

UNIVERSITY OF OKLAHOMA

GRADUATE COLLEGE

DYNAMICAL STRUCTURES AND TRAJECTORY FAMILIES IN THE J_2 —EXTENDED
NEPTUNE—TRITON SYSTEM

A THESIS

SUBMITTED TO THE GRADUATE FACULTY

in partial fulfillment of the requirements for the

Degree of

MASTER OF SCIENCE

By BLAKE T JOHNSON

Norman, Oklahoma

2026

DYNAMICAL STRUCTURES AND TRAJECTORY FAMILIES IN THE J_2 —EXTENDED
NEPTUNE—TRITON SYSTEM

A THESIS APPROVED FOR THE
SCHOOL OF AEROSPACE AND MECHANICAL ENGINEERING

BY THE COMMITTEE CONSISTING OF

Dr. Diogo Sanchez, Chair

Dr. Jingyao Dai

Dr. Yanqing Su

Acknowledgements

This work would not have been possible without the support of those who guide me, sustain me, and inspire me along the way.

To my son, Henry — you are my North Star. Fixed, unwavering, and always there to orient me when I lose my way. You give me a reason to push beyond what I ever imagined I was capable of, and every page of this work carries your name in spirit.

To my wife, Kate — you are my strength, my foundation, and my constant companion on this journey to the stars. You support my dreams before I fully believe in them myself, and you carry more than your share so that I can reach further. None of this happens without you.

To my mother, Cheryl — you teach me how to dream. You nurture my imagination and show me that the stars are not something to admire from a distance but something worth reaching for. The curiosity that drives this work begins with you.

To my father, Michael — you teach me how to work. You show me that reaching the stars is not only about dreaming but about the discipline, persistence, and effort required to get there. Your example is in every late night this thesis demands of me.

To my advisor, Dr. Diogo Sanchez — thank you for three years of guidance, patience, and encouragement. You teach me not only how to navigate the mathematics of multi-body dynamics, but how to maneuver among the stars with confidence and purpose. This work is a reflection of your mentorship.

The computing for this project was performed at the OU Supercomputing Center for Education and Research (OSCER) at the University of Oklahoma (OU).

“It may happen that small differences in the initial conditions produce very great ones in the final phenomena; a small error in the former will produce an enormous error in the latter. Prediction becomes impossible, and we have the fortuitous phenomenon.”

Henri Poincaré
Science and Method (1908)

Contents

Acknowledgements	iv
Abstract	xix
1 Introduction	1
1.1 Motivation	1
1.2 Dynamical Significance of the Neptune–Triton System	2
1.3 Problem Statement	2
1.4 Thesis Contributions	3
1.5 Scope and Assumptions	4
1.6 Thesis Organization	5
2 Dynamical Background and Analysis Methods	7
2.1 The Circular Restricted Three-Body Problem	7
2.1.1 Rotating Frame and Normalized Coordinates	8
2.2 Effective Potential of the CR3BP	9
2.3 Equations of Motion	11
2.4 The Jacobi Integral and Jacobi Constant	12
2.5 Zero-Velocity Surfaces and Hill’s Regions	14
2.6 Lagrange Points	16
2.6.1 Linear Stability of the Lagrange Points	18
2.7 Poincaré Maps	19
2.7.1 Motivation and Dimensional Reduction	20
2.7.2 Definition of the Surface of Section	21
2.7.3 Definition of the Poincaré Map	21
2.7.4 Geometric Interpretation of Map Structures	22
2.7.5 Why the Poincaré Map is Important in the CR3BP	25
2.8 Frequency Analysis of Trajectory Dynamics	26
2.8.1 Time Series Representation of Motion	27
2.8.2 Fourier Representation of the Time Series	28
2.8.3 Dominant Frequency Identification	29
2.8.4 Secondary Frequencies and Sideband Structure	30
2.8.5 Extraction of Secondary Frequencies	31
2.8.6 Frequency Ratio and Resonance Interpretation	33
2.9 Stability Analysis of Periodic Orbits	35
2.9.1 Perturbation and Linearization	35
2.9.2 State Transition Matrix	36
2.9.3 Periodic Orbits and the Monodromy Matrix	37
2.9.4 Floquet Multipliers and Stability Classification	38
2.9.5 Stability Indices	38
2.9.6 Interpretation in Phase Space	39
2.9.7 Summary of Dynamical Methods	40

3	Dynamical Models and Numerical Methods	41
3.1	Reference Frames and Normalization	41
3.2	Planar CR3BP Equations	42
3.3	J_2 -Augmented Model	44
3.4	Equilibrium Structure	45
3.4.1	Lagrange Point Computation	46
3.4.2	Jacobi Constants at the Equilibrium Points	51
3.4.3	Lagrange Point Stability	53
3.5	Lagrange Point Geometry	57
3.5.1	Classical CR3BP Configuration	58
3.5.2	J_2 -Augmented Configuration	59
3.6	Zero-Velocity Structure	60
3.6.1	Zero-Velocity Curves in the Classical CR3BP	61
3.6.2	Critical Energy Levels and Neck Opening Topology	62
3.6.3	Zero-Velocity Curves in the J_2 -Augmented Model	62
3.7	Poincaré Map Construction	64
3.7.1	Surface of Section Definition and Event Detection	65
3.7.2	Initial Condition Construction from the Jacobi Constant	66
3.7.3	Trajectory Propagation and Termination Criteria	67
3.7.4	Data Output Structure and Storage Format	68
3.7.5	Discrete Mapping and Phase-Space Representation	69
3.8	Frequency Ratio Analysis	70
3.8.1	Time Series Construction	70
3.8.2	Spectral Analysis and Frequency Extraction	71
3.8.3	Frequency Ratio Computation	73
3.9	Stability Analysis	75
3.9.1	State Transition Matrix and Monodromy Matrix Formulation	75
3.9.2	Numerical Implementation of Stability Computation	75
3.9.3	Stability Classification via Floquet Multipliers	77
3.10	Summary and Remarks of Dynamical Models and Numerical Methods	77
4	The ORBIT Computational Framework	79
4.1	Motivation for ORBIT	79
4.2	Software Architecture	80
4.2.1	Core Physics Layer	80
4.2.2	Numerical and Data-Generation Layer	81
4.2.3	Analysis and Visualization Layer	82
4.2.4	Graphical Interface and Control Layer	82
4.3	Dynamical Engine Implementation	83
4.3.1	Program Entry and Screen Routing	83
4.4	Poincaré Data Generator	84
4.5	Frequency Analysis Module	85
4.6	Stability Analysis Module	87
4.7	Graphical User Interface	89
4.7.1	Project Creation and Selection	90

4.7.2	System Definition	91
4.7.3	Dashboard and Workflow Control	92
4.7.4	Analysis and Visualization Screens	93
4.7.5	Project Data Organization and Export	94
4.8	Verification and Validation	95
4.8.1	Controlled Verification Case at $\mu = 0.3$	95
4.9	Framework Limitations	99
5	Geometric Phase-Space Exploration	100
5.1	Interpretation of Motion in the Neptune–Triton System	100
5.2	Phase-Space Mapping Results	101
5.2.1	Global Structure of the Phase Space	102
5.2.2	Evolution of Phase-Space Structure with Energy	103
5.3	General Dynamical Signatures in the Poincaré Maps	105
5.4	Directional Organization of the Phase Space	108
5.5	Geometric Synthesis and Transition	109
6	Frequency and Resonance Analysis	111
6.1	Frequency Analysis Framework	111
6.2	Frequency Ratio Plots	112
6.3	Resonance Identification and Family Construction	113
6.3.1	Worked Example: Frequency Extraction and Resonance Classification . . .	114
6.4	Resonance Overlays on Frequency and Phase Space	116
6.5	Individual Trajectory Frequency Behavior	118
6.6	Interpretation of Resonant Structure	120
6.7	Chapter Synthesis	122
7	Candidate Region Identification	125
7.1	Mission Interpretation for a Neptune–Triton Tour	125
7.2	Neptune-Centered Bounded Candidates	126
7.2.1	Compact Resonant Bounded Motion in Region II	126
7.2.2	Extended Prograde Neptune-Centered Motion in Region III	128
7.2.3	Extended Retrograde Neptune-Centered Motion in Region III	130
7.3	Touring-Region Candidates	133
7.3.1	Extended Prograde Island Structure in Region III	134
7.3.2	Retrograde Corridor Touring Candidates	138
7.4	Triton-Centered Bounded Candidates	140
7.4.1	Prograde Triton-Centered Candidates	140
7.4.2	Retrograde Triton-Centered Candidates	142
7.5	Bottleneck-Crossing Transport Candidates	143
7.6	Synthesis of Candidate Structures	145

8	Stability of Candidate Trajectories	147
8.1	Stability Framework	147
8.2	Neptune-Centered Candidate Stability	149
8.3	Touring-Region Candidate Stability	150
8.4	Triton-Centered Candidate Stability	151
8.5	Transport-Candidate Stability and Sensitivity	153
8.6	Stability Synthesis	154
9	Neptune–Triton Tour Architecture	155
9.1	Introduction and Chapter Objectives	155
9.2	Tour Architecture Framework	156
9.3	Neptune-Bounded Orbit Options	157
9.3.1	Nominal Neptune Science Orbit	158
9.3.2	Alternative Neptune-Bounded Orbit	158
9.3.3	Summary of Neptune-Bounded Phase	159
9.4	System-Wide Touring Trajectories	159
9.4.1	Nominal Touring Orbit	160
9.4.2	High-Energy Alternative Touring Orbit	161
9.4.3	Secondary Touring Candidates	161
9.4.4	Summary of Touring Phase	162
9.5	Transfer to the Triton Region	162
9.5.1	Transfer Candidate 1: Retrograde Neck Crossing	162
9.5.2	Transfer Candidate 2: Prograde Neck Crossing	163
9.5.3	Summary of Transfer Phase	164
9.6	Triton-Bounded Orbit Options	164
9.6.1	Nominal Triton Science Orbit	165
9.6.2	Alternative Triton Orbit	165
9.6.3	Summary of Triton-Bounded Phase	166
9.7	Integrated Tour Architecture	166
9.7.1	Sequential Mission Architecture	166
9.7.2	Nominal and Alternative Architectures	167
9.7.3	Key Design Insights	167
9.8	Rejected and Non-Nominal Trajectory Classes	168
9.8.1	Retrograde Trajectory Families	168
9.8.2	High-Energy and Weakly Bounded Prograde Trajectories	168
9.8.3	Interpretation of Rejected Trajectories	169
9.8.4	Implications for Future Work	169
9.9	Mission Architecture Implications and Forward Outlook	170
10	Discussion and Future Work	172
10.1	Scope of the Present Work	172
10.2	Interpretation of Results in a Mission Design Context	173
10.3	Model Limitations and Assumptions	175
10.4	Required Extensions for Real Mission Design	176
10.4.1	Higher-Fidelity Dynamical Models	176

10.4.2	Trajectory Refinement and Targeting	176
10.4.3	Maneuver Design and Δv Budgeting	176
10.4.4	Time-Dependent Mission Analysis	177
10.5	Extended Dynamical Investigations	177
10.5.1	Time-Dependent Multi-Moon Modeling	177
10.5.2	Spatial Extension of the Neptune–Triton Model	178
10.5.3	Transfer Manifold Construction	178
10.5.4	Higher-Order Gravitational Harmonics	178
10.5.5	Neptune Obliquity and Non-Coplanar Geometry	178
10.5.6	Rotating Frame Formulation in Misaligned Systems	179
10.6	Alternative Mission Architectures	179
10.7	Closing Perspective	180
A	Units and Derivations	182
A.1	Canonical Scaling Used in the Thesis	182
A.2	Neptune–Triton Reference Values	182
A.3	Practical Conversion Table	183
A.4	Interpretive Notes	184
B	Data Products and Output Schema	185
B.1	Top-Level Project Organization	185
B.2	Neptune–Triton Data Tree	186
B.3	Per-Dataset Folder Structure	186
B.4	Frequency and Resonance Run Outputs	187
B.5	Key Exported Tables	188
B.6	Naming Conventions and Practical Caveats	190
C	Algorithms	191
C.1	Poincaré Map Construction Workflow	191
C.2	Frequency-Analysis Workflow	192
C.3	Differential Correction and Stability Workflow	193
C.4	Brent Root-Finding in the ORBIT Workflow	195
C.5	Appendix Summary	196

List of Figures

2.1	The primaries m_1 and m_2 remain fixed on the x -axis, the third body m_3 moves under their gravitational influence, and the distances r_1 and r_2 define the instantaneous geometry used in the CR3BP formulation.	9
2.2	Zero-velocity curves in configuration space for a generic CR3BP system with $\mu = 0.3$, shown for multiple Jacobi levels on a single plot. The overlaid contours illustrate how the allowed and forbidden regions evolve as the Jacobi constant changes and provide a compact view of the changing topology of Hill's regions in the rotating frame.	15
2.3	A bottleneck occurs at the L1 Lagrange point of a $\mu = 0.3$ system. At lower C_J values, the bottleneck opens, allowing trajectories to pass between the two primaries.	16
2.4	The equilibria L_1 , L_2 , and L_3 lie on the line connecting the primaries, while L_4 and L_5 form the triangular equilibrium configuration.	17
2.5	Construction of a Poincaré map from a representative trajectory in the rotating frame. The left panel shows the continuous trajectory crossing the section $y = 0$ with $\dot{y} > 0$, while the right panel shows the corresponding sequence of discrete points recorded in the $(x, \dot{x})$ plane at each successive crossing.	22
2.6	Representative Poincaré signatures of regular motion in the $(x, \dot{x})$ plane. The left panel shows the fixed-point structure produced by a periodic orbit, while the right panel shows the closed invariant curve generated by quasi-periodic motion on a surrounding torus.	23
2.7	Poincaré maps in the $(x, \dot{x})$ plane for the CR3BP on the surface of section $y = 0$ with $\dot{y} > 0$. The left panel illustrates the island-chain geometry associated with resonant motion, while the right panel shows the crossing-manifold structure of a hyperbolic trajectory together with the surrounding scattered points indicative of nearby chaotic transport.	25
2.8	Illustrative quasi-periodic time series $s(t)$ exhibiting amplitude modulation produced by the superposition of multiple oscillatory components.	27
2.9	Two complementary views of an illustrative quasi-periodic spectrum. The left panel shows the full Fourier spectrum with the dominant peak and neighboring combination frequencies, while the right panel isolates the dominant frequency ω_1 as the largest-amplitude spectral component.	30
2.10	Symmetric sideband peaks about the dominant frequency ω_1 are separated by the secondary frequency ω_2	32
2.11	The secondary frequency ω_2 is recovered from the spacing of symmetric sidebands about the dominant peak.	33
2.12	Representative low-order rational values are highlighted together with the resonant condition $\omega_2/\omega_1 = p/q$. The bar heights are illustrative and indicate the relative dynamical significance of the low-order resonances.	34

3.1	Lagrange-point geometry in the rotating frame of the classical Neptune–Triton CR3BP. The figure shows the full equilibrium configuration used in the phase-space analysis, with Neptune and Triton fixed on the x -axis and the five equilibrium points plotted at their computed locations.	58
3.2	Zoomed Lagrange-point geometry in the rotating frame of the classical Neptune–Triton CR3BP. The figure resolves the region near Triton and the collinear equilibria, where the local geometry used for comparison with the J_2 -perturbed model is most clearly visible.	59
3.3	Lagrange-point geometry in the rotating frame of the J_2 -perturbed Neptune–Triton system. The figure shows the full equilibrium configuration after Neptune’s oblateness is included through the effective potential used in the perturbed model. .	59
3.4	Zoomed Lagrange-point geometry in the rotating frame of the J_2 -perturbed Neptune–Triton system. The figure highlights the local shifts of the collinear equilibria near Triton that arise from the oblateness correction.	60
3.5	Zero-velocity curves for the classical Neptune–Triton CR3BP in the rotating barycentric frame. The shaded regions indicate forbidden motion where $2\Omega(x, y) - C_J < 0$, and the zero-velocity boundary is defined by $F(x, y; C_J) = 0$	63
3.6	Zero-velocity curves for the Neptune–Triton system with Neptune’s J_2 oblateness included. Panel (a) shows the global accessibility of motion, and panel (b) shows a local zoom near Triton and the collinear equilibrium region. The J_2 perturbation produces small shifts in the zero-velocity geometry while preserving the overall topology.	64
3.7	Schematic illustration of the sideband-based frequency extraction method. Panel (a) shows a representative trajectory segment in the rotating frame. Panel (b) shows the corresponding spectrum, where the dominant peak at f_1 is flanked by symmetric sidebands at $f_1 - f_{2,B}$ and $f_1 + f_{2,B}$. The sideband-derived modulation frequency $f_{2,B}$ equals half the total peak-to-peak sideband separation: $f_{2,B} = \frac{1}{2}(f_+ - f_-)$	73
4.1	High-level architecture of the ORBIT program shown as a computational pipeline. Initial conditions are generated from the selected system parameters and Jacobi constant and are then propagated with the CR3BP or J_2 -augmented equations of motion using a shared numerical integration core. The resulting trajectories feed three parallel analysis branches: Poincaré surface-of-section extraction, frequency analysis using spectral methods, and stability analysis through state-transition and monodromy computations. The outputs of these branches are stored and visualized within ORBIT as maps, frequency ratios, stability metrics, and exported research products.	81
4.2	Data-generation tab of the ORBIT dashboard. The screen provides the controls used to configure the unified trajectory-generation workflow, including the Jacobi level, the initial-condition sweep, the Poincaré crossing settings, and the base frequency-analysis parameters associated with each run.	85
4.3	Poincaré-analysis tab of the ORBIT dashboard. The screen provides dataset filters and access to the plotting and table tools used to inspect the surface-of-section outputs generated by the unified trajectory workflow.	86

4.4	Frequency-analysis tab of the ORBIT dashboard. The screen organizes access to the saved frequency runs and to the plotting tools used to inspect frequency-ratio trends, data tables, and associated spectral products within the active project. . . .	87
4.5	Stability-analysis tab of the ORBIT dashboard. The screen displays the stored target list, the controls used to launch stability runs, and the tabulated results used to classify the candidate trajectories identified in the earlier analysis stages.	88
4.6	Startup screen of the ORBIT graphical user interface for the Neptune–Triton analysis framework. The figure shows the application entry point and project-oriented environment through which the ORBIT workflow is initiated.	89
4.7	System-definition screen reached immediately after creating a new ORBIT project. The interface presents three input pathways: selecting the primary and secondary bodies from the catalog, entering the masses and radii manually, or specifying the reduced mass parameter μ directly for an abstract CR3BP study.	91
4.8	Dashboard screen of the ORBIT graphical user interface for the Neptune–Triton framework. The figure shows the central project-control screen from which model settings, numerical generation steps, and downstream analysis modules are coordinated.	93
4.9	Computed Lagrange-point geometry for the controlled verification case with $\mu = 0.3$. This figure verifies that the equilibrium-point solver reproduces the expected asymmetric CR3BP configuration for the prescribed mass parameter.	96
4.10	Zero-velocity curves in configuration space for the CR3BP at a distinct Jacobi level, showing how the topology of the accessible regions changes as the energy varies. The openings near the collinear equilibria illustrate the transport channels that emerge at critical Jacobi constants.	97
4.11	Verification products for the controlled case with $\mu = 0.3$ using $C_J = 3.88514958412578$, shown using the prograde (blue) family. Left panel (a): Poincaré map in $(x, \dot{x})$, demonstrating the organized surface-of-section structure produced by the trajectory generator and crossing logic. Right panel (b): corresponding frequency-ratio trend extracted from the same prograde family, confirming that the spectral-analysis pipeline recovers a smooth and interpretable dynamical relationship from the propagated trajectories.	98
5.1	Comparison of prograde and retrograde motion in a rotating reference frame. The top panel shows a standard prograde system, while the bottom panel shows the Neptune–Triton retrograde system. The rotating frame (x, y) is used for computation and visualization, while the inertial frame (X, Y) represents the physical motion of the system. In the rotating frame, the retrograde Neptune–Triton configuration produces a reflected orientation relative to the prograde case. The velocity component $\dot{y}$ is shown for reference, where $\dot{y} > 0$ and $\dot{y} < 0$ indicate the direction of motion across the Poincaré map. The sizes of the primary bodies are not to scale and are shown for visual perspective only.	101
5.2	Global Poincaré section of the Neptune–Triton system in the $(x, \dot{x})$ plane at $C_J = 3.018641378$ for the J_2 -perturbed model. Shading indicates the four primary dynamical regions discussed in the text; the boundaries are approximate and are used for interpretive clarity.	102

5.3	Poincaré maps in the $(x, \dot{x})$ plane for the J_2 -perturbed Neptune–Triton system at four Jacobi constants, with the surface of section defined by $y = 0$ and $\dot{y} > 0$. The panels are ordered from higher to lower Jacobi constant and show the expansion of accessible phase space together with the increasing prominence of scattered crossings as the energy decreases.	104
5.4	Example of a smooth invariant curve on the Poincaré map corresponding to a quasi-periodic trajectory confined to an invariant torus in the full phase space. This figure was taken from Region IV of the Poincaré map at $C_J = 3.018641378$. The trajectory repeatedly intersects the surface of section, producing a continuous closed loop that indicates regular, bounded motion. In the Neptune–Triton system, such curves are indicative of stable regions of phase space where trajectories exhibit low sensitivity to perturbations and remain confined over long timescales.	106
5.5	Local view of a resonant island chain in the Neptune–Triton Poincaré map at $C_J = 3.0072570569$, shown over the window $x \in [0.300, 0.315]$ and $\dot{x} \in [-0.25, 0.25]$. The repeated closed islands mark a resonant family in which the section crossings cluster at a discrete set of phases rather than filling a continuous invariant curve. . .	106
5.6	Retrograde-only local view of the Neptune–Triton Poincaré map at $C_J = 3.0072570569$, shown over the window $x \in [0.20, 0.26]$ and $\dot{x} \in [-1.5, 1.5]$. The central chain and the associated crossing structure illustrate the loss of simple invariant-curve organization in a locally unstable region, where nearby resonant structures and separatrix-type crossings produce more complicated phase-space geometry. . . .	107
5.7	Poincaré maps in the $(x, \dot{x})$ plane for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$, shown separately for prograde and retrograde trajectory sets. The surface of section is defined by $y = 0$ with $\dot{y} > 0$, and only the bounded portion of each filtered orbit family is retained. The comparison shows the compact, regular organization of the prograde domain and the broader mixed structure of the retrograde domain.	108
6.1	Frequency ratio plots in the $(x_0, \omega_2/\omega_1)$ plane for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. The full plot shows the global spectral organization of the sampled trajectories, while the zoomed panel highlights a horizontal band structure used to identify candidate resonant neighborhoods.	113
6.2	Configuration-space trajectory for the retrograde orbit at $x_0 = 0.2099207594$ and $C_J = 3.018641378$ under the J_2 -perturbed Neptune–Triton model. The bounded, regularly repeating path is consistent with membership in a structured resonant family.	115
6.3	Normalized frequency spectrum for the retrograde trajectory at $x_0 = 0.2099207594$ and $C_J = 3.018641378$. The dominant peak identifies the primary frequency ω_1 , and the symmetric sideband pair flanking it is used to extract the secondary frequency ω_2 via the sideband spacing method [21].	116

6.4	Overlay of resonance-family classifications on the frequency ratio plot for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. Each color corresponds to a different resonance family identified by ORBIT, while gray points represent trajectories that were not classified as resonant within the selected tolerance. This overlay allows the resonance families to be traced directly in ratio space, revealing how they organize the spectral distribution and where they cluster about rational values.	117
6.5	Overlay of resonance-family classifications on the Poincaré map for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. Each color corresponds to a different resonance family identified by ORBIT, while gray points represent trajectories that were not classified as resonant within the selected tolerance. This overlay allows the resonance families to be traced directly in phase space, revealing how they organize the invariant structures and where they cluster about rational values.	118
6.6	Representative trajectory-level spectra for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. The left panel gives a comparatively clean low-order resonant case whose dominant frequencies are easy to identify, while the right panel shows a trajectory whose denser secondary content makes the spectral interpretation less immediate.	119
6.7	Comparison of frequency-resonance overlays for four representative J_2 -perturbed Neptune–Triton energy levels. Colored points denote trajectories assigned to approximate resonance families by ORBIT, while light gray points remain unclassified within the selected tolerance. Viewed together, these overlays show which horizontal resonance bands persist across the sampled C_J range, which families remain narrow and strongly populated, and which lower-ratio structures appear only as sparse or fragmented classifications.	121
6.8	Comparison of resonance-family projections onto the Poincaré maps for four representative J_2 -perturbed Neptune–Triton energy levels. The same families identified in frequency-ratio space are recolored here in phase space, allowing the persistent spectral bands of Fig. 6.7 to be traced onto the corresponding invariant structures. Gray points indicate trajectories that were not assigned to a resonance family within the selected tolerance.	123
7.1	Region II Poincaré section at $C_J = 3.018641378$. Upper: full-window context; lower: zoomed view of the candidate island used for seed identification.	127
7.2	Extended prograde Neptune-centered structures in Region III at $C_J = 3.018641378$. (a) Regular map, full window $x \in [0, 0.95]$, $\dot{x} \in [-2.5, 2.5]$, prograde in red. (b) Resonance overlay, same window. (c) Zoomed resonance overlay, $x \in [0.06, 0.22]$, for seed-selection context.	129
7.3	Extended retrograde Neptune-centered structures in Region III at $C_J = 3.018641378$. (a) Regular map, full window $x \in [0, 0.95]$, $\dot{x} \in [-2.5, 2.5]$, retrograde in blue. (b) Resonance overlay, same window. (c) Zoomed resonance overlay, $x \in [0.075, 0.115]$, covering candidate island ranges $x \approx 0.081\text{--}0.090$ and $0.091\text{--}0.110$	132

7.4	Interior prograde invariant structures in Region III at $C_J = 3.0107806367$. (a) Regular map, full window $x \in [0, 0.95]$, $\dot{x} \in [-2.5, 2.5]$, prograde in red. (b) Resonance overlay, same window. (c) Zoomed resonance overlay, $x \in [0.30, 0.45]$, $\dot{x} \in [-0.65, 0.65]$, outer island-chain neighborhood.	135
7.5	Zoomed view of the outer resonant island chain identified within the prograde Region III structure (see Fig. 7.4).	136
7.6	Zoomed view of the interior resonant island chain located deeper within the Region III phase space.	137
7.7	Zoomed view of the non-resonant crescent-shaped invariant band within Region III.	137
7.8	Zoomed retrograde Poincaré section at $C_J = 3.0107806367$ with the same resonance-overlay coloring as the ORBIT Poincaré analysis tool (labels from the run <code>full_run.csv</code> ; colors from the shared resonance palette, with unclassified points in light gray). Representative structures discussed in the text include (1) outer interior 1 : 1 island near $x \approx 0.59$, (2) intermediate 1 : 11 crescent near $x \approx 0.72$, (3) embedded 1 : 6 island near $x \approx 0.78$, and (4) deep interior 4 : 15 island near $x \approx 0.87$	139
7.9	Poincaré section at $C_J = 3.018641378$ showing the Triton-centered Region IV structures. The region consists of compact invariant islands located near $x \approx 1.0$, separated from the interior Neptune-centered domain.	141
7.10	Resonance-colored Poincaré section at $C_J = 3.018641378$ for Region IV. The nested structure of the Triton-centered invariant curves is evident, with distinct families corresponding to prograde and retrograde motion.	141
7.11	Frequency ratio map at $C_J = 3.0143741909$. Candidate initial conditions are identified near the L_1 bottleneck region ($x_0 \approx 0.94$ – 0.98) while avoiding clearly defined resonance groupings.	144
7.12	Frequency resonance overlay at $C_J = 3.0143741909$. Resonance-associated trajectories are identified and excluded from the transfer candidate set, leaving non-resonant points near the L_1 region as primary candidates.	144
9.1	Neptune-bounded trajectory for $x_0 = -0.250079$ in the rotating frame. The motion remains tightly confined about Neptune, providing a stable platform for initial mission operations and detailed observation of the primary body. NOTE: Neptune-system moon guides are visual references only. Not included in the propagated force model.	158
9.2	Alternative Neptune-bounded trajectory for $x_0 = 0.20992$. The orbit remains stable and centered about Neptune, but with a larger radial extent and more elongated structure compared to the nominal case.	159
9.3	Nominal system touring trajectory for $x_0 = -0.820079$. The orbit spans a large portion of the Neptune system while maintaining bounded behavior, enabling broad observational coverage.	160
9.4	High-energy touring trajectory for $x_0 = 0.799921$. The orbit provides extensive spatial coverage across the Neptune system, offering broad access to the moon region.	161
9.5	Retrograde bottleneck-crossing trajectory for $x_0 = 0.93992$. This case provides a geometrically meaningful neck crossing and is selected as the primary transfer candidate in the tour architecture.	163

9.6	Prograde transfer trajectory for $x_0 = 0.97992$. This case provides a second geometrically meaningful bottleneck-crossing path and is retained as a secondary transfer option within the tour architecture.	164
9.7	Triton-bounded trajectory options at $x_0 = 1.00992$. The upper panel shows the nominal prograde case, while the lower panel shows the retrograde alternative. Both remain tightly confined about Triton and define viable terminal science configurations within the rotating-frame model.	165
9.8	Representative retrograde trajectories exhibiting structured, large-scale motion within the Neptune–Triton system. While dynamically rich, these trajectories do not satisfy the stability criteria required for inclusion in the tour architecture. . . .	169
9.9	High-energy prograde trajectory for $x_0 = -0.92007$. Although the trajectory spans a large portion of the system, it does not exhibit the bounded, stable behavior required for mission applications.	170

List of Tables

3.1	Neptune–Triton Lagrange-point positions for the classical planar CR3BP.	51
3.2	Neptune–Triton Lagrange-point positions for the planar CR3BP with Neptune J_2 enabled.	51
3.3	Jacobi constants at the Neptune–Triton Lagrange points for the classical planar CR3BP.	53
3.4	Jacobi constants at the Neptune–Triton Lagrange points for the planar CR3BP with Neptune J_2 enabled.	53
3.5	Eigenvalues and stability classification of the Neptune–Triton Lagrange points for the classical planar CR3BP.	56
3.6	Eigenvalues and stability classification of the Neptune–Triton Lagrange points for the planar CR3BP with Neptune J_2 enabled.	57
6.1	Representative trajectories assigned to the 1 : 2 resonance family for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. The listed rows are drawn from the ORBIT resonance-family output and show how the extracted ratio values cluster about the rational target $1/2$ within the selected tolerance.	114
8.1	Stability results for the Neptune-centered bounded candidates at $C_J = 3.018641378$. Reported candidate labels follow the finalized Chapter 7 notation.	149
8.2	Stability results for the touring-region candidates selected from the interior of Region III at $C_J = 3.0107806367$	151
8.3	Stability results for the Triton-centered candidates selected from Region IV at $C_J = 3.018641378$	152
8.4	Stability results for the bottleneck-crossing transport candidates at $C_J = 3.0143741909$	153
A.1	Dimensional Neptune–Triton reference quantities used to construct the canonical CR3BP normalization in this thesis.	183
A.2	Canonical Neptune–Triton scales and normalized system parameters used in the thesis.	183
A.3	Practical conversion factors between nondimensional CR3BP units and dimensional Neptune–Triton values.	184
B.1	Principal top-level folders in the ORBIT project tree used for the Neptune–Triton thesis workflow.	185
B.2	Primary project-level files in the Neptune–Triton data directory.	186
B.3	Representative contents of a Jacobi-level dataset folder.	186
B.4	Principal files produced in a representative frequency-analysis run.	187
B.5	Primary columns in the Neptune–Triton <code>run_catalog.csv</code> file.	188
B.6	Key columns used in the frequency and resonance export tables.	189
B.7	Key columns used in the stability-analysis export tables.	189

Abstract

The Neptune–Triton system represents one of the most dynamically compelling targets for future deep-space exploration, yet the trajectory design infrastructure needed to support a long-duration orbital mission remains largely undeveloped. This thesis addresses that gap by constructing a systematic dynamical framework for identifying, classifying, and evaluating trajectories suitable for a Neptune–Triton tour mission, using tools drawn from nonlinear dynamical systems theory.

The analysis is conducted within the planar Circular Restricted Three-Body Problem (CR3BP), formulated in the rotating barycentric frame of the Neptune–Triton system and extended to include Neptune’s J_2 oblateness perturbation as a correction to the gravitational potential. Large trajectory sets are generated across selected energy levels and analyzed using three complementary tools. Poincaré maps characterize the geometric structure of phase space, frequency analysis identifies resonance families through spectral decomposition of trajectory time histories, and linear stability analysis evaluates periodic-orbit candidates via the monodromy matrix and Floquet multipliers.

All numerical workflows are implemented within ORBIT, a computational framework developed as part of this work that integrates trajectory generation, phase-space mapping, spectral analysis, stability evaluation, and data visualization into a unified and reproducible pipeline.

The analysis identifies trajectory families supporting Neptune-bounded operations, system-wide touring, bottleneck-crossing transfer into the Triton region, and Triton-bounded terminal orbits. These families are synthesized into a conceptual tour architecture demonstrating how the natural phase-space structure of the Neptune–Triton system can be directly leveraged for mission design. The results establish a reproducible dynamical foundation for future trajectory optimization, higher-fidelity modeling, and mission concept development in the Neptune–Triton system.

1. Introduction

The Neptune–Triton system presents a compelling problem at the intersection of planetary science, astrodynamics, and nonlinear dynamical systems. On the scientific side, Neptune remains one of the least explored major planets in the solar system, and Triton remains one of its most dynamically and geologically intriguing satellites. On the mission-design side, any future spacecraft sent to the system would have to operate in a gravitational environment shaped not only by Neptune itself, but also by Triton’s substantial dynamical influence. These two perspectives motivate the central question of this thesis: whether the natural phase-space structure of the Neptune–Triton system can be used to identify long-term trajectories that are compatible with sustained tour operations.

This chapter establishes the motivation, scope, and research direction of the work before the more detailed theoretical and computational developments begin in the later chapters. It first reviews the scientific case for a dedicated return to the Neptune system and explains why the Neptune–Triton subsystem is dynamically significant from a trajectory-design perspective. It then defines the specific problem addressed in the thesis, summarizes the principal contributions of the study, clarifies the scope and assumptions of the adopted model, and concludes by outlining the organization of the full document. In this way, the introduction serves as the bridge between the broader scientific context of outer-planet exploration and the focused dynamical-systems analysis carried out in the remainder of the thesis.

1.1 Motivation

The *Voyager 2* flyby in 1989 provided the first and only *in situ* exploration of the Neptune–Triton system, establishing the initial scientific basis for understanding this distant ice giant and its satellites [1, 2]. In the decades since that encounter, advances in observational capability, theoretical modeling, and computational analysis have significantly expanded knowledge of the system while simultaneously revealing new scientific questions. Modern telescopes, improved detection techniques, and higher-fidelity modeling have shown that the Neptune–Triton system is far more complex and dynamically active than initially suggested, strengthening the motivation for renewed exploration [3].

Continued observations have revealed substantial variability in Neptune’s atmosphere. High-resolution imaging and spectral analysis from both ground-based and space-based observatories have enabled long-term monitoring of atmospheric dynamics, including the evolution of large-scale storm systems on timescales ranging from months to decades [3]. These observations have refined atmospheric models while also exposing limitations in current understanding, reinforcing the need for sustained measurements obtainable only through a dedicated spacecraft mission.

Neptune’s magnetic field presents another major scientific puzzle. *Voyager* measurements revealed an unusually tilted and significantly offset magnetic dipole, producing a highly asymmetric magnetosphere [2, 4]. Advances in dynamo theory suggest that this configuration may provide insight into the internal structure of ice giant planets, yet the limited temporal coverage of the *Voyager* flyby prevents a complete characterization. A modern mission could enable detailed three-dimensional mapping of this environment and its variability.

The Neptune–Triton system’s rings also present unresolved dynamical problems. *Voyager* observations revealed narrow rings containing confined arc structures whose long-term stability cannot be explained by simple gravitational models alone [2]. Subsequent theoretical work has

highlighted the importance of resonance dynamics and satellite–ring interactions, which remain incompletely understood and require direct observational validation [3].

Observational advances have further expanded knowledge of Neptune’s satellite system. The discovery of several irregular satellites with highly eccentric and inclined orbits indicates a complex dynamical history and suggests that the system remains only partially characterized [3].

Among these bodies, Triton stands out as the dominant dynamical component. Its retrograde and inclined orbit strongly suggests a captured origin, likely involving the disruption of a binary object early in solar system history [5]. Voyager imagery revealed active surface processes, including geyser-like plumes, indicating ongoing geological activity [6]. These features, combined with advances in cryovolcanism and ice-shell modeling, make Triton one of the most compelling targets for future exploration.

The growing recognition of the scientific importance of ice giant exploration is reflected in recent strategic planning efforts. The *Neptune Odyssey* flagship mission concept proposes a comprehensive investigation of Neptune and its satellite system [5], while the 2023–2032 Planetary Science and Astrobiology Decadal Survey identifies ice giant exploration as a major priority for advancing planetary science [7, 8].

1.2 Dynamical Significance of the Neptune–Triton System

While the scientific motivation for returning to Neptune is compelling, the feasibility of such a mission depends critically on the ability to design trajectories that efficiently navigate the system’s complex gravitational environment. In contrast to two-body motion, long-duration spacecraft trajectories in the Neptune–Triton system are governed by multi-body interactions, which give rise to intricate phase-space structures including resonances, invariant manifolds, and regions of stability and instability [5, 9–11].

Triton’s relatively large mass and proximity to Neptune create a gravitational subsystem that dominates the local dynamical environment. Although additional moons and ring structures are present, their influence is comparatively weak, and the primary features of spacecraft motion are determined by the interaction between Neptune and Triton. This interaction produces a rich dynamical environment that is well suited for analysis using modern nonlinear dynamical systems techniques [5, 9, 11, 12].

From a dynamical perspective, the Neptune–Triton system exhibits behavior characteristic of multi-body gravitational systems, including quasi-periodic motion, resonance structures, and chaotic transport. These features govern the accessibility and long-term stability of trajectories within the system and therefore play a central role in mission design [5, 9–11, 13].

Understanding the organization of phase space in this environment is essential for identifying trajectories that remain bounded, persist over mission-relevant timescales, and provide access to scientifically important regions. Accordingly, this work focuses on the analysis of dynamical structures within the Neptune–Triton system, with the goal of identifying stable and tour-suitable trajectories that enable sustained exploration of the outer Solar System [5, 9, 14].

1.3 Problem Statement

With a dedicated Neptune–Triton mission identified as a high-priority objective in current planetary science planning, trajectory design emerges as a central element of mission architecture development.

Trajectory selection determines the spacecraft’s accessibility to key scientific targets within the system, including Neptune, Triton, and the surrounding satellite and ring environments. Because propulsion capability and mission duration impose strict operational constraints, trajectories must be designed to efficiently utilize available resources while maximizing scientific return. Consequently, the identification and analysis of viable trajectory architectures represent a critical step in early mission planning and feasibility assessment [15].

Advances in trajectory design methods and numerical modeling capability since the *Voyager 2* era have significantly expanded the range of mission architectures that can be investigated. Modern computational tools allow detailed simulation of multi-body gravitational environments and enable the systematic exploration of dynamical structures that can support long-duration spacecraft motion. Improved knowledge of planetary system parameters and enhanced numerical integration techniques now make it possible to investigate trajectory families that were computationally inaccessible during earlier mission studies [9, 15, 16].

Despite these advances, relatively little research has focused specifically on mission-oriented trajectory design within the Neptune–Triton system. In particular, there remains limited characterization of long-term bounded motion, stability classifications of naturally occurring trajectories, and the role of resonant dynamical structures in enabling sustained exploration of the system. From a mission design perspective, trajectories suitable for extended system exploration must satisfy several fundamental criteria: they must remain dynamically bounded within the region of interest, persist over mission-relevant time scales, and remain controllable under realistic propulsion constraints. The absence of a systematic dynamical framework capable of identifying and classifying such trajectories represents a significant gap in current Neptune–Triton mission analysis.

This thesis addresses this gap by investigating the dynamical structure of the Neptune–Triton system within the framework of the circular restricted three-body problem. Using tools from nonlinear dynamical systems theory—including Poincaré maps, frequency analysis, and Floquet stability assessment—the work identifies trajectory families and phase-space structures that naturally support long-lived motion within the system. By characterizing these features and their associated stability properties, this study aims to provide a foundation for future mission designers seeking to construct efficient and scientifically productive exploration trajectories within the Neptune–Triton system [9, 11, 16].

1.4 Thesis Contributions

This thesis develops a dynamical framework for identifying and evaluating long-lived trajectories within the Neptune–Triton system using tools from nonlinear dynamical systems theory. The primary contributions of this work are summarized as follows:

1. Formulation of the planar Neptune–Triton system within the Circular Restricted Three-Body Problem (CR3BP), including nondimensionalization and incorporation of Neptune’s J_2 oblateness perturbation, establishing the dynamical model used throughout the analysis.
2. Systematic construction of Poincaré maps across ranges of energy levels and initial conditions, enabling identification of invariant phase-space structures including periodic orbits, quasi-periodic invariant curves, and chaotic regions.

3. Extraction of dominant and secondary frequencies from trajectory time series and construction of frequency spectra to characterize the oscillatory structure of motion.
4. Computation of frequency ratios $\nu = \omega_2/\omega_1$ to identify resonant relationships and map resonance structure within the phase space of the Neptune–Triton system.
5. Identification of candidate trajectory regions supporting long-duration bounded motion through combined geometric (Poincaré) and spectral (frequency) analysis.
6. Stability classification of selected periodic orbits through computation of the monodromy matrix and analysis of Floquet multipliers, distinguishing dynamically stable and unstable trajectory families.

Together, these contributions establish a reproducible dynamical systems framework for analyzing trajectory behavior in the Neptune–Triton system and provide a foundation for extending the analysis to higher-fidelity models and mission design applications.

1.5 Scope and Assumptions

The analysis presented in this thesis is conducted within the framework of the planar Circular Restricted Three-Body Problem (CR3BP). In this model, the motion of a massless spacecraft is governed by the gravitational influence of two primary bodies—Neptune and Triton—which move in circular orbits about their common center of mass. The spacecraft is assumed to have negligible mass relative to the primaries and therefore does not influence their motion. These assumptions produce an autonomous dynamical system that captures the dominant gravitational structure of the Neptune–Triton environment while remaining analytically and computationally tractable [9, 12, 17].

The full spatial CR3BP is defined in a six-dimensional phase space,

$$(x, y, z, \dot{x}, \dot{y}, \dot{z}).$$

In this work, the planar CR3BP is adopted by restricting motion to the orbital plane of the primaries through the conditions $z = 0$ and $\dot{z} = 0$, reducing the system to a four-dimensional phase space

$$(x, y, \dot{x}, \dot{y}).$$

This reduction preserves the principal dynamical structures governing motion while enabling efficient application of phase-space analysis tools such as Poincaré maps and frequency analysis [9, 11, 17].

All equations of motion are expressed in the synodic rotating reference frame of the primaries and are formulated using canonical nondimensional units. In this normalization, the distance between the primaries is set to unity, the total mass of the system is normalized to unity, and the mean motion is equal to one. These conventions simplify the governing equations and allow results to be interpreted in a system-independent form. Dimensional quantities can be recovered through appropriate scaling based on the physical parameters of the Neptune–Triton system [12, 17, 18].

The analysis focuses on gravitational dynamics within the CR3BP framework and does not include additional perturbations such as solar gravity, radiation pressure, or non-spherical harmonics beyond the J_2 term. The inclusion of Neptune’s J_2 perturbation provides a first-order correction to

the gravitational potential, capturing the dominant effect of planetary oblateness while maintaining computational efficiency. Higher-order perturbations and full ephemeris models are not considered and are reserved for future work.

1.6 Thesis Organization

This thesis is organized to distinguish clearly between theoretical development, computational implementation, and dynamical application.

Chapter 2 develops the theoretical and methodological foundations used throughout the study. It introduces the CR3BP formulation, derives the governing equations of motion, and presents the dynamical systems tools used to analyze trajectory behavior. These include the Jacobi integral and zero-velocity surfaces, Lagrange points and their stability properties, Poincaré maps as a method for reducing the dimensionality of the flow, and frequency analysis for characterizing quasi-periodic and resonant motion. The chapter focuses on the underlying mathematics and physical interpretation of these methods independent of any specific numerical implementation.

Chapter 3 describes how these methods are applied in the context of the Neptune–Triton system. It details the numerical integration procedures, construction of Poincaré maps, extraction of trajectory time series, implementation of spectral analysis techniques, and computation of frequency ratios. Model parameters, sampling choices, and data-generation procedures are also specified.

Chapter 4 presents the computational framework developed to support this analysis. The architecture of the ORBIT software system is described, including modules for trajectory generation, phase-space mapping, frequency analysis, and stability evaluation, along with tools for visualization and data management.

Chapter 5 presents the geometric phase-space results obtained from the Neptune–Triton exploration. Poincaré maps are used to identify the large-scale organization of the flow, the evolution of invariant structures with Jacobi constant, and the major dynamical regions that later serve as the basis for candidate selection.

Chapter 6 refines that geometric picture using spectral analysis. Frequency-ratio products, resonance overlays, and trajectory-level frequency content are used to distinguish coherent resonance-supported families from more weakly organized or transitional structures, thereby narrowing the broader phase-space survey to a smaller set of dynamical candidates.

Chapter 7 uses the combined geometric and spectral results to identify specific candidate regions and representative trajectory classes. Neptune-centered bounded candidates, touring-region candidates, Triton-centered candidates, and bottleneck-crossing transport candidates are compared across the selected energies and organized according to their likely mission roles.

Chapter 8 evaluates the local dynamical robustness of those shortlisted candidates. Periodic-orbit refinement, monodromy matrices, Floquet multipliers, and scalar stability indices are used to determine which candidates remain dynamically coherent enough to justify follow-on mission interpretation.

Chapter 9 then interprets the surviving candidates at the trajectory level and assembles them into a conceptual Neptune–Triton tour architecture. In this chapter, local bounded motion, broad touring motion, bottleneck-crossing transfer trajectories, and Triton-centered terminal options are considered together in mission-design terms.

Chapter 10 closes the thesis by discussing the scope and limitations of the present framework, interpreting the broader significance of the results, and outlining the modeling, analysis, and

mission-design extensions needed to move from the current idealized dynamical study toward higher-fidelity future work. In particular, it identifies the major next steps required to extend the present analysis toward higher-fidelity ephemeris models, three-dimensional dynamical formulations, and more mission-specific trajectory design constraints.

Together, these chapters progress from foundational dynamical theory through computational implementation to application-driven trajectory analysis in the Neptune–Triton system. The overall organization is intended to provide a rigorous and reproducible framework that supports both the present dynamical investigation and future extensions toward higher-fidelity mission design.

2. Dynamical Background and Analysis Methods

This chapter introduces the dynamical framework and analysis methods used throughout this study. The Neptune–Triton system forms a gravitational environment in which spacecraft motion is strongly influenced by the interaction between Neptune and its largest satellite. Because Triton’s mass dominates the satellite dynamics of the system, motion in the vicinity of Neptune can be modeled effectively using the Circular Restricted Three-Body Problem (CR3BP), in which Neptune and Triton act as the two primary bodies while the spacecraft is treated as a massless third body.

Within this framework, the fundamental dynamical structure of the system is first developed through the Jacobi constant, zero-velocity curves, and the equilibrium (Lagrange) points that arise in the rotating frame. These quantities define the energetic constraints and geometric structure that govern the motion of trajectories in the system.

To analyze the more intricate phase–space behavior that emerges from these dynamics, two complementary analytical tools are introduced. The first is the Poincaré map, which provides a geometric representation of trajectory intersections with a surface of section and reveals periodic, quasi-periodic, and resonant structures in phase space. The second is frequency analysis, which characterizes quasi-periodic motion through the identification of fundamental spectral components in the trajectory time history. Together, these methods provide the conceptual foundation for identifying resonant trajectories and interpreting the underlying dynamical structure of the Neptune–Triton system.

2.1 The Circular Restricted Three-Body Problem

The motion of a spacecraft within a planetary satellite system can often be approximated using the Circular Restricted Three-Body Problem (CR3BP). In this model the gravitational environment is represented by two massive bodies. Collectively, these are referred to as the *primaries*, while individually they are distinguished as the *primary* body and the *secondary* body. A third body of negligible mass moves under their gravitational influence. For spacecraft applications the third body represents the spacecraft itself, whose mass is many orders of magnitude smaller than that of the planetary bodies. As a result, the spacecraft does not influence the motion of the primaries and instead evolves within the gravitational field they generate [12, 17].

The terminology used to describe the CR3BP reflects the simplifying assumptions that define the model. The term *circular* refers to the assumption that the two primaries move in circular orbits about their common barycenter. This assumption removes variations in orbital distance between the primaries and allows their motion to be described by a constant angular velocity. The term *restricted* indicates that the mass of the third body is sufficiently small that it does not affect the motion of the primaries. Consequently, the trajectories of the two massive bodies are known in advance and remain unchanged regardless of the spacecraft trajectory.

These assumptions significantly simplify the full three-body problem. Instead of solving for the motion of three mutually interacting bodies, the CR3BP reduces the problem to determining the motion of a single particle within the gravitational field generated by the two primaries. Despite these simplifications, the CR3BP retains many of the essential dynamical structures that govern motion in multi-body systems, including equilibrium points, periodic orbit families, and resonant phase-space structures [9, 12, 17].

A particularly useful representation of the CR3BP arises when the equations of motion are expressed in a reference frame that rotates with the primaries. In this rotating frame the two bodies remain fixed at constant positions, allowing the trajectory of the spacecraft to be analyzed relative to stationary reference points associated with the primaries. The formulation of the CR3BP in this rotating coordinate system is introduced in Subsection 2.1.1.

2.1.1 Rotating Frame and Normalized Coordinates

In the Circular Restricted Three-Body Problem, the motion of the primaries is most conveniently described in a reference frame that rotates with their orbital motion about the system barycenter. In this rotating frame the line connecting the two primaries remains fixed, allowing the trajectory of the third body to be studied relative to stationary reference points associated with the primaries [12, 17].

Let m_1 denote the mass of the primary body and m_2 the mass of the secondary body. The two bodies revolve in circular orbits about their common barycenter with constant angular velocity. A Cartesian coordinate system (x, y, z) is defined such that the origin lies at the system barycenter, the x -axis lies along the line connecting the two primaries, and the z -axis is perpendicular to the orbital plane of the primaries.

In order to simplify the mathematical formulation of the problem, the CR3BP is commonly expressed using a nondimensional coordinate system. In this formulation the physical quantities that describe the system are scaled using characteristic values associated with the motion of the two primaries. This normalization removes unnecessary physical units from the governing equations and allows the dynamical structure of the system to be studied in a more general form [12, 17].

Three characteristic quantities are chosen to define the nondimensional scaling. The distance between the two primaries is selected as the unit of length, the sum of the two masses defines the unit of mass, and the angular velocity of the rotating frame defines the unit of time. Under these scalings the distance between the primaries, total mass, and mean motion are set to unity:

$$\begin{aligned} a &= 1, \\ n &= 1, \\ m_1 + m_2 &= 1 \end{aligned} \tag{2.1}$$

Expressing the CR3BP in this normalized form significantly simplifies the equations of motion and reveals the fundamental dynamical structure of the problem. In particular, the system can be described using a single dimensionless parameter that characterizes the relative masses of the primaries.

Within this coordinate system the positions of the primaries are determined by the dimensionless *mass parameter*

$$\mu = \frac{m_2}{m_1 + m_2}. \tag{2.2}$$

The parameter μ therefore represents the fraction of the total system mass contained in the secondary body.

Using this definition, the primary body is located at $(-\mu, 0, 0)$ and the secondary body is located at $(1 - \mu, 0, 0)$. The distance between the two bodies is therefore unity in the normalized coordinate system.

CR3BP Rotating-Frame Geometry




Figure 2.1. The primaries m_1 and m_2 remain fixed on the x -axis, the third body m_3 moves under their gravitational influence, and the distances r_1 and r_2 define the instantaneous geometry used in the CR3BP formulation.

With these scalings the dynamical system is fully characterized by the single parameter μ . Expressing the problem in this normalized rotating frame greatly simplifies the mathematical formulation and allows the structure of phase space to be studied independently of the specific physical units of the system [12, 17]. The resulting rotating-frame geometry, including the fixed locations of the primaries and the distances r_1 and r_2 , is shown in Figure 2.1. Before the equations of motion governing this geometry are introduced, however, it is useful to first define the effective potential of the CR3BP, from which the equations of motion follow directly; this is done in Section 2.2, with the full equations of motion developed subsequently in Section 2.3.

2.2 Effective Potential of the CR3BP

In the Circular Restricted Three-Body Problem, the spacecraft is attracted gravitationally by both primaries. However, because the motion is described in a rotating frame rather than an inertial frame, the equations also contain inertial effects associated with the rotation of the coordinate system. For this reason, it is useful to define a single scalar function that combines the gravitational contributions of the two primaries with the centrifugal contribution of the rotating frame [12, 17, 18].

Let (x, y, z) denote the position of the spacecraft in the normalized rotating frame. As introduced in Subsection 2.1.1, the distances from the spacecraft to the two primaries are

$$\begin{aligned} r_1 &= \sqrt{(x + \mu)^2 + y^2 + z^2}, \\ r_2 &= \sqrt{(x - 1 + \mu)^2 + y^2 + z^2}. \end{aligned} \tag{2.3}$$

These distances determine the gravitational attraction exerted by each primary. In normalized units, the gravitational potential associated with the two-body attraction of the primaries may be written as

$$U_g = \frac{1 - \mu}{r_1} + \frac{\mu}{r_2}. \tag{2.4}$$

The first term represents the gravitational contribution of the primary body, while the second term

represents the contribution of the secondary body. If the motion were described in an inertial frame, these two terms alone would define the gravitational potential governing the motion of the spacecraft. In the present formulation, this quantity is written in the positive-sign convention commonly used in CR3BP development, so that it can be combined directly with the centrifugal term in the effective potential.

The CR3BP, however, is formulated in a frame that rotates with the primaries. Because this coordinate system is non-inertial, an additional inertial effect appears in the form of the centrifugal acceleration. This contribution can also be represented by a scalar potential term,

$$U_c = \frac{1}{2}(x^2 + y^2). \quad (2.5)$$

This term depends only on the radial distance of the spacecraft from the rotation axis of the frame. It does not arise from gravity, but rather from the choice to observe the motion in a rotating coordinate system.

It is therefore convenient to combine the gravitational and centrifugal contributions into a single function,

$$\Omega(x, y, z) = \frac{1 - \mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2). \quad (2.6)$$

This function is commonly referred to as the *effective potential* of the CR3BP. In some texts it is also called the *pseudo-potential* [12, 17]. Both names emphasize the same important point: the function Ω is not purely gravitational. Rather, it combines the gravitational potential of the two primaries with the centrifugal potential associated with the rotating frame. For this reason, it is more accurate to interpret Ω as an effective scalar function that governs the structure of motion in the rotating frame, rather than as a true gravitational potential alone.

The utility of Ω becomes clear when its spatial derivatives are computed. The equations of motion in the rotating frame can be written directly in terms of the partial derivatives of this function, allowing the combined gravitational and centrifugal forces to be represented in compact form. Before introducing the full equations of motion, it is helpful to examine these derivatives explicitly.

Taking the partial derivatives of Ω with respect to the coordinates yields

$$\begin{aligned} \Omega_x &= x - \frac{(1 - \mu)(x + \mu)}{r_1^3} - \frac{\mu(x - 1 + \mu)}{r_2^3}, \\ \Omega_y &= y - \frac{(1 - \mu)y}{r_1^3} - \frac{\mu y}{r_2^3}, \\ \Omega_z &= -\frac{(1 - \mu)z}{r_1^3} - \frac{\mu z}{r_2^3}. \end{aligned} \quad (2.7)$$

These expressions show directly how the effective potential combines two distinct contributions. The terms proportional to x and y arise from the centrifugal potential, while the remaining terms arise from the gravitational attraction of the primaries. No centrifugal contribution appears in Ω_z because the rotation occurs in the x - y plane.

The effective potential plays a central role in the CR3BP. Its derivatives govern the equations of motion, its stationary points define the equilibrium solutions of the system, and its value appears

directly in the Jacobi constant. As a result, Ω provides a unifying quantity that links the geometry, energetics, and phase-space structure of the problem. The equations of motion derived from this potential are introduced next in Section 2.3.

