

HYBRID STATION-KEEPING CONTROLLER DESIGN LEVERAGING
FLOQUET MODE AND REINFORCEMENT LEARNING APPROACHES

A Thesis

Submitted to the Faculty

of

Purdue University

by

Andrew B. Molnar

In Partial Fulfillment of the

Requirements for the Degree

of

Master of Science in Aeronautics and Astronautics

December 2020

Purdue University

West Lafayette, Indiana

THE PURDUE UNIVERSITY GRADUATE SCHOOL
STATEMENT OF THESIS APPROVAL

Dr. Kathleen C. Howell, Chair

School of Aeronautics and Astronautics

Dr. Carolin Frueh

School of Aeronautics and Astronautics

Mr. Tim Keeter

Dynetics, Inc.

Approved by:

Dr. Gregory Blaisdell

Associate Head of Graduate Program of Aeronautics and Astronautics

To Mom, Dad, Grandma, & Grandpa

ACKNOWLEDGMENTS

The help that I've received from family, friends, and colleagues has undoubtedly made this endeavor possible. I am forever grateful for the guidance and support you have given me throughout this process. A special thanks is extended to all of the people in my life who have helped me along the way.

I want to thank my mom, whose encouragement has always driven me to better myself, in academics and as a person. It would not have been possible to get to where I am today without your loving support. For as long as I can remember, you have guided me in my pursuit of education. I fondly remember our first trip to Purdue together, the beginning of the journey. I love having you accompany me in my travels, exploring new areas that I never thought I would venture to. I cannot wait to see where we'll go next.

My girlfriend, Maaninee, deserves a special thank you. At every step along the way, you have been there to help me through my struggles. I am so grateful for your constant love and support. Your presence has truly bettered my life. Thank you for always being there; to talk through my research, make me laugh during stressful times, and celebrate all the accomplishments that have been made along the way. I will always be there for you, and hope to make you as happy as you make me.

I also want to thank my advisor, committee members, and research group. Professor Howell, you have always been there to guide me through my research journey. Your suggestions have been invaluable and led me to accomplish more than I ever thought possible. To my committee members, Professor Frueh and Tim Keeter, thank you for taking the time to review my thesis and for providing insightful feedback. To the members of the Multi-Body Dynamics research group, thank you for listening to my research presentations and offering advice that has guided my research progress.

Thank you for providing an environment that fosters innovation and is so welcoming. The memories we have created together will live with me forever.

A final thank you to the organizations that supported my graduate research. Thank you to the College of Engineering Honors for the financial support and unique experiences provided by a teaching assistantship. Finally, thank you to Dynetics for the learning opportunities, and for providing access to their proprietary reinforcement learning toolbox, which greatly accelerated the progress of my research.

TABLE OF CONTENTS

	Page
LIST OF TABLES	viii
LIST OF FIGURES	ix
ABSTRACT	xi
1 INTRODUCTION	1
1.1 Previous Contributions	1
1.1.1 Multi-Body Dynamics	2
1.1.2 Station-Keeping	2
1.1.3 Reinforcement Learning	3
1.2 Thesis Overview	4
2 DYNAMICAL MODEL	6
2.1 The N-Body Problem	6
2.2 The Circular Restricted Three-Body Problem	8
2.2.1 Model Definition	9
2.2.2 Equations of Motion	11
2.2.3 Jacobi Constant	14
2.2.4 Equilibrium Points	15
2.2.5 Zero Velocity Curves	18
2.3 System Parameters	20
3 DYNAMICAL SYSTEMS THEORY	22
3.1 Linear Variational Equations	22
3.2 State Transition Matrix	24
3.3 Differential Corrections	26
3.3.1 Single Shooting	27
3.3.2 Targeting Perpendicular Crossings	29
3.3.3 Natural Parameter Continuation	33
3.3.4 Lyapunov Orbit Families	34
3.4 Stability	38
3.4.1 Stability Index	40
3.5 Bifurcations	40
3.5.1 Types of Bifurcations	41
3.5.2 Halo Families	42
3.6 Invariant Manifolds	45
3.6.1 Invariant Manifold Theory	46

	Page
3.6.2 Manifolds of Equilibrium Points and Periodic Orbits	48
4 FLOQUET MODE STATION-KEEPING	53
4.1 Floquet Theory	53
4.2 Floquet Mode Control Laws	55
4.2.1 Periodic Orbits with a Four-Dimensional Center Subspace	55
4.3 Station-Keeping Simulation	58
4.3.1 Operational Errors	58
4.3.2 Mission Constraints	60
4.4 Station-Keeping Analysis	61
4.4.1 Monte Carlo	62
4.4.2 Performance Metrics	63
4.5 Floquet Mode Control Results	65
4.5.1 Single and Multi-Axis Control of Spatial Orbits	65
4.5.2 Divergence from the Reference Orbit	70
4.5.3 Expanded Operational Errors and Mission Constraints	73
5 HYBRID CONTROLLER DESIGN	78
5.1 Reinforcement Learning Background	78
5.1.1 Introduction to Reinforcement Learning	79
5.1.2 Soft Actor-Critic Algorithm	83
5.2 Reinforcement Learning Framework	86
5.2.1 Observation Space	86
5.2.2 Action Space	88
5.2.3 Reward Definition	90
5.2.4 Hybrid Controller Design	93
5.2.5 Training Process	94
5.3 Hybrid Controller Results	96
5.3.1 Baseline Controller Performance	96
5.3.2 Performance Relative to Floquet Mode Controllers	98
5.4 Comparative Station-Keeping Scenario	101
5.4.1 Trajectory Analysis	103
6 CONCLUSION	107
6.1 Summary	107
6.2 Remarks	108
6.3 Recommendations for Future Work	109
REFERENCES	111
A Parameters for Agent Training Process	114
A.1 Simulation Parameters	114
A.2 Reinforcement Learning Parameters	115

LIST OF TABLES

Table	Page
2.1 Earth-Moon System Parameters.	21
2.2 Sun-EMB System Parameters.	21
3.1 Spatial Periodic Orbit Continuation Algorithm Summary.	44
3.2 Stability of Eigenvalues of Equilibrium Points and Periodic Orbits.	46
4.1 Standard Operational Errors.	60
4.2 Standard Mission Constraints.	61
4.3 Metrics For Classifying Successful Station-Keeping Simulations.	64
4.4 Parameters for the Selected L_1 Halo Orbit.	66
5.1 Observation Space for the Station-Keeping Reinforcement Learning Framework.	88
5.2 Comparison of Expected Baseline Performance Metrics for Floquet Mode and Hybrid Controllers.	99
5.3 Initial Corrective Maneuvers Determined from the Floquet Three-Axis, Hybrid, and Targeting Controllers.	104
A.1 Operational Errors for the Agent Training Process.	114
A.2 Mission Constraints for the Agent Training Process.	114
A.3 Reinforcement Learning Hyperparameters for Agent Training.	115

LIST OF FIGURES

Figure	Page
2.1 N -body System in the Inertial Frame.	7
2.2 Schematic of the Circular Restricted Three-Body Problem.	10
2.3 Equilibrium Points of the CR3BP.	18
2.4 Zero Velocity Curves in the Earth-Moon System.	20
3.1 Perturbed Trajectory Relative to a Reference Trajectory.	24
3.2 Planar Perpendicular Targeting Problem.	30
3.3 L_1 Linear Variational Orbit in the Sun-EMB system.	35
3.4 Corrected L_1 Lyapunov Orbit in the Sun-EMB System.	36
3.5 Evolution of the Sun-EMB L_1 Lyapunov Family.	37
3.6 Sun-EMB L_1 Lyapunov Family.	38
3.7 Types of Bifurcations Resulting from Stability Changes.	42
3.8 Stability Indices of the Sun-EMB L_1 Lyapunov Family.	43
3.9 Axial and Isometric Views of the L_1 Southern Halo Family.	45
3.10 Manifold Structure of the L_1 Libration Point in the Sun-EMB System. . .	50
3.11 Manifold Structure of the L_1 Lyapunov-Halo Bifurcating Orbit.	51
4.1 Projections of the Nominal L_1 Southern Halo Orbit.	67
4.2 Baseline Maintenance Costs for a L_1 Halo Orbit in the Sun-EMB System with x -Axis Control.	68
4.3 Baseline Maintenance Costs for a L_1 Halo Orbit in the Sun-EMB System with Three-Axis Control.	69
4.4 Linear Divergence from a L_1 Halo Orbit in the Sun-EMB System with x -Axis Control.	70
4.5 Position Error of a Spacecraft with x -Axis Control.	71
4.6 Linear Divergence from a L_1 Halo Orbit in the Sun-EMB System with Three-Axis Control.	72

Figure	Page
4.7 Success Rate and Expected Cost of the x -Axis and Three-Axis Controllers for Extended Injection Error.	75
4.8 Three-Axis Controller Performance Relative to the x -Axis Controller. . . .	76
5.1 The Agent-Environment Interaction Process [34], [35].	80
5.2 Action Scaling Function.	90
5.3 Reward Component Signals.	92
5.4 Discounted Returns Per Episode and Moving Average During Training ($\gamma = .99$).	95
5.5 Monte Carlo Results for Baseline Performance Metrics of the Hybrid Controller.	97
5.6 Success Rate and Expected Cost of the Hybrid Controller for Increased Injection Error.	100
5.7 Hybrid Controller Performance Relative to the Floquet Mode Three-Axis Controller.	101
5.8 Trajectories of the Floquet Three-Axis, Hybrid, and Targeting Controllers Prior to the Second Corrective Maneuver.	103
5.9 Complete Spacecraft Trajectories for the Floquet Three-Axis, Hybrid, and Targeting Controllers.	105

ABSTRACT

Molnar, Andrew B. M.S.A.A., Purdue University, December 2020. Hybrid Station-Keeping Controller Design Leveraging Floquet Mode and Reinforcement Learning Approaches. Major Professor: Kathleen C. Howell.

The general station-keeping problem is a focal topic when considering any spacecraft mission application. Recent missions are increasingly requiring complex trajectories to satisfy mission requirements, necessitating the need for accurate station-keeping controllers. An ideal controller reliably corrects for spacecraft state error, minimizes the required propellant, and is computationally efficient. To that end, this investigation assesses the effectiveness of several controller formulations in the circular restricted three-body model. Particularly, a spacecraft is positioned in a L_1 southern halo orbit within the Sun-Earth Moon Barycenter system. To prevent the spacecraft from departing the vicinity of this reference halo orbit, the Floquet mode station-keeping approach is introduced and evaluated. While this control strategy generally succeeds in the station-keeping objective, a breakdown in performance is observed proportional to increases in state error. Therefore, a new hybrid controller is developed which leverages Floquet mode and reinforcement learning. The hybrid controller is observed to efficiently determine corrective maneuvers that consistently recover the reference orbit for all evaluated scenarios. A comparative analysis of the performance metrics of both control strategies is conducted, highlighting differences in the rates of success and the expected propellant costs. The performance comparison demonstrates a relative improvement in the ability of the hybrid controller to meet the mission objectives, and suggests the applicability of reinforcement learning to the station-keeping problem.

1. INTRODUCTION

The human presence in space has expanded from what was once a handful of groundbreaking missions into the dawning of a new age of space exploration. Thousands of organizations and companies, from countries across the globe, are venturing into space-based research and applications. With every passing day, novel missions are proposed that seek to advance the current state of understanding and technology. Each mission presents unique challenges, but any mission that aims to place an object in orbit shares a common requirement: the need for a station-keeping strategy that is reliable, highly accurate, and sustainable. The challenge of station-keeping derives from inherent difficulties associated with modeling the environment that spacecraft operate within. From a theoretical viewpoint, the motion of the spacecraft is governed by a specified dynamical system, e.g., the two-body problem, three-body problem, or a higher fidelity model. However, in actuality, the spacecraft experiences a variety of additional forces such as solar radiation pressure and gravitational perturbations from other celestial bodies. The addition of these perturbing forces results in a natural tendency for the spacecraft to deviate from some desired reference motion. Hence, station-keeping controllers are employed to determine corrective maneuvers that resolve any deviations in the desired position or velocity states.

1.1 Previous Contributions

The formulations and analyses conducted within this investigation build upon existing methodologies and strategies. Without the groundbreaking work of countless researchers, this investigation would not have been possible. Consequently, the history and significance of these contributions are recognized and discussed here, noting three

primary concepts: multi-body dynamics, station-keeping controller formulation, and reinforcement learning.

1.1.1 Multi-Body Dynamics

The study of the dynamics governing celestial motion has a rich history, with contributions from the likes of Johannes Kepler, Isaac Newton, and Leonhard Euler. In the early 1600s, Kepler published the *Astronomia Nova*, deriving mathematical laws to represent the motion of planetary bodies [1]. Later that century, Newton derived gravitational laws to represent the motion of N-bodies in the Law of Universal Gravitation [2]. Despite this advancement, the N-body problem remained unresolved, and was later simplified by Euler into the circular restricted three-body problem (CR3BP) [3]. The conversion of the problem from an inertial frame to a rotating coordinate frame enabled significant advancements in the study of celestial mechanics. Two additional critical developments vastly accelerated insight into this model: the existence of equilibrium solutions proven by Joseph-Louis Lagrange [4], and the discovery of a new integral of motion. Carl Gustav Jacob Jacobi proved that the application of the conservation of energy and angular momentum to the CR3BP determined a new integral of the motion, known as the Jacobi constant [5]. Through this integral, bounding structures that define the accessible regions of a spacecraft are readily determined [6]. As a result of these developments, this model is extensively referenced in literature throughout the 1800s and 1900s, and is widely employed in current research for preliminary analysis.

1.1.2 Station-Keeping

Station-keeping controller formulation is extensively studied for theoretical and practical applications. Consequently, a variety of methods have been formulated which uniquely determine control laws to correct for state error using impulsive maneuvers and continuous low-thrust. Gómez, Llibre, Martínez, and Simó originally

developed a strategy known as Floquet mode, which leverages natural dynamics to determine an impulsive maneuver to correct for state deviation [7]. This technique assesses the eigenstructure, or natural flow, of the current reference state, and determines a corrective maneuver that eliminates the unstable behavior. Additionally, the target point approach, applied to the station-keeping problem by Howell and Pernicka, similarly determines an impulsive corrective maneuver [8]. This approach seeks to minimize a cost function consisting of the corrective maneuver and the deviation in a select number of future target states. Conversely, another branch of station-keeping control focuses on continuous low-thrust techniques. Scheeres demonstrated that the Hamiltonian Structure Preserving (HSP) controller is capable of maintaining reference motion by creating an artificial center manifold [9]. This controller effectively achieves bounded motion about the reference by eliminating the stable and unstable manifolds.

1.1.3 Reinforcement Learning

Reinforcement learning is an emergent subset of machine learning that centers on the concept of learning via trial and error. While this field and associated techniques have only recently been applied to aerospace applications, the fundamental reinforcement learning concept originates from early work in artificial intelligence in the mid-to-late 1900s. Richard Bellman developed an approach to solve optimal control problems using what are now referred to as the Bellman equations [10]. The development of dynamic programming presented a feasible solution to solving stochastic optimal control problems. However, the computational demand grows exponentially as the number of state variables increases, termed by Bellman as “the curse of dimensionality”. Due to limitations in computing resources, research progress was primarily constrained to theoretical analysis. However, with the development of high performance computing, reinforcement learning applications have significantly progressed in recent years.

1.2 Thesis Overview

The primary focus of this investigation centers on station-keeping within the circular restricted three-body problem. More specifically, the traditional Floquet mode controller is introduced and assessed in a controlled analysis to determine performance standards for a L_1 southern halo orbit in the Sun-Earth Moon Barycenter (EMB) system. Thereafter, reinforcement learning is introduced as a means to improve on highlighted failure points associated with the Floquet mode controller.

- **Chapter 2:** The general N-body problem is introduced, demonstrating the complexities and challenges inherent to the problem formulation. Thereafter, simplifying assumptions are applied, which form the basis of the transition from the N-body problem to the circular restricted three-body problem. With the model defined, non-dimensional equations of motion are derived to represent the motion of an infinitesimal spacecraft under the influence of two celestial bodies. To conclude, the Jacobi constant is defined, allowing the introduction of equilibrium solutions and zero velocity curves.
- **Chapter 3:** A variety of dynamical systems techniques are defined that are extensively leveraged throughout the investigation to aid analysis. Initially, the linear variational equations are determined from the linearization of the nonlinear equations of motion. The linear variational equations allow the formulation of the state transition matrix (STM). The significance of the STM is demonstrated with the introduction of a series of differential corrections techniques that are employed to discover families of periodic orbits. Finally, invariant manifold theory is applied to reveal manifold structures that direct flow around periodic orbits.
- **Chapter 4:** Floquet theory is discussed and applied to the station-keeping problem, leading to the derivation of two Floquet mode control laws: x -axis control and three-axis control. Thereafter, the station-keeping problem is defined, with specifications outlined for the operational errors, mission constraints,

conditions for success, and simulation parameters. A specific analysis technique for studying stochastic processes is then defined, which focuses on the use of Monte Carlo simulations to determine statistically significant results for each of the specified performance metrics. The chapter concludes with an in-depth discussion of the performance of each Floquet mode controller. Baseline performance metrics are determined and compared to results listed in prior literature. Additionally, a comprehensive extended analysis highlights the potential for station-keeping failure due to increased state error.

- **Chapter 5:** A new hybrid controller is postulated that leverages Floquet mode and reinforcement learning. The introduction of reinforcement learning necessitates a discussion of the fundamental concepts associated with general reinforcement learning problems and an overview of the specific convergence algorithm used in this investigation. The framework that defines the reinforcement learning process is then introduced, covering the observation space, action space, reward definition, hybrid controller design, and training process. An agent is successfully trained to reliably meet the station-keeping objective. This performance is demonstrated in an analysis identical to the previous Floquet mode controller analysis, directly highlighting differences in performance metrics.
- **Chapter 6:** The body of work is summarized, following by remarks pertaining to significant conclusions. These remarks assess the results collected throughout the investigation, and discuss the applicability of reinforcement learning to the station-keeping problem. Finally, several recommendations for promising future work are suggested.

2. DYNAMICAL MODEL

Dynamical models allow systems of any nature to be represented by a set of mathematical laws that govern motion. Spacecraft tend to operate in complex dynamical fields under the influence of a multitude of forces. To model the motion of spacecraft and celestial bodies, a dynamical model must be defined which adequately represents motion throughout the field. To narrow the focus of this research, a subset of forces is accounted for by applying justifiable assumptions which simplify the equations of motion and allow for the application of many dynamical systems techniques which aid analysis. To examine the breadth of complexity inherent to this system, the general N -body problem is introduced. From this formulation, assumptions are outlined which ultimately yield the dynamical model used throughout this work, the circular restricted three-body problem (CR3BP).

2.1 The N-Body Problem

The general N -body formulation depicts the motion of an arbitrary N number of bodies in a given system. In the context of astrodynamics, a system is composed of N celestial bodies and spacecraft where each body acts on and is acted upon by all other $N-1$ bodies in the system. Illustrated in Figure 2.1, a body denoted by P_i with mass m_i is influenced by the relative position of all other bodies, while simultaneously affecting the motion of those bodies. The position of each body is measured relative to a fixed point O which acts as the origin in the inertial frame. The inertial reference frame is defined as a dextral coordinate system with axes ${}^I\hat{x} - {}^I\hat{y} - {}^I\hat{z}$, where the prescript denotes the inertial frame.




Figure 2.1. N -body System in the Inertial Frame.

The coupled interaction of forces between bodies are modeled by the direct application of Issac Newton's laws of motion. The second law of motion in vector form,

$$\bar{F} = m\ddot{\bar{r}} \quad (2.1)$$

provides a mathematical description that states the force on a body is proportional to the product of the mass and acceleration vector [2]. In Equation 2.1, the force F represents the vector summation of all forces acting on a body P_i . In [2], Newton also formulates the Law of Universal Gravitation for a pair of particles,

$$|F| = \frac{Gm_1m_2}{r^2} \quad (2.2)$$

where G is the gravitational constant. Equation 2.2 is applied for N particles and combined with Equation 2.1 to produce a differentiable equation which represents the motion of a particle P_i .

$$m_i\ddot{\bar{r}}_i = -G \sum_{\substack{j=1 \\ j \neq i}}^N \frac{m_i m_j}{r_{ji}^3} \bar{r}_{ji} \quad (2.3)$$

This differential equation in vector form collects six scalar first-order differential equations for each particle. The number of nonlinear differential equations scales proportionally to the number of particles, with $6N$ equations for a given N -body system.

The simplest case of Equation 2.3 where $N = 2$ is commonly referred to as the two-body problem. The two-body problem has been exhaustively researched with works and observations dating back to Kepler and Brahe [1]. For a two-body scenario, a general assumption asserts that the orbiting body (e.g., a spacecraft) is infinitesimal in comparison to the celestial body. As such, the presence of the smaller body does not affect the motion of the larger, celestial body. This assumption allows for the derivation of analytical solutions where the motion of the spacecraft is illustrated by conic sections. The application of geometrical laws allows for extensive analysis within the two-body problem.

While solutions to the two-body problem offer convenient means to model motion, the merit of analysis is restricted to a limited selection of scenarios. The two-body model isolates the system and therefore, neglects all other external forces that manifest in a higher fidelity model. Depending on the specific situation, the two-body model may provide an adequate representation, but often the inclusion of additional forces is essential to properly model motion and perform valid analysis. The logical next step to increase fidelity is to account for the presence of these additional bodies in the mathematical formulation.

2.2 The Circular Restricted Three-Body Problem

The general three-body problem defines an environment with three gravitational bodies, where each body influences the motion of the remaining two bodies. Consider a spacecraft operating in the vicinity of the Earth and consequently, the Moon. In the previously defined two-body model, the Earth is defined as the primary body influencing the motion of the spacecraft. However, inclusion of perturbing forces exerted by the Moon on the spacecraft may quickly render analysis performed in the

low fidelity two-body model obsolete. To develop a more thorough understanding of the actual motion of the spacecraft, the Moon is included directly in the model. However, the addition of the Moon, or any other body, in Equation 2.3 produces a differential equation with no closed-form solution,

$$m_3\ddot{\bar{r}}_3 = -G\frac{m_3m_1}{r_{13}^3}\bar{r}_{13} - G\frac{m_3m_2}{r_{23}^3}\bar{r}_{23} \quad (2.4)$$

To solve Equation 2.4, the time histories of the position of P_1 and P_2 are required. This information is not readily available as these positions now have a dependency on the state history of P_3 . To completely solve for all three state histories, 18 differential equations must be solved simultaneously which requires 18 integrals of the motion. Conservation of linear momentum, angular momentum, and energy provide six, three, and one integrals of the motion, respectively. The lack of remaining integrals of the motion render this equation to have no analytical solution. Instead, the problem is reformulated by applying appropriate simplifying assumptions which allow for numerical integration of this model.

