicating the sampled orbits used to investigate the behavior of the hyperbolic manifolds along with the regions identified as having distinct behavior.

3.5.7.1 General Behavior

Based on the dynamical structure of the halo orbit family there are two separate center manifolds and two separate hyperbolic manifolds. The quasi-halo orbits inherit the stability properties of the halo orbits in the linear regime and the properties are able to change as the quasi-halo orbits grow larger and move away from the halo orbits. In the stable region of the halo orbits, where there are two center manifolds, small quasi-halos will be stable, however, as it was shown, this behavior does not hold for larger quasi-halos in this region.

Quasi-halo orbits are sampled across the family surface and the unstable half-manifolds W_-^u and W_+^u are computed out to a time of 108.56 days (25 ND time units). The orbits are chosen such that most areas of the family surface have a representative data point in order to see the different types of behaviors for the unstable manifold. From this sampling there appears to be three regions that produce qualitatively different behaviors. The sampled orbits as well as the identified regions

are in Figure 3.23. Region A is the area of CM-A that lies below the 2:1 resonance line, Region B is the area of CM-A that lies above the 2:1 resonance line, and lastly Region C is the area covered by CM-B outside the stable region.

In Region A W_-^u breaks apart with one portion spiraling counterclockwise inward toward the space between the Earth's and Moon's orbital radius, another portion spirals clockwise outward, and another portion stays near the quasi-halo. On the other hand W_+^u spirals clockwise outward maintaining a shape topologically equivalent to the quasi-halo until a portion comes close to the Moon. The portion that comes close to the Moon spirals counterclockwise inward similar to W_-^u while the outer portion continues to spiral outwards.

In Region B W_+^u behaves similarly to W_+^u in Region A except that at some point the manifold appears to get sliced apart from the torus shape. This observation is due to the discretization of the manifold such that it appears to be tearing apart. W_-^u shears away from the quasi-halo such that one part stays near the quasi-periodic orbit while another part departs the area at an increasing rate. The portion of the manifold that departs the quasi-halo spirals clockwise outward and then behaves similarly to W_+^u .

In Region C W_-^u and W_+^u each behave in the same way. This behavior is similar to the way W_-^u behaves in Region A except that most of the manifold spirals outward and the departure from the quasi-halo takes longer. The delay in the departure is due to the fact that these orbits are much less unstable than the orbits in Region A and B.

Examples of unstable manifolds from these three distinct areas are plotted at discrete points in time in Figure 3.24. The left column shows W_-^u and the right column shows W_+^u . The top row uses a quasi-halo with frequencies (1.753081, 0.207659) and a Jacobi constant of 3.056899 from Region A, the middle row uses a quasi-halo with frequencies (2.341753, 1.308604) and a Jacobi constant of 3.010736 from Region B, and the bottom row uses a quasi-halo with frequencies (3.232268, 2.564819) with a Jacobi constant of 3.020008 from Region C.

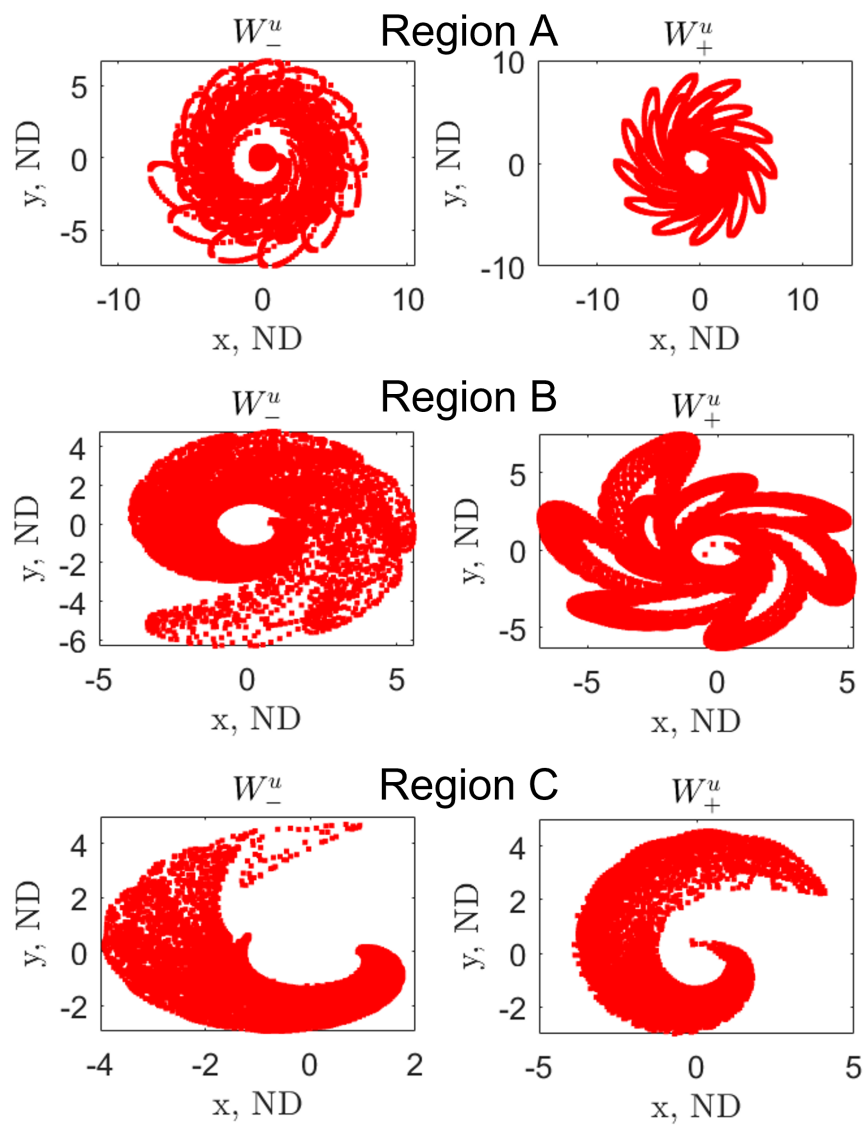


Figure 3.24: W_-^u and W_+^u for three different quasi-halo orbits from the three different regions identified to have distinct behavior.

3.5.7.2 Lunar Intersections With Hyperbolic Manifolds

The application explored here is lunar intersections with the hyperbolic manifolds. The purpose for this is motivated by the Lunar Gateway since quasi-halo orbits exist in the area surrounding the Gateway and can be utilized to support mission operations. Lunar impacts with the hyperbolic manifold give opportunities for transfers between the Moon and quasi-halo orbits. This type of work has been completed before in much greater depth by Bury and McMahon for families of periodic orbits in various systems [43]. After computing the unstable manifold for various quasi-halo orbits that span the CM-A family surface it seems that the quasi-halos in general have hyperbolic manifolds that impact the lunar surface. However, it seems that the quasi-halos from the CM-B family surface need to be large enough before the manifold intersects the Moon. There is a lot of variability in the impact locations on the Moon and four different impact patterns are plotted in Figure 3.25. The impacts shown include lunar intersections with both W_-^u and W_+^u with a maximum propagation time of 108.56 days. To find the impact locations with the stable manifold the transformation from Section 3.2.7 can be made to the unstable manifold impact points and are symmetric about the y -axis. The impact locations on the Moon's surface is fixed since the Moon is tidally locked with the Earth and maintains a fixed position relative to the orbits.

It can clearly be seen from Figure 3.25 that the impact locations appear in an organized fashion and are very different from one another. This suggests that the family surface can be divided into different regions based upon the lunar intersection pattern type which could potentially be accomplished with a clustering algorithm after computing the lunar intersections for a large amount of quasi-halo orbits. This can then be used to pick out orbits to target specific landing areas on the Moon or to know which orbits can be reached from specific locations on the Moon if utilizing the stable manifold. Additionally, at the boundaries of these intersecting regions the manifold lies tangent to the lunar surface, defining a locally optimal place to perform one-impulse maneuvers to enter onto the manifold.




Figure 3.25: Lunar surface intersection points with W_-^u and W_+^u within 108.56 days for four different quasi-halo orbits.

Not every point on a quasi-halo orbit will be a sufficient point to move onto the unstable manifold to achieve a lunar intersection, therefore it is important to know where on the orbit an intersection takes place. This is easily visualized by plotting points according to the angles θ_0 and θ_1 that parameterize the orbit. Figure 3.26 shows a quasi-halo orbit and the points that allow for lunar intersections with W_-^u up to a time of flight (TOF) of 108.56 days. This orbit corresponds to the top left lunar impact pattern of plot (a) in Figure 3.25. The points identified on the orbit are called takeoff points since they leave the orbit on the unstable manifold, and points that arrive on the stable manifold are called arrival points. It should be noted that the locations of the impact points on the quasi-periodic orbit are affected by the magnitude of the step size Δs used to compute the unstable manifold. To get a more accurate representation of the impact points a smaller step size needs to be used along with a longer integration time to ensure the manifold has enough time to reach the lunar surface.

Each takeoff point has a trajectory along the unstable manifold and each of the trajectories

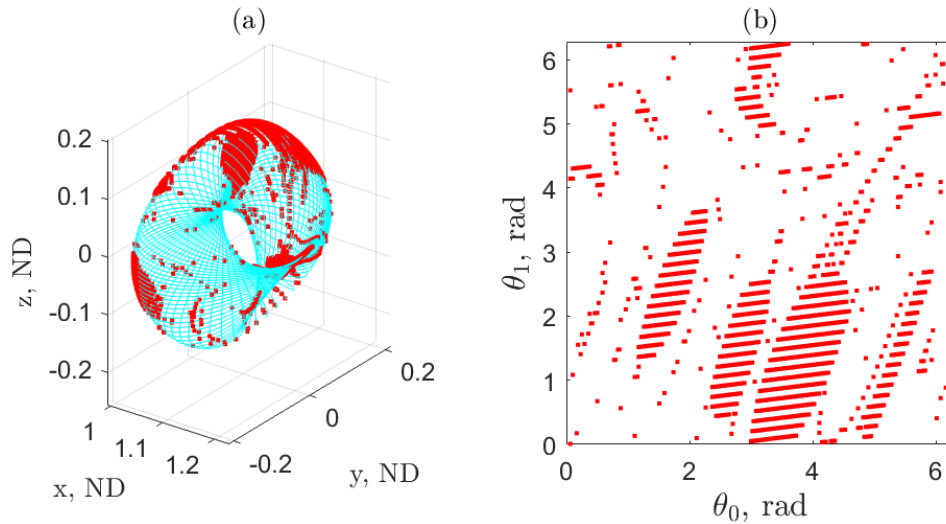


Figure 3.26: A quasi-halo orbit overlaid with the locations on the orbit where W_-^u impacts the lunar surface (left) and the parameterization of the orbit with the W_-^u takeoff points (right).

can be evaluated for different characteristics. Two characteristics are presented here that may be of interest for mission design. The first is the inclination of the Keplerian orbit assumed by each trajectory at the point of impact. This has applications to lunar landers and missions where a spacecraft would go into a low lunar orbit (LLO). This analysis is not constrained to the lunar surface and can easily be adjusted for different altitudes. The second characteristic is the Δv required to land on the surface of the Moon.

These characteristics are plotted in Figure 3.27 and correspond to the lunar intersection points associated with W_-^u of the quasi-halo in Figure 3.26. In plots (a) and (b) the x -axis is θ_0 and the y -axis is θ_1 which parameterize the quasi-periodic orbit, whereas the z -axis in plot (a) is the inclination in degrees and in plot (b) is the Δv in km/s. The points shown are only the W_-^u takeoff points since all of the other points on the orbit do not allow for lunar intersections up to a TOF of 108.56 days, and are colored according to the time for each trajectory to impact the Moon.

From Figure 3.26 there are clearly separate sections on the quasi-halo orbit that have trajectories on W_-^u that impact the lunar surface, and from the top plot in Figure 3.27 it is clear that these sections, and even within the sections, produce qualitatively different approaches to the

Moon. In this case the intersection points cover the entire surface of the Moon with a limited area on the quasi-halo to leave from. By choosing the correct takeoff point a specific landing spot on the Moon can be achieved. The inclinations range between 11 and 169 degrees allowing for nearly equatorial, polar, and retrograde LLOs without the need of an out-of-plane maneuver. Moving attention to the bottom plot in Figure 3.27 it can be seen that all trajectories have roughly the same Δv meaning the speed relative to the Moon is about the same for all impact points, so there is no benefit of choosing a takeoff point to minimize the Δv to land on the Moon. This seems to be the case for all investigated quasi-halo orbits, however a mission designer may want to choose a takeoff point to land at a specific location on the Moon or have a specific approach direction. Moreover, by knowing an impact point on the Moon a range of quasi-halos can be targeted by utilizing the stable manifold intersection points.

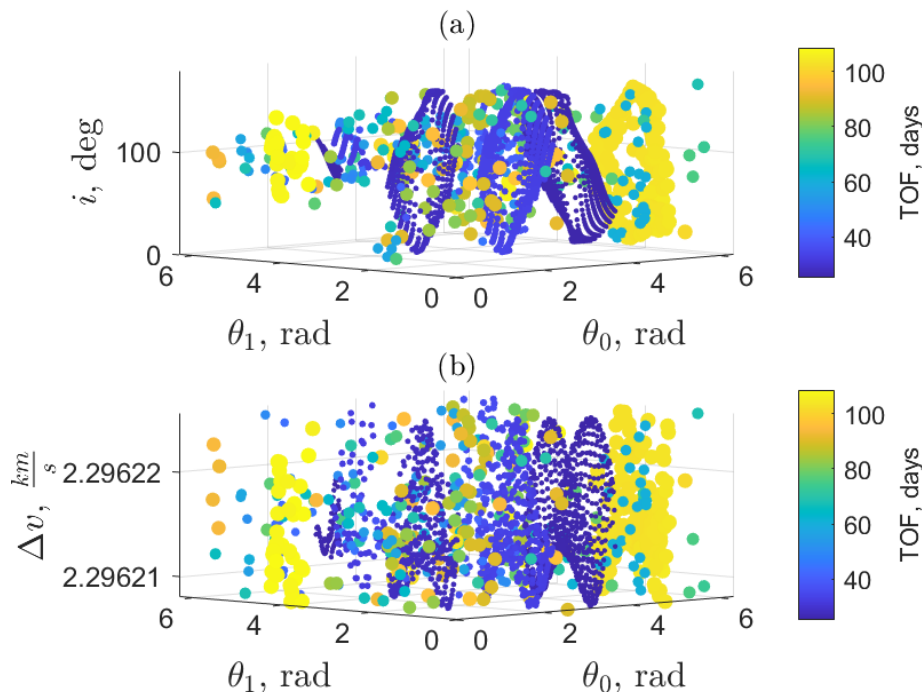


Figure 3.27: The inclination of the instantaneous Keplerian orbit about the Moon at the point of lunar impact (top) and the Δv needed to land on the lunar surface with no velocity (bottom).

These same types of plots can be made to look at other characteristics of the trajectories on W^u that impact the Moon such as the latitude and longitude of the impact sites. Additionally, the radius of impact can be changed to learn about the trajectories of W^u at higher altitudes such as the Keplerian orbital elements and the Δv to circularize the orbit.

3.5.8 9:2 Resonant Synchronous Quasi-Halo Orbits

This section reports findings about the branch of quasi-halo orbits that have the same base frequency ω_0 as the 9:2 halo orbit - the orbit of the Lunar Gateway. The quasi-halo orbits that are synchronous with the 9:2 halo orbit (referred to as 9:2 synchronous quasi-halo orbits) reside in the center manifold (CM-B) emanating from it and form a 1-parameter cantor family. Before talking about the findings of this family an introduction to the 9:2 halo orbit will be made.

The period of the 9:2 halo orbit is about 6.55 days ($\omega_0 = 4.1666$), a perilune of about 3212 kilometers, and a Jacobi constant of about 3.046741. It is pictured in plot (a) of Figure 3.28 with the Moon to scale to show what it looks like and where it is located. The 9:2 resonance is a resonance between the halo orbit's period and the lunar synodic period with respect to the sun and is not a resonance between the halo orbit frequency and the frequency of its center manifold. This orbit was chosen to stage the Lunar Gateway for various reasons. One of which is its ability to avoid most Earth eclipses and has short, infrequent lunar eclipses [169]. The orbit is weakly unstable making station keeping efforts fairly small and infrequent so that the Gateway can stay in orbit for a long time without running out of fuel. Since the orbit is unstable it has a hyperbolic manifold that can be used to arrive at and depart the orbit with little fuel usage.

Upon investigation of the hyperbolic manifold there are lunar intersections, however the intersection points do not cover a wide area of the Moon, and the TOF for the trajectories are at least 80 days long. The intersection points found up to a 108.6 day propagation period are shown in plots (a) and (b) of Figure 3.28, however a zoomed in picture on the Moon in plot (b) clearly shows these points. Additionally, within the propagation time the hyperbolic manifold achieves a minimum distance of about 110,000 kilometers from the Earth which has a TOF of about 107.7

days. A longer propagation time to look for more lunar intersections and closer perigees was not sought out since 108 days is already a long time for a ballistic trajectory to the Moon.

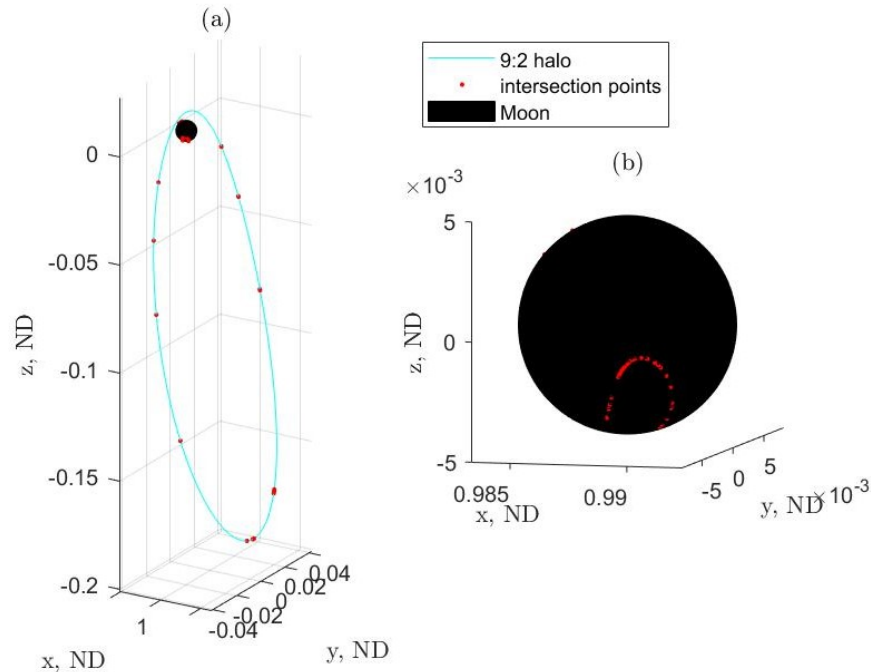


Figure 3.28: 9:2 halo orbit and Moon with lunar intersections (left) with a closeup on the Moon (right).

The 9:2 synchronous quasi-halo family is presented in Figure 3.29. This family of quasi-halos have the same base frequency as the 9:2 halo orbit ($\omega_0 = 4.1666$), so a spacecraft on any of these quasi-halos will oscillate around a spacecraft on the 9:2 halo orbit without secularly drifting apart, making these orbits good candidates for formation flying about the 9:2 halo orbit. This family intersects the 9:2 halo near the crossings of the x -axis, so there are opportunities for the Gateway or other spacecraft to transfer onto any of the family members. Additionally, each member of the family intersects the other members near the more positive x -axis crossing making it possible to directly transfer between any of the family members. This can be seen in plot(c) of Figure 3.29. The frequency ω_1 ranges from 0.540906 to 0.520461 and the Jacobi constant ranges from 3.046741 to 3.045658. Moreover, this family is unstable like the 9:2 halo, however the hyperbolic bundle is

two dimensions greater thereby providing more trajectories to reach the Moon, Earth, and other areas of space.




Figure 3.29: Family of 9:2 synchronous quasi-halo orbits at different views with the 9:2 halo orbit show in plot (d).

The 9:2 synchronous family also has lunar intersections with the hyperbolic manifolds. The intersections look similar to the 9:2 halo orbit lunar intersections at first but as the family grows the intersections take up a larger area on the Moon and changes to a completely different intersection pattern by the end of the family. The progression of the intersections as the family grows are shown in Figure 3.30. Plot (a) is the 9:2 halo lunar intersections and plot (f) is the lunar intersections for the largest computed family member.

For some orbit with $\omega_1 \in [0.536327, 0.536570]$ and $J \in [3.046506, 3.046493]$ a distinct pattern emerges near the north pole of the Moon and is depicted in plot (d). This spot persists and grows larger while the other dispersed intersection points disappear as the family continues to grow. The dispersed intersections points disappear because as the family grows larger the orbits decrease their perilune making the collection of intersection points closer together. For another orbit with $\omega_1 \in [0.5346046, 0.534878]$ and $J \in [3.046400, 3.046415]$ the quasi-halo family obtains a perilune below the lunar surface, thereby directly intersecting the Moon. This is what gives the very distinct and defined intersection pattern in plots (e) and (f). After this point it is possible to intersect the Moon without needing to move onto the unstable manifold, and likewise it is possible to directly

inject into an orbit from the lunar surface without getting onto the stable manifold.




Figure 3.30: Lunar intersections for a sample of the 9:2 synchronous quasi-halo family.

3.5.9 Conclusion

A study of the 2-parameter family of Earth-Moon L_2 quasi-halo orbits have been conducted to provide a global view of the family. Research methodologies and results are presented that allow for quick and initial studies for mission design. The quasi-halos take on a wide range of shapes and sizes providing for unique orbit geometries, and while they emanate from the center manifold of the halo orbits a large amount of the family are not synchronous with the halos. The majority of the family is composed of hyperbolic orbits with hyperbolic manifolds. Their Lyapunov exponents generally decrease as their frequencies move away from the frequencies of the halo orbits and their linear center manifolds. Emanating from stable halo orbits there are stable quasi-halo orbits, but a

bifurcation occurs when the size of the orbits grow large enough and become unstable. Surrounding the stable halo and quasi-halo orbits is a 3-parameter family of quasi-halo orbits diffeomorphic to 3-dimensional invariant tori with more complex motion.

The family of quasi-halo orbits provide opportunities for formation flying about halo orbits and can be transferred onto from halo orbits due to multiple intersection points in coordinate space. The family of halo orbits form a 2-dimensional manifold in phase space while the 2-parameter family of quasi-halo orbits form a 4-dimensional manifold. Additionally, the hyperbolic bundle of the halo orbit family is a 4-dimensional manifold in phase space while the hyperbolic bundle of the family of quasi-halo orbits is a 6-dimensional manifold. This means the quasi-halo family more densely fills phase space than the halo family and their hyperbolic manifolds provide more trajectories to efficiently move around and explore space. This is especially true for transfers to and from the lunar surface via the hyperbolic manifolds. The takeoff and landing points on the lunar surface make distinct patterns, and in a case presented covers the majority of the surface. The quasi-halo orbit family will be of great value to support the Lunar Gateway by providing opportunities to get to cislunar space from Earth, land on the lunar surface from orbit, and be outposts to escape and return to the Earth-Moon system. Additionally, the quasi-halo orbits that are synchronous with the 9:2 halo orbit provide opportunities to place satellites to monitor the external health of the Lunar Gateway and astronauts doing spacewalks.

3.6 The Dynamics in the Vicinity of the Stable Halo Orbits

This section focuses specifically on the quasi-halo orbits which are generated by the stable halo orbits. There are four different types of branches that are initialized from each halo orbit. The first type contains solutions of 2-d quasi-periodic orbits initialized by stepping onto CM-A, the second contains solutions of 2-d quasi-periodic orbits initialized by stepping onto CM-B. Each of these orbit types have two frequencies, however the frequency pair for the second branch type is (ω_0, ω_2) in order to make the labelling of the frequencies consistent with which center manifold they correspond to. The third branch type contains solutions of 3-d quasi-periodic orbits initialized

by stepping onto both center manifolds and holding ω_2 constant. The last branch type contains solutions of 3-d quasi-periodic orbits which hold ω_1 constant. The four branch types are computed from a span of elliptic halo orbits and pieced together to form four 2-parameter families.

The first branch type is called the “constant ω_2 2-d” branch since ω_1 changes throughout the branch while ω_2 is held constant, which effectively mutes motion in CM-B. Likewise, the second branch type is called the “constant ω_1 2-d” branch. The third branch is called the “constant ω_2 3-d” branch, and lastly the fourth branch is called the “constant ω_1 3-d” branch.

To construct branches of solutions Equation (3.33) must be satisfied, however there is a choice for the remaining parametric constraints. In the computations here the stroboscopic time is always chosen to be consistent with the period of the generating halo orbit T^* , corresponding to using the constant ω_0 constraint in Equation (3.38). This suffices as the only parametric constraint for computing branches of the 2-parameter families. One more parametric constraint is needed to compute branches in the 3-parameter family. For the constant ω_2 3-d branch we use a constant ω_2 constraint, while for the constant ω_1 3-d branch we use a constant ω_1 constraint.

Each orbit has been computed with an error tolerance of $7e-11$. The 2-d quasi-halo orbits use $N_1 = 111$ while the 3-d quasi-halos use $N_1 = 15$ and $N_2 = 11$. The reason for using 111 points to represent the invariant curves of the 2-d quasi-halos is because increasing the number of points increases the accuracy of the eigenvalues of the stability matrix [102]. The stability matrix of the 3-d quasi-halos is not necessary to compute as the orbits are inherently stable. The continuation of each branch is run until either: 1. The step size Δs has decreased to an allowable minimum step size. 2. The number of orbits computed reaches the maximum number of allowed orbits. The minimum step size for the 2-d quasi-halo orbits is taken to be $1e-5$, while it is $6e-8$ for the 3-d quasi-halo orbits. The maximum number of orbits is 120 for the 2-d quasi-halo orbits, while it is 100 for the 3-d quasi-halo orbits. Compiling the results from each 1-parameter branch gives four 2-d subsets of the solutions in the vicinity of the stable halo orbits.

3.6.1 Frequencies

The frequencies of the orbits are in Figure 3.31. Recall that each branch is a 1-parameter family of quasi-halo orbits grown from a halo orbit, so the plotted frequencies are lines extending from the frequencies of the halo orbits. Plot (a) of the figure shows the frequencies of the constant ω_2 2-d family, plot (d) is for the constant ω_2 3-d family, plot (b) is for the constant ω_1 2-d family, plot (e) is for the constant ω_1 3-d family, plot (c) is the combination of the constant ω_2 2-d and constant ω_1 2-d families in the three-dimensional frequency space, and plot (f) is the combination of the constant ω_2 3-d and constant ω_1 3-d families in the three-dimensional frequency space.

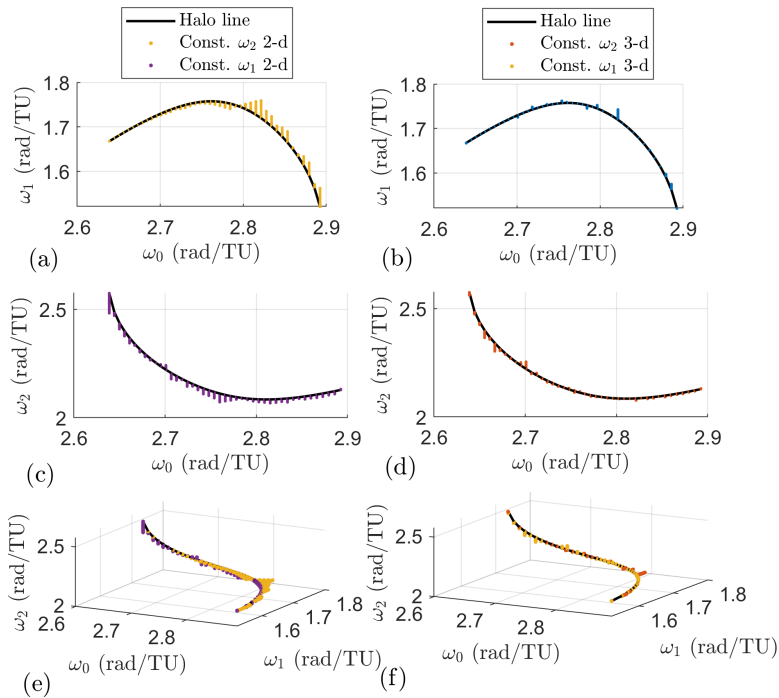


Figure 3.31: Frequencies of the four branches with different views.


The families of 2-d invariant tori fill more area in the frequency space than the families of 3-d invariant tori, showing a wider variety in the rate of motion among the 2-d invariant tori than the 3-d invariant tori. This figure may lead one to believe that there are more 2-dimensional invariant tori than 3-dimensional invariant tori in this region, however this is incorrect. In the 3-dimensional parameter space the 2-parameter families have measure 0 while the 3-parameter family

is non-zero. Moreover, it is much easier to compute 2-d invariant tori than 3-d invariant tori due to computational cost and the interplay between the internal frequencies. In Figure 3.31 branches move upwards and downwards in the subplots. It is interesting to note that we did not tell the continuation method to move “up” or “down”. We initialized a solution in the center manifold and let the continuation method go from there. Attempts to go in the “other” direction were made, but it was found that solutions do not exist on both sides at a given halo orbit.


Looking through the constant ω_2 2-d branch it is found that the frequencies of the invariant tori are not unique, showing the existence of non-KAM tori in the circular restricted three-body problem. Examples of three of these branches are in Figure 3.32. It is evident in these branches ω_1 has a critical point as each computed invariant curve is not repeated. If the continuation procedure had turned around and began backtracking then this would be evident in the invariant curves and orbital characteristics.

3.6.2 Amplitudes


Recall that ω_1 corresponds to motion in CM-A and ω_2 corresponds to motion in CM-B. Continuing with the notation, then A_1 is the amplitude of motion in CM-A and A_2 is the amplitude of motion in CM-B. As such the constant ω_2 2-d family has invariant surfaces with amplitudes A_1 , the constant ω_1 2-d family has invariant curves with amplitudes A_2 , and the constant ω_2 3-d and constant ω_1 3-d families have invariant surfaces with both amplitudes. The amplitudes for each elliptic quasi-halo family are shown in Figure 3.33 in a log-log plot. The constant ω_2 3-d family adheres to the left color bar while the constant ω_1 3-d family adheres to the right color bar. The points are colored according to the ω_0 value of their corresponding quasi-periodic orbit (this also identifies which halo orbit that branch originates from). The amplitudes of the initial quasi-halos for each branch are around $1e-8$ in both dimensions for reference. The distribution of the amplitudes between the two families appears to be symmetric. Many of the orbits grow large in both dimensions while some of them grow small in one of the dimensions. It should be noted that at some point when one of the dimensions gets too small that orbit can no longer be considered a 3-d quasi-halo;



(a)



(b)



(c)

Figure 3.32: Three constant ω_2 2-d branches showing the non-uniqueness of the orbit frequencies.

rather it has degenerated to a 2-d quasi-halo. This degeneracy point should be when an amplitude falls below the error tolerance used for convergence.




Figure 3.33: Distribution of amplitudes for the branches of 3-d quasi-halos.

From Figure 3.33 it can be seen that some constant ω_2 3-d family members have a large A_1 and a small A_2 (bottom right region), while few family members have large A_2 and small A_1 (top left region). This shows a preference for A_1 to grow for the constant ω_2 3-d family. The opposite behavior is seen for the constant ω_1 3-d family, which shows a preference for A_2 to grow while A_1 remains small. Recall that the amplitude A_1 is tied to the magnitude of oscillations in CM-A, while the amplitude A_2 is tied to the magnitude of oscillations in CM-B. Then, the observed behaviors *logically* make sense because freezing ω_2 should limit the growth of the amplitude A_2 in CM-B. Since the initial excitement of CM-B is small, then A_2 *should* remain small. Likewise for freezing ω_1 . However, looking at the top right region of the figure, it is seen that most of the branches contain orbits which grow equally in both amplitudes. This result contradicts the logical argument presented above, so it is unknown why the amplitudes grow the way that they do in these families. Figure 3.33 does not present a relationship between the growth of the frequencies and the growth of the amplitudes, however this will be shown in the next section.




Figure 3.34: Maximum amplitudes within each of the four computed branch types.

A comparison is made to the amplitudes of the 2-d quasi-halos and shown in Figure 3.34. The four marker types correspond to each of the four branch types while each data point is the maximum amplitude within that branch of orbits with fixed ω_0 . It is seen that the maximum size of the 2-d quasi-periodic orbit branches are typically several times larger than the 3-d quasi-periodic orbit branches. In dimensional units the maximum amplitude of: the constant ω_2 2-d family is 20,000 km, the constant ω_1 2-d family is 19,000 km, the constant ω_2 3-d family is 6,700 km, and the constant ω_1 3-d family is 14,000 km. For reference the maximum amplitude of the 2-d quasi-halo orbits in 3.5.4 is 117,800 km. It turns out there is generally a single location along the θ_0 direction where the 2-d quasi-halo orbits in the vicinity of the stable halo orbits attain their maximum or minimum amplitudes, namely the point closest to the secondary ($\theta_0 = \pi$) and the point furthest from the secondary ($\theta_0 = 0$), respectively (Figure 3.35).