2.3 Equations of Motion

The motion of the spacecraft within the CR3BP is governed by Newton's second law applied in the rotating coordinate system introduced earlier. In an inertial frame, the acceleration of the spacecraft would be determined solely by the gravitational attraction of the two primaries. However, because the motion is described in a frame that rotates with angular velocity $n = 1$, additional inertial accelerations appear in the equations of motion. These terms correspond to the Coriolis and centrifugal forces that arise in a rotating reference frame [12, 17].

Let $\mathbf{r} = (x, y, z)$ denote the position vector of the spacecraft in the rotating frame. The acceleration of the spacecraft relative to an inertial frame can be expressed as

$$\ddot{\mathbf{r}}_{inertial} = \ddot{\mathbf{r}} + 2\boldsymbol{\omega} \times \dot{\mathbf{r}} + \boldsymbol{\omega} \times (\boldsymbol{\omega} \times \mathbf{r}), \quad (2.8)$$

where $\boldsymbol{\omega}$ is the angular velocity vector of the rotating frame. The second term represents the Coriolis acceleration and the third term represents the centrifugal acceleration associated with the rotating coordinate system.

In the normalized CR3BP formulation the rotation occurs about the z -axis and the angular velocity has magnitude $n = 1$. Under these conditions the inertial forces take a particularly simple form. When combined with the gravitational forces from the two primaries, the equations governing the motion of the spacecraft can be written in component form as

$$\begin{aligned} \ddot{x} - 2\dot{y} &= x - \frac{(1-\mu)(x+\mu)}{r_1^3} - \frac{\mu(x-1+\mu)}{r_2^3}, \\ \ddot{y} + 2\dot{x} &= y - \frac{(1-\mu)y}{r_1^3} - \frac{\mu y}{r_2^3}, \\ \ddot{z} &= -\frac{(1-\mu)z}{r_1^3} - \frac{\mu z}{r_2^3}. \end{aligned} \quad (2.9)$$

Although the spatial form is written here for completeness, the analysis developed in the remainder of this thesis is restricted to planar motion, for which $z = 0$ and $\dot{z} = 0$.

These expressions describe the motion of a massless particle in the rotating frame under the gravitational influence of the two primaries. The terms containing r_1 and r_2 correspond to the gravitational attraction of the primary and secondary bodies, while the terms proportional to x and y represent the centrifugal acceleration. The Coriolis accelerations appear in the velocity-dependent terms $\pm 2\dot{x}$ and $\pm 2\dot{y}$.

Using the effective potential introduced in Section 2.2, these equations can be written in a more compact form. Because the spatial derivatives of $\Omega(x, y, z)$ contain both the gravitational and centrifugal contributions, the equations of motion become

$$\begin{aligned}
\ddot{x} - 2\dot{y} &= \frac{\partial \Omega}{\partial x}, \\
\ddot{y} + 2\dot{x} &= \frac{\partial \Omega}{\partial y}, \\
\ddot{z} &= \frac{\partial \Omega}{\partial z}.
\end{aligned} \tag{2.10}$$

This formulation highlights the central role of the effective potential in the CR3BP. The derivatives of Ω determine the forces acting on the spacecraft, while the Coriolis terms account for the rotation of the reference frame. Together these equations fully describe the motion of a massless particle in the normalized rotating coordinate system.

An important property of this dynamical system is the existence of a conserved quantity known as the Jacobi constant. This integral of motion constrains the allowed regions of motion within the system and provides valuable insight into the geometric structure of trajectories in the CR3BP. The Jacobi constant and its relationship to the effective potential are introduced in the following section.

2.4 The Jacobi Integral and Jacobi Constant

The equations of motion derived in Section 2.3 possess an important conserved quantity known as the *Jacobi integral*. This quantity plays a central role in the analysis of motion in the Circular Restricted Three-Body Problem because it constrains the regions of space that a spacecraft may access. The value of the Jacobi integral therefore provides direct insight into the geometric structure of the dynamical system [9, 12, 14, 17].

The Jacobi integral can be obtained directly from the equations of motion by considering the scalar product of the velocity vector with the dynamical equations. Starting from the rotating-frame equations

$$\begin{aligned}
\ddot{x} - 2\dot{y} &= \Omega_x, \\
\ddot{y} + 2\dot{x} &= \Omega_y, \\
\ddot{z} &= \Omega_z.
\end{aligned} \tag{2.11}$$

each equation is multiplied by the corresponding velocity component $\dot{x}$, $\dot{y}$, and $\dot{z}$, and the resulting expressions are summed. This procedure yields

$$\dot{x}\ddot{x} + \dot{y}\ddot{y} + \dot{z}\ddot{z} = \dot{x}\Omega_x + \dot{y}\Omega_y + \dot{z}\Omega_z. \tag{2.12}$$

The left-hand side can be recognized as the time derivative of the kinetic energy per unit mass,

$$\frac{d}{dt} \left(\frac{1}{2} (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) \right),$$

while the right-hand side corresponds to the time derivative of the effective potential,

$$\frac{d\Omega}{dt} = \Omega_x \dot{x} + \Omega_y \dot{y} + \Omega_z \dot{z}. \tag{2.13}$$

Combining these results leads to

$$\frac{d}{dt} \left(\frac{1}{2}(\dot{x}^2 + \dot{y}^2 + \dot{z}^2) - \Omega \right) = 0. \quad (2.14)$$

Because the time derivative of this quantity is zero, the expression inside the parentheses must remain constant along any trajectory of the system. Rearranging terms yields the classical form of the Jacobi integral,

$$2\Omega - (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) = C_J. \quad (2.15)$$

The symbol C_J represents a constant value determined by the initial conditions of the spacecraft trajectory. This quantity is commonly referred to as the *Jacobi constant*. It represents the numerical value of the Jacobi integral for a particular trajectory within the CR3BP [9, 12, 14, 17].

It is useful to distinguish between these two related concepts. The *Jacobi integral* refers to the conserved expression derived from the equations of motion,

$$2\Omega - v^2,$$

where $v^2 = \dot{x}^2 + \dot{y}^2 + \dot{z}^2$. The *Jacobi constant*, denoted by C_J , is the specific value that this expression takes for a given orbit. Because the Jacobi integral is conserved, the Jacobi constant remains fixed along the entire trajectory of the spacecraft.

Physically, the Jacobi constant can be interpreted as an energy-like quantity associated with motion in the rotating frame. Although it does not correspond exactly to the classical mechanical energy of the system, it plays an analogous role by constraining the motion of the spacecraft. In particular, the velocity magnitude must satisfy

$$v^2 = 2\Omega - C_J. \quad (2.16)$$

This relationship shows that the kinetic energy of the spacecraft depends directly on the difference between the effective potential and the Jacobi constant. Regions of space where $2\Omega < C$ are therefore dynamically inaccessible, as the resulting velocity squared would be negative [9, 12, 17].

An important feature of the CR3BP formulation is that the usual conserved quantities associated with motion in an inertial gravitational system do not generally exist in the rotating frame. In a classical two-body problem expressed in an inertial reference frame, both mechanical energy and angular momentum are conserved. In the CR3BP, however, the equations of motion are written in a synodic coordinate system that rotates with the primaries. Because this coordinate system is non-inertial, additional inertial forces appear in the equations of motion. These velocity-dependent Coriolis terms and the centrifugal contribution of the rotating frame prevent the existence of separate conserved energy and momentum integrals [9, 12, 17]. As a result, the Jacobi integral represents the only classical integral of motion available in the CR3BP within the present formulation. The existence of this single invariant strongly constrains the motion of the spacecraft and plays a central role in understanding the structure of trajectories within the system.

From a dynamical systems perspective, the Jacobi integral can also be viewed as the Hamiltonian of the CR3BP expressed in the rotating frame, up to a constant sign convention [9, 14]. In Hamiltonian mechanics the Hamiltonian represents the total energy of the system and its conservation follows from the time invariance of the governing equations. Although the CR3BP equations are written in a rotating coordinate system, they remain autonomous, and therefore a Hamiltonian invariant

still exists. The Jacobi constant therefore plays the role of the conserved Hamiltonian quantity that defines the energy surfaces of the system [9, 17].

The geometric implications of the Jacobi constant become especially clear when the velocity is set to zero. Substituting the condition $v = 0$ into the Jacobi integral yields

$$2\Omega(x, y, z) = C. \quad (2.17)$$

This equation defines a set of surfaces in space known as the *zero-velocity surfaces*. When these surfaces are projected into the plane of motion they form the well-known *zero-velocity curves*. These curves partition the configuration space into regions that are dynamically allowed or forbidden for a given value of the Jacobi constant [12, 17].

The topology of these regions changes as the value of the Jacobi constant varies, revealing pathways that connect or isolate different regions of the three-body system. These transitions occur at the critical energy levels associated with the equilibrium points of the CR3BP, where the neck regions between gravitational domains open or close. The geometry of the zero-velocity surfaces and their relationship to the equilibrium points of the CR3BP are examined in Section 2.5 [9, 14, 17].

2.5 Zero-Velocity Surfaces and Hill's Regions

The Jacobi constant provides a direct constraint on the motion of a spacecraft within the CR3BP. As shown in Section 2.4, the velocity magnitude of the spacecraft satisfies

$$v^2 = 2\Omega(x, y, z) - C. \quad (2.18)$$

Because the square of the velocity must be non-negative, this relationship imposes a fundamental restriction on the locations that the spacecraft can occupy within the rotating coordinate system. Specifically, motion is only possible in regions of space where

$$2\Omega(x, y, z) \geq C.$$

Points that violate this condition would require the velocity squared to be negative and are therefore dynamically inaccessible.

The boundary separating accessible and forbidden regions of motion occurs when the velocity is exactly zero. Setting $v = 0$ in the Jacobi integral reproduces Equation (2.17).

This equation defines a set of three-dimensional surfaces known as *zero-velocity surfaces*. These surfaces represent the locations where a particle would momentarily have zero velocity in the rotating frame. Consequently, they form the boundaries of the regions in which motion is dynamically allowed [12, 17].

In many applications the motion of interest occurs primarily within the orbital plane of the primaries, where the three-dimensional zero-velocity surfaces reduce to two-dimensional curves when intersected with the plane $z = 0$. These planar boundaries, known as *zero-velocity curves* (ZVCs), define the accessible and forbidden regions of motion within the configuration plane and are widely used to visualize the dynamical structure of the CR3BP, as illustrated in Figure 2.2. The ZVCs are obtained directly from the zero-velocity condition above by setting $z = 0$ and are therefore fully determined by the Jacobi constant and the system mass parameter μ [9, 12, 17].

The inequality $2\Omega(x, y, z) \geq C$ determines the locations where motion of the spacecraft is

dynamically possible. The portions of space that satisfy this condition define the accessible domains of motion, commonly referred to as *Hill's regions* [12, 17]. Within these regions the spacecraft possesses sufficient energy to move through the gravitational field of the primaries while remaining consistent with the Jacobi constant. Conversely, regions where $2\Omega < C$ are dynamically forbidden because the spacecraft would require negative kinetic energy to enter them.




Figure 2.2. Zero-velocity curves in configuration space for a generic CR3BP system with $\mu = 0.3$, shown for multiple Jacobi levels on a single plot. The overlaid contours illustrate how the allowed and forbidden regions evolve as the Jacobi constant changes and provide a compact view of the changing topology of Hill's regions in the rotating frame.

The geometry of Hill's regions changes as the value of the Jacobi constant varies. For large values of C_J , corresponding to lower kinetic energy, motion is confined to isolated regions surrounding each primary. As the value of C_J decreases, the zero-velocity curves deform and narrow passages begin to open between previously disconnected regions of space, as shown in Figure 2.3. These openings, which first appear near the collinear Lagrange points, allow trajectories to move between the gravitational domains of the primaries and play an important role in transport through multi-body systems [9, 17].

For a spacecraft moving solely under the gravitational influence of the two primaries, the Jacobi constant remains fixed along the trajectory and the associated zero-velocity surfaces therefore remain unchanged. As a result, the spacecraft is confined to the Hill's regions corresponding to its initial value of C_J . In practical spacecraft missions, however, propulsive maneuvers can modify the velocity of the spacecraft and thereby change the value of the Jacobi constant. Such changes effectively alter the topology of the zero-velocity curves, allowing trajectories to access regions of the system that would otherwise remain dynamically forbidden [9, 14].

The transitions between different Hill's region topologies occur at specific values of the Jacobi constant associated with equilibrium solutions of the CR3BP. These equilibrium points occur where the gradient of the effective potential vanishes,

$$\nabla\Omega = 0. \quad (2.19)$$

At these locations the gravitational and centrifugal forces acting on the spacecraft balance exactly in the rotating frame. The corresponding Jacobi constant values define the critical energy levels at

Zero-Velocity Bottleneck




Figure 2.3. A bottleneck occurs at the L1 Lagrange point of a $\mu = 0.3$ system. At lower C_J values, the bottleneck opens, allowing trajectories to pass between the two primaries.

which the zero-velocity curves open or close, revealing connections between different regions of the system. These equilibrium solutions are known as the *Lagrange points* and play a central role in the global dynamical structure of the CR3BP. Their properties and stability characteristics are examined in Section 2.6.

2.6 Lagrange Points

The equilibrium solutions of the Circular Restricted Three-Body Problem are known as the *Lagrange points*. These points correspond to locations in the rotating coordinate system where a massless particle can remain fixed relative to the two primaries. At these locations the gravitational attraction of the primaries and the centrifugal acceleration of the rotating frame balance exactly, producing zero net acceleration in the rotating reference frame [9, 12, 17].

Physically as well as mathematically, equilibrium occurs when both the velocity and acceleration of the particle vanish in the rotating frame. Applying these conditions to the equations of motion yields

$$\begin{aligned}\dot{x} = \dot{y} = \dot{z} &= 0, \\ \ddot{x} = \ddot{y} = \ddot{z} &= 0.\end{aligned}\tag{2.20}$$

Substituting these conditions into the equations of motion leads to the requirement that the spatial derivatives of the effective potential vanish,

$$\Omega_x = 0, \quad \Omega_y = 0, \quad \Omega_z = 0.\tag{2.21}$$

Equivalently, the equilibrium points occur at the stationary points of the effective potential referenced in Equation 2.19. These stationary points represent locations where the combined gravitational and centrifugal potential is locally extremized. In this sense the Lagrange points can be interpreted as

locations where the effective energy landscape of the rotating system possesses critical points.

Solving the equilibrium conditions reveals that the CR3BP possesses five such solutions, known as the Lagrange points and traditionally denoted L_1, L_2, L_3, L_4 , and L_5 [12, 17]. As shown in Figure 2.4, three of these points lie along the line connecting the two primaries and are therefore referred to as the *collinear* Lagrange points. The remaining two points form equilateral triangles with the primaries and are known as the *triangular* Lagrange points as seen in Figure 2.4. This geometric distinction reflects fundamental differences in their dynamical properties, which are examined in the following section.




Figure 2.4. The equilibria L_1, L_2 , and L_3 lie on the line connecting the primaries, while L_4 and L_5 form the triangular equilibrium configuration.

The collinear equilibrium points L_1, L_2 , and L_3 lie along the x -axis of the rotating coordinate system. The point L_1 is located between the two primaries, L_2 lies beyond the secondary body, and L_3 lies on the opposite side of the primary body. These points correspond to saddle points of the effective potential and therefore represent unstable equilibrium configurations.

The triangular equilibrium points L_4 and L_5 are located at the vertices of equilateral triangles formed with the two primaries. In the normalized coordinate system their positions are given by

$$\begin{aligned} x &= \frac{1}{2} - \mu, \\ y &= \pm \frac{\sqrt{3}}{2}. \end{aligned} \tag{2.22}$$

Unlike the collinear points, the triangular Lagrange points correspond to local extrema of the effective potential and can exhibit linear stability depending on the mass ratio of the primaries. In particular, these points are linearly stable when the mass parameter satisfies

$$\mu < \mu_c \approx 0.03852, \tag{2.23}$$

a condition known as the Routh stability criterion [12, 17].

The Lagrange points play a central role in the global dynamical structure of the CR3BP. Each equilibrium point is associated with characteristic energy levels of the Jacobi constant, and the topology of the zero-velocity curves changes as these critical values are crossed. In addition, families of periodic and quasi-periodic trajectories often exist in the vicinity of these equilibrium solutions, forming the foundation for many spacecraft mission design strategies in multi-body systems.

2.6.1 Linear Stability of the Lagrange Points

Once the equilibrium locations of the CR3BP have been determined, the next step is to evaluate their stability. This can be accomplished by examining the behavior of small perturbations to the equations of motion in the vicinity of each equilibrium point. The resulting procedure is known as *linear stability analysis* and is widely used in celestial mechanics to classify the local dynamical structure of equilibrium solutions [9, 12, 17].

Let (x_e, y_e) denote the coordinates of an equilibrium point satisfying

$$\nabla\Omega(x_e, y_e) = 0. \quad (2.24)$$

To investigate the behavior of nearby trajectories, small perturbations are introduced about the equilibrium location,

$$x = x_e + \xi, \quad y = y_e + \eta, \quad (2.25)$$

where ξ and η represent small displacements from the equilibrium point.

Substituting these expressions into the planar equations of motion and retaining only first-order terms in the perturbations yields the linearized equations governing the evolution of the perturbation variables [9, 17]:

$$\begin{aligned} \ddot{\xi} - 2\dot{\eta} &= \Omega_{xx} \xi + \Omega_{xy} \eta, \\ \ddot{\eta} + 2\dot{\xi} &= \Omega_{yx} \xi + \Omega_{yy} \eta. \end{aligned} \quad (2.26)$$

where the second derivatives of the effective potential are evaluated at the equilibrium point.

This system describes the dynamics of infinitesimal deviations from the equilibrium solution. To analyze the behavior of these perturbations, the system is written in first-order form by defining the perturbation state vector

$$\delta\mathbf{X} = [\xi \quad \eta \quad \dot{\xi} \quad \dot{\eta}]^T. \quad (2.27)$$

The linearized system can then be expressed as

$$\delta\dot{\mathbf{X}} = A \delta\mathbf{X}, \quad (2.28)$$

where A is the Jacobian matrix of the dynamical system,

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \Omega_{xx} & \Omega_{xy} & 0 & 2 \\ \Omega_{yx} & \Omega_{yy} & -2 & 0 \end{bmatrix}. \quad (2.29)$$

The stability of the equilibrium point is determined by the eigenvalues of the matrix A . If any eigenvalue has a positive real part, perturbations grow exponentially with time and the equilibrium point is unstable. If all eigenvalues are purely imaginary, perturbations remain bounded and the equilibrium point is linearly stable [9, 11, 12, 17].

For the classical CR3BP, this analysis reveals that the three collinear equilibrium points L_1 , L_2 , and L_3 are linearly unstable for all values of the mass parameter μ . The triangular equilibrium points L_4 and L_5 , however, may exhibit linear stability when the mass ratio of the primaries satisfies the Routh stability criterion Equation 2.23

a result obtained by examining the eigenvalue structure of the linearized system [12, 17].

This linear stability analysis provides a systematic method for determining the local dynamical behavior near each equilibrium point and forms the foundation for studying periodic and quasi-periodic motion in their neighborhood. These local dynamical structures play an important role in the analysis of resonant trajectories and invariant structures within the CR3BP.

2.7 Poincaré Maps

The linear stability analysis of the Lagrange points provides local information about the behavior of trajectories in the immediate vicinity of an equilibrium solution. However, many of the most important dynamical structures in the Circular Restricted Three-Body Problem (CR3BP) are fundamentally global in nature. Periodic orbits, quasi-periodic trajectories, resonant structures, separatrices, and chaotic layers all arise from the nonlinear organization of the phase space away from a single equilibrium point. In order to study these structures, it is necessary to use a tool that captures the geometry of the full flow while reducing the dimensional complexity of the system to a form that can be interpreted more directly [9–11, 19].

One of the most powerful tools for this purpose is the Poincaré map. Introduced by Henri Poincaré in his study *Les méthodes nouvelles de la mécanique céleste*, the Poincaré map transforms a continuous flow into a discrete mapping by recording the successive intersections of a trajectory with a chosen lower-dimensional surface embedded in phase space [10, 11, 20]. Rather than attempting to visualize the full continuous motion of a trajectory in a high-dimensional phase space, the method reduces the problem to the analysis of a sequence of points. This discrete representation preserves the essential dynamical structure of the system while making invariant sets, resonances, and stability transitions much easier to identify.

In the CR3BP, the Poincaré map is especially valuable because the equations of motion define a nonlinear Hamiltonian system whose trajectories are constrained by the Jacobi integral. The resulting phase space contains organized structures that govern transport and long-term behavior, but these structures are difficult to discern from configuration-space trajectories alone. A spacecraft trajectory plotted in the rotating frame may appear geometrically complicated without revealing whether the motion is periodic, quasi-periodic, resonant, or chaotic. The Poincaré map resolves this difficulty by recasting the dynamics in a reduced phase-space representation that exposes the underlying organization of the flow [9–11].

2.7.1 Motivation and Dimensional Reduction

The state of the spatial CR3BP is described by the six-dimensional vector

$$(x, y, z, \dot{x}, \dot{y}, \dot{z}),$$

while the planar CR3BP is described by the four-dimensional state

$$(x, y, \dot{x}, \dot{y}).$$

Even in the planar case, direct visualization of the full phase space is not possible. However, the existence of the Jacobi integral imposes a constraint on the motion that reduces the effective dimensionality of the system. In general, all four variables in the planar CR3BP may vary independently, and the trajectory evolves within this four-dimensional phase space.

The Jacobi constant introduces an additional relation between these variables,

$$C_J = 2\Omega(x, y) - (\dot{x}^2 + \dot{y}^2), \quad (2.30)$$

which must remain constant along any trajectory. This equation is not an equation of motion, but rather an invariant (conservation) of the motion that constrains the allowable combinations of position and velocity.

As a result, the four phase-space variables are no longer independent. For a fixed value of C_J , the state of the system must satisfy the Jacobi relation, which defines a three-dimensional hypersurface embedded within the four-dimensional phase space. The trajectory is therefore restricted to evolve on this hypersurface rather than throughout the full phase space [12, 17].

This reduction can be understood more concretely by solving the Jacobi integral for one of the velocity components. For example,

$$\dot{y} = \pm \sqrt{2\Omega(x, y) - C_J - \dot{x}^2}, \quad (2.31)$$

provided the expression under the square root is non-negative. Thus, once x , y , $\dot{x}$, and C_J are specified, the value of $\dot{y}$ is determined up to a sign. The motion is therefore confined to a three-dimensional energy manifold, often referred to as the *Jacobi surface*, within the four-dimensional phase space.

This constrained structure is fundamental to the analysis of the CR3BP, as it defines the allowable regions of motion and forms the geometric foundation for further dimensional reduction through the introduction of a surface of section [12, 17].

Although this reduction is significant, a three-dimensional phase-space manifold is still difficult to analyze directly. The key idea of the Poincaré method is to reduce the dimensionality by one additional degree through the introduction of a *surface of section*. The trajectory is not examined continuously, but only when it intersects this surface. Since the trajectory is already confined to a fixed-energy manifold, the intersection of that manifold with a chosen section produces a two-dimensional domain on which the dynamics can be represented as a discrete map [10, 11, 19].

This dimensional reduction is central to the utility of the Poincaré map. In the planar CR3BP,

$$4\text{D phase space} \longrightarrow 3\text{D energy manifold} \longrightarrow 2\text{D surface of section}.$$

The resulting map is therefore a two-dimensional representation of the dynamics on a fixed Jacobi level. This reduced representation is rich enough to preserve the essential qualitative behavior of the system, yet simple enough to visualize and interpret.

2.7.2 Definition of the Surface of Section

To construct a Poincaré map, a codimension-one surface must be chosen within the phase space. This surface is called the *surface of section*. In the planar CR3BP, a standard and particularly useful choice is the plane

$$y = 0,$$

which corresponds to the x -axis of the rotating frame. This choice is natural because the x -axis already plays a central geometric role in the problem: it contains the primaries and the collinear Lagrange points, and it partitions the planar configuration space into upper and lower half-planes [12, 17].

However, the condition $y = 0$ alone is not sufficient to define an unambiguous crossing set. A generic trajectory may cross the line $y = 0$ in two opposite directions, once with $\dot{y} > 0$ and once with $\dot{y} < 0$. If both directions are included, the same geometric crossing structure is effectively counted twice, which can obscure interpretation of the map. For this reason, a directional condition is imposed so that only one orientation of crossing is recorded. A common convention is

$$y = 0, \quad \dot{y} > 0, \tag{2.32}$$

which selects only crossings from the lower half-plane to the upper half-plane.

The surface of section may therefore be written formally as

$$\Sigma = \{(x, y, \dot{x}, \dot{y}) \in \mathbb{R}^4 \mid y = 0, \dot{y} > 0, C_J = \text{constant}\}. \tag{2.33}$$

Because the Jacobi constant constrains $\dot{y}$ through Equation (2.30), the value of $\dot{y}$ on the section is not independent once x , $y = 0$, $\dot{x}$, and C_J are specified. Thus, on a fixed energy surface, the state of an intersection can be represented using the coordinates

$$(x, \dot{x}),$$

with $\dot{y}$ determined implicitly from the Jacobi integral. This is why the Poincaré section of the planar CR3BP is usually plotted in the $(x, \dot{x})$ plane [9, 10, 21].

2.7.3 Definition of the Poincaré Map

Once a surface of section has been selected, the continuous flow of the CR3BP can be transformed into a discrete map. Let a trajectory intersect the section Σ at an initial point $X_0 \in \Sigma$. The system then evolves according to the equations of motion until the trajectory next returns to the section. The state of that next intersection defines a new point $X_1 \in \Sigma$. Repeating this process generates a sequence

$$X_0, X_1, X_2, \dots$$

of successive intersections.

The *Poincaré map* is the mapping

$$P : \Sigma \rightarrow \Sigma$$

defined by

$$X_{n+1} = P(X_n), \quad (2.34)$$

where $P(X_n)$ denotes the next return of the trajectory to the section Σ . Thus, the continuous-time flow is replaced by an iterative discrete-time dynamical system on the section [10, 11, 20].

This construction is conceptually important. The original CR3BP is an autonomous system described by differential equations, whereas the Poincaré map is an iterated map defined on a reduced state space. The map does not discard the dynamics; rather, it samples the dynamics at a dynamically meaningful set of events, namely, repeated returns to the same section. In doing so, it captures the recurrence properties of the flow and reveals invariant structures that organize the motion.

For a trajectory that repeatedly intersects the section, the Poincaré map records one point per return. The order in which these points appear corresponds to the time evolution of the trajectory. The map therefore encodes the long-term behavior of the orbit in a compact geometric form.

Poincaré-Map Construction




Figure 2.5. Construction of a Poincaré map from a representative trajectory in the rotating frame. The left panel shows the continuous trajectory crossing the section $y = 0$ with $\dot{y} > 0$, while the right panel shows the corresponding sequence of discrete points recorded in the $(x, \dot{x})$ plane at each successive crossing.

This construction is illustrated in Figure 2.5, which shows a representative trajectory in the rotating frame together with the discrete points generated by its successive crossings of the surface of section.

2.7.4 Geometric Interpretation of Map Structures

The true power of the Poincaré map lies not merely in its construction, but in the interpretation of the patterns formed by the resulting points. Different classes of motion in the continuous system

generate qualitatively distinct structures on the map. These structures provide a direct geometric diagnosis of the underlying dynamics [9–11, 19].

Periodic Orbits

A periodic orbit is a trajectory that returns exactly to its initial state after some period T . If such an orbit intersects the surface of section once per cycle, then every return occurs at the same point in Σ . On the Poincaré map, the orbit therefore appears as a single isolated point satisfying

$$P(X^*) = X^*. \quad (2.35)$$

Such a point is called a *fixed point* of the Poincaré map.

More generally, if the periodic orbit intersects the section q times before repeating, then the map produces a finite set of q repeating points,

$$X_1 \mapsto X_2 \mapsto \cdots \mapsto X_q \mapsto X_1.$$

In this case, the orbit is a period- q orbit of the map, and the points form a closed cycle under successive applications of P . Thus, periodic motion in the continuous system corresponds to discrete periodicity on the section [10, 11].

Periodic and Quasi-Periodic Poincaré Signatures




Figure 2.6. Representative Poincaré signatures of regular motion in the $(x, \dot{x})$ plane. The left panel shows the fixed-point structure produced by a periodic orbit, while the right panel shows the closed invariant curve generated by quasi-periodic motion on a surrounding torus.

Quasi-Periodic Motion and Invariant Curves

Quasi-periodic motion occurs when the trajectory evolves with two or more incommensurate frequencies and does not repeat exactly. In the planar CR3BP, such motion is typically confined to

an invariant torus within the energy manifold. The intersection of that torus with the surface of section produces a closed curve on the Poincaré map.

As the trajectory evolves, each return to the section lands at a new point on this curve. Over time, the points densely populate the curve without ever collapsing to a finite repeating set. The resulting structure is called an *invariant curve*. These curves are one of the clearest signatures of regular quasi-periodic motion and indicate that the trajectory remains confined to an ordered region of phase space [10, 11, 22].

Invariant curves are especially important because they often surround stable periodic orbits. In the neighborhood of a linearly stable periodic point, nearby trajectories do not diverge away but instead remain bounded, producing nested families of closed curves on the section. These closed curves are frequently called *stability islands* or *islands of regular motion*. Their presence indicates a region of phase space in which the dynamics are ordered and structurally persistent. The contrast between periodic and quasi-periodic signatures on the section is shown in Figure 2.6.

Resonant Island Chains

If the fundamental frequencies of a quasi-periodic trajectory satisfy a rational relationship, then the motion is resonant rather than truly quasi-periodic. In this case, the orbit closes after a finite number of returns to the section. On the Poincaré map, such a resonance appears as a chain of distinct islands surrounding a central region.

Suppose the trajectory returns to its original phase after q intersections of the section. Then the resonance appears as a period- q structure on the map, usually in the form of q islands arranged around a central stable region. Each island corresponds to one phase of the resonant orbit. These island chains are among the most important structures in nonlinear Hamiltonian systems because they signal the presence of commensurate frequency relationships and organize the local phase-space geometry [10, 11, 13].

Hyperbolic Points and Manifold Structure

Not all periodic orbits are stable. An unstable periodic orbit corresponds to a *hyperbolic fixed point* of the Poincaré map. Near such a point, trajectories are attracted along one direction and repelled along another. The local geometry is governed by the stable and unstable manifolds associated with the hyperbolic orbit.

When projected onto the section, these manifolds often form crossing patterns that divide distinct regions of motion. In nonlinear systems, the repeated stretching and folding of these manifolds can generate extremely intricate structures and produce transport between different parts of the phase space. Hyperbolic structures are therefore central to the onset of chaotic motion and to the transport mechanisms that connect otherwise separated dynamical regions [9, 11, 23]. Representative resonant and hyperbolic Poincaré structures are shown in Figure 2.7.

Chaotic Regions

Chaotic trajectories differ fundamentally from periodic and quasi-periodic motion. Rather than remaining confined to a fixed point or invariant curve, chaotic trajectories wander irregularly through an extended region of the section. On the Poincaré map, this behavior appears as a cloud or scattered layer of points that does not organize into a smooth geometric structure.

Resonant and Hyperbolic Poincaré Structures




Figure 2.7. Poincaré maps in the $(x, \dot{x})$ plane for the CR3BP on the surface of section $y = 0$ with $\dot{y} > 0$. The left panel illustrates the island-chain geometry associated with resonant motion, while the right panel shows the crossing-manifold structure of a hyperbolic trajectory together with the surrounding scattered points indicative of nearby chaotic transport.

These chaotic layers often arise near separatrices or in regions where invariant tori have broken down under nonlinear interaction. The map is especially effective at revealing such regions because even when the configuration-space trajectory appears bounded, the surface of section may show clearly that the motion lacks the ordered structure characteristic of regular dynamics [10, 19, 23]. In Figure 2.7, this behavior is visible in the scattered points surrounding the hyperbolic structure, which reflect the onset of chaotic transport in the neighboring phase-space region.

2.7.5 Why the Poincaré Map is Important in the CR3BP

The Poincaré map is not merely a visualization tool; it is a fundamental method for organizing the nonlinear dynamics of the CR3BP. By reducing the continuous flow to a discrete return map, it reveals the existence and arrangement of periodic orbits, invariant curves, resonant islands, hyperbolic structures, and chaotic layers. These are precisely the structures that govern long-term trajectory behavior and transport in multi-body dynamical systems [9, 11, 14].

For astrodynamics applications, this information is especially valuable. Regions of ordered motion may indicate families of long-lived bounded trajectories, while hyperbolic structures may indicate transport corridors between dynamical regions. The geometry of the map therefore provides both physical insight and practical guidance for identifying trajectories of interest in mission design problems [9, 21].

Equally important, the Poincaré map provides the geometric foundation for the frequency-based analysis developed in the next section. Invariant curves on the section are the geometric manifestation of quasi-periodic motion, while island chains indicate resonant frequency relationships. Thus, the structures observed on the map are not merely visual patterns; they reflect the underlying oscillatory content of the trajectories. The next section builds on this idea by introducing frequency analysis as

a complementary method for quantifying the dynamical relationships that the Poincaré map reveals geometrically.

2.8 Frequency Analysis of Trajectory Dynamics

The Poincaré map provides a geometric representation of the dynamical structure of the CR3BP by revealing invariant curves, resonant island chains, and chaotic regions. These structures reflect the qualitative organization of the phase space and provide insight into the long-term behavior of trajectories. However, while the Poincaré map identifies these structures visually, it does not directly quantify the underlying motion that produces them.

In particular, the appearance of invariant curves and island chains suggests that the motion is governed by one or more fundamental frequencies. Quasi-periodic trajectories evolve through oscillations that are not immediately apparent from configuration-space plots or even from their geometric representation on a surface of section. To characterize this motion more precisely, it is necessary to examine the time-dependent behavior of the trajectory and extract the frequencies that govern its evolution.

Frequency analysis provides a complementary perspective to the geometric framework established by the Poincaré map. Rather than focusing on where a trajectory intersects a surface, frequency-based methods analyze how the trajectory evolves in time. By decomposing a time series derived from the motion into its spectral components, it is possible to identify the dominant frequencies and determine how they combine to produce the observed dynamics [10, 13, 19, 24].

This approach is particularly powerful for quasi-periodic motion, where trajectories lie on invariant tori in phase space. On such tori, the motion can be described by a small number of incommensurate fundamental frequencies. These frequencies generate a structured spectrum consisting of a primary frequency and a set of secondary components, often referred to as sidebands, which arise from linear combinations of the fundamental frequencies. The relationships among these frequencies provide a direct link between the spectral content of the motion and the geometric structures observed in the Poincaré map [10, 13, 19, 22, 24].

The distinction between *commensurate* and *incommensurate* frequencies is central to the interpretation of this spectral structure. Two frequencies are said to be commensurate when their ratio is rational. In that case, the oscillatory components repeat with a common phase relationship after a finite number of cycles. By contrast, if the ratio is irrational, the frequencies are incommensurate, and the motion does not repeat in this way. This distinction is directly related to resonance: resonant motion occurs when the fundamental frequencies satisfy a commensurate relationship, so that the spectral behavior and the geometry of the trajectory both reflect an underlying integer frequency ratio.

In addition, frequency analysis offers a quantitative means of identifying resonances. When the fundamental frequencies satisfy a rational relationship, the motion becomes resonant, and this condition is reflected both in the spectral structure and in the formation of island chains on the Poincaré section. Thus, frequency ratios serve as an analytical tool for classifying dynamical behavior and for distinguishing between regular and chaotic motion [10, 11, 13, 19].

The goal of this section is to develop the theoretical framework for analyzing trajectory dynamics in terms of their frequency content. This includes the representation of trajectories as time series, the decomposition of these signals into spectral components, the identification of dominant and secondary frequencies, and the interpretation of frequency ratios in the context of dynamical systems.

Together, these tools provide a quantitative counterpart to the geometric insights obtained from the Poincaré map and form the foundation for the frequency-based analysis employed in subsequent sections. The frequency plots in this section are built from a synthetic signal with hard-coded frequencies $\omega_1 = 0.9$ and $\omega_2 = 0.3$, rather than from a spectrum computed from an actual CR3BP trajectory. An illustrative example of the resulting quasi-periodic time series is shown in Figure 2.8.




Figure 2.8. Illustrative quasi-periodic time series $s(t)$ exhibiting amplitude modulation produced by the superposition of multiple oscillatory components.

2.8.1 Time Series Representation of Motion

In order to analyze the frequency content of a trajectory, the continuous motion of the system must first be represented as a time-dependent signal. The equations of motion of the CR3BP define a trajectory as a function of time,

$$\mathbf{x}(t) = (x(t), y(t), z(t), \dot{x}(t), \dot{y}(t), \dot{z}(t)), \quad (2.36)$$

which describes the evolution of the system in phase space. Frequency analysis is performed not on the full state vector directly, but on a scalar or vector-valued function derived from this trajectory.

A *time series* is constructed by selecting a scalar observable $s(t)$ that depends on the system state. Common choices include individual position or velocity components, such as $x(t)$ or $\dot{x}(t)$, or combinations of state variables. The function $s(t)$ therefore represents a projection of the full phase-space trajectory onto a one-dimensional signal that evolves in time.

This representation is justified by the fact that the underlying dynamics are deterministic and governed by a finite number of degrees of freedom. As a result, the essential dynamical information contained in the trajectory is preserved in the time evolution of such observables, and their spectral content reflects the fundamental frequencies of the motion [10, 13, 19].

For the purposes of spectral analysis, the continuous signal $s(t)$ is typically evaluated at a

sequence of discrete time instances,

$$t_k = t_0 + k \Delta t, \quad k = 0, 1, 2, \dots, N-1, \quad (2.37)$$

producing a discrete time series

$$s_k = s(t_k). \quad (2.38)$$

This sampled representation forms the basis for subsequent analysis in the frequency domain. The discrete sequence $\{s_k\}$ captures the temporal evolution of the system over a finite interval and enables the decomposition of the motion into its constituent frequency components [10, 13, 19].

It is important to emphasize that the choice of observable $s(t)$ does not alter the underlying dynamics, but rather determines how the motion is represented for analysis. For quasi-periodic trajectories, any sufficiently smooth observable will contain the same fundamental frequencies, although the relative amplitudes of the spectral components may vary depending on the chosen projection [10, 13].

Thus, the representation of trajectory motion as a time series provides the essential bridge between the continuous dynamical system and its spectral characterization. In Subsection 2.8.2, this time-domain signal is decomposed into its frequency components using Fourier analysis, revealing the structure of the motion in the frequency domain.

2.8.2 Fourier Representation of the Time Series

Once the motion of the system has been represented as a time-dependent signal $s(t)$, the next step is to express this signal in terms of its underlying frequency content. This is accomplished through Fourier analysis, which provides a representation of the signal as a superposition of oscillatory components.

For the quasi-periodic signals of interest here, the time evolution can be expressed as a sum of complex exponentials of the form

$$s(t) = \sum_k A_k e^{i\omega_k t}, \quad (2.39)$$

where A_k are complex amplitudes and ω_k are angular frequencies. Each term in this expansion represents a harmonic component of the motion, and the full signal is obtained by superposing all such contributions.

In the context of dynamical systems, this representation has a direct physical interpretation. The frequencies ω_k correspond to the fundamental oscillatory modes of the motion, while the amplitudes A_k describe the relative contribution of each mode to the overall trajectory. For periodic motion, the spectrum consists of a single fundamental frequency and its harmonics. For quasi-periodic motion, the signal is composed of multiple incommensurate fundamental frequencies and their linear combinations [10, 13, 19].

More generally, a quasi-periodic signal can be written in the form

$$s(t) = \sum_{m_1, m_2, \dots, m_n} A_{m_1, m_2, \dots, m_n} e^{i(m_1\omega_1 + m_2\omega_2 + \dots + m_n\omega_n)t}, \quad (2.40)$$

where $\omega_1, \omega_2, \dots, \omega_n$ are the fundamental frequencies and $m_i \in \mathbb{Z}$ are integer coefficients. This expression shows that the full spectrum consists not only of the fundamental frequencies themselves,

but also of all integer linear combinations of those frequencies. These additional components give rise to the structured spectral patterns observed in quasi-periodic motion.

In practice, the idealized representation above assumes knowledge of the signal over an infinite time interval. However, real trajectory data is available only over a finite time span, so the Fourier representation is necessarily a finite-time approximation of the true spectral content. As a result, the identified frequencies represent approximations to the underlying dynamical frequencies, and the resolution of the spectrum is limited by the duration of the observation window.

When working with a sampled time series $\{s_k\}$, the frequency content is represented in discrete form. The continuous Fourier transform, which applies to signals defined over continuous time, is replaced by a discrete representation constructed from sampled data. In this setting, the continuous spectrum is approximated by a finite set of frequencies determined by the sampling interval and the total observation time. The resulting representation provides a practical means of identifying the dominant frequencies present in the motion and analyzing their relationships.

It is important to emphasize that the Fourier representation is not merely a mathematical transformation, but a reflection of the underlying dynamics of the system. The structure of the spectrum is determined by the geometry of the motion in phase space. In particular, trajectories confined to invariant tori generate discrete frequency spectra composed of a small number of fundamental frequencies and their combinations, while more irregular or chaotic motion produces broader and less structured spectral distributions [10, 13, 16, 24].

Thus, Fourier analysis provides a bridge between the time-domain representation of the trajectory and its dynamical interpretation. By decomposing the motion into its constituent frequencies, it becomes possible to identify the fundamental oscillatory behavior of the system and to relate this behavior directly to the geometric structures observed in the Poincaré map.

2.8.3 Dominant Frequency Identification

Once the time series representation of the motion has been expressed in the frequency domain, the next step is to identify the dominant frequencies that characterize the trajectory. These frequencies correspond to the principal oscillatory modes of the system and provide a compact description of its dynamical behavior.

For a given signal $s(t)$, the Fourier representation consists of a collection of spectral components with associated amplitudes and frequencies. The *dominant frequency* is defined as the frequency associated with the largest amplitude component in the spectrum. Physically, this frequency represents the primary oscillatory mode of the observable and typically corresponds to the most energetic or most prominent contribution to the motion. This identification is illustrated in Figure 2.9, which shows both the full Fourier spectrum and the corresponding dominant peak.

In periodic motion, the dominant frequency coincides with the fundamental frequency of the orbit, and all other spectral components appear as integer multiples (harmonics) of this value. In contrast, for quasi-periodic motion, the spectrum contains multiple fundamental frequencies that are incommensurate, meaning that no integer linear combination of these frequencies vanishes. In such cases, the dominant frequency corresponds to the spectral component with the largest amplitude, but it does not uniquely define the motion, as additional fundamental frequencies are required to fully describe the trajectory [10, 13, 24].

The identification of dominant frequencies is central to the analysis of nonlinear dynamical systems. In the case of motion on invariant tori, the dynamics can be described by a small set of

Fourier Spectrum and Dominant Frequency




Figure 2.9. Two complementary views of an illustrative quasi-periodic spectrum. The left panel shows the full Fourier spectrum with the dominant peak and neighboring combination frequencies, while the right panel isolates the dominant frequency ω_1 as the largest-amplitude spectral component.

fundamental frequencies, and the full spectrum is generated by their integer combinations. The dominant frequency therefore serves as an entry point for reconstructing the frequency structure of the system and for identifying additional spectral components associated with secondary modes of motion [10, 13, 16].

It is important to distinguish between the mathematical definition of dominance, based on spectral amplitude, and the physical interpretation of the corresponding frequency. Depending on the choice of observable $s(t)$, the dominant frequency may represent different aspects of the motion, such as radial oscillations, angular motion, or libration about a reference configuration. While the numerical value of the dominant frequency depends on the chosen projection, the underlying set of fundamental frequencies is an intrinsic property of the trajectory itself [13, 19].

Thus, the process of dominant frequency identification provides the first step in extracting meaningful dynamical information from the spectrum. By isolating the primary oscillatory component—defined by the largest amplitude spectral contribution—one obtains a reference frequency that can be used to organize the remaining spectral structure. In Subsection 2.8.4, this structure is examined in greater detail through the identification of secondary frequency components and their relationships to the dominant mode.

2.8.4 Secondary Frequencies and Sideband Structure

While the dominant frequency provides an initial characterization of the motion, a complete description of quasi-periodic dynamics requires the identification of additional frequency components. These *secondary frequencies* arise from the presence of multiple fundamental oscillatory modes and are essential for understanding the full structure of the motion.

As discussed in Subsection 2.8.3, quasi-periodic motion on an invariant torus is governed by a set of incommensurate fundamental frequencies $\omega_1, \omega_2, \dots, \omega_n$. The time-dependent signal can

therefore be expressed as a superposition of oscillatory components of the form

$$\omega = m_1\omega_1 + m_2\omega_2 + \cdots + m_n\omega_n, \quad m_i \in \mathbb{Z}. \quad (2.41)$$

This relation shows that the spectrum is not composed of isolated frequencies, but rather of a structured set of components generated by integer linear combinations of the fundamental frequencies.

In practical spectral representations, this structure manifests as a dominant peak accompanied by additional components located at frequencies offset from the primary value. These components are commonly referred to as *sidebands*. In the simplest case of two-frequency quasi-periodic motion, the spectrum contains peaks at frequencies that can be written explicitly as

$$\omega = k_1\omega_1 + k_2\omega_2, \quad k_1, k_2 \in \mathbb{Z}, \quad (2.42)$$

and, in particular, at

$$\omega = \omega_1 + m\omega_2, \quad m \in \mathbb{Z}, \quad (2.43)$$

which appear symmetrically around the dominant frequency. The spacing between these peaks is determined by the secondary frequency ω_2 , while their amplitudes reflect the coupling between the underlying modes of motion. This symmetric sideband structure is illustrated in Figure 2.10.

The presence of sidebands is a direct consequence of motion on invariant tori. Each fundamental frequency corresponds to an independent angular variable on the torus, and the observed signal represents a projection of this multi-dimensional motion onto a scalar observable. As a result, the spectral content encodes the geometry of the trajectory in phase space, with the sideband structure revealing the existence and interaction of multiple oscillatory directions [10, 13, 21, 24].

From a dynamical perspective, secondary frequencies often correspond to distinct physical modes of motion. For example, in planar CR3BP trajectories, one frequency may be associated with the primary orbital motion, while another may describe libration about a reference trajectory or oscillations transverse to a guiding path. The interaction of these modes produces the characteristic modulation patterns observed in the time series and gives rise to the sideband structure in the frequency domain [11, 21].

The identification of secondary frequencies is therefore essential for reconstructing the full set of fundamental frequencies that define the motion. While the dominant frequency provides a reference scale, it is the relationship between multiple frequencies that determines the qualitative nature of the trajectory. In particular, these relationships form the basis for defining frequency ratios and identifying resonant behavior, which are examined in the subsequent subsections.

2.8.5 Extraction of Secondary Frequencies

The identification of sideband structure in the frequency spectrum provides a means of recovering the underlying fundamental frequencies that govern the motion. While the dominant frequency establishes a primary reference scale, the secondary frequencies are obtained by analyzing the relationships among neighboring spectral components.

For quasi-periodic motion characterized by two fundamental frequencies ω_1 and ω_2 , the spectrum contains components of the form given previously in Equation (2.43). In this representation, the dominant frequency is typically associated with ω_1 , while the secondary frequency ω_2 determines the spacing between adjacent spectral peaks. Consequently, the difference between neighboring




Figure 2.10. Symmetric sideband peaks about the dominant frequency ω_1 are separated by the secondary frequency ω_2 .

components in the spectrum provides a direct measure of the secondary frequency,

$$\omega_2 = \omega_{k+1} - \omega_k. \quad (2.44)$$

This relationship reflects the underlying structure of the quasi-periodic expansion and does not depend on the specific observable used to construct the time series. Although the amplitudes of the spectral components may vary with the chosen projection, the spacing between peaks is determined solely by the fundamental frequencies of the motion and therefore represents an intrinsic property of the trajectory [10, 13, 24]. The recovery of the secondary frequency from sideband spacing is shown in Figure 2.11.

More generally, when multiple fundamental frequencies are present, the spectrum consists of a lattice of components generated by integer combinations of these frequencies. In such cases, secondary frequencies can be identified by examining consistent spacing patterns within the spectrum and by recognizing families of peaks that share a common linear relationship. This structure reflects the geometry of the invariant torus on which the motion evolves, with each fundamental frequency corresponding to an independent direction on the torus [10, 11].

In practical terms, the extraction of secondary frequencies relies on identifying coherent patterns in the spectral distribution rather than isolated peaks. The presence of symmetric sidebands about a dominant frequency, along with regular spacing between components, provides strong evidence for an underlying two-frequency quasi-periodic structure. These features allow the secondary frequency to be determined as a characteristic separation between spectral components, even when individual peaks vary in amplitude [13, 21].

The recovery of secondary frequencies is a critical step in reconstructing the fundamental frequency set of the motion. Once both the dominant and secondary frequencies have been identified, their relationship can be quantified through a frequency ratio, which provides a concise measure of the dynamical behavior. This ratio forms the basis for identifying resonant motion and classifying

trajectories according to their spectral properties, as discussed in Subsection 2.8.6.

Secondary-Frequency Extraction




Figure 2.11. The secondary frequency ω_2 is recovered from the spacing of symmetric sidebands about the dominant peak.

2.8.6 Frequency Ratio and Resonance Interpretation

Once the fundamental frequencies of the motion have been identified, their relationship can be expressed through a dimensionless quantity known as the *frequency ratio*. For motion characterized by two primary frequencies ω_1 and ω_2 , the frequency ratio is defined as

$$\nu = \frac{\omega_2}{\omega_1}. \quad (2.45)$$

This ratio provides a compact measure of the relative timescales of the underlying oscillatory modes and serves as a key analytical measure for classifying dynamical behavior.

The value of ν determines whether the motion is quasi-periodic or resonant. If the ratio is irrational, the frequencies are incommensurate, and the trajectory densely fills an invariant torus in phase space. In this case, the motion is quasi-periodic, and the corresponding Poincaré map consists of smooth invariant curves. The trajectory never exactly repeats, but instead explores the torus in a structured and non-repeating manner [10, 11, 19].

In contrast, when the frequency ratio is rational,

$$\nu = \frac{p}{q}, \quad p, q \in \mathbb{Z}, \quad (2.46)$$

the frequencies satisfy a commensurability condition,

$$q \omega_2 - p \omega_1 = 0. \quad (2.47)$$

Under this condition, the motion becomes *resonant*, and the trajectory is periodic in a rotating frame

defined by the fundamental frequencies. Rather than densely filling a torus, the motion is confined to a closed orbit, and this change is reflected in both the spectral structure and the geometry of the Poincaré map [10, 11].

Resonant motion gives rise to characteristic structures on the Poincaré section known as *island chains*. A resonance of order $p : q$ produces q distinct islands on the section, corresponding to the q discrete phases of the periodic orbit. These islands represent stable periodic orbits and their surrounding invariant curves, while the regions between them are typically associated with more complex or chaotic behavior [9–11].

From the perspective of frequency analysis, resonance is reflected in the structure of the spectrum. When the frequencies satisfy a rational relationship, the spectral components collapse onto a set of harmonically related peaks, and the sideband structure becomes organized according to the resonance condition. The frequency ratio therefore provides a direct link between the spectral representation of the motion and its geometric manifestation in phase space [13, 21].

In practice, the frequency ratio serves as a powerful tool for mapping the dynamical structure of the system. By evaluating ν for a family of trajectories, it is possible to identify regions of regular motion, locate resonances, and detect transitions to chaotic behavior. Smooth variations in ν are typically associated with invariant tori, while abrupt changes or discontinuities may indicate the presence of resonances or the breakdown of regular motion [13, 21, 24]. The role of low-order rational values in organizing this structure is illustrated schematically in Figure 2.12.

Thus, the frequency ratio provides a unifying framework for interpreting the dynamics of the CR3BP. It connects the time-domain behavior of trajectories, the spectral structure of their frequency content, and the geometric organization revealed by the Poincaré map. This relationship forms the foundation for the analysis of dynamical structures and stability in the system and provides the basis for the computational methods described in Chapter 3.




Figure 2.12. Representative low-order rational values are highlighted together with the resonant condition $\omega_2/\omega_1 = p/q$. The bar heights are illustrative and indicate the relative dynamical significance of the low-order resonances.

2.9 Stability Analysis of Periodic Orbits

The identification of periodic and quasi-periodic structures in the circular restricted three-body problem provides important information about the organization of phase space. However, the existence of a periodic orbit alone does not determine whether it is dynamically useful or physically significant. A trajectory may satisfy the periodicity condition exactly and yet be highly sensitive to small perturbations, so that nearby trajectories diverge rapidly from it. For this reason, the analysis of *stability* is a necessary complement to the geometric and spectral methods introduced in Section 2.7 and Section 2.8 [9–11, 14].

In dynamical systems, stability concerns the response of a reference solution to small deviations in its initial state. If a trajectory that begins near a periodic orbit remains near that orbit for all subsequent time, the orbit is regarded as stable in the local sense. If instead small perturbations grow with time, the orbit is unstable. In the CR3BP, this question is especially important because periodic orbits often serve as the organizing centers of nearby phase-space structure. Stable periodic orbits are typically surrounded by invariant tori and regions of ordered motion, whereas unstable periodic orbits are associated with hyperbolic structures and invariant manifolds that govern transport through the system [10, 11, 14, 21].

The stability of periodic orbits is commonly assessed through *linear stability analysis*. In this approach, the full nonlinear equations of motion are expanded about a reference trajectory, and the evolution of an infinitesimal perturbation is studied. This procedure does not determine the global behavior of the nonlinear system, but it does provide a local characterization of how nearby trajectories behave in the immediate neighborhood of the periodic solution. The resulting framework leads naturally to the variational equations, the state transition matrix, and the Floquet multipliers that are used to classify the stability of periodic motion [9, 11, 14].

2.9.1 Perturbation and Linearization

Let the equations of motion of the dynamical system be written in first-order form as

$$\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}), \quad (2.48)$$

where $\mathbf{x}(t)$ denotes the state vector and $\mathbf{f}$ is the vector field defining the flow. Suppose that $\mathbf{x}(t)$ is a known reference solution of these equations. To examine the stability of this solution, a nearby trajectory is introduced in the form

$$\mathbf{x}(t) + \delta\mathbf{x}(t), \quad (2.49)$$

where $\delta\mathbf{x}(t)$ is an infinitesimal perturbation to the reference state.

Substituting the perturbed trajectory into the equations of motion yields

$$\frac{d}{dt}(\mathbf{x} + \delta\mathbf{x}) = \mathbf{f}(\mathbf{x} + \delta\mathbf{x}). \quad (2.50)$$

The vector field on the right-hand side may then be expanded in a first-order Taylor series about the reference solution:

$$\mathbf{f}(\mathbf{x} + \delta\mathbf{x}) = \mathbf{f}(\mathbf{x}) + \left. \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \right|_{\mathbf{x}(t)} \delta\mathbf{x} + O(\|\delta\mathbf{x}\|^2). \quad (2.51)$$

Since $\mathbf{x}(t)$ is a solution of the equations of motion, the zeroth-order terms cancel identically.

Neglecting second-order and higher-order perturbation terms then yields the *variational equations*,

$$\delta \dot{\mathbf{x}}(t) = \mathbf{A}(t) \delta \mathbf{x}(t), \quad (2.52)$$

where

$$\mathbf{A}(t) = \left. \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \right|_{\mathbf{x}(t)} \quad (2.53)$$

is the Jacobian matrix of the vector field evaluated along the reference trajectory [11, 14, 19].

This linear system governs the evolution of infinitesimal deviations from the reference solution. Its coefficients are generally time-dependent because the Jacobian is evaluated along the trajectory itself rather than at a fixed point in phase space. As a result, even when the original dynamical system is autonomous, the perturbation dynamics along a non-equilibrium solution are typically nonautonomous [10, 11].

The physical meaning of this linearization is important. The variational equations do not replace the full nonlinear equations of motion, nor do they describe the evolution of large deviations from the reference orbit. Instead, they provide a *local tangent-space approximation* to the flow, valid only in the limit of sufficiently small perturbations. Within this approximation, the growth, decay, or bounded evolution of $\delta \mathbf{x}(t)$ provides a first assessment of the local stability of the reference solution [9, 11, 14].