2.2.1 Model Definition

The circular restricted three-body problem is a special case of the general three-body problem, which arises from the application of appropriate simplifying assumptions. These assumptions reduce the complexity of the mathematical equations which govern motion in this environment, allowing for direct integration and analysis. The motion in the restricted case is generally representative of actual motion in a three-body environment. As such, this model has proven to be sufficient in acting as a stepping stone towards a true N -body representation of motion [4].

To define the CR3BP, three simplifying assumptions are made which impact the dynamics of motion. The first assumption concerns the impact of the spacecraft on the dynamical system. The mass of the spacecraft, m_3 , is considered infinitesimal compared to the mass of the primaries, P_1 and P_2 . Due to this assumption, the gravitational effects of P_3 have no impact on the motion of either primary. To reinforce

this assertion, the system is assumed to act as an isolated two-body system with no external forces or perturbations. Within this two-body primary system, the first primary, P_1 , is arbitrarily defined as the celestial body with larger mass, m_1 . Therefore, the second primary, P_2 , has a smaller mass, m_2 . The final assumption asserts that the primaries rotate in circular orbits about their mutual barycenter, B . This assumption is not required, but simplifies the analysis. Other models, such as the elliptic restricted three-body problem, omit this assumption to account for eccentric orbits. However, the scope of this research only concerns assumed circular motion.

The combination of these assumptions defines the CR3BP. In addition, these assumptions impose that the motion of both massive primaries is constrained to a plane. Meanwhile, the massless third primary is "free" to move in three-dimensional space about the primaries. A two-dimensional schematic including the inertial and a new rotating reference frame is depicted in Figure 2.2.




Figure 2.2. Schematic of the Circular Restricted Three-Body Problem.

A rotating reference frame is defined as a dextral coordinate system with axes ${}^R\hat{x} - {}^R\hat{y} - {}^R\hat{z}$, where the prescript denotes the rotating frame. Following the same notation defined by Figure 2.1, three primaries exist, denoted by P_i . The line connecting P_1 and P_2 forms the x axis of the rotating frame. Between these primaries lies the point

B , which is the barycenter of P_1 and P_2 . Both the inertial and the rotating reference frames have an origin at B . The rotating frame is oriented at an angle θ off the inertial x axis, which defines a rotation between the inertial and rotating reference frames. The z axis is shared by both reference frames.

2.2.2 Equations of Motion

To mathematically describe the motion of P_3 in the CR3BP, a set of equations of motion are derived which are numerically integrated to produce trajectories. For convenience, and to aid with numerical computation, the equations of motion are derived in a non-dimensional form. In order to transition between dimensional and non-dimensional units, several characteristic quantities are defined for units of length, mass, and time. The quantities are termed, fittingly, the characteristic length, characteristic mass, and characteristic time, respectively.

$$l^* = |\bar{r}_1| + |\bar{r}_2| \quad (2.5)$$

$$m^* = m_1 + m_2 \quad (2.6)$$

$$t^* = \sqrt{\frac{l^*}{Gm^*}} \quad (2.7)$$

The characteristic length is defined as the distance between primaries, and the characteristic mass is the summation of the masses of both primaries. Characteristic time is selected as a matter of convenience in order to set the value of the non-dimensional gravitational constant equal to one.

To derive the equations of motion, a new quantity is introduced which is defined as μ , the mass ratio of the smaller primary to the total mass of the system.

$$\frac{m_1}{m^*} = \mu \quad (2.8)$$

$$\frac{m_2}{m^*} = 1 - \mu \quad (2.9)$$

Additionally, non-dimensional time τ is defined as,

$$\tau = \frac{t}{t^*} \quad (2.10)$$

Referencing Figure 2.2, the vector which represents the position of P_3 relative to the system barycenter is specified as a new variable in non-dimensional units in the rotating reference frame,

$$\bar{\rho} = \bar{r}_3 = x {}^R\hat{x} + y {}^R\hat{y} + z {}^R\hat{z} \quad (2.11)$$

To produce equations of motion, the derivative of Equation 2.11 is taken twice. To take a derivative relative to both the rotating and inertial reference frames, the Basic Kinematic Equation (BKE) is applied twice. The general form of the BKE is given by,

$$\frac{{}^I d\bar{\rho}}{d\tau} = \frac{{}^R d\bar{\rho}}{d\tau} + {}^I\bar{\omega}^R \times \bar{\rho} \quad (2.12)$$

which requires an angular velocity ${}^I\bar{\omega}^R$. The rotating frame has an angular velocity relative to the inertial frame defined by the non-dimensional mean motion n , which is equal to one due to the choice of t^* .

$${}^I\bar{\omega}^R = n {}^R\hat{z} \quad (2.13)$$

Plugging in Equation 2.13 and Equation 2.11 into Equation 2.12, and taking the derivative with respect to non-dimensional time twice, yields the equation for the inertial acceleration of P_3 expressed in rotating frame coordinates,

$$\frac{{}^I d^2\bar{\rho}}{d\tau^2} = (\ddot{x} - 2n\dot{y} - n^2x){}^R\hat{x} + (\ddot{y} + 2n\dot{x} - n^2y){}^R\hat{y} + \ddot{z} {}^R\hat{z} \quad (2.14)$$

Another expression for the left-hand side of Equation 2.14 is determined by returning to Equation 2.4. The mass of P_3 is removed from this equation and all terms are converted to non-dimensional units using the characteristic quantities,

$$\frac{{}^I d^2\bar{\rho}}{d\tau^2} = -\frac{1-\mu}{d^3}\bar{d} - \frac{\mu}{s^3}\bar{s} \quad (2.15)$$

where the relative vectors \bar{r}_{13} and \bar{r}_{23} are replaced by \bar{d} and \bar{s} , respectively. To determine vector equations for \bar{d} and \bar{s} , the definition of the center of mass is applied to determine that $\bar{r}_1 = -\mu {}^R\hat{x}$ and $\bar{r}_2 = (1-\mu){}^R\hat{x}$. Hence, the relative vectors \bar{d} and \bar{s} are determined,

$$\bar{d} = (x + \mu){}^R\hat{x} + y {}^R\hat{y} + z {}^R\hat{z} \quad (2.16)$$

$$\bar{s} = (x - 1 + \mu)^R \hat{x} + y^R \hat{y} + z^R \hat{z} \quad (2.17)$$

such that \bar{d} is the position vector from P_1 to P_3 and \bar{s} is the position vector from P_2 to P_3 . Finally, equations 2.14, 2.16, and 2.17 are substituted into Equation 2.15 to produce the non-dimensional equations of motion which govern the behavior of an infinitesimal particle in the CR3BP,

$$\ddot{x} - 2n\dot{y} - n^2x = -\frac{(1-\mu)(x+\mu)}{d^3} - \frac{\mu(x-1+\mu)}{s^3} \quad (2.18)$$

$$\ddot{y} + 2n\dot{x} - n^2y = -\frac{(1-\mu)y}{d^3} - \frac{\mu y}{s^3} \quad (2.19)$$

$$\ddot{z} = -\frac{(1-\mu)z}{d^3} - \frac{\mu z}{s^3} \quad (2.20)$$

These equations of motion exist as three, coupled second-order differential equations which describe the motion of P_3 . The autonomous system is transitioned to six first-order differential equations for numerical integration. For completeness, the values of d and s are determined as,

$$d = \sqrt{(x+\mu)^2 + y^2 + z^2} \quad (2.21)$$

$$s = \sqrt{(x-1+\mu)^2 + y^2 + z^2} \quad (2.22)$$

The system formulated here remains a conservative system, which suggests the existence of a potential function to represent the equations of motion. The equations of motion are defined in the rotating frame, and as such, a pseudo-potential function is defined,

$$U^* = \frac{1-\mu}{d} + \frac{\mu}{s} + \frac{1}{2}n^2(x^2 + y^2) \quad (2.23)$$

The equations of motion in Equations 2.18, 2.19, and 2.20 are represented in terms of the newly defined pseudo-potential U^* .

$$\ddot{x} - 2n\dot{y} = \frac{\partial U^*}{\partial x} \quad (2.24)$$

$$\ddot{y} + 2n\dot{x} = \frac{\partial U^*}{\partial y} \quad (2.25)$$

$$\ddot{z} = \frac{\partial U^*}{\partial z} \quad (2.26)$$

The equations of motion defined here provide a mathematical representation for the motion for P_3 due to the forces exerted by P_1 and P_2 . Unfortunately, the lack of known integrals of the motion prevent the derivation of analytical solutions. Instead, numerical integration of Equations 2.18, 2.19, and 2.20 is employed to produce continuous trajectories given any arbitrary initial state. Due to the lack of a closed form solution, this alone is not enough to develop a comprehensive understanding of the dynamics of this environment. In lieu of an analytical solution, many dynamical systems techniques may be applied to gain additional insights and appreciation for the natural dynamics within the circular restricted three-body problem.

2.2.3 Jacobi Constant

The formulation of the equations of motion in the rotating frame, coupled with the fact that the system remains conservative, suggests that an energy-like quantity exists that serves as an integral of the motion. This quantity, termed the Jacobi constant, is a scalar quantity which derives from the rotating velocity vector and the pseudo-potential function. The rotating velocity is represented as,

$$\dot{\rho} = \dot{x} {}^R\hat{x} + \dot{y} {}^R\hat{y} + \dot{z} {}^R\hat{z} \quad (2.27)$$

Equations 2.24 through 2.26 are then dotted with Equation 2.27 to produce three scalar equations. The sum of these scalar equations evaluates to,

$$\dot{x}\ddot{x} + \dot{y}\ddot{y} + \dot{z}\ddot{z} = \frac{\partial U^*}{\partial x}\dot{x} + \frac{\partial U^*}{\partial y}\dot{y} + \frac{\partial U^*}{\partial z}\dot{z} \quad (2.28)$$

where U^* is only a function of position. Therefore, the right side of Equation 2.28 is equal to the derivative of U^* with respect to non-dimensional time τ . This property allows the integration of Equation 2.28, which yields an expression for the Jacobi constant C .

$$\dot{x}^2 + \dot{y}^2 + \dot{z}^2 = 2U^* - C \quad (2.29)$$

The left side of Equation 2.29 is conveniently rewritten as the square of the rotating velocity magnitude, delivering a concise expression for the Jacobi constant,

$$C = 2U^* - \dot{\rho}^2 \quad (2.30)$$

This scalar quantity represents an energy-like parameter within the CR3BP. As an integral of the motion, the Jacobi constant assists in analysis and offers means to categorize orbits, determine energy changes to transfer between states, and to evaluate the error inherent to numerical processes. The availability of even one integral of the motion proves invaluable to advancing the understanding of this dynamical environment.

2.2.4 Equilibrium Points

A common starting place for the analysis of any dynamical system involves the search for equilibrium solutions. In this formulation, which occurs in a rotating frame, such solutions exist where rotating velocity and acceleration are zero. At each of these equilibrium solutions, a spacecraft or object remains in a constant position relative to the rotating frame. Since position remains fixed in the rotating frame and the pseudo-potential is only a function of position, the condition for equilibrium solutions is defined where the gradient of the pseudo-potential equals zero,

$$\frac{\partial U^*}{\partial x} = \frac{\partial U^*}{\partial y} = \frac{\partial U^*}{\partial z} = 0 \quad (2.31)$$

The equations of motion defined by Equations 2.18 through 2.20 are rearranged to meet this condition,

$$\frac{\partial U^*}{\partial x} = 0 = -\frac{(1-\mu)(x_{eq} + \mu)}{d_{eq}^3} - \frac{\mu(x_{eq} - 1 + \mu)}{s_{eq}^3} + n^2 x_{eq} \quad (2.32)$$

$$\frac{\partial U^*}{\partial y} = 0 = -\frac{(1-\mu)y_{eq}}{d_{eq}^3} - \frac{\mu y_{eq}}{s_{eq}^3} + n^2 y_{eq} \quad (2.33)$$

$$\frac{\partial U^*}{\partial z} = 0 = -\frac{(1-\mu)z_{eq}}{d_{eq}^3} - \frac{\mu z_{eq}}{s_{eq}^3} \quad (2.34)$$

where x_{eq} , y_{eq} , and z_{eq} are the positional components of the equilibrium solutions. In order to satisfy Equation 2.34, z_{eq} must be equal to zero. As a result, all equilibrium solutions in this model are confined to the xy plane.

The equilibrium points in this model, also known as the libration points, are categorized into two groups: collinear points and equilateral points. The categorization stems from analysis of Equations 2.32 and 2.33. An immediate conclusion from observing Equation 2.33 is that the equation is satisfied when y_{eq} is equal to zero. When this constraint is true, the equilibrium point(s) lie on the x axis. These solutions are defined as the collinear points. To solve for the x axis positions of the collinear points, Equation 2.32 is evaluated with $y_{eq} = z_{eq} = 0$, and augmented with sign functions,

$$x_{eq} = \mathcal{A} \frac{1 - \mu}{(x_{eq} + \mu)^2} + \mathcal{B} \frac{\mu}{(x_{eq} - 1 + \mu)^2} \quad (2.35)$$

Two sign functions, \mathcal{A} and \mathcal{B} , are defined to account for the loss of sign when simplifying both terms in Equation 2.35,

$$\mathcal{A} = \text{sign}(x_{eq} + \mu) \quad (2.36)$$

$$\mathcal{B} = \text{sign}(x_{eq} - 1 + \mu) \quad (2.37)$$

Three cases of Equation 2.35 exist which define the locations of the collinear libration points. The first libration point, L_1 , is determined from the case where \mathcal{A} is positive and \mathcal{B} is negative. The second point, L_2 , results from the case where \mathcal{A} and \mathcal{B} are both positive. The third and final collinear libration point, L_3 , is found from evaluating Equation 2.35 where \mathcal{A} and \mathcal{B} are both negative. The specific value of x_{eq} for each case is calculated by defining the quantity γ , which represents the displacement from the nearby primary,

$$x_{L_1} = 1 - \mu - \gamma_1 \quad (2.38)$$

$$x_{L_2} = 1 - \mu + \gamma_2 \quad (2.39)$$

$$x_{L_3} = -\mu - \gamma_3 \quad (2.40)$$

Equations 2.38 through 2.40 are separately substituted into Equation 2.35 with the appropriate sign terms. In all three equations, the only unknown quantity is γ_i . However, these equations cannot be algebraically solved for γ_i ; instead, an iterative root-finding algorithm, such as a Newton-Rhapson scheme, is used to numerically determine a value for each γ_i . For the case of L_1 ,

$$\gamma_1^{j+1} = \gamma_1^j - \frac{f(\gamma_1)}{\dot{f}(\gamma_1)} \quad (2.41)$$

defines a simple update equation to iteratively determine γ_1 where,

$$f(\gamma_1) = \frac{1 - \mu}{(1 - \gamma_1)^2} - \frac{\mu}{(\gamma_1)^2} - 1 + \mu + \gamma_1 \quad (2.42)$$

$$\dot{f}(\gamma_1) = \frac{2(1 - \mu)}{(1 - \gamma_1)^3} + \frac{2\mu}{(\gamma_1)^3} + 1 \quad (2.43)$$

A suitable initial guess is required for this formulation to successfully converge to the correct value. The same process is applied for L_2 and L_3 to find the exact locations of all collinear libration points. The general location of each collinear libration point is depicted in Figure 2.3.

Another solution to Equations 2.32 and 2.33 occurs when the values d_{eq} and s_{eq} are equal. Recall that these two quantities are the distances from P_1 and P_2 to P_3 , respectively. Hence, a point located at an equal distance from both primaries and constrained to the xy plane forms an equilateral triangle with the primaries. The two equilibrium solutions corresponding to these equilateral points are simply expressed as a function of the system μ ,

$$x_{L_4, L_5} = \frac{1}{2} - \mu \quad (2.44)$$

$$y_{L_4, L_5} = \pm \frac{\sqrt{3}}{2} \quad (2.45)$$

The two equilateral solutions complete the set of five libration points within this model. Figure 2.3 shows the approximate location of all libration points in the CR3BP. The libration points L_1 and L_2 are always located on either side of P_2 , but are not equidistant from P_2 . Additionally, L_3 is always located to the left of P_1 .




Figure 2.3. Equilibrium Points of the CR3BP.

However, the relative position of all libration points shifts within these constraints as the system mass ratio varies.

2.2.5 Zero Velocity Curves

While equilibrium points stipulate that both the rotating velocity and acceleration must be zero, there exists another condition of interest where solely the rotating velocity is zero. Recall from Equation 2.30, the magnitude of the rotating velocity is expressed by,

$$\dot{\rho}^2 = 2U^* - C \quad (2.46)$$

Mathematically, Equation 2.46 is valid for all values of U^* and C . However, clearly for the condition that $C > 2U^*$, the resulting rotating velocity magnitude is imaginary. For a physical system, this condition is an impossibility. Hence, the special case of $C = 2U^*$, where $\dot{\rho} = 0$, is of particular interest. Evaluating Equation 2.46 for

$\dot{\rho} = 0$ and substituting the expression for U^* yields an expression which defines this particular case,

$$C = x^2 + y^2 + \frac{2(1 - \mu)}{d} + \frac{2\mu}{s} \quad (2.47)$$

Locations in configuration space that satisfy Equation 2.47 for a given C value define the zero velocity surfaces. These surfaces exist in three-dimensional space and act as boundaries restricting the range of potential movement of the body P_3 . However, the zero velocity surfaces offer no insight into how a propagated initial state traverses within the specified boundaries.

For the purposes of this analysis, only cross sections of these surfaces are employed. The resulting planar structure is termed a zero velocity curve. For planar motion, for instance the xy plane, the zero velocity curves create forbidden regions which a spacecraft or object may never enter without a change in the current energy level. An example illustrating the structure of a zero velocity curve for a specific value of C is depicted in Figure 2.4. The zero velocity curves are represented as three distinct black curves; one curve surrounds P_1 , another curve surrounds P_2 , and a third curve surrounds the entire system. The area shaded in gray represents the forbidden regions where a spacecraft cannot reach under natural propagation. For this value of Jacobi constant, two areas of valid motion around both primaries appear. A spacecraft located within the curve surrounding either the Earth or the Moon is unable to escape the vicinity of the respective body, due to a lack in energy. In order to escape either body, a change in energy is required.

In the scenario depicted by Figure 2.4, all libration point gateways are closed. From Equation 2.47, the Jacobi constant value corresponding to each libration point is determined. In order for a gateway to open, the Jacobi constant value corresponding to a spacecraft's state must be less than the Jacobi constant value of the libration point. The order in which the gateways open follows the naming convention of the libration points; L_1 opens first, followed by L_2 , then L_3 , and finally L_4 and L_5 open at the same value of C . Therefore, the knowledge of a spacecraft's state is adequate




Figure 2.4. Zero Velocity Curves in the Earth-Moon System.

to determine if the spacecraft is bounded to a primary, may move between primaries, or even escape the system entirely.

2.3 System Parameters

Throughout this work, two three-body gravitational systems are utilized: the Earth-Moon and the Sun-Earth Moon Barycenter (EMB) systems. Several examples are included in the Earth-Moon system for demonstration purposes. Critical system parameters for this system are listed in Table 2.1. However, the majority of this research is contained within the Sun-EMB system. Table 2.2 lists Sun-EMB specific parameters that are used for all examples throughout this investigation. The inclusion of each table in this section captures all information required to replicate results within

Table 2.1. Earth-Moon System Parameters.

System Parameter	Value	Units
Mass Ratio (μ)	0.0121505842699404	-
Characteristic Length (l^*)	384747.991979046	km
Characteristic Mass (m^*)	$6.04582557449506 \times 10^{24}$	kg
Characteristic Time (t^*)	375699.859037759	s

Table 2.2. Sun-EMB System Parameters.

System Parameter	Value	Units
Mass Ratio (μ)	$3.04042340382006 \times 10^{-6}$	-
Characteristic Length (l^*)	149597892.162912	km
Characteristic Mass (m^*)	$1.98848146179211 \times 10^{30}$	kg
Characteristic Time (t^*)	5022636.3363283	s

the given system. Characteristic quantities are documented to allow for the transition between dimensional and non-dimensional units, and system mass ratios are provided to aid in consistency with processes requiring numerical integration.

3. DYNAMICAL SYSTEMS THEORY

The circular restricted three-body problem is a dynamical model that inherits the complexities and difficulties associated with nonlinear systems. These challenges are due, in part, to the lack of fundamental dynamical systems theory valid for nonlinear systems. Many common dynamical systems techniques are only applicable to linear systems. However, with certain restrictions and limitations, these tools provide many applications which aid analyses and study within the circular restricted three-body problem. Local linearization of the nonlinear equations of motion opens the door to discovering how natural flow propagates within this model. This linearization allows for stability analysis and the implementation of many numerical techniques, such as targeting and continuation methods. Using these tools, natural periodic orbits are generated which extend into families of periodic orbits, revealing fascinating geometric structures which may be leveraged for various mission applications.