3.6.3 Jacobi Constant

The Jacobi constant is plotted as a color gradient for each of the two 2-parameter families in the left column in Figure 3.36. The right column contains the same plots but with a different




Figure 3.35: Histogram of location where maximum and minimum amplitude along each 2-d quasi-periodic orbit occurs (a) and a plot of the invariant curves in configuration space with maximum and minimum amplitude for each 2-d quasi-periodic orbit (b).

color gradient. This color gradient represents the change in Jacobi constant among each branch compared to the Jacobi constant of the halo orbit each branch originates from. The top row shows the constant ω_2 2-d family while the bottom row shows the constant ω_1 2-d family. Figure 3.37 shows the same information but shows the 3-d families instead. The plots in the left columns show that as ω_0 increases the Jacobi constant increases. This is the same behavior the halo orbits themselves follow. Each branch appears to be a single color, showing there is little change in the Jacobi constant as the orbit amplitudes grow larger. The more interesting observations come from the plots in the right columns.

The first observation is that the Jacobi constant does not change very much as the orbits grow larger, however near the ends of many branches the free frequency changes more rapidly as does the Jacobi constant. This seems to suggest that rapid changes in the frequencies and the Jacobi constant compared to the sizes of the orbits indicate branches are nearing the end of the family. The second observation is that the Jacobi constant decreases along some branches while increasing along others. In contrast, Figure 3.12 shows the Jacobi constant typically decreases as the L_2 quasi-halo orbits grow larger while holding either frequency constant. This is largely seen with the quasi-halos emanating from the halo orbits with type *center* \times *saddle*. However, the Jacobi constant exhibits a more complex behavior in the vicinity of the stable halo orbits.

The Jacobi constants computed lie in the range $[3.014788704506776, 3.017204276227027]$ for

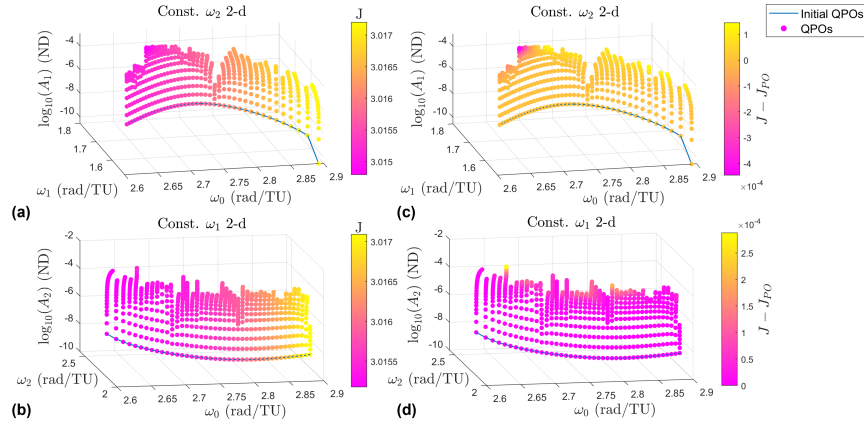


Figure 3.36: Plots displaying the Jacobi constant for the 2-d quasi-halo branches.

the constant ω_2 2-d family, [3.015176655614008, 3.017088403921683] for the constant ω_1 2-d family, [3.015065998407869, 3.017086688496275] for the constant ω_2 3-d family, and [3.015169926957731, 3.017087620465851] for the constant ω_1 3-d family.

3.6.4 Stability Bifurcation

The 3-d quasi-halo orbits are stable, so there are no hyperbolic manifolds to utilize to approach or depart these objects. The benefit of the stability property is that small perturbations will not cause an asymptotic departure from the nominal orbit. At most there will be a bounded secular drift caused by the difference in frequencies between the nominal orbit and the orbit a spacecraft has been perturbed onto. The 2-d quasi-halos in the region come in two stability types: elliptic and partially hyperbolic. The elliptic quasi-halo orbits are stable and behave similarly to the 3-d quasi-halo orbits, however a small perturbation will generally excite the third mode of oscillation creating a 3-d quasi-halo. This is one reason it is important to study the 3-d quasi-halo orbits. The partially hyperbolic orbits have a hyperbolic manifold emanating from them. These orbits, other partially hyperbolic quasi-halo orbits, and unstable halo orbits can be utilized for low-energy transfers to the stable region. Once the spacecraft is close enough a maneuver or some other transfer design (such as in McCarthy and Howell [130]) can lead the spacecraft to a stable orbit.

The method of Jorba in [102] is too numerically unstable to accurately classify the stability

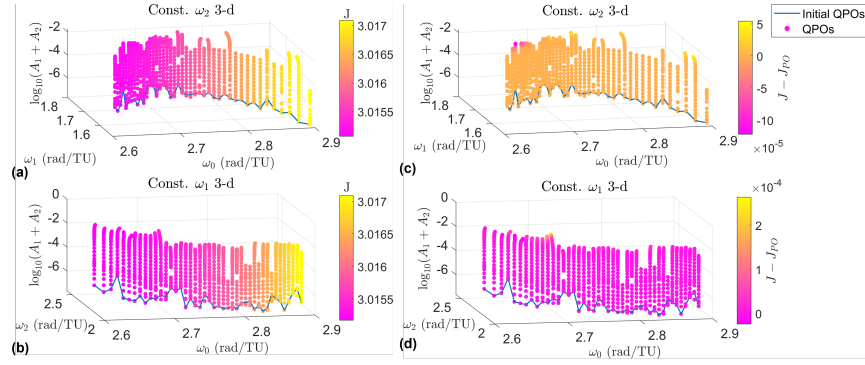


Figure 3.37: Plots displaying the Jacobi constant for the 3-d quasi-halo branches.

of the orbits in the stable region of the halo orbits. In this section small deviations on the order of $1e-10$ are made from the invariant curves of the 2-d quasi-halo orbits and propagated to various points in time. The propagated points are compared to the original invariant curve. This analysis is a visual inspection to distinguish between stable and weakly unstable behavior. After combing through each 2-d quasi-halo branch the points are identified and plotted in Figure 3.38. This figure is the same as Figure 3.36 but with the identified bifurcation points where the branches transition from stable to unstable.

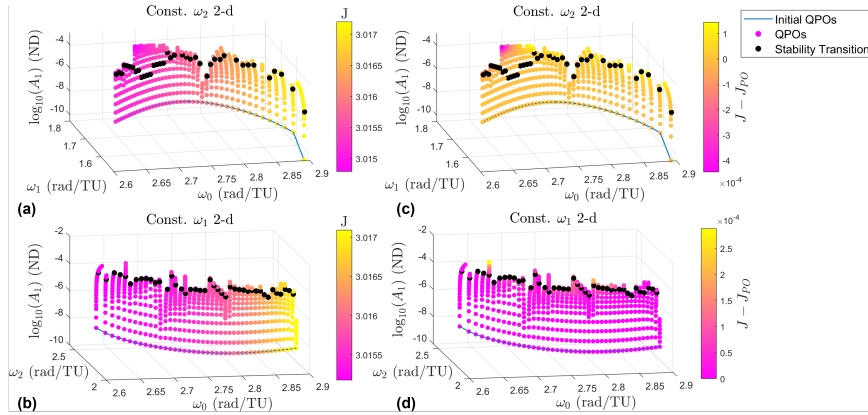


Figure 3.38: Stability transition in the 2-d quasi-halo branches.

The locations of these bifurcation points are generally consistent with the behavior of the eigenvalues. Figure 3.39 shows the eigenvalues of the quasi-halo orbits at the identified bifurcation

points and for orbits before and after the bifurcation point in the continuation process for each constant ω_2 2-d branch. Likewise, Figure 3.40 show the eigenvalues of the constant ω_1 2-d branches around the bifurcation points. The orbits in a branch can be numbered 1, 2, 3, ... in accordance with the order in which they are computed. The number at which the last orbit is stable before becoming unstable in the visual analysis is defined as N^* . Then the notation $N^* \pm n$ is the n^{th} orbit computed after or before the identified transition point N^* . Orbits before the transition point are stable while the orbits after the transition point are unstable. The color of each point in each plot is in accordance with the halo orbit from which each quasi-halo has been generated from.

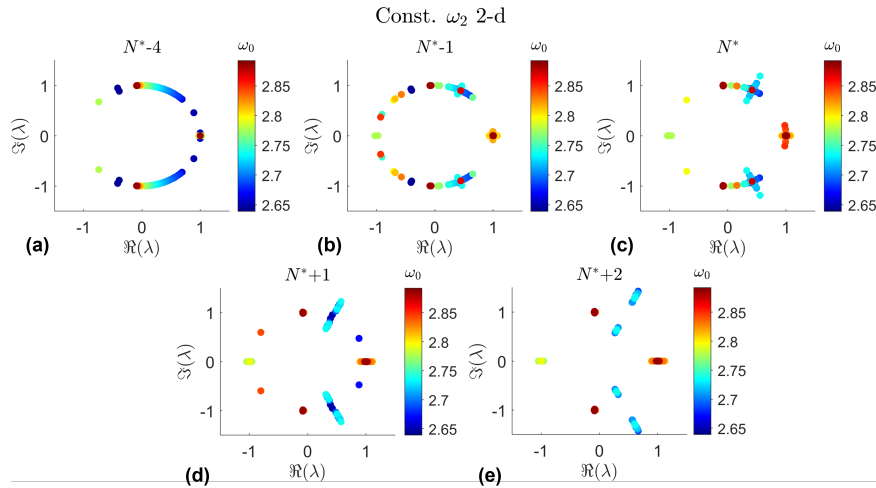


Figure 3.39: Eigenvalues near the identified bifurcation points from the visual analysis for each constant ω_2 2-d branch.

What is seen in Figure 3.39 is the progression of the eigenvalues as the continuation procedure approaches and passes through the bifurcation points for each constant ω_2 2-d branch. The mechanism for the bifurcations are the collisions between eigenvalues that push them off of the unit circle. It is clear that the eigenvalues begin on the unit circle, indicating stability, and break off as they near the identified stability transition point from the visual analysis. There are 3 locations where the eigenvalues break from the unit circle. The eigenvalues that collide at plus or minus 1 break off onto the real axis. The eigenvectors of these eigenvalues give the tangent directions to the stable and unstable manifolds. The eigenvalues that collide on the unit circle away from plus and

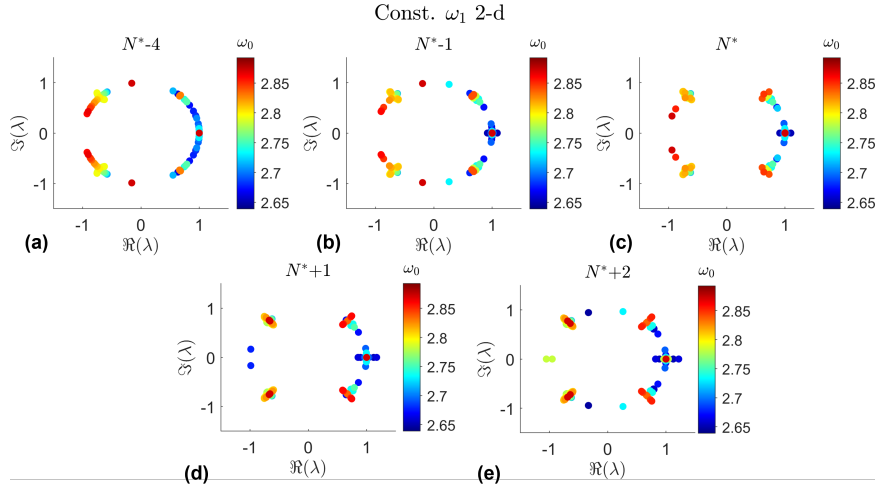


Figure 3.40: Eigenvalues near the identified bifurcation points from the visual analysis for each constant ω_1 2-d branch.

minus 1 break into a complex quadruple, indicating the formation of a complex saddle. A similar behavior is seen for the constant ω_1 2-d family in Figure 3.40.

In Section 3.5.6 it is conjectured that the transition points are thought to bound the sizes of the 3-d quasi-halo orbits. To test this theory the amplitudes of the 2-dimensional elliptic quasi-halos are compared with the amplitudes of the 3-dimensional quasi-halos in the same manner that Figure 3.34 is constructed but with the partially hyperbolic orbits removed. This comparison revealed that there are 3-d quasi-halo orbits with amplitudes larger than the elliptic 2-d quasi-halo orbit amplitudes, so no conclusions can be drawn from the comparison with our data.

3.6.5 Region of Stability

All the orbits in this work are numerically computed with a pseudo-arclength continuation method which terminates when either the maximum number of family members has been computed or when the step size falls below the minimum step size. Outside these conditions the cause of termination of the continuation method is not always certain. The termination could be due to approaching a resonance that the continuation was unable to move past, or it could be due to the branch reaching the end of the family. To check for sizes of potentially missing quasi-periodic

orbits a numerical perturbation analysis is performed. The idea behind this test is that since this region is stable, then any points perturbed from the stable halo orbits should remain in the area, given that the magnitude of the perturbation is small enough. The perturbed points that remain in the area after an amount of time should generically be on 3-d quasi-halo orbits. The size of the perturbation at which orbits begin to depart the area can be used to quantify the size of the stability region. This test will answer the question, given a perturbation magnitude from a stable halo orbit, how likely is it that the point will remain in the area (i.e. in the stable region)?




Figure 3.41: Empirical estimates of the size of the stable region surrounding the stable halo orbits.

The test is performed by initializing points with a given perturbation size and then propagating them forward in time. At the end of the time they are determined to have either departed or remained in the stable region based on their distance from the generating halo orbit. For this test 1,500 points are initialized in the center subspaces of a stable halo orbit from Equation (2.29b) with particular step size Δs_i and various values of θ_0 and θ_1 . Those points are then propagated out to 10 orbital periods of the periodic orbit from which they emanate from. The minimum distance from each point to the halo orbit is computed at the final time. Points with a distance of more than

1.3 times the initial step size are considered to have departed the area. The percentage of departed trajectories is then calculated. A departure percentage is targeted from a bisection method to find the step size required for that departure percentage. For a targeted departure percentage the step size is found for each stable halo orbit and recorded. The targeted departure percentages are 50%, 70%, 90%, and 97%, and the results are in Figure 3.41.




Figure 3.42: Comparison of the 97% line with the computed orbit amplitudes.

An interesting result from this experiment is that as ω_0 increases moving through the stable halo orbits the step size needed to reach the departure percentage continues to increase and is several times larger at the end of the halo orbits than at the beginning. At first the step size increases because the perturbed points are getting further away from the boundary between stable and unstable halo orbits. It was thought that the step sizes would decrease as the perturbed points approach the end of the stable halo orbits, however this is not observed. The general increase in the step size suggests that the hyperbolic manifold after the stable region is weaker than the hyperbolic manifold prior to the stable region. A larger step size is needed in order to push the perturbed points outside the threshold in the allotted amount of time. From this idea it can be inferred

that the actual step size needed to stay in the stable region is smaller than the found step size. The hypothesis about the strength of the hyperbolic manifolds is confirmed from the Lyapunov exponents of the 2-d quasi-halo orbits emanating from outside the stable halo orbit region in Lujan and Scheeres [121].

The 97% line in Figure 3.41 serves as a theoretical limit to the size of the stable region, and hence the maximum size of quasi-halo orbits in this region. This line is then compared with the amplitudes from Figure 3.34 to determine if we have reasonably found the maximum sizes of quasi-halo orbits in this region. The comparison between the region of stability and the orbit amplitudes seem to agree fairly well in Figure 3.42. Some branches of the 2-d quasi-halos have orbits with amplitudes larger than the step sizes that comprise the 97% line while most of the branches lie below this line. Toward the right side of the plot where ω_0 is larger the gap between the orbit amplitudes and the 97% line grows. This could indicate that the 97% line should begin sloping downward to account for the weaker hyperbolic manifold as mentioned above, or it could be that the continuation procedure terminated prematurely and did not find larger orbits.

3.6.6 Geometry

The 2-d quasi-halo orbits are quasi-periodic tori of dimension 2 and have invariant surfaces that are closed curves in the 6-dimensional phase space of the circular restricted three-body problem. As one of these curves evolves in time it extrudes out the shape of the entire quasi-halo orbit forming a 2-dimensional surface. When this object is projected into configuration space the object is still a surface, however it becomes self-intersecting. The motion of a spacecraft on a 2-d quasi-halo orbit lies on the surface of this object.

While we cannot show all the 2-d quasi-halo orbits on a single plot we use the invariant curves to construct a surface encapsulating all the computed invariant curves for each of the two corresponding branch types. This is done by taking the last invariant curve from each branch of a given branch type and constructing a surface that connects the curves to each other. The surface created in this fashion for the constant ω_2 2-d family is in plot (a) of Figure 3.43 while the surface

created in this fashion for the constant ω_1 2-d family is in plot (b). The invariant curves used to construct each surface are plotted and colored according to the value of ω_0 they possess. Coloring them this way allows one to see how the invariant curves change as the generating stable halo orbit changes.

The sharp changes in the surface are a result of the differences in size and shape between the last invariant curve from one branch to another. These differences arise from the partitioning of the stable halo orbits, from dynamical effects from crossing over and getting stuck by resonances, and the numerical reasons of termination of the program.

The 3-d quasi-halo orbits are quasi-periodic invariant tori of dimension 3 and have invariant surfaces that are 2-dimensional surfaces in the 6-dimensional phase space of the circular restricted three-body problem. The extrusion of one of these surfaces in time forms a 3-dimensional surface which fills a volume in any 3-dimensional subspace of the 6-dimensional phase space. The fact that the orbit fills a volume in configuration space makes it difficult to view and understand the motion of a spacecraft on one of these orbits. Additionally, the method of constructing a Poincaré map of a constant Jacobi energy family (such as in [75]) cannot be applied to view and analyze the elliptic quasi-halo orbits. Moreover, the Poincaré maps of the elliptic 2-d quasi-halos orbits are allowed to overlap unlike their hyperbolic counterparts. However, looking at the invariant surfaces individually in configuration space gives geometrical insight into the full orbits.

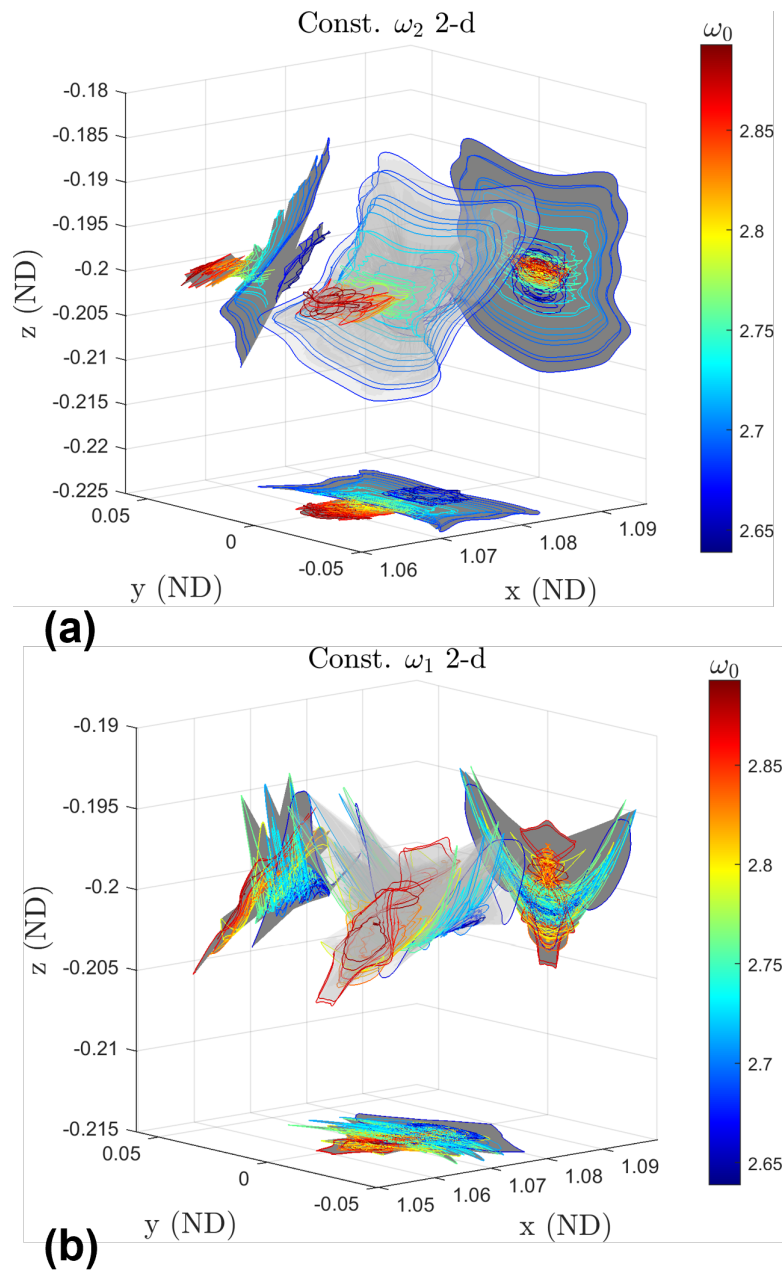


Figure 3.43: Surface (with shadows) made by the last invariant curve of each branch in the constant ω_2 2-d family (a) and in the constant ω_1 2-d family (b).

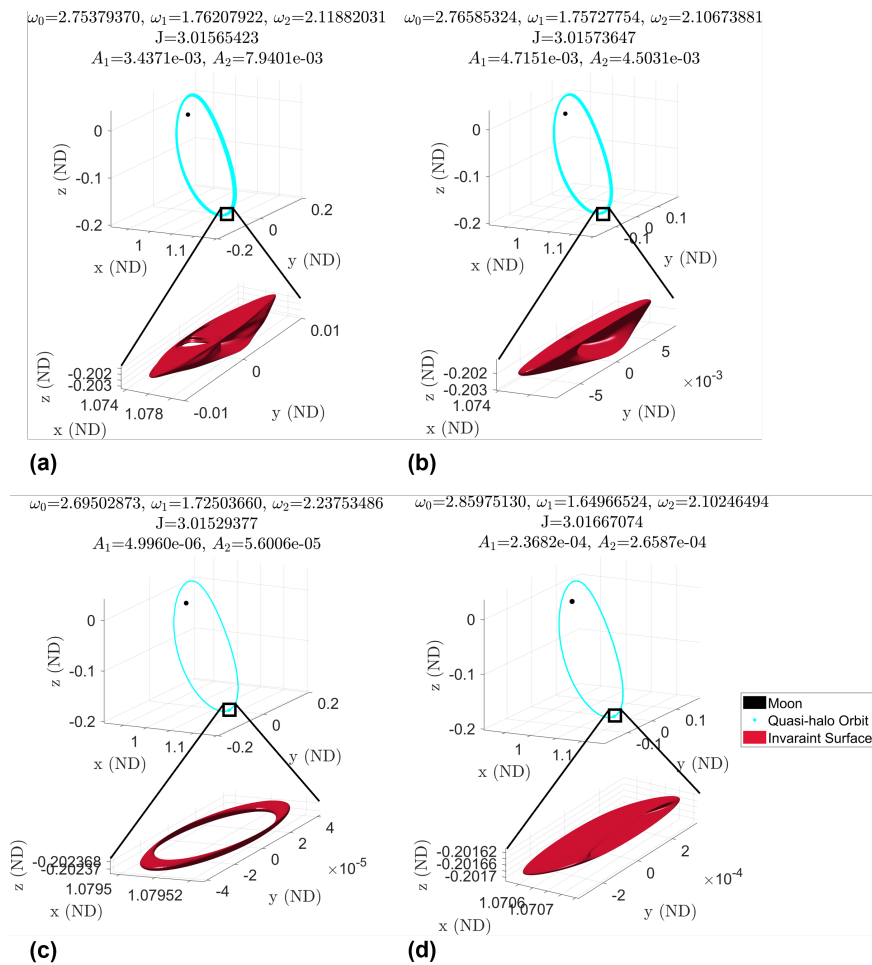


Figure 3.44: Examples of 3-dimensional quasi-halo orbits in configuration space.

Figure 3.44 shows four examples of 3-dimensional quasi-halo orbits in configuration space along with the Moon for a reference on size. Additionally, the invariant surfaces for each orbit are shown. Comparing these orbit amplitudes to Figure 3.33 and Figure 3.34 shows that plots (a) and (b) are among the largest of the 3-dimensional quasi-halo orbits while plots (c) and (d) are among medium-sized orbits. This figure shows just how small the 3-dimensional quasi-halo orbits are compared to the partially hyperbolic orbits.

Figure 3.45 shows the growth and evolution in the continuation procedure of the invariant surface for a constant ω_2 3-d branch and a constant ω_1 3-d branch. In each plot is the points representing the surface and the interpolated surface. A constant ω_2 3-d branch is in plots (a), (b),

and (c). This branch has constant frequencies $\omega_0 = 2.821444$ and $\omega_2 = 2.085261$. Plot (a) is taken near the beginning of the branch, plot (b) is taken from the middle of the branch, and plot (c) is taken toward the end of the branch. Plots (d), (e), and (f) shows the growth and evolution of the invariant surface for a constant ω_1 3-d branch. This branch has constant frequencies $\omega_0 = 2.666571$ and $\omega_1 = 1.698082$. Plot (d) is taken near the beginning of the branch, plot (e) is taken from the middle of the branch, and plot (f) is taken toward the end of the branch. The constant ω_2 3-d branch has an increase in Jacobi constant along the branch, while the constant ω_1 3-d branch has a decrease in Jacobi constant. Figure 3.46 shows other individual invariant surfaces we found interesting.

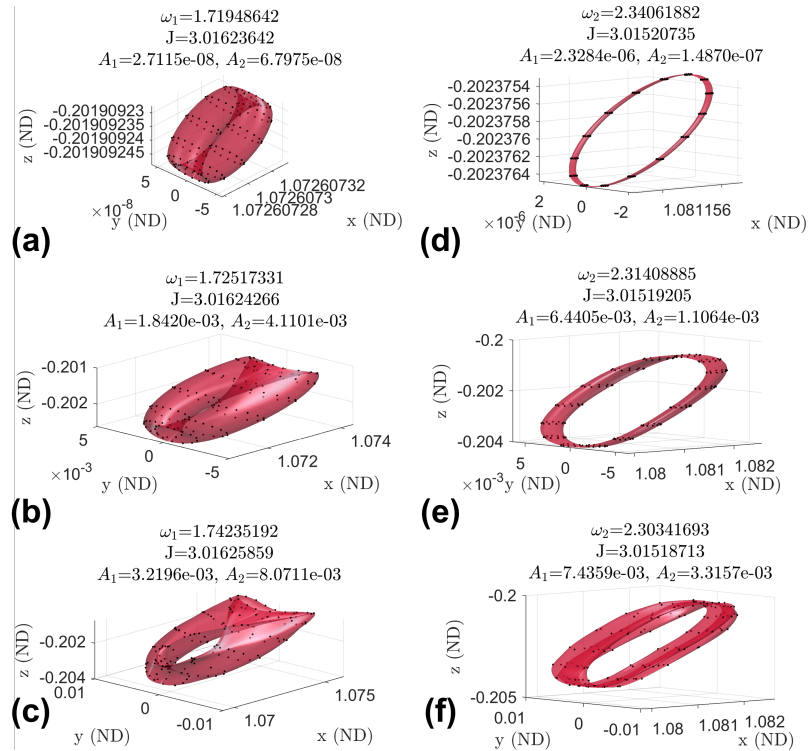


Figure 3.45: Invariant surfaces within a constant ω_2 3-d branch (a-c) and within a constant ω_1 3-d branch (d-f).

It should be noted that most of the invariant surfaces of the 3-d quasi-halo orbits are found to be self-intersecting, and thus the entire orbit is composed of trajectories constantly crossing over

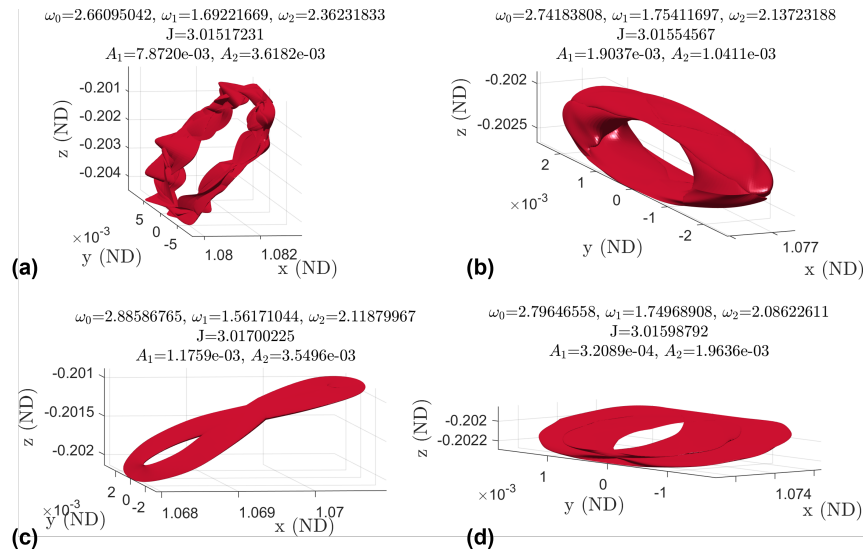


Figure 3.46: Survey of invariant surfaces in the 3-dimensional quasi-halo family.

each other. This observation calls for care and detailed analysis when placing multiple spacecraft on a 3-dimensional quasi-halo orbits to ensure a collision will not occur. However, in light of this, the abundance of intersections means there are boundless opportunities to change the phasing on the orbit.

3.6.7 Relative Motion

The static images of the invariant surfaces give insight into what the orbits look like, however they do not provide information about the motion of the invariant surface in time. To show this behavior the points constructing the invariant surfaces of 415 quasi-halo orbits have been propagated for three stroboscopic times and the relative states with respect to their underlying halo orbit have been examined.

Figure 3.47 shows one example that captures the typical behavior of the relative motion. Plot (a) shows the color map used to identify each point on the invariant surface where each coordinate pair (θ_1, θ_2) has a unique color. Plot (b) shows the invariant surface colored according to the color map. Plot (c) shows in configuration space the relative motion about the stable halo orbit which the orbit was generated from. Plot (d) shows the time history of the distance to the halo orbit for

each point.

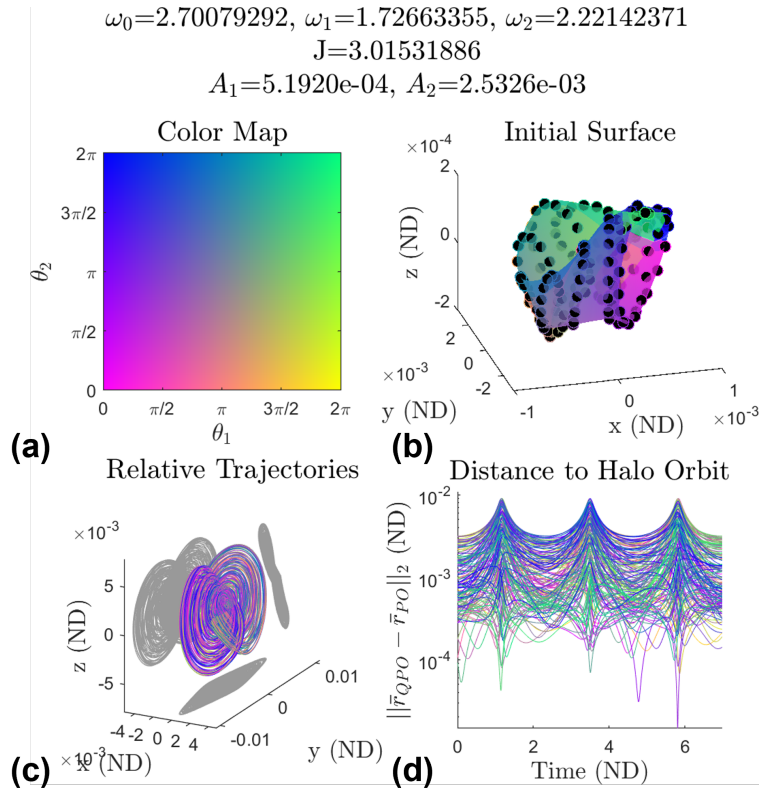


Figure 3.47: Representative example of the relative motion for the 3-d quasi-halo orbits.

The initial invariant surface of plot (b) is embedded in the relative motion in plot (c). The results in plot (c) show that the invariant surface gets warped and stretched in the y and z directions while not much stretching occurs in the x direction. The plot also shows that there is full coverage of the halo orbit by the invariant surface. This shows that these orbits are good to place surveillance satellites to keep watch on an object on the underlying halo orbit.

Combining the information of plot (c) and plot (d) it is surmised that at half the stroboscopic time the maximal stretching occurs while also achieving its closest approach to the halo orbit meaning the invariant surface is largest at this point in time. The smallest spread in distances occurs when the invariant surface is furthest from the Moon meaning it assumes its smallest size. The point in time when the largest spread in distance occurs is when the invariant surface is

making its closest approach to the Moon. And similarly the point in time when the smallest spread in distance occurs is when the invariant surface is furthest from the Moon. This behavior is seen among most of the tested elliptic quasi-halo orbits, and it coincides with where the invariant curves of the 2-d quasi-halo orbits in this region achieve their largest and smallest amplitudes (refer to Figure 3.35).