When the reference solution is a periodic orbit, the variational equations become especially useful because they allow the perturbation to be tracked over one complete period of the motion. The resulting mapping of an initial perturbation to its value after one cycle forms the basis for the monodromy matrix and the Floquet multipliers, which provide the standard tools for classifying the stability of periodic orbits in the CR3BP [9, 14, 21].

2.9.2 State Transition Matrix

The linear system defined by the variational equations admits a fundamental matrix solution that maps perturbations between any two points along the reference trajectory. This mapping is represented by the *state transition matrix* (STM), denoted $\Phi(t, t_0)$, which relates an initial perturbation at time t_0 to its value at a later time t according to

$$\delta \mathbf{x}(t) = \Phi(t, t_0) \delta \mathbf{x}(t_0). \quad (2.54)$$

The state transition matrix is defined as the solution to the matrix differential equation

$$\frac{d}{dt} \Phi(t, t_0) = \mathbf{A}(t) \Phi(t, t_0), \quad \Phi(t_0, t_0) = \mathbf{I}, \quad (2.55)$$

where $\mathbf{A}(t)$ is the Jacobian of the vector field evaluated along the reference trajectory, and $\mathbf{I}$ is the identity matrix [9, 11, 14].

The STM provides a complete linear description of the local flow in the neighborhood of the reference solution. Each column of $\Phi(t, t_0)$ corresponds to the evolution of a basis perturbation direction, and together these columns describe how an infinitesimal neighborhood of initial conditions is deformed by the dynamics over time. In this sense, the STM defines the linearized mapping of phase space in the tangent bundle of the trajectory [10, 11].

Several important properties follow directly from its definition. First, the STM satisfies the composition rule

$$\Phi(t_2, t_0) = \Phi(t_2, t_1) \Phi(t_1, t_0), \quad (2.56)$$

which reflects the sequential evolution of perturbations. Second, the inverse mapping is given by

$$\Phi^{-1}(t, t_0) = \Phi(t_0, t), \quad (2.57)$$

so that the STM is nonsingular for all times along a well-defined trajectory. These properties emphasize that the STM acts as a linear propagator of perturbations along the flow [9, 14].

The significance of the STM for stability analysis lies in its ability to quantify the growth or contraction of perturbations. By examining how $\delta \mathbf{x}(t)$ evolves under the action of $\Phi(t, t_0)$, one obtains a local measure of sensitivity to initial conditions. When applied over a finite time interval, and in particular over one period of a periodic orbit, the STM forms the basis for defining the monodromy matrix and the associated Floquet multipliers used in stability classification [14, 21].

2.9.3 Periodic Orbits and the Monodromy Matrix

When the reference solution $\mathbf{x}(t)$ is a periodic orbit of period T , the linearized dynamics exhibit a corresponding periodic structure. In this case, the state transition matrix evaluated over one full period,

$$\mathbf{M} = \Phi(T, 0), \quad (2.58)$$

is referred to as the *monodromy matrix*. This matrix defines a linear mapping from an initial perturbation at time $t_0 = 0$ to its value after one complete cycle of the orbit [9, 11, 14].

The monodromy matrix therefore represents a discrete-time description of the continuous perturbation dynamics. While the variational equations describe the instantaneous evolution of perturbations along the trajectory, the monodromy matrix captures the cumulative effect of this evolution over one full period. In this sense, it provides a natural bridge between continuous-time dynamics and a stroboscopic mapping defined at intervals of the orbital period [10, 11].

Because the reference orbit is periodic, repeated applications of the monodromy matrix describe the evolution of perturbations over successive cycles:

$$\delta \mathbf{x}(nT) = \mathbf{M}^n \delta \mathbf{x}(0), \quad n \in \mathbb{Z}. \quad (2.59)$$

This formulation highlights the role of the monodromy matrix as the fundamental linear operator governing the long-term behavior of small deviations from the periodic orbit.

The eigenvalues of $\mathbf{M}$, known as *Floquet multipliers*, determine the stability characteristics of the orbit. Each multiplier describes the growth, decay, or bounded evolution of perturbations along a corresponding eigendirection over one period. If the magnitude of a multiplier exceeds unity, perturbations grow from one cycle to the next, indicating instability. If all multipliers lie on the unit circle, perturbations remain bounded and the motion is locally stable in the linear sense [9, 14, 21].

In Hamiltonian systems such as the CR3BP, the structure of the monodromy matrix imposes additional constraints on its eigenvalues. In particular, the multipliers occur in reciprocal pairs $(\lambda, 1/\lambda)$, and complex eigenvalues appear in conjugate pairs. This symmetry reflects the underlying conservation properties of the system and restricts the possible stability types of periodic motion [10, 11].

Thus, the monodromy matrix provides a compact and complete description of the linear stability of periodic orbits. By reducing the continuous perturbation dynamics to a finite-dimensional linear mapping over one period, it enables the classification of orbits based on the spectral properties of a single matrix, forming the foundation for the stability criteria developed in Subsection 2.9.4.

2.9.4 Floquet Multipliers and Stability Classification

Let λ_i denote the eigenvalues of the monodromy matrix $\mathbf{M}$. These eigenvalues, known as *Floquet multipliers*, characterize the evolution of infinitesimal perturbations over one period of the orbit. The magnitude of each multiplier determines whether perturbations grow, decay, or remain bounded under repeated application of the monodromy mapping [9, 11, 14].

The stability of the periodic orbit is therefore determined by the spectral properties of $\mathbf{M}$. If $|\lambda_i| = 1$ for all i , perturbations remain bounded for all time and the orbit is said to be *linearly stable*. If any $|\lambda_i| > 1$, perturbations grow exponentially with successive iterations and the orbit is *unstable*.

It is important to emphasize that this classification is based on the linearized dynamics. While linear stability implies that nearby trajectories remain close to the periodic orbit for sufficiently small perturbations, it does not guarantee nonlinear stability over arbitrarily long times. Nevertheless, it provides a fundamental first-order characterization of the local behavior of the system [10, 19].

In Hamiltonian systems such as the CR3BP, the monodromy matrix possesses a symplectic structure that imposes strong constraints on its eigenvalues. Specifically, the Floquet multipliers occur in reciprocal pairs $(\lambda, 1/\lambda)$, and complex eigenvalues appear in conjugate pairs. As a result, eigenvalues typically appear in sets of four when off the real axis, reflecting the underlying conservation properties of the system [10, 11].

Based on the distribution of the Floquet multipliers in the complex plane, periodic orbits are commonly classified into distinct stability types. In the *elliptic* case, all multipliers lie on the unit circle, so perturbations remain bounded and the motion is locally associated with invariant tori. In the *hyperbolic* case, at least one pair of multipliers lies off the unit circle on the real axis, leading to exponential divergence along the corresponding eigendirections. In the *complex unstable* case, the multipliers form complex conjugate pairs off the unit circle, producing a combination of oscillatory and exponentially growing behavior.

These stability types are directly related to the geometric structures discussed in Section 2.7. Elliptic periodic orbits are typically surrounded by families of invariant curves on the Poincaré map, while hyperbolic periodic orbits are associated with stable and unstable manifolds that organize transport through phase space. Thus, the spectral properties of the monodromy matrix provide a direct link between local linear behavior and the global dynamical structures observed in the CR3BP [9, 11, 14, 21].

2.9.5 Stability Indices

For planar periodic orbits, the spectral information contained in the Floquet multipliers can be reduced to scalar quantities that provide a convenient summary of stability properties. One such measure is the *stability index*, defined for a reciprocal pair of multipliers $(\lambda, 1/\lambda)$ as

$$s = \lambda + \frac{1}{\lambda}. \quad (2.60)$$

This definition arises naturally from the structure of Hamiltonian systems, in which eigenvalues occur in reciprocal pairs. Rather than analyzing the individual multipliers directly, the stability index combines each pair into a single real-valued quantity that preserves the essential information regarding growth or boundedness of perturbations [10, 11].

The interpretation of s follows from the location of the eigenvalues in the complex plane. If the multipliers lie on the unit circle, they may be written in the form $\lambda = e^{i\theta}$, in which case

$$s = e^{i\theta} + e^{-i\theta} = 2 \cos \theta, \quad (2.61)$$

so that s is real and satisfies $|s| \leq 2$. This corresponds to bounded, oscillatory perturbations and therefore linear stability. In contrast, if the multipliers lie off the unit circle, they are real or complex with magnitude different from unity, and the resulting value of s satisfies $|s| > 2$, indicating exponential growth or decay of perturbations [9, 10, 14].

Thus, the stability index provides a simple classification criterion: $|s| < 2$ corresponds to elliptic, linearly stable behavior with bounded perturbations, whereas $|s| > 2$ corresponds to hyperbolic, unstable behavior with exponential divergence.

In practical analysis, stability indices are particularly useful because they reduce the multidimensional eigenvalue problem to a small set of scalar quantities that can be tracked as system parameters vary. This makes them well-suited for studying families of periodic orbits and identifying bifurcations, where transitions in stability occur as the index passes through the critical values $s = \pm 2$ [9, 11, 14].

From a dynamical perspective, these transitions are associated with qualitative changes in the local phase-space structure. Stable periodic orbits, characterized by $|s| < 2$, are typically surrounded by invariant tori and regular motion, while unstable orbits with $|s| > 2$ give rise to hyperbolic structures and invariant manifolds that organize transport. In this way, the stability index provides a compact link between the spectral properties of the monodromy matrix and the geometric structures observed in the CR3BP [9, 11, 21].

2.9.6 Interpretation in Phase Space

The stability properties of periodic orbits have a direct geometric manifestation in phase space. While the preceding analysis characterizes stability through the spectral properties of the monodromy matrix, these results may also be interpreted in terms of the local and global structure of the flow.

In the vicinity of an elliptic (linearly stable) periodic orbit, perturbations remain bounded, and the surrounding phase space is typically foliated by invariant tori. Trajectories on these tori exhibit quasi-periodic motion and appear as closed invariant curves on a Poincaré map. This structure corresponds to regions of ordered dynamics, where motion is regular and confined [9–11].

In contrast, hyperbolic (unstable) periodic orbits are associated with directions of exponential growth and decay in the perturbation dynamics. These directions give rise to stable and unstable invariant manifolds, which extend away from the periodic orbit and form the backbone of phase-space transport. Intersections of these manifolds generate complex structures, including homoclinic and heteroclinic connections, that are fundamental to chaotic dynamics and govern the movement of trajectories between different regions of the system [9, 11, 14].

From this perspective, stability analysis provides a bridge between local linear behavior and global nonlinear dynamics. The eigenvalues of the monodromy matrix determine how perturbations

evolve near a periodic orbit, while the associated geometric structures—such as invariant tori and manifolds—determine how trajectories evolve over longer timescales and across larger regions of phase space [10, 11].

This connection is central to the dynamical interpretation of the CR3BP. The structures identified through stability analysis correspond directly to the features observed in the Poincaré map, where invariant curves indicate stable motion and separatrices and chaotic layers reflect the influence of hyperbolic dynamics. Moreover, these same structures are encoded in the spectral properties of the motion, as discussed in the frequency analysis section. In this way, stability analysis, Poincaré maps, and frequency methods together provide a unified framework for understanding the organization of trajectories in multi-body dynamical systems [9, 14, 21].

2.9.7 Summary of Dynamical Methods

This chapter has established the theoretical framework used to analyze motion in the circular restricted three-body problem. Beginning with the formulation of the equations of motion and the Jacobi integral, the discussion introduced the geometric constraints that define the allowable regions of motion. These constraints provide the foundation for understanding how trajectories evolve within the system.

The Poincaré map was then presented as a method for reducing the continuous dynamics to a discrete representation, revealing the underlying phase-space structure in terms of invariant curves, resonant islands, and chaotic regions. Complementing this geometric perspective, Fourier-based frequency analysis was introduced as a spectral tool for identifying the fundamental oscillatory components of the motion and for quantifying resonant relationships through frequency ratios [10, 13, 21].

Finally, stability analysis was developed through the linearization of the equations of motion, leading to the variational equations, state transition matrix, and monodromy matrix. The resulting Floquet multipliers and stability indices provide a means of classifying periodic orbits and linking local linear behavior to the global phase-space structures observed in the system [9, 11, 14].

Together, these methods form a unified framework for analyzing the dynamics of the CR3BP. The geometric structures observed in Poincaré maps, the spectral signatures identified through frequency analysis, and the stability properties determined from linearization are interconnected representations of the same underlying dynamical behavior.

In Chapter 3, this theoretical framework is carried into the Neptune–Triton problem in explicit computational form. There, the governing models, normalization choices, trajectory-generation procedures, and analysis workflows used in this study are defined in detail, providing the methodological basis for the results presented in the subsequent chapters.

3. Dynamical Models and Numerical Methods

This chapter translates the theoretical framework developed in Chapter 2 into the explicit dynamical models and numerical procedures used to generate the trajectory data analyzed throughout this thesis. While the previous chapter established the CR3BP concepts and analysis methods in general form, the present chapter specifies the exact mathematical formulations and computational procedures implemented for the Neptune–Triton system.

In particular, this chapter establishes the governing equations of motion, the normalization and parameterization of the Neptune–Triton system, and the numerical techniques used for trajectory propagation, event detection, and data generation. Together, these elements define the computational framework used to construct Poincaré maps, perform frequency analysis, and evaluate trajectory stability in the chapters that follow.

The objective of this chapter is therefore twofold: first, to clearly document the precise dynamical models used in this work, including both the baseline CR3BP formulation and its J_2 -augmented extension; and second, to provide sufficient implementation detail to ensure that all numerical results presented in this thesis are fully reproducible.

3.1 Reference Frames and Normalization

The dynamical models and trajectory integrations in this study are formulated in the rotating barycentric reference frame established in Subsection 2.1.1. In that frame, the two primaries remain fixed along the x -axis, eliminating explicit time dependence in their motion and allowing spacecraft accelerations to be evaluated directly during numerical integration. The canonical CR3BP normalization developed in Subsection 2.1.1 is adopted here without modification. The present section specifies the system-specific numerical values and dimensional constants that instantiate that normalization for the Neptune–Triton system used throughout this thesis.

The normalized CR3BP is fully characterized by the single dimensionless mass ratio defined in Equation (2.2). For the present application, m_1 corresponds to Neptune and m_2 to Triton, yielding

$$\mu = 2.09898 \times 10^{-4}.$$

This small value reflects the strongly asymmetric mass distribution of the Neptune–Triton subsystem: Neptune dominates the gravitational field while Triton constitutes a comparatively low-mass but dynamically significant secondary. The degree of this asymmetry is directly relevant to the phase-space structure examined in the later chapters, particularly in the neighborhoods of the collinear bottlenecks, the large Neptune-centered regions, and the region local to Triton.

The characteristic length scale for nondimensionalization is the Neptune–Triton separation. As established by the normalization in Subsection 2.1.1, the total mass and mean motion are set to unity in Equation (2.1). In the present chapter, the normalized Neptune–Triton separation is denoted by r_{NT} , so the corresponding length normalization is

$$r_{NT} = 1. \tag{3.1}$$

All coordinates and initial conditions therefore represent fractions of the physical separation rather than dimensional distances. The physical value of the Neptune–Triton separation used in this study

is approximately 354,759 km, which defines the length scale for dimensional interpretation of all trajectory products. Under this normalization the equations of motion take the standard CR3BP form without explicit appearance of the gravitational constant or dimensional orbital parameters. The dimensional conversion factors corresponding to these choices are collected in Appendix A.

Within the rotating barycentric frame, the positions of the two primaries follow directly from the mass parameter. As shown in Subsection 2.1.1, Neptune and Triton are located at

$$x_1 = -\mu, \quad x_2 = 1 - \mu, \quad (3.2)$$

placing Neptune slightly to the left of the barycenter and Triton slightly to the right. These positions remain constant for the duration of all integrations and serve as the fixed geometric reference points in every trajectory propagation, zero-velocity computation, Poincaré map, and frequency-analysis product generated in this work.

This normalized rotating-frame formulation defines the common coordinate system shared by both the baseline CR3BP model in Section 3.2 and the J_2 -augmented model in Section 3.3. Maintaining a unified nondimensional framework for both models ensures that differences in computed results are attributable to the dynamical perturbation rather than to any inconsistency in units or reference frame.

3.2 Planar CR3BP Equations

This section defines the baseline dynamical model used for trajectory propagation in this study. The planar Circular Restricted Three-Body Problem (CR3BP) is adopted as the unperturbed model governing spacecraft motion in the Neptune–Triton system. The governing equations were developed in Section 2.2 and Section 2.3; the present section states the planar forms of those equations as implemented in the numerical integrator and describes how they are used to generate the trajectory data analyzed throughout this thesis.

In the rotating barycentric frame and canonical normalization described in Section 3.1, the planar motion of the spacecraft is governed by the equations of motion derived in Section 2.3, specifically Equation (2.9) reduced to the plane:

$$\ddot{x} - 2\dot{y} = \Omega_x(x, y), \quad \ddot{y} + 2\dot{x} = \Omega_y(x, y).$$

The effective potential $\Omega(x, y)$ appearing in these equations is the planar form of Equation (2.6) from Section 2.2,

$$\Omega(x, y) = \frac{1 - \mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2} (x^2 + y^2),$$

where the distances from the spacecraft to Neptune and Triton are obtained by setting $z = 0$ in Equations (2.3) of Section 2.2,

$$r_1 = \sqrt{(x + \mu)^2 + y^2}, \quad r_2 = \sqrt{(x - (1 - \mu))^2 + y^2}.$$

The corresponding partial derivatives, which are the planar reductions of Equation (2.7) from

Section 2.2, are

$$\Omega_x = x - \frac{(1-\mu)(x+\mu)}{r_1^3} - \frac{\mu(x-(1-\mu))}{r_2^3}, \quad \Omega_y = y - \frac{(1-\mu)y}{r_1^3} - \frac{\mu y}{r_2^3}.$$

These relations constitute the complete set of working expressions evaluated at each integration step.

For numerical integration, the second-order system must be recast into first-order form. This is an implementation requirement of the adaptive ODE solvers used for trajectory propagation, surface-of-section event detection, and time-history extraction. The state vector is defined as

$$\mathbf{X} = [x \ y \ \dot{x} \ \dot{y}]^T, \quad (3.3)$$

and the first-order system governing its time evolution is

$$\dot{\mathbf{X}} = \begin{bmatrix} \dot{x} \\ \dot{y} \\ 2\dot{y} + \Omega_x(x, y) \\ -2\dot{x} + \Omega_y(x, y) \end{bmatrix}. \quad (3.4)$$

Equations (3.3) and (3.4) define the ordinary differential equation system passed directly to the numerical integrator and form the computational entry point for all trajectory propagation carried out in this work.

The Jacobi constant provides the energy-level parameterization used throughout this work to organize and compare trajectory families. The general form was derived in Section 2.4 as Equation (2.15); restricted to planar motion with $z = \dot{z} = 0$, it reduces to Equation (2.30). For a prescribed value of C_J , the integral constrains the admissible initial conditions and determines the energetically accessible regions of the rotating frame, as established through the zero-velocity analysis of Section 2.5.

Initial conditions are constructed procedurally on a prescribed energy manifold rather than sampled arbitrarily from the full four-dimensional phase space. Positions are selected along the section $y = 0$, and the remaining velocity component is recovered directly from Equation (2.30). When x , $y = 0$, $\dot{x}$, and C_J are specified, the Jacobi constraint yields

$$\dot{y} = \pm \sqrt{2\Omega(x, 0) - C_J - \dot{x}^2},$$

where only states with a nonnegative radicand are retained. This admissibility condition is a direct consequence of the velocity constraint established in Section 2.5, ensuring that every generated initial condition lies within a dynamically accessible region of phase space. The sign of $\dot{y}$ is selected according to the orientation convention required for the relevant trajectory set or surface-of-section analysis.

Constructing initial conditions in this manner ties the numerical data generation directly to the geometric framework of Chapter 2. All trajectory families share a common energy surface and are propagated under a common dynamical model, which ensures that subsequent comparisons across Poincaré maps, frequency spectra, and stability computations are made within a controlled and explicitly defined region of phase space rather than across an arbitrarily sampled distribution of

states.

The planar CR3BP equations presented in this section constitute the baseline, unperturbed dynamical model used throughout this thesis. All trajectory propagation, Poincaré map generation, and frequency analysis are based on this formulation unless otherwise stated. Section 3.3 extends this model by incorporating the J_2 oblateness perturbation of Neptune, enabling examination of how planetary flattening modifies phase-space structure and trajectory stability.

3.3 J_2 -Augmented Model

This section defines the perturbed dynamical model used in this study by extending the baseline CR3BP formulation to include the effect of Neptune’s oblateness. Specifically, the second zonal harmonic coefficient J_2 is incorporated into the equations of motion as a correction to the gravitational potential. This augmented model is implemented alongside the classical CR3BP equations to quantify how planetary flattening modifies trajectory behavior, phase-space structure, and stability characteristics.

In addition to the classical planar CR3BP, this study employs a J_2 -augmented formulation that preserves the rotating-frame structure described in Chapter 2 while introducing the leading-order deviation from spherical symmetry in Neptune’s gravitational field [12,25]. The equations presented in this section define the exact perturbed system of differential equations used in the numerical simulations.

The baseline effective potential is the CR3BP effective potential of Equation (2.6), with r_1 and r_2 defined by Equations (2.3). In the present section this Chapter 2 relation is carried forward into the J_2 -augmented model used for Neptune–Triton trajectory propagation.

The oblateness correction is introduced as an additional body-centered contribution to the effective potential associated with Neptune. In nondimensional form, this term is written as

$$\Omega_{J_2}^{(N)} = -\frac{(1-\mu) J_2 R_{\text{bar}}^2}{2} \left(r_1^{-3} - 3z^2 r_1^{-5} \right), \quad (3.5)$$

where J_2 is Neptune’s second zonal harmonic coefficient, $R_{\text{bar}} = R_N/a$ is the nondimensional equatorial radius of Neptune, and r_1 is the spacecraft distance from Neptune. This formulation is consistent with the canonical CR3BP scaling, in which all lengths are normalized by the primary–secondary separation a [25].

The augmented effective potential used in this study is therefore

$$\Omega_{\text{aug}}(x, y, z) = \Omega(x, y, z) + \Omega_{J_2}^{(N)}. \quad (3.6)$$

Triton is modeled as a point mass, and its oblateness contribution is neglected due to its comparatively small effect.

The corresponding oblateness acceleration is obtained from the gradient of the augmented potential. Defining

$$k = \frac{1}{2}(1-\mu) J_2 R_{\text{bar}}^2, \quad (3.7)$$

the J_2 acceleration components for the Neptune-centered position vector are

$$\begin{aligned} a_{J_2,x} &= \frac{3k(x+\mu)}{r_1^5} - \frac{15k(x+\mu)z^2}{r_1^7}, \\ a_{J_2,y} &= \frac{3ky}{r_1^5} - \frac{15kyz^2}{r_1^7}, \\ a_{J_2,z} &= \frac{9kz}{r_1^5} - \frac{15kz^3}{r_1^7}. \end{aligned} \tag{3.8}$$

These terms are added directly to the classical CR3BP accelerations in the rotating frame, yielding the perturbed equations of motion implemented in the numerical integrator. For the planar case considered in this study, $z = 0$, and the expressions reduce to

$$\begin{aligned} a_{J_2,x} &= \frac{3k(x+\mu)}{r_1^5}, \\ a_{J_2,y} &= \frac{3ky}{r_1^5}, \\ a_{J_2,z} &= 0. \end{aligned} \tag{3.9}$$

The resulting J_2 -augmented equations of motion are therefore

$$\begin{aligned} \ddot{x} - 2\dot{y} &= \Omega_x + a_{J_2,x}, \\ \ddot{y} + 2\dot{x} &= \Omega_y + a_{J_2,y}, \\ \ddot{z} &= \Omega_z + a_{J_2,z}. \end{aligned} \tag{3.10}$$

In this formulation, the J_2 contribution acts as a perturbative modification to the baseline CR3BP dynamics, enabling direct comparison between unperturbed and perturbed trajectory behavior under identical initial conditions and energy levels.

Although the primary trajectory analysis in this thesis is restricted to planar motion, the numerical implementation retains the full three-dimensional formulation of the J_2 model. This ensures internal consistency of the dynamical system and allows the same equations to be used for future three-dimensional extensions without modification.

This augmented formulation defines the perturbed dynamical model used throughout all Neptune–Triton simulations presented in this work. All results labeled as “perturbed” or “ J_2 ” in subsequent chapters are generated using the equations defined in this section, while comparisons to the classical CR3BP are based on the baseline model defined in Section 3.2.

3.4 Equilibrium Structure

Before generating trajectory families and constructing Poincaré maps, it is necessary to characterize the equilibrium structure of the active dynamical model. In the circular restricted three-body problem, equilibrium points correspond to stationary solutions in the rotating barycentric frame and define key geometric and energetic features of the phase space. In particular, the Lagrange points

establish critical Jacobi constant values that govern the opening and closing of zero-velocity regions and therefore play a central role in determining the admissible motion of trajectories [9, 12, 17].

In this work, the equilibrium structure is computed directly from the governing equations of motion described in Section 3.2 and Section 3.3. Two dynamical models are considered: the baseline planar CR3BP and the J_2 -augmented model that incorporates the oblateness of Neptune. All equilibrium quantities are evaluated using the *active dynamical model*, meaning that the effective potential used in the equilibrium conditions is consistent with the model employed for trajectory propagation. As a result, both the equilibrium locations and their associated Jacobi constants differ slightly between the unperturbed and perturbed cases.

The quantities computed in this section are not presented as final dynamical results, but rather as system-specific reference values required for subsequent analysis. In particular, the Jacobi constants associated with the equilibrium points define the energy levels used to generate Poincaré maps, while the equilibrium geometry provides a reference for interpreting the structure of phase space in the Neptune–Triton system.

3.4.1 Lagrange Point Computation

The Lagrange points used in this study are computed numerically from the equilibrium conditions of the active dynamical model. For a given effective potential, equilibrium in the rotating frame occurs at locations where the gradient of the potential vanishes, consistent with Equation (2.24) in Chapter 2 [12, 17]:

$$\Omega_x(x, y) = 0, \quad \Omega_y(x, y) = 0.$$

These two equations define the planar equilibrium conditions for the system.

The equilibrium calculation is performed using the same model selected for trajectory propagation. For the unperturbed case, the effective potential is Equation (2.6) specialized to planar motion:

$$\Omega(x, y) = \frac{1 - \mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2).$$

with distances from Equations (2.3),

$$r_1 = \sqrt{(x + \mu)^2 + y^2}, \quad r_2 = \sqrt{(x - (1 - \mu))^2 + y^2}.$$

In this work, the Neptune–Triton mass ratio is the system-specific evaluation of Equation (2.2),

$$\mu = 2.09898 \times 10^{-4}. \quad (3.11)$$

For perturbed cases, the same equilibrium conditions are solved using the J_2 -augmented potential. In that model, the Neptune oblateness contribution is added directly to the effective potential through Equation (3.5):

$$\Omega_{J_2}^{(N)} = -\frac{(1 - \mu) J_2 R_{\text{bar}}^2}{2} \left(r_1^{-3} - 3z^2 r_1^{-5} \right),$$

where J_2 is Neptune’s second zonal harmonic coefficient and

$$R_{\text{bar}} = \frac{R_N}{a} \quad (3.12)$$

is the nondimensional equatorial radius of Neptune. For the Neptune–Triton system,

$$R_N = 24,764 \text{ km}, \quad a = 354,760 \text{ km}, \quad (3.13)$$

so that

$$R_{\text{bar}} = \frac{24,764}{354,760} \approx 0.0698. \quad (3.14)$$

When the J_2 perturbation is enabled, this augmented potential is used in place of the classical CR3BP potential, so the equilibrium locations and their associated Jacobi constants remain consistent with the active trajectory model.

Collinear Lagrange Points L_1 , L_2 , and L_3

The collinear Lagrange points are the equilibrium solutions that lie on the rotating-frame x -axis. Because these points lie on the line joining Neptune and Triton, their computation is restricted to the condition

$$y = 0. \quad (3.15)$$

In the planar problem, equilibrium occurs where the gradient of the effective potential vanishes. For points on the x -axis, this reduces to the single scalar condition

$$\Omega_x(x, 0) = 0. \quad (3.16)$$

For the classical CR3BP, the effective potential is again the planar form of Equation (2.6),

$$\Omega(x, y) = \frac{1 - \mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2),$$

with the corresponding planar distances from Equations (2.3),

$$r_1 = \sqrt{(x + \mu)^2 + y^2}, \quad r_2 = \sqrt{(x - (1 - \mu))^2 + y^2}.$$

Taking the partial derivative of Ω with respect to x reproduces the x -component derivative given in Equation (2.7), written here in the planar variables used by the collinear root solve:

$$\Omega_x(x, y) = \frac{\partial}{\partial x} \left(\frac{1 - \mu}{r_1} \right) + \frac{\partial}{\partial x} \left(\frac{\mu}{r_2} \right) + \frac{\partial}{\partial x} \left(\frac{1}{2}(x^2 + y^2) \right).$$

Evaluating these terms one by one,

$$\begin{aligned} \frac{\partial}{\partial x} \left(\frac{1 - \mu}{r_1} \right) &= -(1 - \mu)r_1^{-2} \frac{\partial r_1}{\partial x} \\ &= -(1 - \mu)r_1^{-2} \left(\frac{x + \mu}{r_1} \right) \\ &= -\frac{(1 - \mu)(x + \mu)}{r_1^3}, \end{aligned} \quad (3.17)$$

$$\begin{aligned}
\frac{\partial}{\partial x} \left(\frac{\mu}{r_2} \right) &= -\mu r_2^{-2} \frac{\partial r_2}{\partial x} \\
&= -\mu r_2^{-2} \left(\frac{x - (1 - \mu)}{r_2} \right) \\
&= -\frac{\mu (x - (1 - \mu))}{r_2^3},
\end{aligned} \tag{3.18}$$

and

$$\frac{\partial}{\partial x} \left(\frac{1}{2} (x^2 + y^2) \right) = x. \tag{3.19}$$

Collecting these terms yields

$$\Omega_x(x, y) = x - \frac{(1 - \mu)(x + \mu)}{r_1^3} - \frac{\mu (x - (1 - \mu))}{r_2^3}. \tag{3.20}$$

Imposing the collinear condition $y = 0$ yields

$$r_1 = |x + \mu|, \quad r_2 = |x - (1 - \mu)|, \tag{3.21}$$

so the equilibrium equation becomes

$$x - \frac{(1 - \mu)(x + \mu)}{|x + \mu|^3} - \frac{\mu (x - (1 - \mu))}{|x - (1 - \mu)|^3} = 0. \tag{3.22}$$

This is the one-dimensional nonlinear equation solved in this study to obtain the collinear equilibrium locations for the unperturbed Neptune–Triton CR3BP.

For perturbed cases, the same procedure is applied using the J_2 -augmented potential. In that case, the equilibrium condition is still Equation (3.16):

$$\Omega_x(x, 0) = 0,$$

but Ω is replaced by the augmented effective potential of Equation (3.6),

$$\Omega_{\text{aug}}(x, y, z) = \Omega(x, y, z) + \Omega_{J_2}^{(N)}(x, y, z).$$

For the planar case, $z = 0$, so the Neptune oblateness term reduces to

$$\Omega_{J_2}^{(N)}(x, 0, 0) = -\frac{(1 - \mu)J_2 R_{\text{bar}}^2}{2} r_1^{-3}. \tag{3.23}$$

Differentiating this contribution with respect to x produces the additional acceleration term

$$\frac{\partial \Omega_{J_2}^{(N)}}{\partial x} = -\frac{(1 - \mu)J_2 R_{\text{bar}}^2}{2} \frac{\partial}{\partial x} (r_1^{-3}). \tag{3.24}$$

Since

$$\frac{\partial}{\partial x} (r_1^{-3}) = -3r_1^{-4} \frac{\partial r_1}{\partial x} = -3r_1^{-4} \left(\frac{x + \mu}{r_1} \right) = -\frac{3(x + \mu)}{r_1^5}, \quad (3.25)$$

it follows that

$$\frac{\partial \Omega_{J_2}^{(N)}}{\partial x} = \frac{3}{2} (1 - \mu) J_2 R_{\text{bar}}^2 \frac{(x + \mu)}{r_1^5}. \quad (3.26)$$

Using the definition of k from Equation (3.7), this becomes the x -component in Equation (3.9):

$$a_{J_2, x} = \frac{3k(x + \mu)}{r_1^5},$$

Accordingly, the collinear equilibrium equation for the perturbed model becomes

$$x - \frac{(1 - \mu)(x + \mu)}{|x + \mu|^3} - \frac{\mu(x - (1 - \mu))}{|x - (1 - \mu)|^3} + \frac{3k(x + \mu)}{|x + \mu|^5} = 0. \quad (3.27)$$

This is the model-consistent one-dimensional equilibrium equation used when Neptune's J_2 perturbation is enabled.

The Neptune–Triton system parameters used in this computation are the mass-ratio value in Equation (3.11), together with the physical and nondimensional radius scalings in Equations (3.13) and (3.14). For perturbed cases, the same computation also uses Neptune's zonal harmonic coefficient $J_2 = 3.411 \times 10^{-3}$.

Because the collinear equilibrium equation is nonlinear and possesses three distinct roots, it is solved numerically rather than analytically. Separate search intervals are imposed for the three dynamical regions corresponding to L_1 , L_2 , and L_3 . A bracketing procedure is first used to identify intervals containing sign changes in the equilibrium function, and the roots are then refined using Brent's method [26]. This approach is robust for the small mass ratio of the Neptune–Triton system and ensures that the three collinear roots are found in the correct order and labeled consistently.

The result of this computation is the set of equilibrium coordinates

$$(x_{L_1}, 0), \quad (x_{L_2}, 0), \quad (x_{L_3}, 0). \quad (3.28)$$

These locations are then used to compute the corresponding Jacobi constants and to define the critical energy levels used in the subsequent zero-velocity and Poincaré analyses.

Triangular Lagrange Points L_4 and L_5

The triangular Lagrange points correspond to equilibrium solutions that do not lie on the rotating-frame x -axis. Accordingly, the restriction $y = 0$ used for the collinear points cannot be applied, and the full two-dimensional equilibrium system must be solved [12, 17].

Equilibrium is obtained by requiring both components of the gradient of the effective potential to vanish, as in Equation (2.24):

$$\Omega_x(x, y) = 0, \quad \Omega_y(x, y) = 0.$$

Using the same effective potential defined by Equation (2.6) and the planar distances of Equation (2.3), these conditions yield the coupled nonlinear system

$$\begin{aligned} x - \frac{(1-\mu)(x+\mu)}{[(x+\mu)^2 + y^2]^{3/2}} - \frac{\mu(x-(1-\mu))}{[(x-(1-\mu))^2 + y^2]^{3/2}} &= 0, \\ y - \frac{(1-\mu)y}{[(x+\mu)^2 + y^2]^{3/2}} - \frac{\mu y}{[(x-(1-\mu))^2 + y^2]^{3/2}} &= 0. \end{aligned} \quad (3.29)$$

Unlike the collinear case, these equations must be solved simultaneously for both x and y . The triangular equilibria correspond to configurations in which the spacecraft forms an approximately equilateral triangle with the two primaries. For the classical CR3BP, these solutions are the classical coordinates of Equation (2.22):

$$x = \frac{1}{2} - \mu, \quad y = \pm \frac{\sqrt{3}}{2}.$$

In this study, these analytical solutions are used as initial guesses for the numerical solution of the full equilibrium system. Specifically, the nonlinear system is solved using a two-dimensional root-finding method initialized at

$$(x_0, y_0) = \left(\frac{1}{2} - \mu, \pm \frac{\sqrt{3}}{2} \right). \quad (3.30)$$

Because these initial conditions lie in close proximity to the true equilibria, the solver converges rapidly to the corresponding solutions L_4 and L_5 .

When the J_2 perturbation is included, the equilibrium conditions are modified by the addition of the Neptune oblateness acceleration terms. In this case, the system becomes

$$\begin{aligned} x - \frac{(1-\mu)(x+\mu)}{r_1^3} - \frac{\mu(x-(1-\mu))}{r_2^3} + \frac{3k(x+\mu)}{r_1^5} &= 0, \\ y - \frac{(1-\mu)y}{r_1^3} - \frac{\mu y}{r_2^3} + \frac{3ky}{r_1^5} &= 0, \end{aligned} \quad (3.31)$$

where the planar distances are given by Equation (2.3) and k is defined by Equation (3.7).

The classical equilateral coordinates remain valid initial guesses for the perturbed system, but the resulting equilibrium positions are slightly shifted due to the modified gravitational field of Neptune. The numerical solution of this system yields the triangular equilibrium coordinates

$$(x_{L_4}, y_{L_4}), \quad (x_{L_5}, y_{L_5}), \quad (3.32)$$

which remain symmetric about the x -axis.

Together with the collinear equilibria, these points define the complete set of Lagrange points used in the subsequent stability, zero-velocity, and Poincaré analyses.

Following the numerical solution of both the collinear and triangular equilibrium conditions, the complete set of Lagrange point coordinates is obtained. These include the three collinear points L_1 ,

L_2 , and L_3 , and the two triangular points L_4 and L_5 [12, 17].

The computed equilibrium positions for the classical CR3BP are summarized in Table 3.1. These values are obtained by solving the one-dimensional and two-dimensional equilibrium systems described in Equation 3.4.1 and Equation 3.4.1 using the Neptune–Triton mass ratio.

Table 3.1. Neptune–Triton Lagrange-point positions for the classical planar CR3BP.

Point	x_L	y_L
L_1	0.9591556221	0.0000000000
L_2	1.0415562897	0.0000000000
L_3	-1.0000874575	0.0000000000
L_4	0.4997901020	0.8660254038
L_5	0.4997901020	-0.8660254038

When the J_2 perturbation is included, the equilibrium conditions are modified as described previously, resulting in small shifts in the locations of all five Lagrange points. The corresponding equilibrium coordinates for the J_2 -augmented model are summarized in Table 3.2.

Table 3.2. Neptune–Triton Lagrange-point positions for the planar CR3BP with Neptune J_2 enabled.

Point	x_L	y_L
L_1	0.9591525668	0.0000000000
L_2	1.0415538374	0.0000000000
L_3	-1.0000792406	0.0000000000
L_4	0.4997818864	0.8660206605
L_5	0.4997818864	-0.8660206605

As expected, the inclusion of Neptune’s oblateness produces only small deviations from the classical CR3BP equilibrium positions. However, these shifts are sufficient to alter the associated Jacobi constants and, consequently, the energy thresholds used in the subsequent phase-space analyses.

3.4.2 Jacobi Constants at the Equilibrium Points

Following the computation of the equilibrium coordinates, the associated Jacobi constants are evaluated directly from the Jacobi integral of Equation (2.15). In planar form,

$$C_J = 2 \left[\frac{1-\mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2) \right] - (\dot{x}^2 + \dot{y}^2),$$

with planar distances given by Equation (3.21).

At each Lagrange point, the spacecraft is stationary in the rotating frame. Therefore, the velocity components vanish, consistent with the equilibrium specialization of Equation (2.20):

$$\dot{x} = 0, \quad \dot{y} = 0,$$

and the Jacobi constant reduces to a function of position only. Substituting these conditions into the Jacobi integral yields

$$C_L = 2 \left[\frac{1-\mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x_L^2 + y_L^2) \right]. \quad (3.33)$$

Thus, for each equilibrium point, the Jacobi constant is computed by evaluating the distances of Equation (2.3) at the equilibrium coordinates (x_L, y_L) and substituting those values directly into the expression above.

For the collinear Lagrange points, where $y_L = 0$, the same equilibrium expression is evaluated on the x -axis, with the corresponding distance simplification following directly from Equation (2.3). For the triangular Lagrange points, both x_L and y_L are nonzero, and the full expressions for r_1 and r_2 are used without simplification.

When the J_2 perturbation is included, the Jacobi constant must be evaluated using the augmented effective potential of Equation (3.6). In this case, the planar Jacobi integral becomes

$$C_J = 2 \left[\frac{1-\mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2) - \frac{(1-\mu)J_2 R_{\text{bar}}^2}{2} r_1^{-3} \right] - (\dot{x}^2 + \dot{y}^2). \quad (3.34)$$

Applying the equilibrium condition $\dot{x} = \dot{y} = 0$, the Jacobi constant at each Lagrange point in the perturbed model is

$$C_L = 2 \left[\frac{1-\mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x_L^2 + y_L^2) - \frac{(1-\mu)J_2 R_{\text{bar}}^2}{2} r_1^{-3} \right]. \quad (3.35)$$

In this expression, the additional term represents the contribution of Neptune's oblateness to the total energy of the system. Because this term depends on r_1^{-3} , it primarily affects equilibrium points located closer to Neptune, producing small but measurable shifts in the corresponding Jacobi constants.

The Neptune–Triton system parameters used in this computation are the mass-ratio value in Equation (3.11), Neptune's zonal harmonic coefficient $J_2 = 3.411 \times 10^{-3}$, and the nondimensional radius scale in Equation (3.14).

The resulting values C_{L_1} , C_{L_2} , C_{L_3} , C_{L_4} , and C_{L_5} define the critical energy levels associated with the equilibrium structure of the system. In particular, the Jacobi constants of the collinear points correspond to the energy thresholds at which the zero-velocity neck regions near L_1 and L_2 open or close, thereby controlling the accessibility of trajectories between interior and exterior regions.

In the present study, these equilibrium Jacobi constants are used directly to define the energy levels for the Poincaré map computations. Because the governing potential depends on whether the J_2 perturbation is enabled, the values of C_L differ slightly between the classical and perturbed models. For this reason, the Jacobi constants are recomputed for each model configuration to ensure that all subsequent trajectory analyses are performed using energy levels that are consistent with the active dynamical system.

The computed equilibrium Jacobi constants for both the classical CR3BP and the J_2 -augmented model are summarized in Tables 3.3 and 3.4.

Using the equilibrium coordinates obtained in Subsection 3.4.1, the Jacobi constant at each Lagrange point is evaluated directly from Equations (3.33) and (3.35). Because the velocity vanishes

at equilibrium, these values depend only on the position of each point and the active dynamical model.

The resulting Jacobi constants for the classical CR3BP are summarized in Table 3.3. These values define the critical energy levels associated with the equilibrium structure of the unperturbed Neptune–Triton system.

Table 3.3. Jacobi constants at the Neptune–Triton Lagrange points for the classical planar CR3BP.

Point	C_L
L_1	3.0145841138
L_2	3.0143042166
L_3	3.0002098971
L_4	2.9997901460
L_5	2.9997901460

For the J_2 -augmented model, the Jacobi constants are evaluated using the modified effective potential that includes Neptune’s oblateness. The resulting values are summarized in Table 3.4.

Table 3.4. Jacobi constants at the Neptune–Triton Lagrange points for the planar CR3BP with Neptune J_2 enabled.

Point	C_L
L_1	3.0145655093
L_2	3.0142896868
L_3	3.0001934635
L_4	2.9997737185
L_5	2.9997737185

Comparison of Tables 3.3 and 3.4 shows that the inclusion of the J_2 perturbation produces small but systematic shifts in the equilibrium energy levels. These differences are particularly relevant for the collinear points, where the Jacobi constants determine the opening and closing of the zero-velocity neck regions that govern transport between dynamical regions.

These equilibrium Jacobi constants are used directly in this study to select the energy levels for the Poincaré map computations, ensuring consistency between the chosen initial conditions and the governing dynamical model.

3.4.3 Lagrange Point Stability

After the equilibrium coordinates have been computed, the local stability of each Lagrange point is determined by linearizing the planar equations of motion about the equilibrium solution and evaluating the eigenvalues of the resulting Jacobian matrix. In the present study, this procedure is carried out using the active dynamical model so that the stability classification remains consistent with the equations used for trajectory propagation and phase-space analysis. The linearization matrix

used in the implementation is the standard planar CR3BP form of Equation (2.29) [9, 12, 17],

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \Omega_{xx} & \Omega_{xy} & 0 & 2 \\ \Omega_{xy} & \Omega_{yy} & -2 & 0 \end{bmatrix},$$

evaluated at the equilibrium location (x_L, y_L) . The eigenvalues of this matrix are then computed numerically to determine the local stability type of each Lagrange point. This is the same linearization and eigenvalue workflow implemented in the project code.

Classical CR3BP Stability Calculation

For the classical CR3BP, the planar equations of motion are the $z = 0$ specialization of Equation (2.9),

$$\ddot{x} - 2\dot{y} = \Omega_x(x, y), \quad \ddot{y} + 2\dot{x} = \Omega_y(x, y),$$

with effective potential given by the planar form of Equation (2.6),

$$\Omega(x, y) = \frac{1-\mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2),$$

where the distances are the planar reductions of Equation (2.3),

$$r_1 = \sqrt{(x + \mu)^2 + y^2}, \quad r_2 = \sqrt{(x - (1 - \mu))^2 + y^2}.$$

To analyze the local behavior near an equilibrium point (x_L, y_L) , small perturbations are introduced:

$$x = x_L + \xi, \quad y = y_L + \eta, \quad \dot{x} = \dot{\xi}, \quad \dot{y} = \dot{\eta}. \quad (3.36)$$

Substituting Equation (3.36) into the equations of motion and retaining only first-order terms yields the linearized perturbation equations of Equation (2.28),

$$\ddot{\xi} - 2\dot{\eta} = \Omega_{xx} \xi + \Omega_{xy} \eta, \quad \ddot{\eta} + 2\dot{\xi} = \Omega_{xy} \xi + \Omega_{yy} \eta,$$

where the second derivatives are evaluated at (x_L, y_L) .

For the classical CR3BP, these second derivatives are

$$\begin{aligned} \Omega_{xx} &= 1 - (1 - \mu) \left(\frac{1}{r_1^3} - \frac{3(x + \mu)^2}{r_1^5} \right) - \mu \left(\frac{1}{r_2^3} - \frac{3(x - (1 - \mu))^2}{r_2^5} \right), \\ \Omega_{yy} &= 1 - (1 - \mu) \left(\frac{1}{r_1^3} - \frac{3y^2}{r_1^5} \right) - \mu \left(\frac{1}{r_2^3} - \frac{3y^2}{r_2^5} \right), \\ \Omega_{xy} &= 3(1 - \mu) \frac{(x + \mu)y}{r_1^5} + 3\mu \frac{(x - (1 - \mu))y}{r_2^5}. \end{aligned} \quad (3.37)$$

Introducing the first-order perturbation state of Equation (2.27),

$$\delta \mathbf{X} = [\xi \quad \eta \quad \dot{\xi} \quad \dot{\eta}]^T,$$

the linearized system becomes the matrix form of Equation (2.28),

$$\delta \dot{\mathbf{X}} = A \delta \mathbf{X},$$

with (2.29)

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \Omega_{xx} & \Omega_{xy} & 0 & 2 \\ \Omega_{xy} & \Omega_{yy} & -2 & 0 \end{bmatrix}_{(x_L, y_L)}.$$

For each Lagrange point, the equilibrium coordinates (x_L, y_L) obtained in Subsection 3.4.1 are substituted into the expressions for Ω_{xx} , Ω_{xy} , and Ω_{yy} , and the resulting Jacobian matrix is evaluated numerically. The eigenvalues

$$\lambda_1, \lambda_2, \lambda_3, \lambda_4$$

are then computed from

$$\det(A - \lambda I) = 0. \quad (3.38)$$

The stability classification of each Lagrange point is determined directly from the eigenvalues of the linearized system evaluated at the computed equilibrium coordinates. The interpretation of these eigenvalues follows directly from their real and imaginary components.

Eigenvalues that occur as purely imaginary conjugate pairs, $\lambda = \pm i\omega$, correspond to bounded, oscillatory motion about the equilibrium point. In this case, perturbations neither grow nor decay, but instead remain confined, indicating a center-type equilibrium that is neutrally stable in the linear sense.

In contrast, the presence of a real eigenvalue pair, $\lambda = \pm\alpha$ with $\alpha \in \mathbb{R}$, introduces exponential behavior of the form $e^{\pm\alpha t}$. The positive real eigenvalue leads to exponential growth of perturbations, which produces instability in the corresponding direction. When both a real pair and an imaginary pair are present, the equilibrium exhibits a saddle–center structure and is therefore linearly unstable.

Applying this interpretation to the computed eigenvalues, the collinear Lagrange points L_1 , L_2 , and L_3 are found to exhibit one real eigenvalue pair and one purely imaginary pair. This combination corresponds to a saddle–center configuration, confirming that all three collinear points are linearly unstable.

For the triangular Lagrange points L_4 and L_5 , the eigenvalues are found to occur as two purely imaginary conjugate pairs, $\lambda = \pm i\omega_1$ and $\lambda = \pm i\omega_2$, indicating bounded oscillatory motion in all directions. This corresponds to a center-type equilibrium, and therefore the triangular points are neutrally stable in the linear sense.

This behavior is consistent with the classical Routh stability criterion of Equation (2.23) [12, 17]. For the Neptune–Triton system,

$$\mu = 2.09898 \times 10^{-4} \ll \mu_c, \quad (3.39)$$

so stability of L_4 and L_5 is expected. The numerical eigenvalue results presented below confirm this

prediction, providing direct validation of the Routh criterion for the system under consideration.

Eigenvalue Results for the Classical CR3BP

For each Lagrange point, Table 3.5 summarizes the resulting eigenvalue structure in terms of the real pair $\pm\alpha$ and the imaginary pair(s) $\pm i\omega$, together with the corresponding equilibrium type and stability classification. This presentation makes the stability interpretation explicit: purely imaginary pairs correspond to bounded oscillatory motion, while the presence of a real pair indicates exponential growth and therefore instability.

Table 3.5. Eigenvalues and stability classification of the Neptune–Triton Lagrange points for the classical planar CR3BP.

Point	Stability	Type	λ_1	λ_2	λ_3	λ_4
L_1	Unstable	Saddle-center	$2.610705 + 0.000000i$	$-2.610705 + 0.000000i$	$0.000000 + 2.134264i$	$0.000000 - 2.134264i$
L_2	Unstable	Saddle-center	$-2.412258 + 0.000000i$	$2.412258 + 0.000000i$	$0.000000 + 2.013394i$	$0.000000 - 2.013394i$
L_3	Unstable	Saddle-center	$0.000000 + 1.000184i$	$0.000000 - 1.000184i$	$-0.023471 + 0.000000i$	$0.023471 + 0.000000i$
L_4	Stable	Center	$0.000000 + 0.999290i$	$0.000000 - 0.999290i$	$0.000000 + 0.037663i$	$0.000000 - 0.037663i$
L_5	Stable	Center	$0.000000 + 0.999290i$	$0.000000 - 0.999290i$	$0.000000 + 0.037663i$	$0.000000 - 0.037663i$

The computed results show that the collinear Lagrange points L_1 , L_2 , and L_3 each possess one real eigenvalue pair and one purely imaginary pair. This confirms the presence of a saddle–center structure, consistent with linear instability due to exponential divergence along the unstable manifold.

In contrast, the triangular points L_4 and L_5 exhibit two purely imaginary eigenvalue pairs. This indicates bounded oscillatory motion in all directions and corresponds to a center-type equilibrium. In the context of linear analysis, these points are therefore linearly stable.

These numerical results are consistent with the theoretical expectations established by the Routh stability criterion and confirm that the Neptune–Triton system lies well within the stable regime for the triangular Lagrange points.

Eigenvalue Results for the J_2 -Augmented Model

When Neptune’s oblateness is included, the stability analysis follows the same procedure as in the classical CR3BP, but with the augmented equations of motion. In this case, the planar system becomes

$$\ddot{x} - 2\dot{y} = \Omega_x(x, y) + a_{J_2,x}(x, y), \quad \ddot{y} + 2\dot{x} = \Omega_y(x, y) + a_{J_2,y}(x, y),$$

which is the planar specialization of Equation (3.10), where the additional acceleration terms are given by Equation (3.9):

$$a_{J_2,x} = \frac{3k(x + \mu)}{r_1^5}, \quad a_{J_2,y} = \frac{3ky}{r_1^5}.$$

As a result, the linearized system is modified through the second derivatives of the augmented potential. Specifically, the Jacobian matrix (2.29) retains the same structure,

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \Omega_{xx}^{(\text{aug})} & \Omega_{xy}^{(\text{aug})} & 0 & 2 \\ \Omega_{xy}^{(\text{aug})} & \Omega_{yy}^{(\text{aug})} & -2 & 0 \end{bmatrix},$$

but the entries now include contributions from the J_2 perturbation.

These augmented second derivatives can be written explicitly as

$$\Omega_{xx}^{(\text{aug})} = \Omega_{xx} + \frac{\partial a_{J_2,x}}{\partial x}, \quad \Omega_{xy}^{(\text{aug})} = \Omega_{xy} + \frac{\partial a_{J_2,x}}{\partial y}, \quad \Omega_{yy}^{(\text{aug})} = \Omega_{yy} + \frac{\partial a_{J_2,y}}{\partial y}, \quad (3.40)$$

where the additional terms arise from differentiating the J_2 acceleration components. For the planar case ($z = 0$), these derivatives introduce higher-order radial dependence through terms proportional to r_1^{-5} and r_1^{-7} , which modify the local curvature of the effective potential near Neptune.

The eigenvalues are then computed by evaluating this augmented Jacobian matrix at the perturbed equilibrium coordinates (x_L, y_L) and solving the same characteristic equation as in the unperturbed case, Equation (3.38):

$$\det(A - \lambda I) = 0.$$

For each Lagrange point, Table 3.6 summarizes the resulting eigenvalue structure in terms of the real pair $\pm\alpha$ and the imaginary pair(s) $\pm i\omega$, together with the corresponding equilibrium type and stability classification.

The computed results show that the collinear Lagrange points L_1 , L_2 , and L_3 retain a saddle-center structure, characterized by one real eigenvalue pair and one purely imaginary pair. This indicates that the inclusion of the J_2 perturbation does not alter the fundamental instability of the collinear equilibria, although the numerical values of the eigenvalues are shifted.

For the triangular points L_4 and L_5 , the eigenvalues remain purely imaginary, indicating bounded oscillatory motion in all directions. Thus, these points remain center-type equilibria and are linearly stable in the presence of the J_2 perturbation.

Comparison with the classical CR3BP results shows that the primary effect of Neptune's oblateness is to modify the quantitative values of the eigenvalues through changes in both the equilibrium locations and the local curvature of the effective potential. However, for the Neptune–Triton system, these changes are not sufficient to alter the qualitative stability classification of any of the five Lagrange points.

Table 3.6. Eigenvalues and stability classification of the Neptune–Triton Lagrange points for the planar CR3BP with Neptune J_2 enabled.

Point	Stability	Type	λ_1	λ_2	λ_3	λ_4
L_1	Unstable	Saddle-center	$-2.610416 + 0.000000i$	$2.610416 + 0.000000i$	$0.000000 + 2.134094i$	$0.000000 - 2.134094i$
L_2	Unstable	Saddle-center	$-2.412457 + 0.000000i$	$2.412457 + 0.000000i$	$0.000000 + 2.013519i$	$0.000000 - 2.013519i$
L_3	Unstable	Saddle-center	$0.000000 + 1.000208i$	$0.000000 - 1.000208i$	$-0.023471 + 0.000000i$	$0.023471 + 0.000000i$
L_4	Stable	Center	$0.000000 + 0.999315i$	$0.000000 - 0.999315i$	$0.000000 + 0.037662i$	$0.000000 - 0.037662i$
L_5	Stable	Center	$0.000000 + 0.999315i$	$0.000000 - 0.999315i$	$0.000000 + 0.037662i$	$0.000000 - 0.037662i$

3.5 Lagrange Point Geometry

Having computed the equilibrium locations and their associated stability properties, it is instructive to examine the geometric configuration of the Lagrange points in the rotating frame. While Subsection 3.4.1 and Subsection 3.4.3 established the existence and classification of these equilibria through analytical and numerical methods, the present section focuses on their spatial arrangement and how this structure is modified by Neptune's oblateness [12, 17].

In the rotating barycentric frame, the Lagrange points define a fixed geometric structure relative to the two primaries. This structure governs the topology of nearby trajectories and provides the foundational framework for understanding phase-space transport, zero-velocity surfaces, and the organization of periodic and quasi-periodic motion [9, 12, 17].

The following subsections present the equilibrium configurations for both the classical CR3BP and the J_2 -augmented model, followed by a direct geometric comparison. The full-scale view emphasizes the global five-point structure of the system; however, as shown in Subsection 3.5.2, resolving perturbation-induced differences requires localized magnification.