3.1 Linear Variational Equations

For a linear system, the equations of motion contain crucial information which determines local stability and imposes constraints on initial conditions. Therefore, the linearization of the nonlinear equations of motion of the CR3BP has the potential to provide additional information. The nonlinear equations must be linearized relative to a reference point; the equilibrium solutions provide a convenient set of states about which the equations of motion are linearized. However, any reference state is appropriate. Consider any of the libration points, which possess a state defined by $\bar{x}_{eq} = \{x_{eq}, y_{eq}, z_{eq}, 0, 0, 0\}^T$. This state solution is then perturbed from equilibrium, denoted by $\delta\bar{x} = \{\xi, \eta, \zeta, \dot{\xi}, \dot{\eta}, \dot{\zeta}\}^T$. This perturbation is introduced to the equations of motion defined by Equations 2.18, 2.19, and 2.20. A Taylor series expansion about

the equilibrium solution is performed, neglecting the higher order terms, to yield the first order, linear variational equations,

$$\ddot{\xi} - 2\dot{\eta} = U_{xx}^* \xi + U_{xy}^* \eta + U_{xz}^* \zeta \quad (3.1)$$

$$\ddot{\eta} - 2\dot{\xi} = U_{yx}^* \xi + U_{yy}^* \eta + U_{yz}^* \zeta \quad (3.2)$$

$$\ddot{\zeta} = U_{zx}^* \xi + U_{zy}^* \eta + U_{zz}^* \zeta \quad (3.3)$$

Neglecting the higher order terms in the Taylor series expansion is not required, but performed as a matter of convenience. Analysis has shown that neglecting these higher order terms in Equations 3.1 through 3.3 considerably reduces the complexity and computation requirements of these equations. However, the inclusion of additional terms provides another level of fidelity useful to certain applications [11]. The subscripts on the pseudo-potential function U_{ij}^* indicate the second derivative of Equation 2.23 with respect to variables i and j . For reference, the equations for each second partial term are,

$$U_{xx}^* = 1 - \frac{1-\mu}{d^3} - \frac{\mu}{s^3} + \frac{3(1-\mu)(x+\mu)^2}{d^5} + \frac{3\mu(x-1+\mu)^2}{s^5} \quad (3.4)$$

$$U_{yy}^* = 1 - \frac{1-\mu}{d^3} - \frac{\mu}{s^3} + \frac{3(1-\mu)y^2}{d^5} + \frac{3\mu y^2}{s^5} \quad (3.5)$$

$$U_{zz}^* = -\frac{1-\mu}{d^3} - \frac{\mu}{s^3} + \frac{3(1-\mu)z^2}{d^5} + \frac{3\mu z^2}{s^5} \quad (3.6)$$

$$U_{xy}^* = U_{yx}^* = \frac{3(1-\mu)(x+\mu)y}{d^5} + \frac{3\mu(x-1+\mu)y}{s^5} \quad (3.7)$$

$$U_{xz}^* = U_{zx}^* = \frac{3(1-\mu)(x+\mu)z}{d^5} + \frac{3\mu(x-1+\mu)z}{s^5} \quad (3.8)$$

$$U_{yz}^* = U_{zy}^* = \frac{3(1-\mu)yz}{d^5} + \frac{3\mu yz}{s^5} \quad (3.9)$$

Equations 3.1, 3.2, and 3.3 provide a means to gain insight into the motion and stability near a specified reference. In this formulation, the reference is an equilibrium solution. Using these equations, linear approximations for the motion near libration points are readily determined. Additionally, the same approach is equally valid for

any arbitrary reference state, which infers the applicability of this process to the propagation of trajectories.

3.2 State Transition Matrix

At this point, the initial state of a spacecraft in the CR3BP provides a limited amount of information. Using the equations of motion, any initial state may be propagated to produce a time invariant trajectory of position and velocity states. While the knowledge of what will happen given an initial state is important, such information is limited to only that state and provides no other insights. Instead, understanding the natural flow around a particular reference state allows for additional levels of analysis which is leveraged to further explore this dynamical model.

Recall the equations of motion given by Equations 2.18, 2.19, and 2.20. These three coupled second-order differential equations are rewritten as six first-order differential equations of the form $\dot{\bar{x}} = \bar{f}(\bar{x}, t)$, where $\bar{x} = \{x, y, z, \dot{x}, \dot{y}, \dot{z}\}^T$. The state given by \bar{x} represents any set of position and velocity parameters at a given point in time. As such, consider a state which resides on a reference trajectory, denoted by $\bar{x}^*(t)$. At time t , the state is perturbed by a small variation $\delta\bar{x}(t)$, which creates a new, perturbed set of states that now lie on a perturbed trajectory. In the scenario




Figure 3.1. Perturbed Trajectory Relative to a Reference Trajectory.

illustrated by Figure 3.1, the set of initial conditions which previously produced the reference trajectory, now lies on the perturbed trajectory. Mathematically, this is simply expressed as,

$$\bar{x}(t) = \bar{x}^*(t) + \delta\bar{x}(t) \quad (3.10)$$

Similar to the methodology which produced linear variational equations when perturbing an equilibrium state, a Taylor series expansion of Equation 3.10 about the reference state \bar{x}^* produces a linear approximation for $\delta\dot{\bar{x}}$,

$$\delta\dot{\bar{x}} = \left. \frac{\partial \bar{f}}{\partial \bar{x}} \right|_{\bar{x}^*(t)} \delta\bar{x}(t) \quad (3.11)$$

For convenience, Equation 3.11 is rewritten as,

$$\delta\dot{\bar{x}}(t) = A(t)\delta\bar{x}(t) \quad (3.12)$$

where $A(t)$ is the Jacobian matrix of the equations of motion and is evaluated on the reference trajectory. For this formulation, the matrix $A(t)$ is a 6×6 matrix comprised of the second partial derivatives of the pseudo-potential function,

$$A(t) = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ U_{xx}^* & U_{xy}^* & U_{xz}^* & 0 & 2 & 0 \\ U_{yx}^* & U_{yy}^* & U_{yz}^* & -2 & 0 & 0 \\ U_{zz}^* & U_{zy}^* & U_{zz}^* & 0 & 0 & 0 \end{bmatrix} \quad (3.13)$$

The introduction of Equation 3.12 to the numerical integration process allows the calculation of a state transition matrix (STM), which maps variations in $\bar{x}(t)$ due to perturbations applied to $\bar{x}(t_0)$. The STM, expressed as $\Phi(t, t_0)$, is calculated by the differential equation,

$$\dot{\Phi}(t, t_0) = A(t)\Phi(t, t_0) \quad (3.14)$$

with an initial state $\Phi(t_0, t_0) = I_{6 \times 6}$. The STM is evaluated on the reference trajectory by numerically integrating Equation 3.14 alongside the equations of motion to

produce $\Phi(t, t_0)$. Each element of $\Phi(t, t_0)$ is a partial which relates a perturbation in a specific initial state in $\bar{x}(t_0)$, to a variation in specific final state in $\bar{x}(t)$. Therefore, the STM is expressed as a matrix of partials,

$$\Phi(t, t_0) = \begin{bmatrix} \frac{\partial x}{\partial x_o} & \frac{\partial x}{\partial y_o} & \frac{\partial x}{\partial z_o} & \frac{\partial x}{\partial \dot{x}_o} & \frac{\partial x}{\partial \dot{y}_o} & \frac{\partial x}{\partial \dot{z}_o} \\ \frac{\partial y}{\partial x_o} & \frac{\partial y}{\partial y_o} & \frac{\partial y}{\partial z_o} & \frac{\partial y}{\partial \dot{x}_o} & \frac{\partial y}{\partial \dot{y}_o} & \frac{\partial y}{\partial \dot{z}_o} \\ \frac{\partial z}{\partial x_o} & \frac{\partial z}{\partial y_o} & \frac{\partial z}{\partial z_o} & \frac{\partial z}{\partial \dot{x}_o} & \frac{\partial z}{\partial \dot{y}_o} & \frac{\partial z}{\partial \dot{z}_o} \\ \frac{\partial \dot{x}}{\partial x_o} & \frac{\partial \dot{x}}{\partial y_o} & \frac{\partial \dot{x}}{\partial z_o} & \frac{\partial \dot{x}}{\partial \dot{x}_o} & \frac{\partial \dot{x}}{\partial \dot{y}_o} & \frac{\partial \dot{x}}{\partial \dot{z}_o} \\ \frac{\partial \dot{y}}{\partial x_o} & \frac{\partial \dot{y}}{\partial y_o} & \frac{\partial \dot{y}}{\partial z_o} & \frac{\partial \dot{y}}{\partial \dot{x}_o} & \frac{\partial \dot{y}}{\partial \dot{y}_o} & \frac{\partial \dot{y}}{\partial \dot{z}_o} \\ \frac{\partial \dot{z}}{\partial x_o} & \frac{\partial \dot{z}}{\partial y_o} & \frac{\partial \dot{z}}{\partial z_o} & \frac{\partial \dot{z}}{\partial \dot{x}_o} & \frac{\partial \dot{z}}{\partial \dot{y}_o} & \frac{\partial \dot{z}}{\partial \dot{z}_o} \end{bmatrix} \quad (3.15)$$

Each partial provides unique information applicable to a wide variety of numerical methods. However, as a linear approximation, the numerical accuracy of mapping an initial perturbation to a final variation is dependent on the length of the integration and the magnitude of the perturbation. Additionally, inclusion of Equation 3.14 in the integration process increases the number of differential equations by thirty six, requiring significantly more computational effort. Despite these drawbacks, the information provided by the STM considerably expands capabilities and understanding, allowing further exploration of this dynamical model.

3.3 Differential Corrections

Often, a specific initial state is desired which delivers a spacecraft to a location within a set of constraints. Some examples of these constraints include traversing from one location to another in a specified amount of time, arriving at a particular point of interest, or achieving a trajectory with some distinct geometrical structure. A method to reliably meet these requirements relies on a differential corrections technique which iteratively updates an initial state using sensitivity information provided by the state

transition matrix. For this investigation, a single shooting scheme is employed to correct initial conditions to produce periodic orbits.

3.3.1 Single Shooting

One methodology for satisfying a set of constraints is a single shooting scheme. This technique iteratively updates a single set of initial conditions to meet constraints on a desired final state. This process is accomplished by collecting all initial state components that are free to update in a design variable vector, \bar{X} . Then, all boundary conditions are listed in a constraint vector, $\bar{F}(\bar{X})$. For this formulation, the design variable vector,

$$\bar{X} = \begin{bmatrix} X_1 \\ X_2 \\ \vdots \\ X_n \end{bmatrix} \quad (3.16)$$

is a one-dimensional vector of length n . Each design variable X_i represents a position or velocity state component that is iteratively updated. The constraint vector,

$$\bar{F}(\bar{X}) = \begin{bmatrix} F_1(\bar{X}) \\ F_2(\bar{X}) \\ \vdots \\ F_m(\bar{X}) \end{bmatrix} = \bar{0} \quad (3.17)$$

is also a one-dimensional vector, but of length m . Each entry $F_i(\bar{X})$ is an equation or variable which must be equal to the zero constraint at a defined time t . To modify the design variables to meet the constraint equations, an iterative update equation which approximates $\bar{F}(\bar{X})$ is determined by performing a Taylor series expansion of $\bar{F}(\bar{X})$ about the initial \bar{X} ,

$$\bar{F}(\bar{X}^j) + D\bar{F}(\bar{X}^j) \cdot (\bar{X}^{j+1} - \bar{X}^j) = \bar{0} \quad (3.18)$$

The exclusion of higher order terms implies that Equation 3.18 only approximates the value of $\bar{F}(\bar{X})$. As such, Equation 3.18 is applied iteratively until a small tolerance is met,

$$||F(\bar{X}^{j+1})|| < \varepsilon \quad (3.19)$$

The targeting tolerance ε is generally around order 10^{-10} , but varies depending on the specific application. To solve Equation 3.18, the Jacobian matrix $D\bar{F}(\bar{X})$ is defined, which collects the partial derivatives of the constraint vector elements with respect to the design vector elements,

$$D\bar{F}(\bar{X}^j) = \frac{\partial F}{\partial X} = \begin{bmatrix} \frac{\partial F_1}{\partial X_1} & \frac{\partial F_1}{\partial X_2} & \cdots & \frac{\partial F_1}{\partial X_n} \\ \frac{\partial F_2}{\partial X_1} & \frac{\partial F_2}{\partial X_2} & \cdots & \frac{\partial F_2}{\partial X_n} \\ \vdots & \vdots & \ddots & \vdots \\ \frac{\partial F_m}{\partial X_1} & \frac{\partial F_m}{\partial X_2} & \cdots & \frac{\partial F_m}{\partial X_n} \end{bmatrix} \quad (3.20)$$

With each component of Equation 3.18 defined, the equation is rearranged to solve for \bar{X}^{j+1} , noting the number of design variables and constraint equations,

$$\bar{X}^{j+1} = \bar{X}^j - D\bar{F}(\bar{X}^j)^{-1} \bar{F}(\bar{X}^j) \quad (3.21)$$

For Equation 3.21 to be solvable, the number of design variables n must be equal to the number of constraint equations m . When this condition is satisfied, the system is fully constrained with a singular unique solution, and the matrix $D\bar{F}(\bar{X}^j)$ is square and invertible. Another condition exists where the number of design variables n exceeds the number of constraint equations m . For this scenario, infinitely many solutions exist, and $D\bar{F}(\bar{X}^j)$ is no longer square and invertible. Several approaches exist to identify a particular solution for underdetermined systems. One such approach, the minimum norm solution, identifies a solution for \bar{X}^{j+1} that is closest to \bar{X}^j . Consequently, the final solution determined by the minimum norm approach is

dependent on the initial choice of \bar{X} . To solve for \bar{X}^{j+1} , Equation 3.21 is modified to satisfy the minimum norm approach,

$$\bar{X}^{j+1} = \bar{X}^j - D\bar{F}(\bar{X}^j)^T [D\bar{F}(\bar{X}^j) \cdot D\bar{F}(\bar{X}^j)^T]^{-1} \bar{F}(\bar{X}^j) \quad (3.22)$$

Similar to the fully constrained system, Equation 3.22 is iteratively solved until $\|F(\bar{X}^{j+1})\| < \varepsilon$. In general, both Equations 3.21 and 3.22 are converged in a relatively few number of iterations, given an appropriate initial \bar{X} . Accordingly, this approach proves to be a lucrative method of satisfying a variety of mission design requirements.

3.3.2 Targeting Perpendicular Crossings

The targeting scheme most prevalent in this investigation involves a constraint on a trajectory where the velocity in the x direction is zero when the y position is zero. In other words, the trajectory must have a perpendicular intersection with the xz -plane. By targeting sets of initial conditions that meet this constraint, fascinating geometric structures are created, known as periodic orbits. The concept of targeting a perpendicular crossing to create a periodic solution leverages the mirror theorem.

Theorem 3.3.1 (Mirror Theorem) *If n point-masses are acted upon by their mutual gravitational forces only, and at a certain epoch each radius vector from the (assumed stationary) centre of mass of the system is perpendicular to every velocity vector, then the orbit of each mass after that epoch is a mirror image of its orbit prior to that epoch. [12]*

The CR3BP contains structures which exhibit mirror configurations across, most commonly, the plane defined by $y = 0$. Therefore, if a set of initial conditions lie on this plane and depart perpendicularly forward in time, and if the propagated trajectory returns perpendicularly to the plane, a periodic structure arises.

Planar Perpendicular Crossings

Consider trajectories which solely exist on the xy -plane in the CR3BP. To develop a trajectory which is periodic across the plane $y = 0$, defined as the map Σ , a set of initials conditions is chosen which lie on the x axis and depart perpendicularly, denoted as $\bar{x}_0 = \{x_0, 0, 0, 0, \dot{y}_0, 0\}^T$. Illustrated in Figure 3.2, propagating an arbitrary \bar{x}_0 until a return to Σ generally delivers a non-perpendicular final state. To successfully correct \bar{x}_0 such that the return to Σ is perpendicular, an initial guess which generates a *nearly* periodic trajectory is chosen. Given the restrictions on




Figure 3.2. Planar Perpendicular Targeting Problem.

\bar{x}_0 , either x_0 or \dot{y}_0 are updated in order to return perpendicularly to Σ . In this formulation, x_0 is fixed and \dot{y}_0 is iteratively determined. Therefore, the design variable vector for a time independent two-dimensional perpendicular targeter is simply,

$$\bar{X} = \dot{y}_0 \quad (3.23)$$

Other formulations include time t in the design variable vector. However, this approach assumes time as an unconstrained, implicit variable. The initial state is propagated until returning to Σ , where t is simply the amount of propagation time. Then,

at Σ , the requirement which forces perpendicular crossings is collected in the constraint vector,

$$\bar{F}(\bar{X}) = \dot{x}_{\{y=0\}}(\dot{y}_0, y_{\{y=0\}}(\dot{y}_0)) \quad (3.24)$$

Equation 3.24 indicates that at Σ , \dot{x} has functional dependencies on \dot{y}_0 and y . Additionally, y has a functional dependency on \dot{y}_0 . The documentation of the implicit dependencies on each variable in Equation 3.24 is essential to properly evaluate the partials in $D\bar{F}(\bar{X})$. This formulation has one design variable and one constraint equation; therefore, $D\bar{F}(\bar{X})$ reduces to a 1×1 matrix,

$$D\bar{F}(\bar{X}) = \frac{\partial F_1}{\partial X_1} \quad (3.25)$$

Noting the functional dependencies in Equation 3.24, the partial derivative of \dot{x} at Σ with respect to \dot{y}_0 is evaluated by applying the chain rule to produce,

$$D\bar{F}(\bar{X}) = \frac{\partial \dot{x}}{\partial \dot{y}_0} - \frac{\ddot{x}}{\dot{y}} \frac{\partial y}{\partial \dot{y}_0} \quad (3.26)$$

In Equation 3.26, \ddot{x} and \dot{y} are final trajectory states. Both partial terms are calculated from integrating the STM along the trajectory. Therefore, Equation 3.26 is concisely expressed as,

$$D\bar{F}(\bar{X}) = \phi_{45} - \frac{\ddot{x}}{\dot{y}} \phi_{25} \quad (3.27)$$

where ϕ_{ij} references the corresponding element of the STM, listed in Equation 3.15. Finally, an update equation to iteratively determined \dot{y}_0 is created by substituting Equations 3.23, 3.24, and 3.27 into Equation 3.21, yielding,

$$\dot{y}_0^{j+1} = \dot{y}_0^j - (\phi_{45} - \frac{\ddot{x}}{\dot{y}} \phi_{25})^{-1} \cdot \dot{x}^j \quad (3.28)$$

Proper implementation of Equation 3.28 in an iterative algorithm allows the determination of a set of initial conditions which creates a periodic orbit. However, an appropriate guess for the initial state is pertinent to successfully converging a periodic solution.

Spatial Perpendicular Crossings

Similar to the planar perpendicular crossing targeter, a targeting scheme is employed to compute periodic orbits that exist in three-dimensional space. As before, a map is defined as $\Sigma : y = 0$, at which the trajectory must intersect perpendicularly. For spatial orbits, this condition implies that \dot{x} and \dot{z} are zero at Σ . Therefore, consider an initial state $\bar{x}_0 = \{x_0, 0, z_0, 0, \dot{y}_0, 0\}^T$, where either x_0 , z_0 , or \dot{y}_0 is fixed. Then, the remaining two free variables are iteratively updated. The two cases where x_0 and z_0 are fixed are of particular interest to this investigation.

To derive the update equations for the fixed x targeter, the process outlined for planar periodic orbits is replicated. The design variables and constraint equations are collected, noting functional dependencies,

$$\bar{X} = \begin{bmatrix} z_0 \\ \dot{y}_0 \end{bmatrix} \quad (3.29)$$

$$\bar{F}(\bar{X}) = \begin{Bmatrix} \dot{x}_{\{y=0\}}(x_0, \dot{y}_0, y_{\{y=0\}}(x_0, \dot{y}_0)) \\ \dot{z}_{\{y=0\}}(x_0, \dot{y}_0, y_{\{y=0\}}(x_0, \dot{y}_0)) \end{Bmatrix} \quad (3.30)$$

As before, time is optionally included in a fixed x targeter, but excluded in this formulation. The design variable vector and constraint vector are both of length two, resulting in a fully determined system where,

$$D\bar{F}(\bar{X}) = \begin{bmatrix} \frac{\partial F_1}{\partial X_1} & \frac{\partial F_1}{\partial X_2} \\ \frac{\partial F_2}{\partial X_1} & \frac{\partial F_2}{\partial X_2} \end{bmatrix} \quad (3.31)$$

Each partial derivative in Equation 3.31 is evaluated by direct application of the chain rule, and the resulting form is substituted into Equation 3.21 along with Equations 3.29 and 3.30 to produce update equations for the fixed x targeter,

$$\begin{bmatrix} z_0 \\ \dot{y}_0 \end{bmatrix}^{j+1} = \begin{bmatrix} z_0 \\ \dot{y}_0 \end{bmatrix}^j - \begin{bmatrix} \phi_{43} - \frac{\ddot{x}}{\dot{y}}\phi_{23} & \phi_{45} - \frac{\ddot{x}}{\dot{y}}\phi_{25} \\ \phi_{63} - \frac{\ddot{z}}{\dot{y}}\phi_{23} & \phi_{65} - \frac{\ddot{z}}{\dot{y}}\phi_{25} \end{bmatrix}^{-1} \begin{bmatrix} \dot{x}_{\{y=0\}} \\ \dot{z}_{\{y=0\}} \end{bmatrix}^j \quad (3.32)$$

Both update equations in Equation 3.32 are iterated on simultaneously until the targeting tolerance is met, resulting in a converged set of initial conditions that yield a periodic solution. Following the exact same process, a matrix update equation for the fixed z targeter is derived. For reference, the design variable vector and constraint vector are,

$$\bar{X} = \begin{bmatrix} x_0 \\ \dot{y}_0 \end{bmatrix} \quad (3.33)$$

$$\bar{F}(\bar{X}) = \begin{Bmatrix} \dot{x}_{\{y=0\}}(x_0, \dot{y}_0, y_{\{y=0\}}(x_0, \dot{y}_0)) \\ \dot{z}_{\{y=0\}}(x_0, \dot{y}_0, y_{\{y=0\}}(x_0, \dot{y}_0)) \end{Bmatrix} \quad (3.34)$$

The partial matrix $D\bar{F}(\bar{X})$ is similarly evaluated, delivering the update equations for a fixed z targeter,

$$\begin{bmatrix} x_0 \\ \dot{y}_0 \end{bmatrix}^{j+1} = \begin{bmatrix} x_0 \\ \dot{y}_0 \end{bmatrix}^j - \begin{bmatrix} \phi_{41} - \frac{\ddot{x}}{\dot{y}}\phi_{21} & \phi_{45} - \frac{\ddot{x}}{\dot{y}}\phi_{25} \\ \phi_{61} - \frac{\ddot{z}}{\dot{y}}\phi_{21} & \phi_{65} - \frac{\ddot{z}}{\dot{y}}\phi_{25} \end{bmatrix}^{-1} \begin{bmatrix} \dot{x}_{\{y=0\}} \\ \dot{z}_{\{y=0\}} \end{bmatrix}^j \quad (3.35)$$

Depending on the specific geometry, either Equation 3.32 or 3.35 is used to converge a set of initial conditions that results in a periodic orbit. However, the probability of successfully converging a periodic orbit with either equation still relies on the relative proximity of the initial guess.

3.3.3 Natural Parameter Continuation

Natural parameter continuation is one of many numerical continuation techniques which are used to evolve a set of periodic solutions. A recurring difficulty throughout this work and related research revolves around the search for an accurate initial guess required for convergence. As such, natural parameter continuation is implemented in this document as a means to approximate a series of related initial conditions to periodic orbits. Each periodic orbit in the CR3BP exists as a member of a family of periodic solutions [13]. Consequently, once a single periodic solution is discovered, a family of orbits are generated using natural parameter continuation.

Equations 3.28, 3.32, and 3.35 all require relatively accurate approximations for the initial guess used in each respective update equation. If an arbitrary set of conditions are used for the initial guess, the update equations may still converge, but to another solution which differs from the desired result. To account for this behavior when generating a family of periodic orbits, the orbit family is parameterized by some natural property of the orbit, such as the Jacobi constant, orbital period, a position component, or a velocity component. A common, intuitive choice for the characterization of planar orbit families is the initial x position, x_0 . For this scenario, an approximation for the next periodic solution is taken from the previous periodic solution while incrementing x_0 by a step size Δx_0 . The step size, which is generally small, is chosen such that the initial guess for the next periodic solution adequately approximates the true periodic solution.