Plot (d) shows that the distance to the halo orbit remains positive. This result shows that there exists a ball with a radius of about $1e-5$ for which the invariant surface does not penetrate. So spacecraft on this invariant surface will not collide with a spacecraft on the underlying halo orbit. However, if there is a displacement along the θ_0 direction between the invariant surface and a point on the halo orbit, then the analysis will have to be repeated to ensure there are no crossings with the underlying halo orbit. It is shown in Section 3.5.4 that the underlying halo orbits usually penetrate the surfaces of the 2-d quasi-halo orbits, so it is reasonable to assume that the underlying stable halo orbits penetrate the elliptic quasi-halo orbits. The difference between the penetrations is that for the 2-d quasi-halo orbits there are a finite number of penetrations since a line is crossing through an infinitely thin surface. However, the elliptic quasi-halo orbits fill a volume, so there would be an infinite number of penetrations along the crossings.

3.6.8 Error Analysis

As mentioned in Section 3.6 the tolerance level used to compute all the invariant surfaces and curves in this work is $7e-11$. Since each point incurs some amount of error then each point is actually on a different quasi-periodic torus with different frequencies. This means there is a difference between the actual and idealized orbits of the discrete points representing an invariant surface. To measure this difference the points \mathbf{X} of a 3-dimensional quasi-halo invariant surface are propagated forward with the stroboscopic map $x \mapsto \phi_T(x)$. The points $\phi_T(\mathbf{X})$ are rotated by the rotation matrix $R_{-\rho}$ to get the states of where the points should be on the invariant surface. The states should coincide with \mathbf{X} , so the error is the norm of the difference between these quantities. The points $\phi_T(\mathbf{X})$ are then mapped to $\phi_{2T}(\mathbf{X})$ and rotated by the rotation matrix $R_{-2\rho}$. The

error is calculated as the norm of the difference between the newly rotated points and the original points \mathbf{X} . Recall that the vector \mathbf{X} is a vector comprised of points in phase space \mathbf{x}_j that represent points on the invariant surface of a quasi-periodic orbit. Then the equation for the error for point \mathbf{x}_j is given in Equation (3.43). This process is repeated for a total of five mappings and the results for six different quasi-halo orbits are given in Figure 3.48.

$$e_j(n) = \|R_{-n\rho}\phi_{nT}(\mathbf{x}_j) - \mathbf{x}_j\| \quad (3.43)$$

In Equation (3.43) when $n = 1$ then this is equivalent to the quasi-periodicity constraint of Equation (3.3) divided by the number of points representing the surface. When $n = 1$ the error in Figure 3.48 is below the error tolerance $7e-11$, showing that the quasi-periodicity constraint is indeed being satisfied. However, as \mathbf{X} is mapped further in time the errors grow, showing that the actual orbits diverge from the idealized orbits. In plots (a), (b), and (c) the errors grow to about $1e-8$ after two stroboscopic maps and steadily grow for the following mappings. These orbits are from constant ω_2 3-d branches. In plots (d), (e), and (f) the errors grow to about $1e-2$ after two stroboscopic mappings and stay nearly constant for the following mappings. These orbits are from constant ω_1 3-d branches.

3.6.9 Conclusion

This work uses a single-parameter continuation method to explore the dynamical structure in the vicinity of the L_2 stable halo orbits in the Earth-Moon system of the circular restricted three-body problem. The types of solutions explored are the quasi-halo orbits diffeomorphic to 2- and 3-dimensional quasi-periodic tori. Two branches from each types of orbits are computed from the span of stable halo orbits to construct a total of four 2-parameter families. The four branch types are the constant ω_2 2-d branch, the constant ω_1 2-d branch, the constant ω_2 3-d branch, and the constant ω_1 3-d branch. In each branch a parameter is held constant and the value of that parameter is determined by the halo orbit and its monodromy matrix from which the branch is grown.

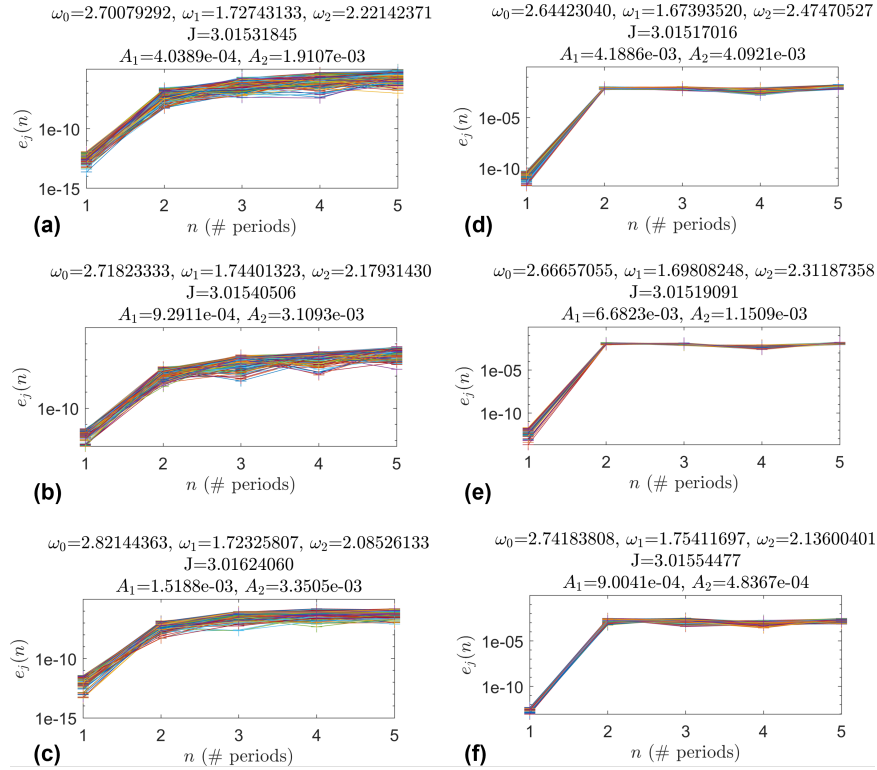


Figure 3.48: The error over time between six computed invariant surfaces and the propagated surfaces.

The 3-d quasi-halo orbits do not have amplitudes as large as the 2-d quasi-halo orbits, but are larger than some 2-dimensional elliptic quasi-halo orbits. However, all the orbits in this region are much smaller than the quasi-halo orbits emanating from the unstable quasi-halo orbits. The amplitudes of all the orbits are compared to an empirical limit to the sizes of orbits in this region.

The Jacobi constant among each of the four branch types has a net change in the range $[-4e-4, 3e-4]$ compared to the value of the Jacobi constant of the underlying halo orbit from which each branch is generated. The direction of change has a turning point within the constant ω_2 2-d and constant ω_2 3-d families. The branches with smaller ω_0 values have a negative change in the Jacobi constant as orbits grow larger. And branches with larger ω_0 values have a positive change in the Jacobi constant as orbits grow larger. The constant ω_1 2-d and constant ω_1 3-d branches exhibit an increase in the Jacobi constant as the orbits grow larger.

The 2-d quasi-halo orbits are comprised of elliptic and partially hyperbolic quasi-halo orbits. Nearly each branch is identified to have a stability bifurcation leading from stable to unstable orbits. All identified bifurcation points are near the ends of branches. The partially hyperbolic orbits do not transition back to elliptic before the end of the branch is reached.

Lastly, the geometry of the invariant surfaces of these quasi-halo orbits are diverse. The invariant surfaces of the elliptic quasi-halos create a volume in phase space when evolved in time resulting in a trajectory with complex behavior. The relative motion of the invariant surfaces with respect to the underlying halo orbit provides full coverage of a point on the halo orbit with the same phasing in θ_0 .

Chapter 4

Resonance Avoidance Methods

“The dynamical properties of an invariant torus with linear flow are very sensitive to the number-theoretical properties of its frequency vector.”

- Broer, Huitema, and Sevryuk [39]

4.1 Introduction

Any work in the area of KAM theory points out the problem of small-divisors which presents itself in KAM type proofs on the existence of invariant tori. It is impractical to present an exhaustive list of references to learn about KAM theory, invariant tori, reducibility, and related material, however the reader is directed to the textbooks and their references in [11, 12, 39, 38, 134, 56]. The reader is also referred to the following research papers and their references in [51, 61, 62, 100, 106, 105, 107, 108, 151, 156, 157, 80].

It is well known that resonances between the internal frequencies of invariant tori play a vital role in the theoretical study and numerical computation of quasi-periodic invariant tori. Greene [77] and MacKay [128] show invariant curves of 2-dimensional invariant tori of flows are most robust when the rotation number of the curve is a noble number. As the rotation number moves away from a noble number and approaches a rational number the robustness diminishes. At a resonance the frequencies of an invariant torus violate the non-resonance condition in Equation (2.2), and the dimension of the torus decreases, resulting in a submanifold of lower-dimensional invariant tori [39].

For a 2-dimensional quasi-periodic invariant torus with frequencies near a resonance the invariant curve no longer covers \mathbb{T} , and resembles an island chain on a Poincaré map [136], a behavior which extends to higher-dimensional quasi-periodic invariant tori. Both phenomena are destructive to the computation of quasi-periodic orbits.

The complications of computing invariant tori near resonances is numerically observed by Gómez and Mondelo in [75] wherein they compute a 2-parameter family of 2-dimensional quasi-periodic orbits. To avoid these challenges, they tactfully choose rotation numbers such that $\frac{2\pi}{\rho}$ is an integer plus the golden number $\frac{1}{2}(1 + \sqrt{5})$, leading to the most irrational set of frequencies. Gómez and Mondelo enforce the continuation to maintain this rotation number to ensure the family of computed quasi-periodic orbits exist in the nonlinear system.

McCarthy and Howell present a method to avoid resonances in a 3-parameter family of 3-dimensional quasi-periodic orbits [130]. They suggest fixing the rotation numbers such that they take values which are not multiples of 2π . Their method ensures the ratios of the frequencies stay away from integer values. They do not address the issue that many resonances occur when the ratios of the frequencies are not integers. Picking the rotation numbers to ensure the frequency ratios avoid rational numbers is a more challenging task. However, avoiding resonances is exactly one of the issues addressed in this chapter.

Both resonance avoidance methods described above seem sufficient for studying large portions of families of quasi-periodic orbits, but they are restrictive in the types of branches which can be computed, and do not allow the parameter space to be explored freely. Utilizing a combination of the methods presented in Chapter 3 allows one to explore parameter spaces with more complicated geometries. The issue of resonances in families of quasi-periodic orbits must be addressed to freely explore the parameter space and is the motivation behind this chapter.

This chapter is organized as follows: Section 4.2 discusses the problem of identifying a resonance between orbit frequencies. Section 4.3 presents a line search method for picking the most irrational set of orbit frequencies within a range of continuation step sizes. Section 4.4 presents a heuristic method for determining when a continuation scheme should increase the step size to

jump beyond a resonance. Section 4.5 analyzes the performance of the methods on a test example. Section 4.6 discusses the impact of resonances on the computations of quasi-periodic orbits and provide advice for dealing with resonances in computing branches of quasi-periodic orbits. Lastly, Section 4.7 concludes the chapter.

4.2 Irrational Numbers in Floating Point Arithmetic

To efficiently compute branches of quasi-periodic orbits in arbitrary directions we devise a methodology to avoid resonances. To avoid resonances it is necessary to identify when the torus frequencies are near a resonance. Identifying a resonance on a computer equates to determining whether the ratio of two frequencies in floating point arithmetic is well approximated by a rational number. Determining the irrationality of a frequency vector on a computer is not a trivial task since all floating point numbers are necessarily rational. While this is not a well-defined problem we leverage research that addresses the issue of detecting resonances on a computer.

A question addressed by Sander and Meiss in [152] is: Given a floating point number x and an interval $I_\delta = [x - \delta, x + \delta]$ what is the rational $\frac{n}{d}$ in this interval with the smallest denominator? They mention there are built in functions, such as in Mathematica and MATLAB, which return an n and d , but do not return the smallest denominator. Therefore, they present an algorithm which does return the smallest denominator such that $\frac{n}{d} \in I$. We use their algorithm, given in Algorithm 4, to determine the approximate rationality between the torus frequencies.

Sander and Meiss perform tests on a set of floating point numbers to determine which values of d make for good approximations to x . They found that $\log_{10}(d)$ has a mean value of $\log_{10}(\delta)/2$, and when $\log_{10}(d)$ deviates from the mean value then x is well approximated by a rational number. Their results lead to the development of the metric

$$\sigma = |\log_{10}(d) + \log_{10}(\delta)/2| \tag{4.1}$$

to rate how “rational” x is [152]. In their work they use $\delta = 10^{-8}$, and deem a point x to be rational when the log deviation σ is greater than 0.3375. We use the same δ value in this thesis, but use

Algorithm 4 Compute the smallest denominator of the rational number approximating a floating point number within an interval

```
1: procedure SMALLDENOM( $x, \delta$ )
2:    $(n, d) = (p_l, q_l) = (0, 1)$ 
3:    $(p_r, q_r) = (1, 0)$ 
4:   while  $|x - \frac{n}{d}| \geq \delta$  do
5:      $(n, d) = (p_l + p_r, q_l + q_r)$ 
6:     if  $x < \frac{n}{d}$  then
7:        $(p_r, q_r) = (n, d)$ 
8:     else
9:        $(p_l, q_l) = (n, d)$ 
10:    end if
11:  end while
12:  return  $(n, d)$ 
13: end procedure
```

different values of σ to determine when an orbit is near a resonance. It is not straightforward to determine what σ should be because each resonance has a region around it where the computation of orbits becomes difficult; the size of this region is different for each resonance.

4.3 Picking Irrational Frequencies

In a continuation method to compute branches of quasi-periodic orbits we have a current solution $\mathbf{z}_k = [\mathbf{X}_k^T, \boldsymbol{\omega}_k^T]^T$ and a family tangent vector \mathbf{z}'_k which governs the prediction of the next solution \mathbf{z}_k . Suppose a nominal step size Δs for the prediction is used, such as Equation (2.70), then the nominal set of frequencies is

$$\boldsymbol{\omega}_{k+1} = \boldsymbol{\omega}_k + \Delta s_k \boldsymbol{\omega}'_k. \quad (4.2)$$

The nominal frequencies may be nearly commensurate, thus making the next torus numerically challenging to compute. However, within any open interval around any number there are irrational numbers. Leveraging this information we conduct a line search to pick a new step size within the interval $[0.9\Delta s_k, 1.1\Delta s_k]$ so that the new predicted set of frequencies is the “most” irrational.

For a given value of $s \in [0.9\Delta s_k, 1.1\Delta s_k]$ we compute a set of predicted frequencies $\boldsymbol{\omega}_p = \boldsymbol{\omega}_k + s\boldsymbol{\omega}'_k$. For each combination of ratios between the frequencies, excluding the inverse of the ratios, we rate the rationality of each ratio using `SMALLDENOM` and Equation (4.1). These values are stored for each combination of ratios and for each set of predicted frequencies, producing a matrix of σ values. The best step size is the one that minimizes the average σ value for a fixed value of s . Picking the best step size this way minimizes the likelihood of being near a resonance, promoting well-conditioned computations of invariant tori. The method to choose the best step size to produce the most irrational set of frequencies is given in Algorithm 5.

We note that Algorithm 4 is not well-behaved over the interval $[0.9\Delta s_k, 1.1\Delta s_k]$ (see Figure 13 in [152]), so locating the optimal step size is reduced to performing a grid search with 1000 values of s equally spaced throughout the interval.

An example of using Algorithm 5 is given in Figure 4.1. The black square represents the

Algorithm 5 Compute best step size to avoid being near a resonance

```

1: procedure IRRATIONALSTEPSize( $\Delta s, \boldsymbol{\omega}, \boldsymbol{\omega}', \delta$ )
2:   for all  $s \in [0.9\Delta s, 1.1\Delta s]$  do
3:      $\boldsymbol{\omega}_p = \boldsymbol{\omega} + s\boldsymbol{\omega}'$ 
4:     for all unique combinations of frequencies  $\omega_{p,i}, \omega_{p,j}$  in  $\boldsymbol{\omega}$  do
5:        $\hat{r} = \frac{\omega_{p,i}}{\omega_{p,j}}$ 
6:        $(n, d) = \text{SMALLDENOM}(\hat{r}, \delta)$ 
7:        $\sigma_{(i,j),s} = |\log_{10}(d) + \log_{10}(\delta)/2|$ 
8:     end for
9:   end for
10:   $\sigma_s = \mathbb{E}_{(i,j)}[\sigma_{(i,j),s}]$  ▷ Take the mean over the  $(i, j)$  dimension of the array.
11:   $\Delta s_{\min} = \arg \min_s \sigma_s$ 
12:  return  $\Delta s_{\min}$ 
13: end procedure

```

frequencies of the current solution. The cyan diamond represents the nominal frequencies using the nominal continuation step size. One thousand predicted frequencies are computed within the interval $[0.9\Delta s, 1.1\Delta s]$. Each pair of predicted frequencies are evaluated using Algorithm 4 and Equation (4.1), and are colored according to their σ value. The blue star is the frequency pair which attains the minimum value of σ in the grid search. The step size corresponding to that frequency pair is the step size chosen for the prediction step. Had the nominal step size been slightly larger the nominal set of frequencies could have been considered rather rational, possibly requiring more Newton iterations to converge.




Figure 4.1: Example of picking step size to minimize chance of stepping near a resonance. Plot (b) is a zoomed in copy of plot (a).

4.4 Hopping Beyond Resonances

With a step size chosen nominally as in Equation 2.70 it is easy for the continuation of quasi-periodic orbits to get stuck near a resonance. As the continuation approaches a strong resonance more iterations will necessarily be needed to converge to a true solution. By taking more iterations to converge the step size will decrease and the frequencies of the quasi-periodic orbits will ever so slowly approach the resonance, leading to stagnation and termination in the continuation procedure.

To circumvent the issue of premature termination due to resonances one needs a way of

identifying when the continuation is stagnating at a particular resonance. From there we can make an informed decision on how to adjust the step size. Tracking the σ values alone to determine the likelihood of being near a resonance is not sufficient. The σ value is only able to detect when the torus frequencies are near a resonance, however the continuation of a branch of quasi-periodic orbits may be progressing sufficiently and just happen to converge with frequency ratios near different rational values.

A method to hop beyond resonances requires not only tracking the σ values, but also requires tracking the ratio of the frequencies to determine progress in the continuation. Therefore, we present the following heuristic method to identify and jump past resonances in Algorithm 6 and Algorithm 7. Let ω_k be the frequencies of the current converged solution with the new nominal step size Δs_k computed according to Equation (2.70). Let \mathbf{r} be a vector which records the value of each combination of frequency ratios, and let \mathbf{c} be a vector of positive integers which keeps a count of the consecutive occurrences when each frequency ratio is near a resonance. At the start of the continuation, for $k = 0$, \mathbf{r} is composed of the values of frequency ratios for the initial quasi-periodic orbit, while \mathbf{c} starts at all zeros. Once the next solution \mathbf{z}_{k+1} has been found then Algorithm 7 is called to determine whether the step size should be increased to move beyond a resonance.

Algorithm 6 Determine if continuation of quasi-periodic orbits are near a resonance

```

1: procedure CHECKRESONANCES( $\mathbf{r}, \mathbf{c}, \omega, \delta$ )
2:    $k = 1$ 
3:   for all unique combinations of frequencies  $\omega_i, \omega_j$  in  $\omega$  do
4:      $\hat{r} = \frac{\omega_i}{\omega_j}$ 
5:      $(n, d) = \text{SMALLDENOM}(\hat{r}, \delta)$ 
6:      $\sigma = |\log_{10}(d) + \log_{10}(\delta)/2|$ 
7:     if  $\sigma > 0.1$  and  $|\mathbf{r}[k] - \frac{n}{d}| < 1e - 3$  then
8:        $\mathbf{c}[k] = \mathbf{c}[k] + 1$ 
9:     else
10:       $\mathbf{c}[k] = 0$ 
11:    end if
12:     $\mathbf{r}[k] = \hat{r}$ 
13:     $k = k + 1$ 
14:  end for
15:  return  $\mathbf{r}, \mathbf{c}$ 
16: end procedure

```

Algorithm 7 first calls Algorithm 6 to update \mathbf{r} and \mathbf{c} . Algorithm 6 goes through each combination of frequency ratios of ω_{k+1} , computes the rational with the smallest denominator using Algorithm 4, and computes σ . For a given frequency ratio we increase the counter associated with that ratio by one if two conditions are met. The first is that σ needs to be sufficiently large to indicate the presence of a resonance. Trial and error has resulted in a threshold value of 0.1. The second condition is the frequency ratio of the current continuation solution must be sufficiently close to the frequency ratio of the previous continuation solution. We let the threshold value be 1e-3. If either of these conditions are not met then the counter associated with that frequency ratio is reset to zero. The new values of the frequency ratios are recorded and compared with the next continuation solution.

Algorithm 7 Decide whether to increase step size to hop a resonance

```

1: procedure HOPRESONANCE( $\mathbf{r}, \mathbf{c}, \omega, \Delta s, \delta$ )
2:    $(\mathbf{r}, \mathbf{c}) = \text{CHECKRESONANCES}(\mathbf{r}, \mathbf{c}, \omega, \delta)$ 
3:   if any element of  $\mathbf{c}$  is equal to 5 then
4:      $\Delta s = 0.05 * \text{sign}(\Delta s)$ 
5:   end if
6:   return  $\mathbf{r}, \mathbf{c}, \Delta s$ 
7: end procedure

```

Algorithm 7 uses \mathbf{c} to decide if the continuation has stagnated at a resonance. If any value in \mathbf{c} has a value of 5, indicating the continuation has been near a particular resonance for 5 consecutive solutions, then we change the step size to be 0.05 multiplied by the sign of the step size. Typical continuation methods only have a positive step size so that the continuation only marches forward along a branch of solutions, but as we will see in Chapter 5, there is use in negative step sizes, so we respect the sign of the step size here. A step size magnitude of 0.05 has been deemed sufficient to be able to move beyond many resonances while also providing for good linear approximations in the prediction step of continuation. However, the value 0.05 should not be taken as absolute and is likely dependent on many factors such as the family of orbits, the dynamical system, and the dimension of the invariant tori. As it has been mentioned the method presented in this section is a heuristic method, and, as such, many values have been chosen to provide favorable performance

for the applications encountered in this thesis.

4.5 Results of Implementation

We tested the performance between the use of both resonance avoidance methods together against using neither method. In the former case we call it the RAv variant, and in the latter case we call it the plain variant. The plain variant uses only the nominal step size update given in Equation (2.70). The nominal step size update is also used in the RAv variant with the inclusion of Algorithm 5 to make little modifications to the step size and with the inclusion of Algorithm 7 to prevent stagnation at a resonance.

For the test we chose 7 quasi-halo orbits from within the Earth-Moon L_2 quasi-halo orbit family to act as initial solutions for continuing branches of solutions. About each initial solution 15 directions equally spaced out on a circle are chosen for the branches to follow. The direction is enforced using the constant slope constraint in Equation (3.41). The initial family tangents are chosen according to the method outlined in Section 3.5.1. Each branch of solutions terminates when getting to a target set of frequencies. The specific method to target the frequencies is given in Chapter 5, and defer details until then, however this does not impact the results of the test.

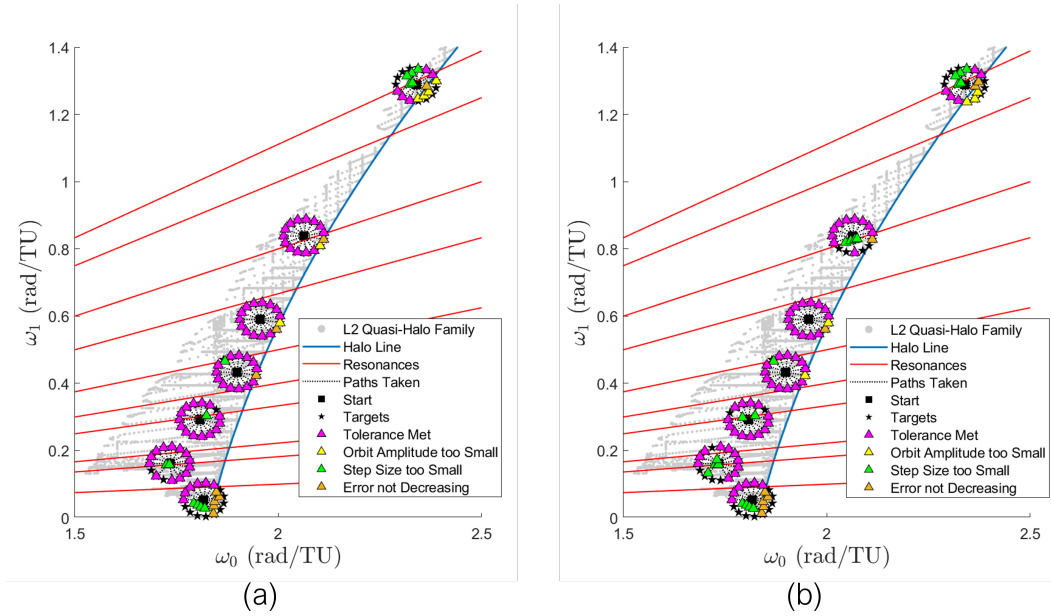


Figure 4.2: Comparing the solutions in frequency space between the resonance avoidance variant (a) to the plain variant (b).

All together there are 105 branches of solutions spanning a large portion of the unstable 2-dimensional quasi-halo orbits for both variants. The branches are faced with resonances to get to the target set of frequencies. We are looking to see how many of the branches successfully reach the target set of frequencies, how much time the branch took to get there, and how many family members were computed along the way. Examining these three quantities gives insight into the performance of the resonance avoidance methods. Statistics relating to these three metrics are in Table 4.1.

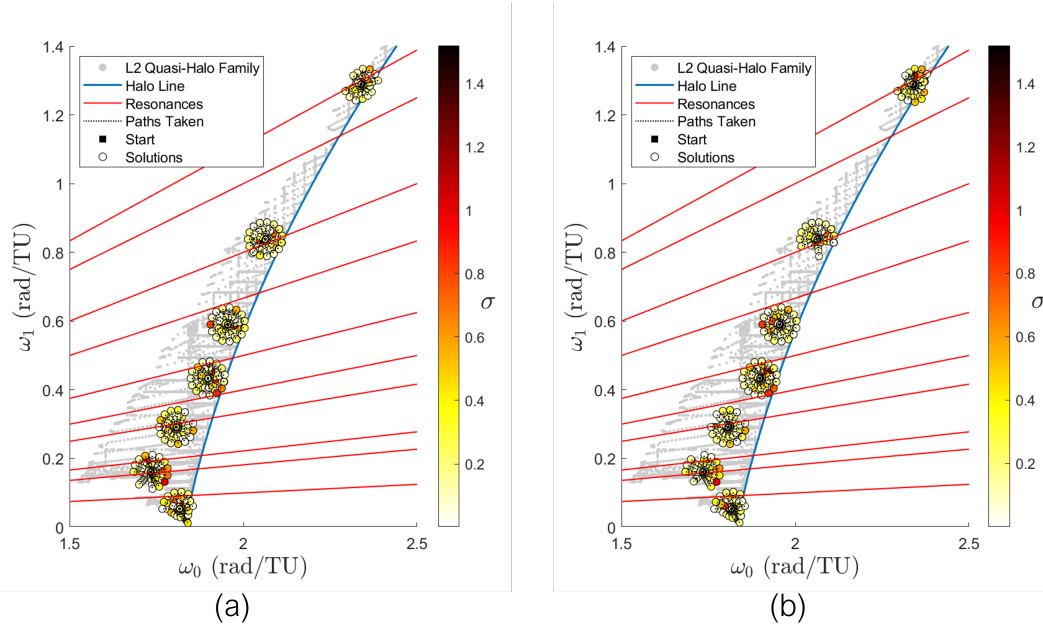


Figure 4.3: Comparing the branches in frequency space between the resonance avoidance variant (a) to the plain variant (b).

Figures 4.2 -4.4 show the performance of the RAv variant compared to the plain variant. The results of the convergence conditions are in Figure 4.2. In frequency space, Figure 4.2 shows the family of quasi-halo orbits, the starting solutions, the target frequencies, and the final solution in each branch. Each final solution is colored according to the reason the branch terminated. Select resonance lines are laid on top of the family to help show interactions between the branches and resonances. Figure 4.3 also shows the quasi-halo family in frequency space with resonance lines laid on top and the starting solutions. In contrast to Figure 4.2, Figure 4.3 shows the individually computed solutions in each branch colored according to the σ value in Equation (4.1).


The RAv variant has 77 branches which converge to solutions with the target set of frequencies, 22 branches which reach the boundary of the family, and 6 branches which are stopped by resonances. Therefore, there are 99 successful branches and 6 unsuccessful branches. The plain variant has 70 convergent branches, 21 which reach the boundary of the family, and 14 that are stopped by resonances, showing the RAv variant performs better at moving past resonances.

Table 4.1: Performance of RAv and Plain variants.

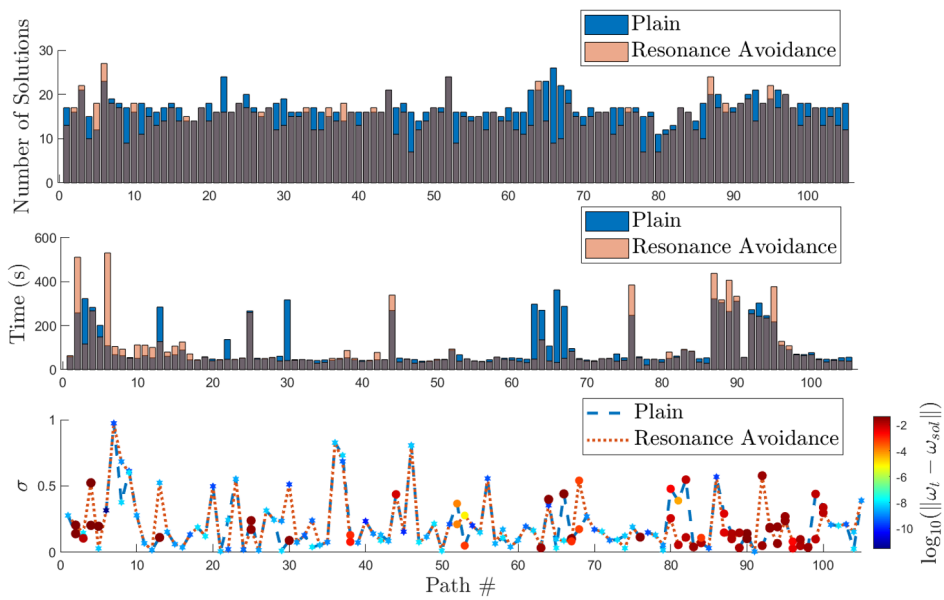
Variant	RAv	Plain
# Hit Converge	77	70
# Hit Boundary	22	21
# Hit Resonance	6	14
Avg. Time per Branch	98 sec	102.5 sec
Avg. # QPOs per Branch	15	17
Avg. Time per QPO	6.5 sec	6 sec

For the RAv variant the average time per branch is 98 seconds, the average number of solutions per branch is 15, and the average time per solution is 6.5 seconds. For the plain variant the average time per branch is 102.5 seconds, the average number of solutions per branch is 17, and the average time per solution is 6 seconds. The results show the RAv variant increases the average time per solution, but decreases the average number of solutions per branch and has an overall decrease in the average time per branch.

A graphical comparison between the resonance avoidance and plain variants are in Figure 4.4. The plot (a) of the figure is a key to identifying the branch number that is on the x -axis of the three subplots of plot (b). The branch numbers start at 9 o'clock relative to the initial quasi-periodic orbits and increase counterclockwise. It is clear from the top subplot of plot (b) that the plain variant tends to compute more solutions compared to the RAv variant. The bottom subplot of plot (b) shows two lines showing the σ value of the returned solution and colored according to the error between the returned solution and the targeted solution. The long dotted line is for the plain variant and the short dotted line is for the RAv variant. The data points are either pentagrams, showing convergence is achieved, or dots whose color and size are dependent on the error. The two lines are aligned for most of the plot, however there are branches which lead to the two lines diverging from one another.



(a)



(b)

Figure 4.4: Comparison of performance between the plain variant and the resonance avoidance variant (b). The x -axis is the branch number. The branch numbers are given in plot (a).

It is interesting to see that many of the high σ value solutions achieved convergence. Moreover, it is interesting to see there are gaps between the two lines at converged solutions, showing a large variability in σ for very small differences in the returned solutions. These observations suggest a high σ value does not measure the strength¹ of a resonance, so in many cases there are no issues computing quasi-periodic orbits with a nearly commensurate set of orbit frequencies. One might wonder why many of the branches have a returned solution with a low σ value, but a high error. If the reader correlates the branch numbers to the key in plot (a) and the convergence results in Figure 4.2 it is evident that the error comes from the fact that the target frequencies are outside the boundary of the family. The hope was there would be a clear correlation between high σ and error, but this is not observed.

4.6 Discussion

There are two main reasons for branches getting stuck at resonances. The first main reason is the target frequencies may be too close to a resonance for the branches to converge. Orbit frequencies close to a resonance pose several complications. First, more Fourier coefficients are needed to increase the accuracy of computations. Second, recall that quasi-periodic orbits near a resonance exhibit island chain behavior, preventing the quasi-periodicity constraint to be satisfied. Lastly, recall that families of quasi-periodic orbits form Cantor families, meaning orbits simply do not exist for every set of frequencies. It may be the case that the target set of frequencies lie in an area that do not allow for the existence of quasi-periodic orbits. For these reasons it is advised to ensure the target frequencies do not lie too close to a low-order resonance.

