Towards Fuel-Efficient Formation Flying of an Observatory and External Occulter at Sun-Earth L2

by

William David Sanchez

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Aeronautics and Astronautics
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
May 2020

© Massachusetts Institute of Technology 2020. All rights reserved.

Author .....
Department of Aeronautics and Astronautics
May 19, 2020

Certified by .....
David W. Miller
Jerome Hunsaker Professor of Aeronautics and Astronautics
Thesis Supervisor

Certified by .....
Sara Seager
Class of 1941 Professor of Physics and Planetary Science
Thesis Supervisor

Certified by .....
Richard Linares
Charles Stark Draper Assistant Professor
Thesis Supervisor

Certified by .....
Oscar Alvarez-Salazar
Principal Engineer, NASA Jet Propulsion Laboratory
Thesis Supervisor

Accepted by .....
Sertac Karaman
Associate Professor of Aeronautics and Astronautics
Chairman, Graduate Program Committee

Towards Fuel-Efficient Formation Flying of an Observatory and External Occulter at Sun-Earth L2

by

William David Sanchez

Submitted to the Department of Aeronautics and Astronautics
on May 19, 2020, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Aeronautics and Astronautics

Abstract

The prolific discovery of habitable zone residing exoplanets via indirect detection methods have spurred many in the astrophysics and space technology community to call for the prioritization of funding for a direct exoplanet imaging space telescope, such as the NASA/JPL proposed HabEx mission. Though the state-of-the-art in optical technology suggests near-term feasibility, successful and efficient high-contrast imaging remains a problem. A promising solution is formation flying an external occulter in front of the observatory to suppress host starlight and allow for imaging of the obscured exoplanet. However, recent analyses have demonstrated that for the required separation distance between the spacecraft, angular slew maneuvers to retarget the formation line-of-sight between stars in a Design Reference Mission (DRM) demand a significant amount of fuel, restricting the potential science yield of a five year mission. It can be found that many of these analyses use traditional, impulsive control solutions to slew the occulter between points in three-dimensional positional space, or attempt exhaustive search methods to find less expensive alternatives. These approaches are uninformed by the rich and complex dynamical six-dimensional phase space in which the spacecraft truly lie. For this work it is assumed that both the observatory and external occulter are operating near Sun-Earth Lagrange point 2 (SEL2). Researchers across celestial mechanics, nonlinear dynamics, chaos theory, and astrodynamics over the last century have made considerable contributions to shedding light on the families and classes of natural trajectories existing in the phase space about Lagrange points. However, it has only been in the last few decades (and still continuing through the present) that it has been revealed how to use these previously elusive pathways in mission design. All of this points to a rich and underutilized design space for crafting naturally existing, or minimally active-control assisted, low-fuel solutions to solve complex motion problems. The difficulty lies in teasing out trajectories of interest in the often times opaque dy-

namical structure. However, history has shown that by understanding the basic classes of motion existing in the phase space through the lens of Dynamical Systems Theory (DST) — which is concerned with qualitatively uncovering the structure of solutions in a system's phase space through the study of its equilibrium points, their stability, sensitivity to parameters, and the vector flow connecting these points — it can be done.

This thesis investigates the use of natural solutions to frame and solve the formation retargeting maneuvers of an observatory/external occulter exoplanet imaging mission. By illuminating the classes of natural motion that can be exploited, fuel costs can be minimized, but more importantly, the set of all available paths contextualized within the dynamical landscape. This provides a baseline from which solutions can be interpreted and mission design trade-offs analyzed. To this end, a Trajectory Design Methodology (TDM) was developed that guides the spacecraft along the natural periodic and quasi-periodic motion of the CR3BP phase space's center manifold. The TDM determines the fuel-minimizing path, under the constraints of the analysis, that passes the formation line-of-sight through the maximum number of stars within an extended time window. Since the framework is dynamically informed, the incremental costs of deviating from this maximal path, to achieve a specific science objective, can be readily considered. A sample mission analysis demonstrating these contributions is provided.

Thesis Supervisor: David W. Miller
Title: Jerome Hunsaker Professor of Aeronautics and Astronautics

Thesis Supervisor: Sara Seager
Title: Class of 1941 Professor of Physics and Planetary Science

Thesis Supervisor: Richard Linares
Title: Charles Stark Draper Assistant Professor

Thesis Supervisor: Oscar Alvarez-Salazar
Title: Principal Engineer, NASA Jet Propulsion Laboratory

Acknowledgments

I would first like to thank my grandparents, Luciano Sanchez, Ana Benito Sanchez, and Maria del Carmen Perez-Carrion, who in the face of peril and uncertainty risked their own security to pursue opportunity and freedom in this country. In many ways, I am living the dream that they sacrificed for. Dreams only partially fulfilled in their lifetimes. Because of their courage and vision for a greater future, I dedicate this work to them.

Though no longer with us, I would also like to thank my dad, Luis Sanchez, who taught me through his twenty-eight years of service to our country the strong work-ethic, gratitude, and sacrifice I needed to succeed in this endeavor. To my mom, Carmen Sanchez — who taught me the power of prayer and so willingly dedicating years of her life for my education — I also express my sincerest gratitude.

Also, a big thank you is in order for my siblings and constant friends, Edric, Elizabeth, Chris, and Daniel — who had a sixth-sense for providing a laugh when I needed it the most — as well as my aunts and uncles Aurora, Joe, Joaquin, and Maria Elena for being some of my proudest and unwavering supporters. As Isaac Newton famously said, "If I have seen further, it is by standing upon the shoulders of giants." To that, I acknowledge that if I have reached further, it is by being pushed forward by the sacrifices of seemingly ordinary, hard-working people. Giants in my life, that just so happen to be my family. Thank you all.

I would also like to express my sincerest gratitude to my thesis committee. My advisor, Professor David Miller, who patiently encouraged me to always press forward and was steadfast in his dedication to my research, even from afar. A special thank you is also in order for Professor Richard Linares, who took me into his research group when Dave was on sabbatical. I would also like to express my gratitude to Professor Sara Seager, Dr. Oscar Alavarez-Salazar, and my thesis readers, Prof. Ryan Russel and Dr. Nicola Baresi, for their time and effort to help make this

the best work it could be, especially during this odd time in history. To the other mentors I've had along the way, Dr. Alvar Saenz-Otero, Dr. Daniel Byrd, and Dr. George Pollock: thank you for your guidance and wisdom.

This research was made possible through the support of NASA Space Technology Research Fellowship NNX16AM72H. Thank you NSTRF team for providing opportunities a younger me never would have dreamed could be possible. I would also like to express my sincere gratitude to The Aerospace Corporation for providing funding for the development of the periodic orbit database detailed in section 3.1 of this thesis (approved for public release, reference 2020-00592).

I would like to end this with a special thank you to my love Becca Massip, thank you for being a constant joy in my life and my light at the end of the tunnel. Your selfless love displayed throughout this journey not only inspired me to never give up, but also to be a better person.

Solo Deo gloria.

Contents

List of Figures

List of Tables

Nomenclature

CONOPS	Concept of Operations
CR3BP	Circular Restricted Three-Body Problem
DST	Dynamical Systems Theory
EML1	Earth-Moon Lagrange Point 1
EML2	Earth-Moon Lagrange Point 2
GMOS	Gómez, Mondelo, Olikara, and Scheeres: QPT solver algorithm named after its originators
JIT	Just-in-Time
KKG	Kolemen, Kasdin, Gurfil: QPT solver algorithm named after its originators
LOS	Line-of-Sight
NRHO	Near Rectilinear Halo Orbit
QPO	Quasi-Periodic Orbit
QPT	Quasi-Periodic Torus
SEL1	Sun-Earth Lagrange Point 1

SEL2	Sun-Earth Lagrange Point 2
TDM	Trajectory Design Methodology

Chapter 1

Introduction

1.1 Motivation

To date, The NASA Exoplanet Archive [1] registers just over four thousand confirmed exoplanet discoveries with about an equal many candidates awaiting final verification. Of these, a fraction are known to lie within a band of distances from their host star that could potentially support liquid water — *the habitable zone*. The culmination of this data, beginning largely with the discovery of 51 Pegasi b in 1995 [74], represents a leap in our understanding of the cosmos. Yet, the boundary of our knowledge is shallow in comparison to estimates of the size of the Universe, which indicate that on average each star hosts a planet, in galaxies consisting of one hundred billion stars, in one hundred billion different galaxies.