3.5.1 Classical CR3BP Configuration

Figure 3.1 shows the geometric arrangement of the five Lagrange points for the classical Neptune–Triton CR3BP in the rotating frame. The primaries are located along the x -axis, with Neptune at $x = -\mu$ and Triton at $x = 1 - \mu$.

The three collinear points L_1 , L_2 , and L_3 lie along the line connecting the primaries. The point L_1 is located between Neptune and Triton, L_2 lies beyond Triton on the exterior side, and L_3 is positioned on the opposite side of Neptune.

The triangular points L_4 and L_5 form an equilateral configuration with the two primaries, located symmetrically above and below the x -axis. This symmetry is a direct consequence of the form of the classical effective potential, which is invariant under reflection about the x -axis.




Figure 3.1. Lagrange-point geometry in the rotating frame of the classical Neptune–Triton CR3BP. The figure shows the full equilibrium configuration used in the phase-space analysis, with Neptune and Triton fixed on the x -axis and the five equilibrium points plotted at their computed locations.

To facilitate direct comparison with the perturbed model, a zoomed view of the classical CR3BP configuration is shown in Figure 3.2. At this scale, the precise locations of the collinear points relative to the secondary are clearly resolved. This localized view serves as a baseline reference for identifying the small geometric shifts introduced by the J_2 perturbation in Subsection 3.5.2.

Classical Zoomed Lagrange-Point Geometry




Figure 3.2. Zoomed Lagrange-point geometry in the rotating frame of the classical Neptune–Triton CR3BP. The figure resolves the region near Triton and the collinear equilibria, where the local geometry used for comparison with the J_2 -perturbed model is most clearly visible.

3.5.2 J_2 -Augmented Configuration

Figure 3.3 shows the equilibrium configuration when Neptune’s oblateness is included through the J_2 perturbation, plotted over the full system scale. At this scale, the geometric structure appears visually indistinguishable from the classical CR3BP configuration.

J_2 -Augmented Lagrange-Point Geometry




Figure 3.3. Lagrange-point geometry in the rotating frame of the J_2 -perturbed Neptune–Triton system. The figure shows the full equilibrium configuration after Neptune’s oblateness is included through the effective potential used in the perturbed model.

This apparent similarity is a consequence of the relative magnitude of the perturbation. Although Neptune's J_2 coefficient is non-negligible, its contribution to the effective potential is scaled by $R_{\text{bar}}^2 = (R_N/a)^2$, resulting in a perturbation that is small compared to the dominant gravitational and centrifugal terms. This is expected otherwise, there will be no orbits around Neptune. As a result, the equilibrium locations experience only minor shifts relative to the system scale.

To resolve these differences, a zoomed comparison near the secondary and collinear equilibrium points is presented in Figure 3.4.

J_2 -Augmented Zoomed Lagrange-Point Geometry




Figure 3.4. Zoomed Lagrange-point geometry in the rotating frame of the J_2 -perturbed Neptune–Triton system. The figure highlights the local shifts of the collinear equilibria near Triton that arise from the oblateness correction.

When compared with Figure 3.2, the small but measurable displacement of the equilibrium points becomes apparent. The collinear points exhibit small but measurable displacements along the x -axis, reflecting the modified gravitational field near Neptune. These shifts arise from the additional radial dependence introduced by the oblateness term, which alters the local balance between gravitational and centrifugal forces.

A similar effect is observed for the triangular points, although their displacement is smaller in relative terms and preserves the symmetry about the x -axis. The persistence of this symmetry is consistent with the planar formulation of the J_2 perturbation, which does not introduce any asymmetry in the y -direction.

Thus, while the global geometric structure of the Lagrange points remains unchanged, the J_2 perturbation introduces localized modifications that are only observable when examined at an appropriate spatial scale. These small geometric shifts play an important role in the detailed dynamical behavior explored in the subsequent phase-space analysis.

3.6 Zero-Velocity Structure

While the Lagrange points are defined by equilibrium conditions in the effective potential, the global accessibility of motion in the rotating frame is governed by the Jacobi integral through a kinematic

constraint on the velocity. This constraint is made explicit through the zero-velocity curves (ZVCs), which define the boundaries at which motion is momentarily arrested in the rotating frame [12, 17]. The governing relation is the Jacobi integral of Equation (2.15), together with the zero-velocity construction discussed in Section 2.5.

Starting from the planar Jacobi relation of Equation (2.30), the velocity magnitude can be written as in Equation (2.16),

$$\dot{x}^2 + \dot{y}^2 = 2\Omega(x, y) - C_J.$$

Because the squared velocity must be non-negative, admissible motion is restricted to regions satisfying

$$2\Omega(x, y) - C_J \geq 0. \quad (3.41)$$

The boundary defined by Equation (2.17),

$$2\Omega(x, y) - C_J = 0,$$

corresponds to points where the velocity is exactly zero in the rotating frame. These locations define the zero-velocity curves. The ZVCs therefore partition the configuration space into allowed regions, where motion is dynamically possible, and forbidden regions, where motion cannot occur.

Although the ZVCs arise purely from the velocity constraint, they are directly linked to the equilibrium structure: the Jacobi constants associated with the Lagrange points,

$$C_{L_1}, \quad C_{L_2}, \quad C_{L_3},$$

define critical values at which the topology of the zero-velocity curves changes. As the Jacobi constant varies, neck regions near the collinear points open or close, regulating transport between the vicinity of Neptune, Triton, and the exterior domain.

In this section, zero-velocity curves are computed and visualized for both the classical CR3BP and the J_2 -augmented model. These results provide a geometric interpretation of the Jacobi constraint and establish the connection between equilibrium structure and global trajectory accessibility.

3.6.1 Zero-Velocity Curves in the Classical CR3BP

For the classical Neptune–Triton CR3BP, the zero-velocity curves are computed from the planar Jacobi integral obtained from Equation (2.30), where the effective potential is the planar specialization of Equation (2.6),

$$\Omega(x, y) = \frac{1 - \mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2),$$

and the distances are the planar reductions of Equation (2.3),

$$r_1 = \sqrt{(x + \mu)^2 + y^2}, \quad r_2 = \sqrt{(x - (1 - \mu))^2 + y^2}.$$

The zero-velocity boundary is obtained by setting $\dot{x} = \dot{y} = 0$, yielding

$$F(x, y; C_J) = 2 \left[\frac{1 - \mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2) \right] - C_J = 0. \quad (3.42)$$

Thus, for a specified Jacobi constant, the ZVC is the zero-level contour of $F(x, y; C_J)$ evaluated on a two-dimensional grid.

Using the Neptune–Triton mass ratio of Equation (3.11),

$$\mu = 2.09898 \times 10^{-4},$$

the equilibrium Jacobi constants define the critical topology transitions:

$$\begin{aligned} C_{L_1} &= 3.0145841138, & C_{L_2} &= 3.0143042166, \\ C_{L_3} &= 3.0002098971, & C_{L_4} = C_{L_5} &= 2.9997901460. \end{aligned}$$

These values determine when transport pathways open:

- $C_J > C_{L_1}$: all necks closed, motion confined to isolated regions
- $C_{L_2} < C_J < C_{L_1}$: L_1 neck opens (Neptune–Triton communication)
- $C_{L_3} < C_J < C_{L_2}$: L_2 neck opens (access to exterior region)
- $C_J < C_{L_3}$: outer barrier removed, global accessibility

Figure 3.5 summarizes the classical zero-velocity topology across the representative energy levels used in this study, combining the global accessibility changes with a local view of the collinear neck region near Triton.

3.6.2 Critical Energy Levels and Neck Opening Topology

The topology of the zero-velocity curves is most clearly understood by examining the transition at the L_1 threshold. Figure 3.5(a) first provides the global configuration after the outer barrier has opened. The local collinear-region progression is then shown in panels (b) through (d): the neck remains closed for $C_J > C_{L_1}$, so motion is confined to isolated regions surrounding each primary; at $C_J = C_{L_1}$, the zero-velocity boundary passes through L_1 , producing the limiting bottleneck geometry; and for $C_{L_2} < C_J < C_{L_1}$, this neck opens and permits transport between the Neptune-centered and Triton-centered domains.


Further decreases in C_J produce analogous transitions at L_2 and L_3 , progressively increasing global accessibility as shown by the combined local and full-system views in Figure 3.5. These transitions demonstrate that the ZVCs act as energy-dependent gateways that regulate transport throughout the system.

3.6.3 Zero-Velocity Curves in the J_2 -Augmented Model

When Neptune’s oblateness is included, the effective potential is modified by the additional J_2 term. Using the augmented potential of Equation (3.6), the Jacobi integral of Equation (3.34) becomes


$$C_J = 2\Omega_{\text{aug}}(x, y) - (\dot{x}^2 + \dot{y}^2), \quad (3.43)$$

Classical Zero-Velocity Overview



a.) $C_J > C_{L_1}$

b.) Zoomed view of bottleneck



c.)
 $C_{L_2} < C_J < C_{L_3}$

d.) $C_J < C_{L_4}$

Figure 3.5. Zero-velocity curves for the classical Neptune–Triton CR3BP in the rotating barycentric frame. The shaded regions indicate forbidden motion where $2\Omega(x, y) - C_J < 0$, and the zero-velocity boundary is defined by $F(x, y; C_J) = 0$.

where the augmented effective potential is the planar specialization of Equation (3.6),

$$\Omega_{\text{aug}}(x, y) = \frac{1 - \mu}{r_1} + \frac{\mu}{r_2} + \frac{1}{2}(x^2 + y^2) - \frac{(1 - \mu)J_2 R_{\text{bar}}^2}{2} r_1^{-3}.$$

The corresponding zero-velocity contour is the J_2 -augmented analogue of Equation (3.42),

$$F_{\text{aug}}(x, y; C_J) = 0. \quad (3.44)$$

Figure 3.6 compares the global and local zero-velocity structure in the presence of the J_2 perturbation.




Figure 3.6. Zero-velocity curves for the Neptune–Triton system with Neptune’s J_2 oblateness included. Panel (a) shows the global accessibility of motion, and panel (b) shows a local zoom near Triton and the collinear equilibrium region. The J_2 perturbation produces small shifts in the zero-velocity geometry while preserving the overall topology.

Although the global topology is preserved, the local deformation of the zero-velocity curves near the primary leads to subtle changes in the geometry of the neck regions. These changes influence transport pathways and become important in the phase-space analyses that follow.

The zero-velocity curves provide a global, energy-dependent description of where motion is possible, but they do not describe how trajectories evolve within these regions. In particular, they do not capture invariant structures, resonance islands, or chaotic transport.

To resolve the internal phase-space organization of trajectories, the next section introduces Poincaré maps, which reveal the dynamical structure within the admissible regions defined by the Jacobi constraint.

3.7 Poincaré Map Construction

This section defines the numerical procedure used to construct Poincaré maps from the propagated trajectory data. While the theoretical foundations of Poincaré maps were introduced in Section 2.7,

the focus here is on their implementation within the computational framework developed for this study.

The Poincaré map is generated by intersecting continuous trajectories with a prescribed surface of section and recording the resulting discrete sequence of phase-space points. In the present work, this process is carried out by combining three primary components: (i) construction of initial conditions at fixed Jacobi constant, (ii) numerical propagation of trajectories under the selected dynamical model, and (iii) detection and extraction of surface-of-section crossings.

Together, these steps produce a discrete mapping that captures the underlying phase-space structure of the system. The resulting datasets form the basis for identifying invariant curves, resonance structures, and regions of chaotic motion, and serve as the primary input for the frequency analysis developed in the following subsection.

3.7.1 Surface of Section Definition and Event Detection

The Poincaré map is constructed by recording intersections of system trajectories with a prescribed surface in phase space. In the present work, the surface of section is defined by the planar condition $y = 0$ [9–11], corresponding to the x -axis in the rotating reference frame. Let the system state be given by

$$(x(t), y(t), \dot{x}(t), \dot{y}(t)).$$

A surface-of-section crossing occurs at a time t^* satisfying the section condition (2.32),

$$y(t^*) = 0.$$

This condition identifies all intersections of the trajectory with the surface but does not distinguish the direction of crossing. Consequently, a single trajectory may intersect the surface multiple times per cycle, producing redundant points in the map.

To ensure a unique and consistent mapping, the directional constraint $\dot{y}(t^*) > 0$ is imposed at the crossing, so that only intersections with positive transverse velocity are retained. This guarantees that each oscillatory cycle contributes a single point to the Poincaré map.

The surface condition $y = 0$ reduces the four-dimensional phase space $(x, y, \dot{x}, \dot{y})$ to the three-dimensional manifold $(x, \dot{x}, \dot{y})$. The Jacobi integral provides the additional constraint of Equation (2.30),

$$C_J = 2\Omega(x, 0) - (\dot{x}^2 + \dot{y}^2),$$

which allows $\dot{y}$ to be expressed from the Jacobi velocity relation (2.31) as

$$\dot{y} = \sqrt{2\Omega(x, 0) - C_J - \dot{x}^2},$$

where the positive branch is selected by the directional condition. For a fixed Jacobi constant, $\dot{y}$ is therefore uniquely determined by $(x, \dot{x})$, reducing the system to a two-dimensional mapping.

At each valid crossing, the state vector is evaluated on the surface $y = 0$ and recorded in the reduced coordinates

$$(x, \dot{x}),$$

which define the Poincaré map.

Numerically, crossings are detected during trajectory integration using event detection within

the ODE solver. The integrator identifies times t^* satisfying $y(t^*) = 0$, reconstructs the crossing state using the dense-output interpolant, and applies the directional constraint $\dot{y}(t^*) > 0$ to retain only valid intersections. In addition, crossings with $|\dot{x}(t^*)|$ below a small numerical tolerance are discarded in order to reject nearly tangential events that do not provide a clean transverse section. The resulting sequence of points $(x, \dot{x})$ defines the discrete mapping of the continuous trajectory in phase space.

3.7.2 Initial Condition Construction from the Jacobi Constant

The construction of initial conditions for the Poincaré map is performed by prescribing values of the Jacobi constant C_J and systematically selecting initial positions along the surface of section. For each chosen value of C_J , initial conditions are generated such that they satisfy both the surface constraint and the energy constraint of the system. This is the Chapter 3 implementation of the section constraints $y = 0$ with $\dot{y} > 0$ and the Jacobi velocity relation of Equation (2.31).

The initial position is defined along the x -axis,

$$y_0 = 0,$$

with a prescribed value of x_0 . The corresponding velocity component $\dot{x}_0$ is selected as

$$\dot{x}_0 = 0,$$

so that initial motion is directed entirely in the y -direction. Under these conditions, the Jacobi integral (2.30) becomes

$$C_J = 2\Omega(x_0, 0) - \dot{y}_0^2. \quad (3.45)$$

Rearranging this expression step by step gives

$$\dot{y}_0^2 = 2\Omega(x_0, 0) - C_J,$$

and therefore the initial velocity (2.31) is

$$\dot{y}_0 = \pm \sqrt{2\Omega(x_0, 0) - C_J}.$$

The existence of real-valued initial conditions requires that the admissibility condition (3.41) hold,

$$2\Omega(x_0, 0) - C_J \geq 0.$$

This condition defines the admissible region of motion and ensures that the selected initial condition lies within the energetically accessible domain of the system.

For each value of x_0 , two possible initial velocities exist, corresponding to the positive and negative branches of $\dot{y}_0$. In the numerical implementation, these branches are organized into prograde and retrograde families relative to a selected reference body. The sign of $\dot{y}_0$ is assigned from the location of the initial point with respect to that reference body, so that each sampled x_0 generates one prograde and one retrograde initial condition with a consistent geometric orientation.

The initial conditions are therefore given by

$$(x_0, 0, 0, \dot{y}_0),$$

with the initial velocity (2.31)

$$\dot{y}_0 = \pm \sqrt{2\Omega(x_0, 0) - C_J}.$$

To generate the Poincaré maps in this study, the Jacobi constant was not selected as a continuous variable but instead sampled systematically between the critical energy levels defined by the Lagrange points. Specifically, energy values were constructed over the intervals $C_{L_1} \rightarrow C_{L_2}$, $C_{L_2} \rightarrow C_{L_3}$, and $C_{L_3} \rightarrow C_{L_4} = C_{L_5}$, which partition the dynamically distinct regions of the system.

Within each interval, the energy range was subdivided into four equal increments. For a given pair of bounding energies C_{upper} and C_{lower} , the step size was defined as

$$\Delta C_J = \frac{C_{\text{upper}} - C_{\text{lower}}}{4}, \quad (3.46)$$

and the sampled energy values were taken as

$$C_J = C_{\text{upper}}, \quad C_{\text{upper}} - \Delta C_J, \quad C_{\text{upper}} - 2\Delta C_J, \quad C_{\text{upper}} - 3\Delta C_J.$$

This procedure was applied sequentially across each interval, with the final lower bound included to complete the set. The sampled values therefore correspond to the quarter-step energy levels generated between C_{L_1} , C_{L_2} , C_{L_3} , and $C_{L_4} = C_{L_5}$ in the numerical pipeline used throughout this study.

For each selected value of C_J , a continuous range of initial positions x_0 was sampled over a prescribed interval along the surface $y = 0$. The corresponding initial velocity $\dot{y}_0$ was then computed from the Jacobi constraint, and each resulting trajectory was propagated to generate a sequence of surface-of-section intersections.

This procedure produces a family of trajectories at fixed energy, allowing the structure of phase space to be examined through the resulting Poincaré map.

3.7.3 Trajectory Propagation and Termination Criteria

Once the initial conditions are constructed, trajectories are propagated under the governing equations of motion using a high-order adaptive Runge–Kutta integrator. In the present work, numerical integration is performed using the DOP853 method implemented through `solve_ivp`, with relative and absolute tolerances selected to ensure accurate preservation of the Jacobi constant over long integration times [26, 27].

For a given initial condition $(x_0, 0, 0, \dot{y}_0)$, the trajectory is integrated forward in time according to the equations of motion of the selected dynamical model (classical CR3BP or J_2 -augmented model). During propagation, event detection is used to identify successive intersections with the surface of section defined by $y = 0$ and $\dot{y} > 0$, as described in the previous subsection.

In the numerical pipeline used in this thesis, each trajectory is propagated over a prescribed time interval $t \in [0, t_{\text{max}}]$. The target number of crossings N_{cross} is still tracked for monitoring purposes, but it is not used as a hard stopping criterion in the trajectory generator. Instead, all valid crossings detected within the full propagation interval are recorded and associated with the corresponding

trajectory.

In addition to the surface-of-section events, the propagation routine also monitors simple geometric safety conditions that flag escape or collision-like behavior. Escape is defined by the condition

$$r > r_{\max},$$

where $r = \sqrt{x^2 + y^2}$, and collision with either primary is detected when the corresponding distance from the spacecraft falls below a prescribed threshold,

$$r_1 \leq r_{1,\min} \quad \text{or} \quad r_2 \leq r_{2,\min}.$$

These thresholds are treated as configurable numerical parameters of the computation and are recorded together with the run metadata.

In the present implementation, these collision and escape conditions are recorded as non-terminal event flags rather than hard termination events. Consequently, a trajectory may be marked as colliding or escaping in the output metadata while the integration itself still proceeds to the prescribed final time $t_{\max}$, unless the numerical solver fails before that point.

For visualization of the Poincaré maps, however, such event-flagged trajectories are not necessarily discarded outright. Instead, the stored surface-of-section crossings may be clipped at the first recorded effective stopping time t_{end} , so that only the bounded portion of the trajectory is retained in the displayed map.

For the Neptune–Triton production datasets analyzed in this thesis, the stored run metadata indicate that the generator was executed with $x_{0,\text{step}} = 0.01$, $N_{\text{cross}} = 1000$, and $t_{\max} = 15000$, while sampling both prograde and retrograde branches relative to the primary-centered reference choice P_1 . The surface of section used the directional constraint $\dot{y} > 0$ (`sos_direction=+1`), and the trajectory generator retained the J_2 -perturbed model with $J_{2,\text{primary}} = 0.003411$, $J_{2,\text{secondary}} = 0$, $R_{\text{primary}} = 0.06940466794$, and $R_{\text{secondary}} = 1$. The corresponding safety thresholds were $r_{1,\min} = 6.940467 \times 10^{-2}$, $r_{2,\min} = 3.812718 \times 10^{-3}$, and $r_{\max} = 244.108693$, with DOP853 relative and absolute tolerances both set to 10^{-12} . These values define the practical numerical configuration behind the Neptune–Triton trajectory sets used throughout Chapters 5 and 6.

The combination of adaptive integration, event detection, dense-output crossing reconstruction, and configurable event flags provides a robust framework for generating trajectory data that remains both dynamically consistent and numerically reliable.

3.7.4 Data Output Structure and Storage Format

The trajectory data generated from numerical integration are recorded in a structured format designed to support both phase-space analysis and frequency-domain computations. For each propagated trajectory, the full time history of the state vector is stored together with the initial-condition parameters, reference-body information, crossing flags, and run metadata needed to reproduce the calculation.

The time-history output is written on a uniform evaluation grid of the form of Equation (2.37),

$$t_k = k \Delta t,$$

constructed through the `t_eval` argument of `solve_ivp`. Thus, while the DOP853 integrator

still uses adaptive internal time steps, the stored trajectory arrays are sampled on a prescribed uniform grid suited for subsequent spectral analysis. In addition to the propagated state variables, the time-series files retain auxiliary quantities such as x_0 , $\dot{y}_0$, C_J , the selected reference body, the effective stopping time used for plotting, and binary flags identifying retained surface-of-section events.

Surface-of-section crossings are written separately as reduced Poincaré tables. These files store the crossing time and crossing state, including x , $\dot{x}$, $\dot{y}$, the reference angular-momentum sign, the generating initial condition, and interpolation metadata associated with the event reconstruction. Although the primary reduced coordinates used in the map visualization are $(x, \dot{x})$, the stored files retain the additional information needed for inspection, candidate selection, and downstream analysis.

The data are organized into a hierarchical directory structure based on the dynamical model and Jacobi constant. The base propagation run writes the time-series files, the Poincaré crossing tables, and the primary frequency products for each dataset. A subsequent ratio and resonance stage then reads the saved frequency tables, assembles the preferred ratio products used for interpretation, and writes dedicated frequency-run outputs, including detailed ratio tables, resonance summaries, and run-level metadata. This structure enables efficient retrieval and filtering of datasets for visualization and analysis while preserving a one-to-one correspondence between each sampled trajectory, its Poincaré representation, and its associated spectral products.

All quantities are stored in nondimensional canonical units consistent with the CR3BP formulation, ensuring compatibility across all components of the computational pipeline. This standardized data format allows seamless integration with the Poincaré map visualization tools, frequency analysis routines, and stability-analysis workflow used in the remainder of the thesis.

3.7.5 Discrete Mapping and Phase-Space Representation

The procedure described in the preceding subsections produces a discrete mapping of the continuous system dynamics. Each trajectory, defined by an initial condition $(x_0, 0, 0, \dot{y}_0)$ at fixed Jacobi constant C_J , generates a sequence of surface-of-section intersections

$$(x_1, \dot{x}_1), (x_2, \dot{x}_2), \dots, (x_n, \dot{x}_n),$$

corresponding to successive crossings satisfying $y = 0$ and $\dot{y} > 0$.

This sequence defines the Poincaré map as an iterative transformation of the form of Equation (2.34),

$$(x_k, \dot{x}_k) \mapsto (x_{k+1}, \dot{x}_{k+1}),$$

which represents the evolution of the system from one surface intersection to the next. The continuous trajectory in four-dimensional phase space is thus reduced to a discrete dynamical system in the two-dimensional variables $(x, \dot{x})$.

For a given Jacobi constant, the collection of all intersection points generated from a range of initial conditions produces a global representation of the accessible phase space. Regular trajectories appear as continuous invariant curves, indicating quasi-periodic motion confined to tori in the full phase space. In contrast, irregular or chaotic trajectories produce scattered or diffuse point distributions, reflecting sensitive dependence on initial conditions.

The Poincaré map therefore provides a compact representation of the dynamical structure of

the system at fixed energy. In the present work, these maps serve as the primary analytical tool for identifying resonant structures and regions of stability, and they provide the foundation for the frequency analysis developed in the following subsection.

3.8 Frequency Ratio Analysis

While Poincaré maps provide a geometric representation of phase-space structure, they do not directly quantify the underlying dynamical frequencies associated with individual trajectories. To obtain a quantitative characterization of motion, frequency analysis is performed using time series extracted from the numerically propagated trajectories. In the present workflow, the primary frequency extraction is carried out as part of the unified generation pipeline, which analyzes the uniformly sampled trajectory histories produced during propagation and stores the resulting spectral quantities together with the dataset. A later ratio and resonance stage then operates on those saved frequency products without requiring the trajectories to be integrated again.

For each trajectory generated at fixed Jacobi constant, the stored time history is used to construct a scalar observable suitable for spectral analysis. In the present implementation, the x -coordinate is used,

$$s(t) = x(t), \quad (3.47)$$

which captures the dominant periodic components of the motion and provides a consistent time series for frequency extraction.

The resulting time series is analyzed in the frequency domain to identify the principal frequencies governing the trajectory. In the current implementation, this analysis combines a Fourier-based global spectral estimate with a local NAFF-like refinement procedure. The Fourier spectrum is used to identify the dominant portion of the signal and to diagnose sideband structure in the sense described by Bosanac [21], while the final extracted component frequencies are refined through local projection-based fitting and residual subtraction. These extracted quantities are written directly by the unified generator and become the source inputs for the later ratio and resonance analysis.

The resulting frequency ratio provides a quantitative measure of the dynamical behavior of the trajectory and enables the identification of resonant structures within the phase space. In the following subsections, the procedures for time-series construction, spectral analysis, and frequency extraction are defined explicitly.

3.8.1 Time Series Construction

Frequency analysis is performed on time series extracted directly from the numerically propagated trajectories. For each initial condition at fixed Jacobi constant, the continuous solution of the equations of motion provides the state variables $(x(t), y(t), \dot{x}(t), \dot{y}(t))$ over the integration interval.

In the present implementation, the scalar observable used for spectral analysis is the x -coordinate, as defined in Equation (3.47),

$$s(t) = x(t),$$

which provides a one-dimensional signal capturing the dominant periodic behavior of the trajectory.

The time series is sampled on a uniform time grid of the form of Equation (2.37),

$$t_k = t_0 + k \Delta t, \quad k = 0, 1, 2, \dots, N,$$

where Δt is a prescribed sampling interval and N is the total number of samples. Although the DOP853 integrator employs adaptive internal time stepping, the solution is evaluated at the specified time points t_k using the `t_eval` parameter of the integrator. This ensures uniform sampling, which is required for consistent frequency-domain analysis.

The sampling interval Δt is selected based on the desired frequency resolution and is controlled directly through the time grid used for evaluation. In the present implementation, the frequency analysis uses a dedicated analysis window of duration T_{FFT} , which need not coincide with the full propagation horizon t_{max} . The quantity Δt_{FFT} denotes the uniform sampling interval applied during frequency analysis; it determines the Nyquist frequency and the spectral bandwidth of the extracted spectrum. The effective spectral resolution is therefore governed by the ratio $1/T_{\text{FFT}}$, with longer analysis windows providing finer frequency discrimination. An additional parameter T_{min} specifies the minimum usable trajectory duration required for a trajectory to be included in the frequency analysis; any trajectory whose effective bounded duration falls below this threshold is excluded from the spectral and ratio computations to prevent poorly resolved records from contaminating the global spectral map.

For the Neptune–Triton datasets used in the present study, the stored metadata record $\Delta t_{\text{FFT}} = 0.01$ and $T_{\text{FFT}} = 15000$. The corresponding frequency-run metadata further indicate that the downstream ratio/resonance analysis was carried out with $T_{\text{min}} = 5000$ and a secondary duration threshold $T_{\text{trajectory,min}}$, which specifies the minimum bounded trajectory length acceptable at the ratio and resonance classification stage; trajectories shorter than $T_{\text{trajectory,min}}$ are excluded from the resonance labeling even if they passed the earlier T_{min} filter. For the Neptune–Triton production runs, $T_{\text{trajectory,min}} = 100$. A resonance tolerance of 0.001 was applied for the production classifications discussed in Chapter 6. These settings define the practical spectral-resolution and filtering limits under which the reported resonance maps were generated.

The resulting uniformly sampled time series serves as the input to the spectral analysis and frequency extraction procedures described in the following subsection.

3.8.2 Spectral Analysis and Frequency Extraction

The uniformly sampled time series $s(t_k)$ constructed in the previous subsection is analyzed in the frequency domain to extract the dominant dynamical frequencies of each trajectory. This is accomplished through a two-stage procedure. First, a discrete Fourier transform (DFT), implemented efficiently via the fast Fourier transform (FFT), is used to obtain the global spectral structure of the signal. Second, the dominant components identified from that spectrum are refined locally through a NAFF-like projection and subtraction procedure applied directly to the time series [13, 21, 24].

Given the sampled signal of Equation (2.38),

$$s_k = s(t_k), \quad k = 0, 1, 2, \dots, N-1,$$

the discrete Fourier transform is defined as

$$\hat{s}_n = \sum_{k=0}^{N-1} s_k e^{-2\pi i k n / N}, \quad n = 0, 1, 2, \dots, N-1. \quad (3.48)$$

The magnitude spectrum

$$|\hat{s}_n|$$

represents the amplitude of the frequency components present in the signal.

The corresponding cycle frequencies are given by

$$f_n = \frac{n}{N\Delta t}, \quad (3.49)$$

where Δt is the uniform sampling interval. The frequency resolution of the spectrum is therefore

$$\Delta f = \frac{1}{N\Delta t} = \frac{1}{T_{\text{FFT}}}, \quad (3.50)$$

indicating that longer analysis windows yield finer frequency resolution.

Prior to computing the spectrum, the sampled signal is multiplied by a Hann window in order to reduce endpoint discontinuities and suppress spectral leakage. The largest nonzero Fourier peak provides an initial estimate of the dominant oscillatory content of the trajectory. Rather than using that bin location directly as the final result, the algorithm then refines the frequency locally by maximizing the projection of the signal onto a sinusoid in a narrow neighborhood of the FFT peak.

After the primary component has been refined, a sinusoidal fit is constructed and subtracted from the signal, and the same procedure is repeated on the residual. This yields a second refined component when the remaining spectral energy is sufficiently distinct from the primary contribution. In this way, the implementation combines the robustness of a global FFT-based view with a local refinement step that reduces bin-locking and improves the extraction of nearby frequencies.

In parallel with this component-extraction process, the primary FFT spectrum is also examined for spectral sideband structure. Local spectral maxima are identified after removal of the zero-frequency component, and a symmetric left-right peak pair about the dominant component is selected by minimizing the asymmetry of the peak spacing while favoring large combined amplitude. Let the selected sideband pair be denoted by f_- and f_+ , with $f_- < f_1 < f_+$. The corresponding sideband spacing is then computed as

$$f_{2,B} = \frac{1}{2} (f_+ - f_-), \quad (3.51)$$

which represents the half-spacing of the symmetric sidebands about the dominant peak. This quantity is not identical in principle to the second component extracted from the NAFF-like residual analysis, but it provides a sideband-based characterization of modulation structure that remains valuable for resonance interpretation.

Accordingly, the implementation retains both a directly extracted second component and, when available, a sideband-based estimate. The output frequency tables record the original second-component ratio, the sideband ratio, and a selected active ratio used for plotting and resonance analysis. When a valid sideband estimate is available, it is preferred; otherwise the ratio formed from the two extracted components is retained.

This sideband interpretation is illustrated in Figure 3.7, where the dominant component f_1 is accompanied by a symmetric peak pair at $f_1 - f_{2,B}$ and $f_1 + f_{2,B}$.

The Hann-windowed signal used in this procedure is defined as

$$s_k^{(w)} = s_k w_k, \quad (3.52)$$

Sideband-Based Frequency Interpretation




Figure 3.7. Schematic illustration of the sideband-based frequency extraction method. Panel (a) shows a representative trajectory segment in the rotating frame. Panel (b) shows the corresponding spectrum, where the dominant peak at f_1 is flanked by symmetric sidebands at $f_1 - f_{2,B}$ and $f_1 + f_{2,B}$. The sideband-derived modulation frequency $f_{2,B}$ equals half the total peak-to-peak sideband separation: $f_{2,B} = \frac{1}{2}(f_+ - f_-)$.

where w_k is the Hann window given by

$$w_k = \frac{1}{2} \left(1 - \cos \left(\frac{2\pi k}{N-1} \right) \right), \quad k = 0, 1, 2, \dots, N-1. \quad (3.53)$$

This window smoothly tapers the signal to zero at the endpoints, reducing discontinuities that would otherwise introduce spurious frequency components in the spectrum. As a result, spectral leakage is suppressed and the dominant peaks, including secondary sideband structures, are more clearly resolved. The Fourier transform is therefore applied to the windowed signal $s_k^{(w)}$ rather than the raw time series, and the resulting spectrum then supplies the initial estimates used by the local refinement procedure.

The resulting cycle frequencies f_1 and f_2 , together with the sideband-based quantity $f_{2,B}$ when available, are used to characterize the dynamical behavior of the trajectory and form the basis for the frequency ratio computation described in the following subsection. In the stored output tables these quantities are retained under the legacy field names `omega1` and `omega2`, but numerically they represent cycle frequencies on the uniform sampling grid rather than angular frequencies. The base dataset also preserves both the directly extracted secondary component and the sideband-corrected quantity so that the subsequent ratio/resonance workflow can construct the preferred ratio without repeating the primary frequency extraction.

3.8.3 Frequency Ratio Computation

Once the dominant and secondary frequencies have been extracted from the spectrum, the dynamical behavior of the trajectory is quantified through the frequency ratio, defined in the form of Equation (2.45),

$$\nu = \frac{f_2}{f_1}.$$

This ratio provides a dimensionless measure of the relationship between the principal frequencies governing the motion. For quasi-periodic trajectories, the motion is characterized by multiple incommensurate frequencies, and the ratio ν varies continuously with the initial condition. In contrast, when the frequencies satisfy the rational relationship of Equation (2.46), the trajectory is said to be in resonance, and the motion exhibits periodic or resonantly structured behavior.

In practice, the frequency ratio is computed from the stored spectral values associated with each trajectory. The downstream ratio/resonance stage reads the saved primary frequency products and then assigns the active ratio according to the available spectral information. When a valid sideband estimate exists, the active ratio is taken as

$$\nu_B = \frac{f_{2,B}}{f_1}, \quad (3.54)$$

whereas otherwise the ratio of the two extracted components is retained,

$$\nu_{\text{orig}} = \frac{f_2}{f_1}. \quad (3.55)$$

The resulting values are then associated with the corresponding initial condition x_0 and Jacobi constant C_J , producing a frequency-ratio dataset over the sampled phase-space domain. In the base frequency plot, these values are displayed in the (x_0, ν) plane and are color-coded by orbit sense so that prograde and retrograde structures may be compared directly in ratio space. This representation is used as an initial spectral analogue of the Poincaré map: smooth bands and ordered clusters in the ratio plot provide a guide to coherent dynamical structure, while diffuse or irregular patterns indicate less ordered spectral behavior.

Resonance identification is then performed in a separate downstream stage using the stored ratio values rather than by repeating the primary frequency extraction. In that stage, each computed ratio is compared against a set of rational targets p/q within a prescribed tolerance. Ratios that fall within the chosen tolerance are flagged as resonant, while those outside the tolerance remain unclassified. This approach enlarges the set of prospective resonant trajectories relative to an exact rational test, but it also admits the possibility of both false positives and false negatives near the tolerance boundary.

Once the resonant trajectories have been flagged, all trajectories assigned to the same resonance label are grouped into a common resonance family. Each family is then assigned a unique display color, and this color assignment is used to construct two complementary overlay products. First, the resonance-family colors are applied to the frequency-ratio plot, so that classified resonances may be distinguished from non-resonant trajectories, which remain gray. Second, the same resonance-family colors are projected back onto the Poincaré map, allowing resonant invariant structures in phase space to be compared directly with their corresponding spectral classification.

In this way, the downstream ratio/resonance stage functions as a classification and visualization pass built on the previously extracted spectral products. The frequency ratios provide a quantitative complement to the Poincaré maps, while the resonance-family overlays establish an explicit link between spectral organization and geometric phase-space structure. These products are used in the present study to identify resonant neighborhoods within the Neptune–Triton system and to guide the later selection of trajectories for further dynamical analysis.

3.9 Stability Analysis

The stability of trajectories identified through Poincaré mapping and frequency analysis is evaluated using a periodic-orbit framework based on the state transition matrix and Floquet multipliers [12, 17]. The underlying variational framework was introduced in Section 2.9 and Subsection 2.9.2. In this study, candidate trajectories are selected from regions of interest in the Poincaré maps and the corresponding frequency-ratio distributions. These trajectories are then analyzed by computing the linearized dynamics along the orbit and evaluating the monodromy matrix over one period.

3.9.1 State Transition Matrix and Monodromy Matrix Formulation

For a trajectory $\mathbf{x}(t) = (x, y, \dot{x}, \dot{y})$ governed by $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, the linearized dynamics are described by the variational system of Equation (2.52) [12, 17]

$$\frac{d}{dt}\delta\mathbf{x}(t) = \mathbf{A}(t) \delta\mathbf{x}(t), \quad \mathbf{A}(t) = \left. \frac{\partial \mathbf{f}}{\partial \mathbf{x}} \right|_{\mathbf{x}(t)}.$$

The associated state transition matrix $\Phi(t)$ satisfies Equation (2.55),

$$\frac{d}{dt}\Phi(t) = \mathbf{A}(t)\Phi(t), \quad \Phi(0) = \mathbf{I},$$

and is integrated simultaneously with the trajectory.

For a periodic orbit with period T , the monodromy matrix (2.58) is obtained as [17]

$$\mathbf{M} = \Phi(T).$$

3.9.2 Numerical Implementation of Stability Computation

The stability computation is implemented by augmenting the equations of motion with the variational equations governing the STM. The resulting system consists of 20 coupled differential equations, including the four state variables and the sixteen elements of the STM.

At each integration step, the Jacobian matrix $\mathbf{A}(t)$ is evaluated along the trajectory. For the planar CR3BP, the Jacobian uses the same structure as Equation (2.29) and is given by

$$\mathbf{A}(t) = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \Omega_{xx} & \Omega_{xy} & 0 & 2 \\ \Omega_{yx} & \Omega_{yy} & -2 & 0 \end{bmatrix},$$

where the second derivatives of the effective potential are evaluated using Equation (3.37),

$$\begin{aligned}\Omega_{xx} &= 1 - (1 - \mu) \left(\frac{1}{r_1^3} - \frac{3(x + \mu)^2}{r_1^5} \right) - \mu \left(\frac{1}{r_2^3} - \frac{3(x - (1 - \mu))^2}{r_2^5} \right), \\ \Omega_{yy} &= 1 - (1 - \mu) \left(\frac{1}{r_1^3} - \frac{3y^2}{r_1^5} \right) - \mu \left(\frac{1}{r_2^3} - \frac{3y^2}{r_2^5} \right), \\ \Omega_{xy} &= \Omega_{yx} = 3(1 - \mu) \frac{(x + \mu)y}{r_1^5} + 3\mu \frac{(x - (1 - \mu))y}{r_2^5}.\end{aligned}$$

When the J_2 perturbation is included, the Jacobian is modified through the addition of the partial derivatives of the perturbation accelerations,

$$A_{ij} \rightarrow A_{ij} + \frac{\partial a_{J_2,i}}{\partial x_j}. \quad (3.56)$$

For the planar model, the acceleration components (3.9) are

$$a_{J_2,x} = \frac{3k(x + \mu)}{r_1^5}, \quad a_{J_2,y} = \frac{3ky}{r_1^5},$$

with the coefficient (3.7)

$$k = \frac{1}{2}(1 - \mu)J_2R_{\text{bar}}^2,$$

leading to

$$\begin{aligned}\frac{\partial a_{J_2,x}}{\partial x} &= 3k \left(\frac{1}{r_1^5} - \frac{5(x + \mu)^2}{r_1^7} \right), \\ \frac{\partial a_{J_2,x}}{\partial y} &= -15k \frac{(x + \mu)y}{r_1^7}, \\ \frac{\partial a_{J_2,y}}{\partial x} &= -15k \frac{(x + \mu)y}{r_1^7}, \\ \frac{\partial a_{J_2,y}}{\partial y} &= 3k \left(\frac{1}{r_1^5} - \frac{5y^2}{r_1^7} \right).\end{aligned} \quad (3.57)$$

The augmented system is integrated using the same DOP853 solver and tolerance settings as the trajectory propagation. For each selected trajectory, an initial estimate of the orbit period is obtained primarily from recurrence information in the Poincaré map and the associated surface-of-section candidate-selection procedure. Spectral information from the dominant oscillatory component may be used as a consistency check on the timescale of the motion, but the working period guess used for differential correction is taken from the recurrence structure of the map itself. This estimate is then refined through a differential-correction procedure in which the initial state and period are adjusted simultaneously so as to satisfy the periodic-orbit closure conditions. The final monodromy matrix is therefore evaluated using the corrected period rather than the initial seed estimate.

At the final time, the STM is reshaped into a 4×4 matrix to obtain the monodromy matrix $\mathbf{M}$.

3.9.3 Stability Classification via Floquet Multipliers

The stability classification is performed using the eigenvalues of the monodromy matrix,

$$\lambda_i, \quad i = 1, 2, 3, 4,$$

referred to as Floquet multipliers [17].

An orbit is classified as linearly stable if all multipliers satisfy

$$|\lambda_i| = 1, \quad (3.58)$$

and unstable if any multiplier satisfies

$$|\lambda_i| > 1. \quad (3.59)$$

Due to the symplectic structure of the CR3BP, the eigenvalues occur in reciprocal pairs $(\lambda, 1/\lambda)$ [17]. Stability is therefore determined by the presence or absence of multipliers outside the unit circle.

For reporting purposes, a scalar stability index is defined as in Equation (2.60), where λ is selected from the nontrivial reciprocal multiplier pair after removing the near-unity pair associated with the periodic orbit itself [17]. This scalar provides a compact summary of the nontrivial linear stability behavior of the orbit and is recorded together with the full set of multipliers.

The procedures defined in this section establish a complete computational framework for the identification and characterization of dynamical structures in the Neptune–Triton system. Trajectories are first sampled and propagated under the selected dynamical model, and their intersections with the surface of section are recorded to construct Poincaré maps. These maps provide a geometric representation of phase-space structure and are used to identify regions of regular and resonant motion.

For selected trajectories, time-series data are extracted and analyzed in the frequency domain to compute characteristic frequencies and their ratios. This enables the identification of resonant relationships and provides a quantitative measure of quasi-periodic behavior. Finally, trajectories of interest are subjected to stability analysis through the computation of the monodromy matrix and its associated Floquet multipliers, yielding a rigorous classification of their dynamical stability.

Together, these components form a unified analysis pipeline in which geometric, spectral, and stability-based analyses are applied consistently to the same set of trajectories. The resulting data products serve as the basis for the analysis and interpretation of dynamical structures presented in the subsequent chapter.

3.10 Summary and Remarks of Dynamical Models and Numerical Methods

This chapter has defined the dynamical models and numerical procedures used to generate and analyze trajectory data in the Neptune–Triton system. The equations of motion, including both the classical CR3BP and the J_2 -augmented model, were implemented in a consistent rotating-frame formulation. Numerical integration, event detection, and data extraction methods were specified in detail to ensure reproducibility of all results.

A complete analysis pipeline was then established, consisting of Poincaré map construction, frequency-domain analysis, and stability evaluation through Floquet multipliers. Each component was formulated explicitly and implemented within a unified computational framework, enabling

consistent comparison of trajectories across different dynamical regimes.

The methods developed in this chapter provide the foundation for the results presented in the following chapter, where the structure, resonance behavior, and stability characteristics of trajectories in the Neptune–Triton system are examined in detail.

4. The ORBIT Computational Framework

The preceding chapters established the theoretical and computational foundation for this study. Chapter 1 introduced the scientific motivation and research objectives, Chapter 2 developed the dynamical systems framework underlying the Circular Restricted Three-Body Problem, and Chapter 3 defined the specific Neptune–Triton models and numerical procedures used to generate the trajectory datasets analyzed throughout this thesis. With those elements in place, the remaining task is to describe how these models and procedures were assembled into a unified research environment that could be used efficiently, consistently, and reproducibly.

That environment is ORBIT (Orbital Research and Body Interaction Tool), the computational system developed for this work to organize system definition, model configuration, data generation, analysis, and visualization within a single program. Rather than treating the underlying physics, numerical solvers, and analysis utilities as isolated scripts, ORBIT connects them through a project-based workflow that allows the user to define a system, select the active dynamical model, generate trajectory data, inspect derived quantities, and export results from a consistent interface. In this sense, ORBIT is the operational layer through which the mathematical framework of Chapter 3 is translated into executable analysis.

This chapter therefore shifts the emphasis from the governing equations themselves to the software framework that implements them. The goal is not to repeat the derivations already presented in the previous chapters, but to show how those derivations are embodied in the architecture of the program, how the major analysis modules are connected, and how the graphical interface guides the user through the full workflow of project creation, system setup, trajectory generation, and post-processing. The discussion is limited to the components of ORBIT that support the physics and mathematical investigations of this thesis; a complete software-engineering treatment of every utility, experimental branch, or archived module in the repository is beyond the scope of the present work.

4.1 Motivation for ORBIT

The numerical and analytical tasks required in this thesis are individually straightforward in principle, but become difficult to manage efficiently when carried out as disconnected scripts and manually coordinated post-processing steps. The Neptune–Triton study developed here requires repeated transitions between physical model definition, parameter normalization, trajectory generation, Poincaré map construction, frequency extraction, zero-velocity analysis, and stability assessment. When these tasks are separated into isolated programs, the workflow becomes vulnerable to inconsistent parameter choices, duplicated effort, fragmented data storage, and reduced reproducibility.

ORBIT was developed to address this problem by providing a single computational framework in which the major stages of the analysis could be linked directly. Its purpose is not merely to provide a graphical wrapper around existing scripts, but to unify the full research workflow so that system definition, model selection, data generation, visualization, and export all operate from a common project state. In practical terms, this means that the same physical parameters and numerical settings used to define a Neptune–Triton configuration are carried consistently into the data products and output plots derived from that configuration.

This unified structure was especially important for the present work because the thesis compares multiple dynamical representations of the same system, including the baseline planar CR3BP and the J_2 -augmented extension introduced in Chapter 3. A modular but coordinated environment was therefore needed to ensure that model variations, energy selections, and output products could be generated and compared without repeatedly rebuilding the workflow by hand. ORBIT provides that coordination layer, allowing the underlying mathematical methods to be used in a way that is systematic, traceable, and efficient for large-scale trajectory studies.

4.2 Software Architecture

ORBIT is organized as a layered research framework rather than as a single monolithic script. This structure was adopted so that the physical models, numerical routines, data-generation workflows, and graphical interface could remain logically distinct while still operating as parts of one consistent system. In practice, this separation makes it possible to refine individual components, such as the governing equations or a plotting module, without having to restructure the entire program each time the workflow evolves.

At a high level, the framework may be viewed as four connected layers. The lowest layer contains the physical and mathematical model definitions, including the normalized equations of motion, Lagrange-point routines, and zero-velocity calculations. Above this sits the numerical and data-generation layer, which carries out trajectory propagation, event detection, frequency extraction, and stability computations. A third layer provides analysis and visualization utilities that transform generated data into tables, computed metrics, and plots. Finally, the graphical interface acts as the coordinating layer through which these capabilities are configured, executed, and reviewed. The following subsections describe these components in the order in which they support the overall system; the resulting pipeline is illustrated in Figure 4.1.

4.2.1 Core Physics Layer

The foundational layer of ORBIT consists of the modules that define the dynamical system itself. These files provide the mathematical descriptions of the normalized Neptune–Triton problem and supply the low-level functions used throughout the rest of the framework for model evaluation. The governing equations are defined in a dedicated file and are called by the higher-level generators and analysis modules whenever dynamical quantities must be evaluated. Physical constants and related reference quantities are stored separately, while additional supporting files handle tasks such as normalization, Lagrange-point computation, and zero-velocity calculations.

This layer is responsible for enforcing consistency between the theoretical development of Chapter 3 and the software environment in which that development is applied. For example, the same mass parameter definitions, rotating-frame equations, and effective-potential relationships used in the Chapter 3 derivations are evaluated here during trajectory propagation, equilibrium-point analysis, and zero-velocity calculations. By concentrating these model definitions in a dedicated physics layer, ORBIT avoids duplicating core mathematical logic across multiple scripts or interface modules.

An additional advantage of this structure is that it separates model definition from model use. The physical equations are not embedded directly within the graphical screens or post-processing utilities; instead, higher-level modules call the shared physics routines whenever system quantities

ORBIT Computational Pipeline




Figure 4.1. High-level architecture of the ORBIT program shown as a computational pipeline. Initial conditions are generated from the selected system parameters and Jacobi constant and are then propagated with the CR3BP or J_2 -augmented equations of motion using a shared numerical integration core. The resulting trajectories feed three parallel analysis branches: Poincaré surface-of-section extraction, frequency analysis using spectral methods, and stability analysis through state-transition and monodromy computations. The outputs of these branches are stored and visualized within ORBIT as maps, frequency ratios, stability metrics, and exported research products.

must be computed. In practice, once the governing equations, constants, and normalization logic were established and verified, this physics layer was treated as fixed for the thesis workflow. Users interact with the framework through controlled inputs at the interface and configuration level rather than by editing the underlying physical model files directly. This reduces the risk of accidental changes to equations or parameter values that could introduce silent errors into the results, while also ensuring that all datasets generated within ORBIT rest on the same mathematical foundation.

4.2.2 Numerical and Data-Generation Layer

Built on top of the physics layer is the numerical and data-generation layer, which carries the mathematical model into executable trajectory studies. This portion of the framework contains the routines that perform orbit propagation, event detection, Poincaré map construction, frequency extraction, candidate selection, and stability computation. If the physics layer defines what system is being studied, the numerical layer determines how that system is sampled, propagated, and converted into the datasets used throughout the thesis.

Within ORBIT, these tasks are organized primarily through dedicated generator and analysis scripts rather than through the graphical interface itself. The interface supplies validated inputs and run settings, but the actual computations are performed in a separate group of numerical modules. This design was adopted intentionally so that the research calculations would remain usable both inside the GUI and as independent computational tools when needed. It also keeps the numerical workflow aligned with the procedures described in Chapter 3, since the same generator routines can

be invoked consistently regardless of how a run is launched.

This layer includes the trajectory-generation pipeline used to build the Poincaré and frequency datasets, along with the auxiliary routines used to derive energy arrays, identify candidate trajectories, and compute stability properties of selected periodic-orbit seeds. In this sense, it serves as the production engine of the system. The graphical environment does not replace these numerical routines; rather, it orchestrates and parameterizes them. By preserving that separation, ORBIT maintains a clear distinction between the underlying scientific computation and the interface used to control it, while ensuring that the outputs remain structured, reproducible, and traceable at the project level.

4.2.3 Analysis and Visualization Layer

Above the numerical engine sits the analysis and visualization layer, which converts generated datasets into forms that can be inspected, interpreted, and exported. This portion of ORBIT includes the routines responsible for rendering trajectory-derived plots, displaying output tables, assembling summary metrics, and providing interactive access to the structures identified in the Neptune–Triton phase space. In practical terms, this is the layer through which the raw outputs of the generators become scientifically useful products.

This layer plays an important role in the workflow because the principal outputs of this thesis are not only numerical trajectories, but also derived structures such as Poincaré maps, zero-velocity curves, Lagrange-point summaries, frequency spectra, and stability classifications. These products require additional processing beyond orbit propagation alone. ORBIT therefore separates the numerical generation of data from the visual and analytical routines that help the user interpret that data. This improves modularity while allowing different views of the same dataset to be produced without rerunning the full propagation each time.

The analysis and visualization layer also supports export and reporting. Tables produced in this layer can be written to structured files, while figures generated here provide direct support for interpretation and, in many cases, for inclusion in the thesis itself. In this way, the layer functions as the bridge between computation and presentation: it does not alter the underlying trajectory solutions, but organizes them into forms that are readable, comparable, and suitable for research documentation.

4.2.4 Graphical Interface and Control Layer

The highest layer of ORBIT is the graphical interface and control system through which the user interacts with the rest of the framework. This layer does not define the physics or perform the low-level numerical calculations itself; instead, it manages project creation, accepts user input, routes configuration choices to the appropriate modules, launches the relevant numerical workflows, and presents the resulting analysis products in an organized way. It is therefore best understood as the coordinating layer of the framework.

This layer is especially important in the context of the present work because the underlying workflow spans many interconnected tasks: defining the physical system, choosing the active model, selecting energy ranges, generating large trajectory sets, reviewing plots and analysis tables, and exporting the resulting products. Managing these steps manually through separate scripts would be possible, but inefficient and prone to inconsistency. The graphical interface provides a structured way

to move through this process while preserving continuity between the selected system configuration and every subsequent computation derived from it.

From the standpoint of research use, the interface layer also provides an important boundary between user-adjustable settings and protected scientific logic. The user can modify project parameters, numerical settings, and model selections through controlled interface elements, while the deeper physics and generator routines remain unchanged beneath that surface. This arrangement makes the framework both usable and robust: it allows flexibility where the study requires it, while preserving the integrity of the underlying methods that support reproducible results.

4.3 Dynamical Engine Implementation

While the software architecture described above explains the logical organization of ORBIT, the implementation of the framework begins with a much more concrete task: creating a stable entry point through which the graphical application can initialize the research environment and route the user into the correct workflow. In ORBIT, this role is handled by a dedicated launch file that initializes the application, ensures the repository structure is accessible to the Python environment, and opens the main Qt window through which the remainder of the framework is accessed.

This entry-stage design is more important than it may first appear. Because ORBIT connects graphical controls to multiple physics, generator, plotting, and analysis modules distributed throughout the repository, the program must begin in a known and consistent state. The launch process therefore establishes the repository root, prepares required runtime paths, and initializes the graphical application before any project-specific interaction begins. This ensures that the modules described in the previous section can be accessed consistently from a single point of entry, rather than being invoked through disconnected or ad hoc execution paths.

Once the program is launched, control is transferred to the main application window, which functions as the routing hub of the framework. Rather than embedding every tool within a single oversized interface, ORBIT organizes the user experience as a collection of distinct screens representing different stages of the workflow. The main window manages these screens, preserves shared application state between them, and allows the user to move between project selection, system definition, dashboard control, and analysis modules without restarting the application or manually reloading the active configuration.

4.3.1 Program Entry and Screen Routing

The entry point of the framework is implemented as a dedicated graphical user interface (GUI) launch script that serves as the only direct executable file in the interface layer. Its responsibility is intentionally narrow: it initializes the runtime environment, applies the graphical style configuration, and instantiates the main application window. All substantive computational behavior is deferred to lower-level modules, ensuring that the launch process remains stable and does not become entangled with trajectory generation or post-processing logic.

From that entry point, the central routing logic is handled by the main window object, which maintains a stacked collection of screens corresponding to the major stages of the ORBIT workflow. These screens include the splash screen, project selector, system definition interface, dashboard, and the dedicated analysis and plotting modules. By organizing the application in this way, ORBIT achieves two goals simultaneously. First, it keeps each screen focused on a clearly defined task

rather than overloading the user with every tool at once. Second, it allows the shared application state to persist across the workflow so that system choices, model settings, and project metadata remain available as the user moves from one stage of the analysis to the next.

This routing structure also reflects the broader design philosophy of the framework. ORBIT is not simply a plotting interface layered on top of a collection of scripts, nor is it a monolithic application in which all computation is embedded directly into the graphical elements. Instead, the GUI acts as a controlled navigation and coordination system that exposes the major research capabilities of the framework in a consistent order. The user enters through a single launch point, proceeds through the screen hierarchy, and interacts with the physics and numerical engines through routed, state-aware interface components rather than through manual command execution.

4.4 Poincaré Data Generator

The Poincaré data generator is one of the central production components of ORBIT because it is responsible for creating the trajectory datasets from which much of the later analysis is derived. In the current implementation of the thesis workflow, this role is carried by the unified generation pipeline, which transforms the mathematical procedures established in Chapter 3 into an executable process that can generate large sets of initial conditions, propagate the corresponding trajectories, identify surface-of-section crossings, and record the resulting map structure in an organized project format.