The second main reason for branches getting stuck at resonances is the path may not have been able to jump past the resonance to reach the target frequencies. This could be that the gap is too large, meaning the family tangent vector does not provide for an accurate prediction of the next quasi-periodic orbit in the continuation. The large gap issue can be circumvented if there are pre-

¹ The quantity σ is purely a number theoretic quantity and does not capture the dynamical aspect of resonances. Each resonance has its own region, around which, invariant tori cease to exist. The larger the region the “stronger” the resonance.

computed quasi-periodic orbits on the other side of the resonance sufficiently close to the predicted set of frequencies. Then the guess of the quasi-periodic orbit states can be replaced with the states of the closest known quasi-periodic orbit while maintaining the predicted set of frequencies.

Lastly, we observe that near-parallel crossings of resonances are difficult, while near-perpendicular crossings of resonances are easier. The direction of crossing makes a difference because there is an area associated with each resonance where computations of quasi-periodic orbits becomes challenging or even impossible. The size of these areas are unknown, but generally lower-order resonances have larger areas. It is evident that crossing in a perpendicular manner reduces the distance that a continuation method has to move from one side of the resonance to the other to continue computing solutions. Leveraging this, one can construct a method to hop past resonances perpendicularly to reduce the chance of branches terminating prematurely.

We present a method to determine the correct family tangent vector to cross perpendicularly in Algorithm 8, and describe how to alter the continuation method to perform the crossing. Let $\mathbf{z} = [\mathbf{X}^T, \boldsymbol{\omega}^T]^T$ be the current solution and V_p be a matrix whose columns contain the principle tangent basis vectors. Let $m_0 > 0$ be the resonance between the frequencies ω_i and ω_j , such that $m_0 \approx |\frac{\omega_i}{\omega_j}|$. The pair of frequencies near a resonance can be determined from Algorithm 6. To cross perpendicularly to the line $y = m_0x$ we need to direct the continuation to follow a line with the slope $m = -\frac{1}{m_0}^2$. We then use Equation (3.42) to pick the correct family tangent vector to move along this line. The last consideration is to move forward in the sense of increasing ω_j or backward in the sense of decreasing ω_j . The direction is determined from the previous solution $\tilde{\mathbf{z}}$. Let $\Delta\omega_j = \omega_j - \tilde{\omega}_j$. If $\Delta\omega_j > 0$ the direction is forward, if $\Delta\omega_j < 0$ the direction is backward, and if $\Delta\omega_j = 0$ then we need to examine $\Delta\omega_i$. If $\Delta\omega_i < 0$ the direction is forward, and if $\Delta\omega_i > 0$ the direction is backward. After the correct family tangent vector has been computed from Algorithm 8 we simply use the constant slope constraint in Equation (3.41) with ω_i taking the place of ω_1 and ω_j taking the place of ω_0 .

After the resonance has been hopped perpendicularly the user can determine what to do

² This result should be familiar from basic algebra.

next. It is likely that before the crossing the user was computing a branch of orbits with a specific characteristic. If the characteristic is the orbit frequencies then the tool developed in Chapter 5 will suffice to get to the desired orbit. However, Chapter 6 presents a tool for getting from an initial solution within a multi-parameter family to a desired solution based on general characteristics.

4.7 Conclusion

Two novel methods for picking step sizes in continuation methods which compute quasi-periodic orbits are developed and tested. One method chooses a step size within an interval centered on a nominal step size to predict a quasi-periodic orbit which has the most irrational frequencies. The other method greatly increases the step size to jump beyond resonances. Both methods leverage number theoretic properties to compute more robust quasi-periodic orbits. These two methods fall into a class called resonance avoidance methods which aim to improve the performance of computing quasi-periodic orbits by avoiding resonances between the internal frequencies of the tori. Both methods focus on the step size rather than the parametric constraint, allowing for flexibility in the computed branches of orbits. The combined results of these new methods show improved performance for the computation of quasi-periodic orbits by decreasing the time to compute branches of orbits and increasing the probability of moving past resonances which would ordinarily halt the continuation process. Lastly, we discuss the impacts of resonances on the computations of quasi-periodic orbits and present ideas to circumvent resonances.

Algorithm 8 Generate family tangent vector to hop over a resonance perpendicularly

```

1: procedure PERPRESHOP( $\omega_i, \omega_j, \tilde{\omega}_i, \tilde{\omega}_j, V_p, \delta$ )
2:    $r = \frac{\omega_i}{\omega_j}$ 
3:    $(n, d) = \text{SMALLDENOM}(r, \delta)$ 
4:    $m = -\frac{d}{n}$ 
5:    $\Delta\omega_j = \omega_j - \tilde{\omega}_j$ 
6:    $\Delta\omega_i = \omega_i - \tilde{\omega}_i$ 
7:   if  $\Delta\omega_j$  equals 0 then
8:     if  $\Delta\omega_i < 0$  then
9:        $\alpha = 1$ 
10:    else
11:       $\alpha = -1$ 
12:    end if
13:  else
14:    if  $\Delta\omega_j < 0$  then
15:       $\alpha = -1$ 
16:    else
17:       $\alpha = 1$ 
18:    end if
19:  end if
20:   $z' = \alpha(V_p[:, j] + mV_p[:, i])$ 
21:  return  $z$ 
22: end procedure

```

Chapter 5

A Retraction on Families of Quasi-Periodic Orbits

5.1 Introduction

Trajectory design of 1-parameter families of solutions is fairly straightforward since there is only one dimension in the space of solutions to search: forward and backward. Multi-parameter families introduce a larger solution space and increase the complexity of finding solutions since there are multiple dimensions to search. Most studies utilizing methods to compute quasi-periodic orbits don't focus on the computation of specific family members, but instead are used to research entire families of orbits. Knowing the characteristics and limitations of each orbit family is the first step in determining which orbit types should be used in designing a mission. The next step is finding explicit members from the candidate families to provide options that meet initial mission constraints. Quasi-periodic orbit families generally exist in multi-parameter families in astrodynamics models, and the ability to search through these higher dimensional spaces to pick out a specific orbit is not trivial.

Research investigating multi-parameter quasi-periodic orbit families usually relies on the use of single-parameter continuation where the continuation is restricted to follow some rule based on the hard coded constraints. In these cases, the computed family members primarily serve to be representative members from the family to gain insight into the family as a whole. The ability to methodically move through the family to find a specific family member is not existent to our knowledge. In 2001 Gómez and Mondelo presented the idea of computing a periodic orbit with given energy level and a prescribed value of a coordinate [75], however the idea is applied to the

refinement of an orbit, so one must have an (approximated) orbit close to the desired one.

In 2012 Mondelo et al. develop a way to interpolate over a 2-parameter family of quasi-periodic orbits allowing them to specify orbits by their energy and rotation number [136]. In many areas of the parameter space they achieve interpolation errors below $1e-8$, and they encounter larger errors near resonances and the boundary of the family as high as 0.01. To improve the errors they take a finite elements approach to reduce the errors down to $1e-10$. Their interpolation method only works when the interpolation nodes are nearly equally spaced in the parameter space. According to their paper, “the rotation numbers have been ‘Diophantised’ in order to be far from low-order resonances.” A grid of 25433 orbits are computed for the L_1 family of Lissajous orbits, and a grid of 73939 orbits are computed for the L_1 family of quasi-halo orbits. The paper neglects many details on their research, such as the implementation of the interpolation, the time needed to compute the grids, how the grid was constructed with continuation methods, and how they deal with the interpolation when the number of Fourier coefficients representing the invariant curves is different amongst the orbits in the grid since they allow this number to vary. Constructing these grids is time-consuming without the use of a computer with many processors and requires the storage of many orbits. Many computers which can be bought off the shelf these days will struggle with RAM requirements to have all the orbits in memory at once. Additionally, any orbit produced from the interpolation will need to be corrected to ensure the invariance is satisfied to a stricter error tolerance.

We consider a different approach to computing orbits with a specified set of parameters, avoiding the computation and interpolation of a dense grid of solutions. Leveraging the implicit function theorem, we develop a retraction (see Section 2.5.3) by augmenting a standard continuation method to travel a specified distance in a specified direction in the parameter space of an implicit manifold \mathcal{M} defined by $\mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}$. The augmentation comes in the form of a line search over step sizes to minimize the difference between the distance in the parameter space which is desired and the distance which has actually been traveled by the continuation method. The retraction takes advantage of the tangent space of \mathcal{M} to continue solutions in the appropriate direction. The

formulation of the retraction is general, showing its versatility to a wide array of problems, however it is validated and applied to the case of single- and multi-parameter families of quasi-periodic orbits in the circular restricted three-body problem.

This chapter is organized as follows: Section 5.2 goes over the development of the retraction. Section 5.3 presents the algorithm for the retraction. Section 5.4 validates the retraction with two examples. The first example computes quasi-halo orbits from a 2-parameter family which have specific sets of orbit frequencies. The second example computes an orbit which is a specific distance from the initial orbit along a 1-dimensional submanifold of constant energy. Section 5.5 ends the chapter with a discussion of the advantages of the retraction and presents various uses for it.

5.2 Solution Development

5.2.1 Manifolds

Let $\mathbf{x} \in \mathbb{R}^{\mathcal{D}}$ be state variables, $\mathbf{q} \in \mathbb{R}^p$ be parameters, and $\mathbf{z} = [\mathbf{x}^T, \mathbf{q}^T]^T$. Consider the system of equations

$$\mathbf{F} : \mathbb{R}^{\mathcal{D}} \times \mathbb{R}^p \rightarrow \mathbb{R}^{\mathcal{D}}, \quad \mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}. \quad (5.1)$$

Assume the Jacobian matrix $D_{\mathbf{x}}\mathbf{F}$ is full rank for all solutions (\mathbf{x}, \mathbf{q}) satisfying $\mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}$. Then the set of points

$$\mathcal{M} = \{(\mathbf{x}, \mathbf{q}) \mid \mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \quad (5.2)$$

is a p -dimensional smooth manifold embedded in $\mathbb{R}^{\mathcal{D}} \times \mathbb{R}^p$. We define the state manifold

$$\mathcal{X} = \{\mathbf{x} \mid \mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \quad (5.3)$$

and the parameter manifold

$$\mathcal{Q} = \{\mathbf{q} \mid \mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \quad (5.4)$$

so that $\mathcal{M} = \mathcal{X} \times \mathcal{Q}$. Moreover, we have $T_{\mathbf{z}}\mathcal{M} = T_{\mathbf{x}}\mathcal{X} \times T_{\mathbf{q}}\mathcal{Q}$. The implicit function theorem (Section 2.5.1) says there is a unique mapping from \mathcal{Q} to \mathcal{X} . We refer to the case in which manifolds are defined by Equation (5.1) as the *unconstrained* case. This is because the tangent space of \mathcal{Q} is

\mathbb{R}^p . Consider $\mathbf{q}^* \in \mathcal{Q}$, then there exists a radius r such that any $\mathbf{q} \in B_r(\mathbf{q}^*)$ implies $\mathbf{q} \in \mathcal{Q}$, where $B_r(\mathbf{q}^*)$ is a ball of radius r centered on \mathbf{q}^* .

Now consider Equation (5.1) with c additional equations, resulting in the system of equations

$$\bar{\mathbf{F}} : \mathbb{R}^D \times \mathbb{R}^p \rightarrow \mathbb{R}^{D+c}, \quad \bar{\mathbf{F}}(\mathbf{x}, \mathbf{q}) = \mathbf{0}, \quad (5.5)$$

where $c \in \mathbb{N}$ and $c \leq p - 1$. Again, assume that the Jacobian matrix of Equation (5.5) is full rank for all solutions (\mathbf{x}, \mathbf{q}) satisfying $\bar{\mathbf{F}}(\mathbf{x}, \mathbf{q}) = \mathbf{0}$. Then the set of points

$$\bar{\mathcal{M}} = \{(\mathbf{x}, \mathbf{q}) \mid \bar{\mathbf{F}}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \quad (5.6)$$

is a $(p - c)$ -dimensional implicit manifold embedded in $\mathbb{R}^D \times \mathbb{R}^p$. We define the state manifold

$$\bar{\mathcal{X}} = \{\mathbf{x} \mid \bar{\mathbf{F}}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \quad (5.7)$$

and the parameter manifold

$$\bar{\mathcal{Q}} = \{\mathbf{q} \mid \bar{\mathbf{F}}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \quad (5.8)$$

so that $\bar{\mathcal{M}} = \bar{\mathcal{X}} \times \bar{\mathcal{Q}}$. It is evident that $\bar{\mathcal{M}} \subset \mathcal{M}$ and the mapping defined from \mathcal{Q} to \mathcal{X} is a valid mapping from $\bar{\mathcal{Q}}$ to $\bar{\mathcal{X}}$ (Figure 5.1). We refer to the case in which manifolds are defined by Equation (5.5) as the *constrained* case. This is because the tangent space of $\bar{\mathcal{Q}}$ is **not** equal to \mathbb{R}^p . Consider $\mathbf{q}^* \in \bar{\mathcal{Q}}$, then there does not exist a ball such that all $\mathbf{q} \in B_r(\mathbf{q}^*)$ implies $\mathbf{q} \in \bar{\mathcal{Q}}$. Therefore, the locally valid directions are constrained to a subspace of \mathbb{R}^p .

5.2.2 Continuation Equations

For simplicity, we let

$$\mathbf{z} = \begin{bmatrix} \mathbf{x} \\ \mathbf{q} \end{bmatrix}.$$

The aim of the tool is to target a solution \mathbf{z}_t satisfying either Equations (5.1) or (5.5) by specifying the target parameters \mathbf{q}_t . To do this we must begin with an initial solution \mathbf{z}_0 from within the family.




Figure 5.1: The parameter and state manifolds defined by Equations (5.1) and (5.5).

In the unconstrained case this is simple. We only require that the line

$$\mathbf{q}(t) = \mathbf{q}_0 + t\delta\mathbf{q}_t, \quad t \in [0, 1] \quad (5.9)$$

be entirely contained in \mathcal{Q} where $\delta\mathbf{q}_t = \mathbf{q}_t - \mathbf{q}_0$. Then it is sufficient to use continuation to compute the branch of solutions lying on $\mathbf{q}(t)$. We append Equation (5.1) with the following $p-1$ parametric constraints, called direction constraints,

$$\begin{aligned} \mathbf{d}(z) &= \frac{\check{\mathbf{q}} - \check{\mathbf{q}}_0}{q - q_0} - \frac{\check{\mathbf{q}}_t - \check{\mathbf{q}}_0}{q_t - q_0} \\ &= \frac{\delta\check{\mathbf{q}}}{\delta q} - \frac{\delta\check{\mathbf{q}}_t}{\delta q_t} \end{aligned} \quad (5.10)$$

where q is the coordinate of \mathbf{q} such that

$$q = \arg \max_{i=1, \dots, p} |q_{0,i} - q_{t,i}|, \quad (5.11)$$

and $\check{\mathbf{q}}$ are the remaining coordinates of \mathbf{q} . Choosing q and $\check{\mathbf{q}}$ in this way ensures there are no issues with singularities. Moreover, we append with a pseudo-arclength equation $s(z)$ so that

$$\mathbf{H}(z) = \begin{bmatrix} \mathbf{F}(z) \\ s(z) \\ \mathbf{d}(z) \end{bmatrix} = \mathbf{0} \quad (5.12)$$

has a full rank Jacobian matrix. Equation (5.12) is used in a continuation method and ensures we compute solutions which lie along the straight line connecting \mathbf{q}_0 to \mathbf{q}_t . However, we need to

modulate the step size so that the continuation can converge to the target solution \mathbf{z}_t and stop computing solutions. We will return to this in Section (5.2.5), and introduce the continuation equations for the constrained case. The reader is pointed to Figure 5.2 to follow the proceeding discussion of the constrained case.




Figure 5.2: Depiction of the direction constraint in the constrained case. The constrained manifold $\bar{\mathcal{Q}}$ shown against the unconstrained manifold \mathcal{Q} in (a). The projection of $\delta \mathbf{q}_t$ onto the tangent space, resulting in $\delta \mathbf{q}_t^{\parallel}$, and the projection of the continuation path onto $\delta \mathbf{q}_t^{\parallel}$ in (b).

Just as in the constrained case we append a pseudo-arclength constraint to Equation (5.5), resulting in $\mathcal{D} + c + 1$ equations. Unlike in the unconstrained case, in the constrained case it is likely that $\mathbf{q}(t) \notin \bar{\mathcal{Q}}$ for all $t \in [0, 1]$, so we must project the vector $\delta \mathbf{q}_t$ into the tangent space $T_{\mathbf{q}_0} \bar{\mathcal{Q}}$. We then want to ensure that the path that the continuation follows on the manifold is aligned with the projected vector $\delta \mathbf{q}_t^{\parallel}$. That is to say the vector pointing from the initial parameters to the current parameters $\delta \mathbf{q} = \mathbf{q} - \mathbf{q}_0$ projected into $T_{\mathbf{q}_0} \bar{\mathcal{Q}}$ is parallel to $\delta \mathbf{q}_t^{\parallel}$. The projection is accomplished similarly to the projection in Equation (2.67). Let V_0 be a $(\mathcal{D} + p) \times (p - c)$ matrix whose columns

are linearly independent vectors spanning $T_{z_0}\bar{\mathcal{M}}$. Furthermore, we let

$$\bar{V}_0 = \begin{bmatrix} \bar{V}_{x_0} \\ \bar{V}_{q_0} \end{bmatrix}, \quad (5.13)$$

where \bar{V}_{q_0} is a $p \times (p - c)$ matrix with linearly independent column vectors¹ and \bar{V}_{x_0} is the corresponding $\mathcal{D} \times (p - c)$ matrix. The projection operator is given as

$$\text{Proj}_{q_0}(\mathbf{v}) = \bar{V}_{q_0} \left[\bar{V}_{q_0}^\dagger \mathbf{v} \right]. \quad (5.14)$$

Now the projected vector we wish to follow along the manifold is

$$\delta \mathbf{q}_t^\parallel = \text{Proj}_q(\delta \mathbf{q}_t). \quad (5.15)$$

The projected vector from the initial parameters to the current parameters in the continuation is given by

$$\delta \mathbf{q}^\parallel = \text{Proj}_{q_0}(\delta \mathbf{q}), \quad (5.16)$$

leading to the following direction constraints

$$\bar{\mathbf{d}}(z) = \frac{\delta \check{\mathbf{q}}^\parallel}{\delta q^\parallel} - \frac{\delta \check{\mathbf{q}}_t^\parallel}{\delta q_t^\parallel}. \quad (5.17)$$

It is important to note the projection operator uses \bar{V}_{q_0} from the initial solution z_0 for the entirety of the computation of the branch.

We do not append all the direction constraints of Equation (5.17). Instead, we only append $p - 1 - c$ constraints. The choice of which equations to discard does not matter. Any $p - 1 - c$ of the scalar equations suffices to construct a full rank Jacobian matrix. Let $\bar{\mathbf{d}}$ be the $p - 1 - c$ direction equations, then the system of equations

$$\bar{\mathbf{H}}(z) = \begin{bmatrix} \bar{\mathbf{F}}(z) \\ s(z) \\ \bar{\mathbf{d}}(z) \end{bmatrix} = \mathbf{0} \quad (5.18)$$

¹ \bar{V}_0 having linearly independent column vectors does not imply any subset of the rows also have linear independence. Thus, we require \bar{V}_{q_0} to also have linearly independent columns.

is used in a continuation method for the constrained case. The original target solution is likely not able to be achieved. Instead, we must settle for a *nearby* solution which satisfies Equation (5.5).

To this end, let

$$d = \|\delta \mathbf{q}_t^\parallel\|. \quad (5.19)$$

Then traveling in the direction of $\delta \mathbf{q}_t^\parallel$ on $\bar{\mathcal{Q}}$ a distance d should provide a nearby solution as long as d isn't too large. In the unconstrained case the projection operator is the identity operator, $d = \|\delta \mathbf{q}_t\|$, $\delta \mathbf{q}_t^\parallel = \delta \mathbf{q}_t$.

5.2.3 Initial Direction

From Equation (5.15) we know the initial direction on $\bar{\mathcal{Q}}$ needed to continue solutions which satisfy the direction constraints of Equation (5.18). However, to initialize the continuation method we need to know the direction for both the states and the parameters. We generate the initial direction \mathbf{z}'_0 to move on $\bar{\mathcal{M}}$ as

$$\begin{aligned} \mathbf{z}'_0 &= \bar{V}_0 \begin{bmatrix} V_{\mathbf{q}_0}^\dagger \delta \mathbf{q}_t \end{bmatrix} \\ &= \begin{bmatrix} \bar{V}_{\mathbf{x}_0} \bar{V}_{\mathbf{q}_0}^\dagger \delta \mathbf{q}_t \\ \bar{V}_{\mathbf{q}_0} \bar{V}_{\mathbf{q}_0}^\dagger \delta \mathbf{q}_t \end{bmatrix} \\ &= \begin{bmatrix} \delta \mathbf{x}_t^\parallel \\ \delta \mathbf{q}_t^\parallel \end{bmatrix}. \end{aligned} \quad (5.20)$$

In the unconstrained case $V_{\mathbf{q}_0}^\dagger = V_{\mathbf{q}_0}^{-1}$, reducing Equation (5.20) to

$$\begin{aligned} \mathbf{z}'_0 &= V_0 \begin{bmatrix} V_{\mathbf{q}_0}^{-1} \delta \mathbf{q}_t \end{bmatrix} \\ &= \begin{bmatrix} V_{\mathbf{x}_0} V_{\mathbf{q}_0}^{-1} \delta \mathbf{q}_t \\ V_{\mathbf{q}_0} V_{\mathbf{q}_0}^{-1} \delta \mathbf{q}_t \end{bmatrix} \\ &= \begin{bmatrix} \delta \mathbf{x}_t \\ \delta \mathbf{q}_t \end{bmatrix}. \end{aligned} \quad (5.21)$$

5.2.4 Stopping Criterion

A standard continuation method computes branches until either the maximum number of solutions has been computed or when a new solution fails to be found. However, to target a specific solution we must introduce a stopping criterion which stops on the correct solution. In the unconstrained case we know the parameters \mathbf{q}_t and can stop when those parameters are met, however in the constrained case the target parameters are likely not in the parameter manifold $\bar{\mathcal{Q}}$. Thus, we introduce a stopping criterion based on distance traveled on the parameter manifold.

The target distance to travel is given in Equation (5.19). Let \mathbf{z}_0 be the initial solution with parameters \mathbf{q}_0 , Δs_0 be the initial step size, and let \mathbf{z}_1 be the first solution from the continuation. Then we define the distance traveled on $\bar{\mathcal{Q}}$ up to this point to be $d_1 = \text{sign}(\Delta s_0) \|\mathbf{q}_1 - \mathbf{q}_0\|$. Extending this to the k^{th} solution leads to

$$d_k = \sum_{j=1}^k \text{sign}(\Delta s_{j-1}) \|\mathbf{q}_j - \mathbf{q}_{j-1}\|. \quad (5.22)$$

In the limit of infinitesimal changes in the parameters we recover a path integral over the path on $\bar{\mathcal{Q}}$ from \mathbf{q}_0 to \mathbf{q}_t . The distance equation ensures that if the target parameters are stepped over in the continuation and the continuation needs to backtrack then the total distance traveled does not increase. A positive step size indicates forward motion on the branch while a negative step size indicates backward motion. Equation (5.22) is also used in the unconstrained case, however the distance computation can be simplified to $d_k = \|\mathbf{q}_k - \mathbf{q}_0\|$. Computing d_k allows us to keep track of the progression of the continuation. Leveraging this knowledge we can intelligently choose step sizes so that $d_k \rightarrow d$.

5.2.5 Step Size Computation

A standard continuation method utilizes some nominal step size update scheme such as the one presented in Equation (2.70). This works well for computing branches of solutions to study families as a whole, but does not perform well when a specific solution from a single- or multi-parameter family is desired. A nominal step size update is adequate until the solutions approach

the targeted solution, at which point the step size needs to be fine-tuned to ensure the targeted solution is not passed over.

The method to choose the step size to converge on the target distance is a line search over step sizes, given in Algorithm 9. Let \mathbf{q}_k and \mathbf{q}'_k be the parameters and their family tangent values, respectively, at the current continuation solution. Additionally, let Δs_k be the current nominal step size, and d_k be the current distance traveled on the parameter manifold. A step size s is chosen from a uniform partition on the interval $[-1.2|\Delta s_k|, 1.2|\Delta s_k|]$ with N points. We choose $N = 1000$, allowing the step size to potentially decrease by 3 orders of magnitude each continuation step. The step size is multiplied by $\|\mathbf{q}'_k\|$ to generate an estimate Δq of how much the parameters will change by on the next continuation solution. Δq is also the predicted distance traveled on the parameter manifold from solution k to solution $k+1$. The distance is added to the current distance traveled d_k to generate a guess of the total distance traveled upon convergence to solution \mathbf{z}_{k+1} . The absolute value of the difference between the target distance d and the predicted total distance traveled gives an error, or the distance left to go before converging to the target distance. We choose the step size s which minimizes this error.

Algorithm 9 Line search to determine the step size to get closest to the target distance

```

1: procedure TARGETSTEPsize( $\Delta s_k, \mathbf{q}'_k, d_k, d, N$ )
2:    $\Delta d_{\min} = |d - d_k|$ 
3:   for  $j \in \{0, 1, \dots, N\}$  do
4:      $s = -1.2|\Delta s_k| + 2.4|\Delta s_k| \frac{j}{N}$ 
5:      $\Delta q = s \|\mathbf{q}'_k\|$ 
6:      $d_j = d_k + \Delta q$ 
7:      $\Delta d_j = |d - d_j|$ 
8:     if  $\Delta d_j < \Delta d_{\min}$  then
9:        $\Delta s_{\min} = s$ 
10:       $\Delta d_{\min} = \Delta d_j$ 
11:     end if
12:   end for
13:   return  $\Delta s_{\min}$ 
14: end procedure

```

We note there are more sophisticated line search methods to pick the optimal step size, however we are not interested in getting the exact optimal step size because the true parameters

will differ from the predicted parameters. Moreover, to truly know the optimal step size it is necessary to solve the BVP for each step size. This is extremely prohibitive, so we rely solely on the predicted error rather than the true error to pick the optimal step size. We comment on the range of step sizes; by allowing the lower bound of the step sizes to be negative we allow the possibility of backtracking in the event the continuation has traveled too far, ultimately allowing for convergence to the target distance. Moreover, we let the upper and lower bounds be larger in magnitude than the nominal step size. From experience the nominal step size update given in Equation (2.70) is conservative when in a well-behaved region of a family of solutions. In practice, Algorithm 9 returns step sizes at the upper bound of the range, until getting close enough to the target distance. Even with the increased step size the Newton's method is able to converge without any more difficulties.

5.3 The Retraction

Thus far we have seen that by projecting the vector from the parameters of the initial solution to the target parameters to the parameter tangent space and using the distance constraints we can control the direction of the continuation. Furthermore, by modulating the step size the continuation can achieve a target solution. Recall from Section 2.5.3 a retraction is a mapping from a manifold's tangent space to down to a specific point on the manifold. By modifying a standard continuation method with the capabilities presented in this chapter we have essentially constructed a continuation method which behaves as a retraction. The ramifications of a retraction may not be fully understood at this point, but a retraction is necessary for the applications in Chapter 6 when it is used to optimize over families of solutions. For now, we present the general algorithm for the retraction here and validate its functionality on simple examples in Section 5.4.

Algorithm 10 takes as inputs the direction to travel in the parameter space, $\delta \mathbf{q}_t$, the initial solution, \mathbf{z}_0 , the initial tangent space, V_0 , the initial step size, Δs_0 , the error tolerance on the distance, ε , and the number of points for the routine TARGETSTEPsize, N . The distance, d , to travel is computed as the norm of $\delta \mathbf{q}_t$. Then the family tangent vector is computed according to

Algorithm 10 Retract from parameter tangent space down to the parameter manifold and find the corresponding solution on the state manifold

```

1: procedure RETRACT( $\delta \mathbf{q}_t, \mathbf{z}_0, V_0, \Delta s_0, \varepsilon, N$ )
2:    $d = \|\delta \mathbf{q}_t\|$ 
3:    $\mathbf{z}'_0 = V_0 \left[ V_{\mathbf{q}_0}^\dagger \delta \mathbf{q}_t \right]$  ▷ Automatically get  $\delta \mathbf{q}_t^\parallel$  from Equation (5.20).
4:    $\mathbf{z}'_0 = \frac{\mathbf{z}'_0}{\|\mathbf{z}'_0\|}$ 
5:    $k = 0$ 
6:    $(\mathbf{q}_k, \mathbf{q}'_k, \Delta s_k, d_k) = (\mathbf{q}_0, \mathbf{q}'_0, \Delta s_0, 0)$ 
7:   while  $|d - d_k| > \varepsilon$  do
8:      $\Delta s_k = \text{TARGETSTEP SIZE}(\Delta s_k, \mathbf{q}'_k, d_k, d, N)$ 
9:      $(\mathbf{z}_{k+1}, V_{k+1}, \Delta s_{k+1}) = \text{SOLVER}(\mathbf{z}_k, \mathbf{z}'_k, \Delta s_k)$ 
10:     $d_{k+1} = d_k + \|\mathbf{q}_{k+1} - \mathbf{q}_k\|$ 
11:     $\mathbf{z}'_{k+1} = \mathbf{z}_{k+1} - \mathbf{z}_k$ 
12:     $k \leftarrow k + 1$ 
13:   end while
14:    $(\mathbf{z}^*, V^*) = (\mathbf{z}_k, V_k)$ 
15:   return  $(\mathbf{z}^*, V^*)$ 
16: end procedure

```

Equation (5.20) and normalized with the appropriate norm $\| \cdot \|$ consistent with the inner product used in the pseudo-arclength equation. The algorithm proceeds to go into the continuation loop. In each iteration the step size is modulated with Algorithm 9, the next solution is computed with the routine SOLVER. The routine SOLVER solves the continuation equations and returns the next solution, its tangent space, and the nominal step size. The cumulative distance in the parameter space is computed. The continuation loop continues until either the distance traveled is within the error tolerance or an updated solution is not found, and returns the final solution.

5.4 Validation

The retraction is tested on both the unconstrained and the constrained case for the application of computing quasi-periodic orbits from within the 2-parameter family of the Earth-Moon L_2 quasi-halo orbit family in the circular restricted three-body problem. For this specific case of computing quasi-periodic orbits we include the resonance avoidance methods from Chapter 4 to gain better performance. The routine IRRATIONALSTEPsize (Algorithm 5) is implemented directly after the routine TARGETSTEPsize, fine-tuning the step size to compute the orbit with the most irrational set of frequencies.

In the unconstrained case we are concerned with computing an orbit with a specific set of orbit frequencies because all dimensions of the frequency space are available. In the constrained case a Jacobi constant constraint is included, constraining the possible set of frequencies. The only way to know the achievable set of frequencies is to compute the entire constant energy branch. Without computing the entire branch we settle for targeting the distance along the branch to produce an orbit close to a desired set of orbit frequencies.