For many, these statistics evoke romanticism for a bygone era of exploration, adventure, and renewal — ideas that are equally encapsulated and motivated by hope: simply, that places unknown may hold something great. In that regard, there will never be a greater physical unknown than the Universe. Philosophically, as human beings we are driven by the desire for community and to know our surroundings — which is why questions relating to our identity and place in the cosmos are as old as our awareness of it. Scientifically, our motivation is more nuanced with questions

related to the characterization of planets, their formation, and the uniqueness of life, which then drive instrument design and rigorous measurements made by spacecraft.

First and foremost, this thesis is motivated by all of these questions. The landscape of scientific and technological advancement still needing to be spanned to bring us to the ultimate discovery of a planet like ours remains large. However, as indicated by the rapid advancements made since the confirmation of 51 Pegasi b and the state-of-the-art, this could realistically occur during our lifetime. It is the hope of the author that the result of this body-of-work would be one out of many others that will contribute to the fulfillment of this age long endeavor.

1.1.1 Science: Direct-Exoplanet Imaging

Nearly 99%[1] of all exoplanet discoveries have been detected and later confirmed via indirect methods. A majority of these discoveries are credited to the successful Kepler Space Telescope [16, 17, 57] which utilized transit photometry to sense and deduce parameters — such as orbital period, radius, and mass[2] — of exoplanets as they eclipsed their host stars. Similar to Kepler, the follow up Transiting Exoplanet Surveying Satellite (TESS) [91] launched in 2018 has already begun successfully making discoveries using this transit method, narrowing the sub-Neptune sized exoplanet candidate list for future science missions. However, through transit, information is only received second-hand from measured disturbances made of the dim planet's much easier to detect stellar host instead of from the planet itself.

Undoubtedly, indirect detection methods have led to unprecedented discoveries in a short span of time. However, the depth of insight gleaned of habitable zone residents stops short of identification of an exoplanet as Earth-like. Ultimately, nothing beats first-hand knowledge. Consequently, direct-imaging exoplanet capa-

[1]NASA Exoplanet Archive [1]: confirmed exoplanets categorized via imaging discovery method over total confirmed exoplanets. All directly imaged exoplanets to date are large and outside the habitable zone due to the current technical challenges.

[2]For more information, refer to Winn [118].

bilities have become the focus of research in a wide-variety of technical disciplines. Beyond providing tangible images of extra-solar planets in the habitable zone, such agency would allow for sampling of light directly from the planet and permit spectroscopic characterization of its atmosphere to reveal its composition and search for biosignatures.

The capacity to directly-image exoplanets is limited by how effectively its host's starlight can be suppressed. A dilemma that has been likened by scientists to finding a firefly next to a lighthouse. From the perspective of an observatory, a given habitable zone residing exoplanet (at a distance of 5-parsecs) has an angular separation from its host star on the order of 200 milliarcseconds. This and their relative brightness require 10^{-10} contrast between the two objects for direct detection [87, 111]. The high-contrast imaging problem described is listed as the root of two out of the four major technical advances needed to enable a direct-imaging exoplanet observatory mission [32].

Significant developments have been made towards two different solutions. One of these is the coronagraph. A popular solution with attraction partly due to its heritage and success in other contrast imaging problems, such as in heliophysics [71] – from which it gets its name. A coronagraph is a mask that lies in the optical pipeline of the telescope and blocks out a disk of unwanted light, allowing a fainter, off-axis object/light source to be observed by the optical instrument. Though coronagraphs, as applied to direct exoplanet imaging, provide the advantage of rapid retargeting between star systems of interest, they are limited. For the level of contrast needed for direct exoplanet imaging, the sensitivity demanded by the coronagraph to maintain the required contrast dictates the observatory requirements (i.e., thermal control, pointing control, optics design, etc.) to the limits of the state-of-the-art, including a 10 picometer level instrument wavefront stabilization [87]. Even if these costly technical hurdles are surpassed, the coronagraph's inherent outer working angle and current required optical design return low throughput of signal, potentially limiting

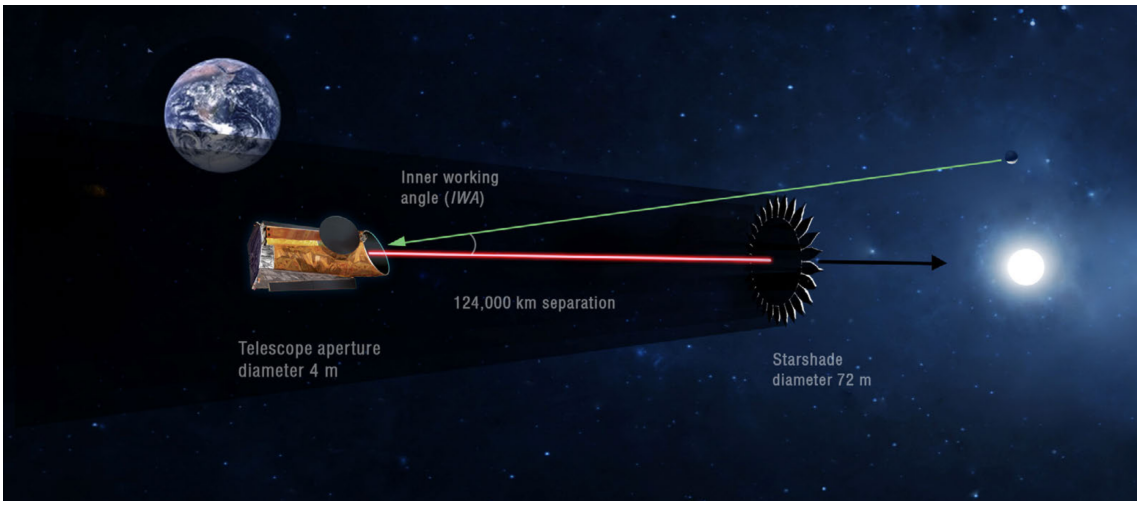

Figure 1-1: Artistic rendering of a reference observatory and starshade occulter [32]

its ability to spectroscopically characterize the exoplanets in question.

Alternatively, researchers in conjunction with the NASA Jet Propulsion Laboratory are developing an external (starshade) occulter, depicted in Figure 1-1. Similar to an internal coronagraph, the objective of the starshade is to suppress host starlight while allowing light from the target exoplanet to pass through to the observatory instrument suite. However, the starshade is external to the telescope and is itself an independent spacecraft designed to fly in formation with the telescope. It's distinctive flower petal/radially apodized shape provides a very effective dark shadow, capable of the 10^{-10} level contrast needed for exoplanet imaging when positioned between a star and the line-of-sight of the observatory. For the external occulter, contrast and inner working angle are independent of the telescope aperture and frees the system from complex, specialized optical design and wavefront correction [97]. Combined with no outer working angle, high-throughput of signal is achieved yielding a solution that is much better suited for spectroscopy than the coronagraph.

In many ways, the external occulter appears to be the most effective and promis-

ing solution to obtaining high-contrast imaging. However, the cost of choosing the external occulter and forgoing the stability burdens levied by the coronagraph on observatory design is the creation of an unprecedented and potentially fuel intensive guidance and control formation flying problem[3]. Stated succinctly, for each star in a given Design Reference Mission (DRM) — i.e., a schedule of target star systems to image — the external occulter, assumed operating in the Sun-Earth L2 environment[4], must slew to align itself between the observatory instrument and target star, maintaining a fixed separation distance of tens of thousands of kilometers axially[5], and within a $1m^2$ area laterally during imaging for the required instrument integration time [32].