Conceptually, this module occupies the point at which the theoretical framework becomes a repeatable computational experiment. Rather than computing a single trajectory in isolation, the generator systematically sweeps across selected initial-condition ranges and Jacobi-constant levels, applies the active dynamical model, and assembles the resulting outputs into structured datasets. These outputs include the time-domain trajectory histories, the reduced surface-of-section data needed for phase-space analysis, and the initial spectral quantities retained at the dataset level. In this way, the generator acts as the main bridge between the dynamical equations of Chapter 3 and the large-scale map-based investigations carried out in the later results chapters.

The dashboard tab through which this generator is configured and launched is shown in Figure 4.2. The screen exposes the principal controls used in the thesis workflow, including the Jacobi level, the x_0 sweep bounds, the surface-of-section generation settings, and the associated base frequency-analysis parameters that are recorded alongside the trajectory outputs.

Within ORBIT, the Poincaré generator is not implemented as an isolated standalone program exposed directly to the user. Instead, it is integrated into the broader workflow through the dashboard and project structure. The user selects the relevant system configuration, specifies the numerical ranges and model options, and then launches the generation process from the interface. Those inputs are passed to the underlying unified generator, which performs the actual propagation, event handling, and dataset assembly before writing the resulting products into the active project directory. This arrangement preserves the scientific integrity of the underlying generator while making it practical to use repeatedly across many runs.

An important feature of this design is that the generator is paired naturally with the rest of the ORBIT analysis pipeline. The same datasets created for Poincaré-map construction also support subsequent frequency analysis, candidate identification, and stability assessment. As a result, the generator is not merely a map-making utility; it is the front end of a broader data-production process that feeds the remainder of the computational framework. By structuring the workflow in this

ORBIT Data-Generation Screen

System Summary

System	Name	x (ND)	m (kg)	J2	Lagrange points (x, y, C)
Primary	Neptune	-0.0002098980	1.024092e+26	3.411000e-03	L1: 0.95915562, 0.00000000, 3.01458411
Secondary	Triton	0.9997901020	2.150000e+22	-	L2: 1.04155629, 0.00000000, 3.01430422
μ (ND)		0.0002098980			L3: -1.00008746, 0.00000000, 3.00020990
Scale (ND)		1 ND = 354,760.000 km			L4: 0.49979010, 0.86602540, 2.99979015
					L5: 0.49979010, -0.86602540, 2.99979015

Use J2 model: ☐ Off Orbit display: System: Retrograde

Generate Data

Data Generation Properties

μ (auto): 0.000209898009092977

CJ: e.g., 3.8801495841 ☐ Auto (Lagrange array)

x0 min: x0 max:

dx (x0 step): 0.05

Poincaré Properties

Crossing target (analysis only): 1000

☐ J2 On

Frequency Properties

t_max: 9000

dt: 0.05

Modes are automatic: prograde + retrograde are both generated. Collision and escape are recorded as flags; the generator always runs to t_max.

Ratio / Resonance Generation

Existing CJ: 2.999790146

Back Change Projects Settings Save Exit

Figure 4.2. Data-generation tab of the ORBIT dashboard. The screen provides the controls used to configure the unified trajectory-generation workflow, including the Jacobi level, the initial-condition sweep, the Poincaré crossing settings, and the base frequency-analysis parameters associated with each run.

way, ORBIT ensures that the information required for later analytical and stability work is created consistently at the moment the trajectories are generated, while still allowing additional downstream frequency-ratio and resonance products to be built later from the stored datasets rather than from ad hoc manual reconstruction.

The corresponding Poincaré-analysis view used to inspect the resulting surface-of-section datasets is shown in Figure 4.3. In practice, this screen provides the immediate bridge between trajectory generation and phase-space interpretation by allowing the user to filter the generated runs and open the associated Poincaré plotting and table tools directly from the same project context.

4.5 Frequency Analysis Module

The frequency analysis module extends the Poincaré generation workflow by extracting the spectral information needed to interpret the underlying structure of the computed trajectories. As described in Chapter 3, frequency analysis provides a complementary view of the dynamics by identifying the dominant oscillatory content of a trajectory and relating that content to resonance structure, quasi-periodic behavior, and candidate periodic-orbit families. Within ORBIT, this analysis now proceeds in two connected stages: an initial spectral pass performed alongside the unified trajectory-generation workflow, followed by a dedicated downstream ratio and resonance stage that operates on the saved frequency products associated with the generated dataset.

ORBIT Poincaré-Analysis Screen




Figure 4.3. Poincaré-analysis tab of the ORBIT dashboard. The screen provides dataset filters and access to the plotting and table tools used to inspect the surface-of-section outputs generated by the unified trajectory workflow.

In practical terms, this module begins with the time-series outputs associated with the generated trajectories and applies the frequency-extraction procedures described in Chapter 3 during the unified generation stage itself. That base stage writes the trajectory histories and the primary spectral products into the project dataset, after which the ratio/resonance workflow can be launched as a separate analysis pass from the dashboard. The downstream pass then reads the saved frequency table for the selected dataset, assembles the preferred ratio products used in the later analysis, and generates the associated resonance tables. This design ensures that the transition from geometric phase-space interpretation to spectral characterization does not require a disconnected manual workflow. Instead, the trajectory data and their associated frequency products remain linked within the same project environment.

The dashboard tab associated with this workflow is shown in Figure 4.4. It provides direct access to the saved frequency runs and to the plotting tools used to review the ratio maps, frequency tables, and spectrum products derived from the generated trajectory datasets.

This integration is especially important because the frequency results are not treated as isolated output values. They serve as the basis for later tasks such as resonance identification, trajectory classification, and candidate selection for stability studies. For this reason, the frequency module is best understood not as a standalone plotting utility, but as a core analytical stage in the ORBIT workflow. It takes the raw dynamical behavior encoded in the propagated trajectories, converts that behavior into stored spectral products during the base generation pass, and then organizes the

ORBIT Frequency-Analysis Screen




Figure 4.4. Frequency-analysis tab of the ORBIT dashboard. The screen organizes access to the saved frequency runs and to the plotting tools used to inspect frequency-ratio trends, data tables, and associated spectral products within the active project.

derived ratios and resonance relationships in a form that can be compared systematically across large datasets.

The module also contributes to the usability of the framework by reducing the amount of manual post-processing required from the user. Rather than exporting time histories and analyzing them in a separate external environment, ORBIT keeps both the base spectral extraction and the later ratio/resonance generation inside the same project-driven workflow. This preserves continuity between initial-condition selection, orbit propagation, Poincaré-map construction, and frequency interpretation, while ensuring that the frequency outputs remain consistent with the physical model and numerical settings associated with each project run.

4.6 Stability Analysis Module

The stability analysis module represents the most selective stage of the ORBIT workflow, focusing not on broad trajectory sweeps but on the detailed evaluation of candidate periodic or near-periodic structures identified in the earlier phases of the analysis. Whereas the Poincaré and frequency modules are designed to survey large regions of phase space and reveal the overall organization of the dynamics, the stability module concentrates on particular trajectories or candidate families and assesses their local dynamical character through periodic-orbit refinement and linear stability analysis.

This module is therefore positioned naturally at the downstream end of the framework. The preceding ORBIT components generate the data needed to recognize potentially important structures in the Neptune–Triton problem, while the stability analysis stage provides the tools needed to examine those structures in greater depth. In the workflow used for this thesis, candidate trajectories are first identified from the geometric and spectral outputs of the earlier modules, and only then are they passed into the stability pipeline for more focused computation. This preserves a logical progression from broad exploration to detailed characterization.

The interface used for this stage is shown in Figure 4.5. The screen collects the saved candidate targets, provides controls for launching stability runs, and displays the resulting classifications and scalar stability measures in a project-linked analysis table.

ORBIT Stability-Analysis Screen




Figure 4.5. Stability-analysis tab of the ORBIT dashboard. The screen displays the stored target list, the controls used to launch stability runs, and the tabulated results used to classify the candidate trajectories identified in the earlier analysis stages.

As established in Chapter 3, the stability analysis itself is based on periodic-orbit targeting, state-transition-matrix propagation, and the evaluation of the resulting Floquet multipliers and scalar stability measures. ORBIT does not reframe these procedures as a separate mathematical model; instead, it packages them into a dedicated module that can be applied systematically to trajectories selected from the larger data products. In this way, the framework maintains continuity between the earlier analytical stages and the later dynamical interpretation of orbit families.

The importance of integrating this capability into ORBIT lies in the fact that stability analysis depends strongly on the provenance of the selected candidates. A useful stability workflow requires

a clear connection between the candidate orbit, the system configuration from which it was drawn, the model options active at the time of generation, and the associated frequency and phase-space analysis. By keeping these elements inside the same project-centered environment, ORBIT reduces ambiguity in the transition from exploratory analysis to targeted stability computation and supports a more traceable interpretation of the final results.

4.7 Graphical User Interface

The graphical user interface is the visible layer through which ORBIT is accessed and controlled. While the preceding sections described the underlying computational modules that carry out the physics, numerical generation, and analysis tasks of the system, the purpose of the interface is to organize those capabilities into a coherent user workflow. In the context of this thesis, the interface is not merely a convenience feature added after the scientific tools were developed; it is the mechanism through which the full Neptune–Triton analysis pipeline is coordinated, repeated, and managed at the project level, as shown in Figure 4.6.

ORBIT Startup Screen



Figure 4.6. Startup screen of the ORBIT graphical user interface for the Neptune–Triton analysis framework. The figure shows the application entry point and project-oriented environment through which the ORBIT workflow is initiated.

The role of the interface is therefore twofold. First, it presents the major framework capabilities in a structured and accessible form, allowing the user to move from project creation to system definition, model configuration, data generation, and post-processing without leaving the program

environment. Second, it acts as a control boundary between user input and protected scientific logic. The user interacts with validated fields, selections, and workflow screens, while the lower-level physics and numerical modules remain insulated from direct modification during ordinary operation. This balance between accessibility and control is one of the principal design features of ORBIT.

Another important aspect of the interface is continuity. The workflow implemented in ORBIT is not a collection of independent windows opened in arbitrary order, but a connected sequence of screens that reflect the logical progression of the research process. The user first establishes or selects a project, then defines the dynamical system, proceeds to the dashboard for generation and analysis control, and finally accesses the dedicated analysis and plotting screens associated with the generated data products. This ordering mirrors the methodology developed in Chapter 3 and helps ensure that each stage of the workflow is grounded in a consistent system configuration and project state.

For this reason, the graphical user interface is best understood as the operational face of the system rather than as a separate software subsystem detached from the scientific work. It provides the means by which the methods of the preceding chapters are carried into repeated computational practice, and it determines how the user experiences the architecture described earlier in this chapter. The following discussion therefore focuses on the major interface screens and the roles they play in linking project organization, model selection, data generation, and analysis into a unified workflow.

4.7.1 Project Creation and Selection

The first functional stage of the ORBIT interface is the project-selection workflow. After the application is launched, the user is directed to a project screen where an existing project may be opened or a new one may be created. This design reflects the fact that ORBIT is organized around persistent research projects rather than around one-off computational sessions. Each project acts as a container for the system definition, numerical settings, generated datasets, and exported analysis products associated with a specific study configuration.

This project-based structure is important for both usability and reproducibility. From the standpoint of usability, it gives the user a clear organizational unit within which all computations and outputs are grouped. Rather than manually managing unrelated files generated at different stages of the workflow, the user works within a single project context that remains active throughout the session. From the standpoint of reproducibility, the project structure helps preserve the connection between a selected physical system, the numerical settings used to analyze it, and the outputs generated from those settings. This becomes especially valuable when multiple Neptune–Triton runs, model variants, or energy regimes are being compared within the same broader investigation.

The project selector therefore serves a more significant role than a simple file-open dialog. It marks the point at which the framework establishes the active working context for the remainder of the analysis. Once a project has been opened or created, the corresponding state is carried forward into the later screens, allowing system definitions, generator settings, and analysis outputs to remain associated with that project without repeated manual re-entry. In this way, the project screen provides the organizational foundation upon which the rest of the graphical workflow is built.

ORBIT System-Definition Screen

System Definition

System Input Mode

- Option 1 — Select primary/secondary bodies from catalog
- Option 2 — Enter masses + radii manually
- Option 3 — Enter μ only (no bodies/masses)

Option 1 Inputs (Catalog)

Primary:

Secondary:

m1 (Neptune): 1.024092e+26 kg

m2 (Triton): 2.150000e+22 kg

R1 (Neptune): 2.462200e+07 m

R2 (Triton): 1.352600e+06 m

μ (computed): 0.0002098980

Option 2 Inputs (Manual)

Primary name:

m1 (kg):

R1 (m):

μ (computed): 0.0002098980

Secondary name:

m2 (kg):

R2 (m):

Option 3 Inputs (μ only)

μ :

Back

Continue

Figure 4.7. System-definition screen reached immediately after creating a new ORBIT project. The interface presents three input pathways: selecting the primary and secondary bodies from the catalog, entering the masses and radii manually, or specifying the reduced mass parameter μ directly for an abstract CR3BP study.

4.7.2 System Definition

Once a project has been established, the next stage of the ORBIT workflow is the system-definition screen (Figure 4.7), where the physical two-body system underlying the restricted three-body model is specified. This screen provides the connection between the abstract framework of the CR3BP and the concrete physical system selected for study. In the present thesis, that system is the Neptune–Triton pair, but the interface is designed so that the underlying framework can also accept other primary-secondary combinations or simplified parameter-based inputs when needed.

The system-definition stage is especially important because it determines the physical inputs from which the normalized model parameters are derived. Rather than forcing the user to work directly with low-level internal quantities alone, the interface provides multiple pathways for defining the system. A user may select bodies from a stored catalog, enter masses and radii manually, or specify the mass parameter directly when a simplified or abstract model is desired. These options allow the framework to support both physically grounded configurations and more generalized theoretical studies while still routing the result into a consistent internal representation.

From the standpoint of the present work, this screen serves as the formal point at which the Neptune–Triton problem is instantiated inside the framework. Once the system definition is accepted, the corresponding derived quantities, such as the mass parameter and associated physical configuration, become part of the active project state. All subsequent stages of the workflow then

inherit these values. This is an important design choice because it reduces the risk that later modules will be run under inconsistent assumptions about the active system. The user does not repeatedly redefine the physical configuration from screen to screen; instead, the selected system propagates forward through the remainder of the analysis pipeline.

This screen also illustrates a recurring design principle of ORBIT: the interface exposes scientifically meaningful choices while shielding the user from unnecessary implementation detail. The user chooses bodies, masses, radii, or the mass parameter in forms that are easy to interpret physically, while the framework handles the translation of those inputs into the internal configuration required by the computational modules. In this way, the system-definition screen provides both accessibility and rigor, allowing the physical problem to be established clearly before the numerical workflow begins.

4.7.3 Dashboard and Workflow Control

After the system has been defined, ORBIT transitions to the dashboard (Figure 4.8), which functions as the primary operational hub of the framework. This screen occupies a central role in the program because it is the point from which the user launches generation routines, selects analysis modes, configures model options, and navigates to the specialized modules described earlier in the chapter. While the project screen establishes the working context and the system-definition screen establishes the active physical model, the dashboard is where the user begins to execute the actual computational workflow.

The importance of the dashboard lies in its ability to connect high-level research questions to the specific generator and analysis tools required to address them. From this screen, the user can define the Jacobi-constant regime of interest, choose trajectory-generation settings, activate or deactivate the relevant model options, and launch the numerical routines that produce the Poincaré, time-series, and base frequency data products. The dashboard also provides access to the later ratio/resonance workflow, allowing an already generated dataset to be revisited and extended without rerunning the full propagation stage. In effect, the dashboard condenses the computational control structure of ORBIT into a single coordinated interface rather than distributing those responsibilities across disconnected scripts or menus.

This screen also plays an important role in maintaining continuity between modules. Because the dashboard remains connected to the active project state, it can pass a consistent set of system parameters, model settings, and numerical controls into the lower-level routines that perform the actual computations. The user therefore interacts with the workflow through a centralized control surface, while the underlying modules inherit the necessary information from the shared state maintained by the application. This arrangement simplifies repeated studies and reduces the likelihood of launching different parts of the framework with mismatched assumptions.

In practical use, the dashboard is the point at which ORBIT becomes most visibly integrated as a research environment. It is where model definition, numerical execution, and downstream analysis begin to converge in a single operational screen. For this reason, the dashboard may be understood as the principal control center of the interface: it does not replace the specialized modules that carry out the detailed work, but it provides the structured access point from which those modules are configured and invoked.

ORBIT Dashboard Screen




Figure 4.8. Dashboard screen of the ORBIT graphical user interface for the Neptune–Triton framework. The figure shows the central project-control screen from which model settings, numerical generation steps, and downstream analysis modules are coordinated.

4.7.4 Analysis and Visualization Screens

Beyond the dashboard, ORBIT provides a set of dedicated analysis and visualization screens through which specific data products can be inspected in greater detail. These screens include the modules associated with Lagrange-point analysis, zero-velocity analysis, and the plotting tools used to visualize phase-space and spectral outputs. Their purpose is to separate detailed interpretation tasks from the general control logic of the dashboard while still preserving continuity with the active project and model state.

This separation is useful because the products generated by the framework differ substantially in character. Some outputs are most naturally interpreted as tables of equilibrium properties or stability quantities, while others are best understood through graphical views of zero-velocity structures, Poincaré maps, or frequency distributions. ORBIT therefore treats analysis and visualization as specialized stages of the workflow, each with its own interface context, rather than forcing all outputs into a single overloaded control screen. This improves readability and allows the user to focus on one class of result at a time without losing access to the broader project environment.

These screens also strengthen the connection between computation and interpretation. The user does not simply launch a generator and then inspect files externally; instead, the generated data can be reviewed within the same framework that produced it. This allows the phase-space structures discussed in Chapters 2 and 3 to be examined directly as part of the computational workflow. In

this sense, the analysis and visualization screens function as the interpretive layer of the interface, translating raw numerical outputs into organized views that support scientific reasoning and later reporting.

Equally important, these screens maintain the same controlled relationship to the underlying computation as the rest of the GUI. They present derived tables, summaries, and plots based on the project data, but they do not alter the protected physical and numerical logic beneath the framework. Their role is therefore observational and analytical rather than definitional. This design preserves consistency while making the framework much more practical for iterative research use, since the user can inspect the outputs of a run, refine the setup if needed, and continue the analysis without leaving the ORBIT environment.

4.7.5 Project Data Organization and Export

An essential feature of ORBIT is that it does not treat generated outputs as temporary by-products of an interactive session. Instead, the framework organizes them as persistent components of the active project. Trajectory data, derived output tables, frequency outputs, and exported figures are written into a structured project directory so that the results of a computational run remain associated with the system definition and numerical settings from which they were produced. This project-centered storage strategy is one of the mechanisms by which the framework supports reproducibility.

The importance of this design becomes clear when a study requires repeated runs across multiple Jacobi constants, model variants, or numerical settings. Without a structured data model, the user would be required to manually track which files correspond to which physical assumptions, a process that is both inefficient and highly vulnerable to error. ORBIT instead preserves this association internally by grouping outputs under the active project and by routing exports through the same organized directory structure. In this way, the project is not only a container for the interface state, but also the persistent record of the computational investigation itself.

This storage model also supports the later phases of the analysis pipeline. Because the frequency and stability modules depend on the provenance of previously generated trajectory data, the framework benefits from a directory structure in which earlier outputs can be located and reused reliably. In the current workflow this includes not only the base trajectory and Poincaré products, but also the later frequency-run products, run catalogs, and resonance summaries generated from those stored datasets. The organized project format therefore supports more than archival convenience; it is a functional requirement for a workflow in which the outputs of one stage become the inputs of the next. By managing this structure internally, ORBIT reduces the burden of manual file handling and helps maintain continuity across the full sequence of computations.

Export functions are integrated into this same organizational approach. Tables, reports, and figures generated within the interface can be written back into the project in a controlled and predictable way, allowing the user to preserve intermediate analyses as well as final products. In the case of figures, the export path is structured to write both PDF and PNG versions of the same plot so that a single analysis product can move directly into thesis preparation, presentation graphics, and general screen-based review without requiring a separate conversion step. This is particularly valuable for a thesis-oriented workflow, where plots and tables often move from exploratory analysis into formal presentation. ORBIT therefore supports not only the execution of the computational study, but also the orderly preservation and extraction of the results needed for documentation and interpretation.

An additional practical feature of the system is that the core generator workflow was not restricted to the local graphical environment alone. For large trajectory campaigns, a reduced batch-execution path was maintained for use on the University of Oklahoma Supercomputing Center for Education & Research (OSCER), where computational loads exceeding the practical limits of a personal workstation could be executed more efficiently. In this workflow, the essential physics and numerical logic were preserved, but the run was launched through a minimal command-line control path rather than through the full graphical interface. The resulting data products were then transferred back into the ORBIT project structure for inspection, plotting, and subsequent analysis. This arrangement made it possible to scale the workflow to larger parameter sweeps while preserving continuity between high-performance batch generation and the interactive project-based environment used for interpretation.

4.8 Verification and Validation

The usefulness of a computational framework such as ORBIT depends not only on its organizational clarity, but also on the degree to which its outputs can be trusted to reproduce the intended mathematics and numerical behavior. For this reason, verification and validation were treated as integral parts of the development process rather than as post hoc checks applied only after the interface had been assembled. In the context of the present work, verification refers primarily to confirming that the software implementation reproduces the governing equations, analytical quantities, and numerical procedures described in Chapter 3, while validation refers more broadly to confirming that the resulting outputs exhibit the expected dynamical structures for the Neptune–Triton problem.

Several aspects of the framework support this verification process directly. The separation of the protected physics layer from the user-facing interface reduces the risk that the governing equations or constants will be altered inadvertently during ordinary use. The project-based workflow preserves the configuration context associated with each run, making it possible to trace a given output back to its defining parameters and numerical settings. In addition, the modular structure of the framework allows individual components, such as Lagrange-point calculations, zero-velocity analysis, frequency extraction, and stability computations, to be examined independently before being integrated into the full interface workflow.

For a research framework such as ORBIT, however, internal consistency alone is not sufficient. It is also necessary to demonstrate that the software can reproduce known CR3BP behavior in controlled cases outside the primary Neptune–Triton application. To address this requirement, two external checks were emphasized during development: a prescribed- μ test case with $\mu = 0.3$, used repeatedly while reproducing the CR3BP structures and frequency-analysis behavior described in the cited prior work [21], and a Sun–Earth benchmark tied to the zero-velocity and Poincaré-map examples presented in *Dynamical Systems, the Three-Body Problem, and Space Mission Design* by Koon *et al.* [9]. These examples do not replace the Neptune–Triton analysis of the thesis, but they do provide concrete evidence that the ORBIT implementation reproduces expected CR3BP structures outside the primary application system.

4.8.1 Controlled Verification Case at $\mu = 0.3$

The first validation example is a controlled CR3BP test problem defined directly by the reduced mass parameter $\mu = 0.3$. This case was used throughout development as a practical verification system

because it allowed the same governing equations, normalization logic, and analysis routines used in the thesis workflow to be exercised without tying the test to a separate catalog-based planetary model. It was also the mass-parameter setting used while recreating the phase-space and frequency-analysis results discussed earlier in the thesis, making it a natural benchmark for checking that the ORBIT implementation remained consistent with the equations and solution procedures adopted for this work.

This $\mu = 0.3$ case provides more than a single numerical spot check. Once the system is defined, ORBIT must compute the equilibrium points, assemble the effective potential, evaluate the Jacobi structure, propagate trajectories, construct the Poincaré section, and extract the dominant frequency information from the resulting time series. Because all of these tasks are driven by the same user-defined mass ratio, the case serves as an end-to-end verification problem for both the protected physics layer and the analysis pipeline built on top of it. Agreement across these outputs is therefore evidence that the software is not merely generating plausible individual plots, but is carrying the same underlying model consistently through each computational stage.

The first pair of products checks the configuration-space quantities derived directly from the equations of motion. The computed Lagrange-point geometry (Figure 4.9) shows the expected asymmetric arrangement associated with a relatively large secondary mass, while the zero-velocity curves (Figure 4.10) reproduce the corresponding partition of the admissible and forbidden regions in the rotating frame. Together, these plots verify that the equilibrium-point solver, effective-potential evaluation, and zero-velocity routines are working coherently for the prescribed value of μ .






Figure 4.9. Computed Lagrange-point geometry for the controlled verification case with $\mu = 0.3$. This figure verifies that the equilibrium-point solver reproduces the expected asymmetric CR3BP configuration for the prescribed mass parameter.

The second pair of products tests the downstream dynamical analysis routines that operate on propagated trajectories. The Poincaré map reveals organized phase-space structure rather than a scattered or numerically unstable cloud, indicating that the trajectory generator, event-detection logic, and section-construction routines are acting consistently. The corresponding frequency-ratio


Mu = 0.3 Zero-Velocity Verification




a.) $C_J = 3.92014958$



b.) $C_J = 3.55641300$



c.) $C_J = 3.29135022$



d.) Combined
Lagrange-point contours

Figure 4.10. Zero-velocity curves in configuration space for the CR3BP at a distinct Jacobi level, showing how the topology of the accessible regions changes as the energy varies. The openings near the collinear equilibria illustrate the transport channels that emerge at critical Jacobi constants.

plot extracted from the same dataset exhibits a smooth, interpretable trend, showing that the spectral analysis module is recovering coherent oscillatory information from the generated trajectories; both products are shown in Figure 4.11. This was especially important because the same frequency-analysis machinery is later used in the Neptune–Triton study and was cross-checked against the phase-space and frequency results from the cited prior work.




Figure 4.11. Verification products for the controlled case with $\mu = 0.3$ using $C_J = 3.88514958412578$, shown using the prograde (blue) family. Left panel (a): Poincaré map in $(x, \dot{x})$, demonstrating the organized surface-of-section structure produced by the trajectory generator and crossing logic. Right panel (b): corresponding frequency-ratio trend extracted from the same prograde family, confirming that the spectral-analysis pipeline recovers a smooth and interpretable dynamical relationship from the propagated trajectories.

Taken together, the $\mu = 0.3$ results provide a direct verification example for the equations, numerical routines, and post-processing steps implemented in ORBIT. They show that a single prescribed mass parameter can be passed consistently through Lagrange-point computation, zero-velocity analysis, trajectory generation, Poincaré construction, and frequency extraction while preserving the expected qualitative CR3BP structure at each stage. For the purposes of this thesis, that consistency is the key validation requirement: it demonstrates that the program is accurately implementing the intended mathematical framework before those same tools are applied to the Neptune–Triton problem.

It is also important to note that, in the present thesis, ORBIT is not being validated as a general-purpose commercial analysis package intended for unrestricted use across arbitrary problems. Rather, it is being documented and assessed as the research environment developed to support the specific Neptune–Triton investigations reported here. Its verification and validation are therefore tied to the internal consistency of the implemented methods, the reproducibility of the generated outputs, and the agreement of those outputs with the theoretical and benchmark expectations established in the earlier chapters and in the reproduced literature examples.

4.9 Framework Limitations

Although ORBIT provides a unified and effective environment for the computational work of this thesis, it is important to define its scope clearly. The framework was developed as a research tool tailored to the needs of the Neptune–Triton study presented here, not as a fully generalized celestial-mechanics platform intended to support every possible restricted multi-body workflow. As a result, some design choices were made in favor of research efficiency and methodological consistency within this thesis rather than maximal generality.

One important limitation is that the framework is built primarily around the planar dynamical models and analysis procedures developed in Chapter 3. While certain components of the codebase are more general in principle, the implemented workflow documented in this chapter is centered on the planar CR3BP, the Neptune-primary J_2 extension, and the associated Poincaré, frequency, and stability analyses used in the present study. The framework therefore should be interpreted in the context of this thesis-specific application rather than as a finished universal environment for all orbit-dynamics investigations.

Another limitation is that the graphical interface is designed as a controlled research workflow rather than as an unrestricted programming environment. This is intentional and, in many respects, advantageous, because it protects the verified physics and numerical logic from accidental modification during ordinary use. However, this also means that extending the framework to incorporate substantially different models, new analysis pipelines, or alternative numerical strategies generally requires development at the code level rather than simple configuration through the interface. ORBIT is thus flexible within its intended domain, but not infinitely configurable from the GUI alone.

Finally, the framework should be understood as one component of the larger research process rather than as a substitute for theoretical interpretation. ORBIT organizes calculations, manages data, and provides structured access to analytical outputs and visualizations, but the scientific meaning of the results still depends on the analytical framework established in Chapters 2 and 3 and on the interpretation developed in the later results chapters. In that sense, ORBIT is best viewed as the computational instrument of the thesis: powerful and carefully structured, but always operating in support of the underlying dynamical analysis rather than replacing it.

Taken as a whole, ORBIT provides the computational bridge between the theoretical and numerical methods developed in the earlier chapters and the dynamical results presented in the chapters that follow. By unifying protected physics models, repeatable generator workflows, integrated analysis modules, and a project-centered graphical interface within a single framework, it makes the Neptune–Triton study both executable and reproducible at scale. With this computational structure now established, the thesis turns next to the phase-space results generated within ORBIT, beginning in Chapter 5 with the dynamical structures revealed by the computed Neptune–Triton trajectory datasets.

5. Geometric Phase-Space Exploration

This chapter presents the phase-space exploration of the Neptune–Triton system within the framework of the planar Circular Restricted Three-Body Problem (CR3BP), with a J_2 perturbation applied to Neptune through the effective potential formulation. Having established the theoretical foundations and computational methods in the preceding chapters, the analysis now focuses on the dynamical behavior that emerges within this specific gravitational environment. The Neptune–Triton system provides a particularly rich setting for this investigation due to its unique mass ratio, retrograde orbital configuration, and resulting phase-space structure.

The primary objective of this chapter is to identify and characterize the organization of phase space through the generation and analysis of large trajectory sets. By examining the resulting motion using geometric tools, including Poincaré maps and direct trajectory visualization, the underlying structure of the system can be revealed. These methods allow for the identification of invariant structures and the classification of trajectories according to their dynamical behavior.

Particular attention is given to the identification of periodic, quasi-periodic, resonant, and chaotic regions of motion, as well as to the transitions between these regimes as system parameters vary. The resulting phase-space structures provide insight into the accessibility, transport pathways, and long-term behavior of trajectories within the Neptune–Triton system. These geometric results establish the framework that is refined in the following chapters through frequency-based classification, candidate selection, and stability analysis.

5.1 Interpretation of Motion in the Neptune–Triton System

In order to interpret the phase-space structures presented in this chapter, it is necessary to briefly consider the relationship between the rotating reference frame used in the CR3BP-based formulation and the inertial frame in which the physical motion of the system occurs. The equations of motion are expressed in a synodic (rotating) reference frame that rotates with constant angular velocity n about the system barycenter. In this frame, the two primary bodies remain fixed along the x -axis, while the motion of the third body is observed relative to this rotating coordinate system. The physical motion, however, occurs in the inertial frame, where the system rotates through an angle nt with respect to the rotating frame. This distinction is important when interpreting trajectory direction and velocity components in the Neptune–Triton system.

A key characteristic of the Neptune–Triton system is that Triton orbits Neptune in a retrograde direction relative to Neptune’s rotation and the general prograde motion of most bodies in the solar system. In the inertial frame, Triton therefore moves opposite to the conventional prograde orientation typically assumed in standard CR3BP examples. As a result, the rotating reference frame defined by the Neptune–Triton system reflects this retrograde motion, and the apparent direction of motion within the rotating frame differs from the intuition developed from prograde systems such as Earth–Moon or Sun–Jupiter. This reversal introduces an important distinction in how trajectory motion and velocity components are interpreted in the phase space.

This difference is particularly relevant when interpreting velocity components in the rotating frame. In planar CR3BP analyses, the direction of motion is often inferred from the sign of the velocity components, particularly $\dot{y}$ when examining crossings of a Poincaré map defined along the x -axis. In prograde systems, a positive $\dot{y}$ is typically associated with motion in the same

direction as the rotating frame, while a negative $\dot{y}$ corresponds to motion in the opposite direction. In the Neptune–Triton system, however, the retrograde motion of Triton effectively reverses this interpretation when viewed relative to inertial motion. Consequently, trajectories that appear to follow standard directional conventions in the rotating frame may correspond to physically retrograde motion in the inertial frame. This distinction is critical when associating velocity sign with physical trajectory direction.

This reversal directly impacts the visualization and interpretation of the phase-space structures presented in this chapter. Poincaré maps and trajectory plots generated in the rotating frame may appear inverted when compared to analogous results from prograde systems. Regions of phase space traditionally associated with prograde motion may instead correspond to retrograde motion in the inertial frame, and vice versa. Care must therefore be taken when assigning physical meaning to the orientation and structure of trajectories observed in the rotating frame. To clarify this distinction, Figure 5.1 illustrates the relationship between inertial and rotating frames for both a standard prograde system and the retrograde Neptune–Triton system.

Retrograde Frame Interpretation




Figure 5.1. Comparison of prograde and retrograde motion in a rotating reference frame. The top panel shows a standard prograde system, while the bottom panel shows the Neptune–Triton retrograde system. The rotating frame (x, y) is used for computation and visualization, while the inertial frame (X, Y) represents the physical motion of the system. In the rotating frame, the retrograde Neptune–Triton configuration produces a reflected orientation relative to the prograde case. The velocity component $\dot{y}$ is shown for reference, where $\dot{y} > 0$ and $\dot{y} < 0$ indicate the direction of motion across the Poincaré map. The sizes of the primary bodies are not to scale and are shown for visual perspective only.

With this distinction established, the phase-space structures presented in the following sections can be interpreted consistently within the context of the Neptune–Triton system.

5.2 Phase-Space Mapping Results

The phase-space structure of the Neptune–Triton system is investigated through the construction of Poincaré maps over a range of Jacobi constants. These maps provide a reduced representation of the system dynamics, allowing complex trajectories to be interpreted in terms of their recurrent intersections with a defined surface of section. By examining the distribution and organization of these intersections, it is possible to identify regions of regular, resonant, and chaotic motion within the system.

In the present work, Poincaré sections are generated using the J_2 -perturbed formulation described in Chapter 3. As a result, the structures observed in the following analysis reflect the influence of oblateness on the underlying dynamical environment. Rather than reproducing the classical circular restricted three-body problem, these results characterize the modified phase-space geometry associated with the Neptune–Triton system under perturbation.

The analysis proceeds by first identifying the global organization of the phase space and introducing a regional decomposition based on observed trajectory behavior. Subsequent sections examine how this structure evolves with Jacobi constant and explore the internal features of each region, including resonant patterns and frequency characteristics.

5.2.1 Global Structure of the Phase Space

The global dynamical structure of the Neptune–Triton system is examined using Poincaré maps constructed in the $(x, \dot{x})$ phase plane. The surface of section is defined by the crossing condition $y = 0$ with $\dot{y} > 0$, producing a one-sided mapping of trajectory intersections with the synodic plane. Figure 5.2 presents a Poincaré section at a Jacobi constant of $C_J = 3.018641378$, computed using the J_2 -perturbed model. All trajectories were generated with collision and escape termination enabled, and event-flagged trajectories are clipped at the first recorded termination time so that only the bounded portion of each orbit is retained. The resulting map provides a clear snapshot of the global phase-space structure at this energy level and serves as the basis for identifying distinct dynamical regions.




Figure 5.2. Global Poincaré section of the Neptune–Triton system in the $(x, \dot{x})$ plane at $C_J = 3.018641378$ for the J_2 -perturbed model. Shading indicates the four primary dynamical regions discussed in the text; the boundaries are approximate and are used for interpretive clarity.

At $C_J = 3.018641378$, the phase space exhibits a consistent and interpretable organization that can be decomposed into four primary dynamical regions. These regions are identified based on qualitative differences in trajectory behavior, density of crossings, and the presence or absence of coherent invariant structures.

1. **Region I:** A sparse outer region characterized by weakly structured motion and a low density of crossings. Trajectories in this region exhibit limited coherence and do not form well-defined invariant structures.
2. **Region II:** A domain dominated by nested quasi-periodic invariant curves, indicating regular, stable motion. This region exhibits clear island structures associated with resonant and near-resonant trajectories.
3. **Region III:** The largest and most dynamically complex region, containing a mixture of coherent invariant structures and irregular, scattered crossings. This region is primarily associated with retrograde motion and exhibits significant variability in trajectory behavior.
4. **Region IV:** A narrow bounded region located near Triton, visible only upon zooming. This region contains tightly clustered crossings corresponding to localized motion near the secondary body.

The boundaries between these regions are approximate and are introduced for interpretive clarity; they do not correspond to exact dynamical separatrices. Nevertheless, this regional decomposition provides a useful framework for understanding the large-scale organization of the system. Across the range of Jacobi constants examined in this work, these four regions persist as identifiable features of the phase space, although their relative size, structure, and degree of coherence vary with energy.

This global classification serves as the foundation for the subsequent analysis, where the internal structure of each region, including resonant features and frequency characteristics, is examined in greater detail.

5.2.2 Evolution of Phase-Space Structure with Energy

The structure of the Poincaré map evolves systematically as the Jacobi constant is varied across the 13 energy levels examined in this study. While the regional classification introduced in Figure 5.2 remains useful for organizing the phase space, the size, shape, and internal structure of each region change with energy, as illustrated across four representative Jacobi constants in Figure 5.3, reflecting shifts in the underlying dynamical accessibility of the system.

At higher Jacobi constants (just below C_{L1}), the Poincaré map is dominated by relatively constrained motion. Region II is compact and well-defined, with tightly organized invariant curves indicating predominantly regular behavior. Region III is present but limited in extent, and the overall phase space is largely partitioned into coherent structures with minimal scattering.

As the Jacobi constant decreases into the range between C_{L1} and C_{L2} , the accessible phase space expands. Region III grows significantly in size, and the boundary between Regions II and III becomes less sharply defined. Scattered crossings begin to appear more frequently, indicating the emergence of increasingly irregular motion. Region I becomes more populated as trajectories gain access to a broader range of spatial configurations.

Energy-Dependent Poincaré Evolution




Figure 5.3. Poincaré maps in the $(x, \dot{x})$ plane for the J_2 -perturbed Neptune–Triton system at four Jacobi constants, with the surface of section defined by $y = 0$ and $\dot{y} > 0$. The panels are ordered from higher to lower Jacobi constant and show the expansion of accessible phase space together with the increasing prominence of scattered crossings as the energy decreases.

Further decreasing the Jacobi constant into the range between C_{L2} and C_{L3} produces a pronounced increase in dynamical complexity. Region III becomes the dominant feature of the Poincaré map, occupying a substantial portion of the accessible phase space. The distribution of crossings in this region becomes more diffuse, and the distinction between coherent and irregular motion becomes increasingly blurred. Region II persists but is reduced in relative size and exhibits deformation in its invariant structures.

At lower Jacobi constants approaching the range between C_{L3} and C_{L4} , the Poincaré map exhibits its most expanded and irregular configuration within the analyzed dataset. Region III remains dominant, with widespread scattering across the phase space. Region I continues to expand outward, while Region II becomes increasingly fragmented. Region IV, the narrow bounded region near Triton, remains present across all energy levels but is consistently confined to a small portion of the phase space and is only clearly resolved under localized zoom.

Across all energy levels considered, the evolution of the Poincaré map reflects a transition from predominantly regular, organized motion at higher Jacobi constants to increasingly complex and irregular behavior as the energy decreases. This progression provides a global view of how accessible phase space and trajectory behavior change within the J_2 -perturbed Neptune–Triton system.

5.3 General Dynamical Signatures in the Poincaré Maps

Having established the global organization of the Neptune–Triton phase space and its evolution with Jacobi constant, the next step is to identify the recurring dynamical signatures that appear throughout the Poincaré maps. These signatures provide the basis for interpreting the large trajectory sets generated in this study. In particular, they allow regions of regular, resonant, transitional, and chaotic motion to be distinguished directly from the geometry of the surface-of-section points. In this chapter, terms such as periodic, quasi-periodic, and chaotic are used to describe the qualitative geometric character of the motion as revealed by the Poincaré map. This usage is distinct from the resonance-based interpretation developed later from the frequency content of the trajectories, and also from the stability classification introduced in the subsequent stability analysis.

The most ordered structures in the maps appear as smooth invariant curves. These curves are produced by trajectories that remain confined to regular quasi-periodic motion on invariant tori in the full phase space. On the surface of section, such motion generates continuous closed loops or nested families of loops, indicating that the trajectory repeatedly returns to the section in a geometrically coherent manner. In the Neptune–Triton results, these structures are most clearly visible in the more regular portions of Region II and in selected localized subregions of Region III, with a representative example shown in Figure 5.4. Their persistence across successive crossings indicates bounded motion with low local sensitivity to perturbation.

Resonant behavior is identified through the appearance of island chains embedded within or adjacent to these regular regions. Rather than filling a continuous closed curve, a resonant trajectory returns to the section through a finite number of preferred phases, producing a repeating chain of discrete islands. In the present dataset, such structures are especially important because they mark the locations where the smooth invariant organization of the map is interrupted by commensurabilities between the underlying frequencies of motion. These island chains, illustrated in Figure 5.5, therefore serve as geometric markers of resonant families that will later be examined using the frequency-ratio analysis.




Figure 5.4. Example of a smooth invariant curve on the Poincaré map corresponding to a quasi-periodic trajectory confined to an invariant torus in the full phase space. This figure was taken from Region IV of the Poincaré map at $C_J = 3.018641378$. The trajectory repeatedly intersects the surface of section, producing a continuous closed loop that indicates regular, bounded motion. In the Neptune–Triton system, such curves are indicative of stable regions of phase space where trajectories exhibit low sensitivity to perturbations and remain confined over long timescales.




Figure 5.5. Local view of a resonant island chain in the Neptune–Triton Poincaré map at $C_J = 3.0072570569$, shown over the window $x \in [0.300, 0.315]$ and $\dot{x} \in [-0.25, 0.25]$. The repeated closed islands mark a resonant family in which the section crossings cluster at a discrete set of phases rather than filling a continuous invariant curve.

Between the most ordered invariant curves and the most diffuse scattered regions lie transitional domains in which the phase-space structure begins to lose coherence. In these areas, invariant curves may become distorted, broken, or interwoven with scattered crossings, and the boundaries between regular and irregular motion become less sharply defined. This behavior is particularly evident as the Jacobi constant decreases and Region III expands, as shown in Figure 5.6. The resulting patterns indicate that the local dynamics are no longer confined to a simple invariant-torus structure and instead reflect the onset of resonance overlap, weak instability, or transport across neighboring regions of phase space.




Figure 5.6. Retrograde-only local view of the Neptune–Triton Poincaré map at $C_J = 3.0072570569$, shown over the window $x \in [0.20, 0.26]$ and $\dot{x} \in [-1.5, 1.5]$. The central chain and the associated crossing structure illustrate the loss of simple invariant-curve organization in a locally unstable region, where nearby resonant structures and separatrix-type crossings produce more complicated phase-space geometry.

The most irregular signatures in the maps appear as diffuse clouds or scattered bands of points that do not organize into smooth curves or isolated island chains. These scattered distributions are interpreted as chaotic or weakly chaotic layers, in which nearby initial conditions separate rapidly and repeated crossings no longer populate a simple geometric object on the section. In the Neptune–Triton system, such regions become increasingly prominent at lower Jacobi constants, where the opening of energetically accessible pathways allows trajectories to explore a broader portion of the phase space. The growth of these scattered layers is one of the clearest visible indicators of structural breakdown as the energy decreases.

Taken together, these signatures provide a practical visual language for interpreting the Neptune–Triton Poincaré maps. Smooth invariant curves indicate regular bounded motion, island chains identify resonant structure, transitional regions reveal the deformation or breakup of ordered behavior, and diffuse scattered layers signal the presence of increasingly irregular dynamics. This classification is used throughout the remainder of the chapter to interpret the prograde and retrograde domains, identify resonant families, and select candidate trajectories for subsequent frequency and stability analysis.

5.4 Directional Organization of the Phase Space

Although the retrograde character of the Neptune–Triton rotating frame was established in the opening section of this chapter, the Poincaré maps themselves show that the phase space is not merely globally retrograde in a descriptive sense. Instead, the maps separate into domains that are dynamically dominated by different senses of motion relative to the selected reference body. This directional decomposition provides an additional layer of interpretation beyond the simple identification of regular and irregular structures, because it clarifies how the large-scale geometry of the map is tied to the actual character of the underlying trajectories.

The clearest directional partition appears between the compact, highly ordered structures of Region II and the broader mixed domain occupying most of Region III, as illustrated in Figure 5.7. Although both regions contain trajectories that remain dynamically associated with Neptune, their distributions on the Poincaré map differ substantially, revealing distinct phase-space organization rather than motion about different central bodies. Region II is dominated by prograde motion and is characterized by nested invariant curves that remain geometrically coherent over a wide range of crossings. These curves indicate trajectories that remain tightly confined and whose motion is organized around relatively regular local families. In contrast, Region III is dominated by retrograde motion and occupies a substantially larger portion of the accessible phase space. This region contains both smooth invariant structures and scattered layers, indicating that retrograde trajectories in the Neptune–Triton system support a much wider range of dynamical behaviors than the more compact prograde family.




Figure 5.7. Poincaré maps in the $(x, \dot{x})$ plane for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$, shown separately for prograde and retrograde trajectory sets. The surface of section is defined by $y = 0$ with $\dot{y} > 0$, and only the bounded portion of each filtered orbit family is retained. The comparison shows the compact, regular organization of the prograde domain and the broader mixed structure of the retrograde domain.

This distinction is significant because it shows that directional classification is not simply a matter of reversing visual intuition from a standard prograde CR3BP example. Rather, the sense

of motion is reflected directly in the geometry of the map. Prograde trajectories are concentrated in a comparatively narrow, well-organized region of phase space, while retrograde trajectories populate the dominant mixed region in which regular, resonant, and irregular structures coexist. This asymmetry is a dynamical consequence of how the spacecraft interacts differently with Triton depending on its direction of motion. Because Neptune lies so close to the system barycenter, its gravitational influence is relatively insensitive to orbital orientation; it is the spacecraft's proximity to and interaction with Triton that drives the directional asymmetry observed in the maps. The resulting structural difference is one of the defining features of the Neptune–Triton phase space and is visible across the full range of Jacobi constants examined in this work.

The boundary between these directional domains is not a sharp separation in a strict analytical sense, but the maps nevertheless show a persistent transition zone in which the coherence of the prograde structures begins to weaken and the broader retrograde domain appears more dominant. It should be noted that both families are shown in the same coordinate frame as a visualization convenience; this apparent shift reflects the greater spatial extent of the retrograde population in the phase plane, not a physical displacement of prograde trajectories. As the Jacobi constant decreases, this transition becomes less sharply defined, and scattered crossings appear more frequently between the most ordered invariant curves. In this sense, the prograde–retrograde decomposition is closely related to the energy-dependent structural breakdown discussed in the preceding section: as the Jacobi constant decreases and the accessible phase space expands toward Triton, its gravitational perturbation grows in relative influence, simultaneously diffusing the boundaries of both the directional domains and the regular invariant structures, leaving them recognizable but with increasingly diffuse internal boundaries as additional transport pathways and irregular layers develop.

Region IV, located near Triton, provides a useful counterpoint to this broader partition. Although narrow in spatial extent, it contains localized bounded structures that remain coherent even when the surrounding global phase space becomes increasingly mixed. These localized structures show that directional organization is not restricted to the two dominant large-scale domains alone. Instead, the maps contain smaller pockets of bounded motion whose directional character must be interpreted in the context of their local geometry rather than only their position within the global map.

For the purposes of the present study, the directional organization of the phase space therefore serves two roles. First, it provides a physically meaningful way to partition the map into broad dynamical domains before examining specific resonant families. Second, it helps identify where different classes of candidate trajectories should be sought. The prograde-dominated region is associated primarily with compact regular families, whereas the retrograde-dominated region contains the richer mixture of resonant and irregular behavior that becomes central to the later trajectory selection and stability analysis. This directional interpretation forms the natural bridge between the general geometric signatures identified above and the more detailed examination of resonant island structure that follows.

5.5 Geometric Synthesis and Transition

The Poincaré maps at various energy levels show that the J_2 -perturbed Neptune–Triton system possesses a rich geometric phase-space structure that can be organized at several levels. At the global scale, the Poincaré maps reveal a persistent regional decomposition whose relative size and coherence evolve with the Jacobi constant. At a more local scale, the maps exhibit the full range of

familiar dynamical signatures, including smooth invariant curves, resonant island chains, transitional breakup zones, and diffuse scattered layers associated with increasingly irregular motion. The directional decomposition into prograde- and retrograde-dominated domains adds an additional level of structure by showing that different portions of the phase space support different classes of bounded motion even when all trajectories remain associated with the same Neptune–Triton dynamical environment.

These geometric results provide the essential first stage of the Neptune–Triton analysis, but they do not by themselves supply a complete dynamical classification of the observed structures. In particular, the appearance of an island chain or a coherent region in the Poincaré map does not uniquely determine the associated frequency ratio, confirm a resonance assignment, or establish whether the trajectory family identified in that region is dynamically robust under more refined analysis. The most useful structures in the Poincaré maps are not the broadly scattered or weakly organized regions, but the compact island chains and smooth invariant neighborhoods that suggest repeatable and geometrically coherent motion. These structures provide the strongest geometric foundation for identifying trajectory families that may later prove useful as candidate pathways in the Neptune–Triton tour design problem. The geometric phase-space exploration therefore serves as a structured map of where important dynamical features are located, while the subsequent chapters use frequency-based analysis and targeted candidate examination to determine which of those features correspond to the most significant and practically useful trajectory families.

6. Frequency and Resonance Analysis

The geometric phase-space exploration presented in Chapter 5 established the large-scale organization of the Neptune–Triton system and identified the principal domains of regular, resonant, transitional, and chaotic motion visible in the Poincaré maps. Geometry alone, however, does not reveal the underlying dynamics responsible for those structures. The positions and shapes of invariant curves and island chains in the map reflect the spacecraft’s interactions with both primaries, and the frequency commensurabilities that govern those interactions cannot be inferred from visual inspection alone. In particular, smooth invariant curves, resonant island chains, and scattered layers may be recognized visually, but the associated commensurabilities and frequency families must be examined directly through spectral analysis.

This chapter therefore turns to the frequency-analysis capabilities of ORBIT and applies them to the same bounded trajectory sets used in the preceding phase-space study. The objective is to characterize the internal frequency structure of the Neptune–Triton system, identify approximate resonances, and interpret how those resonances organize the phase-space features observed in the Poincaré maps. The resulting frequency-ratio products provide a complementary analytical framework that extends the geometric interpretation developed in Chapter 5 and supports the later selection of candidate trajectory families for more detailed analysis.

6.1 Frequency Analysis Framework

The theoretical basis for frequency analysis was established in Section 2.8, where quasi-periodic motion was described in terms of a small number of fundamental frequencies and resonance was related to commensurate frequency relationships. In the present chapter, those ideas are applied directly to the Neptune–Triton trajectory sets generated by ORBIT. Rather than treating the Poincaré map and the frequency spectrum as separate analytical tools, the analysis developed here uses them as complementary descriptions of the same underlying motion. The Poincaré map reveals the geometry of the trajectories on a surface of section, while the frequency products quantify the oscillatory relationships responsible for that geometry.

As discussed in Section 2.8, two frequencies are said to be commensurate when their ratio is rational. In principle, this provides a direct way to identify resonance from the extracted frequency ratio. In practice, however, requiring an exactly rational ratio would return very few cases because the computed frequencies are affected by numerical error, finite sampling, and spectral resolution. For this reason, the analysis must also admit ratios that are approximately rational. This provides a larger set of prospective resonant targets, but it also introduces the possibility of false positives and false negatives. The frequency ratios therefore provide a compact way to organize large trajectory sets and to relate smooth invariant curves, island chains, and transitional structures in the Poincaré map to corresponding spectral behavior.

Within ORBIT, the frequency workflow begins with the time-series data generated during the unified trajectory propagation process described in Chapter 3. Those stored trajectories are used to compute dataset-level frequency products, including the extracted ratio values and the derived resonance classifications used throughout the present analysis. These products are then used to construct the frequency maps, resonance overlays, and family classifications examined in the remainder of this chapter. In this sense, the frequency-analysis module is not a separate

dynamical model, but an analytical layer built on top of the same Neptune–Triton trajectories already interpreted geometrically in Chapter 5.

It is important, however, to distinguish the interpretation of resonance from the numerical procedure used to identify it. The resonance products developed here are constructed from extracted frequency ratios and finite tolerance thresholds, so the resulting labels should be interpreted as numerical guides rather than confirmed dynamical classifications. Accordingly, the resonance computations performed by ORBIT are used primarily to identify regions of interest and organize the phase space for further analysis, rather than to replace geometric inspection or later stability assessment.

6.2 Frequency Ratio Plots

The first global visualization product used in this chapter is the frequency ratio plot. This plot is constructed in the $(x_0, \omega_2/\omega_1)$ plane, where the horizontal axis corresponds to the sampled initial-condition coordinate and the vertical axis represents the extracted frequency ratio associated with each trajectory. The base plot is color-coded by trajectory type, so that prograde and retrograde trajectories may be distinguished directly in ratio space. This representation is used to examine how the large trajectory sets organize spectrally across the sampled domain before any explicit resonance labeling is applied.

The value of this plot lies in the fact that it reveals structure that is not immediately obvious from the Poincaré map alone. Ordered trajectory families tend to form coherent bands or clustered patterns in ratio space, whereas more irregular spectral behavior tends to produce scattered or poorly organized distributions. In particular, horizontal or nearly horizontal groupings provide a useful visual indication of families whose spectral behavior varies slowly across a range of initial conditions. For this reason, the frequency ratio plot serves as a practical guide for identifying regions that merit closer resonance analysis, especially near low-order rational ratios.

The representative case shown in Figure 6.1 is drawn from the J_2 -perturbed Neptune–Triton runs at $C_J = 3.018641378$, consistent with the full set of computations used throughout this project. For the dataset used to construct these figures, the underlying trajectories were propagated to $t_{\max} = 15000$ nondimensional time units and sampled for frequency analysis with $\Delta t_{\text{FFT}} = 0.01$ over a fast Fourier transform (FFT) window of $T_{\text{FFT}} = 15000$ (see Subsection 3.7.3 and Subsection 3.8.1 for the formal definitions of these parameters). The downstream frequency-classification run applied $T_{\min} = 5000$ and $T_{\text{trajectory},\min} = 100$. Taken together, these settings ensured that the spectral products were built from long, uniformly sampled records while still excluding trajectories whose usable duration was too short for reliable comparison. The long propagation and analysis windows were needed to preserve frequency resolution and keep the horizontal band structure in the ratio plot sharply defined, whereas the $T_{\text{trajectory},\min} = 100$ cutoff acted as a data-quality filter that prevented short or poorly resolved cases from contaminating the global spectral map.

Within the Neptune–Triton workflow, the frequency ratio plot is therefore interpreted as a spectral analogue to the geometric phase-space maps of Chapter 5. It does not by itself confirm resonance, but it helps reveal where the ratio distribution is smooth, where it breaks sharply, and where families of trajectories appear to cluster about potentially significant values. These visual patterns provide the first step in narrowing the large dataset to a smaller set of candidate resonant neighborhoods.

Global Frequency-Ratio Structure




Figure 6.1. Frequency ratio plots in the $(x_0, \omega_2/\omega_1)$ plane for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. The full plot shows the global spectral organization of the sampled trajectories, while the zoomed panel highlights a horizontal band structure used to identify candidate resonant neighborhoods.

6.3 Resonance Identification and Family Construction

The role of resonance identification can be illustrated clearly using the working case $C_J = 3.018641378$. At this Jacobi constant, the frequency ratio plot already reveals that the trajectories do not populate ratio space randomly. Instead, the points organize into visible bands and clusters, some of which appear to gather near low-order rational values. These visual patterns provide the first indication that the dataset may contain structured resonance families, but they do not yet assign a formal resonance label to any individual trajectory.

That assignment is carried out in a separate downstream stage of the ORBIT workflow, where the extracted ratio values are compared against candidate rational ratios using a prescribed numerical tolerance. For each trajectory with a valid extracted ratio, the program compares the computed value of ω_2/ω_1 to a library of rational targets of the form of Equation (2.46),

$$\frac{p}{q}, \quad p, q \in \mathbb{Z},$$

subject to the selected maximum denominator and tolerance threshold. If the computed ratio lies within the prescribed tolerance of one of these rational values, the trajectory is flagged as resonant and assigned the corresponding label. Ratios that do not satisfy any admissible match remain unclassified. In the $C_J = 3.018641378$ example, this step converts the visually suggestive groupings in the ratio plot into explicit candidate resonances that can be tracked systematically rather than interpreted by eye alone.