3.3.4 Lyapunov Orbit Families

The concepts presented in the previous sections are used in conjunction to produce families of planar periodic orbits. In particular, a family of planar orbits about the L_1 libration point is discovered, called the L_1 Lyapunov family. In order to determine a series of orbits that represent this family, a starting point is required. To begin, the linear variational equations listed in Equations 3.1 through 3.3 are employed to determine a set of initial conditions that produce a linear periodic orbit about the L_1 libration point. These initial conditions are then used as the initial guess in a targeting scheme governed by Equation 3.28 to converge to a set of initial conditions that generate a nonlinear periodic orbit. After the first nonlinear periodic orbit is discovered, successive orbits along the family are found using natural parameter continuation.

To produce a linear orbit in the Sun-EMB system about the L_1 libration point, a *step* is taken away from the libration point. The size of the step is an arbitrary distance in a specified direction off of the libration point. While the exact distance

is not pertinent, behavior in the vicinity of the libration is better represented by the linear variational equations. Consequently, the step size should be reasonable for the system. For the Sun-EMB system, a step size of about 75000 *km* in the negative *x* direction is selected. Using the linear variational equations, a linear periodic orbit is produced, illustrated in Figure 3.3, where the arrow indicates the direction of flow along the orbit.




Figure 3.3. L_1 Linear Variational Orbit in the Sun-EMB system.

Integrating the initial conditions of this linear variational orbit with the nonlinear equations of motion produces a non-periodic orbit. Numerically correcting for the initial velocity in the *y* direction converges to a new set of initial conditions that produce a nonlinear periodic orbit. In Figure 3.4, the corrected L_1 Lyapunov orbit is shown overlaid the linear variational orbit. Geometrically, both orbits look relatively similar. This observation reinforces the idea that near the libration points,




Figure 3.4. Corrected L_1 Lyapunov Orbit in the Sun-EMB System.

linear analysis adequately approximates the nonlinear behavior. However, as the distance from the libration point increases, the structure of the nonlinear periodic orbits changes drastically from a linear counterpart.

The family is now continued in the negative x direction using natural parameter continuation. Figure 3.5 shows the initial evolution of the L_1 Lyapunov family. As previously discussed, each successive orbit is used as a stepping stone to reach the next orbit in the family. Without this approach, or another continuation method, it would be impractical to determine an initial guess for the initial conditions of an orbit that is no longer near the L_1 libration point. This process is repeated to continually grow the family until correcting for the next family member is no longer numerically feasible. In the Sun-EMB system and other systems with equivalent or smaller mass ratios, it becomes numerically challenging to correct for periodic orbits when operating in




Figure 3.5. Evolution of the Sun-EMB L_1 Lyapunov Family.

the vicinity of the primaries. Both primaries act essentially as singularities, leading to complex and rapidly changing dynamics that are less than ideal for numerical applications. Nevertheless, a representative view of a portion of the Sun-EMB L_1 Lyapunov family is depicted in Figure 3.6. At this point, attempts to continue to grow the size of the family result in numerical failure, where the targeting tolerance ε is no longer achievable. In order to further continue the L_1 Lyapunov family in the Sun-EMB system, another corrections approach, such as multiple-shooting, must be employed.

It is worth noting that a family of orbits makes up a continuous structure, and therefore, an infinite number of orbits exist within the family. By displaying a subset of members spaced throughout the family, a general geometric portrait of the family is produced. However, the orbits shown in Figure 3.6 only highlight a portion of the




Figure 3.6. Sun-EMB L_1 Lyapunov Family.

L_1 Lyapunov family, and are not an exhaustive representation of the entire family. In fact, in other systems such as the Earth-Moon system, a L_1 Lyapunov family that covers the entirety of the expanse between both primaries is easily generated using the simple single shooting algorithm implemented here.

3.4 Stability

The knowledge of how natural flow reacts to unaccounted perturbations in the vicinity of a specific periodic orbit is a critical insight that furthers the understanding of the CR3BP. This notion poses the need to define a new concept, known as stability, that aids in the determination of the boundedness of a periodic orbit. In this sense, stability is defined as a means to quantify the likelihood that a periodic orbit will

remain bounded over a fixed time interval. This stability information is directly obtained from the STM, defined in Section 3.2. Recall that the STM maps variations in the final state due to perturbations in the initial state. In particular, the STM after exactly one revolution of a periodic orbit, defined as the monodromy matrix $\Phi(\mathbb{P} + t_0, t_0)$ or $\Phi(\mathbb{P}, 0)$, contains orbit specific stability information.

The eigenvalues of the monodromy matrix classify the stability of a periodic orbit. As the STM is a discrete, linear mapping of variations, the boundary that distinguishes stable and unstable behavior is the unit circle in the complex plane. Consequently, there are three classifications for the linear stability of each eigenvalue: stable, marginally stable, and unstable.

- **Stable:** If all $|\lambda_i| < 1$, the orbit is stable and perturbations will asymptotically approach the reference periodic orbit over time.
- **Marginally Stable:** If all $|\lambda_i| = 1$, the orbit is marginally stable and perturbations remain bounded around the reference periodic orbit.
- **Unstable:** If all $|\lambda_i| > 1$, the orbit is unstable and perturbations will tend to depart the vicinity of the reference periodic orbit over time.

As the monodromy matrix is a 6×6 matrix, the stability of the periodic orbit is governed by the collection of all six eigenvalues. For example, if the monodromy matrix has a single $|\lambda_i| > 1$, the periodic orbit is considered unstable. Additionally, because the CR3BP system as a whole is time-invariant, Lyapunov's Theorem is applicable,

Theorem 3.4.1 (Lyapunov's Theorem) *If λ is an eigenvalue of the monodromy matrix $\Phi(\mathbb{P}, 0)$ of a time invariant system, then λ^{-1} is also an eigenvalue, with the same structure of elementary divisors. [14]*

For truly periodic solutions, one of the eigenvalues of the monodromy matrix is purely real and has unit magnitude. As a consequence of Theorem 3.4.1, another eigenvalue is also real and equal to one. This pair of unit eigenvalues is commonly referred to as

the trivial pair. The other two pairs of eigenvalues may be real or complex. A pair of complex eigenvalues exist as conjugates and takes the form $a \pm bi$. Consequently, the values of the two non-trivial pairs of eigenvalues determine the linear stability properties of a periodic orbit.

3.4.1 Stability Index

Another means of representing the stability of a periodic orbit is the stability index ν . The stability index is simply a function of the eigenvalues of the monodromy matrix. However, by leveraging the reciprocity property, a stability index formula is defined which yields an equivalent value for pairs of eigenvalues [15],

$$\nu_i = \frac{1}{2}(\lambda_i + \frac{1}{\lambda_i}) \quad (3.36)$$

For the six eigenvalues, there exist up to three unique stability indices. The stability index corresponding to the trivial pair of eigenvalues always has a value of one. The non-trivial eigenvalues equate to stability indices which offer insight into the orbit stability. If the magnitude of both non-trivial stability indices is less than one, the periodic orbit is stable. For cases where the magnitude of either stability indices is greater than one, the periodic orbit is unstable.

Compressing the stability information contained in the six eigenvalues to two metrics eases the analysis of orbit stability throughout a family. A plot of the stability indices as a function of some family parameter, such as the initial x position, gives immediate insight into the evolution of periodic orbit stability within a family.

3.5 Bifurcations

The analysis of stability across a set of periodic solutions leads to a phenomenon known as bifurcations. For general dynamical systems, bifurcations occur when the structure of the stability characteristics changes form [16]. In the CR3BP, the presence of a bifurcation may indicate the existence of a new family of periodic orbits, which

intersect at the point of bifurcation. At the location where the bifurcation occurs, both families share a common, bifurcating orbit.

3.5.1 Types of Bifurcations

In dynamical systems theory, a variety of bifurcation types exist for various types of stability changes. However, three distinct classifications [17] are of particular interest and leveraged in this research: tangent bifurcations, period doubling bifurcations, and secondary Hopf bifurcations.

- **Tangent Bifurcations:** A tangent bifurcation occurs when a pair of non-trivial eigenvalues meet at unity: $\lambda_i = \lambda_j = +1$. The presence of a tangent bifurcation suggests that a new family of periodic orbits may intersect the bifurcation point, but does not guarantee existence. To determine the structure that results from a tangent bifurcation, three subcategories are defined which designate behavior: cyclic folds, pitchfork, and transcritical bifurcations [18].
- **Period Doubling Bifurcations:** A period doubling bifurcation occurs when a pair of non-trivial eigenvalues meet at the negative of unity: $\lambda_i = \lambda_j = -1$. As the name suggests, the occurrence of a period doubling bifurcation results in a new orbit family, with an orbital period twice as large the original family.
- **Secondary Hopf Bifurcations:** A secondary Hopf bifurcation occurs when a pair of non-trivial eigenvalues meet at a location on the unit circle, other than ± 1 , and depart into the complex plane.

A graphical depiction of these three types of bifurcations in the complex plane is illustrated in Figure 3.7. Using this schematic, along with the stability diagram introduced in Section 3.4.1, bifurcation types and locations within families are quickly identified. The resulting behavior allows for the expansion of study within the CR3BP and leads to the discovery of families of spatial periodic orbits.




Figure 3.7. Types of Bifurcations Resulting from Stability Changes.

3.5.2 Halo Families

The halos are spatial periodic orbit families which originate from tangent bifurcations in the Lyapunov families. For this research, the L_1 halo family is of particular interest. A subset of this family orbits the L_1 libration point. These orbits with low out-of-plane amplitudes are of high interest for mission applications, especially in the Sun-EMB system [19]. A spacecraft located in one of these orbits has continuous contact with an Earth based communication station during daytime. As the family extends, orbit members gradually shift away from the L_1 libration point towards the vicinity of P_2 . A subset of these orbits, known as the near rectilinear halo orbits (NRHOs), have been proposed for numerous applications in the Earth-Moon system [17].

The strategy employed to discover the initial L_1 Lyapunov orbit in Section 3.3.4 is impractical for determining an initial L_1 halo orbit. At the closest L_1 halo orbit to the L_1 libration point, the linearized dynamics no longer provide an adequate representation of the nonlinear dynamics. Additionally, this location remains a mystery without

other knowledge to guide the search for the L_1 halo family. Instead, bifurcations in the L_1 Lyapunov family are studied, which provide a starting point for the L_1 halo family. For every family member in the L_1 Lyapunov family included in Figure 3.6, the stability indices are calculated and displayed as a function of the initial x position of the respective orbit. This relationship is illustrated in Figure 3.8, with indications for the locations of each bifurcation at $\nu = \pm 1$. The bifurcation which branches to




Figure 3.8. Stability Indices of the Sun-EMB L_1 Lyapunov Family.

the L_1 halo family is located at an initial x position of approximately 1.479×10^8 km in the Sun-EMB rotating frame. This orbit, shared by the L_1 Lyapunov and L_1 halo families, acts as the starting point for a continuation scheme.

In Section 3.3.4, the L_1 Lyapunov family was produced using natural parameter continuation in conjunction with a targeting scheme for planar orbits. For a spatial family, Equations 3.32 and 3.35 are utilized to converge initial conditions for

continuation in the x and z directions, respectively. Thus, a variable continuation algorithm is implemented which continues in either one of these directions. To determine which direction is incremented by a fixed step size, the rate of change of x_0 and z_0 is considered,

$$\text{step variable} = \begin{cases} x & \text{if } |x_0^i - x_0^{i-1}| \geq |z_0^i - z_0^{i-1}| \\ z & \text{else} \end{cases}$$

The other direction, along with the initial y velocity, is targeted to produce a spatial periodic orbit. This process is repeated to produce the L_1 halo family. Additionally, the continuation algorithm is automated by determining the step direction, i.e., in the positive or negative direction, from the sign of the respective rate of change. A summary of the continuation algorithm is included in Table 3.1.

Table 3.1. Spatial Periodic Orbit Continuation Algorithm Summary.

Step Variable	Step Direction	Update Equation
x_0	$\text{sign}(x_0^i - x_0^{i-1})$	Equation 3.32
z_0	$\text{sign}(z_0^i - z_0^{i-1})$	Equation 3.35

The halo families are each further divided into two sub-categories: the northern and southern halos. Starting from the bifurcating orbit indicated in Figure 3.8, continuation in the positive z direction leads to the creation of the L_1 northern halos. Inversely, continuing in the negative z direction produces the L_1 southern halos. These two sub-categories are symmetrical in configuration space across the xy -plane. Axial projections of members of the L_1 southern halo family are included in Figure 3.9, along with an isometric view. These projections capture the size and structure this portion of the L_1 southern halo family. In the isometric view of Figure 3.9, each orbit member is colored by the respective Jacobi constant value. Evident by this sub-figure, the Jacobi constant value changes as the family evolves.




Figure 3.9. Axial and Isometric Views of the L_1 Southern Halo Family.

3.6 Invariant Manifolds

Analysis of a single solution offers insight and understanding of an individual trajectory. However, knowledge of the surrounding fundamental behavior provides an extended level of analysis. The concept of invariant manifolds, which govern the local and global behavior of flow around a particular reference, is applied to equilibrium points and periodic orbits. The invariant manifolds of each reference solution provide a representative model of the phase space, which offers information that is leveraged for a variety of mission applications. The exploitation of invariant manifold theory has directly benefited transfer design [20] and station keeping strategies [21].

3.6.1 Invariant Manifold Theory

The n eigenvalues λ_i of matrix $A(t)$ and of the monodromy matrix $\Phi(\mathbb{P}, 0)$ represent the characteristic multipliers of equilibrium points and periodic orbits, respectively. Highlighted in Section 3.4, the characteristic multipliers determine local linear stability and suggest how the natural flow propagates in the vicinity of a reference solution. For the consideration of an equilibrium solution or a periodic solution, three possibilities occur which categorize stability, denoted in Table 3.2. For n distinct

Table 3.2. Stability of Eigenvalues of Equilibrium Points and Periodic Orbits.

Stability	Equilibrium Point	Periodic Orbit
Asymptotically Stable	$\Re(\lambda_i) < 0$	$ \lambda_i < 1$
Unstable	$\Re(\lambda_i) > 0$	$ \lambda_i > 1$
Center	$\Re(\lambda_i) = 0$	$ \lambda_i = 1$

eigenvalues of $A(t)$ or $\Phi(\mathbb{P}, 0)$, let n_S , n_U , and n_C be the count of asymptotically stable, unstable, and oscillatory eigenvalues, respectively,

$$n = n_S + n_U + n_C \quad (3.37)$$

Additionally, each eigenvalue has an associated linearly independent eigenvector in \mathbb{R}^n . The collection of all n eigenvectors completes a set of invariant, linear subspaces that span \mathbb{R}^n . The stable E^S , unstable E^U , and center E^C subspaces, or eigenspaces, provide a basis to represent \mathbb{R}^n .

The existence of a linear eigenspace that represents flow allows the application of manifold theory, with explicit distinctions for stable, unstable, and center manifolds. Consider a general hyperbolic equilibrium solution. The equilibrium solution is hyperbolic if all λ_i possess non-zero real components. In this case, the dimension of the center eigenspace E^C is zero. Here, the Stable Manifold Theorem [22] describes how the stable and unstable eigenspaces are associated to the local invariant manifolds.

Theorem 3.6.1 (Stable Manifold Theorem) *Suppose $\dot{\bar{x}} = \bar{f}(\bar{x})$ possesses a hyperbolic equilibrium point \bar{x}_{eq} . Then, there exist local stable and unstable manifolds $W_{loc}^S(\bar{x}_{eq})$, $W_{loc}^U(\bar{x}_{eq})$ of the same dimension n_S , n_U as that of the eigenspaces E^S and E^U of the linearized system, and tangent to E^S and E^U at \bar{x}_{eq} .*

The local stable and unstable manifolds provide information into the natural flow towards and away from the equilibrium solution, respectively. Each manifold is extended in configuration space to provide a physical representation of the global manifold behavior. A spacecraft or object, placed anywhere on the global stable or unstable manifold, remains on the manifold for all times. This object naturally flows towards the respective equilibrium solution in positive time if placed on the stable manifold. Conversely, an object on the unstable manifold naturally flows towards the equilibrium solution in negative time.

In the event of a non-hyperbolic equilibrium solution, where n_S , n_U , and n_C are all greater than zero, a center manifold exists in addition to the stable and unstable manifolds. In this scenario, the Center Manifold Theorem [22] is applicable.

Theorem 3.6.2 (Center Manifold Theorem) *Let \bar{f} be a vector field on \mathbb{R}^n vanishing at the origin $\bar{f}(\bar{x}_{eq}) = \bar{0}$ and $A = D\bar{f}(\bar{x}_{eq})$. Divide the spectrum of A into stable n_S , unstable n_U , and center n_C components where,*

$$\Re(\lambda) = \begin{cases} < 0 & : \lambda \in n_S \\ = 0 & : \lambda \in n_C \\ > 0 & : \lambda \in n_U \end{cases}$$

Let the generalized eigenspaces be E^S , E^C , and E^U , respectively. Then, there exist stable and unstable manifolds W^S and W^U tangent to E^S and E^U at \bar{x}_{eq} and a center manifold W^C tangent to E^C at \bar{x}_{eq} . The manifolds W^S , W^U , and W^C are all invariant for the flow \bar{f} . The stable and unstable manifolds are unique, but the center manifold need not be.

The presence of one or more center eigenspaces implies the existence of some oscillatory behavior about the equilibrium solution. In the CR3BP, this implication links

to the previously discovered periodic orbits presented in Section 3.3.4 and Section 3.5.2. Additionally, quasi-periodic orbits, which are bounded to a surface, exist in the vicinity of periodic orbits [23].

Both the Stable Manifold Theorem and Center Manifold Theorem are equally valid for fixed points that make up periodic orbits, noting the change in stability bounds outlined in Table 3.2. However, there is an important and non-trivial distinction between manifolds of equilibrium points and periodic orbits. For an equilibrium point, a single trajectory exists which is entirely representative of the structure of either the stable or unstable manifold. For periodic orbits, the manifolds exist as surfaces of natural flow. To view these structures, a series of fixed points along a periodic orbit are used to generate trajectories which lie on the stable and unstable manifold surfaces. The collection of these trajectories, when viewed in configuration space, offers a visual representation of the respective manifold surface that corresponds to the periodic orbit.

3.6.2 Manifolds of Equilibrium Points and Periodic Orbits

Consider the L_1 libration point in the Sun-EMB system. This equilibrium point has a four-dimensional center eigenspace, one-dimensional unstable eigenspace, and one-dimensional stable eigenspace. Proven by Theorem 3.6.2, a stable and unstable manifold exist which are tangent to the stable and unstable eigenspaces at \bar{x}_{L_1} . Consider, for the moment, only the stable manifold W^S . To determine the flow for W^S , the eigenvalues and eigenvectors of $A(\bar{x}_{L_1})$ are determined. The eigenvector corresponding to the stable eigenvalue is selected, and is of the form,

$$\bar{v}_S = [x_S \ y_S \ z_S \ \dot{x}_S \ \dot{y}_S \ \dot{z}_S]^T \quad (3.38)$$

The stable eigenvector in Equation 3.38 is then normalized solely by position,

$$\bar{v}^{W^S} = \frac{\bar{v}_S}{\sqrt{x_S^2 + y_S^2 + z_S^2}} \quad (3.39)$$

rather than by the magnitude of the entire eigenvector, such that a step in the direction of \bar{v}^{W^S} represents a physical distance. This decision allows for an intuitive understanding of a step-off distance from \bar{x}_{L_1} . To produce the stable manifold of the L_1 libration point, a set of initial conditions are determined by taking a step of magnitude d off of \bar{x}_{L_1} in the direction of \bar{v}^{W^S} ,

$$\bar{x}_{S\pm} = \bar{x}_{L_1} \pm d \cdot \bar{v}^{W^S} \quad (3.40)$$

Now, the state represented by $\bar{x}_{S\pm}$ is located a physical distance d away from \bar{x}_{L_1} . This distance is important, as the selected value must be close enough to the libration point to properly approximate the manifold structure. However, the distance must not be too close, as velocity approaches zero as the distance to the libration point decreases. Hence, if $\bar{x}_{S\pm}$ is too close to \bar{x}_{L_1} , an increase in integration steps is required to escape the vicinity of the libration point, which leads to a build up of numerical error. Additionally, the \pm in Equation 3.40 refers to the direction of the normalized eigenvector. Any scalar multiple of \bar{v}_S is an equally valid representation of the stable eigenspace. Consequently, the stable eigenspace is arbitrarily categorized into a positive and negative direction in congruence with the sign of the component x_S . Finally, to produce trajectories that are emulative of the positive and negative stable half manifolds, the states $\bar{x}_{S\pm}$ are integrated in negative time.

The series of equations used to create the stable manifold are directly applicable to the computation of the unstable manifold by simply substituting the unstable eigenvector in place of the stable eigenvector. The only difference pertains to how the states corresponding to the unstable half manifolds, $\bar{x}_{U\pm}$, are propagated. Natural flow along the unstable manifold departs the L_1 libration point. Therefore, the unstable half manifolds are produced by integrating $\bar{x}_{U\pm}$ in positive time.

With knowledge of the eigenspaces, along with Equation 3.39 and 3.40, the stable and unstable manifolds are easily produced, given an appropriate step-off distance d . Generally, an appropriate dimensional value for d is determined empirically and varies with system mass ratio. Koon et al. recommend a step distance order of magnitude of around 10^{-6} non-dimensional units, approximately 150 *km* in the Sun-EMB system

[24]. Using this value as the step-off distance, both the stable and unstable manifolds of the L_1 libration point in the Sun-EMB system are produced and illustrated in Figure 3.10. Clearly, in the immediate vicinity of L_1 , both manifolds are tangent




Figure 3.10. Manifold Structure of the L_1 Libration Point in the Sun-EMB System.

to the respective eigenspaces. As the manifolds are extended, the linear eigenspaces no longer approximate the manifold structures. Continued propagation only furthers this behavior, but allows for visualization of the global manifold structure. However, it is important to remember that this visual is only a two-dimensional projection in configuration space. Each manifold in this planar example is actually a four-dimensional structure.