Current estimates claim that the retargeting (angular slew) maneuvers at the required large separation distance between the spacecraft make the external occulter a fuel limited system, and ultimately capable of imaging fewer targets (even if more effectively) than the coronagraph[6]. As an example, initial studies for starshade integration with the planned WFIRST mission demonstrate severe fuel restrictions [112]. However, many of these analyses use traditional impulsive control solutions to slew the occulter between points in three-dimensional positional space [58,100,112]. This approach is uninformed by the rich and complex six-dimensional phase space in

[3]It should be noted that while formation flight is consider one of the major technical hurdles engineers will need to overcome to enable an external occulter based direct-imaging mission, it is not the only one. The Starshade Technology Development Activity (S5) documentation lists at least three other key technical gaps [117]. These are related to validating (through the continued development of high-fidelity models and demos) the external occulter's starlight suppression technology (note, this also includes reducing sunlight scattering around the edges of the starshade petals), sensing and control of the formation's lateral displacement during imaging, manufacturing the petal shapes under the tolerances required by imaging, and robust deployment mechanisms to maintain the integrity of that shape during operation. Though important to note, these additional technical gaps are not addressed in this thesis.

[4]This is a fair assumption since this is an ideal operating environment for observatories. Major mission concepts for future direct exoplanet imaging telescopes are projected to operate in this regime including HabEx [32] and LUVOIR [110].

[5]The exact value is dependent on the external occulter size — which is on the order of tens of meters

[6]Due to both large slew times and potentially shorter mission lifetime as the HabEx mission predicts the external occulter running out of fuel in five years [32]

which the spacecraft truly lie. Researchers across celestial mechanics, nonlinear dynamics, chaos theory, and astrodynamics over the last century have made considerable contributions to shedding light on understanding the families and classes of natural trajectories existing in the phase space about Lagrange points. However, it has only been in the last few decades (and is still continuing to the present) that it has been revealed how to use these previously elusive pathways in mission design [47]. All of this points to a rich and underutilized design space for crafting naturally existing, or minimally control assisted, low-fuel solutions to solve complex motion problems. The difficulty lies in teasing out trajectories of interest in the often times opaque dynamical structure. However, history has shown that by understanding the basic classes of motion existing in the phase space through the lens of Dynamical Systems Theory (DST) — which is concerned with qualitatively uncovering the structure of solutions in a system's phase space through the study of its equilibrium points, their stability, sensitivity to parameters, and the vector flow connecting these points — it can be done.

1.1.2 Engineering: Opening the Trajectory Design Space

In addition to the scientific motivation of addressing a critical need of future direct exoplanet imaging missions, secondary engineering motivations abound. Primarily, these lie within the field of multi-body astrodynamics and trajectory design. Historically, the problem of characterizing the motion of a particle under the attraction of multiple massive bodies has been opaque. Despite progress over the last century in moving beyond two-body analytical solutions to mapping out the Circular Restricted Three-Body Problem (CR3BP) phase space (Section 2.2), understanding of how to maximize utilization of the set of available, natural trajectories is still developing. The building of trajectory design methodologies, as demonstrated in this thesis, that further employ different classes of solutions/stability regimes in the dynamical phase space for a larger variety of missions is of merit in itself.

1.1.3 Additional Applications

Other applications exist beyond astrophysics/observatory pointing and high-contrast imaging. The results developed in this thesis will in general extend to any three-body Lagrange point environment. In the current political climate, significant attention is being paid to the cis-lunar (specifically Earth-Moon L1) region of space. It may be of interest for manned exploration applications to place a spacecraft to observe and characterize the operating environment from an EML1 vantage point. Strategically manipulating positioning along the center manifold, including changing location along the surface of invariant tori guided by its vector flow, could prove advantageous in these situations.

1.2 Literature Review

This thesis intersects three disciplines of study: direct-exoplanet imaging, Dynamical Systems Theory as applied to trajectory design, and formation flying guidance and control. Foundational work and the state-of-the-art of each field are highlighted in the following sections. Many of these studies were instrumental in inspiring the research proposed in this document, and consequently builds on their results. These cases are emphasized.

1.2.1 Direct Exoplanet Imaging

The first reported exoplanet discovery by direct imaging was of a gas giant, several times larger in mass than Jupiter, about the young brown dwarf 2M1207 in 2004. As described by authors Chauvin et al. [20], the finding was made using infrared wavefront sensing and the adaptive optics instrument NACO of the European Southern Observatory's Very Large Telescope (VLT). In 2009 the Hubble Space Telescope, with onboard coronograph instrument, made the first direct detection of an exo-

planet in the visible spectrum (a gas giant three times the mass of Jupiter), and furthermore, orbiting a bright star (Fomalhaut) [54]. Four more gas giants were discovered using infrared direct imaging orbiting HR8799; three of these in 2008 [72] and the fourth in 2010 [73].

Though these findings are encouraging, they are rare and have occurred only for special classes of planets. Only about 1% of the over 4,000 confirmed exoplanets have been directly imaged [1]. Of these, all have been giant planets far from their host star. Additionally they have either been young — such that they have retained sufficient heat from formation, making them either bright enough for detection in the visible range (e.g., Fomalhaut) or hot enough for the infrared [98] (e.g., HR8799) — or orbiting brown dwarfs, which are markedly dim compared to the sun (e.g., 2M1207). These special cases and similar have a lower bar of required imaging contrast that happen to be within the current capabilities of telescope suppression techniques. Research effort has been spent on furthering search capabilities and narrowing candidates for imaging such classes of planets, both under the lens of current [79][7] [19, 46][8] [62][9] and future [18][10] capabilities.

Studies suggest that atmospheric distortions will limit contrast for ground based telescopes to 10^{-8}, even when combining extreme adaptive optics and next generation extremely large telescopes (ELT) [23, 105]. Thus, the most promising path of obtaining regular direct imaging of exoplanets is through space-based telescopes. However, as described by Crill and Siegler [23], in-space, starlight suppression technology is still orders of magnitude in contrast away from achieving direct imaging of rocky, Earth-like planets in the habitable zone. The authors provide an overview of the two technologies in development to overcome these deficiencies: the aforementioned coronagraph and external occulter.

[7]Gemini Planet Imager – searching for giant planets and brown dwarfs.
[8]Adaptive optics and coronagraphy for ground based exoplanet detection.
[9]Direct detection via optical interferometry.
[10]Extreme adaptive optics for 30m class telescopes and giant planet detection.

Dating back to 1939, the coronagraph was initially designed by Bernard Lyot to observe the solar corona [71]. As discussed in Chapter 1, the coronograph lies internal to the observatory optics, requiring active wavefront control via deformable mirrors to achieve the stability required to handle speckling caused by diffraction along support structures and optical imperfections [23, 97, 106]. As previously described, this method has been used successfully on ground and space based telescopes (i.e., Hubble) for giant exoplanet imaging. However, achieving the stability characteristics needed to attain 10^{-10} contrast would drive heavy requirements throughout the observatory design, including picometer class wavefront control [87]. Despite this, Trauger [111] demonstrated in lab, the capability to attain 1.2×10^{-10} contrast with a Hybrid Lyot Coronagraph.

The external occulter was first envisioned by Spitzer in 1962 [104]. It has since been studied and developed into the JPL starshade concept [97] and provides high-contrast while removing the burden of complex optical/stability design from the telescope. At full-scale, the starshade is tens of meters in diameter and requires tens of thousands of kilometers of separation from the observatory for operation. Thus, testing occurs at a reduced scale [42, 99]. The level of contrast is limited only by the manufacturing precision of the approximated radially apodized shape as well as the occulter size and corresponding spacecraft separation distance.