In the present study, the resonance matching was carried out with a tolerance of 0.001. This value was chosen as a practical compromise between two competing requirements. If the tolerance is chosen too small, trajectories that are genuinely associated with a resonance may fail to be identified because the extracted ratio is affected by finite sampling, spectral resolution, and numerical error. If

the tolerance is chosen too large, trajectories that are not truly associated with the resonance may be grouped into that class. The adopted value of 0.001 was therefore used to capture the visibly coherent ratio bands in plots such as Figure 6.1 without letting the classification spread so broadly that nearby non-resonant trajectories were absorbed indiscriminately. For this reason, the resonance labels produced by ORBIT are interpreted as approximate numerical classifications rather than exact dynamical truths. In the present example, as in the full dataset, the classification should therefore be understood as a practical guide for locating promising resonance neighborhoods, not as a final proof that every flagged trajectory belongs to a rigorously confirmed resonant structure.

Once the resonant trajectories have been identified, the program groups all trajectories sharing the same assigned resonance into a common resonance family. Each family is then given a unique display color, allowing the classified results to be visualized consistently across multiple plotting environments. In the $C_J = 3.018641378$ case, this step transforms the broad ratio plot from a collection of spectrally organized points into a set of color-coded families associated with specific approximate resonances. This family structure is especially useful because it creates a direct bridge from numerical ratio matching to the overlay products discussed in the next section.

For the present working example, the 1:2 resonance family is used as an illustrative example of the ORBIT family-classification process. A subset of the trajectories assigned to this family is listed in Table 6.1. Even in this reduced view, the table shows the practical meaning of the ORBIT classification workflow: trajectories drawn from different parts of the sampled x_0 range, and from both prograde and retrograde orbit types, can still be grouped into the same approximate resonance family when their extracted frequency ratios fall within the selected tolerance of the rational value 1/2.

Table 6.1. Representative trajectories assigned to the 1:2 resonance family for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. The listed rows are drawn from the ORBIT resonance-family output and show how the extracted ratio values cluster about the rational target 1/2 within the selected tolerance.

x_0	Orbit Type	ω_2/ω_1	Approx. Ratio	Resonance Error
−0.510079	Retrograde	0.499980	1/2	2.00×10^{-5}
−0.500079	Retrograde	0.499971	1/2	2.90×10^{-5}
−0.920079	Prograde	0.499962	1/2	3.85×10^{-5}
0.219921	Retrograde	0.500117	1/2	1.17×10^{-4}
0.209921	Retrograde	0.500125	1/2	1.25×10^{-4}
−0.900079	Retrograde	0.499706	1/2	2.94×10^{-4}

6.3.1 Worked Example: Frequency Extraction and Resonance Classification

To make the resonance identification procedure concrete, this subsection traces the trajectory at $x_0 = 0.2099207594$ through the complete frequency-extraction and classification pipeline. This retrograde orbit at $C_J = 3.018641378$ under the J_2 -perturbed model corresponds directly to the fifth row of Table 6.1 and is used here to show explicitly how the numerical entries in that table are produced.

Figure 6.2 shows the trajectory in configuration space over the full propagation interval $t \in [0, t_{\max}]$ (see Subsection 3.7.3). The orbit remains bounded throughout the integration and exhibits the characteristic looping structure associated with Neptune-centered retrograde motion.

The regular, repeating geometry of the path is a qualitative indicator that the trajectory belongs to a structured dynamical family, though it does not by itself identify the resonance.

Worked Example: Retrograde Trajectory at $x_0 = 0.2099$




Figure 6.2. Configuration-space trajectory for the retrograde orbit at $x_0 = 0.2099207594$ and $C_J = 3.018641378$ under the J_2 -perturbed Neptune–Triton model. The bounded, regularly repeating path is consistent with membership in a structured resonant family.

To extract the dynamical frequencies, the x -coordinate time series $s(t_k) = x(t_k)$ is sampled on a uniform grid with $\Delta t_{\text{FFT}} = 0.01$ over an analysis window of $T_{\text{FFT}} = 15000$ as defined in Subsection 3.8.1. The resulting normalized frequency spectrum is shown in Figure 6.3.

The dominant spectral peak identifies the primary frequency:

$$\omega_1 = 0.159227. \quad (6.1)$$

The secondary frequency is extracted using the spectral sideband spacing procedure [21], which locates the symmetric pair of spectral peaks flanking ω_1 and computes ω_2 as half their total separation:

$$\omega_2 = 0.079633. \quad (6.2)$$

The frequency ratio is then computed directly from Equation (2.45):

$$\frac{\omega_2}{\omega_1} = \frac{0.079633}{0.159227} = 0.500125. \quad (6.3)$$

Comparing this value to the candidate rational ratio $p/q = 1/2 = 0.500000$, the resonance error is

$$\varepsilon = \left| \frac{\omega_2}{\omega_1} - \frac{1}{2} \right| = 1.25 \times 10^{-4}, \quad (6.4)$$

which falls well within the adopted tolerance of 0.001. The trajectory is therefore classified as a member of the 1:2 resonance family. This result corresponds directly to the fifth row of Table 6.1, confirming that each table entry is the direct numerical product of the pipeline described in Chapter 3.




Figure 6.3. Normalized frequency spectrum for the retrograde trajectory at $x_0 = 0.2099207594$ and $C_J = 3.018641378$. The dominant peak identifies the primary frequency ω_1 , and the symmetric sideband pair flanking it is used to extract the secondary frequency ω_2 via the sideband spacing method [21].

This trajectory will appear again in the resonance overlay products of the following section, where it is recolored according to its 1 : 2 family assignment and projected back onto both the frequency ratio plot and the Poincaré map. Having traced this single example from initial condition to resonance label, the reader can use that familiarity to interpret the overlay figures with a concrete understanding of the classification step they represent.

The resonance families therefore serve as the bridge between raw frequency-ratio data and the overlay products examined in the next section. The ratio plot reveals where structure exists, the tolerance-based classification determines which trajectories are treated as resonant, and the family construction organizes those trajectories into a form that can be projected back onto both the spectral and geometric maps. In this way, the working example at $C_J = 3.018641378$ illustrates the central interpretive step of the Chapter 6 workflow before the family colors are applied back onto the frequency and phase-space plots.

6.4 Resonance Overlays on Frequency and Phase Space

Once resonance families have been assigned, ORBIT uses those classifications to generate overlay products that connect the numerical resonance labels directly to the visual structures identified in the results. These overlays are among the most useful outputs of the frequency-analysis workflow because they allow the same resonance family to be traced across both spectral and geometric representations. Rather than viewing the resonance table as an isolated numerical artifact, the overlay process projects its classifications back into the plotting environment used to interpret the dynamics.

The first overlay product is applied to the frequency ratio plot itself. In this view, trajectories that have been assigned to a resonance family are recolored according to the family label, while trajectories that remain unclassified are shown in gray. The resulting classification is shown in

Figure 6.4 and differs from the base ratio plot of Section 6.2 by replacing the broad prograde–retrograde distinction with an explicit resonance-family classification. This makes it possible to see which portions of the ratio distribution belong to the same resonance family, how widely each family extends in the sampled domain, and where non-resonant trajectories remain intermixed with the classified resonant structures.




Figure 6.4. Overlay of resonance-family classifications on the frequency ratio plot for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. Each color corresponds to a different resonance family identified by ORBIT, while gray points represent trajectories that were not classified as resonant within the selected tolerance. This overlay allows the resonance families to be traced directly in ratio space, revealing how they organize the spectral distribution and where they cluster about rational values.

The second overlay product applies the same resonance-family colors to the Poincaré map. This step is particularly important because it links the spectral classifications back to the invariant curves, island chains, and transitional regions introduced in Chapter 5. When the family colors are projected onto the surface-of-section points, as shown in Figure 6.5, resonant structures that were previously interpreted only geometrically can now be associated with specific approximate resonance families. In this way, the Poincaré map no longer displays only geometric organization; it also becomes a resonance-classified phase-space map.

Taken together, these two overlay products provide a direct comparison between frequency space and phase space. The frequency overlay shows how resonance families are distributed in $(x_0, \omega_2/\omega_1)$, while the Poincaré overlay shows how those same families occupy invariant structures in $(x, \dot{x})$. Agreement between the two strengthens the interpretation of coherent resonant organization, while disagreement or fragmentation may indicate transitional behavior, tolerance sensitivity, or incomplete family assignment. The overlay framework therefore serves as a critical interpretive bridge between the numerical classifications generated by ORBIT and the physical dynamical structures discussed throughout this thesis.




Figure 6.5. Overlay of resonance-family classifications on the Poincaré map for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. Each color corresponds to a different resonance family identified by ORBIT, while gray points represent trajectories that were not classified as resonant within the selected tolerance. This overlay allows the resonance families to be traced directly in phase space, revealing how they organize the invariant structures and where they cluster about rational values.

6.5 Individual Trajectory Frequency Behavior

While the frequency ratio plots and resonance overlays provide a global view of the spectral organization of the Neptune–Triton dataset, they do not show how the frequency content is assembled for any individual trajectory. For that reason, ORBIT also provides trajectory-level frequency plots that allow a single case to be examined in greater detail. These products are used to inspect the dominant spectral peaks, evaluate the local clarity of the extracted ratio, and determine whether a candidate resonance classification is supported by an organized underlying frequency structure.

This trajectory-level view is particularly useful when interpreting points that lie near a rational ratio in the global frequency plot. In some cases, the corresponding spectrum exhibits a clean and well-separated set of peaks, making the associated ratio straightforward to interpret. In other cases, the spectrum is more diffuse, contains competing peaks of comparable amplitude, or exhibits sideband structure that complicates the extraction. The individual trajectory plots therefore provide an important intermediate view between the large-scale resonance maps and the later dynamical interpretation of specific families.

For the present study, these trajectory-level products are not treated as standalone results in the same sense as the global ratio plots or the resonance overlays. Instead, they are used as supporting analytical views when a candidate family, horizontal ratio band, or resonance-classified structure requires closer inspection. In this role, they help distinguish between trajectories whose frequency behavior is coherent enough to support the resonance interpretation and trajectories whose extracted ratio may be more sensitive to local numerical or spectral ambiguity.

Accordingly, the individual frequency plots serve as a verification layer within the Chapter 6 workflow. The global maps identify where structure exists, the tolerance-based classification organizes that structure into candidate families, and the trajectory-level spectral plots provide a closer look at how particular cases realize those classifications. This makes them especially useful when moving from broad resonance mapping toward the more focused selection of candidate regions and families developed in the following chapters.

Trajectory-Level Frequency Spectra




Figure 6.6. Representative trajectory-level spectra for the J_2 -perturbed Neptune–Triton system at $C_J = 3.018641378$. The left panel gives a comparatively clean low-order resonant case whose dominant frequencies are easy to identify, while the right panel shows a trajectory whose denser secondary content makes the spectral interpretation less immediate.

The comparison in Figure 6.6 illustrates how the trajectory-level spectra were used in practice

throughout this project. In the easier case shown in the left panel, the dominant peak may be identified cleanly as the primary frequency f_1 , and the secondary structure is separated well enough that the associated sideband spacing or second component can be read with comparatively little ambiguity. In that situation, the corresponding ratio is straightforward to interpret, and the trajectory can be associated confidently with the band structure seen in the global frequency-ratio plot.

The harder case in the right panel shows why this analytical step is necessary. Although the dominant peak remains visible, the surrounding spectral content is much denser, with multiple nearby peaks and a broader cluster of competing secondary features. For trajectories of this type, the ratio cannot be assigned responsibly by eye from the global map alone. Instead, the local spectrum must be inspected to determine whether a coherent secondary spacing is actually present and whether the extracted ratio should be trusted as evidence of organized quasi-periodic behavior or treated more cautiously as a numerically ambiguous case.

In this way, the individual spectra provide the practical reading rule for the Chapter 6 resonance workflow: first identify the dominant peak f_1 , then inspect the nearby secondary peaks or sideband structure used to define the secondary frequency contribution, and finally compare the resulting ratio with the rational values highlighted in the global maps and family classifications. When the spectrum is clean, this procedure strengthens the resonance interpretation. When the spectrum is crowded or diffuse, it signals that the corresponding ratio and resonance label should be interpreted with more caution.

6.6 Interpretation of Resonant Structure

Taken together, the frequency-resonance overlays in Figure 6.7 show that the resonance organization is not reconfigured dramatically over the present C_J range. Instead, the dominant result is one of qualitative persistence. Comparing the cases at $C_J = 3.018641378, 3.0143741909, 3.0107806367$, and 2.999790146 , the classified points continue to concentrate into a small number of recurring horizontal spectral groupings rather than dispersing uniformly across the $(x_0, \omega_2/\omega_1)$ plane. This indicates that the resonance-identification procedure is isolating a stable set of preferred frequency relationships throughout the sampled domain. The principal change with energy is therefore not the appearance of an entirely different resonance architecture, but rather a change in how densely and continuously those same resonance neighborhoods are populated.

The clearest evidence of that persistence is provided by the strong high-ratio structures. In each of the sampled cases, a prominent classified band remains visible near $\omega_2/\omega_1 \approx 0.89$, and a smaller set of classified points also recurs near $\omega_2/\omega_1 \approx 1.0$. These repeated concentrations are important because they show that the most coherent resonance families identified by ORBIT are not confined to a single isolated energy level. Instead, they persist across the sampled C_J values as robust spectral features. In practical terms, these high-ratio families appear to define the clearest and most consistently recovered part of the resonance picture developed in this chapter: they are narrow, repeatedly recovered, and visually distinct from the surrounding unclassified background. This makes them stronger candidates for physically meaningful resonance structure than isolated single points or loosely scattered classifications.

By contrast, the lower-ratio portion of the overlays is much less coherent. Below roughly $\omega_2/\omega_1 \approx 0.3$, the classified points appear only as sparse clusters or isolated detections, and the intermediate bands near ratios such as 0.33, 0.40, 0.50, and 0.67 are present but comparatively weak and fragmented. These families are not absent, but they do not organize the ratio plane with the




Figure 6.7. Comparison of frequency-resonance overlays for four representative J_2 -perturbed Neptune–Triton energy levels. Colored points denote trajectories assigned to approximate resonance families by ORBIT, while light gray points remain unclassified within the selected tolerance. Viewed together, these overlays show which horizontal resonance bands persist across the sampled C_J range, which families remain narrow and strongly populated, and which lower-ratio structures appear only as sparse or fragmented classifications.

same clarity as the dominant high-ratio structures. The overlay at $C_I = 3.0107806367$ is especially useful in this regard because it appears slightly sparser overall, reinforcing the point that the weaker resonance families are more sensitive to energy level and less continuously populated than the dominant bands. At this stage, the frequency-resonance overlays establish the spectral side of the Chapter 6 interpretation: a few strong horizontal families persist across the sampled energies, while the lower-ratio population remains fragmented and should be interpreted more cautiously. The next step is to project these same family labels back onto the Poincaré maps and examine where the strongest spectral resonances are realized geometrically in phase space.

The Poincaré resonance overlays in Figure 6.8 show that the dominant spectral families identified in ratio space do project onto coherent geometric structures rather than appearing as disconnected numerical labels. In the higher and intermediate energy cases, the strongest families trace large portions of the principal lobes and their bounding invariant envelopes, especially on the right-hand side of the map where the family structure is richest. This is an important confirmation of the Chapter 6 interpretation, because it shows that the recurring horizontal bands observed in frequency space correspond to organized regions of phase space rather than to isolated classifications with no clear geometric realization.

At the same time, the Poincaré overlays show that the resonance-family coverage is not uniform across the section. Some geometrically ordered regions remain largely gray, while others contain several interwoven family colors rather than one single dominant classification. This is especially noticeable in the lower-energy cases, where the extended left-hand structures and the more intricate internal arcs of the right-hand lobe appear more fragmented. That geometric fragmentation is consistent with the spectral result that the lower-ratio and intermediate-ratio resonance populations are weaker and less continuously populated than the dominant high-ratio bands. Taken together, the frequency and Poincaré resonance overlays therefore support the same conclusion: a small number of strong resonance families persist across the sampled energies and map onto recognizable invariant structures, while the weaker resonance population remains more diffuse and should be interpreted with greater caution.

Taken together, the frequency-resonance overlays and their projections onto the Poincaré maps show that only a subset of the identified resonance families remain both spectrally persistent and geometrically coherent across the sampled Neptune–Triton energy levels. In the present results, the strongest candidates for later study are the narrow, repeatedly recovered high-ratio families, since these appear most consistently in both frequency space and phase space. By contrast, the weaker lower-ratio and intermediate-ratio populations are more fragmented and energy-sensitive, and therefore provide less reliable targets for immediate follow-on analysis. The combined spectral and geometric agreement developed in this chapter thus serves as a screening step for the later selection of trajectory families and candidate pathways relevant to the Neptune–Triton tour design problem.

6.7 Chapter Synthesis

This chapter has extended the geometric phase-space interpretation of Chapter 5 by showing that the bounded Neptune–Triton trajectory sets also possess an interpretable spectral organization. Rather than appearing as an unstructured collection of extracted ratios, the trajectories organize into recurring resonance-supported groupings, only some of which remain coherent enough to matter for later analysis. The important result is therefore not simply that resonances can be labeled

Poincaré Resonance Comparison Across Energy




Figure 6.8. Comparison of resonance-family projections onto the Poincaré maps for four representative J_2 -perturbed Neptune–Triton energy levels. The same families identified in frequency-ratio space are recolored here in phase space, allowing the persistent spectral bands of Fig. 6.7 to be traced onto the corresponding invariant structures. Gray points indicate trajectories that were not assigned to a resonance family within the selected tolerance.

numerically, but that the spectral products distinguish between strong families that persist across energy and weaker families that appear only in more fragmented form.

The central result of the chapter is that only a subset of the identified resonance families remain both spectrally persistent and geometrically coherent across the sampled J_2 -perturbed energy levels. In particular, the strongest high-ratio families recur across multiple C_J values as narrow, clearly populated bands in frequency space and as recognizable invariant structures in the Poincaré overlays. By contrast, the lower-ratio and intermediate-ratio populations appear more fragmented, more weakly populated, and more sensitive to energy level. The trajectory-level frequency spectra reinforce this distinction by showing that some classifications are supported by clean local spectral structure, while others arise in cases whose secondary frequency content is more ambiguous.

Taken together, these results provide the spectral screening step for the remainder of the thesis. They show which structures identified geometrically in Chapter 5 are supported by persistent resonance organization, which ones remain too fragmented to interpret confidently as robust families, and which candidate regions deserve closer follow-on study. In this way, Chapter 6 narrows the large Neptune–Triton trajectory database to a smaller set of resonance-family candidates that can be examined more directly in the later candidate-selection, stability-analysis, and mission-architecture stages of the thesis.

7. Candidate Region Identification

The results of Chapters 5 and 6 established that the J_2 -perturbed Neptune–Triton system contains a wide range of bounded, resonant, transitional, and transport-capable structures. Taken in aggregate, however, those results remain broader than what is needed for the design logic of a Neptune–Triton tour. A practical mission concept does not require an exhaustive catalog of every trajectory family present in the phase space; it requires a smaller and more defensible set of candidate regions and candidate trajectories whose behavior is coherent enough to support repeated operations, or sufficiently accessible to support motion between dynamical neighborhoods. The purpose of the present chapter is therefore to move from global dynamical characterization to candidate identification.

For that reason, Chapter 7 is not yet intended to declare which trajectories should be accepted as mission-ready solutions. That judgment depends on the stability analysis of Chapter 8 and the broader mission interpretation developed in Chapter 9. Instead, the task here is to determine which portions of the previously analyzed phase space are promising enough to merit continued study. The chapter therefore organizes the search around four mission-oriented candidate classes: Neptune-centered bounded candidates, structured candidates in the broader touring region, Triton-centered bounded candidates, and bottleneck-crossing transport candidates. Within each class, the comparison across the selected Jacobi levels is used to distinguish persistent, well-organized structures from those that are either too fragmented or too weakly expressed to justify follow-on analysis.

7.1 Mission Interpretation for a Neptune–Triton Tour

In the present chapter, the term *candidate region* denotes a portion of the Neptune–Triton phase space whose geometry and spectral structure suggest that it may support a recognizable mission function. One such function is that of a bounded observation region, in which trajectories remain confined near Neptune or Triton and therefore provide repeated, geometrically coherent motion about one of the primaries. Regions of this type are of interest because they may support local tour phases, repeated viewing opportunities, or staging behavior that remains dynamically organized over many section crossings. A second function is that of a transport region, in which trajectories approach or cross the bottleneck separating major dynamical neighborhoods. These trajectories are not valuable because they are maximally regular, but because they may provide access between regions that remain disconnected at higher Jacobi constants.

From the standpoint of tour design, these two roles should be viewed as complementary rather than competing. A practical Neptune–Triton campaign would likely require both locally coherent trajectories that support repeated observation and transfer-type trajectories that permit motion between distinct neighborhoods of the system. The mission interpretation developed here is therefore not limited to asking which structures are the most regular in a purely mathematical sense. Instead, it asks which structures are most useful as local tour segments, staging regions, or inter-region pathways within the larger dynamical environment.

7.2 Neptune-Centered Bounded Candidates

The first mission-relevant candidate class considered here consists of trajectories that remain bounded in the Neptune-centered portion of the phase space. These candidates are attractive because they provide the clearest foundation for local observation segments, repeatable viewing geometries, and other mission phases that require motion to remain dynamically organized over many section crossings. Unlike the bottleneck-crossing trajectories examined later, the aim of this section is not to identify access between regions, but to isolate coherent neighborhoods of motion that remain confined near Neptune and can therefore serve as plausible local tour candidates.

Within the four reference cases considered throughout this chapter,

$$C_J = 3.018641378, \quad 3.0143741909, \quad 3.0107806367, \quad 2.999790146.$$

the Neptune-centered bounded structures separate naturally into two useful groups. Region II contains the most compact resonance-supported islands, while Region III extends that bounded behavior into larger-amplitude prograde and retrograde motion that remains clearly organized in the Poincaré maps. The discussion therefore begins with the compact 8:9 family in Region II before moving outward to the broader bounded structures in Region III. This progression is deliberate: it establishes the most tightly confined Neptune-centered candidates first, and then expands the candidate set to include larger-amplitude families that preserve bounded motion while covering more of the Neptune-centered domain.

7.2.1 Compact Resonant Bounded Motion in Region II

The most clearly defined Neptune-centered bounded candidates are found within the compact invariant structure of Region II in the Poincaré maps. Across the full set of Jacobi constants considered in this chapter,

$$C_J = 3.018641378, \quad 3.0143741909, \quad 3.0107806367, \quad 2.999790146.$$

this region exhibits a persistent island family centered near

$$x \approx -0.25, \quad \dot{x} \approx 0.$$

surrounded by a set of nested invariant curves that form a coherent and well-isolated bounded neighborhood. Relative to the other Neptune-centered structures identified later in the chapter, this family is the most compact, the most regularly organized, and the most consistently recovered across the sampled energy range. For that reason, it provides a natural starting point for the candidate-identification process.

At this stage, candidate selection is performed at the level of *invariant structures* observed in the Poincaré section, rather than directly from initial conditions. Each candidate is therefore defined in two steps: first by identifying a specific invariant island or torus in the $(x, \dot{x})$ section, and then by associating that structure with the corresponding initial-condition seed x_0 that generates it. Because the mapping between x_0 and the Poincaré section is not one-to-one, the location of an island in $(x, \dot{x})$ does not directly correspond to the numerical value of x_0 . This distinction is important for the interpretation that follows, since the candidate is the invariant family itself, while the listed seed is

simply the initial condition used to reproduce it.

Figure 7.1 shows the Region II Poincaré section at $C_J = 3.018641378$, with the wider context alongside a zoomed view of the candidate island region.




Figure 7.1. Region II Poincaré section at $C_J = 3.018641378$. Upper: full-window context; lower: zoomed view of the candidate island used for seed identification.

The central portion of this structure corresponds to a resonance-supported island associated with the 8:9 family identified in Chapter 6. In the Poincaré sections with resonance overlays, this family appears as a compact, symmetric set of closed invariant curves whose geometry remains largely unchanged as the Jacobi constant varies. This persistence indicates that the structure represents a robust class of Neptune-centered bounded motion and, more importantly for the present chapter, that it can be treated as a repeatable candidate family rather than an isolated feature that appears only at one energy.

Two representative island candidates are selected from within this 8:9 resonance family:

- **Primary island (central torus):** The compact central island located near $(x, \dot{x}) \approx (-0.26, 0)$ corresponds to the core of the 8:9 resonance family. The trajectory that generates this invariant structure is obtained from the initial condition

$$x_0 = -0.250079240596.$$

- **Secondary island (outer torus):** A larger invariant curve within the same resonance family intersects the Poincaré section near $(x, \dot{x}) \approx (-0.285, 0)$. This island represents a higher-amplitude member of the same dynamical family and is generated by the initial condition

$$x_0 = 0.209920759404.$$

These seeds do not correspond directly to the $(x, \dot{x})$ coordinates of the island centers; rather, they are the initial conditions whose trajectories produce the observed invariant structures. At each Jacobi

level, the corresponding seeds are selected so that they generate the same resonance-supported island family, even as the precise geometry of the curves deforms slightly with energy. In this way, the candidate set remains tied to a persistent dynamical family rather than to a single fixed numerical value of x_0 .

From a dynamical standpoint, the defining feature of this region is the presence of smooth, closed invariant curves indicating quasi-periodic motion confined to a compact neighborhood around Neptune. The Region II islands remain well separated from surrounding mixed layers and do not exhibit significant fragmentation as the energy varies, which makes them among the most coherent candidates identified in the chapter. From a mission-design perspective, this compact resonance-supported structure is particularly attractive because it offers tightly confined motion in both configuration and velocity space, making it a natural candidate for repeated local coverage or observation-class behavior. The persistence of the structure across the selected Jacobi range is therefore the key reason it is carried forward: not because it is already proven stable in the Chapter 8 sense, but because it provides one of the clearest bounded Neptune-centered candidate families available for that later analysis.

7.2.2 Extended Prograde Neptune-Centered Motion in Region III

In addition to the compact prograde candidates identified in Region II, the Poincaré maps reveal a second class of Neptune-centered prograde bounded motion in Region III. This region differs from Region II in that it contains mixed-sense dynamics, with both prograde and retrograde structures occupying different parts of the phase space. For that reason, the prograde component of Region III is treated separately here from the retrograde structures that occupy the broader outer envelope. This distinction is also useful from the standpoint of candidate identification, since the prograde and retrograde branches play different roles in the Neptune-centered bounded set.

Across the selected Jacobi values,

$$C_J = 3.018641378, \quad 3.0143741909, \quad 3.0107806367, \quad 2.999790146.$$

the prograde contribution to Region III includes a set of large-amplitude invariant curves located on the left side of the Neptune-centered phase space. These structures form a broad, coherent envelope of closed trajectories that remain clearly bounded and dynamically organized, indicating that they represent an extended class of Neptune-centered motion rather than a transport or escape layer. In comparison with Region II, the relevant change is not a loss of organization, but an increase in amplitude and spatial reach.


Figure 7.2 shows the prograde Region III structure at $C_J = 3.018641378$, combining the full-window map with a zoomed resonance overlay for seed selection.

The dominant prograde structure in this region corresponds to a large-amplitude 1:1 resonance family located near the left boundary of the accessible phase space. This family appears in the Poincaré section as a broad set of smooth, symmetric invariant curves spanning approximately

$$x \approx 0.078 \text{ to } 0.083, \quad |v_x| \lesssim 2.$$


and forms the outer prograde envelope of the Neptune-centered bounded motion in Region III. The

Region III Prograde Bounded Candidates



a.) Regular Poincaré map

b.) Poincaré map with resonance overlay



c.) Zoomed view of island chain

Figure 7.2. Extended prograde Neptune-centered structures in Region III at $C_J = 3.018641378$. (a) Regular map, full window $x \in [0, 0.95]$, $\dot{x} \in [-2.5, 2.5]$, prograde in red. (b) Resonance overlay, same window. (c) Zoomed resonance overlay, $x \in [0.06, 0.22]$, for seed-selection context.

trajectory generating this invariant structure is obtained from the initial condition

$$x_0 = 0.799920759404.$$

This candidate is retained because it captures the largest clearly bounded prograde structure in the Neptune-centered domain without merging into the transition-layer behavior considered later in the chapter.

A second candidate is selected from a nearby bounded invariant curve located slightly interior to the primary 1:1 family. This structure appears in the Poincaré section as a thin, well-defined curve intersecting the section near

$$x \approx 0.0994 \text{ to } 0.1017.$$

and represents a smaller-amplitude Neptune-centered bounded trajectory distinct from the dominant outer envelope. The corresponding trajectory is generated by the initial condition

$$x_0 = 0.820079240596.$$

Retaining this second seed makes it possible to compare the outer and inner portions of the same extended bounded regime, rather than relying on a single representative orbit.

As in Region II, the invariant structures are identified directly from the Poincaré section, while the associated seeds x_0 are determined from the trajectories that generate these curves. The numerical values of x_0 do not coincide with the locations of the islands in $(x, \dot{x})$, but instead serve as the parameters that produce the observed invariant families.

From a dynamical standpoint, the defining feature of the Region III prograde structures is their large spatial extent combined with their regular, closed geometry. Unlike the compact Region II islands, these trajectories span a broader portion of phase space while remaining clearly bounded and well separated from nearby mixed layers. This indicates that the motion remains quasi-periodic and Neptune-centered, but operates in a higher-amplitude regime. From a mission-design perspective, these extended prograde trajectories provide a natural complement to the compact Region II candidates. Whereas Region II supplies the most localized bounded motion, the Region III prograde family offers broader coverage of the Neptune-centered environment while retaining enough dynamical coherence to remain credible as a candidate class. For that reason, these structures are carried forward not as replacements for the Region II family, but as a second Neptune-centered bounded option with a different operational scale.

7.2.3 Extended Retrograde Neptune-Centered Motion in Region III

In addition to the prograde bounded family identified in Region III, the Poincaré maps show that this region also contains a broad retrograde Neptune-centered component. This retrograde contribution is dynamically distinct from both the compact prograde islands of Region II and the more interior prograde family of Region III. Rather than appearing as a tightly clustered set of nested islands, the retrograde structures form the dominant large-amplitude envelope in the outer portion of Region III and occupy a substantially larger area of phase space. They therefore represent the most expansive Neptune-centered bounded trajectories identified in the present chapter.

Across the selected Jacobi values,

$$C_J = 3.018641378, \quad 3.0143741909, \quad 3.0107806367, \quad 2.999790146.$$

the retrograde Region III structures appear as extended closed invariant curves that wrap around the outer Neptune-centered domain. When the retrograde component is isolated, it becomes clear that much of the broad outer envelope seen in the full Region III map belongs to this family. These trajectories remain bounded and quasi-periodic, but their geometry is more elongated and their spatial extent is significantly larger than that of the corresponding prograde candidates. This makes them especially relevant as the outer limit of what still qualifies, at this stage, as a Neptune-centered bounded candidate.

Figure 7.3 shows the retrograde counterpart at the same energy, following the panel layout of Fig. 7.2, with panel (c) zoomed to the candidate island neighborhood.

As shown in Fig. 7.3 (c), the left portion of the retrograde map in that zoom resolves into a set of smooth, non-resonant invariant curves, indicating bounded Neptune-centered motion. Panel (b) shows that the strongest resonance labeling is concentrated elsewhere in the retrograde domain, while the zoomed candidate region remains dominated by geometrically coherent bounded curves. Within this retrograde envelope, two distinct invariant structures are selected as representative Neptune-centered candidates. Their role is not to sample every ring of the envelope exhaustively, but to capture its inner and more spatially extended portions with a minimal candidate set.

The first candidate corresponds to a bounded invariant curve located on the left side of the retrograde envelope, intersecting the Poincaré section approximately over the range

$$x \approx 0.0809 \text{ to } 0.0904.$$

This structure forms part of the inner portion of the retrograde family and represents a moderately sized, well-defined invariant curve. The trajectory generating this structure is obtained from the initial condition

$$x_0 = -0.080079240596.$$

The second candidate corresponds to a neighboring invariant structure located slightly further interior, intersecting the section over the range

$$x \approx 0.0914 \text{ to } 0.1104.$$

This structure represents a larger-amplitude member of the same retrograde family, with a broader spatial extent while remaining clearly bounded and dynamically organized. The corresponding trajectory is generated by the initial condition

$$x_0 = 0.809920759404.$$

As in the preceding subsections, the invariant structures are identified directly from the Poincaré section, while the associated seeds x_0 are determined from the trajectories that generate these curves. The values of x_0 therefore do not coincide with the visible locations of the structures in $(x, \dot{x})$, but instead define the initial conditions that produce the observed invariant families.

From a dynamical standpoint, the retrograde Region III candidates represent the outer Neptune-

Region III Retrograde Bounded Candidates




Figure 7.3. Extended retrograde Neptune-centered structures in Region III at $C_J = 3.018641378$. (a) Regular map, full window $x \in [0, 0.95]$, $\dot{x} \in [-2.5, 2.5]$, retrograde in blue. (b) Resonance overlay, same window. (c) Zoomed resonance overlay, $x \in [0.075, 0.115]$, covering candidate island ranges $x \approx 0.081\text{--}0.090$ and $0.091\text{--}0.110$.

centered bounded regime. Although these trajectories are significantly larger in amplitude than the prograde Region II and Region III families, they remain smooth, closed, and well separated from strongly mixed layers, indicating persistent quasi-periodic motion over the section. At the same time, their proximity to the outer phase-space boundary suggests that they occupy a less tightly confined regime than the compact inner families. From a mission-design perspective, these retrograde trajectories are therefore valuable for two reasons. First, they extend the Neptune-centered candidate set to include large-amplitude bounded motion with substantial spatial coverage. Second, because they lie near the outer boundary of the Neptune-centered domain, they provide a natural dynamical bridge between tightly bound observation-class motion and the transition-layer structures examined later in the chapter. For that reason, the retrograde Region III family is retained as part of the Neptune-centered bounded set while also serving as a conceptual precursor to the transport-focused candidates.

7.3 Touring-Region Candidates

While the bounded families identified in the preceding section provide well-defined Neptune-centered structures suitable for sustained observation, they do not fully exploit the broader dynamical pathways available within Region III. In particular, those families remain tied to comparatively narrow spatial envelopes, which limits their usefulness for trajectories intended to traverse multiple dynamical subregions. To address this limitation, the present section considers candidate trajectories drawn from the interior resonant architecture of Region III, where the phase space exhibits richer structure and increased connectivity across a wider range of Neptune-centered distances.

At the selected energy level,

$$C_J = 3.0107806367.$$

the Region III Poincaré sections reveal a complex mixture of bounded islands, elongated resonant chains, and partially chaotic transport regions. Compared to the higher-energy bounded cases emphasized earlier, the structures at this energy retain sufficient coherence to identify distinct families while simultaneously exhibiting greater radial extent and stronger interaction between neighboring portions of phase space. This balance makes the configuration particularly well suited for identifying trajectories that can support multi-phase motion, including repeated excursions between inner and outer Neptune-centered domains.

The focus of this section is therefore not on the most tightly bounded or geometrically regular families, but rather on those structures that demonstrate extended spatial coverage while maintaining recognizable dynamical organization. In practice, this corresponds to island chains and resonant arcs located within the interior of Region III, especially those that span a range of x -values and exhibit layered or nested behavior in the Poincaré maps. These features are indicative of trajectories capable of accessing multiple dynamical regimes, making them strong candidates for tour-like motion.

To better characterize these candidates, the analysis is separated into prograde and retrograde components. Although both senses of motion are present in the mixed Poincaré sections, the underlying structures differ significantly in their geometry and extent. The prograde families tend to retain stronger connections to the outer bounded envelopes, while the retrograde families more prominently occupy the interior resonant corridors and exhibit broader coverage of the Region III domain. Accordingly, candidate trajectories are identified and evaluated separately within each class.

The touring-region candidates are therefore distinguished according to sense of motion, reflecting the fundamentally different organization of the underlying phase space. The prograde structures remain more closely connected to the outer bounded envelopes and also support interior island chains and corridor-like motion at this energy, while the retrograde structures populate complementary interior subregions and can display strong layering in other parts of Region III. These differences are not merely geometric, but directly influence the accessibility, coverage, and dynamical character of the resulting trajectories. By examining each class independently, it becomes possible to isolate the distinct mechanisms that support extended motion within Region III and to identify the subset of families that most plausibly support a tour-like role.

7.3.1 Extended Prograde Island Structure in Region III

In addition to the large-amplitude prograde envelope identified previously in Region III, the Poincaré sections reveal a set of finer invariant structures embedded within the interior of the Neptune-centered phase space. These structures appear as distinct island chains and quasi-periodic bands located interior to the dominant outer 1:1 family, and represent a secondary class of bounded prograde motion with reduced spatial extent but increased structural richness. Their importance in the present chapter lies in the fact that they occupy the portion of Region III where bounded behavior begins to overlap with broader radial access.

Across the selected Jacobi values,

$$C_J = 3.018641378, \quad 3.0143741909, \quad 3.0107806367, \quad 2.999790146.$$

these interior structures persist as coherent features of the phase space, indicating that they correspond to invariant tori rather than transient or purely chaotic behavior. Unlike the outer envelope, which is dominated by a single large-amplitude resonance family, the interior region contains multiple candidate structures, including both resonance-supported island chains and smooth non-resonant bands. This diversity is precisely what makes the region useful for candidate identification, since it provides several distinct trajectory classes with different combinations of coherence and spatial reach.

As in the Region II analysis, candidate selection is performed at the level of invariant structures observed in the Poincaré section. Individual trajectories are then obtained by identifying the initial-condition seeds x_0 that generate these structures. Because the mapping between x_0 and the $(x, \dot{x})$ section is not one-to-one, the selected seeds correspond to trajectories that reproduce the observed islands rather than to the geometric centers of the structures themselves.

Figure 7.4 shows the interior prograde structure at $C_J = 3.0107806367$, using the same three-panel layout as Fig. 7.2, with panel (c) zoomed to the outer resonant island-chain region rather than the seed-selection neighborhood.

The first candidate structure is selected from an outer resonant island chain located approximately within

$$x \approx 0.30 \text{ to } 0.35.$$

This structure appears as a sequence of well-separated elliptical islands embedded within a larger quasi-periodic band. The trajectory generating this invariant family is obtained from the initial condition

$$x_0 = -0.900079240596.$$

Region III Interior Prograde Context




Figure 7.4. Interior prograde invariant structures in Region III at $C_J = 3.0107806367$. (a) Regular map, full window $x \in [0, 0.95]$, $\dot{x} \in [-2.5, 2.5]$, prograde in red. (b) Resonance overlay, same window. (c) Zoomed resonance overlay, $x \in [0.30, 0.45]$, $\dot{x} \in [-0.65, 0.65]$, outer island-chain neighborhood.




Figure 7.5. Zoomed view of the outer resonant island chain identified within the prograde Region III structure (see Fig. 7.4).

This island chain exhibits strong geometric coherence and clear separation from neighboring trajectories, indicating that it corresponds to a robust resonance-supported invariant structure. Its position near the outer portion of the interior region suggests that trajectories initialized here will explore a relatively large spatial extent while remaining dynamically organized. It is therefore retained as a prograde touring candidate because it combines recognizable resonance support with broader interior coverage than the more compact bounded families.

A second candidate is selected from a distinct interior island chain located approximately within

$$x \approx 0.52 \text{ to } 0.57.$$

The corresponding trajectory is generated by the initial condition

$$x_0 = -0.930079240596.$$

This deeper interior chain is shown in Fig. 7.6.

This structure differs from the outer chain in both geometry and location. The islands are more elongated and lie deeper within the central region of the phase space, closer to the transition between strongly ordered motion and weakly chaotic behavior. As a result, trajectories associated with this structure are expected to exhibit more complex phase-space evolution while still retaining the coherence characteristic of resonance-supported motion. Its inclusion broadens the prograde touring set beyond the most obvious outer chain and helps demonstrate how candidate behavior changes across the interior of Region III.

In addition to these resonance-supported structures, the Poincaré sections reveal a smooth, crescent-shaped invariant band spanning approximately

$$x \approx 0.50 \text{ to } 0.65.$$




Figure 7.6. Zoomed view of the interior resonant island chain located deeper within the Region III phase space.

The trajectory associated with this structure is obtained from the initial condition

$$x_0 = -0.940079240596.$$

The geometry of this non-resonant band is shown in Fig. 7.7.




Figure 7.7. Zoomed view of the non-resonant crescent-shaped invariant band within Region III.

Unlike the previously selected island chains, this structure does not correspond to a clearly

identifiable low-order resonance. However, it remains a well-defined and bounded feature of the phase space, indicating quasi-periodic motion on invariant tori. The absence of a dominant resonance does not preclude its usefulness; rather, it highlights that coherent Neptune-centered motion can arise both within and outside of strict commensurability conditions. This makes it a useful comparison case within the touring set.

From a dynamical standpoint, the inclusion of this non-resonant structure provides a useful contrast to the resonance-supported candidates. Resonant islands correspond to trajectories whose frequencies satisfy a commensurate ratio, producing discrete and well-separated structures in the Poincaré section. In contrast, non-resonant invariant bands arise from quasi-periodic motion with incommensurate frequencies, leading to smoother and more continuous phase-space coverage. This distinction is relevant for trajectory design, as non-resonant motion may provide greater flexibility in certain mission scenarios.

For this reason, the non-resonant structure is retained as an additional candidate for further investigation. Its inclusion allows for a direct comparison between resonance-supported and non-resonant Neptune-centered trajectories, and helps establish that the selection of touring candidates in Region III is not limited to resonance families alone. In the context of Chapter 7, that broader comparison is useful because it prevents the candidate set from being restricted too narrowly to only the most visually obvious resonant chains.

7.3.2 Retrograde Corridor Touring Candidates

The interior of Region III at $C_J = 3.0107806367$ also supports a layered corridor of retrograde structures composed of nested invariant bands and embedded resonant islands. Unlike the broad outer retrograde envelope identified earlier in the chapter, these interior structures occupy a corridor-like region that spans a wide range of Neptune-centered distances while maintaining coherent geometric organization in the Poincaré section. This configuration makes the retrograde interior particularly well suited for identifying candidate trajectories capable of extended motion through multiple dynamical subregions.

At

$$C_J = 3.0107806367$$

the retrograde Poincaré section reveals several distinct island chains embedded within a broader non-resonant background. Rather than selecting candidates from a single resonance family, the present analysis focuses on a set of representative structures that span the interior retrograde region from outer to inner domains. This approach ensures that the selected candidates capture the diversity of dynamical behavior present in the corridor, including both low-order and higher-order resonant motion, as well as structures embedded within the surrounding quasi-periodic band. In other words, the goal here is not to privilege one resonance label, but to identify a small set of retrograde candidates that collectively represent the dynamical layering of the corridor.

The four retrograde touring candidates considered in this subsection are identified in the zoomed Poincaré section shown in Fig. 7.8. These candidates were selected based on three criteria: (i) clear geometric coherence in the Poincaré map, (ii) placement within the interior retrograde corridor rather than the outer bounded envelope, and (iii) representation of distinct resonance families or dynamical roles within the region.

The first candidate corresponds to the outer interior retrograde island located near $x \approx 0.59$,

Region III Retrograde Corridor Candidates




Figure 7.8. Zoomed retrograde Poincaré section at $C_J = 3.0107806367$ with the same resonance-overlay coloring as the ORBIT Poincaré analysis tool (labels from the run `full_run.csv`; colors from the shared resonance palette, with unclassified points in light gray). Representative structures discussed in the text include (1) outer interior 1 : 1 island near $x \approx 0.59$, (2) intermediate 1 : 11 crescent near $x \approx 0.72$, (3) embedded 1 : 6 island near $x \approx 0.78$, and (4) deep interior 4 : 15 island near $x \approx 0.87$.

with initial condition

$$x_0 = 0.589920759404,$$

and associated with the 1 : 1 resonance. This structure represents the outer boundary of the interior retrograde corridor and provides a reference configuration for comparison with more deeply embedded candidates. Its relatively large size and regular geometry indicate stable, repeatable motion while still allowing excursions across a broader radial range than the outer bounded families.

The second candidate is the elongated crescent-shaped island located near $x \approx 0.70$ – 0.75 , with

$$x_0 = -0.720079240596,$$

corresponding to a higher-order 1 : 11 resonance. Compared to the 1 : 1 island, this structure exhibits a more stretched geometry aligned with the interior corridor, suggesting increased radial variation and a stronger connection to the surrounding transport region. This makes it a representative intermediate structure bridging the outer and inner portions of the retrograde domain.

The third candidate, located near $x \approx 0.78$ with

$$x_0 = 0.779920759404,$$

belongs to the 1 : 6 resonance and is embedded within a broader non-resonant invariant band. This configuration is particularly important, as it demonstrates that organized resonant motion persists within the surrounding quasi-periodic region. As a result, this candidate provides insight into how coherent trajectories may be sustained while traversing the interior corridor.

The fourth candidate is a compact island near $x \approx 0.85$, defined by

$$x_0 = 0.870079240596,$$

and associated with the 4:15 resonance. This structure lies deeper within the interior of Region III and represents a localized pocket of organized motion. Although smaller in spatial extent than the other candidates, it provides a useful contrast by illustrating how higher-order resonant islands persist in the inner portion of the retrograde corridor.

Together, these four candidates span the interior retrograde region from outer to inner domains and represent a range of resonance classes and geometric structures. This selection provides a systematic basis for evaluating retrograde touring trajectories, capturing both the hierarchical organization of the phase space and the diversity of dynamical behaviors available within Region III. It also gives the later stability analysis a deliberately varied retrograde candidate set rather than a cluster of nearly redundant seeds.

7.4 Triton-Centered Bounded Candidates

In addition to the Neptune-centered and interior Region III structures identified in the previous sections, the Poincaré maps also reveal a distinct set of invariant structures located near Triton. These structures form Region IV and are characterized by compact, highly organized island families centered near the Triton primary. Unlike the broader and more complex phase-space architecture observed in Region III, the Region IV structures exhibit a relatively simple and hierarchical organization consisting of nested invariant curves. For the purposes of Chapter 7, this simplicity is advantageous, since it allows a concise Triton-centered candidate set to be defined without the same degree of ambiguity present in the interior Neptune-centered regions.

At the energy level

$$C_J = 3.018641378.$$

the Triton-centered region is clearly separated from the interior Neptune-centered domain. The corresponding Poincaré section is shown in Fig. 7.9, where both prograde and retrograde structures are visible. This separation is one of the main reasons the region is attractive as a candidate class: it represents localized motion about Triton that is dynamically distinct from the broader Neptune-centered environment.

To better distinguish the individual families, Fig. 7.10 shows the same region with resonance-based coloring applied, highlighting the nested structure of the Triton-centered islands.

7.4.1 Prograde Triton-Centered Candidates

The prograde structures in Region IV consist of two primary invariant islands located near $x \approx 1.0$. These islands form a compact pair of concentric families and represent the only clearly identifiable prograde structures within this region. Their small number and regular geometry make the prograde Triton-centered set comparatively straightforward to define.

The inner prograde island is defined by the initial condition

$$x_0 = 0.989920759404,$$




Figure 7.9. Poincaré section at $C_J = 3.018641378$ showing the Triton-centered Region IV structures. The region consists of compact invariant islands located near $x \approx 1.0$, separated from the interior Neptune-centered domain.




Figure 7.10. Resonance-colored Poincaré section at $C_J = 3.018641378$ for Region IV. The nested structure of the Triton-centered invariant curves is evident, with distinct families corresponding to prograde and retrograde motion.

while the outer prograde island corresponds to

$$x_0 = 1.009920759404.$$

These two islands provide a representative pair of prograde Triton-centered candidates, capturing both the inner and outer extent of the prograde family. Their compact geometry and close proximity to the Triton primary distinguish them from the larger-scale structures observed in the Neptune-centered regions. In the candidate-selection framework of this chapter, they therefore serve as the most direct analog of localized observation-class motion near Triton.

7.4.2 Retrograde Triton-Centered Candidates

The retrograde structures in Region IV exhibit a nested, multi-ring configuration centered near Triton. Unlike the prograde case, the retrograde region contains multiple concentric invariant curves corresponding to distinct but closely related families. This richer layering makes the retrograde set somewhat broader, but it still remains far more compact and orderly than the interior Region III structures.

From these structures, three representative retrograde candidates are selected to span the region from inner to outer extent:

$$\begin{aligned} x_0 &= 1.009920759404 && \text{(center ring),} \\ x_0 &= 0.989920759404 && \text{(intermediate ring),} \\ x_0 &= 0.979920759404 && \text{(outer ring).} \end{aligned}$$

These selections capture the hierarchical structure of the retrograde families while avoiding redundancy among closely spaced rings. Although additional outer rings are present, their geometric similarity to the selected candidates suggests that they belong to the same nested family and do not introduce qualitatively distinct structures at this stage of the analysis. The purpose here is therefore not to exhaust every visible ring, but to retain a minimal set that spans the region in a representative way.

It is noted that similar values of x_0 may correspond to both prograde and retrograde structures, depending on the sense of motion associated with the initial condition. The classification of each candidate is therefore determined not only by its location in the Poincaré section, but also by the direction of motion at the surface of section. This distinction is especially important in Region IV, where the geometric compactness of the structures can otherwise make the prograde and retrograde families appear deceptively similar at first glance.

Together, the prograde and retrograde candidates identified in this section provide a concise representation of the Triton-centered phase-space structures in Region IV. As in the previous sections, these candidates are selected on the basis of their geometric organization and phase-space location, with further dynamical characterization deferred to subsequent analysis. The resulting set gives the later chapters a compact but meaningful Triton-centered comparison group alongside the larger Neptune-centered and touring-region candidates.

7.5 Bottleneck-Crossing Transport Candidates

In contrast to the bounded and touring candidates identified in the previous sections, transfer candidates are associated with trajectories that pass between dynamically distinct regions of the system through the narrow openings in the zero-velocity surfaces. These openings, or bottlenecks, first appear when the Jacobi constant satisfies

$$C_{L1} < C_J < C_{L2}. \quad (7.1)$$

at which point motion through the L_1 region becomes accessible while the L_2 gateway remains closed. This regime therefore provides a natural setting for identifying candidate trajectories that may traverse between the Neptune-centered and Triton-centered domains without escaping the system. The selection logic in this section is therefore intentionally different from that used for the bounded families: here the goal is not maximal confinement, but credible inter-region access.

For the present analysis, the energy level

$$C_J = 3.0143741909.$$

is selected, which lies within the interval $C_{L1} < C_J < C_{L2}$. At this energy, the L_1 bottleneck is open, and the associated region of the Poincaré section near

$$x_{L1} \approx 0.95915257.$$

corresponds to the phase-space location where passage through the neck first becomes possible.

Rather than identifying candidates from well-defined resonant island chains, transfer candidates are selected by examining the local neighborhood of the L_1 bottleneck in the frequency-ratio map. The frequency ratio (2.45)

$$\frac{\omega_2}{\omega_1}$$

provides a classification of the underlying dynamical behavior, with rational values corresponding to identifiable resonance families. In this subsection, attention is focused on trajectories in the L_1 neighborhood that do not lie on clearly defined resonance branches. Such trajectories are not associated with strongly organized invariant structures and therefore represent natural candidates for transfer-type motion. In other words, the transfer set is identified partly by what it is *not*: it is not drawn from the most coherent bounded families, but from the neighborhood where confinement begins to weaken and passage becomes possible.

The frequency-ratio distribution for this energy level is shown in Fig. 7.11.

Figure 7.12 shows the same data with resonance classifications overlaid, distinguishing the organized families from the non-resonant transfer candidates.

Based on this selection procedure, three candidate initial conditions are identified in the vicinity of the L_1 bottleneck:




Figure 7.11. Frequency ratio map at $C_J = 3.0143741909$. Candidate initial conditions are identified near the L_1 bottleneck region ($x_0 \approx 0.94$ – 0.98) while avoiding clearly defined resonance groupings.




Figure 7.12. Frequency resonance overlay at $C_J = 3.0143741909$. Resonance-associated trajectories are identified and excluded from the transfer candidate set, leaving non-resonant points near the L_1 region as primary candidates.

$$\begin{aligned}
x_0 = 0.94 \quad (\text{retrograde}), \quad & \frac{\omega_2}{\omega_1} \approx 0.3, \\
x_0 = 0.94 \quad (\text{prograde}), \quad & \frac{\omega_2}{\omega_1} \approx 0.4, \\
x_0 = 0.98 \quad (\text{retrograde}), \quad & \frac{\omega_2}{\omega_1} \approx 0.025.
\end{aligned}$$

These candidates are located near the L_1 bottleneck and are not associated with clearly identifiable resonance families in the frequency-ratio map. As such, they represent a set of initial conditions that are well positioned within the phase-space region where inter-region motion becomes accessible, while remaining distinct from the bounded and resonance-supported structures considered in previous sections. Their inclusion ensures that the Chapter 7 candidate set is not limited to observation-class or touring-class motion, but also contains trajectories chosen specifically for their transport relevance. The dynamical behavior and practical usefulness of these transfer candidates are examined more rigorously in the chapters that follow.

7.6 Synthesis of Candidate Structures

The preceding sections have established a structured set of candidate trajectories across the primary dynamical regions of the Neptune–Triton system. Beginning with the Neptune-centered bounded families, extending through the interior Region III structures, and concluding with the Triton-centered and bottleneck-accessible configurations, the analysis has identified representative invariant sets that capture the dominant mission-relevant organization of the phase space at the selected energy levels.

Each class of candidates was selected using criteria appropriate to its dynamical context. In the Neptune-centered regions, candidate trajectories were drawn from clearly defined bounded island families that remain coherent across the sampled energies. Within Region III, the selection emphasized interior structures spanning a range of resonance classes and geometric configurations, so that both resonant and non-resonant touring behavior would be represented. Transfer candidates were identified through the L_1 bottleneck using frequency-based filtering to isolate trajectories not strongly associated with dominant resonance families. Finally, the Triton-centered region provided a compact set of nested invariant curves representing localized motion near the secondary body.

Taken together, these selections form a coherent set of phase-space representatives that reflect the hierarchical structure of the system while also preserving the mission-oriented logic introduced at the beginning of the chapter. Importantly, the identification process has remained strictly within the framework of the Poincaré map and frequency analysis, ensuring that all candidates are defined in terms of their geometric and spectral properties rather than by prematurely judging their fully realized trajectories. That separation is important for the thesis as a whole, because it keeps Chapter 7 focused on candidate identification rather than allowing it to collapse into the stability analysis reserved for Chapter 8.

While this phase-space characterization provides a necessary foundation, it does not by itself determine the dynamical behavior or practical utility of the selected candidates. In particular, the stability of the associated motion and the response of each candidate to perturbations cannot be inferred solely from the structure of the Poincaré section. A more rigorous assessment requires examination of the underlying periodic or quasi-periodic dynamics and their associated linearized behavior.

Accordingly, the following chapter focuses on the dynamical validation of the candidate set through stability analysis. This includes the computation of periodic orbit representations where applicable, the construction of the state transition matrix over one period, and the evaluation of the corresponding monodromy matrix. The resulting Floquet multipliers provide a quantitative classification of each candidate, distinguishing between stable and unstable configurations and establishing the dynamical properties necessary for further application.

8. Stability of Candidate Trajectories

Chapter 7 narrowed the broad Neptune–Triton phase space to a smaller set of candidate structures whose geometry, resonance support, and mission relevance justified continued study. That selection process, however, was intentionally based on the organization of the Poincaré sections and the associated frequency information rather than on a direct assessment of the long-term dynamical behavior of the candidates themselves. As a result, the Chapter 7 candidate set should be understood as a physically motivated shortlist, not as a final statement about which trajectories are dynamically robust enough to support later mission interpretation.

The purpose of the present chapter is therefore to evaluate the stability properties of the candidate trajectories identified in Chapter 7. This requires moving beyond geometric and spectral classification to examine the behavior of the associated periodic or quasi-periodic motion under linearized perturbations. In practical terms, the analysis developed here is used to determine which candidate families remain sufficiently organized under perturbation to warrant continued consideration, which candidates appear strongly unstable, and which occupy an intermediate regime that may still be useful despite limited instability. The role of Chapter 8 is thus to refine the candidate set on local dynamical grounds; the separate question of how those surviving candidates function in the full mission architecture is taken up in Chapter 9.