The identical process is applied to a periodic orbit, but for a series of points along the orbit rather than a single fixed point. In the Sun-EMB system, the orbit which bifurcates the L_1 Lyapunov and L_1 Halo family is selected. Using 60 points, spaced

equidistant around the orbit, individual trajectories which lie on the stable and unstable manifolds are created. The combination of these trajectories, illustrated in Figure 3.11, provide a two-dimensional visual representation of the four-dimensional manifold tubes. The manifolds shown in Figure 3.11 are integrated for approximately 260




Figure 3.11. Manifold Structure of the L_1 Lyapunov-Halo Bifurcating Orbit.

days to produce a visual that captures the appearance of the manifolds in the vicinity of the periodic orbit. Both the stable and unstable manifolds are extended to provide global manifold structures that traverse throughout the Sun-EMB system. With this knowledge, a spacecraft placed anywhere on the surface of the stable manifold will eventually reach this periodic orbit.

On the right side of P_2 in Figure 3.11, the manifolds are not connected. Both manifolds are simply propagated until reaching the x axis. However, the intersection of stable and unstable manifolds does occur in the CR3BP. This type of intersection,

known as a homoclinic connection, refers to a situation where a trajectory on the stable manifold of a periodic orbit connects to a trajectory on the unstable manifold of the same periodic orbit. Alternatively, heteroclinic connections occur when a trajectory on either manifold of a periodic orbit connects to a trajectory on the opposing manifold of a different periodic orbit. Both homoclinic and heteroclinic connections require that the trajectories are connected in position and velocity space.

4. FLOQUET MODE STATION-KEEPING

The ability to maintain a nominal orbit for a specified duration of time is a fundamental requirement of the mission design process. The process of maintaining a nominal orbit, commonly referred to as a reference orbit, is known as station-keeping. A variety of unique methodologies for station-keeping strategies, such as the target point and Floquet mode, have been extensively studied for mission applications. The target point approach attempts to minimize a cost function composed of the corrective maneuvers and state errors by targeting future states [25], [26]. The Floquet mode approach focuses on leveraging natural dynamics to eliminate the unstable directions of flow to achieve bounded motion about a reference orbit [7]. Both methodologies use a heuristic based approach to calculate corrective maneuvers that will, in theory, prevent the spacecraft from departing the reference orbit. However, each method has independent strengths and innate weaknesses, depending on the specific application. Consequently, this investigation seeks to expand the understanding and applicability of the Floquet mode station-keeping implementation.

4.1 Floquet Theory

In Section 3.2, the state transition matrix, which allows for linear mappings of variations in flow, is introduced. Recall that the STM is numerically calculated during integration,

$$\dot{\Phi}(t, t_0) = A(t)\Phi(t, t_0) \quad (4.1)$$

using the Jacobian matrix of the equations of motion. For a periodic orbit, $A(t)$ is a T -periodic matrix, where T is equivalent to the period of the orbit, \mathbb{P} . As a result of the periodicity of $A(t)$, the Floquet Theorem [27] allows for the decomposition of the STM as a function of the eigenstructure at time t .

Theorem 4.1.1 (Floquet Theorem) *If $A(t)$ is a continuous T -periodic matrix, then for all $t \in \mathbb{R}^n$, the STM can be written in the form,*

$$\Phi(t, t_0) = F(t)e^{B(t-t_0)}F^{-1}(0)$$

where $F(t)$ is a nonsingular, differentiable, T -periodic matrix and B is a constant matrix.

Recall that the monodromy matrix is the STM of a periodic orbit at exactly one revolution, i.e., one orbit period \mathbb{P} . Theorem 4.1.1 is applied to develop an expression for the monodromy matrix,

$$\Phi(\mathbb{P}, 0) = F(\mathbb{P})e^{\mathcal{J}\mathbb{P}}F^{-1}(0) \quad (4.2)$$

In Equation 4.2, the Floquet modal matrix $F(\mathbb{P})$ is a periodic function of time, such that $F(\mathbb{P}) = F(0)$. Each column of $F(\mathbb{P})$, or $F(0)$, is an eigenvector of $\Phi(\mathbb{P}, 0)$. Additionally, the matrix $e^{\mathcal{J}\mathbb{P}}$ is a diagonal matrix comprised of the corresponding eigenvalues of the monodromy matrix.

Previously, the monodromy matrix $\Phi(\mathbb{P}, 0)$ of a periodic orbit was found by integrating a set of initial conditions for exactly \mathbb{P} time. However, at any point along the orbit, a different monodromy matrix exists, $\Phi(\mathbb{P} + t_0, t_0)$, that is calculated by integrating for one revolution from the desired point. The periodicity of the orbit ensures that the eigenvalues of the monodromy matrix are constant regardless of the starting point along the orbit [28]. This property allows Equation 4.2 to be expressed for any arbitrary time t ,

$$\Phi(t, 0) = F(t)e^{\mathcal{J}t}F^{-1}(0) \quad (4.3)$$

where $e^{\mathcal{J}t} = e^{\mathcal{J}\mathbb{P}}$. Using the eigenvectors of the monodromy matrix $\Phi(\mathbb{P}, 0)$ to produce $F(0)$, Equation 4.3 is rearranged to easily solve for the Floquet modal matrix,

$$F(t) = \Phi(t, 0)F(0)e^{-\mathcal{J}t} \quad (4.4)$$

at a specific time t using the STM at that point. This relationship eases the computational burden required for the determination of the Floquet modal matrix at a

desired location by removing superfluous numerical integration. Similar to $F(0)$, each column of $F(t)$ is an eigenvector that corresponds to the respective eigenvalue along the diagonal of $e^{\mathcal{J}t}$. These eigenvectors provide a basis of eigenspaces that represent the local linear behavior of natural flow.

4.2 Floquet Mode Control Laws

The relationships provided by Floquet theory lead to the development of the station-keeping control strategy known as Floquet mode. Due to unmodeled perturbations and the unstable nature of the orbits of interest to this investigation, spacecraft will tend to drift from the reference orbit. To avoid this undesirable behavior, the Floquet mode station-keeping strategy attempts to remove the unstable component of flow via a corrective maneuver.

At any point along a trajectory, the state error between the spacecraft and the reference orbit is represented by $\delta\bar{x}(t)$. At this point, the Floquet modal matrix is determined by Equation 4.4. The columns of this matrix provide a basis to represent the natural flow dynamics. Therefore, the state error is represented as a linear combination of the Floquet mode,

$$\delta\bar{x}(t) = \delta\bar{x}_1(t) + \delta\bar{x}_2(t) + \delta\bar{x}_3(t) + \delta\bar{x}_4(t) + \delta\bar{x}_5(t) + \delta\bar{x}_6(t) \quad (4.5)$$

where $\delta\bar{x}_i(t)$ is the projection of $\delta\bar{x}(t)$ along the eigenvector \bar{v}_i . Therefore, a controller is designed such that a corrective burn will eliminate the components of $\delta\bar{x}(t)$ associated with unstable behavior.

4.2.1 Periodic Orbits with a Four-Dimensional Center Subspace

Historically, significant amounts of station-keeping research has involved the investigation into periodic orbits with one unstable, one stable, and four oscillatory eigenspaces [21], [29], [30]. For periodic orbits possessing this eigenstructure, Floquet mode delivers a corrective burn that removes the component of state error along the

unstable eigenspace, arbitrarily assigned to $\delta\bar{x}_1(t)$. The component of state error in the direction of the stable eigenspace is associated with $\delta\bar{x}_2(t)$, and the four remaining projections are associated with center subspaces. A corrective burn law is postulated,

$$\delta\bar{x}(t) + \Delta\bar{V} = \sum_{i=2}^6 \alpha_i \cdot \delta\bar{x}_i(t) \quad (4.6)$$

such that the addition of $\Delta\bar{V}$ serves to place the spacecraft in stable, bounded motion, about the reference orbit [21]. Additionally, each projection of error $\bar{x}_i(t)$ is augmented with a scalar coefficient α_i . Equation 4.6 is rewritten into a matrix form to incorporate the corrective maneuver as,

$$\delta\bar{x}(t) = \begin{bmatrix} \delta\bar{x}_2(t) & \delta\bar{x}_3(t) & \delta\bar{x}_4(t) & \delta\bar{x}_5(t) & \delta\bar{x}_6(t) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & -1 \end{bmatrix}_{6 \times 8} \begin{pmatrix} \alpha_2 \\ \alpha_3 \\ \alpha_4 \\ \alpha_5 \\ \alpha_6 \\ \Delta V_x \\ \Delta V_y \\ \Delta V_z \end{pmatrix} \quad (4.7)$$

which is condensed into a shorthand notation,

$$\delta\bar{x}(t) = F_{6 \times 8}^* \cdot \bar{\alpha}_{8 \times 1}^* \quad (4.8)$$

Equation 4.7 contains eight free variables: the five α_i coefficients and the three directions of the corrective maneuver. The resulting form is an underdetermined set of constraint equations, suggesting that infinitely many solutions exist for $\bar{\alpha}_{8 \times 1}^*$ that satisfy the system of equations.

Due to the infinite solution space of Equation 4.7, the resulting underdetermined system of equations acts as a branching point for various research areas and is applied differently depending on the specific mission application. It is possible to add additional constraints to restrict corrective maneuvers to defined axes or planes [31].

Alternatively, a strategy which zeros certain coefficients α_i relating to a subset of the center subspaces is utilized to achieve formation keeping [32]. For this investigation, single-axis and three-axis control laws are employed for station-keeping of periodic orbits with four-dimensional center subspaces. A control law which delivers corrective maneuvers solely in the x direction is achieved by adding two constraints on the corrective maneuver, $\Delta V_y = \Delta V_z = 0$. These constraints are incorporated into Equation 4.7 to reduce the number of free variables to six, creating a fully constrained system. The new system of equations is contained by the matrix $F_{6 \times 6}^*$, which is a square and invertible matrix. Hence, an x -axis maneuver is directly calculated by,

$$\begin{pmatrix} \alpha_2 \\ \alpha_3 \\ \alpha_4 \\ \alpha_5 \\ \alpha_6 \\ \Delta V_x \end{pmatrix} = \begin{bmatrix} & & & & & 0 \\ & & & & & 0 \\ & & & & & 0 \\ \delta \bar{x}_2(t) & \delta \bar{x}_3(t) & \delta \bar{x}_4(t) & \delta \bar{x}_5(t) & \delta \bar{x}_6(t) & 0 \\ & & & & & -1 \\ & & & & & 0 \\ & & & & & 0 \end{bmatrix}_{6 \times 6}^{-1} \cdot \delta \bar{x}(t) \quad (4.9)$$

similarly written in a shorthand notation,

$$\bar{\alpha}_{6 \times 1}^* = (F_{6 \times 6}^*)^{-1} \cdot \delta \bar{x}(t) \quad (4.10)$$

Evaluating Equation 4.9 results in a unique solution for the single-axis controller. Conversely, if a three-axis controller is desired for a particular application, Equation 4.7 is revisited. While this equation has infinitely many solutions, one solution is identified via a minimum norm solution. Applying this approach, a corrective maneuver with components in all three directions is determined,

$$\bar{\alpha}_{8 \times 1}^* = (F_{6 \times 8}^*)^T \cdot [F_{6 \times 8}^* \cdot (F_{6 \times 8}^*)^T]^{-1} \cdot \delta \bar{x}(t) \quad (4.11)$$

The minimum norm solution returns $\bar{\alpha}_{8 \times 1}^*$ as the solution to $\delta \bar{x}(t) = F_{6 \times 8}^* \cdot \bar{\alpha}_{8 \times 1}^*$ that minimizes the magnitude of $\bar{\alpha}_{8 \times 1}^*$. However, this approach does not guarantee optimal performance, i.e., a minimum $|\Delta \bar{V}|$ cost solution.

4.3 Station-Keeping Simulation

A numerical simulation that adequately models the trajectory of a spacecraft over the time span of a mission is required. In reality, spacecraft operate under the influence of numerous perturbing forces induced from celestial bodies. A few examples of these external forces include solar radiation pressure, gravitational forces, and atmospheric drag. In a numerical simulation, the potential exists to account for the effects of these forces by appropriate modeling in the equations of motion. However, increasing the fidelity of the modeled dynamics escalates the complexity of the simulation and limits the applicability of some dynamical systems theory. Consequently, this research solely focuses on the inclusion of forces contained within the CR3BP, i.e., gravitational influences from the primary celestial bodies. In addition to external perturbations, the simulation must account for limitations imposed by the infrastructure used to support real world missions. Therefore, several mission constraints are included in the simulation that emulate the restrictions that accompany actual mission operations.

4.3.1 Operational Errors

Given a station-keeping mission plan, the trajectory over a series of revolutions is easily modeled in a numerical simulation. The simulation environment has the ability to place a spacecraft in orbit with no error in the initial states, maintain omniscient knowledge of the position and velocity states of the spacecraft, and deliver precise corrective maneuvers. However, for an actual mission, the functional systems which track the spacecraft and perform maneuvers possess inherent error. The accuracy of the injection maneuver and future corrective maneuvers is limited by capabilities of the spacecraft tracking infrastructure and the spacecraft engine. To account for unmodeled operational errors, three classifications of error are introduced and added to the simulation environment: orbit injection error, spacecraft tracking error, and corrective maneuver error.

- **Orbit Injection Error:** This investigation focuses on the station-keeping portion of a mission plan. However, in order for a spacecraft to reach the desired orbit, a transfer trajectory is designed to place the spacecraft in orbit. Upon reaching the desired orbit, an orbit injection maneuver is executed which aims to match the nominal velocity. Of course, this injection maneuver is not free of error, resulting in an initial orbit state that differs from the nominal reference state.
- **Spacecraft Tracking Error:** Once a spacecraft is placed along a reference orbit, the position and velocity states are tracked at regular intervals, also known as an orbit update. Again, due to the limitations of tracking instruments, the measured states are deviated from the true states of the spacecraft. Subsequently, the deviated tracking data is utilized to calculate corrective maneuvers.
- **Corrective Maneuver Error:** With the passage of time, the spacecraft will naturally deviate from the reference orbit, due to the aforementioned unmodeled perturbations and other mission operation errors. As the deviations increase in magnitude, corrective maneuvers are required to keep the spacecraft in bounded motion about the reference orbit. Using the measured tracking data, corrective maneuvers are calculated and applied to the current state. However, due to engine thrusting limitations, the executed maneuver differs from the desired maneuver in magnitude and direction.

All three operational errors are accounted for in the numerical simulation via the introduction of Gaussian error distributions. The standard deviation of the error distribution is determined by analyzing the accuracy of the instruments used on a mission-by-mission basis. For a baseline simulation, a set of standard operational errors are outlined in Table 4.1. Using this set or a similar set of standard deviations, a perturbation is readily sampled and applied to the nominal values to appropriately model the inaccuracy in real world applications. Essentially, the deterministic quantities are converted into nominal-centered Gaussian distributions. For injection

Table 4.1. Standard Operational Errors.

Error Type	Position [km]			Velocity [mm/s]				
	σ_x	σ_y	σ_z	$\sigma_{\dot{x}}$	$\sigma_{\dot{y}}$	$\sigma_{\dot{z}}$	σ_{mag}	σ_{dir}
Orbit Injection	1.0	2.0	5.0	1.0	1.0	3.0	-	-
Spacecraft Tracking	1.0	2.0	5.0	1.0	1.0	3.0	-	-
Corrective Maneuver	-	-	-	-	-	-	1.0%	1.0°

error, each state that makes up the nominal initial conditions is independently perturbed. The position and velocity components of each tracked state throughout the mission are perturbed in a similar fashion. Finally, the nominal corrective maneuver is isolated into two components: the magnitude and direction of the maneuver. The magnitude is perturbed by a percentage of the nominal maneuver magnitude, while the direction vector is perturbed by an offset angle.

4.3.2 Mission Constraints

In addition to mission operation errors, constraints are imposed which mimic realistic mission capabilities related to the spacecraft operational systems, ground tracking infrastructure, and mission requirements. The following constraints are included to emulate these limitations:

- **Minimum Time Between Maneuvers:** A minimum time between successive maneuvers is included as a means to model the amount of time required to obtain sufficient post maneuver orbit determination data, and to allow for the operation of scientific instruments without maneuver induced perturbations.
- **Spacecraft Tracking Rate:** Limitations in the availability of ground based tracking stations restricts the frequency of updated spacecraft tracking data.

The spacecraft states are returned at a fixed tracking rate, e.g., once every 24 hour period.

- **Minimum Maneuver Magnitude:** A restriction on the minimum allowable maneuver magnitude is included to prevent maneuvers that are not realistically feasible by the spacecraft propulsion system.
- **Divergence in Position Error Restriction:** Corrective maneuvers are only executed when the spacecraft is actively departing the reference orbit in configuration space. This constraint prevents the execution of wasteful corrective maneuvers while the spacecraft is approaching the reference orbit.

A list of quantitative values for each mission constraint is contained in Table 4.2. This table serves as a set of reference constraints used for baseline simulations and analysis. Once all constraints for the given simulation are satisfied, a corrective maneuver is implemented.

Table 4.2. Standard Mission Constraints.

Constraint Type	Value
Minimum Time Between Maneuvers	30 days
Spacecraft Tracking Rate	1.0 days
Minimum Maneuver Magnitude	25.0 mm/s
Position Error Divergence	True

4.4 Station-Keeping Analysis

The ability to model the trajectory of a spacecraft allows for rapid testing of various mission scenarios. Each simulation, or series of simulations, is analyzed to obtain information related to the propulsive costs and other performance metrics, subject to

the specified operational errors and mission constraints. The information obtained from this analysis is vital for determining the feasibility of a Floquet mode station-keeping approach. Consequently, the importance of obtaining relevant benchmarks facilitates the need for clearly defined performance metrics supported by statistical analysis. The overall metrics are subsequently used to determine the relative performance of differing implementations.

4.4.1 Monte Carlo

In Section 4.3.1, several operational errors are introduced in the form of random perturbations. The addition of random error sampling converts the simulation to a stochastic process, i.e., two simulations with identical parameters generally yield different results. To obtain an accurate estimate of any performance metric, a series of simulations is required. The process of repeated simulations to determine an estimate for the mean of a metric is known as a Monte Carlo simulation.

The number of trials n , or simulations, is an unknown factor when performing a Monte Carlo simulation. If the number of trials is too small, the estimated mean is generally a poor representation of the true mean. Conversely, performing too many trials leads to excessive computational effort that does not significantly improve the estimate of the mean. Often, a Monte Carlo simulation is performed for an arbitrarily large number of simulations until the sample mean has *stabilized*, without the inclusion of quantitative evidence to support such results. Rather than aimlessly performing simulations, the Monte Carlo is executed until the estimated mean meets a desired confidence and precision level [29]. A summary of this methodology, which determines the total number of trials required to meet a specified statistical significance, is outlined here.

Consider a series of simulations that each produce a performance metric, with mean $\bar{\mathcal{X}}$ and standard deviation \mathcal{S} . A confidence interval is constructed,

$$\left[\bar{\mathcal{X}} - z^* \frac{\mathcal{S}}{\sqrt{n}}, \bar{\mathcal{X}} + z^* \frac{\mathcal{S}}{\sqrt{n}} \right] \quad (4.12)$$

where z^* is a measure of standard deviations and corresponds to the desired confidence level. Additionally, a limit on the half-width of the confidence interval in Equation 4.12 is imposed using a relative precision tolerance ε . The allowable error is formatted such that the true mean falls within $\bar{\mathcal{X}} \pm (\varepsilon/100)\bar{\mathcal{X}}$. Using this relationship and the half-width of the confidence interval in Equation 4.12, an estimate of the required number of trials is computed [29],

$$n = \left(100 \cdot \frac{z^* \mathcal{S}}{\varepsilon \bar{\mathcal{X}}} \right)^2 \quad (4.13)$$

Unless otherwise stated, a z^* value of approximately 2.58 is employed which corresponds to a 99% confidence interval. For this confidence level, the half-width should not exceed 1% of the estimated mean. Each Monte Carlo is terminated once the total number of simulations exceeds n .

4.4.2 Performance Metrics

To quantify the capabilities of station-keeping algorithms, key performance metrics are identified here. These metrics offer concise representations of the merits of each station-keeping methodology and allow for the direct comparison of different orbital maintenance strategies. For this investigation, two performance metrics are identified and defined: the rate at which a station-keeping strategy successfully maintains the reference orbit, and the total cost of all corrective maneuvers required for station-keeping purposes.

The success of a simulation is a function of the control strategy, celestial system, reference orbit, operational errors, and mission constraints. This research aims to investigate the reliability of various station-keeping strategies over a range of simulation parameters. Consequently, the knowledge of how successfully these strategies maintain the reference orbit is of critical importance. Therefore, conditions which objectively differentiate between failed and successful simulations are defined by applying limitations on three attributes: total maneuver cost, positional deviation from the nominal orbit, and linear divergence of position error. Threshold values for each

attribute are determined via empirical studies and are declared in Table 4.3. For any simulation to be classified as successful, all conditions must hold true. The restriction

Table 4.3. Metrics For Classifying Successful Station-Keeping Simulations.

Attribute	Limit
Total Maneuver Cost	1000 m/s
Positional Error	100000 km
Linear Position Divergence	+5000 km/year

on the maximum maneuver cost is imposed simply to terminate simulations where the maneuver magnitudes become excessively large due to increased state error. Some simulations may recover to the reference orbit via unrealistically high maneuvers that are not practical in real world applications. Additionally, the max positional error denotes a cutoff value beyond which the reference orbit generally becomes unrecoverable. The final constraint on the max linear position divergence is included to require that station-keeping simulations are sustainable over increased time spans. Occasionally, for the given simulation time interval, e.g., five years, both the maneuver cost and position error constraints are satisfied. However, a closer inspection of the actual trajectory reveals divergent behavior that will lead to a station-keeping failure if the simulation is extended. A truly successful algorithm will maintain the reference orbit beyond the original mission life span.

The total station-keeping cost addresses the practicality of a station-keeping strategy. During mission planning, propulsion budgets are allocated for each mission phase. A desirable control strategy serves to minimize the required cost for the station-keeping portion, potentially yielding propellant to other mission operations. The total cost for a station-keeping simulation is defined such that,

$$|\Delta \bar{V}_{Total}| = \sum_{i=1}^m |\Delta \bar{V}_i| \quad (4.14)$$

where $|\Delta\bar{V}_i|$ is the magnitude of the i^{th} corrective maneuver. Subsequently, the total cost from each simulation is utilized in Equation 4.13 to determine the required number of trials to achieve a statistically significant representation of the expected total station-keeping cost. At this point, the rate of success of the station-keeping algorithm is also recorded. Intuitively, an ideal station-keeping algorithm requires minimal total maneuver cost with maximal success probability. The expected total maneuver cost should be reasonable for the specific mission scenario. However, solely focusing on the minimization of propellant usage neglects operational success. Therefore, the success rate is considered an equally important metric for assessing the quality of each station-keeping methodology.

4.5 Floquet Mode Control Results

Floquet mode is implemented to maintain spacecraft in spatial periodic orbits. Initially, baseline simulations are analyzed and performance is compared relative to prior literature. This preliminary analysis serves as a reference point, which is then extended via the exploration of operational errors and control strategies.