The rapid advancements in high-contrast imaging technology have led to both direct exoplanet imaging add-on capabilities for planned missions as well as detailed studies for near-future, direct-imaging-focused, flagship observatories. For example, The James Webb Space Telescope will be carrying a Mid-Infrared instrument for characterization of young giant exoplanets [24]. Largely as a technical demonstration, WFIRST is slated to be the first observatory to carry a high-contrast coronagraph. Furthermore, studies are taking place to evaluate the feasibility of starshade integration with WFIRST [23, 112]. Two primary candidates in competition to become the next generation great observatory — and currently at in-depth

study phase — are the Large UV Infrared Surveyor (LUVOIR) [110] and the Habitable Exoplanet Observatory (HabEx) [32, 33]. Both cite direct exoplanet imaging as a primary objective. While LUVOIR is designed for segmented aperture coronagraphy, the main HabEx design incorporates both a coronagraph instrument [27] and starshade integration. For HabEx, the two would work in tandem, with the coronagraph — less capable of spectral characterization, but more nimble — would search for exoplanets during long starshade slew maneuvers. However, a recent tangential study — named HabEx Lite [87] — is exploring the benefit of neglecting the coronagraph altogether. This starshade only option greatly reduces the stability requirements and complexity of optical design on the observatory, resulting in a telescope of half the mass as the original configuration. The only question of concern is the perceived reduction in volume of star systems that can be imaged due to current estimates of the starshade's fuel consumption and slew times.

1.2.2 Dynamical Systems Theory and Trajectory Design

This survey will be limited to the study of the Circular Restricted Three-Body Problem (CR3BP) phase space and trajectory design.

Obtaining qualitative understanding of the classes of solutions provided by a system of differential equations is the first step in trajectory design. Usually, insight into the phase space begins with finding equilibrium points and classifying their stability. This is no different for the Circular Restricted Three-Body Problem. Despite their namesake, the three collinear libration (equilibrium) points were discovered by Leonard Euler (1767) [28], while the two off-axis points (L4, L5) were found by Lagrange (1772) [63]. While others, such as Jacobi (1836) [51] and Hill (1877) [45] made important contributions along the way, it was Poincaré's (1899) [85] methods to obtain qualitative understanding of the phase space that later broke through the otherwise largely opaque problem.

Interest in the three-body problem gained a resurgence in the 1960s due to the

space race and the advent of modern computing. At this time, researchers were attracted to the advantages of placing a spacecraft in the vicinity of Sun-Earth (or Earth-Moon) L1 and L2. However, the leading strategies of attack were numerical analyses made recently approachable by computers, as opposed to the Dynamical Systems Theory (DST) understanding of the phase space as proposed by Poincaré. The objective at the time could be summarized as searching for highly sensitive initial conditions that yield useful trajectories for libration point missions. To that end, Szebehely [109] compiled much of the known information of the three-body problem at the time and obtained some of the first numerical solutions. Famously, Farquhar [29] is credited with finding halo orbits to be viable trajectories for a libration point mission. Due to the highly sensitive nature of the phase space, differential corrector algorithms[11] were employed to converge higher order analytical approximations of halo orbit initial conditions found using perturbation methods. Richardson [90] developed the widely used third-order solution used for seeding a halo orbit initial condition solver. Several flight mission about Lagrange points spawned out of this work (ISEE-C [115], SOHO [25], WIND [80], ACE [107]). All trajectories for these missions were very similar. Largely, this is due to the fact that strictly numerical approaches do not incorporate substantial theoretic understanding of the classes of solutions available, nor understanding of how to leverage relationships and transitions between trajectories available in the design space.

Though they had gained less traction at the time, researchers such as Conley [21] and Llibre et al. [66], had in fact been concurrently exploring the multi-body astrodynamics problem from the perspective Dynamical Systems Theory. Several studies were released during the last quarter of the twentieth century illuminating the invariant manifold structure of the CR3BP phase space and the natural pathways connecting the classes of motion (e.g., [34, 36, 38, 52]).

Howell, Barden, and Lo [47] made a significant contribution with their 1997

[11]For more detail, refer to section 3.1.1.

paper "Applications of Dynamical Systems Theory to Trajectory Design for a Libration Point Mission," summarizing the dense insights gained on the CR3BP phase space into a clearly executable trajectory design methodology. This was partly built on Barden's work on the utility of invariant manifolds [6, 7] and shortly thereafter led to further insight into the characterization and use of the center manifold in mission design [8]. These methodologies eventually led to the successful, low fuel trajectory design for the Genesis mission [67, 70]. Koon, Lo, Marsden, and Ross both summarized and extended this work in *Dynamical Systems, The Three-Body Problem, and Space Mission Design*, [61] leading to larger scale understanding of natural transport phenomena through the solar system (the Interplanetary Transport Network) [69, 93].

These breakthroughs have spurred numerous publications from academia; far too many to list. However, a notable contribution integral to this work has been the planar periodic orbit database recently made public by Restrepo and Russell [88], which includes thousands of initial conditions for planar periodic orbits existing in our solar system[12]. Additionally, the work of Olikara and Scheeres [81] and Baresi et al. [13], on the analysis and development of quasi-periodic solutions on the center manifold invariant tori, were critical for the success of this thesis. As will be highlighted in the next section, though there has been some academic interest in DST applied to spacecraft formation flight, there has been little to no application that meet the very specific needs of an observatory/external occulter mission.

1.2.3 Spacecraft Formation Flying Guidance and Control

Many texts exist on the subject of spacecraft formation flight [22]. Murray and Durmott [78] provide a particularly comprehensive derivation of Hill's equation for the relative motion of a spacecraft near the secondary celestial body in a three-body system, starting from the CR3BP formulation and assuming small mass parameter

[12]Several three-body systems were considered.

μ. The well known Hill-Clohessy-Wiltshire equations of relative motion for a spacecraft quickly emerge by neglecting mutual gravitational attraction terms (i.e., the secondary body becomes a spacecraft of negligible mass).

Hosts of research problems concerning formation flight have been addressed and solved over the decades, such as the development of fuel-minimizing control laws that overcome J2 perturbation for earth orbiting spacecraft [2, 95]. The focus of this section is specifically on understanding the state-of-the-art in formation flight within the CR3BP phase space and analyses directly applied to external occulter operations.

Barden and Howell [8, 9] investigated the center manifold of the CR3BP phase space to categorize classes of solutions and bifurcations between regions available for trajectory design. These include planar periodic, three-dimensional periodic, and quasi-periodic motion. The authors note that quasi-periodic solutions exist on the surface of tori surrounding the three-dimensional periodic solutions, and could potentially be used for constellations or interferometric formation flying missions. The periodicity of configurations of shapes on the torus were studied by analyzing the stability and time evolution of points on the surface. The intuition gained from these papers will be used through the remainder of this thesis. Other works describing formation flying control of interferometers are prevalent [37, 41, 53]. However, it should be noted that this problem is quite different than the one in question. Interferometers attempt to maintain a stable relative shape among all spacecraft in orbit for the duration of the imaging mission. However, the challenge of an observatory/external occulter mission is not to maintain a specific shape, but how to slew/maneuver between different target points in space in the most natural/fuel efficient way possible. However, these techniques may prove to be useful when considering holding formation during extended sensor integration (imaging) time[13].

[13]Though important, this problem is considered beyond the scope of this thesis and is left for future work.

Of particular interest are the contributions made by Baresi [11, 14] who studied the use of quasi-periodic tori for formation flying applications. The author's objective was to obtain and characterize bounded relative orbital motion between spacecraft operating about planets and small bodies. Though the contributions made to understanding the computation of quasi-periodic tori were widely used in this thesis, the formation flying applications analyzed, in general, do not extend to the observatory/external occulter retargeting needs pursued. Henry and Scheeres [44] also studied the use of quasi-periodic tori for passive formation flight (bounded relative motion), developing an analysis tool called the ellipse map to aide in the design of such relative trajectories.