5.4.1 Targeting Orbit Frequencies in a Family of Quasi-Periodic Orbits

5.4.1.1 Problem Formulation

Recall from Section 3.2 the solution vector \mathbf{z} is composed of discretized states over the invariant surface, \mathbf{X} , the stroboscopic time, T , the rotation numbers, $\boldsymbol{\rho}$, and the orbit frequencies, $\boldsymbol{\omega}$. Furthermore, the orbit frequencies parameterize the solutions, so the state variables

$$\mathbf{x}(\boldsymbol{\omega}) = \begin{bmatrix} \mathbf{X}(\boldsymbol{\omega}) \\ T(\boldsymbol{\omega}) \\ \boldsymbol{\rho}(\boldsymbol{\omega}) \end{bmatrix} \quad (5.23)$$

are functions of the orbit frequencies. For 2-dimensional quasi-periodic orbits we have $\boldsymbol{\rho} = \rho_1$ and $\boldsymbol{\omega} = (\omega_0, \omega_1)$. For a 2-parameter family of quasi-periodic orbits in the circular restricted three-body problem the generating system of equations for the family in a single-shooting setting² is

$$\mathbf{F}(\mathbf{z}) = \mathbf{F}(\mathbf{X}, T, \rho_1, \boldsymbol{\omega}) = \begin{bmatrix} R_{-\rho_1} \varphi_T(\mathbf{X}) - \mathbf{X} \\ \left\langle \mathbf{X} - \tilde{\mathbf{X}}, \frac{\partial \tilde{\mathbf{X}}}{\partial \theta_0} \right\rangle \\ \left\langle \mathbf{X} - \tilde{\mathbf{X}}, \frac{\partial \tilde{\mathbf{X}}}{\partial \theta_1} \right\rangle \\ T\omega_0 - 2\pi \\ T\omega_1 - \rho_1 \end{bmatrix}, \quad (5.24)$$

generating the following 2-dimensional Cantor manifolds

$$\mathcal{M} = \{(\mathbf{x}, \boldsymbol{\omega}) \mid \mathbf{F}(\mathbf{x}, \boldsymbol{\omega}) = \mathbf{0}\} \quad (5.25)$$

$$\mathcal{X} = \{\mathbf{x} \mid \mathbf{F}(\mathbf{x}, \boldsymbol{\omega}) = \mathbf{0}\} \quad (5.26)$$

$$\Omega = \{\boldsymbol{\omega} \mid \mathbf{F}(\mathbf{x}, \boldsymbol{\omega}) = \mathbf{0}\}. \quad (5.27)$$

The system of equations used in the continuation to compute branches of solutions needs an additional parametric constraint to constraint the continuation to a 1-dimensional curve. To reduce the dimension from 2 to 1, a direction constraint from Equation (5.10) is appended to \mathbf{F} , along

² See Section 3.2 for the notation and the extrapolation of the equations to multiple-shooting.

with a pseudo-arclength constraint to trace out the branch of solutions. This leads to the following system of equations for the continuation method

$$\mathbf{H}(z) = \begin{bmatrix} \mathbf{F}(z) \\ \frac{1}{\mathcal{N}} \langle \mathbf{X} - \tilde{\mathbf{X}}, \tilde{\mathbf{X}}' \rangle + (T - \tilde{T})\tilde{T}' + (\rho_1 - \tilde{\rho}_1)\tilde{\rho}'_1 + \langle \boldsymbol{\omega} - \tilde{\boldsymbol{\omega}}, \tilde{\boldsymbol{\omega}}' \rangle - \Delta_s \\ \frac{\delta \tilde{\boldsymbol{\omega}}}{\delta \boldsymbol{\omega}} - \frac{\delta \tilde{\boldsymbol{\omega}}_t}{\delta \boldsymbol{\omega}_t} \end{bmatrix}. \quad (5.28)$$

We choose an initial quasi-halo orbit with frequencies

$$\boldsymbol{\omega} = (1.946982196701564, 0.590131700668313),$$

target frequencies

$$\boldsymbol{\omega}_t = (1.978018524093039, 0.622782717360435),$$

and a distance tolerance of $\varepsilon = 1e-8$. Therefore, we have

$$\delta \boldsymbol{\omega}_t = (0.031036327391475, 0.032651016692122)$$

and $d = 0.045048224260009$. Based on $\delta \boldsymbol{\omega}_t$ we have $\delta \tilde{\boldsymbol{\omega}} = \boldsymbol{\omega}_0$ and $\delta \boldsymbol{\omega} = \boldsymbol{\omega}_1$. Additionally, we use $\mathcal{N} = N_1 = 41$ to discretize the invariant curves.

5.4.1.2 Results

Algorithm 10 produces a quasi-halo with frequencies

$$\boldsymbol{\omega}^* = (1.978018520706167, 0.622782713797358).$$

The results are given in Figure 5.3. Plots (a) and (b) show the continuation in frequency space, while plot (c) shows the distance error versus the continuation solutions. The initial step size is $1e-3$ while the maximum nominal step size is $5e-3$, leading to small steps at the beginning of the continuation. The step size drastically increases between solution 14 and 15 due to the resonance avoidance method from Section 4.4 reducing the distance error by orders of magnitude. In plot (b) it is seen that the solution jumps past the target solution and begins to backtrack. From there it takes only a few more solutions to hone in on a quasi-halo orbit with frequencies satisfying the

tolerance. The entire algorithm computes 19 orbits and takes 105 seconds to run. The computations were performed in the Julia language and the integrations were not run in parallel. The runtime can be significantly reduced if the integrations are run in parallel.

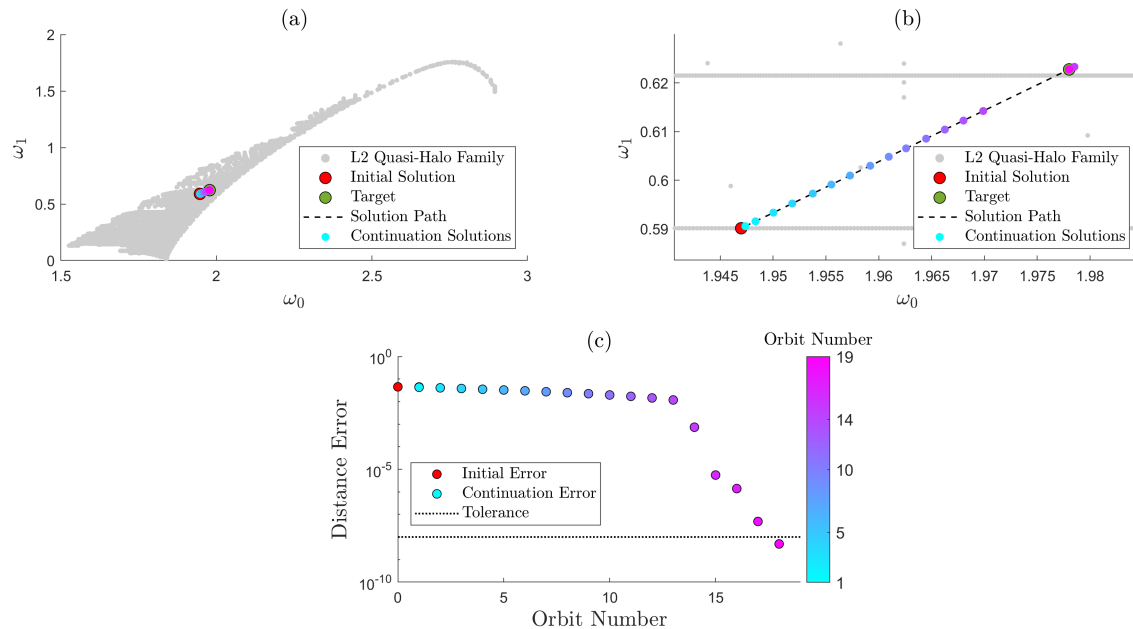


Figure 5.3: Targeting the frequencies of a quasi-periodic orbit.

Plot (c) shows a monotonically decreasing error as the orbits are successively computed. This behavior is expected for larger step sizes, but as the step size gets small enough it is not unexpected for the distance error to oscillate. The oscillation can occur because the predicted frequencies do not exactly determine the true frequencies. The Newton's method can converge to an orbit further away from the target than expected based on the predicted frequencies.

5.4.2 Targeting Distance on a Subset of a Family of Quasi-Periodic Orbits

5.4.2.1 Problem Formulation

For this test case a Jacobi constant constraint is included in the generating system of equations \bar{F} which constructs 1-dimensional Cantor manifolds. No direction constraints are necessary since

the continuation method can trace out a 1-parameter branch. The generating system of equations, the manifolds defined by the generating system of equations, and the continuation equations are presented below.

$$\bar{\mathbf{F}}(\mathbf{z}) = \bar{\mathbf{F}}(\mathbf{X}, T, \rho_1, \boldsymbol{\omega}) = \begin{bmatrix} R_{-\rho_1} \varphi_T(\mathbf{X}) - \mathbf{X} \\ \langle \mathbf{X} - \tilde{\mathbf{X}}, \frac{\partial \tilde{\mathbf{X}}}{\partial \theta_0} \rangle \\ \langle \mathbf{X} - \tilde{\mathbf{X}}, \frac{\partial \tilde{\mathbf{X}}}{\partial \theta_1} \rangle \\ T\omega_0 - 2\pi \\ T\omega_1 - \rho_1 \\ \bar{J}(\mathbf{X}) - J^* \end{bmatrix} \quad (5.29)$$

$$\bar{\mathcal{M}} = \{(\mathbf{x}, \boldsymbol{\omega}) \mid \bar{\mathbf{F}}(\mathbf{x}, \boldsymbol{\omega}) = \mathbf{0}\} \quad (5.30)$$

$$\bar{\mathcal{X}} = \{\mathbf{x} \mid \bar{\mathbf{F}}(\mathbf{x}, \boldsymbol{\omega}) = \mathbf{0}\} \quad (5.31)$$

$$\bar{\Omega} = \{\boldsymbol{\omega} \mid \bar{\mathbf{F}}(\mathbf{x}, \boldsymbol{\omega}) = \mathbf{0}\} \quad (5.32)$$

$$\bar{\mathbf{H}}(\mathbf{z}) = \begin{bmatrix} \bar{\mathbf{F}}(\mathbf{z}) \\ \frac{1}{N} \langle \mathbf{X} - \tilde{\mathbf{X}}, \tilde{\mathbf{X}}' \rangle + (T - \tilde{T})\tilde{T}' + (\rho_1 - \tilde{\rho}_1)\tilde{\rho}'_1 + \langle \boldsymbol{\omega} - \tilde{\boldsymbol{\omega}}, \tilde{\boldsymbol{\omega}}' \rangle - \Delta s \end{bmatrix} \quad (5.33)$$

The same initial quasi-halo orbit, target frequencies, distance error tolerance, and number of points discretizing the invariant curve are used as in Section 5.4.1. The only difference is the use of the Jacobi constant constraint.

5.4.2.2 Results

Algorithm 10 produces a quasi-halo with frequencies

$$\boldsymbol{\omega}^* = (1.978018524093039, 0.622782717360435).$$

The results are given in Figure 5.3. Plots (a) and (b) show the continuation in frequency space, while plot (c) shows the distance error versus the continuation solutions, and plot (d) shows the

error in the prescribed value of the Jacobi constant for each computed solution. The initial step size is $1e-3$ while the maximum nominal step size is $5e-3$. Keeping the maximum step size small in the constrained case cultivates accurate distance computations. For continuation curves with more curvature the maximum step size should be kept small while for continuation curves that resemble straight lines the maximum step size can be larger. The number $5e-3$ is conservative since constant Jacobi energy lines in frequency space are nearly straight³. No large steps are taken because a resonance is not detected along this path, indicating the detected resonance in the previous test case is likely a high-order resonance with little effect on the continuation. Therefore, the distance between consecutive solutions remains fairly constant near the maximum step size. Though it cannot be seen in plot (b) the solution jumps past the target solution and begins to backtrack. The entire algorithm computes 23 orbits and takes 134 seconds to run.

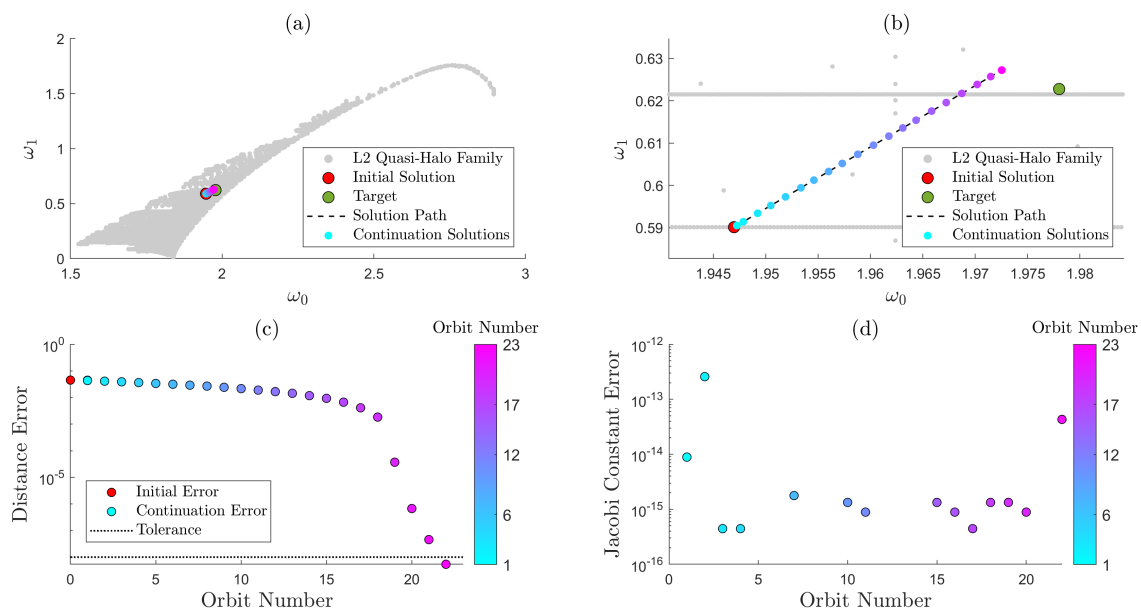


Figure 5.4: Targeting the distance along a constant Jacobi energy line.

Plot (c) shows the distance error decreasing monotonically again while plot (d) shows the computed orbits share the same Jacobi constant. Plot (b) shows the orbit a distance d along

³ See Figure 3.12

this constant Jacobi energy branch. If we compare the distance between the initial and the final parameters on Ω we get a distance of 0.04504817, but the distance on $\bar{\Omega}$ is 0.04504822. The distance on Ω does not quite meet the distance error tolerance of 1e-8, while the distance on $\bar{\Omega}$ does.

5.5 Conclusion

Mondelo et al. compute tens of thousands of quasi-periodic orbits to construct grids in an energy-rotation number space [136]. They interpolate the family to essentially construct a continuous representation of the family. Their paper does not specify how quickly the interpolator produces an invariant torus, however they will need to correct the interpolated solution to become a truly invariant torus. With the retraction presented in this chapter the desired quasi-periodic orbit automatically satisfies the invariance condition to a tight tolerance, computed in a matter of minutes, and does not require the computation of thousands of orbits - only tens of orbits. With the retraction no knowledge of the family is needed. If the specified orbit frequencies or the continuation path lies outside Ω or $\bar{\Omega}$ then the continuation will simply terminate upon reaching the boundary of the family or at a low-order resonance. Even if the returned orbit does not satisfy the frequency or distance requirements it will still satisfy the invariance condition and give insight into the structure of the family.

The retraction function can be used to construct grids of solutions over the parameter space. In this chapter we looked at its use for computing quasi-periodic orbits parameterized by their orbit frequencies. The retraction can be used to compute a grid of quasi-periodic orbits in frequency space or in the stroboscopic time-rotation number space. Moreover, the retraction is applicable to families of periodic orbits. However, the true power of the retraction is explored in Chapter 6. With the tools developed later on, a grid of quasi-periodic orbits can be constructed in a variety of spaces and a quasi-periodic orbit can be targeted based on general characteristics rather than just the parameters. Furthermore, the retraction allows for optimization over families of quasi-periodic orbits.

The work of Chapter 4 was motivated and developed concurrently with this chapter. To

target any set of orbit frequencies the continuation needs the ability to move in arbitrary directions in frequency space. Naturally, these directions will cross over resonances. Therefore, methods to avoid resonances without restricting the direction in which branches are computed were sought. The method of Gómez and Mondelo directly restricts the continuation to compute constant rotation number branches [75]. The method of McCarthy and Howell impose a generalization of the method of Gómez and Mondelo to higher dimensional invariant tori by fixing all but one rotation number to ensure the ratios between the frequencies stay constant [130]. The purpose of both of these methods is to stay away from resonances altogether, restricting the direction of continuation, whereas the methods developed in Chapter 4 do not restrict the continuation direction and seek to improve the performance of the continuation of invariant tori in the face of resonances.

If a target set of frequencies cannot be reached from an initial quasi-periodic orbit due to a resonance, and orbits have been computed across the family, then the target quasi-periodic orbit may be reached by changing the initial quasi-periodic orbit. An idea we propose is if one knows the target frequencies and has a grid of computed quasi-periodic orbits, then all the possible paths to the target frequencies from the pre-computed frequencies can be analyzed in frequency space ahead of time to choose the path, and hence the initial quasi-periodic orbit, that will have the least interaction with resonances on its way to the target frequencies. This initial quasi-periodic orbit should be the best option to reach the quasi-periodic orbit with the target set of frequencies.

Chapter 6

Optimization Over Families of Quasi-Periodic Orbits

6.1 Introduction

A family of quasi-periodic orbits satisfying the Diophantine conditions has been shown to be a nearly continuous family called a Cantor family (Cantor manifold) in [110, 39, 114]. A Cantor manifold has gaps in it where the torus frequencies do not satisfy the Diophantine conditions. It has been shown that a Cantor manifold can be treated as a smooth manifold where smoothness is defined in the sense of Whitney [143, 168, 56, 114, 39]. The smoothness of families of quasi-periodic orbits is a powerful notion as it allows one to draw ties to smooth manifolds defined by the implicit function theorem, and utilize all the rich theory and tools developed for smooth manifolds [116, 33]. One such tool is the gradient; an indispensable tool enabling gradient-based optimization over families of orbits, adding to the list of available tools for trajectory designers. This chapter focuses on the optimization of cost functions over a family of quasi-periodic orbits.

Optimization can be done on smooth and non-smooth manifolds using non-gradient-based methods, such as genetic algorithms and particle swarm optimizers [60], however in general, it is beneficial to utilize the topology and smoothness of a space when available. The exceptions to this are when the cost of evaluating the gradient is prohibitive, the gradient is not well-behaved over the space, or when there are many local minima. As seen in Chapter 3 the L_2 quasi-halo family is well-behaved, meaning it is well-suited for gradient-based optimization.

The standard optimization formulation in Problem (2.77) treats the orbit states \mathbf{X} and the orbit frequencies $\boldsymbol{\omega}$ as independent variables which need to satisfy equality constraints defining

a family of quasi-periodic orbits, equivalently the parametric constraints in Equations (3.34) or (3.35). Instead, we reformulate the optimization problem as

$$\begin{aligned} \min_{\boldsymbol{\omega} \in \Omega} f(\boldsymbol{x}, \boldsymbol{\omega}) \\ \text{subject to } h_i(\boldsymbol{x}, \boldsymbol{\omega}) \leq b_i, \quad i = 1, \dots, s, \end{aligned} \tag{6.1}$$

where \boldsymbol{x} is defined in Equation (5.23), to leverage the fact that each orbit is uniquely parameterized by the orbit frequencies $\boldsymbol{\omega}$, such that the optimization occurs over Ω , the set of orbit frequencies in a family of quasi-periodic orbits. Optimizing over the orbit frequencies greatly reduces the optimization variables. The negative gradient of a cost function, $-\nabla f$, dictates the direction in Ω to move locally to minimize f . We use the retraction developed in Chapter 5 to explore the descent direction and compute the corresponding quasi-periodic orbits. Using the retraction derives several benefits. First, and foremost, the retraction ensures the invariance equations and any additional parametric constraints are always satisfied. Moreover, the retraction is specialized to handle the subtleties of computing quasi-periodic orbits, such as avoiding resonances, efficiently computing quasi-periodic orbits. Lastly, the retraction uses continuation which is leveraged to perform line searches in a gradient descent method.

Optimizing over families of quasi-periodic orbits, and even periodic orbit, closely relates to a class of problems called PDE-constrained optimization (PDECO). PDECO problems occur when there are constraints governed by physical laws described by partial differential equations (PDEs). PDECO is a well-developed and active field which supports research in a variety of disciplines. One popular use is for topology (shape) optimization which aims to find the optimal shape of a surface given constraints within a dynamical environment. The most well-known use of topology optimization in aerodynamics might be to design aircraft. Other applications of PDECO include crystal growth [137], cooling of electronic components [144], drug transport [50, 160, 10], image denoising [58], and uncertainty [70]. The reader is directed to the recent book by Antil and Leykekhman in [9] for a deeper dive into the subject of PDECO and its applications. The following papers contain additional applications [3, 8, 99]. While no PDEs are used here, the form of the

resulting problem is similar. The equations defining a family of quasi-periodic orbits serve the same purpose as PDEs with the primary constraint being the quasi-periodicity constraint involving the solution flow of an ODE. The process to compute families of solutions in ODEs involves the use of continuation. The idea of using continuation to solve optimization problems seems to have originated in the work of Kernévez and Doedel [112].

Kernévez and Doedel devise an optimization routine to find the solution pair $(\boldsymbol{\lambda}_0, \mathbf{u}_0)$ of control variables and state variables which maximizes a cost functional $f(\boldsymbol{\lambda}, \mathbf{u})$ such that the state equation $\mathbf{g}(\boldsymbol{\lambda}, \mathbf{u}) = \mathbf{0}$ is satisfied and subject to inequality constraints $h_i(\boldsymbol{\lambda}, \mathbf{u}) \leq \mathbf{0}$. They're approach freezes all $\boldsymbol{\lambda}$ except one element, employs continuation to compute a branch of solutions to $\mathbf{g}(\boldsymbol{\lambda}, \mathbf{u}) = \mathbf{0}$, and then pick the solution from the branch which maximizes $f(\boldsymbol{\lambda}, \mathbf{u})$. They repeat this process in succession and change which control variable is left unfrozen. Their approach is applied to a family of dynamical systems which depend on $\boldsymbol{\lambda}$ and experience bifurcations in these variables. The successive use of continuation is called successive continuation and has been used in a variety of disciplines [170, 164, 53, 2, 117], however they find optimal control variables which change the dynamical system. In essence, they are solving for a dynamical system which provides favorable behaviors to optimize their processes. The application here considers optimizing over families of solutions within a single dynamical system using successive continuation.

This chapter is organized as follows: Section 6.2 defines the optimization problem. Section 6.4 details a modified Newton's method devised for computing solutions with specified characteristics. Several examples in the mindset of mission design are presented. Section 6.5 details a modified gradient descent algorithm which recasts an equality constrained problem into an unconstrained optimization problem. Examples of optimization problems are presented. Section 6.6 uses the modified gradient descent algorithm in an augmented Lagrangian method (ALM) to perform constrained optimization with inequality constraints. And again, examples of optimization problems are presented. Lastly, Section 6.7 ends the chapter with concluding remarks.

6.2 Problem Statement

We proceed with a general statement of the problem to be solved. Let $\mathbf{x} \in \mathbb{R}^D$ be state variables, $\mathbf{q} \in \mathbb{R}^p$ be parameters, and $\mathbf{z} = [\mathbf{x}^T, \mathbf{q}^T]^T$. Let $\mathbf{F} : \mathbb{R}^D \times \mathbb{R}^p \rightarrow \mathbb{R}^D$ be a system of equations defining a p -parameter family of solutions, and $\mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{F}(\mathbf{z})$. Recall from Section 5.2.1 this is the unconstrained case. We drop the notion of the constrained and unconstrained case to avoid confusion with constrained and unconstrained optimization problems, however the difference between the two cases are important and will be recognized in Section 6.3.

Assume that $D_{\mathbf{x}}\mathbf{F}$ is full rank for all \mathbf{z} satisfying $\mathbf{F}(\mathbf{z}) = \mathbf{0}$, then by the implicit function theorem $\mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}$ defines the following p -dimensional smooth manifolds

$$\mathcal{M} = \{(\mathbf{x}, \mathbf{q}) \mid \mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \quad (6.2)$$

$$\mathcal{X} = \{\mathbf{x} \mid \mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \quad (6.3)$$

$$\mathcal{Q} = \{\mathbf{q} \mid \mathbf{F}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\}, \quad (6.4)$$

with $\mathcal{M} = \mathcal{X} \times \mathcal{Q}$. Moreover, there is a mapping from \mathcal{X} to \mathcal{Q} such that $\mathbf{x} = \mathbf{x}(\mathbf{q})$.

We are interested in finding a solution $\mathbf{z}^* \in \mathcal{M}$ such that \mathbf{z}^* minimizes a cost functional $f : \mathbb{R}^D \times \mathbb{R}^p \rightarrow \mathbb{R}$ subject to equality and inequality constraints. The optimization problem can be stated as

$$\begin{aligned} & \min_{\mathbf{z} \in \mathbb{R}^{D+p}} f(\mathbf{z}) \\ & \text{subject to } \mathbf{F}(\mathbf{z}) = \mathbf{0} \\ & \quad g_i(\mathbf{z}) = a_i, \quad i = 1, \dots, r \\ & \quad h_j(\mathbf{z}) \leq b_j, \quad j = 1, \dots, s. \end{aligned} \quad (6.5)$$

In Problem (6.5), \mathbb{R}^{D+p} is the space over which the optimization occurs, treating \mathbf{z} as the optimization variables and \mathbf{x} and \mathbf{q} as independent quantities. From the implicit function theorem we know \mathbf{x} and \mathbf{q} are not independent. Each \mathbf{x} is uniquely parameterized by \mathbf{q} and the parameters

\mathbf{q} lie on the manifold \mathcal{Q} . We recast Problem (6.5) as the following optimization problem

$$\begin{aligned} \min_{\mathbf{q} \in \mathcal{Q}} \quad & f(\mathbf{x}, \mathbf{q}) \\ \text{subject to} \quad & g_i(\mathbf{x}, \mathbf{q}) = a_i, \quad i = 1, \dots, r \\ & h_j(\mathbf{x}, \mathbf{q}) \leq b_j, \quad j = 1, \dots, s. \end{aligned} \tag{6.6}$$

In Problem (6.6), \mathcal{Q} is the space over which the optimization occurs. For any $\mathbf{q} \in \mathcal{Q}$ the state variable \mathbf{x} is uniquely determined. Therefore, the dependence on \mathbf{x} is removed from the problem. We can further remove the remaining equality constraints by considering the fact that the equality constraints perform the same function as parametric constraints, leading to the new system of equations $\bar{\mathbf{F}} : \mathbb{R}^D \times \mathbb{R}^p \rightarrow \mathbb{R}^{D+r}$

$$\bar{\mathbf{F}}(\mathbf{z}) = \begin{bmatrix} \mathbf{F}(\mathbf{z}) \\ g_1(\mathbf{z}) \\ \vdots \\ g_r(\mathbf{z}) \end{bmatrix} = \mathbf{0}. \tag{6.7}$$

If $D_{\mathbf{x}}\bar{\mathbf{F}}$ is full rank for all \mathbf{z} satisfying $\bar{\mathbf{F}}(\mathbf{z}) = \mathbf{0}$ then each set

$$\bar{\mathcal{M}} = \{(\mathbf{x}, \mathbf{q}) \mid \bar{\mathbf{F}}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \tag{6.8}$$

$$\bar{\mathcal{X}} = \{\mathbf{x} \mid \bar{\mathbf{F}}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \tag{6.9}$$

$$\bar{\mathcal{Q}} = \{\mathbf{q} \mid \bar{\mathbf{F}}(\mathbf{x}, \mathbf{q}) = \mathbf{0}\} \tag{6.10}$$

define a $(p - r)$ -dimensional smooth manifold. It is evident that $0 \leq r \leq (p - 1)$, since for $r = p$ there is at most a single solution in the space to optimize over, and $r > p$ results in no solutions. Then Problem (6.6) reduces to

$$\begin{aligned} \min_{\mathbf{q} \in \bar{\mathcal{Q}}} \quad & f(\mathbf{x}, \mathbf{q}) \\ \text{subject to} \quad & h_i(\mathbf{x}, \mathbf{q}) \leq b_i, \quad i = 1, \dots, s. \end{aligned} \tag{6.11}$$

Now each $\mathbf{q} \in \bar{\mathcal{Q}}$ satisfies \mathbf{F} and the equality constraints. Problem (6.6) effectively removes the equality constraints by considering a space of parameters which identically satisfy the equality constraints. We note that the equality constraints decrease the dimension of the feasible set of




Figure 6.1: Equality and inequality constraints reduce the feasible set of solutions.

solutions while the inequality constraints reduce the measure of the feasible set of solutions (Figure 6.1).

6.3 Derivative With Respect to the Parameters

Gradient-based methods are used to solve Problem (6.11). When there are no inequality constraints we call Problem (6.11) an unconstrained problem because the parameters are free to explore all of \mathcal{Q} or $\bar{\mathcal{Q}}$. The unconstrained problem is readily solved with gradient descent. The inclusion of inequality constraints leads to a constrained problem, and gradient descent alone is not enough to find the optimal solution if the optimal solution lies outside the set of feasible solutions. In the constrained problem we use an ALM to solve for the optimal solution. In the optimization problem we find \mathbf{z}^* minimizing a cost function f . Moreover, we consider the case of finding \mathbf{z}^* , the roots of a vector function \mathbf{g} . Solving for the roots of a vector function is closely related to minimizing a cost function with Newton's method. In the latter case, \mathbf{g} can be thought of as the gradient of an unknown cost function. In both cases we take derivatives with respect to the parameters. We proceed with an explanation of the derivative of \mathbf{g} with respect to \mathbf{q} .

Let $\mathbf{g} : \mathbb{R}^D \times \mathbb{R}^p \rightarrow \mathbb{R}^t$ with $t \geq 1$. Then

$$\begin{aligned} \frac{d\mathbf{g}}{d\mathbf{q}} &= \frac{\partial \mathbf{g}}{\partial \mathbf{x}} \frac{\partial \mathbf{x}}{\partial \mathbf{q}} + \frac{\partial \mathbf{g}}{\partial \mathbf{q}} \frac{\partial \mathbf{q}}{\partial \mathbf{q}} \\ &= \begin{bmatrix} \frac{\partial \mathbf{g}}{\partial \mathbf{x}} & \frac{\partial \mathbf{g}}{\partial \mathbf{q}} \end{bmatrix} \begin{bmatrix} \frac{\partial \mathbf{x}}{\partial \mathbf{q}} \\ \mathbb{I}_p \end{bmatrix} \\ &= \frac{d\mathbf{g}}{d\mathbf{z}} \frac{d\mathbf{z}}{d\mathbf{q}}. \end{aligned} \quad (6.12)$$

Here, we make two observations: First, $\frac{d\mathbf{g}}{d\mathbf{q}}$ is the standard Euclidean gradient. It is defined on \mathcal{Q} and not necessarily on $\bar{\mathcal{Q}}$. Second, $\frac{d\mathbf{z}}{d\mathbf{q}}$ facilitates how \mathbf{z} changes given unit changes in each parameter. The matrix $\frac{d\mathbf{z}}{d\mathbf{q}}$ is formed from a particular set of basis vectors which span $T_{\mathbf{z}}\mathcal{M}$. We encountered this matrix in Section 3.5 and called it the principle tangent basis. Here we give how to compute the principle tangent basis V_p . Let $\mathbf{z} \in \mathcal{M}$ and

$$V = \begin{bmatrix} V_{\mathbf{x}} \\ V_{\mathbf{q}} \end{bmatrix} \quad (6.13)$$

be a matrix whose columns are linearly independent vectors spanning $T_{\mathbf{z}}\mathcal{M}$. Furthermore, we require the $p \times p$ matrix $V_{\mathbf{q}}$ be full rank. Then the principle tangent basis is found as

$$V_p = V [V_{\mathbf{q}}^{-1} \mathbb{I}_p]. \quad (6.14)$$

Column i of V_p gives how \mathbf{z} changes with unit changes in q_i , or equivalently gives $\frac{d\mathbf{z}}{dq_i}$. Therefore, we have

$$V_p = \frac{d\mathbf{z}}{d\mathbf{q}}, \quad (6.15)$$

and the Euclidean derivative of \mathbf{g} with respect to \mathbf{q} is

$$\frac{d\mathbf{g}}{d\mathbf{q}} = \frac{d\mathbf{g}}{d\mathbf{z}} V_p \quad (6.16)$$

the matrix multiplication between the Euclidean derivative of \mathbf{g} , with respect to \mathbf{z} , and V_p . The derivative $\frac{d\mathbf{g}}{d\mathbf{z}}$ is readily obtained through analytical or numerical means, such as finite differencing or automatic differentiation. The principle tangent basis is readily obtained from the tangent space of \mathcal{M} . The tangent space is an output of the retraction.

We still require a step to transform the Euclidean derivative $\frac{d\mathbf{g}}{d\mathbf{q}}$ defined on \mathcal{Q} to a manifold derivative defined on $\bar{\mathcal{Q}}$. The transformation of the derivative is achieved by projecting the derivative onto $T_{\mathbf{z}}\bar{\mathcal{M}}$ through the projection operator. Now, we require $\mathbf{z} \in \bar{\mathcal{M}}$, and let \bar{V} be a matrix whose columns are linearly independent basis vectors spanning $T_{\mathbf{z}}\bar{\mathcal{M}}$. Then the projection operator of a matrix $A \in \mathbb{R}^{p \times a}$, $a \geq 1$, onto $T_{\mathbf{q}}\bar{\mathcal{Q}}$ is given as

$$\text{Proj}_{\mathbf{q}}(A) = \bar{V}_{\mathbf{q}} \left[\bar{V}_{\mathbf{q}}^{\dagger} A \right], \quad (6.17)$$

and the manifold derivative of \mathbf{g} with respect to the parameters is

$$\frac{d\bar{\mathbf{g}}}{d\mathbf{q}}^T = \text{Proj}_{\mathbf{q}} \left(\frac{d\mathbf{g}}{d\mathbf{q}}^T \right), \quad (6.18)$$

where $\bar{\mathbf{g}} = \mathbf{g}|_{\bar{\mathcal{M}}}$ is the restriction of \mathbf{g} to $\bar{\mathcal{M}}$. When \mathbf{g} is a scalar function, g , the Euclidean gradient of g with respect to \mathbf{q} is given by

$$\nabla_{\mathbf{q}} g(\mathbf{z}) = \frac{dg}{d\mathbf{q}}^T. \quad (6.19)$$

The manifold gradient is then

$$\nabla_{\mathbf{q}} \bar{g}(\mathbf{z}) = \text{Proj}_{\mathbf{q}} (\nabla_{\mathbf{q}} g(\mathbf{z})). \quad (6.20)$$

In the constrained and unconstrained optimization problems the choice of derivative type is dependent on whether equality constraints are included. As we have seen, equality constraints reduce the dimension of the search space, so the Euclidean gradient may point in a direction off of the manifold. To target solutions with specific characteristics, only Euclidean derivatives need to be computed. This is because we consider solutions in the entire p -parameter family, so all dimensions of the parameter space can be explored locally.