4.5.1 Single and Multi-Axis Control of Spatial Orbits

Certain spatial periodic orbits offer geometries that are highly desired for mission applications. In particular, a low out-of-plane amplitude L_1 halo orbit is selected which has been extensively studied in previous work [21], [31]. The initial conditions of the exact orbit used in this investigation are included in Table 4.4. All initial conditions are supplied in non-dimensional units for consistency with the non-dimensional equations of motion. Each condition is easily transitioned to dimensional units using the characteristic quantities for the Sun-EMB system listed in Table 2.2.

The initial conditions of this periodic orbit are directly propagated using the nonlinear equations of motion for one orbit period to produce geometric views of the orbit. The planar projections and an isometric view of this L_1 halo orbit are illustrated

Table 4.4. Parameters for the Selected L_1 Halo Orbit.

Component	Value [nd]
x_0	0.98885243515679
y_0	0
z_0	-0.0015
\dot{x}_0	0
\dot{y}_0	0.00918524634855
\dot{z}_0	0
\mathbb{P}	3.05847982227289

in Figure 4.1. Clearly, the resulting trajectory forms a periodic structure within the specified targeting tolerance. However, due to numerical error, if the initial conditions are propagated for extended periods of time, the trajectory will eventually depart the periodic solution. For this orbit, the periodic geometry is lost after approximately four revolutions. In addition to operational errors, the station-keeping strategies will have to account for the error inherent in numerical processes.

A simulation is designed to maintain a spacecraft in this orbit for 10 revolutions, which is approximately 4.87 *years*. For a baseline simulation, the operational errors and mission constraints listed in Tables 4.1 and 4.2 are employed, respectively. Following the process outlined in Section 4.4.1, the required maintenance cost is determined using two control strategies: single-axis control and three-axis control. For single-axis control, all corrective maneuvers are constrained to only apply changes in velocity in the x direction. The magnitude of the corrective maneuver is calculated using the control law defined by Equation 4.10. The resulting total corrective maneuver costs are depicted in Figure 4.2. Each data point refers to the required $|\Delta \bar{V}_{Total}|$ to maintain the reference orbit for a single simulation. The blue line corresponds to the moving average of the total corrective maneuver cost for all simulations that are




Figure 4.1. Projections of the Nominal L_1 Southern Halo Orbit.




Figure 4.2. Baseline Maintenance Costs for a L_1 Halo Orbit in the Sun-EMB System with x -Axis Control.

performed up to that point. For this Monte Carlo, 613 simulations were performed to achieve a stabilized expected corrective maneuver cost. The Monte Carlo concludes that, with 99% confidence, the true cost to maintain the reference orbit with x -axis control is within 1% of 2.798 m/s . Additionally, every simulation in this trial met the requirements listed in Table 4.3, suggesting that the x -axis control strategy is able to maintain the reference orbit without failure for this set of simulation parameters.

Alternatively, three-axis control allows corrective maneuvers to modify all three velocity components. Equation 4.11 defines a control law for three-axis maneuvers using a minimum norm solution. Similar to the analysis for single-axis control, a Monte Carlo is performed to determine an interval for the true cost to maintain the same reference orbit with the same parameter sets. Maintenance costs for individual simulations are illustrated in Figure 4.3, along with the moving average of the expected cost. The Monte Carlo for three-axis control required 372 simulations to stabilize the expected corrective maneuver cost within the desired confidence level and relative precision. For this set of operational errors and mission constraints, the true cost to maintain the reference orbit with three-axis control is within 1% of 0.720 m/s ,




Figure 4.3. Baseline Maintenance Costs for a L_1 Halo Orbit in the Sun-EMB System with Three-Axis Control.

at a confidence level of 99%. It is noteworthy that this methodology successfully maintained the reference orbit for 100% of the simulations performed.

The station-keeping costs determined in this investigation for single and three-axis control are consistent with prior works that seek to maintain similar halo orbits for comparable lengths of time [21], [30]. An initial observation simply concludes that both methods successfully maintain the reference orbit, while three-axis does so at a reduced cost. However, other factors are considered for a more comprehensive analysis. One strength of x -axis control is the lack of orientation requirements from one maneuver to the next. Given a spacecraft oriented with the thrust axis aligned along the x direction, the orientation required for a corrective maneuver is always satisfied. The added complexity of providing appropriate orientation for three-axis maneuvers is non-trivial. Consequently, for specific mission applications, the x -axis station-keeping approach has the potential to be satisfactory. On the other hand, the three-axis strategy consistently maintains the reference orbit with a reduction in the total corrective maneuver cost. Therefore, a mission which requires a minimal

station-keeping budget, but has the capacity to meet orientation demands, would benefit from a three-axis control strategy.

4.5.2 Divergence from the Reference Orbit

While the propulsive requirement is an integral factor when considering a station-keeping strategy, other factors such as success rate and orbit sustainability are equally important. Consequently, the concept of divergence is introduced, which defines a metric to represent how the spacecraft trends towards or away from the reference orbit over the course of a simulation. The spacecraft position error at each state along the trajectory has been previously employed in the Floquet mode station-keeping formulation. Here, a linear trendline is best fit to these position error magnitudes, which indicates if the spacecraft is diverging from the reference orbit. This trendline gives a rate of linear position divergence over a fixed time span.

For the same Monte Carlo performed in Section 4.5.1, the expected linear divergence is calculated to obtain an estimate of the true divergence over time. The linear divergence for each simulation and moving average is illustrated in Figure 4.4. The




Figure 4.4. Linear Divergence from a L_1 Halo Orbit in the Sun-EMB System with x -Axis Control.

data collected from these trials suggests a divergence rate of 24.159 km/year for the x -axis controller. This positive divergence rate suggests that, on average, the spacecraft will drift away from the reference orbit over the course of a simulation. As an example, a simulation with a linear divergence rate approximately equal to the expected divergence rate is selected from Figure 4.4. The position error history for this simulation and a linear trendline are depicted in Figure 4.5. The oscillatory motion of the




Figure 4.5. Position Error of a Spacecraft with x -Axis Control.

position error illustrates how the x -axis Floquet mode controller maintains bounded motion about the reference orbit. However, visual observation of the position error and the linear trend clearly indicates divergence from the desired orbit. Over time, as the deviation in position continues to increase, maintaining the orbit with an x -axis control approach has the potential to become unsustainable. However, additional analysis is required to properly determine conclusions on the sustainability of the x -axis controller. Meanwhile, the three-axis Floquet mode controller maintains the spacecraft in consistently bounded motion about the reference orbit, with an order of magnitude reduction in divergence rate. Evident by Figure 4.6, a spacecraft utilizing the three-axis controller is expected to diverge at a rate of 1.137 km/year . Relative

to the x -axis controller, the three-axis controller reduces drift from the reference orbit at a reduction in propulsive cost.




Figure 4.6. Linear Divergence from a L_1 Halo Orbit in the Sun-EMB System with Three-Axis Control.

This preliminary analysis suggests that a three-axis controller exceeds the performance of the x -axis controller for this halo orbit in terms of total corrective maneuver cost and the divergence from the reference orbit. This conjecture relies on the capability to deliver multi-directional corrective maneuvers. Therefore, the choice of controller is still dependent on specific mission requirements. Despite the higher cost, the x -axis controller is still desirable due to the simplification of orientation requirements. Additionally, when considering divergence rate, the size of the reference orbit is an important factor. This halo orbit in the Sun-EMB system has in-plane and out-of-plane amplitudes of approximately 690000 km and 180000 km , respectively. Accordingly, a divergence rate of approximately 24 $km/year$ may be acceptable for certain mission scenarios.

4.5.3 Expanded Operational Errors and Mission Constraints

A primary focus of this investigation centers on the robustness of station-keeping algorithms. In this context, a robust station-keeping strategy continues to maintain the spacecraft in the vicinity of the reference orbit despite unexpected operational errors and changes in mission constraints. In Sections 4.5.1 and 4.5.2, baseline performance metrics for x -axis and three-axis control strategies were determined using the operational parameters listed in Tables 4.1 and 4.2. This preliminary investigation highlighted the differences in performance of both control strategies in an environment with relatively small operational errors and mission constraints that allowed for frequent corrective maneuvers. However, the narrow focus of this analysis prevents general conclusions regarding the overall performance of both control strategies. In this section, an extended analysis is performed which restricts the frequency of corrective maneuvers and considers a variety of increased injection errors in order to provide a comprehensive assessment of the capabilities of both station-keeping strategies.

The baseline problem setup allowed for corrective maneuvers to be performed every 30 *days*. This rate is consistent with previous literature, allowing the direct comparison of preliminary results [21], [31]. In addition, other station-keeping strategies have attempted to maintain similar L_1 southern halo orbits in the Earth-Moon system via corrective maneuvers solely at x -axis crossings [33]. To emulate a similar restriction, the corrective maneuver cooldown is extended to 90 days, effectively limiting the number of corrective maneuvers to two per revolution. The updated minimum time between consecutive maneuvers is subsequently imposed with all other mission constraints listed in Table 4.2. Additionally, the injection errors are significantly varied in order to investigate the sensitivity of each station-keeping controller in response to unexpected errors in the initial state. Rather than a constant $1\text{-}\sigma$ error for each state component, a range of $1\text{-}\sigma$ position errors between 1.0 *km* and 2000 *km*, and a range of $1\text{-}\sigma$ velocity errors between 0.001 *m/s* and 2.0 *m/s* are employed. These

injection error ranges are then implemented in conjunction with the tracking and corrective operational errors previously defined in Table 4.1.

For this updated mission scenario, two performance metrics are of interest: the total expected propulsive cost and the expected rate of success to maintain the spacecraft in orbit. A successful station-keeping simulation keeps the spacecraft in bounded motion about the reference orbit. The conditions for success are defined by the parameter limits listed in Table 4.3. Additionally, the expected total cost is the summation of the magnitudes of all corrective maneuvers required to keep the spacecraft in orbit. However, only the results of successful station-keeping simulations are included in the expected cost. Failed simulations generally result in unrealistically large corrective maneuvers due to significant state error. Including the excessive maneuvers results in a skewed expected total cost, which offers little insight into the performance of the controller. By excluding outliers, the resulting expected cost represents the propulsive requirements on the condition that the simulation is successful.

A series of Monte Carlo runs are performed to determine a comprehensive depiction of the performance of the x -axis and three-axis controllers. For each Monte Carlo simulation, a position error and velocity error permutation is sampled and independently applied as a perturbation to the appropriate injection state component. Simulations are then performed until the expected cost stabilizes. To meet this condition, Equation 4.13 is employed to determine the required number of trials for a 99% confidence interval and a relative precision of 1%. Once terminal, the expected cost to maintain the orbit and the rate of success are recorded. This process is repeated over the range of all position and velocity errors for both controllers. However, each Monte Carlo simulation requires a significant amount of computational effort. For an accurate yet computationally feasible analysis, the position and velocity errors are discretized into 100 km and 0.1 m/s increments, respectively. This selection of errors results in 441 injection error permutations, each requiring a Monte Carlo simulation to determine performance metrics.

The resulting performance maps for the rate of success and expected total propulsive cost of both controllers are illustrated in Figure 4.7. Each subfigure contains data points for all Monte Carlo simulations, colored by the respective performance metric. For visualization purposes, the gaps between the discretized points are colored using two-dimensional linear interpolation. The resulting plots effectively display how




Figure 4.7. Success Rate and Expected Cost of the x -Axis and Three-Axis Controllers for Extended Injection Error.

the success rate and expected cost vary as a function of the position and velocity

injection errors. The x -axis controller successfully maintains the spacecraft in the desired orbit between 23.4% and 100% of attempted simulations, while the three-axis controller ranges between 72.2% and 100%. Both controllers show a gradual decrease in the success rate as the position and velocity errors increase. However, clearly, the ability to maintain the orbit using the x -axis controller diminishes at a significantly higher rate than the three-axis controller. Meanwhile, the expected cost is observed to proportionally increase with either injection error type. Relative to the three-axis controller, the x -axis controller requires less propellant to maintain the orbit for successful simulations, with the disparity increasing as injection error increases. To clearly highlight the differences in both performance metrics, a direct comparison of three-axis controller relative to x -axis controller is illustrated in Figure 4.8. The




Figure 4.8. Three-Axis Controller Performance Relative to the x -Axis Controller.

relative success rate shown in Figure 4.8(a) increases rapidly with the injection error until reaching a plateau around 40%. Likewise, the relative propulsive cost gradually increases throughout the range of tested injection errors. This relationship suggests that the additional cost delivered from the three-axis formulation is related to higher rates of success.

The previous analysis in Section 4.5.1 for baseline controller performance suggested that both x -axis and three-axis controllers reliably maintain the reference orbit, noting that the three-axis controller does so at a reduced cost. Meanwhile, this section imposes a restriction on the frequency of corrective maneuvers and permutes over a series of increased injection errors. A clear observation of this analysis pertains to the significant reduction in the rate at which both controllers meet the defined success criteria. As a result, mission plans that limit the number of corrective maneuvers or desire station-keeping strategies robust to operational errors, potentially require a reevaluation of how the control strategy is implemented.

5. HYBRID CONTROLLER DESIGN

The exploration of station-keeping techniques in literature has uncovered a variety of methodologies which share the common goal of maintaining a spacecraft in orbit. The various station-keeping control laws aim to correct for unaccounted errors, generally via derived heuristics that leverage the dynamics of the system. Inherently, shortcomings or failure points exist when exploring the versatility of the different controllers. In Chapter 4, Floquet mode was introduced and the performance of the controller was demonstrated. For the scope of that analysis, the three-axis Floquet mode controller generally maintained bounded motion about a reference orbit. However, as mission constraints and operational errors were increased, performance degradation was exhibited, resulting in higher failure rates. In these events, the three-axis controller provided inadequate corrective maneuvers to avoid the rapid divergence of the spacecraft from the reference orbit. This observation results in the manifestation of a requirement to develop a robust controller which delivers corrective maneuvers that avoid undesirable divergence behavior. To that end, a new controller is proposed that leverages the three-axis Floquet mode control law, augmented with a corrective maneuver perturbation determined by a reinforcement learning (RL) agent. This hybrid controller (RLFM) exploits the highly successful nature of the three-axis Floquet mode controller for small state errors, while enhancing the ability of the controller to recover and maintain the reference orbit in the presence of significant operational errors and restrictive mission constraints.

5.1 Reinforcement Learning Background

Reinforcement learning is a branch of machine learning in which an agent, i.e., a controller, is trained to accomplish a defined task or set of tasks. The agent “learns”

via a training process in which the agent directly interacts with an environment. The process of learning through interaction, a trial and error based approach, is conceptually similar to the human or animal learning process. For instance, Sutton discusses a gazelle calf struggling to walk moments after birth, but within a short period of time, the calf manages to learn how to run [34]. In this scenario, the calf experiments with different actions that result in favorable, neutral, and detrimental outcomes. Instinctively, negative behavior is discouraged, e.g., falling over when attempting to stand, and positive behavior is encouraged. Via this action-reward cycle, the calf eventually discovers the proper combination of actions to successfully run.

5.1.1 Introduction to Reinforcement Learning

The primary components of the reinforcement learning structure are the *agent* and the *environment* in which the agent operates. The agent is viewed as a controller which relates a set of inputs, i.e., *observations*, to a set of *actions*. The observations supplied to the agent are either a partial or full view of the current environment state. Once the actions are applied to the environment, the state is propagated using the dynamics of the environment. At the updated state, a *reward* is determined to inform the agent of the impact of the chosen actions, and to represent the quality of the current state. This cyclic procedure is iterated from an initial state until a terminal condition is reached, as represented by the schematic in Figure 5.1 [34], [35]. The completion of this process from the initial to final state is known as an *episode*. Once an episode terminates, the environment is reset to a new initial state, and the process is repeated. Over a series of episodes, the *policy* is updated in order to maximize the cumulative reward, i.e., the *return*, per episode. The policy dictates the actions that are determined, and is often viewed as the “brain” of the agent. For clarity, an overview of each of the aforementioned RL components and notations is outlined below.




Figure 5.1. The Agent-Environment Interaction Process [34], [35].

Observation Space

The observation space, which contains observations o_t , captures the agent's perspective of the current state of the environment. On the other hand, the state space, \mathcal{S} , containing states s_t , refers to a complete depiction of the environment. The observation space potentially encompasses a smaller subset of the full state. In this event, the environment is considered partially observed. Conversely, when the observation space captures all knowledge of the state, the environment is fully observed.

Action Space

The action space, \mathcal{A} , consists of the set of available actions that an agent is able to perform that inevitably affect the environment. The action space is divided into two subcategories: discrete and continuous action spaces. Discrete action spaces are employed in environments limited to a finite number of actions. For continuous action spaces, the actions are represented as real-valued vectors. The type of action space is generally driven by the environment. For instance, reinforcement learning frameworks utilizing discrete action spaces have achieved super-human performance in strategy games such as chess and shogi [36]. Likewise, continuous action space controllers have been incorporated for environments requiring precise control for transfer trajectory design [35].

Policy

The actions are determined as a function of the current state using the agent's policy. Essentially, the policy contains rules which govern actions, and is either deterministic or stochastic, denoted by μ and π , respectively. Furthermore, the policy is parameterized by a set of weights and biases which are refined during the training process. Therefore, the actions determined by the policy are represented by the policy type and parameter set,

$$a_t = \mu_\theta(s_t) \quad (5.1)$$

$$a_t \sim \pi_\theta(a_t|s_t) \quad (5.2)$$

where θ denotes the parameterization. For stochastic policies, each action is sampled from a Gaussian distribution denoted by a mean value and standard deviation. The randomness aids in the exploration of new actions during the training process. Once training concludes, the stochastic policies are generally converted to deterministic variants where the mean action value is applied to the environment.

Return

The reward is defined as a function of the current state and is potentially extended to include the current action and/or the next state. Each reward, r_t , determined at the various states throughout an episode, is culminated to calculate a return which is indicative of episodic performance. Several methods are used to calculate the return; commonly, a finite-horizon undiscounted return and an infinite-horizon discounted return are employed. The finite-horizon undiscounted return is simply the summation of all rewards over a fixed number of state transitions. Conversely, the infinite-horizon discounted return is the sum of all rewards, each augmented by a time decaying discount factor γ ,

$$\mathcal{R}(\tau) = \sum_{t=0}^{\infty} \gamma^t r_t \quad (5.3)$$

where τ denotes an episode. The discount factor γ is bounded in the range $[0, 1]$ and controls the relative influence of current rewards versus delayed rewards. In essence, the discount factor is an adjustable, problem specific parameter, that allows intuitive weighting of the ratio between immediate and future rewards.

Value Functions

The return directly relates to the concept of value functions. Within most reinforcement learning algorithms, value functions are incorporated as means of estimating the expected return. One type of value function, the state value function,

$$V^\pi(s_t) = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1}(s_k) \right] \quad (5.4)$$

estimates the expected return from a specific state, assuming the current policy is followed indefinitely. Additionally, the state-action value function is defined,

$$Q^\pi(s_t, a_t) = \mathbb{E}_\pi \left[\sum_{k=0}^{\infty} \gamma^k r_{t+k+1}(s_k, a_k) \right] \quad (5.5)$$

which similarly estimates the expected return. However, the state-action value function estimates the expected return by incorporating an arbitrary action taken at the current state, which is not necessarily from the current policy. Thereafter, the current policy is followed indefinitely.

Markov Decision Process

A Markov Decision Process (MDP) is defined as a discrete time stochastic control process, and represents the mathematical formulation of the reinforcement learning problem [37]. Each MDP is formulated as a tuple $(\mathcal{S}, \mathcal{A}, \mathcal{T}, \mathcal{R}, \gamma)$, where \mathcal{T} is the state-transition probability function. Within the MDP, the environment is fully observable, i.e., the observation space and state space are equivalent. Additionally, the Markov property, which suggests that future states are only reliant on the current state, rather than the prior time dependent state history, must hold true at each state [37].

5.1.2 Soft Actor-Critic Algorithm

Within the domain of off-policy reinforcement learning, numerous algorithms exist which uniquely attempt to develop an ideal policy that maximizes the expected return. Most off-policy approaches employ Q-Learning, where an approximator represents the optimal action-value function. As the optimization occurs off-policy, sample data collected throughout the training process is used for policy updates. Conversely, on-policy approaches rely on policy optimization techniques in which the policy parametrization, θ , is optimized using the current version of the policy. This investigation leverages soft actor-critic (SAC), an off-policy maximum entropy reinforcement learning algorithm [38]. This algorithm is relatively similar to an earlier algorithm, deep deterministic policy gradient (DDPG) [39]. Both algorithms leverage the sampling-efficiency innate to off-policy learning while mitigating hyperparameter sensitivity and instability. However, the SAC formulation incorporates entropy maximization and adjusts the policy iteration process, which ultimately enhances convergence stability and provides demonstrably improved relative performance [38].

The most general reinforcement learning framework optimizes the policy by maximizing the expected return. For the maximum entropy framework, the objective function is augmented with the expected entropy of the policy [40]. The resulting objective function is further modified for the infinite horizon discounted problem,

$$J(\pi) = \sum_{t=0}^{\infty} \mathbb{E}_{(s_t, a_t) \sim \rho_{\pi}} \left[\sum_{k=t}^{\infty} \gamma^{k-t} \mathbb{E}_{s_t \sim \mathcal{T}, a_t \sim \pi} \{ r(s_t, a_t) + \alpha \mathcal{H}[\pi(\cdot | s_t)] | s_t, a_t \} \right] \quad (5.6)$$

where ρ_{π} corresponds to the probability of the state-action tuple [41]. Additionally, the temperature parameter, denoted α , controls the stochasticity of the optimal policy by weighting the entropy relative to the reward. The objective function in Equation 5.6 essentially weights the discounted expected return and entropy by the corresponding state-action probability for the current policy.