As a final note, Hsiao and Scheers [50, 96] produced formation-flying, stabilizing controllers about unstable trajectories. The application was directed toward interferometers about a halo orbit, using Hill's equations. The work relied on linearized relative motion dynamics to produce a nonlinear, continuous controller that eliminates motion projected onto the stable and unstable manifolds of the corresponding periodic orbit. Fuel consumption, assuming low-thrust ion engines, are reported to be on the order of meters per second over a five year span (1000 km separation). However, there is still much work to be done before something similar could be applied to an observatory/external occulter formation flying mission. For instance, how to manipulate the motion of the formation — or transition between solutions — in a manner that could cover different parts of the celestial sphere on command is yet to be determined. Additionally, further understanding of the limits of the linear analysis (i.e., at what separation distance do the assumptions break down and the controller fail) and stability guarantees would need to be developed. Ultimately, this direction was not pursued in this thesis.

Observatory and External Occulter Formation Flight

The final point to consider will be the state-of-the-art in formation flight as applied directly to an observatory/external occulter system[14]. As far as the author is aware, little consideration has been made into the role of Dynamical Systems Theory or how classes of solutions naturally existing in the CR3BP phase space can be utilized to help solve this problem. Applications to starshade integration with WFIRST [113] — where the Design Reference Mission, and thus imaging schedule, is fixed — cite concerning levels of fuel consumption[15]. However, the authors note that fuel costs for the same maneuvers can be reduced by several factors when performed at more favorable times of the year (i.e., when they align more closely with the natural gravitational tides provided by the CR3BP). This highlights the potential consequences of applying control independently of knowledge of the dynamics of the system — that is, maximizing exploitation of nature versus fighting it. Millard and Howell [76] arrive at a similar conclusion in their study of controlling formation flying spacecraft imaging arrays through a set DRM. These studies all show fuel consumption on the order of thousands of kilograms. Leitner [64] and Folta [31] consider a more general starshade mission and cite similar numbers. Folta [31] also solved several fixed DRMs using low-thrust propulsion, reducing the total fuel consumed. However, trajectory design methods providing insight into how to interpret solutions or exploit the natural dynamics were not developed

The current state-of-the-art in formation flight specific to the external occulter is credited to Soto [100, 101]. Soto's contribution focused on the complex task of optimal scheduling of imaging the exoplanets — which includes managing integration time, keep-out zones due to bright celestial objects, mission lifetime, and other imposed timing constraints. However, control and maneuvering is determined by

[14]Though not explicitly mentioned, for those interested in sensing between observatory and external occulter for feedback control, refer to Scharf [94]

[15]This is using simple, impulsive ΔV maneuvers to arrive from point A to point B on schedule.

impulsive maneuvers uninformed by the natural dynamics. Kolemen [58, 59] did provide a methodology to design the set of formation retargeting sequences using low-thrust-propulsion-based, minimum-fuel optimal control transfers (with ordering framed and solved as a Traveling Salesman Problem). However, neither knowledge of dynamical structures nor the natural flow of the phase space/solutions were explicitly exploited or included in the design. Thus, any low-fuel trajectories selected were found by means of an exhaustive search with no mechanism to interpret them via natural phenomena.

1.3 Research Gap and Problem Statement

From the literature survey it is apparent that the observatory/external occulter formation retargeting problem has not yet been considered from the perspective of Dynamical Systems Theory. Instead, current methods are uninformed by the underlying dynamics and either neglect the development of a trajectory design methodology altogether (for the case of studies with a fixed DRM) or rely on exhaustive search methods. The results are paths that are difficult to contextualize within the dynamical landscape. Alternatively, by rooting a trajectory design methodology from the onset in DST, the design process can become informed by the natural flow and structure of the phase space, allowing for the efficient search of solutions aligning with natural phenomena. Additionally, by knowing how these paths pass through the phase space and which dynamical structures can be exploited, generated "fuel-optimal" trajectories can quickly be adapted to included specific stars of interest. More importantly, this could be accomplished through understanding of the incremental cost of adding these diversions from the optimal path — providing a baseline from which informed, design trade-off decisions can be made. Finally, this approach would also provide a foundation for future work, opening an immense design space of natural solutions that can continue to be developed in potentially

limitless ways. This thesis, with overarching objective summarized in the following problem statement, seeks to provide a framework that begins to address these research gaps[16].

Thesis Problem Statement

> To develop a trajectory design methodology — for the retargeting maneuvers of an observatory/external occulter formation flying mission at Sun-Earth L2 — rooted in Dynamical Systems Theory, that maximizes exploitation of natural solutions, thereby decreasing fuel costs while minimizing sacrifice to science efficiency.

1.3.1 List of Contributions

Provided is a list of the contributions made by this thesis to address the technical gap summarized by the problem statement and advance the state-of-the-art.

Contributions:

- A trajectory design methodology for the retargeting maneuvers of an observatory/external occulter formation flying mission operating at Sun-Earth L2, rooted in Dynamical Systems Theory, that exploits the natural motion provided by the center manifold of the Circular Restricted Three-Body Problem phase space. This includes:
  - A search method to find the set of natural trajectories on the quasi-periodic torus (for the external occulter) that naturally pass the formation line-of-sight through alignment with all reachable stars in a given time window

[16]Note, the terminal constraints on fuel costs and science efficiency are added to ground the work back to ultimate objective of an observatory/external occulter formation flying mission: to image as many candidate exoplanets as possible.

  - A method to determine and order the transfers between natural solutions on the torus that yields the maximum number of intercept between the formation line-of-sight and the set of available target stars within a given time window.
  - A method to connect the natural solutions on the torus — through minimum-fuel, low-thrust transfers — into a continuous trajectory, passing the formation line-of-sight through the maximum number of target stars in a given time window
  - A method for combining the results of individual time window analyses into a global solution across an extended time horizon.

- A queryable database of three-dimensional periodic and quasi-periodic orbits existing in the CR3BP phase space. This includes a large infrastructure of solvers and tools for mission designers to further develop and interact with the database — including rapid retrieval of the set of orbits that achieve a given mission objective. The database includes halo orbits, surrounding quasi-periodic tori, butterfly orbits, near-rectilinear halo orbits, and natural connections between select family members.

1.4 Thesis Outline

This thesis presents a trajectory design methodology for the retargeting maneuvers of an observatory/external occulter formation flying mission operating near Sun-Earth L2. The first chapter provides an introduction and motivation of the research problem, a survey of relevant literature and the state-of-the-art, a research gap analysis and problem statement, and a list of contributions made by this thesis that begin to address that gap. The remaining chapters are outlined below:

Chapter 2: Background Theory and Concept of Operations

Relevant background theory is provided including an introduction to Dynamical Systems Theory, the Circular Restricted Three-Body Problem phase space, and more specifically, the center manifold existing near the collinear libration points. A Concept of Operations for an observatory/external occulter formation flying mission that inspired the Trajectory Design Methodology is provided, based on exploitation of the center manifold, extending a design database of periodic and quasi-periodic solutions in the phase space, and optimal methods of transfer between these database solutions.

Chapter 3: Database Development

This chapter describes the development of the three-dimensional periodic and quasi-periodic torus database first introduced in the Concept of Operations. Solution methods for both cases are described as well as details of the software infrastructure developed for creating and interfacing with the database. Included in the discussion are details on stability analysis, bifurcation analysis, and computation of stable and unstable manifold solutions.

Chapter 4: Trajectory Design Methodology

This chapter outlines the Trajectory Design Methodology (TDM) developed to plan and solve the sequence of retargeting maneuvers for an observatory/external occulter formation flying mission. The methodology exploits the previously cataloged[17] periodic and quasi-periodic natural motion provided by the center manifold of the CR3BP phase space near SEL2. The TDM is divided into three steps: finding solutions on quasi-periodic tori that match formation alignment conditions with target stars, ordering the solutions along the dynamical flow of the phase space, and transferring between the ordered solutions. Simple examples are provided for each of

[17]That is, in the periodic and quasi-periodic torus database of Chapter 3

these steps and/or details of their software implementation.