6.4 Targeting Characteristics

In this section we consider the root-finding problem to find a solution within a p -parameter family of solutions with desired characteristics. We do this by modifying a Newton's method to solve for the roots of a function $\mathbf{g} : \mathcal{M} \rightarrow \mathbb{R}^t$, $1 \leq t \leq p$. The function \mathbf{g} can be thought of as equality constraints or parametric constraints where each g_i specifies a characteristic we wish the




Figure 6.2: Example of two parametric constraints on a 2-dimensional manifold, leading to a single solution \mathbf{z}^* .

solution to have. Specifying p parametric constraints for a p -dimensional manifold generally leads to a single solution on the manifold (Figure 6.2). Other cases include multiple, isolated solutions, no solutions, and infinite solutions. For $1 \leq t < p$ there are generally an infinite number of solutions.

A Newton's method is used to update \mathbf{q} rather than \mathbf{z} . The updated parameters are used with the retraction to compute the solution at the updated parameter values. The reason for updating \mathbf{q} instead of \mathbf{z} is Newton's method is only able to converge on the solution, \mathbf{z}^* , satisfying

$$\mathbf{H}(\mathbf{z}) = \begin{bmatrix} \mathbf{F}(\mathbf{z}) \\ \mathbf{g}(\mathbf{z}) \end{bmatrix} = \mathbf{0}$$

if the initial guess, \mathbf{z}_0 , is reasonably close to the true solution. We call this the refinement of a solution. We leverage the retraction to remove \mathbf{F} , focusing solely on \mathbf{g} . The retraction ensures $\mathbf{F}(\mathbf{z}) = \mathbf{0}$ is satisfied at all times, and uses continuation which allows \mathbf{z}_0 to be far from \mathbf{z}^* .

Section 6.4.1 presents and describes the algorithm used to target the characteristics of a solution within a p -parameter family of solutions. We then validate the algorithm with various examples in Section 6.4.2, testing it on the 2-parameter family of quasi-halo orbits computed previously in Chapter 3.

6.4.1 Algorithm

We wish to find a solution $\mathbf{z}^* \in \mathcal{M}$ such that $\mathbf{g}(\mathbf{z}^*) = \mathbf{0}$. The algorithm to target the characteristics of a solution from within a p -parameter family of solutions is given in Algorithm 11. Given an initial solution $\mathbf{z}_0 \in \mathcal{M}$ and a matrix V_0 whose columns span the tangent space $T_{\mathbf{z}_0}\mathcal{M}$, the principle tangent basis is computed according to Equation (6.14), providing the derivative of \mathbf{z}_0 with respect to the parameters, $\frac{d\mathbf{z}}{d\mathbf{q}}|_{\mathbf{z}_0}$. The derivative of \mathbf{g} with respect to \mathbf{z} is computed and evaluated at \mathbf{z}_0 , $\frac{d\mathbf{g}}{d\mathbf{z}}|_{\mathbf{z}_0}$. The derivative of \mathbf{g} with respect to \mathbf{q} is computed with Equation (6.12) and evaluated at \mathbf{z}_0 .

A Newton update is computed to provide a change in parameters $\delta\mathbf{q}_1$. If the number of scalar equations composing \mathbf{g} is equal to p then the inverse in the Newton update is the usual inverse. If the number of scalar equations is less than p then the inverse is the Moore-Penrose inverse. The change in parameters dictates a direction and a distance to move in \mathcal{Q} , directly providing the set of parameters \mathbf{q}_1 of the solution \mathbf{z}_1 . The updated solution \mathbf{z}_1 and its tangent space V_1 is computed with the retraction in Algorithm 10 of Chapter 5. The actual updated parameters are not exactly equal to the guess provided by the Newton step, so \mathbf{q}_1 is taken from the returned solution \mathbf{z}_1 . For brevity, we leave out certain parameters from Algorithm 11 and assume the appropriate parameters are passed into Algorithm 10.

The iterations continue until either $\|\mathbf{g}(\mathbf{z}_k)\| \leq \varepsilon$ or $\|\mathbf{q}_{k+1} - \mathbf{q}_k\| \leq \varepsilon$, where ε is the error tolerance on the characteristics. The first termination condition guarantees the solution has characteristics close enough to the targeted characteristics. The second termination condition indicates convergence to a local minimum of \mathbf{g} .

Algorithm 11 IFT Newton's Method: Find a root of a function defined on an implicit manifold

```

1: procedure IFTNM( $\mathbf{z}_0, V_0, \mathbf{g}, \varepsilon$ )
2:    $k = 0$ 
3:    $(\mathbf{z}_k, V_k) = (\mathbf{z}_0, V_0)$ 
4:   while  $\|\mathbf{g}(\mathbf{z}_k)\| > \varepsilon$  do
5:      $\frac{d\mathbf{z}}{d\mathbf{q}}|_{\mathbf{z}_k} = V_k \left[ V_{\mathbf{q},k}^{-1} \mathbb{I}_p \right]$ 
6:      $\frac{d\mathbf{g}}{d\mathbf{q}}|_{\mathbf{z}_k} = \left[ \frac{d\mathbf{g}}{d\mathbf{z}} \frac{d\mathbf{z}}{d\mathbf{q}} \right]_{\mathbf{z}_k}$ 
7:      $\delta \mathbf{q}_{k+1} = - \left[ \frac{d\mathbf{g}}{d\mathbf{q}} \right]_{\mathbf{z}_k}^\dagger \mathbf{g}(\mathbf{z}_k)$ 
8:      $(\mathbf{z}_{k+1}, V_{k+1}) = \text{RETRACT}(\delta \mathbf{q}_{k+1}, \mathbf{z}_k, V_k)$ 
9:     if  $\|\mathbf{q}_{k+1} - \mathbf{q}_k\| \leq \varepsilon$  then
10:      return  $(\mathbf{z}_{k+1}, V_{k+1})$ 
11:    end if
12:     $k \leftarrow k + 1$ 
13:  end while
14:   $(\mathbf{z}^*, V^*) = (\mathbf{z}_k, V_k)$ 
15:  return  $(\mathbf{z}^*, V^*)$ 
16: end procedure

```

Table 6.1: Test cases for targeting solution characteristics

Case #	$\mathbf{g}(\mathbf{z})$	Run Time (min)	# Iterations
1	$A(\mathbf{X}) - A_1^*$	3.7	4
2	$\bar{J}(\mathbf{X}) - J^*$	2.4	3
3	$\begin{bmatrix} A(\mathbf{X}) - A_1^* \\ \bar{J}(\mathbf{X}) - J^* \end{bmatrix}$	3.2	4
4	$\begin{bmatrix} A(\mathbf{X}) - A_1^* \\ \omega_0 - \omega_0^* \end{bmatrix}$	4.1	4
5	$\lambda_u(\mathbf{z}) - \lambda_u^*$	12	3
6	$\begin{bmatrix} \bar{J}(\mathbf{X}) - J^* \\ \lambda_u(\mathbf{z}) - \lambda_u^* \end{bmatrix}$	19.2	4

6.4.2 Validation

Algorithm 11 is tested on the 2-parameter family of quasi-halo orbits, so we have

$$\mathbf{x} = \begin{bmatrix} \mathbf{X} \\ T \\ \rho_1 \end{bmatrix}, \quad \mathbf{q} = \begin{bmatrix} \omega_0 \\ \omega_1 \end{bmatrix}, \quad \mathcal{Q} = \Omega.$$

The maximum number of characteristics we can choose an orbit to have is 2. For a single characteristic equation there are an infinite number of solutions which can be targeted. As long as the two characteristics generate intersecting contours on Ω there is a solution \mathbf{z}^* such that $\mathbf{g}(\mathbf{z}^*) = \mathbf{0}$.

We test the algorithm on six different characteristic functions. The cases, run time, and number of Newton iterations are given in Table 6.1. Case 1 targets a solution with a specified orbit amplitude A_1 . The amplitude is computed by Algorithm 3. Case 2 targets a solution with a specified Jacobi constant. Case 3 targets a solution with a specified amplitude size and Jacobi constant. Case 4 targets a solution with a specified orbit amplitude and frequency ω_0 . Case 5 targets a solution with a specified value of the unstable eigenvalue. Lastly, case 6 targets a solution with a specified Jacobi constant and unstable eigenvalue. Cases 1, 2, and 5 have an infinite number of solutions, while cases 3, 4, and 6 have a single unique solution. In cases 3, 4, and 6 we choose characteristic values from pre-computed quasi-halo orbits, so there is a truth target orbit to compare to. In all cases we use an error tolerance, ε_0 , of 1e-7 for Algorithm 11 and 1e-8 for Algorithm 10.

The results of each case are given in Figures 6.3-6.8. Plots (a) and (b) show the iterates $\boldsymbol{\omega}_k$.

Plot (c) shows the error of each characteristic function versus the iteration number. Plot (d) shows the invariant curve of each solution \mathbf{X}_k . Cases 3, 4, and 6 show the truth target invariant curve in plot (d). All cases converge in 3 or 4 iterations. Cases 1 through 4 converge in 4 minutes or less, while cases 5 and 6 take longer. The increased time of cases 5 and 6 are due to the computation time required to compute the derivative of the unstable eigenvalue with respect to \mathbf{z} . There is no analytical means to compute the derivative of the unstable eigenvalue, so we must resort to numerical derivatives. Finite differencing is prohibitively slow, so we use automatic differentiation to perform the numerical derivatives.

Cases 1 through 4 find solutions within the error tolerance. Cases 5 and 6 terminate before the error tolerance is reached because the difference between the current and updated frequencies are within tolerance. Even with the premature termination of the algorithm the last iteration of case 6 is nearly identical to the truth target solution as seen in plot (d) of Figure 6.8. In plot (c) of Figure 6.6 the error in ω_0 reached tolerance on the first iteration, and exactly matched the targeted value in the subsequent iterations. In cases 1, 2, and 5 a solution is still found even though $\frac{d\mathbf{g}}{d\mathbf{q}}$ is not invertible in the usual sense.

6.4.3 Discussion

Algorithm 11 converges on the target solution quickly even with initial solutions not near the truth solution. Treating $\boldsymbol{\omega}$ as the free vector and \mathbf{x} as a dependent vector to be found by the retraction shows improved performance over a standard Newton's method which treats \mathbf{z} as the free vector.

With Algorithm 11 we can construct grids of solutions over many sets of characteristics at desired resolutions. From the grid points one can quickly and easily find the desired solution with great accuracy without the need for a nearby solution or a dense grid. Interpolation methods, such as in [136], require dense grids of solutions to produce approximate solutions which have to be refined.

The characteristic values do not need to be known exactly for Algorithm 11 to work. Consider

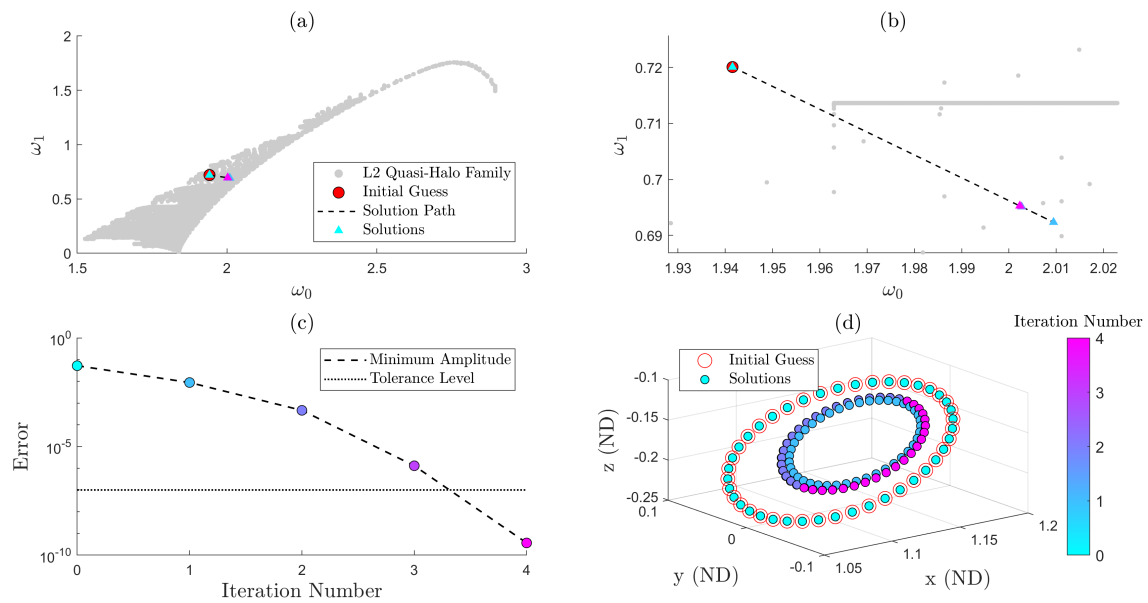


Figure 6.3: Results of case 1.

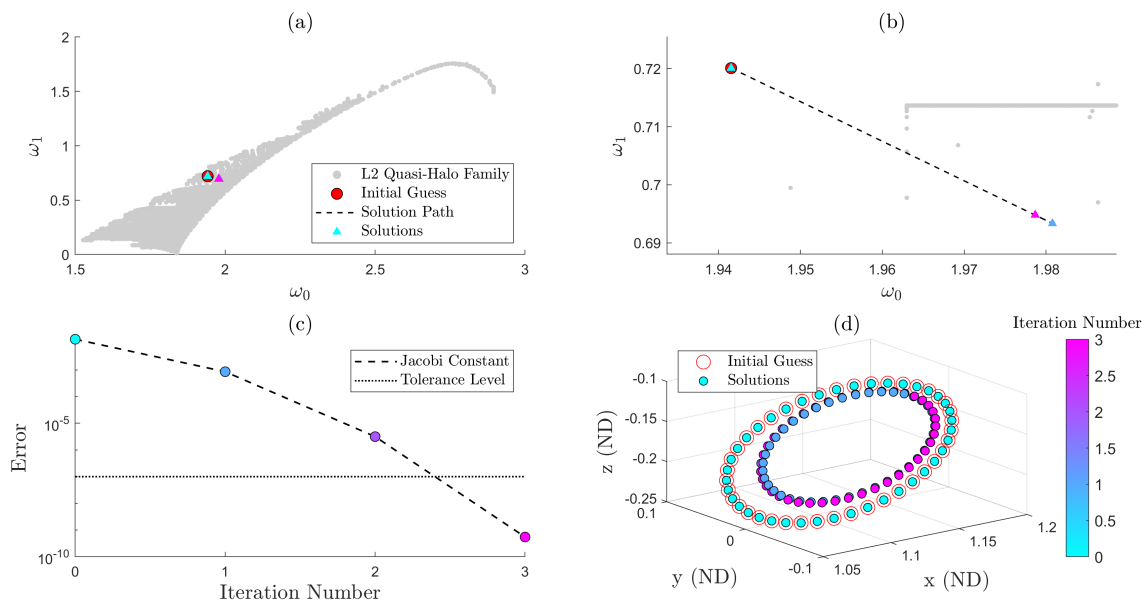


Figure 6.4: Results of case 2.

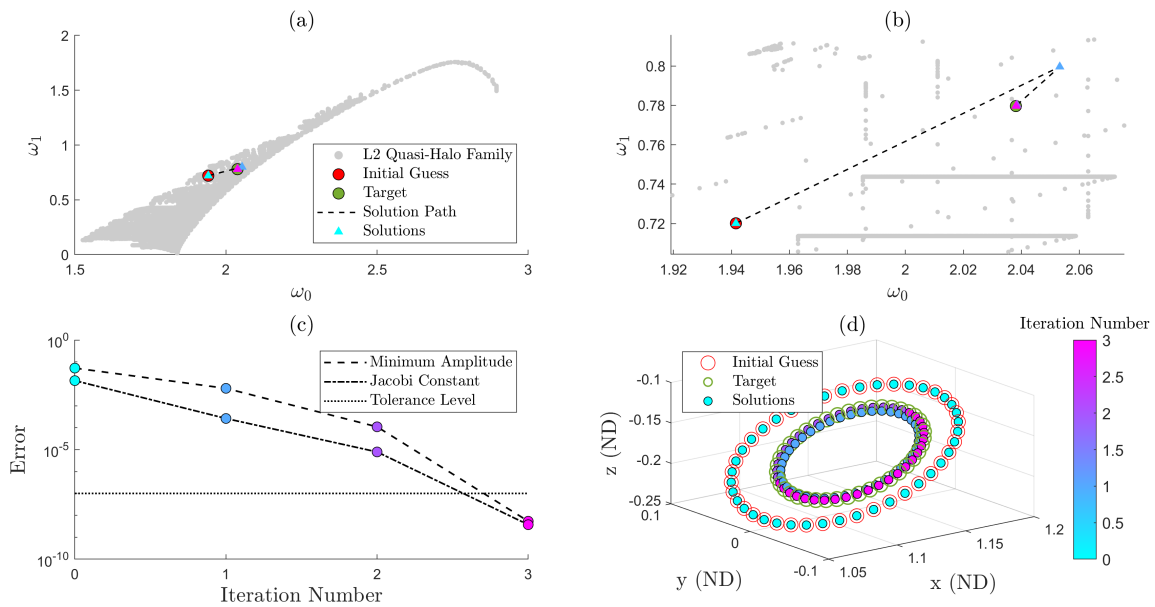


Figure 6.5: Results of case 3.

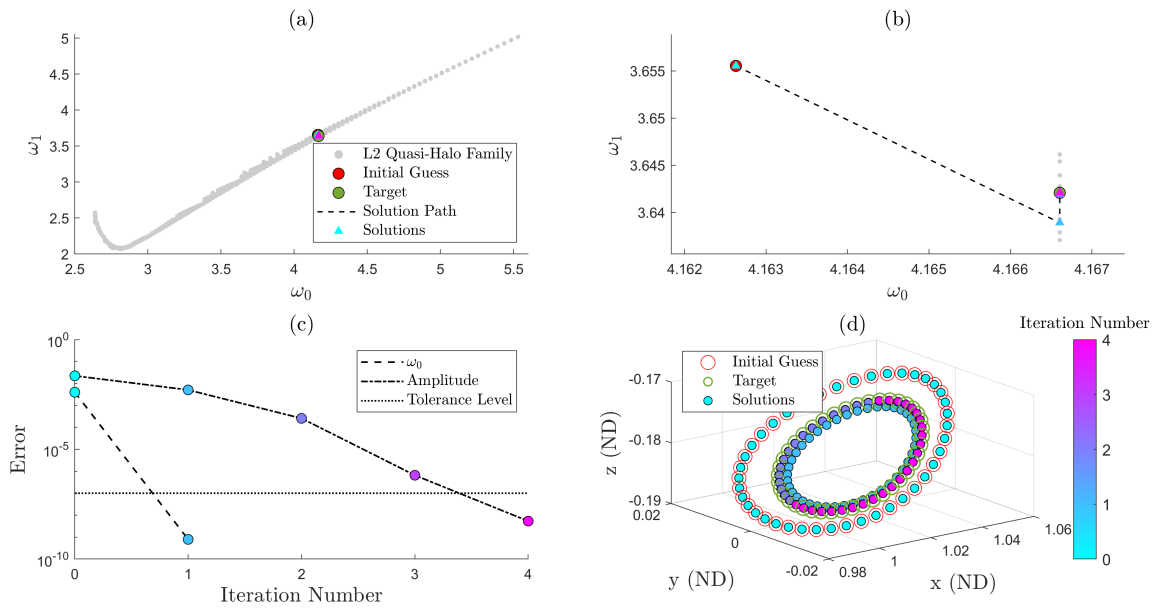


Figure 6.6: Results of case 4.

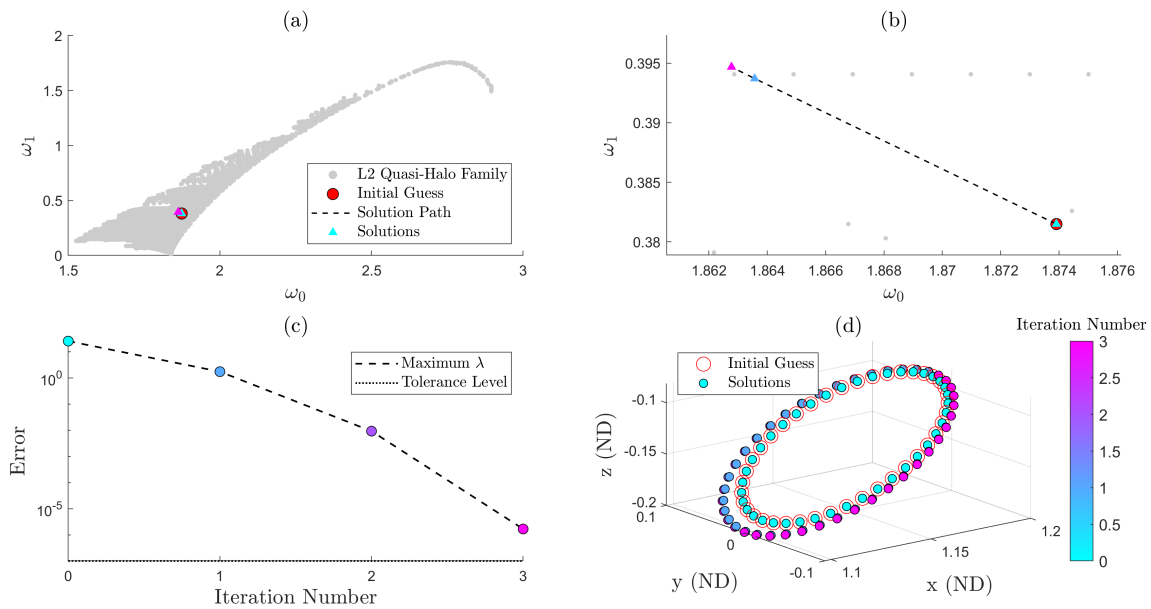


Figure 6.7: Results of case 5.

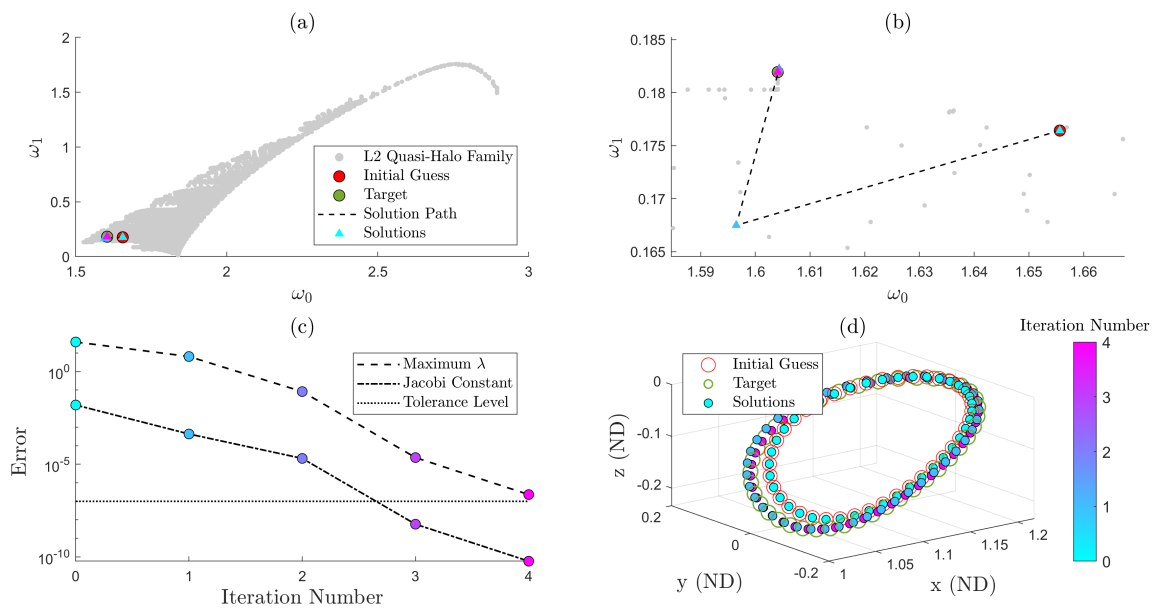


Figure 6.8: Results of case 6.

we know the values exactly, and they correspond to a solution \mathbf{z}^* displayed in Figure 6.2. Suppose there are errors in the specified characteristic values. Graphically, this corresponds to shifting the contours in Figure 6.2 slightly, resulting in a nearby solution. Suppose the characteristic values do not correspond to any solution in the family of interest, meaning there will be no real roots of \mathbf{g} . Algorithm 11 can handle this case without issues. Rather than finding a root of \mathbf{g} , Algorithm 11 will either find a solution which is a local minimum of \mathbf{g} or will reach the boundary of the family and stop. The returned solution is likely to be the solution with the characteristics closest to the desired ones.

Algorithm 11 can be used to construct advanced resonance avoidance methods. Suppose we are continuing a branch of quasi-periodic orbits within a multi-parameter family and encounter a resonance. From Chapter 4 we suspect the best way to avoid the resonance is to hop over it perpendicularly. Doing so results in a solution which is not on the original branch of solutions. Algorithm 11 can be used to find a solution on the original branch of solutions which is known to be on the appropriate side of the resonance, so that the continuation of the branch of solutions can be reinitialized.

6.5 Unconstrained Optimization

The standard gradient descent method presented in Section 2.7.1 assumes all the variables in \mathbf{z} are independent optimization variables and optimizes over a Euclidean space. For solutions $\mathbf{z} \in \bar{\mathcal{M}}$ only \mathbf{q} are the independent optimization variables, and the space $\bar{\mathcal{Q}}$ is not Euclidean, requiring manifold derivatives and the retraction presented in Chapter 5. Equality constraints are absorbed by the retraction, reducing the search space to the only space the optimization is aware of. Therefore, including equality constraints still results in the unconstrained optimization problem

$$\min_{\mathbf{q} \in \bar{\mathcal{Q}}} f(\mathbf{x}, \mathbf{q}). \quad (6.21)$$

In Section 6.5.1 we present, what we call, the IFT gradient descent algorithm. In Section 6.5.2 we validate it with two test cases and optimize over the family of quasi-halo orbits.

6.5.1 Algorithm

The IFT gradient descent is similar to the Euclidean gradient descent except the gradient is taken only with respect to the parameters \mathbf{q} and the gradient is a manifold gradient when there are equality constraints, \mathbf{g} . The IFT gradient descent aims to find a local minimum solution \mathbf{z}^* of the function $f : \bar{\mathcal{M}} \rightarrow \mathbb{R}$. Beginning with an initial solution $\mathbf{z}_0 \in \bar{\mathcal{M}}$ and \bar{V}_0 we compute the gradient of f with respect to the parameters \mathbf{q}_0 with Equation (6.20). The initial descent direction, \mathbf{d} , is given as

$$\mathbf{d} = -\nabla_{\mathbf{q}} \bar{f}(\mathbf{z}_0).$$

The descent direction provides the direction to search on $\bar{\mathcal{Q}}$ to find solutions which reduce the value of the cost function. While the magnitude of \mathbf{d} provides a distance to travel, the descent direction only provides local information. Traveling exactly a distance $d = \|\mathbf{d}\|$ may result in a solution with a cost larger than the initial solution. To ensure the next solution, \mathbf{z}_1 , has a cost less than the initial solution a line search is employed. The descent direction is computed for \mathbf{z}_1 and the process continues until either $\|\nabla_{\mathbf{q}} \bar{f}(\mathbf{z}_k)\| \leq \varepsilon$ or $\|\mathbf{z}_{k+1} - \mathbf{z}_k\| \leq \varepsilon$.

Ideally, we want the next iterate to be the solution in the direction of \mathbf{d} which has the lowest cost associated with it. That is we want to find the solution which attains the minimum of the set of points

$$L_\alpha = \{f(\mathbf{z}) \mid \mathbf{z} = \text{RETRACT}(\gamma \mathbf{d}, \mathbf{z}_k, \bar{V}_k), \gamma \in [0, \alpha]\} \quad (6.22)$$

for a sufficiently large $\alpha > 0$. Finding the exact optimal solution of L_α is computationally expensive since each solution must be found with the retraction function. The retraction uses continuation, so we leverage this to construct L_α . A small initial step size, Δs_0 , is chosen so that many solutions are computed near \mathbf{z}_k , giving a fine resolution near \mathbf{z}_k and less resolution further away. We don't want to compute too many solutions in the continuation, and we don't want to pass over the optimal solution to Problem 6.21 when \mathbf{z}_k is near the optimal solution. The cost is computed for each solution in the continuation and the solution with the minimum cost is chosen as the next iterate in the gradient descent algorithm.

Algorithm 12 IFT Gradient Descent: Find minimum of a cost function defined over an implicit manifold

```

1: procedure IFTGD( $f, \mathbf{z}_0, \bar{V}_0, \varepsilon$ )
2:    $k = 0$ 
3:    $(\mathbf{z}_k, \bar{V}_k) = (\mathbf{z}_0, \bar{V}_0)$ 
4:   while  $\|\nabla_q \bar{f}(\mathbf{z}_k)\| > \varepsilon$  do
5:      $\mathbf{d} = -\nabla_q \bar{f}(\mathbf{z}_k)$ 
6:      $(\mathbf{z}_{k+1}, \bar{V}_{k+1}) = \text{LINESEARCH}(f, \mathbf{d}, \mathbf{z}_k, \bar{V}_k)$ 
7:     if  $\|\mathbf{z}_{k+1} - \mathbf{z}_k\| \leq \varepsilon$  then
8:       return  $(\mathbf{z}_{k+1}, \bar{V}_{k+1})$ 
9:     end if
10:     $k \leftarrow k + 1$ 
11:  end while
12:   $(\mathbf{z}^*, \bar{V}^*) = (\mathbf{z}_k, \bar{V}_k)$ 
13:  return  $(\mathbf{z}^*, \bar{V}^*)$ 
14: end procedure

```

Table 6.2: Test cases for unconstrained optimization

Case #	$f(\mathbf{z})$	$\mathbf{g}(\mathbf{z})$	Run Time (min)	# Iterations
1	$\frac{1}{2} (\bar{J}(\mathbf{X}) - J^*)^2$	-	29.3	-
2	$\frac{1}{2} (\bar{J}(\mathbf{X}) - J^*)^2$	$\omega_0 - \omega_0^*$	4	-

6.5.2 Validation

We test Algorithm 12 on the family of quasi-halo orbits to find an orbit which minimizes the distance squared to the desired Jacobi constant value of $J^* = 3.055460$. The first case does not include equality constraints, having two significant effects. First, we only need to take a Euclidean gradient. Second, there are an infinite number of solutions to Problem 6.21 having a Jacobi constant with a specific value. In the second case we include a constant ω_0 equality constraint, identifying a single solution to Problem 6.21. The two cases are given in Table 6.2 along with the run time. Due to the specific implementation of Algorithm 12 we did not record the number of gradient descent iterations needed for convergence.

The results of case 1 and 2 are in Figure 6.9 and 6.10, respectively. In Figure 6.9 we see the descent direction lies perpendicular to the Jacobi constant contour, showing the continuation path follows the path of steepest descent. It is not clear why the run time is nearly 30 minutes, however this time can surely be decreased by tuning the line search. We note the use of Algorithm 11 is likely to be more efficient to find a solution with the desired Jacobi constant.

In case 2 we constrain the optimal solution to have an ω_0 value of 1.947982. The initial solution does not meet this constraint, so Algorithm 11 is first used to find a solution with the desired ω_0 to initialize Algorithm 12. The gradient of f is projected to the tangent space of the 1-dimensional manifold $\bar{\mathcal{M}}$, pointing parallel to the constant ω_0 line. The continuation path follows the constant ω_0 constraint until the gradient vanishes at the optimal solution on the Jacobi constant line. Case 2 has a 7x speed-up over case 1, showing that equality constraints speed up the optimization.

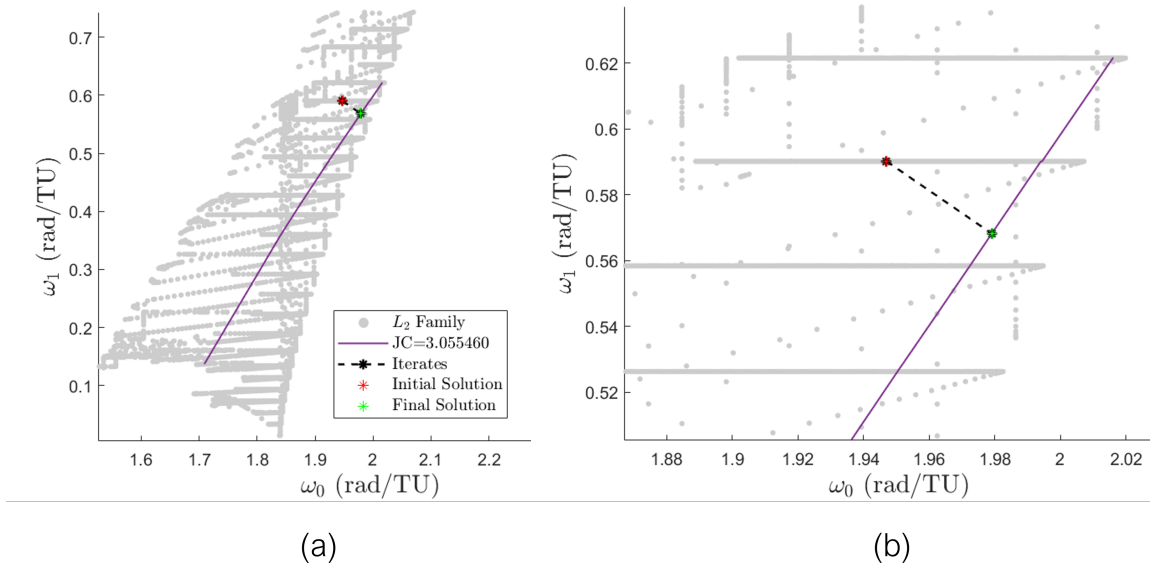


Figure 6.9: Results of case 1. Plot (a) is a zoomed out version of plot (b).