The chapter proceeds in four parts. It first defines the stability framework used for the candidate set, including the periodic-orbit representation adopted where applicable and the associated state transition and monodromy matrices. It then applies that framework to the Neptune-centered, touring-region, and Triton-centered candidates identified in Chapter 7. A separate section is used for the bottleneck-crossing transport candidates, since their value is tied less to long-term confinement and more to their ability to provide access between dynamical regions. The chapter concludes by summarizing which candidates remain most viable from the standpoint of stability and by identifying the subset carried forward to the mission-level interpretation of Chapter 9.

8.1 Stability Framework

The stability methodology used in this chapter was developed in detail in Chapter 3 and implemented in the ORBIT framework described in Chapter 4. The purpose of the present section is therefore not to repeat that derivation in full, but to briefly restate the key steps needed to interpret the candidate-level results that follow. In the present chapter, each candidate selected from Chapter 7 is treated as a trajectory family of interest whose local dynamical behavior can be assessed through a periodic-orbit representation when such a representation is well defined. This allows the geometric and spectral candidates identified earlier to be tested in a more rigorous way using the linearized dynamics along the orbit.

The starting point of the analysis is the same J_2 -perturbed Neptune–Triton model used throughout the thesis. For a candidate trajectory that admits a periodic approximation, the equations of motion are augmented by the associated variational equations so that the nonlinear trajectory and the linearized perturbation dynamics are propagated simultaneously. As summarized in Chapter 3, the linearized behavior is governed by the state transition matrix $\Phi(t)$, which maps an initial perturbation in the state to its value at a later time. In practice, the propagated system therefore contains both the four state variables of the planar problem and the sixteen elements of the 4×4

state transition matrix.

For each candidate considered in this chapter, an initial estimate of the orbit period is taken from the recurrence structure identified in the Poincaré map and the associated candidate-selection workflow established in Chapter 7. Where useful, the dominant frequency information from Chapter 6 provides a consistency check on the expected timescale of the motion, but the working period estimate remains tied to the recurrence structure of the section. That initial estimate is then refined through the same differential-correction procedure described in Chapter 3 so that the orbit closes as accurately as possible under the adopted periodic representation. The corrected period is important because the stability classification is based on the orbit after correction, not on the unrefined seed taken directly from the map.

Once the corrected trajectory and the augmented variational system have been propagated over one period T , the final state transition matrix is reshaped into the monodromy matrix (2.58),

$$\mathbf{M} = \Phi(T).$$

The eigenvalues of $\mathbf{M}$ are the Floquet multipliers and provide the primary linear stability measure used in this chapter. As discussed previously, one reciprocal pair is associated with the nontrivial stability behavior of the orbit, while the near-unity pair is related to the periodic nature of the solution itself. The location of the nontrivial multipliers relative to the unit circle is used to determine whether the candidate behaves as a linearly stable, marginal, or unstable configuration in the neighborhood of the corrected orbit.

In addition to the full multiplier set, this chapter also reports the scalar stability index introduced in Chapter 3. That quantity is derived from the nontrivial reciprocal multiplier pair after removing the near-unity pair, and it provides a compact way to compare candidates across different dynamical classes. The use of both the full Floquet multipliers and the scalar index is important here: the multiplier set retains the most complete information about the local linearized behavior, while the scalar index makes it easier to summarize candidate quality across the Neptune-centered, touring-region, Triton-centered, and transport-oriented groups.

One final point is worth emphasizing before turning to the candidate-specific results. The stability analysis in this chapter is not intended to replace the geometric and spectral interpretation developed in the preceding chapters, but to refine it. A candidate that appears coherent in the Poincaré section or well organized in the resonance analysis is not automatically retained unless its local dynamical behavior is also defensible under perturbation. Conversely, a candidate that is not strongly stable in the strict linear sense may still remain relevant if its degree of instability is limited and its mission role is tied to access or transfer rather than long-term confinement.

For this reason, the tables below distinguish between successful elliptic-stable cases and candidates labeled as ‘Invalid orbit.’ In the present chapter, ‘Invalid orbit’ does not mean that the original seed is meaningless or that no nearby motion exists in a geometric sense. Rather, it means that the periodic-orbit correction and validation workflow did not produce a defensible corrected periodic representative for that seed under the adopted stability framework. In practice, this outcome occurs when the corrector fails to converge satisfactorily, when the corrected state does not meet the closure or consistency tolerances required for a usable periodic orbit, or when the resulting orbit is so poorly behaved under the monodromy-based test that a stability classification is not considered reliable. The sections that follow therefore use the same stability framework throughout, but interpret the results in light of the intended function of each candidate class.

8.2 Neptune-Centered Candidate Stability

The Neptune-centered candidates form the most natural starting point for the stability results because they were also the most geometrically coherent candidates identified in Chapter 7. These include the compact 8:9 resonance family of Region II together with the larger-amplitude bounded families selected from Region III. Taken together, they test whether the bounded structures that appeared most promising in the Poincaré and resonance analyses also admit well-behaved periodic-orbit representations under the correction and stability workflow described above.

Across this candidate class, the results are broadly favorable. Most of the Neptune-centered seeds converge under differential correction, and the converged cases are classified as elliptic stable. This is the strongest positive result in the chapter, since it shows that the compact Region II family and much of the larger Region III bounded structure are not merely visually regular in the Poincaré section, but also dynamically well behaved under the linearized periodic-orbit analysis. At the same time, the Neptune-centered set is not uniformly successful: a small number of candidates either fail to converge or exhibit invalid corrected solutions, indicating that even within the bounded Neptune-centered domain the stability properties remain family-dependent.

For consistency with the finalized candidate definitions of Chapter 7, the stability results in this chapter are reported using the Chapter 7 candidate labels and seed values. Where the exported stability file contained labeling inconsistencies, the tables below follow the finalized Chapter 7 notation so that the candidate definitions remain uniform across the thesis.

Table 8.1. Stability results for the Neptune-centered bounded candidates at $C_J = 3.018641378$. Reported candidate labels follow the finalized Chapter 7 notation.

Candidate	x_0	$ \lambda _{\max}$	Stability index	Result
Region II central torus	-0.250079	1.000×10^0	1.535	Elliptic stable
Region II outer torus	0.209921	1.000×10^0	1.110	Elliptic stable
Region III prograde outer envelope	0.799921	1.000×10^0	0.925	Elliptic stable
Region III prograde inner curve	0.820079	1.000×10^0	-1.690	Elliptic stable
Region III retrograde candidate A	-0.080079	1.469×10^1	-14.754	Invalid orbit
Region III retrograde candidate B	0.089921	1.000×10^0	0.452	Elliptic stable

The Region II results are especially encouraging. Both members of the 8:9 family converge cleanly and remain extremely close to neutrality in the largest multiplier, with values only slightly displaced from unity. In practical terms, this indicates that the compact resonance-supported family identified earlier retains the strongest case for continued consideration. Among the Neptune-centered candidates, Region II therefore provides the clearest example of agreement between geometric organization, spectral coherence, and linearized stability.

The Region III bounded families preserve much of this positive picture, but with more variation. Both prograde Region III candidates converge and are classified as elliptic stable, indicating that

the larger-amplitude Neptune-centered bounded regime remains dynamically tractable under the periodic-orbit framework. This result is important because it shows that increasing spatial extent does not automatically destroy local stability. The extended Region III prograde families therefore remain credible mission candidates, particularly where broader coverage of the Neptune-centered environment is desirable.

The retrograde Region III pair is more mixed. One candidate converges to an elliptic-stable solution with a multiplier still close to unity, while the other fails to produce a valid corrected orbit and exhibits a very large multiplier magnitude. This split result is consistent with the geometric interpretation of Chapter 7, where the retrograde outer envelope appeared to occupy the least tightly confined portion of the Neptune-centered bounded domain. In Chapter 8 terms, the retrograde Region III family should therefore be regarded as only partially retained: it contains at least one strong candidate, but it is less uniformly robust than the compact Region II family or the better-behaved Region III prograde branch.

8.3 Touring-Region Candidate Stability

The touring-region candidates were selected in Chapter 7 precisely because they offered more spatial reach and richer phase-space access than the tightly bounded Neptune-centered families. For that reason, a less favorable stability outcome is not surprising. These candidates were never expected to outperform the compact Region II islands in terms of linearized robustness; instead, the point of this section is to determine how much dynamical quality is retained once the analysis shifts from tightly confined motion to extended, region-spanning structures in the interior of Region III.

The results confirm that the touring set is substantially more mixed than the Neptune-centered bounded set. Several of the prograde and retrograde touring candidates fail to converge to valid corrected periodic representatives, and the unsuccessful cases are generally associated with much larger multiplier magnitudes and correspondingly larger scalar stability indices. At the same time, the section is not purely negative. A smaller subset of the touring candidates remains elliptic stable after correction, which means that the touring-region concept is not lost entirely; rather, it becomes more selective and more sensitive to the exact structure from which the candidate is drawn. This is precisely the kind of outcome that justifies separating Chapter 8 from Chapter 9: stability narrows the touring set, but does not by itself determine how the surviving trajectories should ultimately be used.

Table 8.2. Stability results for the touring-region candidates selected from the interior of Region III at $C_J = 3.0107806367$.

Candidate	x_0	$ \lambda _{\max}$	Stability dex	in-	Result
Prograde outer island chain	-0.900079	1.000×10^0	1.759		Elliptic stable
Prograde interior island chain	-0.930079	1.000×10^0	1.030		Invalid orbit
Prograde non-resonant band	-0.940079	1.000×10^0	1.542		Elliptic stable
Retrograde outer interior island	0.589921	4.212×10^0	4.449		Invalid orbit
Retrograde crescent candidate	-0.720079	7.474×10^0	7.608		Invalid orbit
Retrograde embedded island	0.779921	1.551×10^1	15.572		Invalid orbit
Retrograde deep interior island	0.870079	1.000×10^0	-1.386		Elliptic stable

The prograde touring set illustrates the basic pattern well. The outer island-chain candidate and the non-resonant band both converge to elliptic-stable solutions, indicating that extended motion in the interior of Region III can still support locally well-behaved periodic representations. However, the interior resonant chain fails to converge to a valid corrected orbit. This suggests that the prograde touring regime does not become uniformly unstable, but instead contains alternating pockets of dynamical tractability and regions where the periodic approximation is more difficult to sustain.

The retrograde touring results are more severe. Of the four retrograde touring candidates, only the deep interior island converges to an elliptic-stable solution, while the remaining three produce invalid corrected orbits with substantially larger multiplier magnitudes. This is a meaningful result rather than a computational inconvenience. The retrograde touring corridor was selected because it offered layered access through the interior of Region III, but that same richness appears to come with a higher degree of sensitivity under correction. In practical terms, the retrograde touring region still contains at least one viable candidate, but the class as a whole is much less robust than the Neptune-centered bounded families.

Viewed together, the touring-region results support a clear interpretation. The interior Region III structures do provide access to broader spatial coverage and more varied dynamical behavior, but that broader reach comes at the cost of reduced stability yield. Chapter 8 therefore refines the Chapter 7 touring candidate set in an important way: the touring region remains relevant, but only a subset of its candidates survives as strong stability-supported options.

8.4 Triton-Centered Candidate Stability

The Triton-centered candidates occupy a different dynamical role from the Neptune-centered and touring-region structures. Their geometry in Chapter 7 was notably more compact and more localized near the secondary body, which made them attractive as possible local Triton-phase candidates. The central question in the present section is whether that geometric compactness is matched by favorable local stability properties, or whether the Triton-centered families remain more

fragile than their appearance initially suggests.

The stability results show a mixed picture. The candidate set contains successful elliptic-stable cases in both senses of motion, but it also contains several invalid corrected orbits. This outcome places the Triton-centered candidates between the clearly favorable Neptune-centered bounded set and the more difficult touring-region set. Region IV therefore remains mission-relevant, but the stability analysis indicates that only part of the originally selected Triton-centered family should be carried forward with confidence.

Table 8.3. Stability results for the Triton-centered candidates selected from Region IV at $C_J = 3.018641378$.

Candidate	x_0	$ \lambda _{\max}$	Stability dex	in-	Result
Prograde inner island	0.989921	1.237×10^0	1.888		Invalid orbit
Prograde outer island	1.009921	1.000×10^0	1.795		Elliptic stable
Retrograde center ring	1.009921	1.000×10^0	1.729		Elliptic stable
Retrograde intermediate ring	0.989921	1.000×10^0	1.842		Invalid orbit
Retrograde outer ring	0.979921	1.000×10^0	1.615		Invalid orbit

Within the prograde Triton-centered pair, only the outer island converges to an elliptic-stable solution. The inner island fails to produce a valid corrected orbit, which indicates that even within the compact Triton-centered set the local stability properties can vary noticeably over a relatively small geometric range. This is a useful caution against reading the nested visual organization of Region IV as a guarantee of uniformly favorable dynamical behavior.

The retrograde Triton-centered set behaves similarly. The center ring converges and remains elliptic stable, while the intermediate and outer rings do not yield valid corrected solutions. That result suggests that the Triton-centered retrograde family is not wholly unsuitable, but also not uniformly robust across its visible nested structure. The stability-supported Triton-centered set is therefore smaller than the geometric candidate set originally identified in Chapter 7.

For transparency, the failed Triton-centered entries should be understood in the same sense as the framework definition given above. These cases were not discarded because they belonged to the wrong Jacobi level or because the underlying geometric structures were obviously spurious. Instead, the issue is that the selected seeds did not validate as usable corrected periodic representatives under the present workflow. In other words, the corresponding Region IV structures remain visible and dynamically interesting in the maps, but this chapter cannot claim a defensible local periodic-stability result for those specific seeds. That distinction is important because it preserves the geometric meaning of the candidate while still being honest about the limits of the periodic-orbit validation.

From a thesis-level standpoint, the Triton-centered results are still valuable. They show that localized motion near Triton is not eliminated by the stability analysis, but it does require additional filtering. The stable prograde outer island and the stable retrograde center ring remain plausible candidates for later mission interpretation, whereas the failed cases are better viewed as geometrically interesting structures that do not survive the stricter periodic-orbit test in a useful way.

8.5 Transport-Candidate Stability and Sensitivity

The transport candidates require a different reading from the other candidate classes. These trajectories were never selected because they appeared to be the most tightly organized or the most likely to yield long-lived bounded motion. Instead, they were chosen because they lie in the neighborhood of the L_1 bottleneck, where motion between the Neptune-centered and Triton-centered regions first becomes accessible. As a result, the relevant question here is not simply whether the candidates are linearly stable in the same sense as a compact bounded family, but whether the stability analysis supports their interpretation as sensitive transfer-type trajectories.

For consistency with the finalized Chapter 7 candidate definition, the three bottleneck cases are reported here at

$$C_J = 3.0143741909,$$

which is the transfer energy used in the candidate-selection discussion and in the corresponding Chapter 9 mission-architecture interpretation.

Table 8.4. Stability results for the bottleneck-crossing transport candidates at $C_J = 3.0143741909$.

Candidate	x_0	$ \lambda _{\max}$	Stability in- dex	Result
Retrograde bottleneck candidate	0.939921	1.000×10^0	-0.260	Elliptic stable
Prograde bottleneck candidate	0.939921	7.987×10^3	7987.407	Invalid orbit
Outer prograde bottleneck candidate	0.979921	1.008×10^3	1008.378	Invalid orbit

The results are strongly asymmetric. The retrograde bottleneck candidate converges and is classified as elliptic stable, with a largest multiplier essentially at unity. By contrast, both prograde bottleneck candidates fail to produce valid corrected periodic solutions and are associated with extremely large multiplier magnitudes. This is one of the sharpest contrasts in the entire table and indicates that the bottleneck neighborhood is highly sensitive to both the chosen initial condition and the sense of motion.

This outcome is dynamically plausible. The transport candidates were deliberately selected from the portion of phase space where confinement weakens and inter-region access becomes possible. In such a regime, one should not expect the same rate of success achieved for the compact Region II families. The two invalid prograde bottleneck cases therefore do not simply indicate failed candidates in a narrow computational sense; they also reinforce the interpretation that these trajectories occupy a highly sensitive transfer regime rather than a smoothly bounded one.

At the same time, the successful retrograde bottleneck case is important because it shows that the transport class is not uniformly unusable. At least one candidate in the L_1 neighborhood admits a corrected periodic representative with favorable local linearized behavior. That does not automatically make it the best mission-transfer solution, since Chapter 9 must still test whether the full realized trajectory actually performs the required transfer role, but it does justify carrying the transport class forward to the mission-interpretation stage with selective emphasis on the strongest local stability result.

8.6 Stability Synthesis

The stability analysis refines the candidate set of Chapter 7 in a decisive way. The strongest overall performance is found in the Neptune-centered bounded families, especially the compact Region II 8:9 resonance and the better-behaved Region III bounded branches. These candidates not only appeared coherent in the phase-space and frequency analyses, but also remain well behaved under the periodic-orbit and monodromy-based stability framework. As a result, the Neptune-centered bounded class emerges from Chapter 8 as the most consistently defensible candidate group.

The touring-region and Triton-centered candidates produce a more selective outcome. In both classes, some structures survive the periodic-orbit correction and remain elliptic stable, while others fail to converge to valid corrected solutions. This split is informative. It shows that neither class should be discarded wholesale, but it also confirms that the broad geometric candidate sets identified in Chapter 7 require further filtering before they can support mission-level interpretation. The touring-region set, in particular, remains valuable primarily through a smaller subset of successful cases rather than through its full original extent.

The transport results sharpen that conclusion further. The bottleneck neighborhood remains important because it is the only portion of the candidate set selected specifically for inter-region access, but its local linearized behavior is clearly more sensitive than that of the bounded Neptune-centered families. The retrograde bottleneck case provides the strongest local stability result within this class, while the prograde bottleneck seeds do not converge successfully under the periodic-orbit framework. This means that the transfer class remains relevant, but only in a much narrower and more delicate form than the broader bounded classes. As Chapter 9 makes clear, even this strongest local result must still be checked against the full realized transfer geometry before it can be accepted architecturally.

Taken together, the Chapter 8 results support a three-tier interpretation of the candidate set. The first tier consists of the strongest locally stability-supported candidates: the Region II bounded family, the successful Region III bounded representatives, the most successful Triton-centered cases, and the best-performing bottleneck result under the periodic-orbit framework. The second tier consists of marginal but still informative candidates, especially those touring-region cases that preserve some degree of dynamical structure but do not perform as consistently as the bounded families. The third tier consists of candidates that fail to produce valid corrected periodic representatives and are therefore not retained as primary stability-supported options. This tiered result provides the final dynamical screening step needed before the thesis turns to the broader mission-level interpretation of the surviving trajectory classes in Chapter 9.

9. Neptune–Triton Tour Architecture

The preceding chapters established the dynamical foundation required to interpret the Neptune–Triton system in mission terms. Chapter 5 identified the large-scale phase-space structure, Chapter 6 clarified the resonance organization of that structure, Chapter 7 narrowed the search to a focused candidate set, and Chapter 8 evaluated which of those candidates remain dynamically defensible under the adopted stability framework. Taken together, those results provide more than a catalog of trajectories: they provide a filtered collection of dynamical building blocks from which a conceptual tour architecture can be assembled.

The purpose of the present chapter is to interpret those building blocks in a mission-design context. The goal is not to present a fully optimized mission with complete maneuver design, timeline constraints, or high-fidelity ephemeris modeling. Instead, the chapter asks how the surviving classes of trajectories can be arranged into a coherent exploration strategy, including Neptune-centered operations, broader system touring, transfer into the Triton neighborhood, and localized motion near Triton itself. In this way, Chapter 9 serves as the final interpretive step that connects the dynamical analysis of the earlier chapters to a physically meaningful tour concept.

All trajectory visualizations in this chapter remain tied to the CR3BP-based modeling framework used throughout the thesis and are shown in the rotating reference frame. Consequently, the figures should be read as dynamical representations rather than exact geometric mission maps. The sizes of Neptune and Triton are not necessarily to scale, zoomed panels may distort relative distances, and apparent intersections between trajectories and bodies in the rotating frame should not be interpreted as physical collisions. Likewise, the additional major Neptune moons included in selected trajectory figures are shown only as visual references and are not incorporated as gravitating bodies in the equations of motion.

9.1 Introduction and Chapter Objectives

The dynamical analysis developed in Chapters 7 and 8 has reduced the broad Neptune–Triton phase space to a smaller set of candidate trajectory classes with clear dynamical roles. Some of these candidates are most naturally interpreted as local bounded or staging trajectories, while others are better viewed as touring or transfer-enabling structures whose value lies in access rather than long-term confinement. The objective of this chapter is to assemble those distinct dynamical roles into a mission-oriented framework that remains faithful to the results of the earlier analysis.

Rather than introducing new dynamical constructs, the discussion here synthesizes previously identified trajectory classes into a coherent exploration architecture for the Neptune–Triton system. The emphasis is therefore on mission function: which trajectories are best suited for Neptune-centered observation, which provide broader regional coverage, which support transfer through the bottleneck, and which are most natural for local Triton operations. In that sense, the chapter marks the transition from dynamical screening to mission interpretation, treating the selected structures as building blocks for tour design.

It is important to emphasize that the tour architecture presented here is conceptual in nature. The trajectories are derived from the CR3BP formulation and are not the result of a fully constrained optimization process. No maneuver sequences, timing constraints, or higher-fidelity perturbations are incorporated at this stage. Instead, the purpose is to demonstrate how the stable and mission-relevant

candidates surviving Chapter 8 may be combined to support distinct phases of an exploration strategy, while also noting where marginal or transfer-oriented candidates may still play a useful role. The trajectory figures should therefore be interpreted in the rotating-frame sense described at the start of the chapter rather than as final geometric mission layouts.

The chapter is organized according to the major dynamical phases of the proposed tour architecture. Section 9.2 introduces the overall mission framework and defines the sequence of trajectory classes considered. Sections 9.3, 9.4, 9.5, and 9.6 present the trajectory options associated with each phase of the mission. Finally, Section 9.8 discusses selected candidate types that were not retained as primary architectural elements, thereby clarifying the selection logic and preserving the broader dynamical context of the study.

9.2 Tour Architecture Framework

The construction of a coherent tour architecture within the Neptune–Triton system requires the organization of dynamically distinct trajectory classes into a structured sequence that reflects both the underlying phase-space geometry and the practical progression of a mission. Based on the dynamical analysis presented in the preceding chapters, the trajectory families considered in this work may be grouped into four primary operational regimes: Neptune-bounded motion, system-wide touring trajectories, transfer trajectories into the Triton region, and Triton-bounded motion.

These regimes are not arbitrarily defined, but instead emerge naturally from the topology of the CR3BP. At higher Jacobi constant values, motion is constrained to regions surrounding Neptune, resulting in stable, bounded trajectories that are well-suited for initial mission operations. As the energy level decreases, the zero-velocity curves open and permit trajectories that extend across a broader portion of the system, enabling access to the region populated by Neptune’s major moons. Further decreases in the Jacobi constant allow for transport through the neck regions near the collinear Lagrange points, facilitating transition into the dynamical domain of Triton. Finally, within the Triton region, stable bounded trajectories provide opportunities for localized operations analogous to those available in the Neptune-centered regime.

In this chapter, these dynamical regimes are interpreted as sequential phases of a conceptual mission architecture. The resulting framework is composed of the following stages:

1. **Neptune-Bounded Operations:** Stable, bounded trajectories centered about Neptune that support initial science operations and system characterization.
2. **System Touring Phase:** Extended trajectories that traverse a large portion of the Neptune system, providing geometric access to multiple moons and broader spatial coverage.
3. **Transfer to the Triton Region:** Trajectories that enable transition between Neptune-dominated motion and the dynamical environment surrounding Triton through the neck regions of the zero-velocity surface.
4. **Triton-Bounded Operations:** Stable, bounded trajectories centered about Triton that support detailed investigation of the moon and its immediate environment.

Each phase of this architecture is populated using trajectory candidates that have already been identified and evaluated on dynamical grounds. In particular, the nominal architectural elements are

drawn from the candidates retained through the stability analysis of Chapter 8. Trajectories that do not survive that filtering step are not treated as primary components of the architecture and are instead reserved for the later discussion of rejected or secondary options in Section 9.8.

It is important to note that the transitions between these phases are conceptual and do not imply direct, continuous trajectories without maneuvering. The CR3BP provides a valuable framework for identifying natural pathways and dynamical connections within the system, but the realization of a complete mission would require the incorporation of maneuver strategies, timing considerations, and higher-fidelity gravitational models. The objective here is therefore to establish a dynamically informed structure that defines how different classes of motion may be combined, rather than to prescribe a fully optimized trajectory sequence.

The sections that follow examine each phase of the tour architecture in detail, beginning with the identification of stable Neptune-bounded orbit options and progressing through the system touring, transfer, and Triton-bounded phases.

9.3 Neptune-Bounded Orbit Options

The initial phase of the proposed tour architecture is defined by trajectories that remain bounded in the vicinity of Neptune. These orbits represent the most dynamically reliable class identified in the preceding analysis and therefore provide the most natural starting point for mission operations. Within the CR3BP framework, such trajectories are associated with higher Jacobi constant values, where the zero-velocity curves restrict motion to the region surrounding the primary body. As a result, these trajectories exhibit regular, repeatable behavior and are well suited for long-duration observation and early system characterization.

The Neptune-bounded regime serves several important roles within the overall architecture. First, it provides a comparatively controlled environment for initial spacecraft checkout and calibration following system arrival. Second, it enables sustained observation of Neptune itself, including its atmosphere, magnetosphere, and ring system. Finally, these trajectories act as a staging region from which the spacecraft may later transition into broader system-exploration phases.

Two candidate trajectories are selected from the stable Neptune-centered solutions identified in Chapter 8 to represent this regime. Both exhibit bounded motion centered about Neptune, but they differ in radial extent and geometric character. In that sense, they provide not only two viable operational options, but also two different ways of balancing confinement against spatial coverage within the Neptune-centered domain.

It is important to distinguish between *dynamical classification* and *mission-role interpretation*. In Chapters 7 and 8, a broader set of Region II and Region III trajectories was retained under the Neptune-centered bounded classification because those candidates remained dynamically associated with the Neptune-centered domain and admitted useful periodic-orbit representations. When the full trajectories were then examined in physical space for mission interpretation, however, not all of those stable Neptune-centered candidates behaved in the same operational way. The more compact cases remained natural local Neptune orbits, while the larger-amplitude cases traced broader paths through the near-Neptune moon region and were therefore more appropriately reassigned to the touring phase of the architecture. The present section is thus intentionally restricted to the compact Neptune-centered options, while the larger Neptune-centered stable trajectories are discussed later as system-touring candidates.

9.3.1 Nominal Neptune Science Orbit

The first candidate, corresponding to an initial condition of $x_0 = -0.250079$, is selected as the nominal Neptune-bounded orbit. This trajectory, shown in Figure 9.1, exhibits tightly confined motion about Neptune, maintaining a consistent radial distance and demonstrating a high degree of regularity in its evolution. The resulting geometry is well-suited for sustained observation, as the spacecraft remains in close proximity to the primary body while avoiding large excursions into the broader system.




Figure 9.1. Neptune-bounded trajectory for $x_0 = -0.250079$ in the rotating frame. The motion remains tightly confined about Neptune, providing a stable platform for initial mission operations and detailed observation of the primary body.

NOTE: Neptune-system moon guides are visual references only. Not included in the propagated force model.

The stability of this trajectory, as established in Chapter 8, supports its role as the nominal option for the initial mission phase. Its relatively compact structure simplifies interpretation, keeps the spacecraft close to Neptune, and provides a clear dynamical baseline against which the broader touring and transfer trajectories introduced later in the chapter may be compared.

9.3.2 Alternative Neptune-Bounded Orbit

A second candidate, corresponding to $x_0 = 0.20992$, is included as an alternative Neptune-bounded orbit. While this trajectory, shown in Figure 9.2, remains clearly centered about Neptune, it exhibits a slightly larger radial extent and a more elongated geometry than the nominal case. This increased spatial coverage may provide advantages in terms of observational diversity, allowing the spacecraft to sample a broader region of the Neptune environment while still remaining within a bounded operational regime.

As with the nominal trajectory, this orbit satisfies the stability criteria established in Chapter 8, indicating that it also represents a viable option for sustained operations. The choice between these two Neptune-centered trajectories would therefore depend less on basic dynamical viability and

Neptune-Bounded Alternative Trajectory




Figure 9.2. Alternative Neptune-bounded trajectory for $x_0 = 0.20992$. The orbit remains stable and centered about Neptune, but with a larger radial extent and more elongated structure compared to the nominal case.

more on mission-specific priorities, such as desired observational geometry, coverage requirements, and operational constraints.

9.3.3 Summary of Neptune-Bounded Phase

The Neptune-bounded trajectories presented in this section define a stable and well-understood operational regime that serves as the foundation of the tour architecture. Both candidates provide long-duration bounded motion and maintain close association with Neptune, making them suitable for early mission phases centered on system arrival, checkout, and sustained primary-body observation. Just as importantly, they establish a reference regime against which the more expansive and dynamically complex trajectory classes considered in the following sections can be interpreted. The transition from these compact Neptune-centered options to broader touring and transfer trajectories is therefore not just a change in geometry, but a change in mission function.

9.4 System-Wide Touring Trajectories

Following the initial Neptune-bounded phase, the tour architecture expands into a class of trajectories that traverse a significantly larger portion of the Neptune system. These system touring trajectories arise at lower Jacobi constant values, where the zero-velocity curves open sufficiently to permit motion that extends beyond the immediate vicinity of Neptune. As a result, the spacecraft is no longer confined to a localized region, but instead follows trajectories that span a broad spatial domain, enabling access to the region populated by Neptune’s major moons.

Within the context of mission design, these trajectories represent the primary mechanism for system-wide exploration. Rather than maintaining a fixed orbital radius, the spacecraft naturally evolves along paths that sample multiple regions of the system, providing diverse observational geometry over time. It is important to emphasize, however, that the positions of the moons shown in

the visualizations are included as reference markers only. While these trajectories may pass near the indicated locations of the moons in the rotating frame, no direct gravitational interaction with the moons is modeled in the CR3BP formulation used here.

The touring phase also provides the clearest example in this chapter of why the final architectural interpretation cannot be read directly from the earlier dynamical labels alone. Two of the trajectories that survived the Neptune-centered bounded filtering in Chapters 7 and 8, namely the cases associated with $x_0 = 0.799920$ and $x_0 = -0.820079$, are dynamically Neptune-centered in the phase-space sense but behave quite differently when plotted as full trajectories. Rather than remaining as compact local Neptune orbits, both sweep broadly across the near-Neptune moon region and provide repeated geometric access to the four largest inner moons used here as visual reference bodies. For mission interpretation, those trajectories are therefore better understood as touring trajectories than as purely local Neptune-bounded science orbits. This section makes that reassignment explicit and uses operational geometry, rather than phase-space label alone, to define the touring role.

9.4.1 Nominal Touring Orbit

The primary touring trajectory selected for this study corresponds to an initial condition of $x_0 = -0.820079$. This orbit, shown in Figure 9.3, exhibits a broad spatial extent while maintaining a coherent and bounded structure, making it well suited for sustained system-level exploration. Unlike the compact Neptune-centered science orbits, it does not remain confined to a nearly circular neighborhood about Neptune. Instead, it sweeps repeatedly across the near-Neptune moon region, providing favorable flyby geometry and broad system coverage while still preserving a dynamically organized path.




Figure 9.3. Nominal system touring trajectory for $x_0 = -0.820079$. The orbit spans a large portion of the Neptune system while maintaining bounded behavior, enabling broad observational coverage.

The stability characteristics of this trajectory, as established in Chapter 8, indicate that it retains bounded motion over long time intervals. This combination of stability and spatial coverage makes it the most suitable candidate for the primary touring phase of the mission. In practical terms, it is

the strongest example of a trajectory that remains dynamically acceptable while also providing a mission-useful sequence of moon-region flybys.

9.4.2 High-Energy Alternative Touring Orbit

An alternative touring trajectory is identified for $x_0 = 0.799921$. This orbit, shown in Figure 9.4, also provides extensive spatial coverage and repeated access to the near-Neptune moon region, making it dynamically attractive from the standpoint of broad exploration. In trajectory space, it behaves much more like a fluid moon-tour path than a compact Neptune science orbit, which is why it is considered here rather than in the preceding section.




Figure 9.4. High-energy touring trajectory for $x_0 = 0.799921$. The orbit provides extensive spatial coverage across the Neptune system, offering broad access to the moon region.

Despite its favorable geometric properties, this trajectory introduces a significant operational concern. The path approaches Neptune at a very low periapsis altitude, bringing it uncomfortably close to the planetary radius in a way that would complicate its interpretation as a practical nominal tour trajectory. As a result, while the orbit is dynamically attractive and mission-relevant as a moon-flyby path, it is better treated as a high-risk or aggressive alternative rather than as the preferred touring option.

9.4.3 Secondary Touring Candidates

Additional trajectories were examined within this regime, including the case corresponding to $x_0 = -0.900079$. While this trajectory exhibits extended motion across the system, its geometry is less favorable for comprehensive exploration than the preferred -0.820079 case. In particular, it does not provide the same level of balanced moon-region coverage and is therefore better interpreted as a secondary touring or targeted-observation option rather than as the nominal system tour.

Due to these limitations, such trajectories are not selected as primary components of the tour architecture, but are noted here for completeness and comparison.

9.4.4 Summary of Touring Phase

The system touring trajectories represent a transition from localized Neptune-centered motion to broad exploration of the Neptune system. Among the candidates considered, the trajectory corresponding to $x_0 = -0.820079$ provides the most balanced combination of stability, moon-region access, and operational plausibility, making it the preferred choice for the nominal touring phase. Alternative trajectories, especially the case at $x_0 = 0.799920$, confirm that the larger Neptune-centered stable families can indeed support a moon-tour interpretation, but they also show that not every geometrically attractive tour path is equally practical once close Neptune approaches are taken into account.

These results highlight the importance of considering both dynamical structure and operational constraints when selecting trajectories for mission design. The touring phase serves as a critical bridge between the initial Neptune-bounded regime and the transfer into the Triton region, and therefore plays a central role in the overall architecture.

9.5 Transfer to the Triton Region

The transition from Neptune-centered motion to the dynamical environment of Triton represents a critical stage in the tour architecture. Within the CR3BP framework, this transition is governed by the topology of the zero-velocity surfaces, which define the permissible regions of motion for a given Jacobi constant. As the energy level decreases, the zero-velocity curves open near the collinear Lagrange points, creating neck regions that allow trajectories to pass between previously disconnected domains. These openings provide the dynamical pathways through which a spacecraft may transition from the Neptune-dominated region into the vicinity of Triton.

While the existence of these neck regions suggests a natural mechanism for transport, not all trajectories that approach or traverse them result in meaningful transfer into a Triton-bounded orbit. The geometry of the trajectory, as well as its dynamical compatibility with the target region, plays a critical role in determining whether a viable transition is achieved. Among the candidates examined in this study, two exhibit genuine neck-crossing behavior and are retained as transfer elements in the architecture: the retrograde case at $x_0 = 0.93992$, which serves as the primary transfer candidate, and the prograde case at $x_0 = 0.97992$, which is retained as a secondary alternative.

One point of notation deserves explicit clarification. In this thesis, the seed value x_0 identifies the location of the trajectory on the chosen section, but it does not by itself determine the direction of motion. A prograde and a retrograde trajectory may therefore share the same x_0 value while still representing different physical initial conditions because the sign convention associated with the direction of motion is different. The two transfer candidates below share neither their x_0 values nor their directional sense and thus represent genuinely distinct dynamical realizations.

9.5.1 Transfer Candidate 1: Retrograde Neck Crossing

The primary transfer candidate is the retrograde case at $x_0 = 0.93992$. This trajectory, shown in Figure 9.5, passes through the neck region of the zero-velocity surface and therefore represents a geometrically meaningful transfer path into the Triton-accessible domain. In the final architectural interpretation adopted here, it is treated as the preferred transfer option because it demonstrates clear bottleneck-crossing behavior while maintaining a dynamically coherent path through the transition

region.




Figure 9.5. Retrograde bottleneck-crossing trajectory for $x_0 = 0.93992$. This case provides a geometrically meaningful neck crossing and is selected as the primary transfer candidate in the tour architecture.

This candidate highlights the fact that transfer usefulness cannot be judged by stability alone. In the architecture context, a trajectory that actually crosses the bottleneck and reaches the Triton-accessible region is the essential requirement for a transfer role. The retrograde $x_0 = 0.93992$ case satisfies this requirement and is therefore retained as the nominal transfer option in the mission-level interpretation.

9.5.2 Transfer Candidate 2: Prograde Neck Crossing

The second transfer candidate is the prograde case at $x_0 = 0.97992$. This trajectory, shown in Figure 9.6, also provides genuine neck-crossing behavior, reaching the portion of the phase space relevant to transfer into the Triton-centered regime. Within the final transfer set it is better interpreted as a secondary alternative to the preferred retrograde candidate rather than as a co-equal nominal transfer path, but it remains a useful dynamical reference for understanding the range of viable transfer geometries available within the model.

Together, the retrograde $x_0 = 0.93992$ case and the prograde $x_0 = 0.97992$ case define the transfer options that remain relevant after the geometry of the full trajectories is taken into account. The retrograde case is retained as the preferred transfer trajectory, while the prograde case is preserved as an informative alternative. Neither should be interpreted as a final optimized transfer solution, but together they provide the most useful dynamical evidence in the present study that passage from the touring regime into the Triton neighborhood can be achieved within the structure of the model.




Figure 9.6. Prograde transfer trajectory for $x_0 = 0.97992$. This case provides a second geometrically meaningful bottleneck-crossing path and is retained as a secondary transfer option within the tour architecture.

9.5.3 Summary of Transfer Phase

The transfer trajectories examined in this section illustrate the importance of combining stability interpretation with full-trajectory geometry. While the opening of the zero-velocity surfaces provides the necessary pathways for transport, a candidate must not only be dynamically interesting but must also demonstrate actual bottleneck-crossing behavior to serve a transfer role in the architecture.

The transfer phase is represented by two candidates that satisfy this geometric requirement: the retrograde $x_0 = 0.93992$ path, which serves as the primary transfer element, and the prograde $x_0 = 0.97992$ path, which is retained as a secondary alternative. The comparison reinforces a broader lesson of the thesis: mission architecture cannot be built from stability information alone, but must also account for the realized geometry of the trajectories themselves.

9.6 Triton-Bounded Orbit Options

The final phase of the tour architecture is defined by trajectories that remain bounded in the vicinity of Triton. Following transfer through the neck region, the spacecraft enters a dynamical regime in which motion is governed primarily by the gravitational influence of Triton within the CR3BP framework. As in the Neptune-centered case, this regime supports stable, repeatable trajectories that are well-suited for sustained observation and detailed investigation of the target body.

From a mission perspective, the Triton-bounded phase represents the terminal science orbit, where the spacecraft transitions from system-level exploration to focused operations within the immediate environment of Triton. These trajectories provide opportunities for high-resolution observation of Triton's surface, atmosphere, and potential geological activity, while maintaining bounded motion that supports long-duration stability.

Two candidate trajectories are selected to represent this regime, corresponding to prograde and retrograde solutions at $x_0 = 1.00992$. Both cases exhibit tightly confined motion about Triton, as

compared in Figure 9.7, and satisfy the stability criteria established in Chapter 8, indicating that they are viable options for terminal operations.

As in the transfer section, the repeated appearance of the same seed value does not indicate a typographical error. The prograde and retrograde Triton-centered cases share the same section location $x_0 = 1.00992$, but differ in their directional sense and therefore correspond to distinct trajectory realizations within the rotating-frame dynamics.




Figure 9.7. Triton-bounded trajectory options at $x_0 = 1.00992$. The upper panel shows the nominal prograde case, while the lower panel shows the retrograde alternative. Both remain tightly confined about Triton and define viable terminal science configurations within the rotating-frame model.

9.6.1 Nominal Triton Science Orbit

The nominal Triton-bounded trajectory is selected as the prograde solution at $x_0 = 1.00992$. This orbit exhibits a compact and regular structure centered about Triton, maintaining close proximity to the moon while avoiding large excursions into the surrounding region. The resulting geometry is well-suited for detailed observation and provides a stable platform for extended scientific operations.

The stability of this trajectory ensures that it remains bounded over long time intervals, making it an appropriate endpoint for the tour architecture. Its compact structure also facilitates clear interpretation and supports sustained observation of Triton's local environment.

9.6.2 Alternative Triton Orbit

An alternative Triton-bounded trajectory is provided by the retrograde solution at $x_0 = 1.00992$. While this orbit shares many similarities with the nominal case in terms of boundedness and overall structure, it exhibits a reversed direction of motion in the rotating frame, resulting in a distinct geometric configuration.

This alternative trajectory satisfies the same stability criteria as the prograde case and may offer different observational geometry or operational advantages depending on mission requirements.

The inclusion of both prograde and retrograde solutions highlights the diversity of stable motion within the Triton region and provides flexibility in the design of the terminal mission phase.

9.6.3 Summary of Triton-Bounded Phase

The Triton-bounded trajectories presented in this section define the final stage of the tour architecture, where the spacecraft transitions into sustained, localized operations about the target moon. Both the prograde and retrograde solutions at $x_0 = 1.00992$ provide stable, bounded motion and are suitable for long-duration scientific investigation.

This phase represents the culmination of the dynamical progression established throughout the chapter, completing the sequence from Neptune-centered motion, through system-wide exploration and transfer, to focused operations within the Triton environment.

9.7 Integrated Tour Architecture

The trajectory classes developed in the preceding sections may be combined to form a coherent, dynamics-informed tour architecture within the Neptune–Triton system. Each class represents a distinct dynamical regime, and when arranged sequentially, these regimes define a structured progression of motion that supports both system-level exploration and targeted investigation of Triton. The resulting architecture is not a single continuous trajectory, but rather a framework composed of dynamically compatible phases that may be connected through appropriate transfer strategies.

9.7.1 Sequential Mission Architecture

The proposed tour architecture is constructed as a sequence of four primary phases, each corresponding to a specific class of motion identified through the dynamical analysis:

1. **Neptune-Centered Operations:** The architecture begins with a compact Neptune-centered bounded orbit, represented by the nominal case at $x_0 = -0.250079$, or alternatively by the broader local Neptune option at $x_0 = 0.20992$. This phase supports arrival, checkout, and sustained observation of Neptune before the spacecraft is committed to broader excursions through the system.
2. **System Touring Phase:** The spacecraft then transitions to a system-spanning trajectory, represented nominally by the touring orbit at $x_0 = -0.820079$. In this phase, the trajectory traverses a large portion of the near-Neptune system and provides repeated geometric access to the region associated with the major moons. The higher-amplitude case at $x_0 = 0.799920$ remains available as a more aggressive alternative, but is not treated as nominal because of its very close Neptune approach.
3. **Transfer Phase:** The spacecraft next utilizes a bottleneck-crossing trajectory to move from Neptune-dominated motion into the Triton-accessible domain. In the final transfer interpretation of this thesis, the preferred candidate is the retrograde case at $x_0 = 0.93992$, with the prograde case at $x_0 = 0.97992$ retained as a secondary alternative. Both candidates exhibit genuine neck-crossing behavior and are retained on that geometric basis.

4. **Triton-Bounded Phase:** The tour concludes with insertion into a stable, bounded trajectory about Triton, such as the nominal prograde orbit at $x_0 = 1.00992$ or the corresponding retrograde alternative at the same seed value. This phase supports sustained, high-resolution investigation of Triton and represents the terminal stage of the mission.

Taken together, these phases define a progression from localized Neptune-centered motion to system-wide exploration, then onward through a narrow transfer gateway into focused operations within the Triton environment. This structure reflects the natural organization of the phase space and demonstrates how distinct dynamical regimes may be leveraged to support different mission objectives within a single dynamical architecture.

9.7.2 Nominal and Alternative Architectures

While the sequence described above defines a nominal tour architecture, alternative configurations may be constructed by substituting trajectory candidates within each phase. For example, the Neptune-bounded phase may be initialized using the alternative orbit at $x_0 = 0.20992$, which provides a larger radial extent and increased observational diversity. Similarly, the system touring phase may be shifted from the nominal case at $x_0 = -0.820079$ to the higher-amplitude path at $x_0 = 0.799920$, trading greater moon-region coverage for increased operational risk due to the very low Neptune approach.

In the transfer phase, the retrograde $x_0 = 0.93992$ case and the prograde $x_0 = 0.97992$ case represent two distinct ways of entering the Triton-accessible domain. The former is the preferred architectural choice as the primary neck-crossing transfer candidate, while the latter remains a secondary alternative. In the Triton-bounded phase, both the prograde and retrograde solutions at $x_0 = 1.00992$ provide viable options with distinct geometric characteristics. The availability of these alternatives highlights the flexibility of the architecture and demonstrates that multiple mission profiles may be constructed from the same underlying dynamical framework.

However, the selection of specific trajectories must balance dynamical properties with operational considerations. As discussed earlier in the chapter, trajectories that provide extensive spatial coverage may also introduce undesirable conditions, such as close approaches to Neptune, while trajectories that appear favorable in local stability analysis may still fail the geometric requirements of a transfer. The nominal architecture therefore represents a compromise between coverage, dynamical defensibility, and operational plausibility.

9.7.3 Key Design Insights

The construction of the tour architecture from dynamically derived trajectory classes provides several important insights into the design of missions within the Neptune–Triton system.

First, stability serves as a fundamental constraint on trajectory selection. The filtering process established in Chapters 7 and 8 reduces the broad geometric candidate set to a smaller number of trajectories that remain dynamically defensible under the periodic-orbit framework. This highlights the importance of stability analysis as a prerequisite for mission design, even when stability alone is not the final architectural criterion.

Second, geometric accessibility alone is insufficient for trajectory selection. As demonstrated in the transfer analysis, trajectories that appear to provide access to a desired region may fail to

produce meaningful or sustained motion within that region. Dynamical compatibility, rather than visual proximity, determines the success of a transfer.

Third, trajectories that exhibit the most extensive spatial coverage are not necessarily the most suitable for mission implementation. High-energy trajectories may provide broad access to the system, but can introduce operational challenges that limit their practical use. This reinforces the need to consider both dynamical behavior and physical constraints when evaluating candidate solutions.

Finally, the CR3BP framework provides a powerful tool for identifying the underlying structure of motion within the system, but does not, by itself, yield a complete mission design. The architecture developed in this chapter should therefore be interpreted as a dynamically informed framework that defines how different classes of motion may be combined, rather than as a fully optimized trajectory sequence.

The integration of these trajectory classes into a coherent architecture represents the culmination of the dynamical analysis presented in this work, demonstrating how phase-space structure, resonance identification, and stability characterization may be directly translated into mission-relevant design concepts.

9.8 Rejected and Non-Nominal Trajectory Classes

In addition to the trajectory classes incorporated into the tour architecture, a number of dynamically interesting solutions were identified during the analysis of Chapters 7 and 8. While these trajectories exhibit complex geometric structure and, in some cases, extensive spatial coverage, they do not satisfy the combined stability and mission-role criteria required for inclusion in the primary framework. The purpose of this section is to document these non-nominal trajectory classes, thereby clarifying the selection process and preserving the broader dynamical context of the Neptune–Triton system.

9.8.1 Retrograde Trajectory Families

Several retrograde trajectory families were identified that display highly structured and often symmetric motion within the rotating frame. Representative examples include the cases corresponding to $x_0 = -0.720079$ and $x_0 = 0.599921$, both retrograde, which exhibit large-scale excursions and visually coherent patterns, as shown in Figure 9.8, suggestive of underlying resonance structure.

Despite their visual appeal and apparent dynamical organization, these trajectories were not selected for use in the tour architecture because they failed to survive the combined stability and mission-role filtering adopted in the thesis. As shown in Chapter 8, retrograde solutions within these regions often fail to produce satisfactory periodic-orbit representations or remain too sensitive to support a primary mission role. They therefore remain valuable for dynamical interpretation, but not as retained architectural elements.

9.8.2 High-Energy and Weakly Bounded Prograde Trajectories

In addition to the retrograde families, certain prograde trajectories were identified that exhibit large excursions across the system but lack the bounded structure required for practical use. An example of this behavior is observed in the case corresponding to $x_0 = -0.92007$ (prograde), shown in




Figure 9.8. Representative retrograde trajectories exhibiting structured, large-scale motion within the Neptune–Triton system. While dynamically rich, these trajectories do not satisfy the stability criteria required for inclusion in the tour architecture.

Figure 9.9, which traverses a wide spatial domain but does not maintain a coherent or repeatable pattern over extended time intervals.

These trajectories are characterized by their extensive spatial reach but lack the dynamical coherence necessary for sustained operations. As such, they are not incorporated into the tour architecture, which prioritizes stability and repeatability.

9.8.3 Interpretation of Rejected Trajectories

The trajectory classes presented in this section highlight the distinction between dynamical richness and mission suitability. While the CR3BP admits a wide variety of complex and visually compelling trajectories, only a subset of these satisfy the combined requirements of stability, boundedness, and operational utility.

The exclusion of these trajectories from the tour architecture is therefore not indicative of a lack of dynamical significance, but rather reflects the application of the filtering criteria established in Chapters 7 and 8. In particular, the retained architecture favors candidates that combine useful trajectory geometry with a defensible degree of stability or transfer relevance. Trajectories that fail one side of that test may still be scientifically interesting, but they are not carried forward as primary mission elements.

9.8.4 Implications for Future Work

Although these trajectories are not utilized within the present tour architecture, they may represent viable candidates under alternative mission scenarios or extended modeling frameworks. For example,

Rejected High-Energy Prograde Trajectory




Figure 9.9. High-energy prograde trajectory for $x_0 = -0.92007$. Although the trajectory spans a large portion of the system, it does not exhibit the bounded, stable behavior required for mission applications.

retrograde trajectories could be reconsidered for short-duration missions or targeted observation campaigns, where long-term stability is not a primary requirement. Similarly, higher-fidelity models that incorporate additional gravitational influences may alter the stability characteristics of these solutions.

As such, the trajectory classes presented here provide a basis for future investigation and highlight the richness of the dynamical environment within the Neptune–Triton system beyond the subset of trajectories selected for this study.

9.9 Mission Architecture Implications and Forward Outlook

The tour architecture developed in this chapter represents a direct synthesis of the dynamical structures identified throughout this work. By combining stable Neptune-bounded trajectories, system-spanning touring orbits, dynamically viable transfer pathways, and stable Triton-centered motion, a coherent framework for exploration of the Neptune–Triton system has been established. This framework demonstrates that the complex phase-space structure of the CR3BP can be interpreted not only as an abstract dynamical system, but as a set of practical pathways that may be leveraged for mission design.

A key outcome of this analysis is the identification of a structured progression of motion that aligns naturally with mission objectives. The transition from localized Neptune-centered operations to system-wide exploration, followed by transfer into the Triton region and eventual insertion into a Triton-bounded orbit, reflects the intrinsic organization of the dynamical environment. This progression is not imposed externally, but rather emerges from the topology of the zero-velocity surfaces and the stability properties of the associated trajectory families.

At the same time, the results of this chapter highlight several important limitations. The trajectories presented here are derived from the CR3BP formulation and therefore do not account for additional gravitational influences, including the effects of Neptune’s moons, solar perturbations, or

non-uniform mass distributions. Furthermore, the architecture does not include maneuver strategies, timing constraints, or optimization of transfer trajectories. As such, the framework should be interpreted as a first-order approximation of the available dynamical pathways, rather than a finalized mission design.

These limitations do not diminish the value of the results, but instead define the scope of their applicability. The architecture provides a structured foundation upon which higher-fidelity models and mission-specific constraints may be incorporated. In particular, the trajectory classes identified here may serve as initial guesses for numerical optimization, guide the selection of transfer strategies, and inform the design of observation sequences within the system.

The broader implication of this work is that the methods employed—including Poincaré analysis, frequency-based resonance identification, and stability classification through periodic orbit refinement—form a unified framework for exploring and exploiting the dynamical structure of multi-body systems. While applied here to the Neptune–Triton system, these methods are directly extensible to other planetary and satellite systems, offering a general approach to trajectory design in complex gravitational environments.

With the dynamical architecture now established, the remaining questions center on how this framework may be extended, refined, and applied within a more complete mission-design context. These broader implications and the most important next steps are taken up in the following chapter.

10. Discussion and Future Work

The preceding chapters have developed a complete dynamical framework for the analysis and interpretation of motion within the Neptune–Triton system. Building on the CR3BP-based model established earlier in the thesis, together with the J_2 extension used for Neptune, the analysis progressed through phase-space characterization using Poincaré maps, frequency-based resonance identification, candidate-region selection, and periodic-orbit stability analysis. These tools were then synthesized in Chapter 9 to construct a coherent mission architecture, demonstrating how the underlying dynamical structure of the system can be directly leveraged for trajectory design.

It is important to emphasize that the results presented in this work are not isolated numerical artifacts, but rather manifestations of the intrinsic geometric and dynamical organization of the Neptune–Triton system. The identification of stable periodic orbits, quasi-periodic invariant regions, resonance-supported families, and transport pathways reflects the existence of structured phase-space features that persist across a range of energy levels and modeling assumptions. As such, the trajectory classes and mission architecture developed herein should be interpreted as first-order dynamical frameworks that guide, rather than replace, higher-fidelity mission design.

At the same time, the present work is intentionally constrained in scope. The analysis is conducted within a rotating reference frame under the assumptions of the CR3BP, with optional inclusion of Neptune’s J_2 term, and is primarily restricted to planar motion. Additional perturbations, including the gravitational influence of other Neptunian moons, solar effects, and higher-order gravitational harmonics, are not included. Furthermore, the trajectories analyzed are not optimized with respect to maneuver cost, timing constraints, or operational considerations. These simplifications are deliberate, as they allow the fundamental dynamical structure of the system to be isolated and understood without the confounding effects of higher-order complexities.

Accordingly, the purpose of this chapter is twofold. First, it provides a structured interpretation of the results within the broader context of mission design, clarifying what aspects of the developed framework are directly applicable and what aspects still require refinement. Second, it establishes a clear and technically grounded set of directions for future work, outlining the steps required to transition from the idealized dynamical models employed in this thesis to a more complete mission-design capability. These extensions include both higher-fidelity modeling efforts and deeper dynamical investigations, each of which builds directly upon the foundation established in the preceding chapters.

In this sense, the present work should be viewed not as a terminal analysis, but as a foundational layer in a larger research program aimed at enabling systematic exploration of multi-body dynamical environments. The results obtained here define the structure of that environment; the work that follows will determine how it can be navigated in practice.

10.1 Scope of the Present Work

The analysis presented in this thesis is conducted within a deliberately constrained dynamical framework in order to isolate and characterize the fundamental structures governing motion in the Neptune–Triton system. As established in Chapters 2 and 3, that framework is the synodic-frame CR3BP adapted to the Neptune–Triton problem, with the third body treated as massless and the two primaries providing the governing gravitational field.

An extension to this model is introduced through the inclusion of Neptune’s oblateness via a J_2 perturbation, modifying the effective potential and, consequently, the phase-space structure of the system. This augmentation provides a first-order correction to the spherical primary assumption and enables examination of how deviations from idealized gravitational symmetry influence equilibrium structure, invariant regions, and trajectory stability. However, no additional gravitational harmonics beyond J_2 are considered, and Triton is treated as a point mass without internal structure or higher-order gravitational effects. Within that scope, the thesis is concerned primarily with the identification of dynamical structure: where bounded motion persists, where resonant organization appears, where transport channels open, and where locally stable periodic representatives can be found.

The analysis is primarily restricted to planar motion, with trajectories confined to the rotating-frame x – y plane. As a result, vertical dynamics and out-of-plane stability characteristics are not captured in the present study. All trajectory classifications, including Neptune-bounded, touring, transfer, and Triton-bounded motion, are therefore defined within this planar context. While this restriction significantly simplifies the dynamical system and facilitates visualization and interpretation through Poincaré maps, it necessarily excludes three-dimensional trajectory families that may play an important role in practical mission design.

All results are expressed in the rotating reference frame, and no direct transformation to inertial-frame trajectories is performed for mission realization. Consequently, temporal phasing, absolute trajectory timing, and alignment with external bodies are not addressed. The trajectories analyzed in this work should therefore be interpreted as geometric and dynamical constructs, rather than fully specified mission trajectories with defined launch windows, encounter sequences, or navigation timelines.