Chapter 5: Sample Mission Analysis

A sample mission analysis is performed — demonstrating the TDM developed in Chapter 4 — that is based on the Habitable Exoplanet Observatory (HabEx) study's target star list. After acknowledging relevant assumptions, the mission is analyzed by applying the TDM to a series of quarter-period time windows. The local maximal paths solved for each time window are then combined to determine the global maximal path for the formation across an extended time horizon — such as a full mission lifetime. Estimates of fuel and science efficiency are performed and the resulting framework contextualized within the state-of-the-art. It is shown that by illuminating the classes of natural motion that can be exploited, fuel can potentially be minimized, but more importantly, the set of all available paths understood within the dynamical landscape — ultimately providing a baseline from which trajectories can be interpreted and mission design trade-offs analyzed.

Chapter 6: Conclusion

A summary of the thesis and research contributions is presented as well as suggestions for future work.

Chapter 2

Background Theory and Concept of Operations

This chapter presents background theory pertinent to the design and interpretation of trajectories that exploit the natural solutions of a nonlinear system's phase space. Section 2.1 provides a primer on Dynamical Systems Theory (DST) — the framework in which this thesis is rooted. Section 2.2 introduces the phase space of interest, defined by the Circular Restricted Three-Body Problem (CR3BP), while section 2.3 focuses on its center manifold subspace — which contains the families of bounded motion near the collinear Lagrange points. Finally section 2.4 sketches a Concept-of-Operations (CONOPS) for how these naturally occurring structures can be used to design the sequence of retargeting maneuvers for an observatory/external occulter formation flying mission. This CONOPS frames the approach of this thesis and guides the remaining chapters.

2.1 Dynamical Systems Theory

Consider a set of ordinary differential equations describing the vector flow of a system

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), t; \mu) \tag{2.1}$$

where $\mathbf{x} \in \mathbb{R}^n$ is its state — that is, a minimum set of independent variables that together describe the system (e.g., position and velocity), and from which all future states can be predicted given knowledge of all future inputs. Furthermore, it is made explicit that the system depends on some set of parameters $\mu \in \mathbb{R}^m$. Adopting the Newtonian convention $\dot{\mathbf{x}} = \frac{d\mathbf{x}}{dt}$, the function $f(\mathbf{x}(t), t; \mu)$ describes the time rate of change of the state, and is, in general, nonlinear and explicitly depends on time. The system considered in this work — of a small body in a gravitational field — is known to be both nonlinear and autonomous, that is, described by differential equations that do not explicitly depend on time[1]. All future analyses are restricted to this case.

The word *trajectory* and *solution* are used interchangeably, and both refer to a unique time history $\mathbf{x}(t)$ that satisfies equation 2.1. The set of all possible states make up the *phase space* of the system. Assuming $f(\cdot)$ is Lipschitz continuous, then existence and uniqueness of solutions stemming from real valued initial conditions guarantee that trajectories are dense in the phase space [92]. A well developed phase portrait depicting these trajectories — plotted as a vector field for lower-dimensional systems — can give great qualitative insight into the classes of solutions available. Often times, the phase space spans $\mathbb{R}^n$, but on occasion — usually in association with periodic motion — can be more naturally represented through cylindrical, spherical, or toroidal topological spaces [116].

[1]Note: the state of the system $\mathbf{x}$ is still time dependent. However, while the differential equation is dependent on a time varying state, the structure of the differential equation itself is not changing with time. This is what is meant by an autonomous system

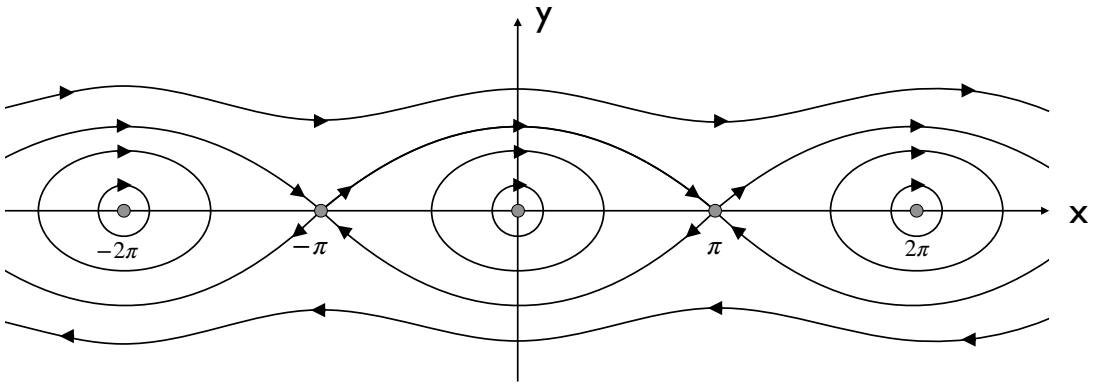

Figure 2-1: The phase portrait of a simple pendulum.

Dynamical Systems Theory is concerned with uncovering the qualitative structure of solutions in a system's phase space, as opposed to solving for individual, analytical (or numerical) trajectories[2]. This is especially true for nonlinear systems where analytical solutions may be either highly elusive or nonexistent. Through DST, $n-$dimensional complex, nonlinear system can be studied, and their topology — including equilibrium points, their stability, sensitivity to parameters, and the vector flow connecting these points — revealed. When combined, whole families of solutions (such as limit cycles and other forms of nonlinear oscillators) can be discovered that together describe natural transport through the system.

This philosophy is best exemplified by the simple pendulum, described through the following system of nonlinear differential equations and phase portrait depicted in Figure 2-1.

$$\dot{x} = y \tag{2.2}$$

$$\dot{y} = -\frac{g}{L}\sin x \tag{2.3}$$

All expected classes of motion are captured in this graphic, including stable

[2]However, when combined, qualitative structure through DST provides context for individual solutions that may otherwise be obtained.

equilibrium points[3] at $x = 0$ and multiples of 2π[4], unstable saddle points at multiples of π[5], simple harmonic motion about the stable equilibrium points (closed curves), and in the case of a high energy system, continuous spinning (top-most and bottom-most trajectories). Compare this result to the analytical solution obtained by taking a small angle approximation about $x = 0$, which fails for large-amplitude deviations from the equilibrium point and yields no knowledge or intuition outside its individual scope.

For comprehensive texts on the subject of Nonlinear Dynamics, Dynamical Systems Theory, and Chaos Theory refer to Wiggins [116], Strogatz [108], and Koon, Lo, Marsden, and Ross [61].

2.2 The CR3BP Phase Space

This section introduces the Circular Restricted Three-Body Problem (CR3BP) system of equations and summarizes the classes of natural solutions existing in its phase space. The theory and findings presented are the culmination of decades of research on the application of Dynamical Systems Theory to multi-body astrodynamics that have proven to be seminal to the field of study. Several resources exist in the literature that comprehensively describe these developments [6–8, 34, 47, 49, 61]. The reader is referred to the referenced texts for any details that may have either been omitted in this review or require further clarification.

2.2.1 Deriving the CR3BP System of Equations

Astrodynamics is the study of the natural motion of man-made objects in space. Primarily, this falls under the modeling and prediction of a spacecraft's state under the influence of a gravitational field and trajectory design. Consider a spacecraft

[3] These are states x_e such that $\dot{x} = f(x_e; \mu) = 0$.
[4] This corresponds to a downward pendulum position.
[5] This corresponds to an upward pendulum position.

operating in our solar system far from any planet. Neglecting perturbations, such as solar radiation pressure, it's motion is governed by the gravitational pull exerted by all massive bodies in the Universe. This system of differential equations describing this motion is modeled via Newtonian mechanics as follows:

$$m\ddot{\mathbf{r}} = -\sum_{i=1}^{n} \frac{GmM_i}{||\mathbf{r}_i - \mathbf{r}||^3}(\mathbf{r}_i - \mathbf{r}) \tag{2.4}$$

where m and $\mathbf{r}$ are the mass and position of the spacecraft respectively and M_i and $\mathbf{r}_i$ are the mass and position of the i^{th} massive body. Including all celestial bodies is not feasible, even when restricting to just those in the solar system, as doing so yields a system of equations that are intractable and proven to have no closed form solution [85]. Thankfully, gravitational forces drop off by $\mathcal{O}(r^{-2})$. Due to the large distances between bodies in the solar system, often times it is reasonable to neglect all but the influence of the closest (or most massive) celestial body — such as when operating in orbit about a planet or when using a patched conic approach for interplanetary travel. This two-body assumption is attractive due to it yielding well-understood, closed-form, analytical solutions[6]. However, for some regions of space, this assumption strips away dense sets of useful families of trajectories from mission design as well as valuable intuition of how natural objects transit the solar system.