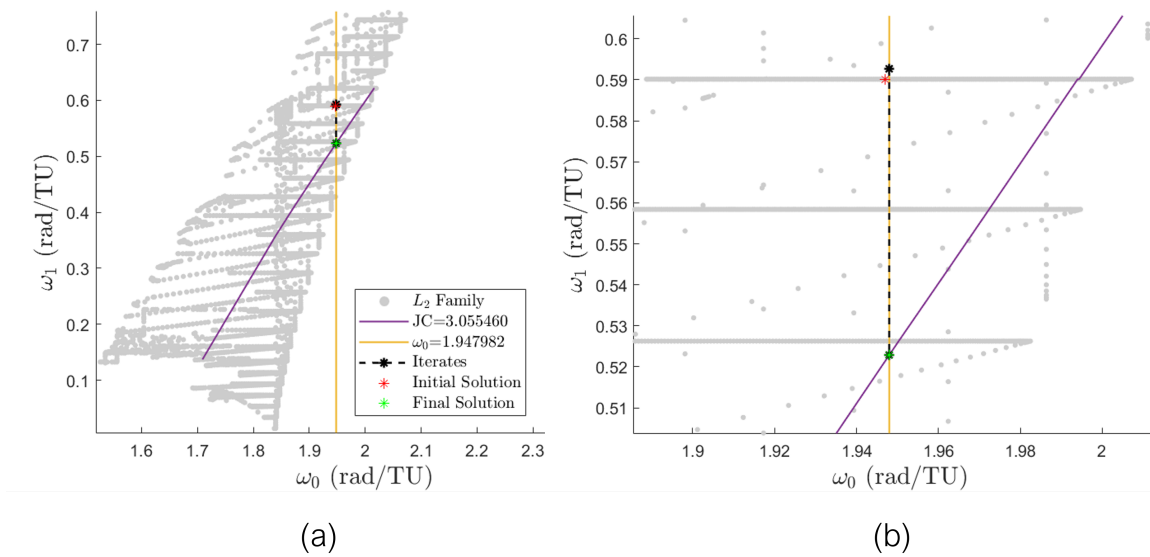


Figure 6.10: Results of case 2. Plot (a) is a zoomed out version of plot (b).

6.6 Constrained Optimization

An ALM constructs a modified cost function called the augmented Lagrangian given by

$$\mathcal{L}_\rho(\mathbf{x}, \boldsymbol{\eta}, \boldsymbol{\mu}) = f(\mathbf{x}) + \frac{\rho}{2} \left(\sum_{i=1}^r \left[g_i(\mathbf{x}) + \frac{\eta_i}{\rho} \right]^2 + \sum_{i=1}^s \left[\max \left(0, h_i(\mathbf{x}) + \frac{\mu_i}{\rho} \right) \right]^2 \right), \quad (6.23)$$

where $\rho > 0$ is a penalty parameter, $\boldsymbol{\eta} \in \mathbb{R}^r$, and $\boldsymbol{\mu} \in \mathbb{R}_+^s$ are Lagrange multipliers. An ALM solves the unconstrained subproblem

$$\min_{\mathbf{x} \in \mathbb{R}^d} \mathcal{L}_\rho(\mathbf{x}, \boldsymbol{\eta}, \boldsymbol{\mu}) \quad (6.24)$$

successively while updating ρ , $\boldsymbol{\eta}$, and $\boldsymbol{\mu}$ after each solve. The method to solve the subproblem can be any method to solve unconstrained optimization problems. We call the method to solve the subproblems the SUBSOLVER routine. We refer the reader to the book by Birgin and Martinez [27] and to Chapter 17 in [138] to learn more about ALMs.

Since all equality constraints are pulled into the retraction we only need to treat the inequality constraints, making $\boldsymbol{\eta} \equiv \mathbf{0}$, and the unconstrained subproblem becomes

$$\min_{\mathbf{q} \in \mathcal{Q}} \mathcal{L}_\rho(\mathbf{z}, \boldsymbol{\mu}). \quad (6.25)$$

6.6.1 Algorithm

Before introducing the algorithm we present the clip operator defined by

$$\text{clip}_{[a,b]}(x) = \max(a, \min(b, x)), \quad (6.26)$$

define $\mathcal{I} = 1, \dots, s$ to be the set of indices for the inequality constraints, and let

$$\mathbf{h}(\mathbf{z}) = [h_1(\mathbf{z}), \dots, h_s(\mathbf{z})]^T.$$

The algorithm used to solve the constrained optimization problems in this thesis is given in Algorithm 13 and pulled from Liu and Boumal in [118], however slight modifications have been made to it.

The variable $\boldsymbol{\mu}_0$ is the initial values of the Lagrange multipliers. We set these initial values to 1e-8. The variables $\boldsymbol{\mu}_{\min}$ and $\boldsymbol{\mu}_{\max}$ are minimum and maximum bounds to the Lagrange multipliers,

and set $\mu_{\min} = 1e-12$ and $\mu_{\max} = 1e+9$. The variable ρ_0 is an initial penalty parameter, and we set $\rho_0 = 1e-4$. The multiplier $\theta_\rho > 1$ determines the rate of growth of the penalty parameter, and we set $\theta_\rho = 20$. The variables ε_0 , ε_{\min} , and $\theta_\varepsilon \in (0, 1)$ play a role in determining how many iterations are necessary before convergence can be declared. We set $\varepsilon_0 = 1$, $\varepsilon_{\min} = 1e - 8$, and $\theta_\varepsilon = 0.005$. The multiplier θ_σ determines a threshold for the relative amount of movement in the inequality constraints in order to adjust the penalty parameter. Essentially, if the values of \mathbf{h}_{k+1} do not change much from the previous iteration then the penalty parameter is increased. We set this value to be 0.8. The constant d_{\min} is an error bound. When the distance between two consecutive solutions is less than this amount then it is likely that the algorithm is converging to an optimal solution. However, the convergence decision is balanced with the number of iterations which have satisfied $\|\mathbf{z}_{k+1} - \mathbf{z}_k\| < d_{\min}$, thereby increasing the confidence that the returned solution is actually an optimal solution to Problem (6.11). We set $d_{\min} = 1e-6$. Finally, N_{\max} limits the number of iterations, so the algorithm cannot continue forever. We set $N_{\max} = 50$.

Algorithm 13 Augmented Lagrangian Method

```

1: procedure ALM( $f, \mathbf{h}, \mathbf{z}_0, \bar{V}_0, \mu_0, \mu_{\min}, \mu_{\max}, \rho_0, \theta_\rho, \varepsilon_0, \varepsilon_{\min}, \theta_\varepsilon, \theta_\sigma, d_{\min}, N_{\max}$ )
2:    $(\mathbf{z}_k, \bar{V}_k) = (\mathbf{z}_0, \bar{V}_0)$ 
3:   for  $k = 0, 1, \dots, N_{\max}$  do
4:      $(\mathbf{z}_{k+1}, \bar{V}_k) = \text{SUBSOLVER}(\mathcal{L}_{\rho_k}(\mathbf{z}, \boldsymbol{\mu}_k), \mathbf{z}_k, \bar{V}_k)$ 
5:     if  $\|\mathbf{z}_{k+1} - \mathbf{z}_k\| < d_{\min}$  and  $\varepsilon_k \leq \varepsilon_{\min}$  then
6:       return  $\mathbf{z}_{k+1}$ 
7:     end if
8:      $\mu_{k+1}^i = \text{clip}_{[\mu_{\min}^i, \mu_{\max}^i]}(\mu_k^i + \rho_k h_i(\mathbf{z}_{k+1}))$ , for  $i \in \mathcal{I}$ 
9:      $\sigma_{k+1}^i = \max\left(h_i(\mathbf{z}_{k+1}), -\frac{\mu_k^i}{\rho_k}\right)$  for  $i \in \mathcal{I}$ 
10:    if  $\|\mathbf{z}_{k+1} - \mathbf{z}_k\| < d_{\min}$  and  $\rho > 1$  then
11:       $\varepsilon_{k+1} = \max(\varepsilon_{\min}, \theta_\varepsilon \varepsilon_k)$ 
12:    end if
13:    if  $k = 0$  or  $\max_{i \in \mathcal{I}}(|\sigma_{k+1}^i|) \leq \theta_\sigma \max_{i \in \mathcal{I}}(|\sigma_k^i|)$  then
14:       $\rho_{k+1} = \rho_k$ 
15:    else
16:       $\rho_{k+1} = \theta_\rho \rho_k$ 
17:    end if
18:  end for
19:   $\mathbf{z}^* = \mathbf{z}_{k+1}$ 
20:  return  $\mathbf{z}^*$ 
21: end procedure

```

After Algorithm 13 returns a solution \mathbf{z}^* we check the gradient of $\mathcal{L}_{\rho^*}(\mathbf{z}^*, \boldsymbol{\mu}^*)$ to determine whether the returned solution is optimal. If $\|\nabla \mathcal{L}_{\rho^*}(\mathbf{z}^*, \boldsymbol{\mu}^*)\| \leq \varepsilon_{\min}$ and each $h_i(\mathbf{z}^*) \leq 0$ for $i \in \mathcal{I}$, then \mathbf{z}^* is declared to be a solution to Problem (6.11).

6.6.2 Validation

We validate Algorithm 13, using Algorithm 12 in place of SUBSOLVER, with the three test cases. The cases are given in Table 6.3 along with the run time and the number of iterations for the ALM to converge. In cases 1 and 2 we minimize the distance squared to the desired Jacobi constant value of $J^* = 3.055460$. Case 1 does not have an equality constraint, while case 2 has a constant ω_0 equality constraint. Both cases constrain ω_1 to be less than 0.518223. In case 3 we minimize the square of the logarithm of the unstable eigenvalue and place constraints on the minimum and maximum values of both ω_0 and the Jacobi constant. The inequality constraints nearly form a quadrilateral in frequency space. Equality constraints make the problem easier by reducing the dimension of the search space, so we optimize with only inequality constraints in case 3.

In all three cases we know a priori where the optimal solution lies. In case 1, any quasi-halo orbit with the desired Jacobi constant and $\omega_1 \leq 0.518223$ is an optimal solution. In case 2, there is a single optimal solution. The optimal solution is the quasi-halo with orbit frequencies $\boldsymbol{\omega} = (1.947982, 0.518223)$. In case 3, the optimal solution lies on the boundary of the feasible set at the intersection of the lines

$$\begin{aligned}\bar{J}(\mathbf{X}) &= 3.030083 \\ \omega_0 &= 1.884444.\end{aligned}$$

The location of the optimal solution can be inferred from Figure 3.21.

The results of case 1 through 3 are in Figures 6.11-6.13, respectively. The black stars are solution to each subproblem in the ALM, beginning with the initial solution and ending with the optimal solution. In case 1 the feasible set of solutions lies below the $\omega_1 = 0.518223$ line. In Figure

Table 6.3: Test cases for constrained optimization

Case #	$f(\mathbf{z})$	$\mathbf{g}(\mathbf{z})$	$\mathbf{h}(\mathbf{z})$	Run Time (hr)	# Iterations
1	$\frac{1}{2} (\bar{J}(\mathbf{X}) - J^*)^2$	-	$\omega_1 - \omega_1^*$	2.1	9
2	$\frac{1}{2} (\bar{J}(\mathbf{X}) - J^*)^2$	$\omega_0 - \omega_0^*$	$\omega_1 - \omega_1^*$	0.4	20
3	$\frac{1}{2} \log_{10}(\lambda_u(\mathbf{z}))^2$	-	$\begin{bmatrix} \omega_0 - \omega_0^* \\ \hat{\omega}_0 - \omega_0 \\ \bar{J}(\mathbf{X}) - J^* \\ \hat{J} - \bar{J}(\mathbf{X}) \end{bmatrix}$	24.7	18

6.11 we see the first iteration of the ALM finds a solution near the desired Jacobi constant because the penalty parameter and Lagrange multipliers are very small. At the first solution the inequality constraint is not satisfied, so the penalty parameter and Lagrange multipliers are modified to construct a new augmented Lagrangian with a different vector field. As the iterations continue the penalty parameter and Lagrange multipliers continue to be modified until an optimal solution is found.

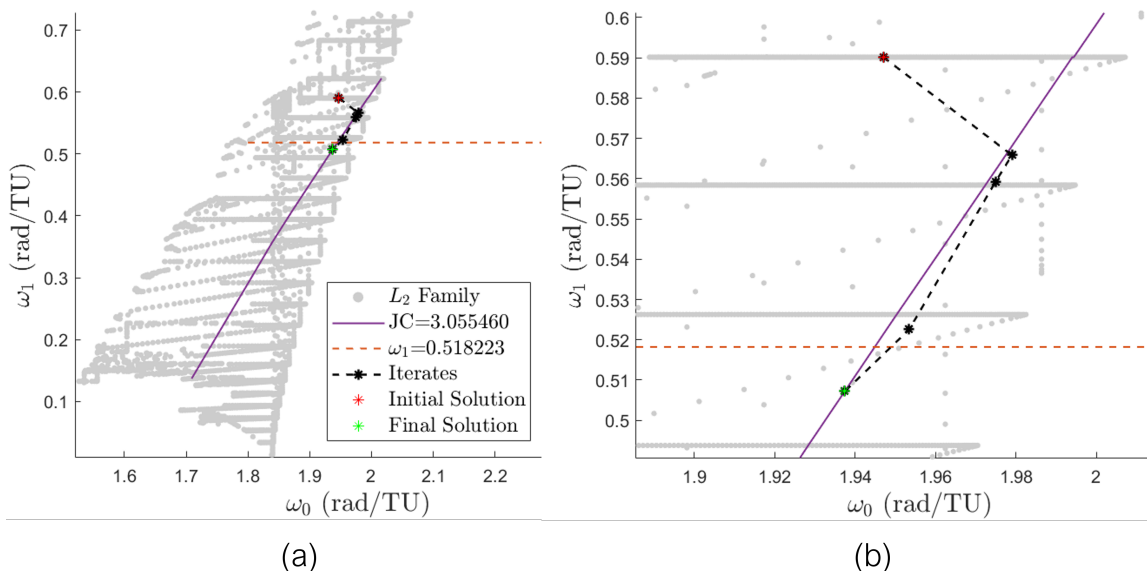


Figure 6.11: Orbit frequency iterates to minimize the difference in the Jacobi constant with an ω_1 inequality constraint.

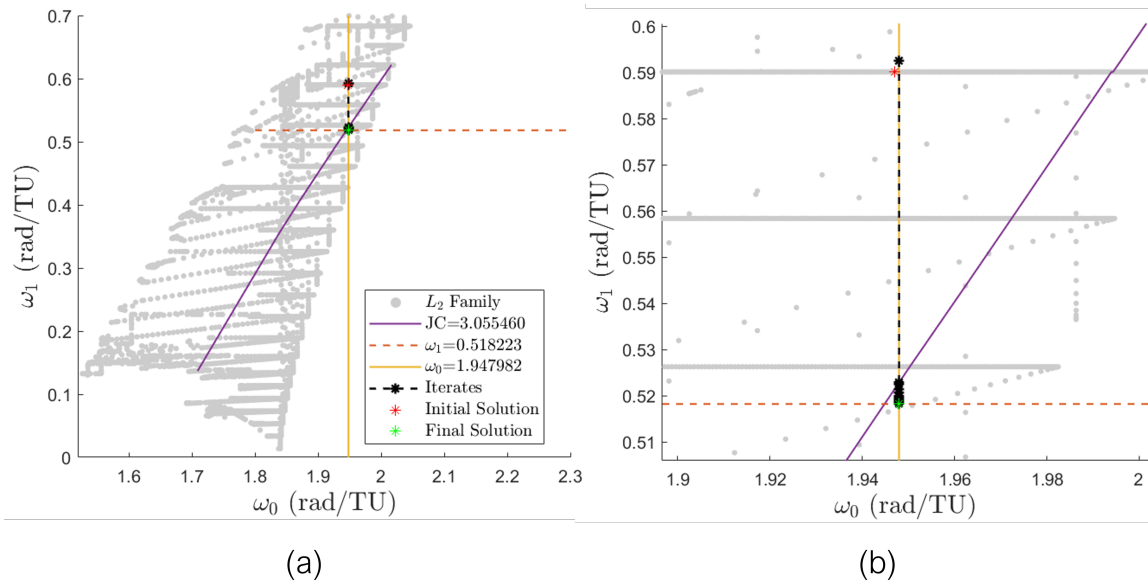


Figure 6.12: Orbit frequency iterates to minimize the difference in the Jacobi constant with an ω_0 equality constraint and ω_1 inequality constraint.

Case 2 begins with the same initial solution as case 1 and optimizes the same cost function, however case 2 has an equality constraint. The feasible set of solutions is the set of all quasi-halo orbits with $\omega_0 = 1.947982$ and $\omega_1 \leq 0.518223$. In Figure 6.12 we see the initial quasi-halo does not satisfy the equality constraint, so Algorithm 11 is used to initialize Algorithm 13. The first iteration of the ALM finds a quasi-halo orbit near the desired Jacobi constant. The remaining iterates are found with ever smaller ω_1 values as the penalty parameter and the Lagrange multipliers are modified until the final solution satisfies the inequality constraint such that $\omega_1 = 0.518223$.

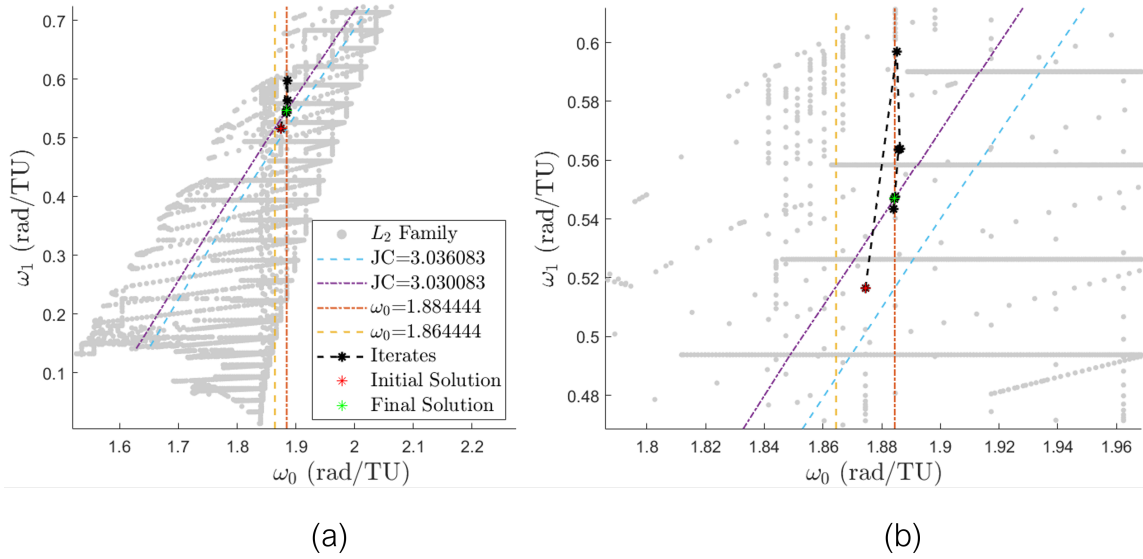


Figure 6.13: Orbit frequency iterates to minimize the unstable eigenvalue with Jacobi constant and ω_0 inequality constraints.

For case 3 we adjust the initial value of the penalty parameter to be equal to 1000 because we know the optimal solution to the unconstrained problem is far from the feasible set of solutions. The feasible set of solutions lies on the interior and boundary of the area encompassed by the four inequality constraints. The time to find the optimal solution is quite high, at 24 hours, due to the high cost of evaluating the gradient of the augmented Lagrangian. There are other ways to compute the optimal solution in this situation. We could have used Algorithm 11 to compute the solution with $\omega_0 = 1.884444$ and a Jacobi constant of 3.030083. However, in other cases the location of the optimal solution may not be available. The family may not be known in great detail ahead of time, or the location of the optimal solution may not be easy to determine.

6.7 Conclusion

This chapter utilizes the retraction of Chapter 5 to construct tools to optimize over families of solutions. Three main tools were developed: a tool to precisely target the characteristics of a solution, a tool to perform unconstrained optimization, and a tool to perform constrained optimization. The tools apply to manifolds defined by the implicit function theorem. Each tool was

tested on a variety of test cases within the 2-parameter family of Earth-Moon L_2 quasi-halo orbits, displaying novel capabilities for mission design. The cumulative results from all the cases show the retraction is effective at methodically moving through multi-parameter families to find desired orbits. Limited knowledge of an orbit family is needed to apply the tools, indicating a dense grid of solutions is not needed, saving computation time, memory, and storage space.

Chapter 7

Concluding Remarks

7.1 Conclusion

Quasi-periodic invariant tori and their invariant manifolds organize the dynamics in astrodynamics. These invariant structures unlock complex and efficient mission designs, enabling greater scientific returns and the advancement of human civilization. Thus, it is of great importance to develop the knowledge, tools, and framework to utilize quasi-periodic invariant tori in astrodynamics. The work in this thesis aims to better understand quasi-periodic orbits, the families they compose, and how to compute orbits meeting mission constraints. These goals were accomplished by presenting in great detail a versatile single-parameter continuation method to compute d -dimensional quasi-periodic invariant tori, computing and studying the Earth-Moon L_2 quasi-halo orbit family, and developing tools to precisely target orbits within single- and multi-parameter families.

This thesis is meant to be readable for a wide range of audiences with limited background. Explaining complex material in a simple way without completely sacrificing the rigor provides a good entry point into the field of quasi-periodic invariant tori. Many of the tools developed here are applicable to fields outside of astrodynamics, so we attempted to keep discussions general. Each chapter builds upon the previous chapters, so it is important to read each chapter in the order they are presented. Chapter 2 sets the stage for the entire thesis. We provide the mathematical framework on which this work is built, and define many of the terms used throughout the thesis. We began with simple concepts, and gradually build up to more complicated ideas, introducing commonly used terminology used throughout the literature.

Chapter 2 begins with the ideas of quasi-periodic tori and moves to dynamical systems theory. We explain the relevant types of invariant dynamical structures, allowing us to define quasi-periodic invariant tori - the centerpiece of this thesis. From there we describe the circular restricted three-body problem - the setting in which we compute quasi-periodic invariant tori. In this dynamical system quasi-periodic invariant tori form 2- and 3-parameter families of quasi-periodic orbits. These families of orbits construct Cantor manifolds which are regarded as smooth manifolds, thanks to Whitney [168]. We explain the necessary concepts of smooth manifolds so we can appropriately talk about the computations of families of quasi-period orbits. The computation of a single quasi-periodic orbit involves solving a boundary value problem of an ordinary differential equation with shooting methods. The computation of a family of a one-parameter family of quasi-periodic orbits uses single-parameter continuation. Multiple-parameter families are compiled from many branches of solutions. Lastly, we discuss the ideas of optimization theory. All the ideas of Chapter 2 come together in Chapter 6, while most ideas are needed immediately beginning with Chapter 3.

Chapter 3 presents an already known method for continuing families of n -dimensional quasi-periodic invariant tori in great detail with emphasis on numerical implementation. This method was developed by Olikara and Scheeres with an emphasis on the theoretical development [140]. The purpose of presenting the method in the way that we did is to allow researchers to easily implement the method themselves, promoting the research of quasi-periodic invariant tori. We use this algorithm to study the usefulness of various parametric constraints and employ a novel parametric constraint. This parametric constraint generalizes the commonly used constant rotation number parametric constraint to maintain a particular direction in frequency space. Furthermore, we use a variety of parametric constraints to explore a 2- and 3-parameter family of quasi-periodic orbits in the circular restricted three-body problem. Motivated by NASA's Lunar Gateway, we computed the Earth-Moon L_2 quasi-halo orbit family. This family has been partially computed in prior work by Gómez and Mondelo in [75]. We explore a larger portion of the family, and we compute for the first time the 2- and 3-dimensional quasi-periodic invariant tori generated by the stable halo orbits. In astrodynamics, engineers are often concerned with the amplitude of orbits,

whether for mission constraints or for geometrical reference. We make a quantitative definition for the amplitudes of quasi-periodic orbits irrespective of any reference orbits, and develop an algorithm to compute the amplitudes. The algorithm is applicable to periodic orbits as well.

The analysis of the Earth-Moon L_2 quasi-halo orbit family goes beyond the work of Gómez and Mondelo to provide insight into the global orbital characteristics of the family relevant for mission design. We compute characteristics such as the Jacobi constant, the stability characteristics, perilune and apolune, amplitudes, and box size. We discover a stability bifurcation in the 2-parameter family of quasi-periodic orbits generated by the stable halo orbits. The quasi-halo orbits begin as elliptic and transition to partially elliptic just before the end of the computed family. We also examine the unstable manifolds across the quasi-halo orbits generated by the unstable halo orbits and observe three distinct behaviors. Moreover, we investigate transfers between the quasi-halo orbits and the lunar surface. Surface transfers have been explored in great detail for periodic orbits by Bury and McMahon in [43], though we show the extension to quasi-periodic orbits provides larger coverage of the surface.

Resonances between the internal frequencies of a torus destroy the quasi-periodic structure, reducing the dimension of the torus. Computing families of quasi-periodic orbits necessarily involves the difficulties encountered with resonances. Previous work develop ways to avoid resonances by introducing parametric constraints which guide the direction of the continuation, restricting the exploration of multi-parameter families of quasi-periodic orbits. We take a different approach to avoid resonances by choosing appropriate step sizes during the continuation. The step size computation is informed by the current frequencies of the torus and the change in frequencies according to the family tangent vector. Chapter 4 presents two methods to choose step sizes to avoid resonances. Both methods are built from the ideas of Sander and Meiss who define a measure of the irrationality of floating point number [152]. The first method chooses step sizes to predict torus frequencies away from resonant values. The second method tracks the frequencies during the continuation of tori to determine if the continuation is stuck near a resonance. The step size is greatly increased in an attempt to hop beyond resonances. The benefit of these methods is the

direction of the continuation is not restricted, allowing for freedom in exploring multi-parameter families of quasi-periodic orbits.

In Section 3.5.5 we saw the quasi-periodic orbits generated from the period-doubling halo orbit family are qualitatively different from the quasi-halo orbits, showing the period-doubling halo orbit family is different from the periodic orbits that the quasi-halo orbits collapse to at the 2:1 resonance. An example of a quasi-halo near the 3:1 resonance was shown in the bottom left plot of Figure 3.15. In this plot there are dark bands that give a preview of what periodic orbits on the 3:1 resonance line look like. It would also be of interest to study the 1-parameter families of periodic orbits at specific resonances which cut through the quasi-halo family. These resonant periodic orbits help structure the quasi-halo family and could provide insight into the overall structure. Moreover, these resonant periodic orbits could be key to stepping beyond resonances or even generate new families of quasi-periodic orbits.

The precise computation of quasi-periodic orbits is enabled by the retraction in Chapter 5 which leverages the implicit function theorem. The retraction modifies the continuation method which computes quasi-periodic orbits in two main ways. The first alteration is to leverage the newly developed constant slope parametric constraint from Chapter 3 to control the direction of the continuation. The second alteration keeps track of the distance traveled in frequency space and conducts a line search to target a specified distance. Together, these two alterations provide precise control over the mapping from the tangent space of a family of quasi-periodic orbits about a given orbit to the family of quasi-periodic orbits. The retraction can be used in a variety of settings wherein the family of orbits is an implicitly defined manifold, including periodic orbits. The retraction enables the development of other tools to precisely target quasi-periodic orbits - the focus of Chapter 6.

Chapter 6 imports tools from optimization, bringing forth novel capabilities for the design of orbits within multi-parameter families. Gradient-based optimization uses the gradient to determine search directions to minimize functions. The gradient of a function with respect to the orbit frequencies provides a direction in frequency space for a continuation method to compute or-

bits. The retraction is paired with a gradient descent method to optimize functions over the space of quasi-periodic orbits. The optimization problem is similar to a class of optimization problems called PDE-constrained optimization. We use the retraction to absorb equality constraints, requiring the use of manifold gradients, making the optimization problem formulation lie somewhere between PDE-constrained optimization and optimization on smooth manifolds. The modified gradient descent method is sufficient to solve unconstrained and equality constrained optimization problems. To solve inequality constrained problems, the modified gradient descent method is combined with an augmented Lagrangian method. We also develop a modified Newton's method to target quasi-periodic orbits with specified orbital characteristics. Similar to the optimization tools, the modified Newton's method takes derivatives with respect to the orbit frequencies to predict the frequencies of the orbit which will attain the desired characteristics. The retraction then computes the quasi-periodic orbit with the updated frequencies.

7.2 Avenues for Future Work

There are several interesting future investigations based on the work developed and presented in this thesis. To date, a full 3-parameter family of quasi-periodic orbits has not been computed. In this thesis we examined two, 2-parameter family slices of the 3-parameter family. A general purpose multi-parameter continuation method has been developed by Henderson [85], and successfully used by Henry and Scheeres to compute 3-parameter families of heteroclinic connections between families of quasi-periodic orbits [89]. We attempted to use the same multi-parameter continuation tool for the 3-parameter family of quasi-halo orbits with little success. It would be of interest to develop a multi-parameter continuation tool which leverages the dynamics and the subtleties of quasi-periodic invariant tori to study the 3-parameter family of quasi-halo orbits as a whole, and 2-parameter families with the same Jacobi energy. A tool of this sort would be advantageous as many families of quasi-periodic orbits in astrodynamics are 2- and 3-parameter families.

Chapter 3 shows the quasi-halo orbits are self-intersecting, family-intersecting, and intersect halo orbits. The intersections provide numerous opportunities for impulsive transfers to either

change phase or change orbits, however they also provide numerous opportunities for spacecraft to collide with each other. The opportunities for collision allude to the necessity of a framework to ensure spacecraft on these orbits do not collide. With this we also need to study the effects of uncertainty in the torus states on the resulting trajectories. Henry and Scheeres have started the framework for these analyses in [88] where they create expansion maps and find certain phase offsets where the spacecraft would crash into one another at a future point in time.

The resonance avoidance methods developed in this thesis perform well to mitigate the effects of resonances on the continuation of families of quasi-periodic orbits, however we believe there is further room for improvement to increase the freedom in exploring multi-parameter families of quasi-periodic orbits. In Chapter 4 we proposed a way to hop over resonances in a perpendicular sense to increase the likelihood of computing an orbit on the other side of the resonance. For general exploration of a family this is sufficient, but if the original branch of solutions is of interest, then hopping a resonance this way results in the computation of an orbit off the original branch. Algorithm 11 from Section 6.4 can be used to move back to the original branch to reinitialize the branch on the appropriate side of the resonance.

Algorithm 11 was tested on a 2-parameter family to find a quasi-halo orbit satisfying $\mathbf{g}(\mathbf{z}) = \mathbf{0}$ with $\mathbf{g} : \mathbb{R}^{\mathcal{D}+2} \rightarrow \mathbb{R}^t$ for the case $t = 1$ and $t = 2$. For $t = 1$ there were an infinite number of solutions, and for $t = 2$ there was a single, unique solution. It would be interesting to test the over constrained case $t > 2$ where there does not exist a solution \mathbf{z}^* such that $\mathbf{g}(\mathbf{z}^*) = \mathbf{0}$. We speculate the algorithm is will converge to a local minimum of \mathbf{g} so that for any $\mathbf{z} \in B_r(\mathbf{z}^*)$, $\|\mathbf{g}(\mathbf{z})\| \geq \|\mathbf{g}(\mathbf{z}^*)\|$.

In PDE-constrained optimization, a retraction-like function is typically put inside the cost function, equality constraints, and inequality constraints, so the optimizer is only aware of the parameters as the only variables in the problem. Putting the retraction inside the cost function and constraint functions allows for the use of off-the-shelf optimization packages at the expense of repeated computations due to the repeated, independent calls to the retraction inside the constraint functions. The optimization routines in this thesis were custom-made. This was done with

the intention of creating numerically efficient algorithms and avoid repeated computations at the expense of a less robust optimizer. It would be interesting to restructure the problem and the code to use an off-the-shelf optimizer such as IPOPT or NLOpt and compare the performance to the approach taken in this thesis. The better approach can be turned into an easy-to-use publicly available tool for mission designers to use.

Lastly, we mention all the tools developed or presented in this thesis can come together to create a comprehensive and versatile trajectory design tool to efficiently construct databases of quasi-periodic orbits, look up orbit families and specific orbits, and display families in various views. The idea complements the work of Folta et al., Cox et al., and Guzzetti et al. which have created a similar tool for rapid trajectory design in the circular restricted three-body problem [52, 67, 68, 79]. The work in this thesis provides the capabilities to compute higher-dimensional quasi-periodic invariant tori in autonomous and non-autonomous dynamical systems, provides different views to study the families, enables the precise computation of specific orbits given various constraints. Combining their pre-existing trajectory design tool with the tools in this thesis will result in a stronger mission design tool.