Furthermore, the analysis does not incorporate maneuver modeling or trajectory optimization. Transitions between trajectory classes are identified based on dynamical accessibility within the phase space, but no Δv costs, targeting strategies, or control laws are computed. Stability is assessed through periodic-orbit identification and linear stability analysis using Floquet multipliers, providing insight into local behavior near periodic solutions but not addressing nonlinear stability, long-term robustness, or sensitivity to perturbations.

These assumptions and limitations are not deficiencies, but rather intentional design choices that allow the underlying dynamical structure of the Neptune–Triton system to be examined in a controlled and interpretable manner. By constraining the problem in this way, the present work establishes a clear baseline from which more complex and higher-fidelity models can be systematically developed. That baseline is valuable precisely because later refinements can be evaluated against a well-understood reference problem rather than introduced into an already ambiguous design space.

10.2 Interpretation of Results in a Mission Design Context

The dynamical structures identified throughout this work provide a framework for interpreting motion in the Neptune–Triton system from a mission design perspective. Rather than viewing trajectories as isolated numerical solutions, the phase-space structures revealed through Poincaré maps, frequency analysis, and stability characterization define organized regions of motion that can be directly associated with distinct operational roles within a mission architecture.

Stable periodic orbits, identified through differential correction and classified via Floquet

multipliers, serve as foundational elements within this framework. These orbits represent dynamically consistent, repeatable motion that can act as long-duration staging locations or observation platforms. In the context of the Neptune–Triton system, such orbits provide natural candidates for sustained proximity operations around Neptune or Triton, depending on their spatial localization and stability characteristics.

Surrounding these periodic solutions are quasi-periodic invariant regions, which appear in the Poincaré maps as coherent island structures. These regions correspond to bounded, non-repeating motion that remains confined over long timescales. From a mission design standpoint, these quasi-periodic trajectories define operational envelopes that offer flexibility in trajectory selection while preserving stability. They may be interpreted as families of trajectories that enable extended exploration within a bounded region without the strict constraints of periodic motion.

In contrast, regions of phase space characterized by instability or transition behavior serve a different, but equally important, role. These regions are associated with trajectories that traverse between dynamically distinct zones, such as those connecting Neptune-bounded motion to Triton-bounded motion. Although such trajectories are not suitable for long-term residence, they represent natural transport pathways within the system. When properly understood and utilized, these pathways provide the mechanism by which a spacecraft may transition between mission phases with minimal control input.

The classification of trajectories into Neptune-bounded, touring, transfer, and Triton-bounded categories, as developed in Chapter 9, emerges directly from this underlying phase-space structure. Neptune-bounded trajectories correspond to motion confined near the primary, while Triton-bounded trajectories represent motion localized around the secondary. Touring trajectories occupy intermediate regions that enable repeated traversal of the system, and transfer trajectories facilitate transitions between these regimes. This classification is not imposed externally, but rather arises naturally from the geometry of the dynamical system.

It is important to emphasize that this interpretation is inherently geometric and dynamical, rather than operational in the traditional sense of trajectory design. The structures identified here do not prescribe specific maneuver sequences or mission timelines. Instead, they define a set of dynamically consistent pathways and regions that can be leveraged in subsequent design stages. In this way, the present analysis establishes a conceptual bridge between the abstract phase-space structures of the CR3BP and the practical requirements of mission architecture development.

The value of this framework lies in its ability to reduce the complexity of trajectory design by organizing the space of possible motions into a small number of interpretable categories. By identifying where stable motion exists, where transitions are possible, and how these regions are connected, the analysis provides a systematic basis for constructing mission concepts that exploit, rather than oppose, the natural dynamics of the Neptune–Triton system.

Just as importantly, the thesis shows that no single analytical lens is sufficient on its own. Resonant structure, local linear stability, and realized trajectory geometry each provide different information about candidate usefulness. Some trajectories that appear attractive in resonance space prove less compelling once their full geometry is examined, while some locally stable candidates must still be screened against mission-role requirements such as bottleneck crossing, spatial coverage, and proximity to the primary. The resulting mission architecture is therefore strongest when these layers of interpretation are used together rather than in isolation.

10.3 Model Limitations and Assumptions

This section does not restate the full model definition introduced earlier in the thesis. Instead, it identifies the practical consequences of the modeling assumptions summarized in Section 10.1 and clarifies how those assumptions constrain the interpretation of the present results.

The primary modeling assumption is the use of the Circular Restricted Three-Body Problem (CR3BP), in which Neptune and Triton are treated as point masses in circular orbits about their common barycenter, and the third body is assumed to have negligible mass. This formulation neglects the gravitational influence of all additional bodies, including the Sun and other Neptunian moons, thereby removing time-dependent perturbations that would otherwise modify the structure of the phase space. While this simplification enables a clear identification of invariant structures, it limits the direct applicability of the results to real mission environments where such perturbations are non-negligible.

An extension to the classical CR3BP is included through the addition of Neptune's J_2 oblateness term. This provides a first-order correction to the gravitational potential of the primary; however, higher-order zonal harmonics and any non-axisymmetric contributions are neglected. Additionally, Triton is modeled as a purely spherical point mass, and its physical characteristics beyond mass are not represented. As a result, the gravitational model remains an approximation, particularly in regions where higher-order effects may influence long-term trajectory behavior.

The analysis is further restricted to planar motion within the rotating reference frame. All trajectories are confined to the x - y plane, and no out-of-plane dynamics are considered. This eliminates vertical motion, inclination effects, and three-dimensional stability behavior from the analysis. While the planar approximation captures many of the dominant features of the system, it necessarily excludes entire classes of spatial trajectories, including halo and quasi-halo orbits, which may be of interest for mission applications. The tour architecture developed in Chapter 9 should therefore be interpreted as a planar dynamical template rather than a final three-dimensional mission design.

All results are expressed in the synodic (rotating) frame, and no direct mapping to the inertial frame is performed for trajectory realization. Consequently, the analysis does not account for absolute timing, phasing with external bodies, or the evolution of the system with respect to the Sun. This limits the ability to directly translate the identified trajectories into mission timelines or to assess their feasibility under real ephemeris conditions.

The trajectory analysis is conducted without the inclusion of maneuver modeling or control. Transitions between trajectory classes are identified based on dynamical accessibility within the phase space, but no Δv costs, thrust profiles, or targeting strategies are computed. As such, the results describe what is dynamically possible, rather than what is operationally optimal or achievable under propulsion constraints. This distinction is especially important in the transfer regime, where geometric access through a bottleneck does not automatically imply a practical transfer once timing, control effort, and safety margins are introduced.

Finally, stability is assessed through the identification of periodic orbits and the computation of their associated Floquet multipliers. This provides a measure of linear stability in the vicinity of periodic solutions but does not capture nonlinear effects, finite-time behavior, or sensitivity to perturbations over extended durations. Trajectories classified as stable within this framework may still exhibit complex behavior when subjected to higher-order perturbations or extended mission timescales. The distinction made in Chapter 9 between local stability and full mission usefulness is

therefore an important practical reminder of the limits of the present analysis.

These limitations are inherent to the modeling approach and are not indicative of deficiencies in the analysis. Rather, they define the scope within which the results are valid and provide a clear set of conditions under which the identified dynamical structures can be interpreted. The following sections build upon this foundation by outlining the extensions required to bridge the gap between the present idealized framework and a fully realized mission design methodology.

10.4 Required Extensions for Real Mission Design

The dynamical framework developed in this work provides a structured understanding of motion within the Neptune–Triton system; however, additional modeling and analysis are required to transition from this idealized representation to a fully realizable mission design. The extensions outlined in this section define the necessary steps for bridging that gap, progressing from simplified dynamical constructs to operationally viable trajectories.

10.4.1 Higher-Fidelity Dynamical Models

The CR3BP, even with the inclusion of Neptune’s J_2 term, represents an idealized approximation of the Neptune–Triton system. A critical next step is the incorporation of additional gravitational influences, including the Sun and the major Neptunian moons, within an N -body framework. These additions introduce time-dependent perturbations that modify the structure of invariant regions, alter stability properties, and influence long-term trajectory evolution.

Such models enable assessment of whether the phase-space structures identified in the CR3BP persist under more realistic conditions, and whether trajectories derived from the simplified model remain bounded or require correction. The inclusion of ephemeris-based motion for secondary bodies further allows for evaluation of encounter timing, resonance effects, and mission sequencing constraints.

10.4.2 Trajectory Refinement and Targeting

Trajectories identified within the CR3BP framework must be refined to satisfy the constraints of higher-fidelity models. This process typically involves the use of differential correction techniques and continuation methods to adjust initial conditions such that trajectories remain consistent when additional perturbations are introduced.

Refinement also includes the transformation of trajectories from the rotating frame to the inertial frame, enabling the specification of absolute position and velocity states as functions of time. This step is essential for aligning trajectories with mission events, such as planetary encounters, flybys, or insertion phases.

10.4.3 Maneuver Design and Δv Budgeting

The trajectories analyzed in this work are purely dynamical and do not account for propulsion requirements. To translate these trajectories into mission concepts, it is necessary to determine the maneuver sequences required to insert a spacecraft into the desired dynamical structures, transition between trajectory classes, and maintain bounded motion over extended durations.

This includes the computation of Δv requirements for key mission phases, such as Neptune system insertion, transfers between Neptune-bounded and Triton-bounded regimes, and potential orbit maintenance strategies. Optimization techniques may be employed to minimize fuel consumption while preserving the dynamical advantages identified in the present analysis.

10.4.4 Time-Dependent Mission Analysis

A fully realized mission design must account for the temporal evolution of the system. This includes the phasing of Neptune, Triton, and additional moons, as well as the position of the Sun relative to the system. Time-dependent analysis enables the identification of feasible launch windows, encounter opportunities, and sequencing of mission phases.

Incorporating real ephemerides allows for the evaluation of how the dynamical structures identified in the CR3BP evolve over time and whether trajectories remain viable under changing geometric conditions. This step is essential for determining whether the conceptual mission architectures developed in this thesis can be implemented within realistic mission timelines.

Collectively, these extensions transform the present dynamical framework into a practical design tool. The CR3BP-based analysis provides the structural foundation, while higher-fidelity modeling, trajectory refinement, maneuver design, and time-dependent analysis supply the necessary detail to realize a mission within the Neptune–Triton system. In that sense, the present thesis should be viewed as defining the architecture of the problem, while these future extensions define the path toward an implementable solution.

10.5 Extended Dynamical Investigations

Beyond the extensions required for immediate mission realization, the results of this thesis naturally motivate a set of deeper dynamical investigations aimed at further understanding the structure and behavior of the Neptune–Triton system. These investigations are not strictly necessary for constructing a baseline mission architecture; however, they represent critical avenues for advancing both the theoretical understanding of the system and the robustness of trajectory design methodologies. They also provide a means of testing which conclusions drawn in the current idealized framework are genuinely persistent and which are more model-dependent.

10.5.1 Time-Dependent Multi-Moon Modeling

The present analysis considers only the Neptune–Triton system, neglecting the gravitational influence of additional Neptunian moons. A natural extension is the incorporation of these bodies as time-dependent gravitating masses within an N -body framework. This introduces periodic and quasi-periodic perturbations that can significantly influence trajectory evolution, particularly for touring and transfer trajectories that traverse large regions of the system.

Such modeling enables the investigation of repeated flyby opportunities, resonance interactions with additional moons, and the persistence or deformation of invariant structures identified in the CR3BP. It also provides a pathway toward evaluating the feasibility of multi-moon tour architectures within the broader Neptunian system.

10.5.2 Spatial Extension of the Neptune–Triton Model

The restriction to planar motion excludes a wide class of dynamical behavior associated with out-of-plane motion. Extending the analysis to the full spatial three-body problem allows for the identification and characterization of three-dimensional trajectory families, including halo and quasi-halo orbits, as well as vertically stable and unstable structures.

This extension also enables the study of vertical stability properties of periodic orbits and the coupling between in-plane and out-of-plane dynamics. Such effects may play a significant role in long-duration missions, particularly in the presence of perturbations that drive motion out of the plane.

10.5.3 Transfer Manifold Construction

While the present work identifies trajectory classes and their relative accessibility within the phase space, it does not explicitly construct the dynamical pathways that connect them. The computation of invariant manifolds associated with periodic orbits provides a systematic method for identifying these pathways.

Stable and unstable manifolds define natural transport channels that can be used to transition between Neptune-bounded, touring, transfer, and Triton-bounded regimes. By explicitly computing and analyzing these manifolds, it becomes possible to design dynamically consistent transfer sequences that exploit the underlying structure of the system, potentially reducing maneuver requirements and increasing mission efficiency.

10.5.4 Higher-Order Gravitational Harmonics

The inclusion of Neptune’s J_2 term represents a first-order correction to the gravitational potential; however, higher-order zonal harmonics may also influence trajectory behavior, particularly over long timescales or in regions close to the primary. Extending the gravitational model to include additional harmonics enables assessment of how these terms modify equilibrium locations, stability properties, and phase-space topology.

Such analysis provides insight into the sensitivity of the identified dynamical structures to deviations from idealized gravitational symmetry and informs the level of model fidelity required for accurate mission design.

10.5.5 Neptune Obliquity and Non-Coplanar Geometry

The present formulation assumes a coplanar configuration in which the rotating frame and the orbital plane of Triton are aligned. In reality, Neptune’s axial obliquity introduces a misalignment between its equatorial plane, Triton’s orbital plane, and the Sun–Neptune plane. Incorporating this obliquity introduces a fundamentally three-dimensional and potentially time-varying geometric configuration.

This extension enables the investigation of how axial tilt influences long-term trajectory behavior, modifies invariant structures, and affects the stability of both periodic and quasi-periodic motion. It also provides a more physically accurate representation of the system, particularly for missions with long operational durations.

10.5.6 Rotating Frame Formulation in Misaligned Systems

Closely related to the inclusion of obliquity is the question of how the rotating reference frame should be defined in non-coplanar configurations. The classical CR3BP assumes a single well-defined orbital plane about which the frame rotates; however, in the presence of axial tilt and external perturbations, this assumption may no longer be strictly valid.

A reformulation of the rotating frame for misaligned systems is therefore warranted, with the goal of determining how frame definition influences the interpretation of bounded motion, invariant structures, and transfer pathways. This investigation addresses both a conceptual and practical aspect of dynamical modeling, ensuring that the coordinate system used for analysis remains consistent with the underlying physical geometry.

Taken together, these extended investigations provide a roadmap for advancing the present work beyond its current scope. Each direction builds upon the dynamical structures identified in this thesis, offering the potential to deepen understanding, improve model fidelity, and enhance the applicability of the results to complex mission design scenarios.

10.6 Alternative Mission Architectures

The mission architecture developed in Chapter 9 represents a structured and dynamically consistent approach to exploring the Neptune–Triton system. However, it is not unique. The phase-space structure identified throughout this work admits multiple viable interpretations, and alternative mission concepts can be constructed by emphasizing different dynamical regions or operational priorities. Exploring these alternatives provides additional context for the selected architecture and highlights the flexibility inherent in the system.

One class of alternative architectures focuses on short-duration, targeted missions that prioritize specific scientific objectives over extended system exploration. In this context, trajectories may be designed to rapidly transition from Neptune-bounded motion to Triton-bounded encounters, minimizing time spent in intermediate touring regions. Such architectures may rely more heavily on transfer trajectories and may accept higher instability in exchange for reduced mission duration or simplified operational planning.

Another approach emphasizes repeated flyby sequences, in which a spacecraft remains within a touring regime and exploits natural dynamical pathways to achieve multiple encounters with Triton or other bodies within the system. In this case, the objective is not to achieve long-term bounded motion around a single body, but rather to maximize observational opportunities through carefully structured traversal of the phase space. The feasibility of such architectures depends strongly on the persistence of quasi-periodic regions and the accessibility of transport pathways identified in the present analysis.

A third class of architectures involves staging the mission into distinct phases that are dynamically decoupled. For example, a mission may consist of an initial Neptune-centered exploration phase followed by a dedicated Triton phase, with a transfer segment connecting the two. In contrast to the continuous tour architecture presented in Chapter 9, this approach treats each phase as a separate dynamical problem, potentially simplifying design and operational constraints at the expense of reduced dynamical continuity.

It is also possible to consider architectures that exploit regions of phase space that were not selected in the primary analysis. Certain trajectory families identified in the Poincaré maps

and frequency analysis were not pursued due to unfavorable stability characteristics or limited accessibility within the constraints of the present framework. However, under different mission priorities or with the inclusion of additional control strategies, these trajectories may become viable candidates. For example, trajectories exhibiting mild instability could be stabilized through active control, enabling access to regions of phase space that offer unique observational geometries or transfer opportunities.

These alternative architectures underscore the fact that the Neptune–Triton system does not admit a single optimal solution, but rather a spectrum of dynamically feasible mission concepts. The architecture developed in this thesis represents one coherent realization of the underlying dynamical structure, selected based on stability, interpretability, and continuity. The alternatives presented here demonstrate how different priorities and constraints can lead to distinct, yet equally valid, mission designs. This is an important conclusion in its own right, since it suggests that the main value of the present framework may be less in producing a single definitive tour than in organizing the space of plausible tours.

The ability to generate and evaluate such alternatives is a direct consequence of the phase-space framework established in this work. By organizing the system into identifiable regions of motion and understanding the connections between them, it becomes possible to systematically explore a wide range of mission concepts, each grounded in the natural dynamics of the system.

10.7 Closing Perspective

The analysis presented in this thesis demonstrates that the Neptune–Triton system, when viewed through the lens of dynamical systems theory, exhibits a high degree of intrinsic organization that can be directly leveraged for mission design. The structures identified—including periodic orbits, quasi-periodic invariant regions, resonance-supported families, and transport pathways—are not isolated phenomena, but rather components of a coherent phase-space framework that governs motion throughout the system.

By systematically developing this framework, the present work establishes a clear progression from fundamental dynamical modeling to applied mission architecture. The use of Poincaré maps and frequency analysis provides a means of visualizing and classifying the phase space, while stability analysis offers a mechanism for identifying robust trajectory candidates. These tools, when combined, enable the construction of mission concepts that align with the natural dynamics of the system, reducing the reliance on continuous control and enabling more efficient trajectory design.

The Neptune–Triton tour architecture developed in this thesis represents one such application. It demonstrates that it is possible to define a structured sequence of motion that transitions between Neptune-bounded, touring, transfer, and Triton-bounded regimes in a manner that is both dynamically consistent and conceptually coherent. This result highlights the practical value of dynamical systems analysis in informing mission design, particularly in complex multi-body environments. At the same time, the thesis also shows that the final architectural interpretation emerges only after geometric structure, resonance classification, local stability, and full-trajectory behavior are considered together.

At the same time, the work presented here is not intended to provide a complete mission solution. Rather, it establishes a foundational layer upon which more detailed and higher-fidelity analyses can be built. The limitations of the current model, along with the extensions and investigations outlined in this chapter, define a clear path forward for advancing the study of the Neptune–Triton system and translating its dynamical structure into operational capability.

In this sense, the primary contribution of this thesis is not limited to the specific trajectories or architectures identified, but lies in the methodology used to reveal and interpret the underlying structure of the system. The approach developed here is broadly applicable to other multi-body environments and provides a systematic framework for exploring and exploiting complex dynamical systems in the context of space mission design.

The Neptune–Triton system serves as a compelling example of how order emerges from complexity. By identifying and understanding that order, it becomes possible to move beyond ad hoc trajectory design and toward a more principled approach in which mission concepts are derived from, and guided by, the fundamental dynamics of the system itself. That broader lesson extends beyond the specific planetary system studied here: careful dynamical interpretation is not merely a tool for analyzing candidate trajectories after they are found, but a framework for discovering which classes of trajectories are worth pursuing in the first place.

A. Units and Derivations

This appendix summarizes the nondimensionalization used throughout the thesis and provides a dimensional reference table for the Neptune–Triton system. The governing derivations were introduced in Chapter 2 and applied to the Neptune–Triton model in Chapter 3; the purpose of the present appendix is not to repeat those chapters in full, but to collect the most useful scaling relations and dimensional conversions in one place for reference. This is particularly useful because all trajectory products, Poincaré maps, resonance plots, and stability calculations in the thesis are expressed in canonical CR3BP units rather than in dimensional mission-design units.

A.1 Canonical Scaling Used in the Thesis

The Neptune–Triton analysis adopts the standard CR3BP normalization. Let the primary and secondary masses be m_1 and m_2 , and let a denote the Neptune–Triton separation used as the characteristic length. The canonical length, time, velocity, and energy-per-unit-mass scales are then

$$L^* = a, \quad T^* = \sqrt{\frac{(L^*)^3}{G(m_1 + m_2)}}, \quad V^* = \frac{L^*}{T^*}, \quad E^* = (V^*)^2,$$

where G is the gravitational constant. In this normalization, the dimensionless mass parameter is

$$\mu = \frac{m_2}{m_1 + m_2}.$$

The rotating-frame CR3BP is then written in terms of nondimensional variables. Dimensional quantities may be recovered from the canonical values through

$$x_{\text{dim}} = L^* x_{\text{nd}}, \quad t_{\text{dim}} = T^* t_{\text{nd}}, \quad v_{\text{dim}} = V^* v_{\text{nd}}, \quad \mathcal{E}_{\text{dim}} = E^* \mathcal{E}_{\text{nd}}.$$

The corresponding inverse relations are obtained by dividing dimensional quantities by the appropriate characteristic scale.

One useful consequence of this normalization is that the dimensionless mean motion of the primaries is unity. The physical orbital period of the secondary is therefore not T^* , but

$$P = 2\pi T^*.$$

This distinction matters when converting nondimensional integration times back into familiar physical units.

A.2 Neptune–Triton Reference Values

The Neptune–Triton scaling used in this thesis was computed directly from the project constants and nondimensionalization utilities in the ORBIT codebase. The primary values used for the normalization are summarized in Table A.1. These are the dimensional system values from which the canonical scales were constructed.

Table A.1. Dimensional Neptune–Triton reference quantities used to construct the canonical CR3BP normalization in this thesis.

Quantity	Symbol	Value
Neptune mass	m_1	1.02413×10^{26} kg
Triton mass	m_2	2.14×10^{22} kg
Neptune mean radius	R_1	24,764.0 km
Triton mean radius	R_2	1,353.4 km
Neptune–Triton separation	a	354,759.0 km

Substituting these values into the canonical scaling relations above yields the Neptune–Triton nondimensionalization summarized in Table A.2. These are the actual scales underlying the normalized quantities used throughout the thesis.

Table A.2. Canonical Neptune–Triton scales and normalized system parameters used in the thesis.

Derived quantity	Symbol	Value
Characteristic length	L^*	354,759.0 km
Characteristic time	T^*	80,811.7213 s
Characteristic time	T^*	0.935321 days
Secondary orbital period	$2\pi T^*$	5.876794 days
Characteristic velocity	V^*	4.389945 km/s
Characteristic energy per unit mass	E^*	$19.271616 \text{ km}^2/\text{s}^2$
Mass parameter	μ	2.09898×10^{-4}
Normalized Neptune radius	$\bar{R}_1 = R_1/L^*$	0.0698051
Normalized Triton radius	$\bar{R}_2 = R_2/L^*$	0.0038150

Several points are worth noting. First, the characteristic time T^* is a dynamical scaling time, not the full orbital period of Triton. The physical orbital period is recovered through the factor of 2π , as shown above. Second, the normalized Neptune radius is not small compared with the unit separation, which helps explain why close approaches to the primary are still geometrically significant in the canonical plots even though the primaries themselves are often shown schematically. Third, the small mass parameter confirms that the Neptune–Triton CR3BP is a strongly asymmetric system, with Triton acting as a comparatively low-mass secondary embedded in the gravitational field of a dominant primary.

A.3 Practical Conversion Table

For practical use, it is often helpful to convert the most common nondimensional quantities directly into dimensional values without reworking the full normalization each time. Table A.3 provides that reference. These conversions are especially useful when interpreting trajectory plots, integration times, and orbit sizes in Chapter 9.

Table A.3. Practical conversion factors between nondimensional CR3BP units and dimensional Neptune–Triton values.

Nondimensional quantity	Dimensional equivalent	Interpretation
1 length unit	354,759.0 km	Neptune–Triton separation
0.1 length units	35,475.9 km	Moderate regional scale near a primary
0.01 length units	3,547.59 km	Localized orbit-scale distance
1 time unit	80,811.7213 s	Canonical dynamical time
1 time unit	22.4477 hr	Canonical dynamical time
2π time units	5.876794 days	Triton orbital period
1 velocity unit	4.389945 km/s	Canonical speed scale
0.1 velocity units	0.438995 km/s	Representative local speed scale
1 energy unit	19.271616 km ² /s ²	Specific energy scale

The most common conversions used in practice are therefore straightforward. A trajectory integrated for $t = 100$ nondimensional units corresponds to approximately 93.53 days of physical time, while a structure spanning 0.05 nondimensional length units corresponds to approximately 17,738 km. Likewise, a characteristic speed of 0.2 in the normalized equations corresponds to approximately 0.878 km/s. These simple scale conversions allow the canonical plots and trajectory families discussed in the main body of the thesis to be translated back into familiar mission-design magnitudes.

A.4 Interpretive Notes

The use of canonical units is one of the main reasons the same CR3BP framework can be transferred from one system to another with only a change in the underlying physical constants. Once the Neptune–Triton masses and separation are specified, the resulting unit system fixes the interpretation of all nondimensional trajectory products. This is useful mathematically because it preserves the compact form of the equations of motion, but it is also useful practically because it provides a clean separation between the invariant dynamical structure and the later task of dimensional mission interpretation.

For this reason, the dimensional table given here should be read as a reference layer rather than as a replacement for the canonical analysis itself. Chapters 5 through 9 rely on nondimensional variables because they are the natural language of the CR3BP and its associated phase-space structures. The present appendix simply provides the conversion bridge needed to interpret those results more directly in Neptune–Triton mission-design terms.

B. Data Products and Output Schema

This appendix documents the project-level data organization used by the ORBIT workflow for the Neptune–Triton study. Its purpose is practical rather than software-engineering oriented. The ORBIT codebase was developed as a research tool to support the analysis performed in this thesis, not as a polished public software package, and the present appendix should therefore be read as a guide to the data products actually used in the thesis rather than as a formal application programming interface or a frozen external specification.

For that reason, the appendix emphasizes file and folder organization, the purpose of the principal exported products, and the meaning of the most important output columns. It does not attempt to document every internal helper script or every intermediate file generated during development. The goal is reproducibility at the project-data level: a future reader should be able to understand where the main Neptune–Triton outputs were stored, how the major files relate to one another, and which exported tables were used to construct the figures and results discussed in the main body of the thesis.

B.1 Top-Level Project Organization

The ORBIT workflow used in this thesis was maintained in the project tree rooted at

`/Volumes/neptune_triton_system/THESIS`

Within that tree, the most important top-level folders are the following:

Table B.1. Principal top-level folders in the ORBIT project tree used for the Neptune–Triton thesis workflow.

Folder	Purpose
<code>physics/</code>	Governing equations, constants, nondimensionalization, zero-velocity computations, and related physical-model utilities.
<code>data_generators/</code>	Batch-generation scripts for trajectory propagation, Poincaré data, frequency products, resonance tables, and stability inputs.
<code>stability/</code>	Candidate-selection and periodic-orbit stability-analysis utilities used in Chapter 8.
<code>plots/</code>	Plotting and export routines used to generate visual products from the stored datasets.
<code>gui/</code>	The ORBIT graphical interface and its associated configuration and display helpers.
<code>data/</code>	Project output tree containing the saved Neptune–Triton datasets, metadata files, and run catalogs used throughout the thesis.
<code>config/</code>	Shared configuration files, including resonance configuration and metadata-handling utilities.
<code>core/</code>	Shared file-handling and plotting support utilities used across the workflow.

Not all of these folders are equally important to a thesis reader. For the purpose of reproducibility, the `data/` tree is the most important because it contains the actual exported data products used to create the Poincaré, frequency, resonance, stability, and trajectory results reported in the thesis.

B.2 Neptune–Triton Data Tree

The main project output directory for this thesis is

`/Volumes/neptune_triton_system/THESIS/data/Neptune_Triton`

This folder contains the master metadata for the Neptune–Triton study together with energy-specific output directories. At the project level, the most important files are:

Table B.2. Primary project-level files in the Neptune–Triton data directory.

File	Purpose
<code>project.json</code>	Stores the system definition and top-level numerical settings used for the project.
<code>physics_metadata.json</code>	Stores key physical parameters for the Neptune–Triton model, including μ , J_2 , and normalized body radii.
<code>run_catalog.csv</code>	Master catalog of dataset-level runs, including dataset name, Jacobi constant, coordinate range, selected J_2 mode, and file locations for key outputs.
<code>CJ_<value>_J2_<mode>/</code>	Energy-specific dataset folder containing per- C_J outputs for one Jacobi constant and one gravitational-model setting.

The energy-specific naming convention is one of the most important organizational rules in the ORBIT output tree. A folder name of the form

`CJ_3.0186413780_J2_primary`

indicates that the dataset was generated at Jacobi constant $C_J = 3.0186413780$ with the primary-body J_2 term enabled. This convention is used extensively throughout the Neptune–Triton project and provides the main bridge between the physical model selection and the saved output products.

B.3 Per-Dataset Folder Structure

Each Jacobi-level dataset directory stores both metadata and the major derived products associated with that run. A representative dataset folder contains the following principal items:

Table B.3. Representative contents of a Jacobi-level dataset folder.

Item	Purpose
<code>run_metadata.json</code>	Dataset-level metadata, including input settings, geometric reference values, algorithm settings, and per-orbit summaries.
<code>postprocess_summary.json</code>	Postprocessing summary reporting counts of expected, present, or missing files after dataset generation.
<code>poincare/</code>	Surface-of-section outputs and trajectory time-series files for the sampled x_0 values.
<code>frequency/</code>	Base frequency products associated with the dataset prior to the more detailed resonance-run stage.
<code>frequency_runs/</code>	One or more downstream frequency-analysis runs, each with its own metadata and derived products.

Within the `poincare/` folder, two file types appear repeatedly:

- `x0_<value>_<mode>_hybrid_timeseries.dat.gz`, which stores the propagated time-series history for one sampled seed;
- `sos_x0_<value>_<mode>_hybrid.dat.gz`, which stores the associated surface-of-section crossings used in the Poincaré analysis.

These filenames encode the sampled initial condition x_0 , the direction or mode of motion, and the fact that the stored outputs belong to the hybrid event-detection and time-series pipeline used in the project. A future reader does not need to know the internal implementation details of that pipeline to use the files, but the naming convention is useful because it preserves the direct connection between a specific seed and its corresponding trajectory products.

B.4 Frequency and Resonance Run Outputs

The downstream spectral analysis products are organized under the `frequency_runs/` subdirectory of each dataset. A representative frequency-analysis run folder contains:

Table B.4. Principal files produced in a representative frequency-analysis run.

File	Purpose
<code>analysis_metadata.json</code>	Metadata for the frequency-analysis run, including thresholds, tolerances, row counts, and source dataset information.
<code>frequency_data.csv</code>	Main per-trajectory frequency extraction table for the run.
<code>ratio_table.csv</code>	Ratio-focused export used for resonance interpretation and plotting.
<code>resonance_assignment.csv</code>	Row-level resonance assignments produced by the tolerance-based matching stage.
<code>resonance_families.csv</code>	Family-level summary table reporting the identified resonance families and their extent in x_0 .
<code>resonance_candidates.csv</code>	Candidate resonance list used as an intermediate organizational product.
<code>stability_targets.csv</code>	Filtered list of trajectories promoted from the resonance stage to later candidate/stability consideration.
<code>full_run.csv</code>	Broader merged export retaining much of the row-level information from the frequency-analysis stage.

The naming convention for these run folders is timestamp-based. A folder such as

`freqrun_20260327_205257`

indicates one specific downstream frequency-analysis pass applied to a previously generated Jacobi-level dataset. The accompanying `run_catalog.csv` inside `frequency_runs/` provides a compact index of those run folders and their associated configuration choices.

B.5 Key Exported Tables

Although the ORBIT project produces many files, only a small subset is central to the reproducibility of the thesis narrative. The most important exported tables fall into three groups:

1. project-level catalog files;
2. frequency and resonance analysis tables;
3. stability-analysis result tables.

The most important columns of the project-level `run_catalog.csv` are summarized in Table B.5.

Table B.5. Primary columns in the Neptune–Triton `run_catalog.csv` file.

Column	Meaning
<code>dataset_name</code>	Human-readable dataset label.
<code>dataset_dir</code>	Absolute path to the Jacobi-level dataset directory.
<code>CJ</code>	Jacobi constant for the dataset.
<code>mu</code>	Mass parameter used in the run.
<code>x0_min</code> , <code>x0_max</code> , <code>x0_step</code>	Sampled initial-condition range and step size.
<code>J2_mode</code>	Selected J_2 configuration for the run.
<code>poincare_dir</code>	Path to the Poincaré output folder.
<code>index_csv</code> , <code>frequency_csv</code>	Pointers to key derived tables used later in the workflow.
<code>row_count</code>	Number of entries associated with the dataset.
<code>created_utc</code>	Dataset creation timestamp.

The most important columns in the per-trajectory frequency and resonance exports are summarized in Table B.6. The full files contain more fields than are shown here, but the listed columns are the ones most relevant to the thesis workflow.

Table B.6. Key columns used in the frequency and resonance export tables.

Column	Meaning
x0, x0_seed	Sampled initial-condition coordinate and the seed value retained for later reference.
mode, sense, dir_init	Labels describing the direction or orbit-sense convention used for the trajectory.
CJ, C	Jacobi constant fields retained in the export products.
omega1, omega2	Legacy stored frequency fields; numerically interpreted as cycle-frequency quantities in this thesis (see Section 3.8).
ratio	Preferred frequency ratio f_2/f_1 (cycle frequencies) used in the resonance analysis.
ratio_method	Indicates how the preferred ratio was chosen from the spectral data.
valid_fft, fft_reason	Whether the frequency extraction was treated as usable and the associated reason code.
crossings_recorded	Number of recorded crossings available for that trajectory.
collision_flag, escape_flag	Indicators of collision or escape behavior during propagation.
termination_reason	Text label for how the trajectory integration terminated.
is_resonant, p, q, resonance_label	Output of the tolerance-based resonance classification stage.
resonance_error, resonance_family	Numerical closeness to the assigned rational ratio and the grouped family label.

At the stability stage, the principal thesis-facing export is the `stability_results.csv` style table used in Chapter 8. The most important columns in that output are summarized in Table B.7.

Table B.7. Key columns used in the stability-analysis export tables.

Column	Meaning
x0, vy0, CJ	Initial-condition values used to identify the candidate trajectory.
mode, j2_mode	Orbit-sense and selected gravitational-model setting.
period, corrected_period	Initial and corrected period values used in the stability workflow.
x0_corrected, y0_corrected, vx0_corrected, vy0_corrected	Corrected periodic-orbit initial state.
floquet1– floquet4	Floquet multipliers of the corrected periodic representative.
largest_floquet_ multiplier	Largest multiplier magnitude used for summary interpretation.
stability_index	Scalar stability summary derived from the nontrivial multiplier pair.
classification	Thesis-facing stability label such as elliptic stable or invalid orbit.
correction_converged, correction_iterations	Whether the differential corrector converged and how many iterations were required.
closure_error, final_residual_norm	Residual measures used to assess periodic-orbit validity.

B.6 Naming Conventions and Practical Caveats

The ORBIT project uses descriptive filenames and directory names rather than a tightly standardized relational schema. This is helpful during exploratory research because file purposes are often recognizable directly from the filename, but it also means that the output tree should be interpreted with some care. Several practical points are therefore worth stating explicitly.

First, Jacobi-level datasets are named consistently enough to be reliable reference points. Folder names of the form `CJ_<value>_J2_<mode>` are the backbone of the project organization and are the most important naming rule in the output tree. Second, timestamped downstream folders such as `freqrun_<timestamp>` identify a specific postprocessing pass applied to an existing dataset. Third, many row-level filenames encode the sampled x_0 value and the orbit sense directly in the filename, which makes it possible to trace a figure or selected candidate back to the corresponding saved trajectory products.

At the same time, the file structure should not be interpreted as a fully curated public release. Some exported tables preserve internal working fields that are more detailed than a typical end-user would need. Some naming conventions evolved during development, and therefore certain files contain overlapping fields retained for convenience rather than for elegance. In addition, the project tree contains internal caches, helper files, and development artifacts that are useful to the workflow but not essential to the interpretation of the thesis results.

For these reasons, Appendix B should be read as a practical guide to the data products that support the thesis rather than as a promise that every internal file or field is part of a stable long-term interface. The important point for reproducibility is that the major dataset folders, metadata files, run catalogs, resonance tables, and stability tables are organized consistently enough to allow a future reader to reconstruct how the principal results were generated and how the candidate trajectories discussed in the thesis were traced through the ORBIT workflow.

C. Algorithms

This appendix provides a compact reference for the principal numerical workflows used throughout the thesis. Its purpose is not to reproduce the full implementation details of ORBIT or to restate the theoretical development already presented in Chapters 2 through 4. Instead, it collects the core algorithmic logic in one place so that a reader can see how the major trajectory products, spectral outputs, and stability quantities were generated at the workflow level.

The descriptions below are intentionally practical. Each section identifies the purpose of the algorithm, the main inputs and outputs, and the sequence of computational steps used in the thesis. In this way, the appendix serves as a reproducibility aid for future students or readers who wish to understand how the principal dynamical products were assembled without treating the thesis as a software manual.

C.1 Poincaré Map Construction Workflow

The Poincaré map is one of the primary geometric tools used throughout the thesis. In ORBIT, it is constructed by propagating a family of trajectories and recording the state each time a chosen surface-of-section event is satisfied. For the Neptune–Triton applications considered here, the section was defined by

$$y = 0, \quad \dot{y} > 0,$$

so that only crossings with the selected orientation were retained. The resulting map provides a reduced representation of the phase-space structure in the $(x, \dot{x})$ plane.

Inputs:

- system definition and active dynamical model,
- Jacobi constant or energy-selection rule,
- sampled initial-condition line or seed set,
- integration time horizon and numerical tolerances,
- surface-of-section event definition.

Outputs:

- stored crossing states for each trajectory,
- assembled Poincaré point cloud in $(x, \dot{x})$,
- trajectory-level metadata used later in the frequency and stability workflows.

The algorithm used in the thesis may be summarized as follows.

1. Define the physical system, active model, and Jacobi level to be studied.
2. Generate the initial-condition set along the prescribed sampling line, typically by varying x_0 and solving the corresponding velocity constraint from the Jacobi integral.

3. Propagate each trajectory forward in time with the selected numerical integrator.
4. Monitor the event condition $y = 0$ during propagation and detect sign changes indicating a section crossing.
5. Refine the crossing time and state so that the section is located accurately rather than only at the nearest output step.
6. Accept the crossing only if the directional condition $\dot{y} > 0$ is satisfied.
7. Record the corresponding $(x, \dot{x})$ values and any associated trajectory metadata.
8. Repeat the procedure until the propagation terminates or the required number of crossings has been collected.
9. Aggregate all accepted crossings into the final Poincaré map for the selected energy level.

Several practical features of the workflow are worth noting. First, the event-detection and crossing-refinement stages are essential because the plotted points should represent true section intersections rather than coarse output samples. Second, the directional filter prevents the map from mixing the two passage directions across the same section. Third, the trajectory metadata produced at this stage are preserved because the same trajectories are later revisited during the frequency-ratio and resonance workflows of Chapter 6.

C.2 Frequency-Analysis Workflow

The frequency-analysis procedure described in Chapter 6 was applied to the same trajectory families first organized geometrically in Chapter 5. The purpose of the workflow is to convert stored time histories into a compact spectral description that can be used to identify broad frequency bands, approximate commensurabilities, and candidate resonance families.

Inputs:

- stored trajectory time histories,
- analysis window length and sampling interval,
- minimum usable trajectory duration,
- ratio tolerance and resonance-library settings.

Outputs:

- dominant frequency estimates for each trajectory,
- ratio products such as ω_2/ω_1 ,
- frequency-ratio plots,
- approximate resonance labels and family assignments.

At the workflow level, the procedure used in the thesis is:

1. Retrieve the stored trajectory time histories produced during the propagation stage.
2. Discard trajectories whose usable duration is too short for reliable spectral interpretation.
3. Extract a uniformly sampled analysis segment from each retained trajectory.
4. Compute the trajectory spectrum and identify the dominant frequency peak.
5. Inspect the remaining spectral structure to identify the second relevant frequency contribution or sideband spacing used in the ratio calculation.
6. Form the ratio ω_2/ω_1 for the trajectory.
7. Store the extracted frequency quantities and the associated initial-condition metadata.
8. Compare the computed ratio against a library of low-order rational targets within a selected tolerance.
9. Assign an approximate resonance label when a satisfactory match is found; otherwise leave the trajectory unclassified.
10. Group trajectories with the same assigned resonance into common resonance families for later plotting and interpretation.

This workflow is intentionally conservative in interpretation. The extracted ratios are numerical products derived from finite time windows, finite spectral resolution, and practical matching tolerances. For that reason, the resonance labels generated by ORBIT are used in the thesis as organizing devices rather than as stand-alone proofs of exact resonance. Their main value is that they provide a systematic bridge between the trajectory-level spectra, the global frequency-ratio plots, and the resonance overlays later projected back onto the Poincaré maps.

C.3 Differential Correction and Stability Workflow

The stability analysis of Chapter 8 required representative periodic-orbit solutions rather than only uncorrected seeds selected from the Poincaré maps. To obtain these, ORBIT applied a differential-correction procedure to candidate seeds and then evaluated the resulting periodic orbit through the associated state-transition dynamics.

Inputs:

- selected seed state on the surface of section,
- target symmetry or periodicity condition,
- correction tolerance and iteration limit,
- variational-equation model for the state-transition matrix.

Outputs:

- corrected periodic-orbit initial condition,
- corrected period,
- monodromy matrix,
- Floquet multipliers and derived scalar stability index,
- convergence or invalid-orbit classification.

The numerical workflow is summarized below.

1. Select a seed from the candidate family identified in Chapter 7.
2. Propagate the seed to the next symmetry point, return point, or closure condition appropriate to the chosen periodicity formulation.
3. Evaluate the residual measuring failure of the periodic-orbit condition.
4. Simultaneously propagate the variational equations to obtain the local sensitivity of that residual to changes in the initial condition.
5. Build a linearized correction step from the residual and the relevant state-transition information.
6. Update the seed and repeat the propagation-and-correction loop until the residual falls below the selected tolerance or the iteration limit is reached.
7. If the corrector converges, propagate the corrected periodic orbit for one full period while integrating the state-transition matrix.
8. Form the monodromy matrix from the propagated state-transition matrix over one period.
9. Compute the Floquet multipliers and derive the stability measures reported in Chapter 8.
10. If convergence fails, the orbit does not close adequately, or the resulting periodic representative is not reliable enough for stability classification, record the case as `Invalid orbit`.

This last point is important for interpreting the Chapter 8 tables. In the present thesis, an `Invalid orbit` classification does not mean that the original seed is physically meaningless or that the surrounding phase-space structure does not exist. It means only that the adopted periodic-orbit validation workflow did not produce a defensible corrected periodic representative from that seed under the selected tolerances and iteration limits. The invalid classification is therefore a numerical-stability result, not a denial of the broader geometric structure identified earlier in the thesis.

C.4 Brent Root-Finding in the ORBIT Workflow

Several parts of ORBIT require the solution of one-dimensional nonlinear equations for which a robust bracketing method is more appropriate than an unconstrained Newton iteration. In those cases, the framework uses Brent’s method. This algorithm combines bisection, the secant method, and inverse quadratic interpolation, providing the reliability of a bracketing scheme together with faster local convergence when the function behaves smoothly.

Within the thesis workflow, Brent’s method was used in two especially important contexts. The first is event and crossing refinement, where a bracketed sign change indicates that a trajectory has crossed the chosen surface of section. The second is equilibrium and auxiliary scalar calculations, such as solving for collinear Lagrange-point locations or other one-dimensional closure conditions that arise naturally in the CR3BP analysis.

Inputs:

- scalar function $f(s)$,
- lower and upper bracket values s_a and s_b such that $f(s_a)f(s_b) < 0$,
- stopping tolerance,
- maximum iteration count.

Outputs:

- bracketed root estimate,
- convergence flag and solver status.

The practical Brent workflow used by the framework is:

1. Identify a scalar bracket over which the target function changes sign.
2. Initialize the algorithm with the two bracket endpoints and their function values.
3. Attempt an interpolation-based step when the recent iterates indicate that a higher-order estimate is safe and useful.
4. Fall back to a secant or bisection step when the interpolation estimate is not trustworthy or would leave the bracket.
5. Replace the bracket endpoint whose sign matches the new trial point, thereby preserving the sign-change condition.
6. Continue until the bracket width or function value falls below the requested tolerance, or until the maximum iteration count is reached.

The advantage of Brent’s method in this context is not merely speed, but reliability. Event localization and equilibrium-point calculations are foundational to the later phase-space and stability products, so a method that preserves a valid bracket whenever possible is preferable to one that converges quickly only when the initial guess is already close. For this reason, Brent’s method is a natural numerical component of the ORBIT workflow wherever a one-dimensional root must be located accurately and robustly.

C.5 Appendix Summary

The four workflows summarized here correspond to the principal computational stages that support the thesis results: geometric phase-space construction, spectral classification, periodic-orbit correction and stability evaluation, and robust scalar root-finding. Together they provide the algorithmic bridge between the theoretical development of Chapters 2 and 3 and the applied Neptune–Triton results of Chapters 5 through 9. The present appendix is therefore intended as a compact computational roadmap: it does not replace the detailed scientific discussion in the main text, but it identifies the principal numerical procedures through which those results were produced.

Bibliography

- [1] “Voyager 2 - NASA Science,” Dec. 2017.
- [2] “30 Years Ago: Voyager 2’s Historic Neptune Flyby - NASA,” Aug. 2019.
- [3] “Neptune: Exploration - NASA Science,” Nov. 2017.
- [4] “30 Years Ago: Voyager 2 Explores Neptune - NASA,” Aug. 2019.
- [5] A. M. Rymer, K. D. Runyon, B. Clyde, J. I. Núñez, R. Nikoukar, K. M. Soderlund, K. Sayanagi, M. Hofstadter, L. C. Quick, S. A. Stern, T. Becker, M. Hedman, I. Cohen, F. Cray, J. J. Fortney, J. Vertesi, C. Hansen, I. De Pater, C. Paty, T. Spilker, T. Stallard, G. B. Hospodarsky, H. T. Smith, H. Wakeford, S. E. Moran, A. Annex, P. Schenk, M. Ozimek, J. Arrieta, R. L. McNutt, A. Masters, A. A. Simon, S. Ensor, C. T. Aplan, J. Bruzzi, D. A. Patthoff, C. Scott, C. Campo, C. Krupiarz, C. J. Cochrane, C. Gantz, D. Rodriguez, D. Gallagher, D. Hurley, D. Crowley, E. Abel, E. Provornikova, E. P. Turtle, G. Clark, J. Wilkes, J. Hunt, J. H. Roberts, J. Rehm, K. Murray, L. Wolfarth, L. N. Fletcher, L. Spilker, E. S. Martin, M. Parisi, M. Norkus, N. Izenberg, R. Stough, R. J. Vervack, K. Mandt, K. B. Stevenson, S. Kijewski, W. Cheng, J. D. Feldman, G. Allen, D. Prabhu, S. Dutta, C. Young, and J. Williams, “Neptune Odyssey: A Flagship Concept for the Exploration of the Neptune–Triton System,” *The Planetary Science Journal*, vol. 2, p. 184, Oct. 2021.
- [6] “Tritons Dark Plume.” <https://www.jpl.nasa.gov/images/pia14448-tritons-dark-plume/>.
- [7] Committee on the Planetary Science and Astrobiology Decadal Survey, Space Studies Board, Division on Engineering and Physical Sciences, and National Academies of Sciences, Engineering, and Medicine, *Origins, Worlds, and Life: A Decadal Strategy for Planetary Science and Astrobiology 2023-2032*. Washington, D.C.: National Academies Press, Oct. 2023.
- [8] Committee on the Planetary Science and Astrobiology Decadal Survey, Space Studies Board, Division on Engineering and Physical Sciences, and National Academies of Sciences, Engineering, and Medicine, *Origins, Worlds, and Life: Planetary Science and Astrobiology in the Next Decade*. Washington, D.C.: National Academies Press, Nov. 2023.
- [9] W. S. Koon, M. W. Lo, J. E. Marsden, and S. D. Ross, *Dynamical Systems, the Three-Body Problem and Space Mission Design*, vol. 32 of *Interdisciplinary Applied Mathematics*. New York: Springer, 2011.
- [10] A. J. Lichtenberg, M. A. Lieberman, and A. J. Lichtenberg, *Regular and chaotic dynamics*. No. v. 38 in *Applied mathematical sciences*, New York: Springer-Verlag, 2nd ed ed., 1992.
- [11] S. Wiggins, *Introduction to applied nonlinear dynamical systems and chaos*. No. 2 in *Texts in applied mathematics*, New York Berlin Heidelberg: Springer, 2. ed., [softcover version of original hardcover ed. 2003] ed., 2010.

- [12] C. D. Murray and S. F. Dermott, *Solar system dynamics*. Cambridge ; New York: Cambridge University Press, 1999.
- [13] J. Laskar, “Frequency analysis for multi-dimensional systems. Global dynamics and diffusion,” *Physica D: Nonlinear Phenomena*, vol. 67, pp. 257–281, Aug. 1993.
- [14] G. Gómez, ed., *Dynamics and mission design near libration points*. No. vol. 2-5 in World scientific monograph series in mathematics, Singapore ; River Edge, NJ: World Scientific, 2001.
- [15] M. A. Aguirre, *Introduction to Space Systems: Design and Synthesis*. No. 27 in Space Technology Library, New York, NY: Springer, 2013.
- [16] N. Bosanac, *Leveraging natural dynamical structures to explore multi-body systems*. PhD thesis, Purdue University, United States – Indiana, 2016.
- [17] V. G. Szebehely, *Theory of Orbit*. Place of publication not identified: Academic Press, 1967.
- [18] H. D. Curtis, *Orbital mechanics for engineering students*. Elsevier aerospace engineering series, Oxford [England] ; Cambridge, MA: Butterworth-Heinemann, fourth edition ed., 2020.
- [19] S. H. Strogatz, *Nonlinear dynamics and Chaos: with applications to physics, biology, chemistry, and engineering*. Studies in nonlinearity, Reading, Mass: Addison-Wesley Pub, 1994.
- [20] H. Poincaré, *Les méthodes nouvelles de la mécanique céleste: Méthodes de MM. Newcomb, gylden, lindstedt et bohlin. Tome II*, vol. 2 of *Les méthodes nouvelles de la mécanique céleste*. Gauthier-Villars, 1893.
- [21] N. Bosanac, *Exploring the Influence of a Three-Body Interaction Added to the Gravitational Potential Function in the Circular Restricted Three-Body Problem: A Numerical Frequency Analysis*. PhD thesis, Purdue University, West Lafayette, Indiana, Dec. 2012.
- [22] J. Moser, *On invariant Curves of area-preserving mappings of an annulus*. Nachrichten der akademie der wissenschaften zu göttingen, mathematisch-physikalische klasse, Vandenhoeck & Ruprecht, 1962.
- [23] E. Ott, *Chaos in dynamical systems*. Cambridge: Cambridge University Press, 2. ed ed., 2002.
- [24] J. Laskar, “The chaotic motion of the solar system: A numerical estimate of the size of the chaotic zones,” *Icarus*, vol. 88, pp. 266–291, Dec. 1990.
- [25] G. Beutler, L. Mervart, and A. Verdun, *Methods of celestial mechanics*. Astronomy and astrophysics library, Berlin ; New York: Springer, 2005.
- [26] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0

- Contributors, “SciPy 1.0: Fundamental algorithms for scientific computing in Python,” *Nature Methods*, vol. 17, no. 3, pp. 261–272, 2020.
- [27] J. Dormand and P. Prince, “A family of embedded Runge-Kutta formulae,” *Journal of Computational and Applied Mathematics*, vol. 6, pp. 19–26, Mar. 1980.
- [28] J. Laskar, “Frequency map analysis of an hamiltonian system,” *AIP Conference Proceedings*, vol. 344, pp. 130–159, Sept. 1995.
- [29] “Neptune: Facts - NASA Science,” Nov. 2017.
- [30] “Voyager - NASA Science,” Dec. 2023.
- [31] C. Aydin, “Comet-type periodic motions and their out-of-plane bifurcations in the Earth-Moon CR3BP: a computational symplectic analysis,” Feb. 2026.
- [32] A. Binder and D. Arnas, “Reliable and Repeatable Transit Through Cislunar Space Using 2:1 Resonant Spatial Orbits,” *Journal of Guidance, Control, and Dynamics*, vol. 47, pp. 1973–1979, Sept. 2024.
- [33] N. Bosanac, K. C. Howell, and E. Fischbach, “EXPLORING THE IMPACT OF A THREE-BODY INTERACTION ADDED TO THE GRAVITATIONAL POTENTIAL FUNCTION IN THE RESTRICTED THREE-BODY PROBLEM,” in *AAS/AIAA Space Flight Mechanics Meeting*, (San Diego, California), 2013. AAS 13-490.
- [34] B. W. Carroll and D. A. Ostlie, *An Introduction to Modern Astrophysics*. Cambridge University Press, 2 ed., Sept. 2017.
- [35] B. Fiedler, K. Gröger, and J. Sprekels, *Equadiff 99: (In 2 Volumes)*. World Scientific Publishing Company, Sept. 2000.
- [36] D. Grebow, *GENERATING PERIODIC ORBITS IN THE CIRCULAR RESTRICTED THREE-BODY PROBLEM WITH APPLICATIONS TO LUNAR SOUTH POLE COVERAGE*. Thesis, Purdue University, West Lafayette, Indiana, 2006.
- [37] J. Guckenheimer and P. Holmes, *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer Science & Business Media, Nov. 2013.
- [38] D. D. Nolte, *Introduction to Modern Dynamics: Chaos, Networks, Space, and Time*. Oxford University Press, 2 ed., Sept. 2019.
- [39] G. M. Oliveira, A. F. Prado, D. M. Sanchez, and V. M. Gomes, “Traveling Between the Earth-Moon Lagrangian Points and the Earth,” in *SpaceOps 2016 Conference*, (Daejeon, Korea), American Institute of Aeronautics and Astronautics, May 2016.
- [40] University of Oklahoma Graduate College, “Thesis/dissertation instruction packet.” University of Oklahoma Graduate College, 2025. Formatting guidance document.

- [41] D. M. Sanchez, A. F. Prado, and T. Yokoyama, “On the effects of each term of the geopotential perturbation along the time I: Quasi-circular orbits,” *Advances in Space Research*, vol. 54, pp. 1008–1018, Sept. 2014.
- [42] D. M. Sanchez, R. Deienno, A. F. B. A. Prado, and K. C. Howell, “Perturbation Maps and the ring of Haumea,” *Monthly Notices of the Royal Astronomical Society*, vol. 496, pp. 2085–2097, Aug. 2020.
- [43] L. B. T. Santos, A. K. de Almeida, P. A. Sousa-Silva, M. O. Terra, D. M. Sanchez, S. A. A. F. B. A. Prado, and F. Monteiro, “Optimal transfers from Moon to L_2 halo orbit of the Earth-Moon system,” *Advances in Space Research*, vol. 70, pp. 3362–3372, Dec. 2022.
- [44] T. Yokoyama, E. Vieira Neto, O. C. Winter, D. M. Sanchez, and P. I. d. O. Brasil, “On the Evection Resonance and Its Connection to the Stability of Outer Satellites,” *Mathematical Problems in Engineering*, vol. 2008, no. 1, p. 251978, 2008.
- [45] E. Hairer, S. P. Nørsett, and G. Wanner, *Solving Ordinary Differential Equations I: Nonstiff Problems*, vol. 8 of *Springer Series in Computational Mathematics*. Berlin, Heidelberg: Springer, 2 ed., 1993.
- [46] “Chapter 24. Fourier Transform — Python Numerical Methods.” <https://pythonnumericalmethods.studentorg.berkeley.edu/notebooks/chapter24.00-Fourier-Transforms.html>.
- [47] “DOP853 — SciPy v1.17.0 Manual.” <https://docs.scipy.org/doc/scipy/reference/generated/scipy.integrate.DOP853.html>.