For many important operating environments, a good compromise is the addition of a third body to the dynamical system. The three-body problem, with a few additional simplifying assumptions, adds enough complexity to qualitatively capture the richness of the full multi-body phase space, while also remaining approachable for qualitative analysis through Dynamical Systems Theory[7]. The aforementioned assumptions together make up the Circular Restricted Three-Body Problem (CR3BP)

[6]For more information on the two-body problem and patched conic methods refer to Prussing and Conway [22].

[7]Note, the three-body problem was also proven by Poincaré to possess no analytical solutions.

system of equations, which are derived through the remainder of this section.

The CR3BP formulation is constructed by assuming a three-body system where the mass of two of these bodies are much greater than the third — $M_1, M_2 \gg m_3$ Furthermore, if $M_1 > M_2$, then the body of mass M_1 is named the primary and M_2 the secondary. The body of mass m_3 is assumed to be a particle, and for this application, the spacecraft. It is further assumed that the orbital motion of the primary and secondary are each circular about the system center-of-mass — the barycenter — and are coplanar. The coordinate system of reference is defined as a rotating frame $x - y$ (refer to Figure 2-2) with origin placed at the barycenter. The frame rotates with a constant angular rate equivalent to that of the circular orbital motion. Thus in the rotating frame, M_1 and M_2 are fixed at their respective distances from the origin on the x-axis, as compared to their perceived rotational motion from an inertial frame $X - Y$ with the same origin.

The system is non-dimensionalized by mass unit $M^* = M_1 + M_2$, unit of length L^* equal to the distance between M_1 and M_2, and unit of time T^* such that M_1 and M_2 complete one revolution in the $X - Y$ frame in a non-dimensional time of 2π. In

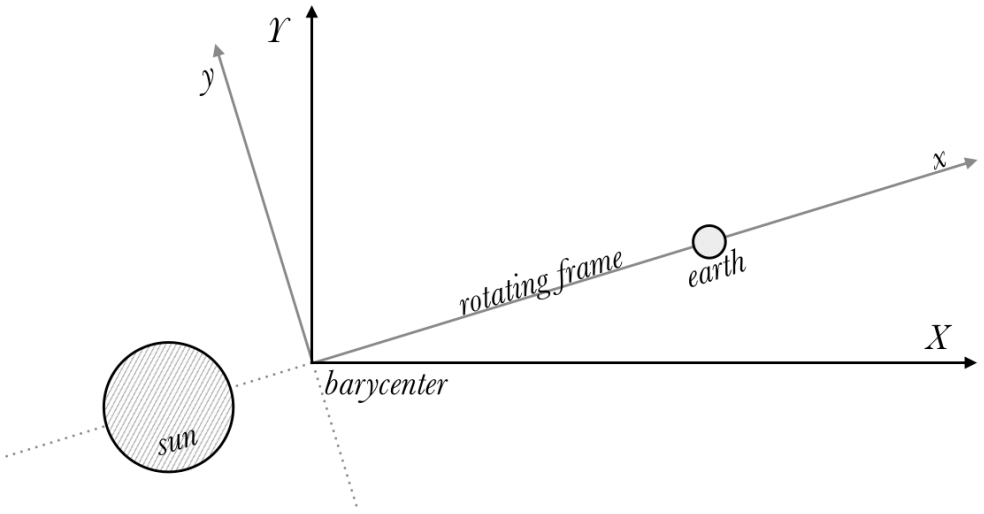

Figure 2-2: Rotating and inertial frames where M_1 is the Sun and M_2 is the Earth

these units, $G = 1$ where G is the constant of Newton's universal law of gravitation. The mass parameter is defined as follows:

$$\mu = \frac{M_2}{M_1 + M_2} \tag{2.5}$$

yielding

$$\mu_1 = \frac{M_1}{M_1 + M_2} = 1 - \mu \qquad \mu_2 = \mu \tag{2.6}$$

In the rotating frame, the secondary is located at a distance μ_1 along the $x-$axis and the primary at $-\mu_2$. The equations of motion for the CR3BP are derived starting from Newton's second law expressed in the inertial frame (X-Y)

$$\mathbf{F}_3 = m_3\ddot{\mathbf{X}}_3 = -\frac{GM_1m_3}{R_{13}^3}\mathbf{R}_{13} - \frac{GM_2m_3}{R_{23}^3}\mathbf{R}_{23} \tag{2.7}$$

where R_{13} is the position of m_3 relative to the primary and R_{23} is the position of m_3 relative to the secondary. The CR3BP system of equations are obtained after converting equation 2.7 to the rotating frame via Coriolis' Theorem and non-dimensionalizing:

$$\ddot{x} - 2\dot{y} - x = -\frac{\mu_1(x + \mu_2)}{r_1^3} - \frac{\mu_2(x - \mu_1)}{r_2^3} \tag{2.8}$$

$$\ddot{y} + 2\dot{x} - y = -\frac{\mu_1 y}{r_1^3} - \frac{\mu_2 y}{r_2^3} \tag{2.9}$$

$$\ddot{z} = -\frac{\mu_1 z}{r_1^3} - \frac{\mu_2 z}{r_2^3} \tag{2.10}$$

where

$$r_1 = \sqrt{(x - \mu_2)^2 + y^2 + z^2}\,, \qquad r_2 = \sqrt{(x - \mu_1)^2 + y^2 + z^2}$$

The CR3BP possess a well-known integral of motion called the Jacobi constant

(equation 2.11) and is a measure of the energy of the system.

$$J_c = x^2 + y^2 + 2\left(\frac{\mu_1}{r_1} + \frac{\mu_2}{r_2}\right) - \dot{x}^2 - \dot{y}^2 - \dot{z}^2 \tag{2.11}$$

2.2.2 Classifying Solutions

The CR3BP equations yield five equilibrium points — named Lagrange (libration) points. Three of these (L1, L2, and L3) are collinear with M_1 and M_2 — along the rotating frame x-axis — and are unstable. Points L4 and L5 are stable and exist symmetrically off the x-axis, forming the points of equilateral triangles with the primary and secondary bodies. Intuitively, these equilibrium points represent the balance of gravitational and centrifugal forces as observed in the rotating frame.