Starting from this maximum entropy framework, the SAC algorithm is derived by first applying the soft policy iteration method. Thereafter, several approximators are defined which are readily implemented in a practical algorithm that is computationally

feasible. A summary of the derivation is included in this investigation, but is described in detail in the original work [38]. Initially, the value of the policy is determined according to Equation 5.6, where the soft Q-value function is deduced by iterative application of a modified Bellman backup operator, \mathcal{B}^π ,

$$\mathcal{B}^\pi Q^\pi(s_t, a_t) = r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim \mathcal{T}}[V^\pi(s_{t+1})] \quad (5.7)$$

with a corresponding soft state value function,

$$V^\pi(s_t) = \mathbb{E}_{a_t \sim \pi} [Q^\pi(s_t, a_t) - \log \pi(a_t | s_t)] \quad (5.8)$$

Additionally, during the policy improvement step, the policy is updated towards the exponential of the new Q-function, which guarantees policy improvement in terms of the soft value [38]. The Kullback-Leibler divergence, which relates the differences between two probability distributions, is employed to project the new policy,

$$\pi_{new} = \arg \min_{\pi' \in \Pi} D_{KL} \left(\pi'(\cdot | s_t) \middle| \frac{\exp Q^{\pi_{old}}(s_t, \cdot)}{Z^{\pi_{old}}(s_t)} \right) \quad (5.9)$$

where Π is a set of policies that restricts the policy domain and $Z^{\pi_{old}}(s_t)$ serves to normalize the distribution. The projection is an essential step to account for the added constraint on the policy domain, and is selected as a matter of convenience [38]. This process of evaluating the policy and improving the policy is iterated, and generally converges to an optimal policy. However, the formulation is limited to the tabular form, preventing direct implementation for continuous domains. This limitation facilitates the need to develop several function approximators for the Q-function and the policy. Additionally, rather than switching between policy evaluation and improvement, each approximator is optimized using stochastic gradient descent.

Three approximators are defined specifically for the SAC formulation: the state value function $V_\psi(s_t)$, the soft Q-function $Q_\theta(s_t, a_t)$, and the policy $\pi_\phi(s_t, a_t)$. The parameters of each network-based approximator are denoted by ψ , θ , and ϕ , respectively. The state value function is used to estimate the soft value. During training,

this approximator seeks to minimize the squared residual error, specified by the objective function,

$$J_V(\psi) = \mathbb{E}_{s_t \sim \mathcal{D}} \left[\frac{1}{2} (V_\psi(s_t)) - \mathbb{E}_{a_t \sim \pi_\phi} [Q_\theta(s_t, a_t) - \log \pi_\phi(a_t | s_t)]^2 \right] \quad (5.10)$$

where \mathcal{D} signifies the distribution of sampled states and actions. Meanwhile, the soft Q-function network parameters are optimized in order to minimize the soft Bellman residual,

$$J_Q(\theta) = \mathbb{E}_{(s_t, a_t) \sim \mathcal{D}} \left[\frac{1}{2} \left(Q_\theta(s_t, a_t) - \hat{Q}(s_t, a_t) \right)^2 \right] \quad (5.11)$$

where $\hat{Q}(s_t, a_t)$ is readily determined,

$$\hat{Q}(s_t, a_t) = r(s_t, a_t) + \gamma \mathbb{E}_{s_{t+1} \sim \mathcal{T}} [V_{\bar{\psi}}(s_{t+1})] \quad (5.12)$$

The target value network is denoted by $V_{\bar{\psi}}$ and is updated using either an exponential moving average of the value network weights, or by matching the value function weights. Lastly, the policy parameters are updated by minimizing the expectation of Equation 5.9, where the policy is reparameterized via a neural network transformation $a_t = f(\epsilon_t, s_t)$. Here, ϵ_t is input noise, represented as a vector and sampled from a fixed distribution. Noting this reparameterization, the objective function for the policy is determined,

$$J_\pi(\phi) = \mathbb{E}_{s_t \sim \mathcal{D}, \epsilon_t \sim \mathcal{N}} [\log \pi_\phi(f_\phi(\epsilon_t, s_t) | s_t) - Q_\theta(s_t, f_\phi(\epsilon_t, s_t))] \quad (5.13)$$

This combination of objective functions is implemented in an iterative algorithm to train the parameters of each network. As a final remark, performance studies indicate that simultaneously training two Q-function approximators tends to mitigate positive bias and results in accelerated convergence. Additionally, these studies have generally taken a single environment step followed by one or more gradient steps. The results of these studies, hyperparameters, and relative performance to other algorithms are discussed in the original work [38].

5.2 Reinforcement Learning Framework

A reinforcement learning framework is composed of all critical algorithmic components, such as the observation space, action space, and reward definition. Each component is finely tailored to a specific application; this investigation operates exclusively in the CR3BP environment. The design of each component is governed in part by standard operating principles and best practices. However, considerable freedom in design is available, with seemingly minor design changes contributing to significant differences in performance. Therefore, proper care is taken during the algorithm design phase to establish appropriate framework components.

5.2.1 Observation Space

For the scope of this analysis, the agent is assumed to possess omniscient knowledge of the environment, i.e., the state space and observation space are identical. The observation space is adequately defined to properly capture all necessary environmental information. In fact, for a valid MDP, the agent must be aware of all present information which impacts future states. For this application, the nonlinear equations of motion which govern natural dynamics in the CR3BP are dependent on current position and velocity states. Therefore, the inclusion of the three position and three velocity states is essential to satisfy the Markov property. However, to capture additional environmental information, the acceleration of the spacecraft is included in the observation space. The acceleration is determined via Equations 2.18 through 2.20. The acceleration could arguably be excluded due to the dependency on position and velocity. However, directly including the acceleration removes uncertainty and provides additional insight into the dynamics of the spacecraft.

Each aforementioned dynamical state provides sufficient information to the agent regarding how the spacecraft will react during an environmental step. However, the observation space is augmented to capture additional information imperative to the station-keeping application. Chapter 4 defines the reference orbit which represents

the six-dimensional nominal spacecraft state. Currently, the agent has no awareness of this desired position and velocity. Consequently, the nominal position and velocity states are appended to the observation space in the form of relative error. The relative state error is defined in Section 4.2 and is a critical component of the Floquet mode station-keeping formulation.

The addition of relative state error is sufficient in the quest to inform the agent of the station-keeping objective. However, additional information exists to assist the agent in determining successful corrective maneuvers. To that end, the nominal three-axis Floquet mode maneuver determined via Equation 4.11 is included in the observation space. The Floquet mode corrective maneuver acts as linear reference that generally maintains the spacecraft in the vicinity of the reference orbit. Finally, the running total cost of all corrective maneuver magnitudes is included in the observation space as a scalar quantity. The total cost to maintain the spacecraft in orbit is critically important for mission design due to budgetary constraints, as previously discussed in Section 4.4.2. Given the desire to minimize this quantity, including the total cost allows the agent to relate potential negative or positive rewards to mission performance.

This combination of observations completes the observation space defined in this investigation. Each observation, along with notation and dimensionality, is concisely included in Table 5.1 for reference. However, the designed observation space is supplied to the agent in a vector form. Therefore, the observation space, which is equivalent to the state space in this evaluation, is more formally written,

$$\mathcal{S} = [\bar{\rho} \quad \dot{\bar{\rho}} \quad \ddot{\bar{\rho}} \quad \delta\bar{x}(t) \quad \Delta\bar{V}_{FM3} \quad |\Delta\bar{V}_{Total}|]^T \quad (5.14)$$

All observations listed in Equation 5.14 serve to inform the agent of necessary information without providing redundant or extraneous data. Including unnecessary information in the observation space is generally not detrimental, but potentially hampers training rate, as the agent must discover the insignificance of the extra observations. To assess the significance of each observation, methodologies exist that assign relative weightings to each observation parameter, allowing for the discovery

Table 5.1. Observation Space for the Station-Keeping Reinforcement Learning Framework.

Observation	Notation	Dimension
Spacecraft Position	$\bar{\rho}$	3
Spacecraft Velocity	$\dot{\bar{\rho}}$	3
Spacecraft Acceleration	$\ddot{\bar{\rho}}$	3
Relative State Error	$\delta\bar{x}(t)$	6
Floquet Mode Three-Axis Burn	$\Delta\bar{V}_{FM3}$	3
Total Corrective Maneuver Cost	$ \Delta\bar{V}_{Total} $	1

of observations that are essentially neglected during the training process [42]. These weights approximate the contribution of each parameter towards determining a specific action. However, this analysis is not included in this investigation, and remains a topic for future work.

5.2.2 Action Space

The action space must be appropriately defined for the specific application. For station-keeping purposes, the range of actions is generally restricted to impulsive or continuous low-thrust corrective maneuvers. This investigation focuses solely on impulsive corrective maneuvers to prevent spacecraft from departing reference orbits. These maneuvers are implemented via the propulsion system, adjusting the current velocity state. As a result, the spacecraft is placed in bounded motion about the reference orbit.

Despite the restriction on the set of potential actions for this environment, several options exist for how the agent action space is defined. For example, the agent could determine a corrective maneuver magnitude and direction separately, which are then combined and implemented. Alternatively, and perhaps more intuitively, the agent

may simply output scalars which define a corrective maneuver vector. This type of implementation has demonstrated promising performance, and is leveraged by this investigation [43]. Therefore, the action space is defined as a set of three actions, where each action is a vector component of the corrective maneuver,

$$\mathcal{A} = [\Delta V_x \ \Delta V_y \ \Delta V_z] \quad (5.15)$$

Consequently, the actions determined by the agent are directly applied to the spacecraft velocity states. The updated states are then propagated during the environment step for a fixed time span.

The actions that are output from the agent are governed by the output activation function. Activation functions incorporate nonlinearity into neural network models, and the output activation function bounds the action. This investigation employs the hyperbolic tangent function as the activation function. Therefore, each action is bounded such that $a_t \in [-1, 1]$. This action is scaled in order to determine a realistic corrective maneuver component. The CR3BP is inherently a highly sensitive system, where a small modification to the velocity states may lead to drastic downstream effects. Consequently, the action must be appropriately scaled to deliver a significant impact, but not be overwhelming to the point of continuously destructive behavior. Empirical trials lead to the implementation of the following vector scaling function,

$$\Delta \bar{V}_{RL} = \text{sign}(\mathcal{A}) \cdot (\exp(2.25\mathcal{A}^2) - 1) \quad (5.16)$$

which is graphically illustrated in Figure 5.2. This scaling function is chosen specifically to accommodate the sensitivity associated with the working environment. Approximately half of the unscaled action relates to a minuscule portion of the entire scaled action range. This more easily allows for small-scale changes in the applied corrective maneuver. However, Floquet mode station-keeping analysis has demonstrated that as the state error increases, the magnitude of the corrective maneuver increases accordingly. Therefore, the spectrum captured by the action signal must also include a portion of corrective maneuver magnitudes capable of delivering impactful alterations to the current velocity state.




Figure 5.2. Action Scaling Function.

5.2.3 Reward Definition

The reward serves to inform the agent of the current state of the environment with respect to the station-keeping objective. To that end, a reward is devised which concisely quantifies the current standing of the spacecraft. Clearly, a spacecraft that has no state error is located on the reference orbit and matches the desired velocity. This ideal situation is quantified by two rewards to incentivize the agent for states that minimize position error and velocity error. Additionally, two rewards are applied at terminal states to provide incentive for minimizing the total cost and to promote success. Therefore, the total reward for any given state is determined by culminating all four individual reward components,

$$r_t = r_{\delta\bar{\rho}(t)} + r_{\delta\dot{\rho}(t)} + r_{|\Delta\bar{V}_{Total}|} + r_{Success} \quad (5.17)$$

With the exception of the reward for successful episodes, all other rewards are determined as a function of unbounded quantities, e.g., position error $|\delta\bar{\rho}(t)| \in [0, \infty]$. Consequently, bounded functions are employed to clearly define a range for each unbounded reward component. This concept is necessitated mathematically, but also makes practical sense. For example, the difference between a position error of 10^6 km

and $\infty \text{ km}$ is irrelevant in terms of successful station-keeping. Both cases are classified as failures, and as a result, should receive no reward. This conceptualization leads to the definition of a bounding function for the position error reward,

$$r_{\delta\bar{\rho}(t)} = -\tanh(4000 \cdot |\delta\bar{\rho}(t)|) + 1 \quad (5.18)$$

where the position error $\delta\bar{\rho}(t)$ is supplied in non-dimensional units. Similarly, the velocity error is bounded using hyperbolic tangent,

$$r_{\delta\dot{\bar{\rho}}(t)} = -\tanh(2500 \cdot |\delta\dot{\bar{\rho}}(t)|) + 1 \quad (5.19)$$

where $\delta\dot{\bar{\rho}}(t)$ is also non-dimensional. Unlike position and velocity error rewards, which are applied for all state-action pairs, the rewards for minimal total cost and success are only applied at terminal state-action pairs. This discontinuity is represented mathematically as piece-wise functions. The reward for total cost,

$$r_{|\Delta\bar{V}_{Total}|} = \begin{cases} -\tanh(100 \cdot |\Delta\bar{V}_{Total}|) + 1 & : (\mathcal{S}, \mathcal{A}, \mathcal{T} = 0, \mathcal{R}, \gamma) \\ 0 & : (\mathcal{S}, \mathcal{A}, \mathcal{T} > 0, \mathcal{R}, \gamma) \end{cases} \quad (5.20)$$

is a function of the non-dimensional total corrective maneuver magnitude of all non-terminal MDPs for the episode. The reward signals for the position error, velocity error, and total cost are illustrated in Figure 5.3. The position error, velocity error, and total cost are dimensionalized in each subplot to provide perspective into each signal. However, all terms in Equations 5.18 through 5.20 are evaluated using non-dimensional quantities. Finally, the bonus reward for state-action pairs that terminate as successful episodes is similarly defined as a discrete piece-wise function,

$$r_{Success} = \begin{cases} 3 & : (\mathcal{S}, \mathcal{A}, \mathcal{T} = 0, \mathcal{R}, \gamma) \text{ and } p_\tau = 1 \\ 0 & : \text{else} \end{cases} \quad (5.21)$$

where p_τ is defined as a Boolean that indicates whether an episode meets the conditions for success, as defined by Table 4.3.

Looking at each reward signal, clearly the bonus for successful episodes dominates the net reward. By design, successful episodes are regarded as the most important




Figure 5.3. Reward Component Signals.

objective that the agent seeks to accomplish. The emphasis on success has the effect of indirectly aiding propellant minimization. This observation results from the minimization of state error. As state error decreases, especially in early episodic states, the downstream corrective maneuvers are significantly reduced in magnitude. Additionally, if too much priority is applied to cost reduction, the agent may quickly converge to a less-than ideal high reward solution: terminate the episode via a single maneuver that leads to immediate failure. In this investigation, this exploitation of the reward structure is avoided by reducing the relative weighting. However, the exploitation could equally be removed by adding additional reward constraints that prevent positive reinforcement for such events.

5.2.4 Hybrid Controller Design

A station-keeping controller is designed to maintain a spacecraft in desired motion, modeled by a reference orbit. Chapter 4 details an in depth analysis of a traditional Floquet theory based controller, known as Floquet mode. Relevant literature has similarly investigated the merits of Floquet mode and compared performance to a reinforcement learning controller [43]. This investigation also seeks to study how reinforcement learning is applicable to the station-keeping problem. However, rather than replacing the Floquet mode controller outright with a reinforcement learning agent, a new type of hybrid controller is formulated. This controller seeks to leverage the generally successful performance of the three-axis Floquet mode controller, while using reinforcement learning to enhance the corrective maneuver. The net maneuver is then commanded to the spacecraft. To that end, the corrective maneuver of the hybrid controller is determined,

$$\Delta\bar{V}_{RLFM} = \Delta\bar{V}_{FM3} + \Delta\bar{V}_{RL} \quad (5.22)$$

where $\Delta\bar{V}_{FM3}$ is the three-axis Floquet mode maneuver determined from Equation 4.11 and $\Delta\bar{V}_{RL}$ is the scaled action calculated using Equation 5.16. Essentially, the agent seeks to augment the Floquet mode corrective maneuver in an effort to meet the objectives defined by the reward definition. While the three-axis controller presented in Section 4.5.3 reliably succeeds in the station-keeping problem, performance decay is noted due to increased operational errors and mission constraints. This degradation is attributed, in part, to the three-axis formulation. The corrective maneuver is determined via a control law that derives from a linear analysis of the nonlinear system. The linear approximation is generally representative of the nonlinear dynamics, resulting in corrective maneuvers that tend to successfully recover the reference orbit. However, as the state error increases, the linear corrective maneuvers become increasingly unreliable. This observation has led to a technique which uses the Floquet mode corrective maneuver as an initial guess in a targeting scheme. While this control scheme results in superior corrective maneuvers, the computational requirements

are significantly increased. Therefore, this investigation seeks to deliver similarly adjusted corrective maneuvers without the additional computational burden.

5.2.5 Training Process

The agent is trained using the reinforcement framework detailed throughout this section. However, to fully define the training process, episodic details including the initial environment setup, operational errors, mission constraints, and termination conditions are clarified here. A full tabulation of all training parameters is included in Appendix A for reference.

The agent is trained through mass iteration of episodes. Each episode is initialized with a set of randomized initial position and velocity states. These initial conditions are determined by applying random injection errors to the nominal initial state, dictated by the reference orbit. For the training process in this investigation, the reference orbit previously introduced in Section 4.5.1 is selected. An environmental step is then performed for a length of time equal to the minimum time between corrective maneuvers. The agent does not perform an action at the initial state, as this state is assumed to be delivered via an orbit insertion maneuver. After the environmental step, random tracking error is applied to the current state. At this point, Equation 4.11 is employed to determine the portion of the corrective maneuver from Floquet mode. An observation of the current state is then performed, which is fed to the agent. The agent then determines an action in the form of a corrective maneuver. The two corrective maneuver vectors are combined and commanded to the spacecraft. However, prior to updating the current velocity state, random corrective maneuver error is applied to $\Delta\bar{V}_{RLFM}$. The perturbed corrective maneuver is then implemented, and another environmental step is taken. This cyclic process is repeated until a terminal condition is reached. For training, an episode is terminated in success if the final simulation time of approximately *4.87 years* is reached, which corresponds to 10 revolutions of the reference orbit. Conversely, an episode is ter-

minated in failure if at any state, the total maneuver cost exceeds 1000 m/s or the position error surpasses 100000 km . After the episode concludes, the environment is reinitialized and the episodic process is repeated.

After each episode terminates, the discounted return is recorded. These returns are used to create a training curve which represents the performance of the agent. In Figure 5.4, the discounted return for each episode is denoted by a data point. Additionally, the moving average of the discounted returns is displayed using a sliding window size of 1000 episodes. This figure gives immediate insight into the current




Figure 5.4. Discounted Returns Per Episode and Moving Average During Training ($\gamma = .99$).

performance of the agent. An increase in the moving average corresponds to a positive trend in performance, as more episodes have terminated with higher discounted returns. Additionally, the curve indicates the point of maximum performance and when

performance gains have plateaued. For this training session, maximum performance was reached at episode number 28147, with a discounted return moving average of 3.217. The trend appears to show that additional training would lead to increases in performance. However, this training session performed nearly 50000 episodes, and no increase was observed past the range depicted in Figure 5.4. Therefore, after the training process is concluded, the agent state corresponding to episode 28147 is reinstated. This agent is subsequently employed for all further analysis and results depicted in this chapter.

5.3 Hybrid Controller Results

An analysis of the hybrid controller formulated in Section 5.2.4 is conducted to determine performance characteristics. These evaluations include studies into the expected cost, sustainability, and reliability of the hybrid controller, with respect to the station-keeping objective. The controller is assessed using baseline simulation parameters similar to prior literature, allowing for direct performance comparison. Thereafter, an extended analysis mirroring Section 4.5.3 is examined, which imposes increased mission constraints and permutes injection error.

5.3.1 Baseline Controller Performance

The baseline performance of the controller is assessed in a study that seeks to prevent the spacecraft from departing a L_1 halo orbit in the Sun-EMB system. Projections of this reference orbit and non-dimensional initial conditions are provided in Figure 4.1 and Table 4.4, respectively. The motion of the spacecraft is then simulated for 4.87 *years*, equivalent to 10 revolutions of the reference orbit, using the baseline operational errors and mission constraints from Tables 4.1 and 4.2, respectively. Clearly, the introduction of stochastic operational error suggests that controller performance will vary across simulations. Therefore, a Monte Carlo is performed to accurately capture the desired performance metrics. Section 4.4.1 details a process

for determining the required number of simulations to meet a desired confidence level and relative precision of a specified performance metric. This technique is leveraged to ascertain the expected corrective maneuver cost with 99% confidence at a relative precision of 1%. In other words, there is 99% confidence that the true cost to maintain the orbit is within 1% of the cost returned from the Monte Carlo simulation.

The results of the Monte Carlo analysis are illustrated in Figure 5.5, where each data point is the cost of a single simulation. Meanwhile, the curve represents the moving average of all simulations performed up to the current simulation number. This




Figure 5.5. Monte Carlo Results for Baseline Performance Metrics of the Hybrid Controller.

analysis required 394 simulations to stabilize the expected cost within the specified statistical requirements. At this point, the Monte Carlo is terminated and the performance metrics for total cost, linear divergence rate, and success rate are recorded.

Concludingly, the true cost to maintain the reference orbit using the hybrid controller, at 99% confidence, is within 1% of 0.660 m/s for the specified mission constraints and operational errors. Additionally, a spacecraft utilizing this controller is expected to diverge from the reference orbit at a rate of 0.594 km/year . Finally, in all 394 simulations performed, the station-keeping objective was successful, as defined by the metrics for success in Table 4.3.

5.3.2 Performance Relative to Floquet Mode Controllers

This investigation places emphasis on three primary concerns that are generally associated with station-keeping controllers: the expected cost to perform all necessary station-keeping operations, the sustainability of maintaining the reference orbit throughout and beyond the scope of the simulation, and the reliability of the controller. Clearly, the expected cost determined via the Monte Carlo speaks directly to the first concern. To address the second concern, the linear divergence is considered. This metric is determined by fitting a linear trendline to the position state error history. Preferably, the controller tends to keep spacecraft in bounded motion about the reference orbit, resulting in a position state error history that does not increase over time. Therefore, the linear divergence rate concisely differentiates between bounded and divergent behavior. Additionally, a possibility exists where a simulation terminates and meets all conditions for success, but one or more constraints are on the verge of failure. In this event, any additional propagation will result in failure. This undesirable behavior is also captured by the divergence rate. Finally, the rate of success is presented as a means of quantifying the reliability of the controller.

The three performance metrics are subsequently used to represent the quality of station-keeping controllers. Consequently, the metrics for various controllers are comparable, assuming that the analysis of each controller utilizes the same simulation parameters, operational errors, and mission constraints. The baseline analyses for both Floquet mode controllers in Section 4.5.1 and the hybrid controller in Section

5.3.1 were conducted under identical simulation conditions. Therefore, a comparison of the performance metrics is acceptable. The expected cost, linear divergence rate, and success rate for the baseline performance of each controller are listed in Table 5.2. Clearly, each controller type successfully performs station-keeping operations for

Table 5.2. Comparison of Expected Baseline Performance Metrics for Floquet Mode and Hybrid Controllers.