Bibliography

- [1] R. A. Adomaitis. The trouble with spurious eigenvalues. International Journal of Bifurcation and Chaos, 17(4):1375–1381, 2007.
- [2] Z. Ahsan, H. Dankowicz, and J. Sieber. Optimization along families of periodic and quasiperiodic orbits in dynamical systems with delay. Nonlinear Dynamics, 99, 2020.
- [3] G. Allaire, F. Jouve, and A. M. Toader. A level-set method for shape optimization. Comptes Rendus Mathématique, 334(12):1125–1130, 2002.
- [4] I. Amidror. Mastering the Discrete Fourier Transform in One, Two or Several Dimensions: Pitfalls and Artifacts. Springer Publishing Company, Incorporated, 1st edition, 2015.
- [5] European Space Agency. Athena assessment study report, 2011.
- [6] European Space Agency. Ariel assessment study report, 2017.
- [7] European Space Agency. Plato definition study report, 2017.
- [8] H. Antil, M. Heinkenschloss, R. H. W. Hoppe, C. Linsenmann, and A. Wixforth. Reduced order modeling based shape optimization of surface acoustic wave driven microfluidic biochips. Mathematics and Computers in Simulation, 82(10):1986–2003, 2012.
- [9] H. Antil and D. Leykekhman. Frontiers in PDE-Constrained Optimization. Springer New York, 2018.
- [10] H. Antil, R. H. Nochetto, and P. Venegas. Optimizing the kelvin force in a moving target subdomain. Mathematical Models and Methods in Applied Sciences, 28(01):95–130, 2018.
- [11] V. I. Arnold. Mathematical Methods of Classical Mechanics. Springer New York, NY, 2 edition, 1989.
- [12] V. I. Arnold, V. V. Kozlov, and A. I. Neishtadt. Mathematical Aspects of Classical and Celestial Mechanics. Springer Berlin, Heidelberg, 2 edition, 1997.
- [13] U. M. Ascher, R. M. M. Mattheij, and R. D. Russell. Numerical Solution of Boundary Value Problems for Ordinary Differential Equations. Society for Industrial and Applied Mathematics, 1995.
- [14] B. T. Barden and K. C. Howell. Fundamental motions near collinear libration points and their transitions. The Journal of the Astronautical Sciences, 46(4):361–378, 1998.

- [15] N. Baresi. Spacecraft Formation Flight on Quasi-Periodic Invariant Tori. Phd thesis, University of Colorado Boulder, 2017.
- [16] N. Baresi and L. Dell’Elce. Periodic and quasi-periodic orbits near close planetary moons. Journal of Guidance, Control, and Dynamics, 46(4):680–694, 2023.
- [17] N. Baresi, X. Fu, and R. Armellin. A high-order taylor polynomials approach for continuing trajectories in restricted three-body problems. In AAS/AIAA Astrodynamics Specialist Conference, South Lake Tahoe, California, 2020.
- [18] N. Baresi, Z. P. Olikara, and D. J. Scheeres. Fully numerical methods for continuing families of quasi-periodic invariant tori in astrodynamics. Journal of Astronautical Science, 65(2):157–182, 2018.
- [19] N. Baresi and D. J. Scheeres. Quasi-periodic invariant tori of time-periodic dynamical systems: Applications to small body exploration. In 67th International Astronautical Congress, Guadalajara, Mexico, 2016.
- [20] N. Baresi and D. J. Scheeres. Bounded relative motion under zonal harmonics perturbations. Celestial Mechanics and Dynamical Astronomy, 127(4):527–548, 2017.
- [21] N. Baresi and D. J. Scheeres. Design of bounded relative trajectories in the earth zonal problem. Journal of Guidance, Control, and Dynamics, 40(12):3075–3087, 2017.
- [22] N. Baresi, D. J. Scheeres, and H. Schaub. Bounded relative orbits about asteroids for formation flying and applications. Acta Astronautica, 123:364–375, 2016.
- [23] N. Baresi, D. A. Dei Tos, H. Ikeda, and Y. Kawakatsu. Orbit design and maintenance in the elliptical hill problem with applications to the phobos sample return mission mmx. In 70th International Astronautical Congress, Washington DC, United States, 2019.
- [24] R. H. Battin. An Introduction to the Mathematics and Methods of Astrodynamics. American Institute of Aeronautics and Astronautics, Inc., 1999.
- [25] A. Beck. Introduction to Nonlinear Optimization. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, 2014.
- [26] F. J. Beron-Vera, M. J. Olascoaga, M. G. Brown, H. Koçak, and I. I. Rypina. Invariant-tori-like Lagrangian coherent structures in geophysical flows. Chaos: An Interdisciplinary Journal of Nonlinear Science, 20(1), 01 2010.
- [27] E. G. Birgin and J. M. Martínez. Practical Augmented Lagrangian Methods for Constrained Optimization. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2014.
- [28] S. Bonasera and N. Bosanac. Computing natural transitions between tori near resonances in the earth–moon system. Journal of Guidance, Control, and Dynamics, 46(3):443–454, 2023.
- [29] N. Bosanac. Bounded motions near resonant orbits in the earth-moon and sun-earth systems. In AAS/AIAA Astrodynamics Specialist Conference, Snowbird, UT, August 2018.
- [30] K. Boudad, K. C. Howell, and D. C. Davis. Heliocentric access from cislunar space within the context of the bicircular restricted four-body problem. In AAS/AIAA Astrodynamics Specialist Conference, Lake Tahoe, CA, August 2020.

- [31] K. Boudad, K. C. Howell, and D. C. Davis. Energy and phasing considerations for low-energy transfers from cislunar to heliocentric space. In 31st AAS/AIAA Space Flight Mechanics Meeting, Virtual, February 2021.
- [32] K. K. Boudad, K. C. Howell, and D. C. Davis. Departure and escape dynamics from the near rectilinear halo orbits in the earth-moon-sun system. The Journal of the Astronautical Sciences, 69(4), 2022.
- [33] N. Boumal. An introduction to optimization on smooth manifolds. Cambridge University Press, 2023.
- [34] N. Boumal, P. A. Absil, and C. Cartis. Global rates of convergence for nonconvex optimization on manifolds. IMA Journal of Numerical Analysis, 39(1):1–33, feb 2018.
- [35] S. Boyd and L. Vandenberghe. Convex Optimization. Cambridge University Press, 2004.
- [36] N. Bradley, Z. P. Olikara, S. Bhaskaran, and B. Young. Cislunar navigation accuracy using optical observations of natural and artificial targets. Journal of Spacecraft and Rockets, 57(4), 2020.
- [37] J. V. Breakwell and J. V. Brown. The 'halo' family of 3-dimensional periodic orbits in the earth-moon restricted 3-body problem. Celestial Mechanics, 20(4):389–404, 1979.
- [38] H. W. Broer, B. Hasselblatt, and F. Takens, editors. Handbook of Dynamical Systems, volume 3. Elsevier, 2010.
- [39] H. W. Broer, G. B. Huitema, and M. B. Sevryuk. Quasi-Periodic Motions in Families of Dynamical Systems. Springer-Verlag Berlin Heidelberg, 1996.
- [40] S. Broschart, M-K. J. Chung, S. J. Hatch, J. H. Ma, T. H. Sweetser, S. S. Weinstein-Weiss, and V. Angelopoulos. Preliminary trajectory design for the artemis lunar mission. In AAS/AIAA Astrodynamics Specialist Conference, Pittsburg, Pennsylvania, August 9-13, 2009.
- [41] E. R. Burnette. Novel Dynamics and Control Formulations for Multi-Spacecraft Formation Flying, Rendezvous, and Proximity Operations. Phd thesis, University of Colorado Boulder, 2021.
- [42] L. Bury. Low-Energy Secondary-Body Landing Trajectories in the Three-Body Problem. Phd thesis, University of Colorado Boulder, 2021.
- [43] L. Bury and J. McMahon. Landing trajectories to moons from the unstable invariant manifolds of periodic libration point orbits. In AAS/AIAA Astrodynamics Specialist Conference, Lake Tahoe, CA (Virtual), 2020.
- [44] X. Cabré, E. Fontich, and R. de la Llave. The parameterization method for invariant manifolds i: Manifolds associated to non-resonant subspaces. Indiana University Mathematics Journal, 52:283–328, 2003.
- [45] X. Cabré, E. Fontich, and R. de la Llave. The parameterization method for invariant manifolds ii: Regularity with respect to parameters. Indiana University Mathematics Journal, 52(2):329–360, 2003.

- [46] M. Canadell and À. Haro. Computation of quasi-periodic normally hyperbolic invariant tori: Algorithms, numerical explorations and mechanisms of breakdown. Journal of Nonlinear Science, 27(6), 2017.
- [47] E. Castellà and À. Jorba. On the vertical families of two-dimensional tori near the triangular points of the bicircular problem. Celestial Mechanics and Dynamical Astronomy, 76:35–54, 2000.
- [48] A. Celletti and A. Giorgilli. On the stability of the lagrangian points in the spatial restricted problem of three bodies. Celestial Mechanics and Dynamical Astronomy, 50:31–58, 1991.
- [49] A. Celletti, G. Pucacco, and D. Stella. Lissajous and halo orbits in the restricted three-body problem. Journal of Nonlinear Science, 25, 2015.
- [50] S. P. Chakrabarty and F. B. Hanson. Optimal control of drug delivery to brain tumors for a distributed parameters model. In Proceedings of the 2005, American Control Conference, 2005., pages 973–978 vol. 2, 2005.
- [51] C. Chavaudret. Reducibility of quasiperiodic cocycles in linear lie groups. Ergodic Theory and Dynamical Systems, 31(3):741–769, 2011.
- [52] A. D. Cox, N. Bosanac, D. Guzzetti, K. C. Howell, and D. C. Folta. An interactive trajectory design environment leveraging dynamical structures in multi-body regimes. 2016.
- [53] G. D’Avino, S. Crescitelli, P. L. Maffettone, and M. Grosso. On the choice of the optimal periodic operation for a continuous fermentation process. Biotechnology Progress, 26(6):1580–1589, 2010.
- [54] D. C. Davis, K. K. Boudad, S. M. Phillips, and K. C. Howell. Disposal, deployment, and debris in near rectilinear halo orbits. In 29th AAS/AIAA Space Flight Mechanics Meeting, Ka’anapali, Maui, Hawaii, January 2019.
- [55] D. C. Davis, E. M. Zimovan-Spreen, R. J. Power, and K. C. Howell. Cubesat deployment from a near rectilinear halo orbit. In AIAA SCITECH 2022 Forum, San Diego, CA, January 2022.
- [56] R. de la Llave. A tutorial on kam theory. 2003.
- [57] R. de la Llave and J. D. Mireles James. Parameterization of invariant manifolds by reducibility for volume preserving and symplectic maps. Discrete and Continuous Dynamical Systems, 32(12):4321–4360, 2012.
- [58] J. Carlos de los Reyes and C. B. Schönlieb. Image denoising: Learning the noise model via nonsmooth pde-constrained optimization. Inverse Problems and Imaging, 7:1183–1214, 2013.
- [59] D. W. Dreisigmeyer. Equality constraints, riemannian manifolds and direct search methods. 2006.
- [60] H. Aguiar e Oliveira. Evolutionary Global Optimization, Manifolds and Applications. Springer Publishing Company, Incorporated, 1st edition, 2015.
- [61] L. H. Eliason. Reducibility and point spectrum for linear quasi-periodic skew-products. In International Congress of Mathematics, pages 779–787, Berlin, Germany, 1998.

- [62] L. H. Eliasson. Reducibility for linear quasi-periodic differential equations. Winter School, St Etienne de Tinée, February 7-11, 2011.
- [63] R. Farquhar and A. A. Kamel. Quasi-periodic orbits about the translunar libration point. Celestial Mechanics, 7(4):458–473, 1973.
- [64] Robert W. Farquhar. Station-keeping in the vicinity of collinear libration points with an application to a lunar communications problem. 1966.
- [65] Robert W. Farquhar. The utilization of halo orbits in advanced lunar operations. 1971.
- [66] J. Figueras and A. Haro. Different scenarios for hyperbolicity breakdowns in quasiperiodic area preserving twist maps. Chaos, 2015.
- [67] D. C. Folta, N. Bosanac, D. Guzzetti, and K. C. Howell. An earth-moon system trajectory design reference catalog. Acta Astronautica, 110, 2015.
- [68] D. C. Folta, C. M. Webster, N. Bosanac, A. D. Cox, D. Guzzetti, and K. C. Howell. Trajectory design tools for libration and cis-lunar environments. 2016.
- [69] F. Gabern and À. Jorba. Effective computation of the dynamics around a two-dimensional torus of a hamiltonian system. Journal of Nonlinear Science, 15:159–182, 2005.
- [70] S. Garreis, T. M. Surowiec, and M. Ulbrich. An interior-point approach for solving risk-averse pde-constrained optimization problems with coherent risk measures. SIAM Journal on Optimization, 31(1):1–29, 2021.
- [71] J. Gimeno, À. Jorba, B. Nicolás, and E. Olmedo. Numerical computation of high-order expansions of invariant manifolds of high-dimensional tori. SIAM Journal of Applied Dynamical Systems, 21(3):1832–1861, 2022.
- [72] J. Gimeno, À. Jorba, M. Jorba-Cuscó, N. Miguel, and M. Zou. Numerical integration of high-order variational equations of odes. Applied Mathematics and Computation, 442:127743, 2023.
- [73] G. Gómez, M. Marcote, and J. M. Mondelo. The invariant manifold structure of the spatial hill’s problem. Dynamical Systems, 20(1):115–147, 2005.
- [74] G. Gómez, J. Masedemont, and C. Simó. Quasihalo orbits associated with libration points. Journal of Astronautical Science, 46(2):135–176, 1998.
- [75] G. Gómez and J. M. Mondelo. The dynamics around the collinear equilibrium points of the rtbp. Physica D: Nonlinear Phenomena, 157(4):283–321, 2001.
- [76] J. A. Greaves and D. J. Scheeres. Observation and maneuver detection for cislunar vehicles. Journal of Astronautical Sciences, 68:826–854, 2021.
- [77] J. M. Greene. A method for determining a stochastic transition. Journal of Mathematical Physics, 20(6):1183–1201, 07 1979.
- [78] A. Griewank and A. Walther. Evaluating Derivatives: Principles and Techniques of Algorithmic Differentiation. Society for Industrial and Applied Mathematics, USA, 2 edition, 2008.

- [79] D. Guzzetti, N. Bosanac, A. Haapala, K. C. Howell, and D. C. Folta. Rapid trajectory design in the earth-moon ephemeris system via an interactive catalog of periodic and quasi-periodic orbits. Acta Astronautica, 126:439–455, 2016.
- [80] H. Hanßmann. Non-degeneracy conditions in kam theory. Indagationes Mathematicae, 22(3):241–256, 2011. Devoted to: Floris Takens (1940–2010).
- [81] À. Haro and R. de la Llave. Manifolds on the verge of a hyperbolicity breakdown. Chaos: An Interdisciplinary Journal of Nonlinear Science, 16:013120, November 2005.
- [82] À. Haro and R. de la Llave. A parameterization method for the computation of invariant tori and their whiskers in quasi-periodic maps: explorations and mechanisms for the breakdown of hyperbolicity. SIAM Journal on Applied Dynamical Systems, 6(1):142–207, 2007.
- [83] À. Haro and J. M. Mondelo. Flow map parameterization methods for invariant tori in hamiltonian systems. Communications in Nonlinear Science and Numerical Simulation, 101:105859, oct 2021.
- [84] À. Haro, M. Canadell, and J-M. Mondelo J-L. Figueras, A. Luque. The Parameterization Method for Invariant Manifolds. Springer Cham, 2016.
- [85] M. E. Henderson. Multiple parameter continuation: Computing implicitly defined k -manifolds. International Journal of Bifurcation and Chaos, 12(3):451–476, 2002.
- [86] M. Hénon. Generating Families in the Restricted Three-Body Problem. Springer Berlin, Heidelberg, 1997.
- [87] J. Henrard. Lyapunov’s center theorem for resonant equilibrium. Journal of Differential Equations, 14(3):431–441, 1973.
- [88] D. B. Henry and D. J. Scheeres. Expansion maps: Designing relative trajectories on quasi-periodic orbits. Journal of Guidance, Control, and Dynamics, 44(3):457–468, 2021.
- [89] D. B. Henry and D. J. Scheeres. A survey of heteroclinic connections in the earth-moon system. In 73rd International Astronautical Congress, Paris, France, 2022.
- [90] D. B. Henry and D. J. Scheeres. Quasi-periodic orbit transfer design via whisker intersection sets. Journal of Guidance, Control, and Dynamics, 2023.
- [91] D. B. Henry and D. J. Scheeres. Formation control on frequency matched quasi-periodic orbits. In International Workshop on Satellite Constellations and Formation Flying, Milan, Italy, June 2022.
- [92] D. B. Henry and D. J. Scheeres. Impulsive spacecraft formation control on quasi-periodic orbits. In 2022 American Control Conference, Atlanta, Georgia, June 8-10, 2022.
- [93] S. Hosseini, W. Huang, and R. Yousefpour. Line search algorithms for locally lipschitz functions on riemannian manifolds. SIAM Journal on Optimization, 28(1):596–619, 2018.
- [94] K. C. Howell. Three-dimensional, periodic, ‘halo’ orbits. Celestial Mechanics, 32(1):53–71, 1984.

- [95] K. C. Howell, B. T. Barden, and M. W. Lo. Application of dynamical systems theory to trajectory design for a libration point mission. The Journal of Astronautical Sciences, 45(2), 1997.
- [96] K. C. Howell and J. V. Breakwell. Almost rectilinear halo orbits. Celestial Mechanics, 32(1):29–52, 1984.
- [97] K. C. Howell and H. J. Pernicka. Numerical determination of lissajous trajectories in the restricted three-body problem. Celestial Mechanics, 41(1-4):107–124, 1987.
- [98] W. Huang. Optimization Algorithms on Riemannian Manifolds with Applications. PhD thesis, Florida State University, 2013.
- [99] Antony Jameson. Aerodynamic shape optimization using the adjoint method. 2003.
- [100] R. A. Johnson and G. R. Sell. Smoothness of spectral subbundles and reducibility of quasi-periodic linear differential equations. Journal of Differential Equations, 41(2):262–288, 1981.
- [101] À. Jorba. ON PRACTICAL STABILITY REGIONS FOR THE MOTION OF A SMALL PARTICLE CLOSE TO THE EQUILATERAL POINTS OF THE REAL EARTH-MOON SYSTEM, pages 197–213. 2000.
- [102] À. Jorba. Numerical computation of the normal behavior of invariant curves of n -dimensional maps. Nonlinearity, 14(5):943–976, 2001.
- [103] À. Jorba and J. Masdemont. Dynamics in the center manifold of the collinear points of the restricted three body problem. Physica D: Nonlinear Phenomena, 132(1-2):189–213, 1999.
- [104] À. Jorba and E. Olmedo. On the computation of reducible invariant tori on a parallel computer. SIAM Journal on Applied Dynamical Systems, 8(4):1382–1404, 2009.
- [105] À. Jorba and C. Simó. On quasi-periodic perturbations of elliptic equilibrium points. SIAM Journal on Mathematical Analysis, 27(6):1704–1737, 1996.
- [106] À. Jorba and C. Simó. On the reducibility of linear differential equations with quasiperiodic coefficients. Journal of Differential Equations, 98(1):111–124, 1992.
- [107] À. Jorba and J. Villanueva. On the normal behavior of partially elliptic lower-dimensional tori of hamiltonian systems. Nonlinearity, 10, 1997.
- [108] À. Jorba and J. Villanueva. On the persistence of lower dimensional invariant tori under quasi-periodic perturbations. Journal of Nonlinear Science, 7(1):427–473, 1997.
- [109] À. Jorba and J. Villanueva. Numerical computation of normal forms around some periodic orbits of the restricted three-body problem. Physica D: Nonlinear Phenomena, 114(3):197–229, 1998.
- [110] À. Jorba and J. Villanueva. The fine geometry of the cantor families of invariant tori in hamiltonian systems. In European Congress of Mathematics, pages 557–564, Basel, 2001. Birkhäuser Basel.
- [111] A. Katok and B. Hasselblatt. Introduction to the Modern Theory of Dynamical Systems. Cambridge University Press, 1995.

- [112] J. P. Kernévez and E. J. Doedel. Optimization in Bifurcation Problems using a Continuation Method, pages 153–160. Birkhäuser Basel, Basel, 1987.
- [113] E. Koleman, N. J. Kasdin, and P. Gurfil. Multiple poincare sections method for finding quasi-periodic orbits of the restricted three body problem. Celestial Mechanics, 112(1):47–74, 2012.
- [114] S. Kuksin and J. Pöschel. Invariant cantor manifolds of quasi-periodic oscillations for a nonlinear schrödinger equation. Annal of Mathematics, 143(1), 1996.
- [115] B. Kumar, R. L. Anderson, and R. de la Llave. Rapid and accurate methods for computing whiskered tori and their manifolds in periodically perturbed planar circular restricted 3-body problems. Celestial Mechanics and Dynamical Astronomy, 134(3), 2022.
- [116] J. M. Lee. Introduction to Smooth Manifolds. Springer New York, NY, 2012.
- [117] M. Li and H. Dankowicz. Staged construction of adjoints for constrained optimization of integro-differential boundary-value problems. SIAM Journal on Applied Dynamical Systems, 17(2):1117–1151, 2018.
- [118] C. Liu and N. Boumal. Simple algorithms for optimization on riemannian manifolds with constraints. Applied Mathematics & Optimization, 82(3), 2020.
- [119] M. W. Lo, B. G. Williams, W. E. Bollman, D. Han, Y. Hahn, J. L. Bell, E. A. Hirst, R. A. Corwin, P. E. Hong, K. C. Howell, B. Barden, and R. Wilson. Genesis mission design. The Journal of Astronautical Sciences, 49(1), 2001.
- [120] D. G. Luenberger and Y. Ye. Linear and Nonlinear Programming. Springer Cham, 5 edition, 2021.
- [121] D. Lujan and D. J. Scheeres. Earth-moon L_2 quasi-halo orbit family: Characteristics and manifold applications. Journal of Guidance, Control, and Dynamics, 45(11):2029–2045, 2022.
- [122] D. Lujan and D. J. Scheeres. Dynamics in the vicinity of the stable halo orbits. Journal of the Astronautical Sciences, 70(20), 2023.
- [123] D. Lujan and D. J. Scheeres. Frequency structure of the nrho family in the earth-moon system. In 2019 AAS/AIAA Astrodynamics Specialist Conference, Portland, Maine, August 11-15, 2019. AAS.
- [124] D. Lujan and D. J. Scheeres. Global L_2 quasi-halo family and their characteristics. In 31st AAS/AIAA Spaceflight Mechanics Meeting, Virtual, February 1-3, 2021. AAS.
- [125] D. Lujan and D. J. Scheeres. Method to target quasi-periodic orbit frequencies within multi-parameter families. In 33rd AAS/AIAA Space Flight Mechanics Meeting, Austin, Texas, January 15-19, 2023. AAS.
- [126] D. Lujan and D. J. Scheeres. The earth-moon L_2 quasi-halo orbit family: Characteristics and manifold applications. In 32nd AIAA/AAS Spaceflight Mechanics Meeting, San Diego, California, January 3-7, 2022. AIAA.
- [127] D. Lujan and D. J. Scheeres. Dynamics in the vicinity of the stable halo orbits. In 73rd International Astronautical Congress, Paris, France, September 18-22, 2022.

- [128] R. S. MacKay. A renormalization approach to invariant circles in area-preserving maps. Physica D: Nonlinear Phenomena, 7(1):283–300, 1983.
- [129] C. Marchal. The Three-Body Problem. Elsevier Science, 1990.
- [130] B. P. McCarthy and K. C. Howell. Leveraging quasi-periodic orbits for trajectory design in cislunar space. Astrodynamics, 5(2):139–165, 2021.
- [131] B. P. McCarthy and K. C. Howell. Four-body cislunar quasi-periodic orbits and their application to ballistic lunar transfer design. Advances in Space Research, 71(1):556–584, 2023.
- [132] D. J. McComas, E. R. Christian, N. A. Schwadron, N. Fox, J. Westlake, F. Allegrini, D. N. Baker, D. Biesecker, M. Bzowski, G. Clark, and et al. Interstellar mapping and acceleration probe (imap): A new nasa mission. Space Science Reviews, 214(8), 2018.
- [133] J. D. Meiss. Differential Dynamical Systems. Society of Industrial and Applied Mathematics, Philadelphia, PA, 2017.
- [134] K. R. Meyer and D. C. Offin. Introduction to Hamiltonian Dynamical Systems and the N-Body Problem. Springer, Cham, 3 edition, 2017.
- [135] W. Ming, C. Yang, and H. Zhang. Family of resonant quasi-periodic distant retrograde orbits in cislunar space. In The 28th International Symposium on Space Flight Dynamics, Beijing, China, September 2022.
- [136] J. M. Mondelo, E. Barrabés, G. Gómez, and M. Ollé. Fast numerical computation of lissajous and quasi-halo libration point trajectories and their invariant manifolds. In 63rd International Astronautical Congress, Naples, Italy, 2012.
- [137] J. Ng and S. Dubljevic. Optimal boundary control of a diffusion–convection–reaction pde model with time-dependent spatial domain: Czochralski crystal growth process. Chemical Engineering Science, 67(1):111–119, 2012. Dynamics, Control and Optimization of Energy Systems.
- [138] J. Nocedal and S. J. Wright. Numerical Optimization. Springer New York, NY, 2 edition, 2006.
- [139] Z. P. Olikara. Computation of Quasi-Periodic Tori and Heteroclinic Connections in Astrodynamics Using Collocation Techniques. Phd thesis, University of Colorado Boulder, 2016.
- [140] Z. P. Olikara and D. J. Scheeres. Numerical method for computing quasi-periodic orbits and their stability in the restricted three-body problem. Advances in Astronautical Science, 145:911–930, 01 2012.
- [141] N. L. Parrish, M. J. Bollinger, E. Kayser, M. R. Thomspson, J. S. Parker, B. W. Cheetham, D. C. Davis, and D. J. Sweeney. Near rectilinear halo orbit determination with simulated dsn observations. In 30th AIAA/AAS Space Flight Mechanics Meeting, Orlando, FL, January 2020. AIAA.
- [142] L. T. Peterson and D. J. Scheeres. Defining the fundamental frequencies of quasi-periodic invariant tori. In AAS/AIAA Spaceflight Mechanics Meeting, Virtual, February 1-3, 2021. AAS.

- [143] J. Pöschel. Integrability of hamiltonian systems on cantor sets. Communications in Pure and Applied Mathematics, 35(5):653–696, 1982.
- [144] N. Queipo, R. Devarakonda, and J. A. C. Humphrey. Genetic algorithms for thermosciences research: application to the optimized cooling of electronic components. International Journal of Heat and Mass Transfer, 37(6):893–908, 1994.
- [145] K. R. Rao, D. N. Kim, and J. J. Hwang. Fast Fourier Transform - Algorithms and Applications. Springer Dordrecht, 2010.
- [146] D. L. Richardson. Analytic construction of periodic orbits about the collinear points. Celestial Mechanics, 22, 1980.
- [147] D. L. Richardson. A note on a lagrangian formulation for motion about the collinear points. Celestial Mechanics, 22, 1980.
- [148] J. J. Rosales, À. Jorba, and M. Jorba-Cuscó. Families of halo-like invariant tori around L_2 in the earth-moon bicircular problem. Celestial Mechanics and Dynamical Astronomy, 133(16), 2021.
- [149] J. J. Rosales, À. Jorba, and M. Jorba-Cuscó. Invariant manifolds near L_1 and L_2 in the quasi-bicircular problem. Celestial Mechanics and Dynamical Astronomy, 135(2), 2023.
- [150] A. E. Roy. The Foundations of Astrodynamics. CRC Press, New York, 2005.
- [151] R. J. Sacker and G. R. Sell. A spectral theory for linear differential equations. Journal of Differential Equations, 27(3):320–358, 1978.
- [152] E. Sander and J. D. Meiss. Birkhoff averages and rotational invariant circles for area-preserving maps. Physica D, 411:132569, 2020.
- [153] H. Schaub and J. L. Junkins. Analytical Mechanics of Space Systems. American Institute of Aeronautics and Astronautics, Inc., 4 edition, 2018.
- [154] D. J. Scheeres. Orbital Motion in Strongly Perturbed Environments: Applications to Asteroid, Comet and Planetary Satellite Orbiters. Springer, Heidelberg, 2012.
- [155] F. Schilder, H. M. Osinga, and W. Vogt. Continuation of quasi-periodic invariant tori. Journal of Applied Dynamical Systems, 4(3):459–488, 2005.
- [156] G. R. Sell. Hyperbolic almost periodic solutions and toroidal limit sets. Proc Nat Acad Sci, USA, 74, 1977.
- [157] G. R. Sell. The structure of a flow in the vicinity of an almost periodic motion. Journal of Differential Equations, 27(3):359–393, 1978.
- [158] R. Seydel. Practical Bifurcation and Stability Analysis. Springer, New York, NY, 2010.
- [159] C. L. Siegel and J. K. Moser. Lectures on Celestial Mechanics. Springer Berlin, Heidelberg, 1995.
- [160] M. R. Somayaji, M. Xenos, L. Zhang, M. Mekarski, and A. A. Linninger. Systematic design of drug delivery therapies. Computers & Chemical Engineering, 32(1):89–98, 2008. Process Systems Engineering: Contributions on the State-of-the-Art.

- [161] V. Szebehely. Theory of Orbits: The Restricted Problem of Three Bodies. Academic Press, New York and London, 1967.
- [162] Joan Pau Sánchez, David Morante, Pablo Herмосin, Daniel Ranuschio, Alvaro Estalella, Dayana Viera, Simone Centuori, Geraint Jones, Colin Snodgrass, Anny Chantal Levasseur-Regourd, and et al. Esa f-class comet interceptor: Trajectory design to intercept a yet-to-be-discovered comet. Acta Astronautica, 188:265–277, 2021.
- [163] M. R. Thompson, A. J. Zara, C. Ott, M. Bolliger, E. Kayser, and D. C. Davis. Comparisons of filtering algorithms for orbit determination in near rectilinear halo orbits. In 2022 Astrodynamics Specialist Conference, 2022.
- [164] J. O. Toilliez and A. J. Szeri. Optimized translation of microbubbles driven by acoustic fields. The Journal of the Acoustical Society of America, 123(4):1916–1930, 04 2008.
- [165] D. A. Vallado. Fundamentals of Astrodynamics and Applications. Springer New York, NY, 4 edition, 2013.
- [166] D. Villegas-Pinto, N. Baresi, S. Locoche, and D. Hestroffer. Resonant quasi-periodic near-rectilinear halo orbits in the elliptic-circular earth-moon-sun problem. Advances in Space Research, 71(1):336–354, 2023.
- [167] R. J. Whitley and R. M. Martinez. Options for staging orbits in cislunar space. In 2016 IEEE Aerospace Conference, Big Sky, Montana, March 5-12, 2016. IEEE.
- [168] H. Whitney. Analytic extensions of differentiable functions defined in closed sets. Transactions of the American Mathematical Society, 36(1):63–89, 1934.
- [169] J. Williams, D. E. Lee, R. J. Whitley, K.A. Bokelmann, D. C. Davis, and C. F. Berry. Targeting cislunar near rectilinear halo orbits for human space exploration. In NASA Center for AeroSpace Information, Hampton, Virginia, February 5, 2017. NASA/Langley Research Center.
- [170] M. Wyczalkowski and A. J. Szeri. Optimization of acoustic scattering from dual-frequency driven microbubbles at the difference frequency. The Journal of the Acoustical Society of America, 113(6):3073–3079, 05 2003.
- [171] D. Wysham. Reducibility, Manifolds, and Bifurcations of Invariant Tori in Dynamical Systems. PhD thesis, University of Colorado, 2000.
- [172] Y. Yao and R. de la Llave. Computing the invariant circle and the foliation by stable manifolds for a 2-d map by the parameterization method: Numerical implementation and results, 2021.
- [173] S. Yun, K. Tuggle, R. Zanetti, and C. D’Souza. Sensor configuration trade study for navigation in near rectilinear halo orbits. Journal of Astronautical Sciences, 67:1755–1774, 2020.
- [174] E. M. Zimovan-Spreen. Dynamical structures nearby nrhos with applications to transfer design in cislunar space. The Journal of the Astronautical Sciences, 69(3), 2022.

Appendix A

Publicly Available Code

Many of the functions created during the work on this thesis are publicly available on Github at <https://github.com/dlujan17/QPOs>. At this location one will find functions to compute n -dimensional quasi-periodic invariant tori and various other routines mentioned throughout this thesis. The work of Chapter 3 was done in Matlab, while the work of Chapters 4, 5, and 6 was done in Julia. A Python version of the function which continues families of 2-dimensional quasi-periodic invariant tori was created to satisfy obligations for a grant, though we never worked in Python. Therefore, the most developed codebase is in Julia. The functions used to complete the work in Chapters 5 and 6 are available upon request by emailing david.lujan@colorado.edu.