Controller Type	Total Cost [m/s]	Divergence Rate [$km/year$]	Success Rate [%]
Floquet x -Axis	2.798	24.159	100
Floquet Three-Axis	0.720	1.137	100
Hybrid (RLFM)	0.660	0.594	100

all evaluated simulations. However, relative to the Floquet mode x -axis controller, the hybrid controller offers a significant reduction in the total cost and expected divergence rate, similar to the Floquet mode three-axis controller. Meanwhile, the hybrid controller offers marginal reductions in total cost and divergence rate compared to the Floquet mode three-axis controller.

In addition to the baseline analysis, Section 4.5.3 evaluates the performance of both Floquet mode controllers in response to increased injection error and limited corrective maneuvers. The aforementioned study extends the minimum time between maneuvers to 90 *days*, which approximately applies a constraint of two maneuvers per revolution. Additionally, a series of injection error permutations are studied, determined from 1- σ position errors ranging between 1.0 *km* and 2000 *km*, and 1- σ velocity errors ranging between 0.001 *m/s* and 2.0 *m/s*. For each permutation, the position and velocity error are constant across all directions, i.e., $\sigma_x = \sigma_y = \sigma_z$ and $\sigma_{\dot{x}} = \sigma_{\dot{y}} = \sigma_{\dot{z}}$. In this section, the hybrid controller undergoes an identical analysis, offering insight into the robustness and adaptability of the controller. Performance contour maps for the expected rate of success and expected total cost as a function of

the injection error are displayed in Figure 5.6. Immediately noticeable is the relatively




Figure 5.6. Success Rate and Expected Cost of the Hybrid Controller for Increased Injection Error.

high rate of success for all injection error permutations. The success rate remained steady around 99%, with a spike in performance at the minimum injection error permutation, and a dip at the maximum permutation. Additionally, the expected total cost increases proportionally to both position and velocity injection error.

A direct comparison of the performance contour maps for the hybrid controller in Figure 5.6 and the Floquet mode three-axis controller in Figure 4.7 is illustrated in Figure 5.7. This figure displays the relative difference of the hybrid controller compared to the three-axis controller. Clearly, the hybrid controller outperforms the three-axis controller in terms of success rate for all error permutations in this study. The hybrid controller met the criteria for success between 98% and 100% of the time, depending on the injection error permutation. In contrast, the expected success rate of the Floquet mode three-axis controller varied between 72% and 100%. However, this gain in reliability is offset by an increase in required propulsive costs. A significant increase in expected cost is observed for the minimum injection error permutation, and as the injection error increases, the gap in performance diminishes.




Figure 5.7. Hybrid Controller Performance Relative to the Floquet Mode Three-Axis Controller.

In conclusion, the hybrid controller achieves nearly optimal performance in terms of success rate, with a trade-off in expected cost relative to the Floquet mode three-axis controller. However, as discussed in Section 5.2.3, the agent prioritizes succeeding in the station-keeping objective over minimizing the expected cost. To that end, the hybrid controller formulation successfully delivers a robust station-keeping methodology that is capable of reliably managing unexpected operational errors and restrictive mission constraints. While the controller type is generally selected to satisfy a set of requirements, a controller which demonstrably minimizes failure points is highly desirable for mission applications. Therefore, the observations drawn for this analysis suggest that the application of reinforcement learning to the station-keeping problem offers significant benefits, and further investigation is warranted.

5.4 Comparative Station-Keeping Scenario

A sample scenario is assessed to demonstrate the differences in resulting spacecraft trajectories when employing the Floquet mode three-axis controller and the hybrid

controller. However, the station-keeping problem formulated in this investigation relies on stochastic errors. Hence, Monte Carlo simulations are required to determine expected performance metrics. As a result of this stochasticity, the results from any one simulation, such as this example, may not be representative of true performance. Therefore, care is taken to not draw any conclusions from this example, and simply observe the resulting behavior.

In addition to the Floquet mode three-axis controller and the hybrid controller, a new control strategy is introduced that leverages the differential corrections techniques introduced in Section 3.3 to deliver a corrective maneuver that eliminates position error. This controller uses a simple fixed point, fixed time targeting scheme using the corrective maneuver determined from the Floquet mode three-axis controller as an initial guess. To that end, a position state along the reference orbit that is a specified amount of time downstream is targeted. For this analysis, the fixed time is equivalent to the minimum time between corrective maneuvers. This targeting scheme effectively removes all position error for the next state at which a maneuver is allowable. At this point, a second maneuver simply corrects for velocity error, enabling the full recovery of a reference state.

Each control type is subsequently evaluated in an independent, identical trial. Accordingly, the stochastic operational errors are replaced with seeded random error. Therefore, for each controller evaluation, the spacecraft is initialized with identical perturbed initial conditions, and the resulting propagated trajectory is similarly perturbed with identical tracking error. Finally, once a corrective maneuver is required, the nominal maneuver is perturbed by identical corrective error. For this scenario, injection errors of 2000 km and 2.0 m/s are applied, along with the standard tracking and corrective errors from Table 4.1. Additionally, all baseline mission constraints from Table 4.2 are imposed with the exception of the minimum time between successive maneuvers, which is extended to 90 days .

5.4.1 Trajectory Analysis

Due to the differences in the corrective maneuver determined from each controller, the spacecraft follows differing trajectories over the 4.87 *year* simulation time. However, the operational errors are seeded across controller evaluations, leading to a common initial trajectory segment prior to the first corrective maneuver. As a result, the state of the spacecraft immediately prior to the initial corrective maneuver is equivalent for all three controllers. Thereafter, the velocity of the spacecraft is updated according to the controller, resulting in unique trajectories. This behavior is demonstrated in Figure 5.8, which contains in-plane and out-of-plane projections of all three trajectories. Clearly, all trajectories match until the execution of the




Figure 5.8. Trajectories of the Floquet Three-Axis, Hybrid, and Targeting Controllers Prior to the Second Corrective Maneuver.

first corrective maneuver, which serves as a branching point. Each trajectory is then propagated until the location of the second corrective maneuver. At this point, the Floquet three-axis controller produced a spacecraft state with the most significant position error, while the targeting controller unsurprisingly reduced the position error to approximately zero.


In addition to observing the trajectory, the initial corrective maneuver determined by each controller is of interest. This maneuver, for each controller, is listed in Table 5.3. The maneuver determined via targeting requires the least amount of propellant

Table 5.3. Initial Corrective Maneuvers Determined from the Floquet Three-Axis, Hybrid, and Targeting Controllers.


Controller	$\Delta \bar{\mathbf{V}}_1$ [m/s]	$ \Delta \bar{\mathbf{V}}_1 $ [m/s]
Floquet Three-Axis	[93.203 0.5457 7.8931]	93.539
Hybrid (RLFM)	[90.280 0.0621 7.2833]	90.574
Fixed Point Targeting	[88.761 1.5455 5.4786]	88.944

and propagates to a state with nearly zero position error, where a second corrective maneuver is executed to resolve the discrepancy in the velocity states. Immediately following this corrective maneuver, a reference state is completely recovered, i.e., the spacecraft state has approximately zero position and velocity error. Consequently, additional propagation leads to minimal developments in state error, where corrective maneuvers are only required to account for the numerical error and the imposed tracking and corrective operational errors. Meanwhile, both other controllers determine maneuvers that fail to fully resolve the state error 90 *days* after the initial maneuver. However, the hybrid controller results in less significant position and velocity errors, compared to the Floquet mode three-axis controller.


The spacecraft are further propagated until the final simulation time is reached, or until a failure criterion is met. The full state histories for all three scenarios are displayed in configuration space in Figure 5.9. Clearly, each controller resulted in vastly different trajectories that vary in terms of boundedness. The Floquet three-axis controller was unable to determine corrective maneuvers that delivered bounded motion. For the majority of the simulation, the in-plane error was relatively minimal. However, with each successive revolution, the out-of-plane error is exacerbated, ultimately resulting in failure shortly after the 8th revolution. Meanwhile, the hybrid




(a) Floquet Three-Axis In-Plane.




(b) Floquet Three-Axis Out-of-Plane.




(c) Hybrid In-Plane.



(d) Hybrid Out-of-Plane.



(e) Targeting In-Plane.



(f) Targeting Out-of-Plane.

Figure 5.9. Complete Spacecraft Trajectories for the Floquet Three-Axis, Hybrid, and Targeting Controllers.

and targeting controllers successfully met all the station-keeping objectives. However, Figure 5.9 clearly illustrates the differences in the state histories. While both controllers achieve bounded motion, the magnitude of the position errors is significantly higher for the hybrid controller. Concludingly, the targeting controller achieves relatively improved behavior in terms of state error and propellant cost. This observation is not unexpected as targeting methods are formulated to provide refined corrective maneuvers that meet a set of constraints. However, targeting algorithms are computationally expensive, often involving a significant number of large matrix inverse calculations. The addition of this computational burden is a significant factor when considering the implementation of the controller into on-board flight software.

6. CONCLUSION

6.1 Summary

This investigation considers a mission in which a spacecraft is placed in orbit about the L_1 libration point that exists within the circular restricted three-body problem. In particular, the spacecraft is in a L_1 southern halo orbit within the Sun-EMB system that is considered useful for scientific applications [21]. Unfortunately, the subset of low-amplitude L_1 southern halo orbits are unstable and highly susceptible to unmodeled external forces. For instance, the spacecraft motion is affected by solar radiation pressure and gravitational perturbations from other celestial bodies such as Jupiter and Venus [29]. As a result of the inherent instability and these unaccounted perturbations, the spacecraft will naturally diverge from the reference halo orbit, rendering the mission a failure. To avoid this clearly undesirable behavior, a station-keeping controller is employed to determine corrective maneuvers that place the spacecraft in bounded motion about the reference orbit.

Two station-keeping control strategies are of particular interest to this investigation: traditional Floquet mode and a new hybrid controller which augments Floquet mode via reinforcement learning. The Floquet mode controller leverages the natural dynamics of the model environment by eliminating the unstable direction of motion with a corrective maneuver. Thereafter, the updated spacecraft state naturally flows towards the reference orbit. The effectiveness of this control strategy is demonstrated with a comprehensive analysis of the x -axis and three-axis Floquet mode controllers, concluding that both controllers are capable of maintaining the reference orbit for baseline scenarios. However, the Floquet mode formulation derives from a linear analysis of the nonlinear dynamical system. While the linear approximation is generally representative of the nonlinear dynamics, increases in state error lead to increasingly

ineffective corrective maneuvers. As a result, the controller performance begins to degrade, with a larger percentage of simulations failing the station-keeping objective. This discovery gives rise to the formulation of a new controller which, in similar scenarios, is capable of consistently recovering the reference orbit. A hybrid controller is postulated that uses the traditional Floquet mode corrective maneuver as a baseline solution. This maneuver is subsequently perturbed by a reinforcement learning agent, and then commanded to the spacecraft. An identical analysis is performed for the hybrid controller, demonstrating an increase in success rate for all evaluated scenarios relative to the Floquet mode controller. Concludingly, the hybrid controller reliably succeeds in the station-keeping objective without a significant increase in propellant cost or computational requirements.

6.2 Remarks

Current station-keeping algorithms are faced with the challenge of minimizing the computational demand of the controller, while maximizing other performance metrics. Consequently, the analyses performed throughout this investigation seek to demonstrate a station-keeping control strategy that is effective in terms of reliability and computational efficiency. The primary conclusions and remarks that are drawn from these analyses are summarized in this section.

Both the x -axis and the three-axis Floquet mode controllers are highly desirable for mission applications. This investigation highlighted situations in which the controllers failed to maintain the reference orbit. However, for state errors that are reasonably expected in an actual mission scenario, the controllers are adequate for orbit maintenance algorithms. The control laws are relatively simplistic, suggesting that the controllers could feasibly be implemented for on-board guidance due to the computational efficiency. Additionally, in scenarios that allow for increased computational effort, the Floquet mode corrective maneuvers are excellent candidates for the initial guess.

The hybrid controller, which leverages Floquet mode and reinforcement learning, demonstrably increases the success rate without the additional computational burden that a targeting scheme imposes. The additional computational demand is offloaded to the training session, allowing the final agent to efficiently calculate the required maneuvers. Additionally, the reinforcement learning framework is easily adjusted, enabling the training of new agents that meet other mission objectives. Considering this flexibility, the observed increase in reliability, and the computational efficiency, reinforcement learning appears to be a promising venture for future exploration of the station-keeping problem.

6.3 Recommendations for Future Work

Throughout the exploration of material in this investigation, numerous branching points are encountered that warrant further study. In particular, this work exclusively pertains to the Sun-EMB system and a singular L_1 southern halo orbit. However, the analysis of other systems with varying mass ratios offers the potential for expanded applicability to a wider variety of mission scenarios. Additionally, the study of orbits with differing geometries is essential for demonstrating comprehensive performance standards. To that end, various orbits that are located in the vicinity of the primaries, have expansive geometries that traverse systems, or have different eigenstructures are of particular interest.

In addition to expanding the dynamical analysis, the proposed hybrid controller acts as a foundation to be built upon. The design of any reinforcement learning framework is highly subjective, suggesting that permutations of the framework link to varying levels of performance. Two aspects of the framework are considered for additional analysis: the observation space and the control law. The observation space included in this formulation seeks to capture all necessary information. However, further study of the importance of each observation is suggested to determine if extraneous information is supplied to the agent. Additionally, the postulated hybrid

control law considers the nominal maneuvers determined via Floquet theory and from the agent. Rather than directly incorporating both contributions, the use of weighting coefficients to modify the relative significance of either component has the potential to alter overall performance.

Finally, an apprehension common to reinforcement learning based solutions concerns the predictability of the controller successfully responding to unforeseen situations. In an attempt to minimize this uncertainty, further experimentation is warranted to investigate the response of the controller in mission scenarios outside the scope of the training environment. Within the CR3BP, the controller should be evaluated for variations in other simulation parameters. For example, analysis that includes modifications to the tracking error, corrective error, and maneuver constraints would further limit uncertainty surrounding controller performance. Thereafter, to transition the applicability of the controller from the modeled environment to a more realistic scenario, benchmarks of the controller performance should be obtained within higher fidelity models, such as a Sun-Earth-Moon ephemeris model.

REFERENCES

- [1] Johannes Kepler and Tycho Brahe. *Astronomia nova Aitiologetos*. G. Voegelinus, 1609.
- [2] Isaac Newton. *Philosophiae naturalis principia mathematica*. Jussu Societatis Regiae ac Typis Josephi Streater, 1687.
- [3] J Barrow-Green. *Poincare and the Three Body Problem, History of Mathematics*, volume 11. American Mathematical Society, 1997.
- [4] Victor Szebehely. Theory of orbits: The restricted problem of three bodies. *American Journal of Physics*, 1609.
- [5] Carl D. Murray and Stanley F. Dermott. *Solar System Dynamics*. Cambridge University Press, 2000.
- [6] G. W. Hill. *Reasearches into the Lunar Theory*. American Journal of Mathematics, 1878.
- [7] G. Gómez, J. Llibre, R. Martínez, and C. Simó. Station keeping of libration point orbits: Final report. *ESA Contract Report, Technical Report*, 1985.
- [8] Henry Pernicka. *The Numerical Determination of Nominal Libration Point Trajectories and Development of a Station-Keeping Strategy*. Ph.d. dissertation, Purdue University, West Lafayette, Indiana, 1990.
- [9] D. Scheeres, F. Y. Hsiao, and N. Vinh. Stabilizing motion relative to an unstable orbit: applications to spacecraft formation flight. *Journal of Guidance, Control, and Dynamics*, 26, 2003.
- [10] Richard Bellman. The Theory of Dynamic Programming. *Bulletin of the American Mathematical Society*, 1954.
- [11] Christopher M. Spreen. *Automated Patch Point Placement Capability for Hybrid Trajectory Targeting*. Ph.d. dissertation, Purdue University, West Lafayette, Indiana, 2017.
- [12] A. E. Roy and M. W. Ovenden. On the occurrence of commensurable mean motions in the solar system: The mirror theorem. *Monthly Notices of the Royal Astronomical Society*, 1955.
- [13] Michel Hénon. *Generating Families in the Restricted Three-Body Problem*, volume 1. Springer-Verlag, 1997.
- [14] Thomas S. Parker and Leon O. Chua. Practical numerical algorithms for chaotic systems. *Mathematics of Computation*, 1989.

- [15] Daniel Grebow. Generating periodic orbits in the circular restricted three-body problem with applications to lunar south pole coverage. M.s. thesis, Purdue University, West Lafayette, Indiana, 2006.
- [16] Rüdiger Seydel. *Practical Bifurcation and Stability Analysis: From Equilibrium to Chaos*. New York: Springer-Verlag, 1994.
- [17] Emily M. Zimovan-Spreen and Kathleen C. Howell. Dynamical structures nearby nrhos with applications in cislunar space. *AAS 19-808*, 2019.
- [18] Natasha Bosanac. *Leveraging Natural Dynamical Structures to Explore Multi-Body Systems*. Ph.d. dissertation, Purdue University, West Lafayette, Indiana, 2016.
- [19] G. Gómez, J. Llibre, R. Martínez, and C. Simó. *Dynamics and Mission Design Near Libration Points Vol. I Fundamentals: The Case of Collinear Libration Points*. World Scientific Monograph Series in Mathematic Vol. 2, 2001.
- [20] Brian Barden. Using stable manifolds to generate transfers in the circular restricted problem of three bodies. M.s. thesis, Purdue University, West Lafayette, Indiana, 1994.
- [21] Tim Keeter. Station-keeping strategies for the libration point orbits: Target point and floquet mode approaches. M.s. thesis, Purdue University, West Lafayette, Indiana, 1994.
- [22] John Guckenheimer and Philip Holmes. Nonlinear oscillations dynamical systems, and bifurcations of vector fields. *Journal of Applied Mechanics*, 1983.
- [23] Brian McCarthy. Characterization of quasi-periodic orbits for applications in the sun-earth and earth-moon systems. M.s. thesis, Purdue University, West Lafayette, Indiana, 2018.
- [24] Wang Sang Koon, Martin W. Lo, Jerrold E. Marsden, and Shane D. Ross. *Dynamical Systems, the Three-Body Problem and Space Mission Design*. California Institute of Technology, 2011.
- [25] NP Dwivedi. Deterministic optimal maneuver strategy for multi-target missions. *Journal of Optimization Theory and Applications*, 17, 1975.
- [26] Henry. J. Pernicka and Kathleen C. Howell. Stationkeeping method for libration point trajectories. *Journal of Guidance Control and Dynamics*, 16, 1993.
- [27] Carmen Chicone. *Ordinary Differential Equations with Applications*. Springer-Verlag, 1999.
- [28] Sergio Bittanti and Patrizio Colaneri. *Periodic Systems, Filtering and Control*. Springer, 2009.
- [29] Vivek Muralidharan. Orbit maintenance strategies for sun-earth/moon libration point missions: Parameter selection for target point and cauchy-green tensor approaches. M.s. thesis, Purdue University, West Lafayette, Indiana, 2017.
- [30] S. Soldini, C. Colombo, and S. J. I. Walker. Comparison of hamiltonian structure-preserving and floquet mode station-keeping for libration-point orbits. *AIAA/AAS Astrodynamics Specialist Conference*, 2014.

- [31] Kiarash Tajdaran. Incorporation of mission design constraints in floquet mode and hamiltonian structure-preserving orbital maintenance strategies for libration point orbits. M.s. thesis, Purdue University, West Lafayette, Indiana, 2015.
- [32] Belinda G. Marchand. *Spacecraft Formation Keeping Near the Libration Points of the Sun-Earth/Moon System*. Ph.d. dissertation, Purdue University, West Lafayette, Indiana, 2004.
- [33] Diane C. Davis, Sean M. Phillips, Kathleen C. Howell, Srianish Vutukuri, and Brian P. McCarthy. Stationkeeping and transfer trajectory design for spacecraft in cislunar space. *AAS 17-826*, 2017.
- [34] Richard G. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. The MIT Press, 2014.
- [35] Nicholas B. LaFrage. Autonomous guidance for multi-body orbit transfers using reinforcement learning. M.s. thesis, Purdue University, West Lafayette, Indiana, 2020.
- [36] David Silver, Thomas Hubert, Julian Schrittwieser, and et. al. A general reinforcement learning algorithm that masters chess, shogi and go through self-play. *Science*, 2018.
- [37] Vincent François-Lavet, Peter Henderson, Riashat Islam, Marc G. Bellemare, and Joelle Pineau. *An Introduction to Deep Reinforcement Learning*. Foundations and Trends in Machine Learning: Vol. 11, No. 3-4, 2018.
- [38] Tuomas Haarnoja, Aurick Zhou, Pieter Abbeel, and Sergey Levine. Soft actor-critic: Off-policy maximum entropy deep reinforcement learning with a stochastic actor. In *35th International Conference on Machine Learning, ICML 2018*, 2018.
- [39] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, and Daan Wierstra. Continuous control with deep reinforcement learning. In *4th International Conference on Learning Representations, ICLR 2016 - Conference Track Proceedings*, 2016.
- [40] Brian Ziebart. Modeling Purposeful Adaptive Behavior with the Principle of Maximum Causal Entropy. *Thesis*, 2010.
- [41] Philip Thomas. Bias in natural actor-critic algorithms. In *Proceedings of Machine Learning Research*, volume 32, pages 441–448. PMLR, 2014.
- [42] Scott M. Lundberg and Su In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, 2017.
- [43] Davide Guzzetti. Reinforcement Learning and Topology of Orbit Manifolds for Stationkeeping of Unstable Symmetric Periodic Orbits. *AAS Astrodynamics Specialists Conference*, 2019.

A. Parameters for Agent Training Process

A.1 Simulation Parameters

Table A.1. Operational Errors for the Agent Training Process.

Error Type	Position [km]			Velocity [mm/s]			σ_{mag} σ_{dir}	
	σ_x	σ_y	σ_z	$\sigma_{\dot{x}}$	$\sigma_{\dot{y}}$	$\sigma_{\dot{z}}$		
Injection	2000.0	2000.0	2000.0	2000.0	2000.0	2000.0	-	-
Tracking	1.0	2.0	5.0	1.0	1.0	3.0	-	-
Corrective	-	-	-	-	-	-	1.0%	1.0°

Table A.2. Mission Constraints for the Agent Training Process.

Constraint Type	Value
Minimum Time Between Maneuvers	90 days
Spacecraft Tracking Rate	1.0 days
Minimum Maneuver Magnitude	25.0 mm/s
Position Error Divergence	True

A.2 Reinforcement Learning Parameters

Table A.3. Reinforcement Learning Hyperparameters for Agent Training.

Hyperparameter	Value
Learning Rate	$3 \cdot 10^{-4}$
Discount Factor γ	.99
Target Smoothing Coefficient	$5 \cdot 10^{-3}$
Number of Gradient Steps	1
Target Update Rate	1
Max Buffer Size	10^6
Number of Hidden Layers	2
Number of Neurons (All Layers)	256
Output Activation	tanh