Qualitative understanding of the solutions available in the CR3BP phase space is at first best gained by studying the phase space near the equilibrium points — specifically the collinear points L1 and L2. The CR3BP system is both conservative and Hamiltonian[8]. Recall that the phase space is six-dimensional, thus linearizing about one of the unstable libration points and examining the eigendata yields six eigenvalues. Of these, one is stable, one is unstable, and four are center — the latter existing as complex conjugate pairs. This type of equilibrium point is classified as having center $\times$ center $\times$ saddle stability[9]. One would expect from the saddle classification that in general, trajectories depart hyperbolically along the direction of the unstable eigenvector, while a single pair of trajectories asymptotically approach the libration point along the stable eigenvector direction. Locally, this is the case. Furthermore, from the center pairs, one would predict the existence of periodic motion. In fact, for a Hamiltonian system such as this, Lyapunov center theorem ensures their existence in dense families [34]. These orbits exist in the immediate

[8]For references on Hamiltonian systems see Moser [77] and Wiggins [116]

[9]Many texts exist in the literature that describe how this result is obtained. For a particularly comprehensive resource, refer to Koon, Lo, Marsden, and Ross [61]

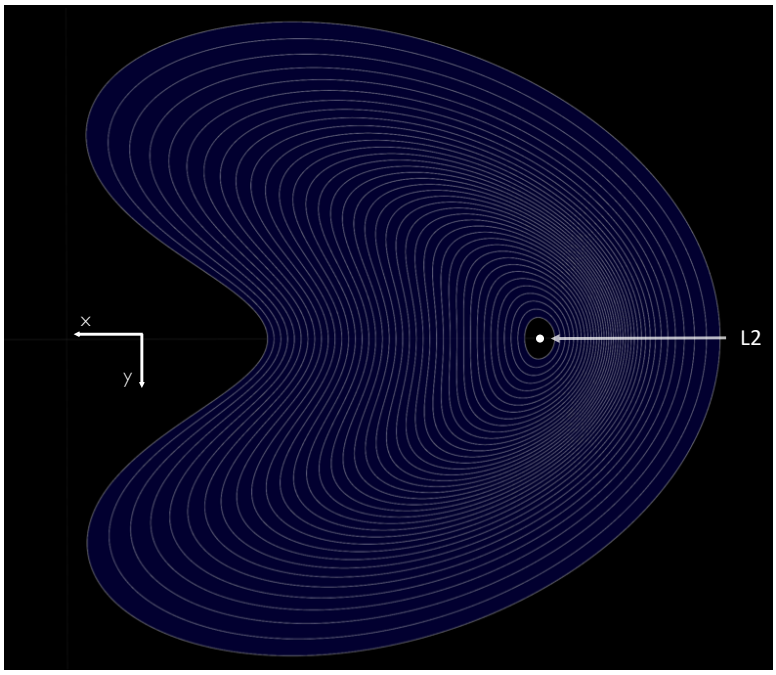

Figure 2-3: Planar Lyapunov family of orbits about Earth-Sun L2

vicinity of the equilibrium point in two varieties, planar Lyapunov orbits (which are fixed to the rotating $x - y$ plane) and the out-of-plane vertical Lyapunov orbits. Figure 2-3 depicts a dense family of these planar periodic solutions existing near Sun-Earth L2.

Prior analyses have revealed that for a critical orbital amplitude along the planar Lyapunov family, a phase space bifurcation[10] occurs (i.e., a change in the structure of solutions) [43]. Three-dimensional periodic trajectories — *halo orbits* — emerge out of the planar case at this condition (refer to Figure 2-4). These trajectories — characterized by bounded, out-of-plane periodic motion — have been the foundation of trajectory design for libration point missions.

Similar to its neighboring Lagrange point, these orbits are nominally unstable[11]. This can be demonstrated by intersecting a point along the orbit with a surface of

[10]For a more detailed discussion on phase space bifurcations, see section 3.1.3

[11]Note, that there is a small subset near the end of the halo orbit family (within the near-rectilinear halo orbit sub-classification) that are in fact stable. However, these will not be of focus in this thesis.

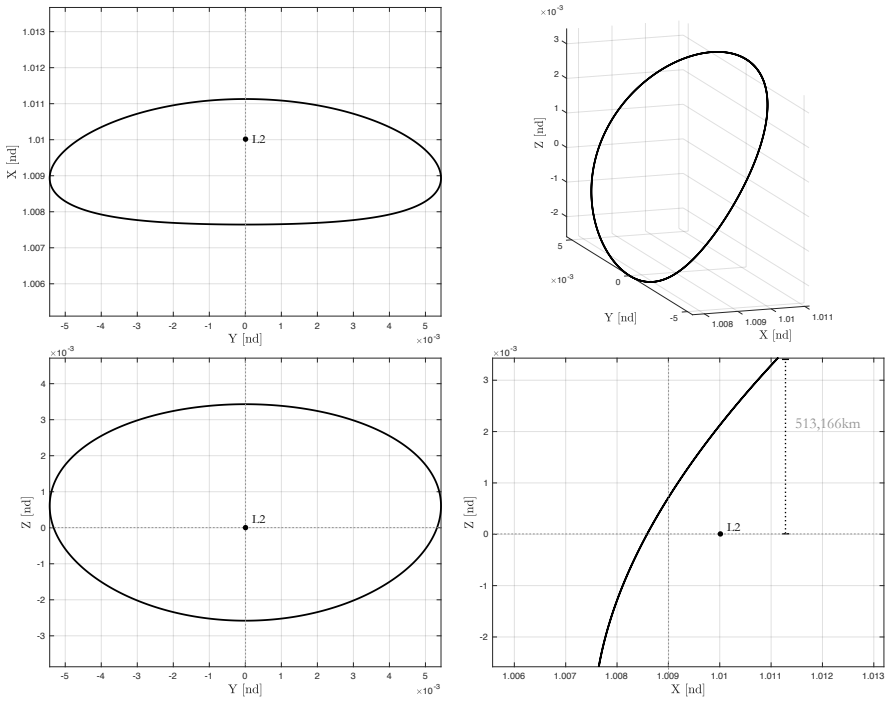

Figure 2-4: Projection of a halo orbit about SEL2 onto the rotating frame orthogonal planes

dimension $(n-1)$ — where n is the dimension of the phase space — normal to the flow. The stability of the perioidic orbit can be determined by analyzing how a perturbed trajectory re-intersects the surface after each period of revolution, where — as depicted in Figure 2-5 — each of these intersections will appear as a point. A discrete map capturing this behavior is called a Poincaré (or first return) map ($\Sigma : z_n \mapsto z_{n+1}$). Thus, if z_p is the point of intersection of a periodic orbit on this surface then $z_p = \Sigma(z_p)$. If the orbit is stable/unstable, the points on the the surface will converge/diverge from z_p.

For the CR3BP system, the state transition matrix $\Phi(t, t_0)$ along a trajectory can be computed by numerically solving the following differential equation:

$$\dot{\Phi}(t, t_0) = D_x f(x; \mu)\Phi(t, t_0), \qquad \Phi(t_0, t_0) = I_6 \tag{2.12}$$

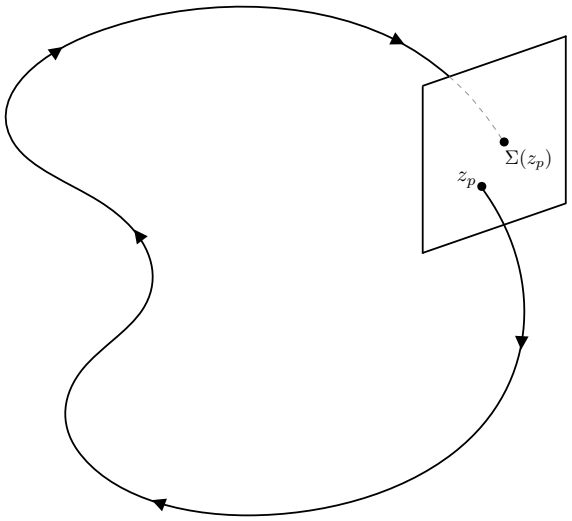

Figure 2-5: Reference 2D Poincaré surface embedded in 3D space

where the $D_x f(x;\mu)$ is the Jacobian of the CR3BP system of equations, I_6 is the 6×6 identity matrix, and $t > t_0$[12]. The monodromy matrix — $M \equiv \Phi(T,0)$ — is obtained by selecting a point on the periodic orbit and integrating equation 2.12 over a full period: T. This serves as a local, linear approximation of the Poincaré map for the surface of section at that point [61]. The eigendata of M reveals the stability of the orbit. It is known that for halo orbits, the eigenvalues take the following form (as presented by Howell, Barden, and Lo [47]):

$$\lambda_1 > 1, \qquad \lambda_2 = (1/\lambda_1) < 1, \qquad \lambda_3 = \lambda_4 = 1,$$

$$\lambda_5 = \lambda_6^* = 1/\lambda_6 \qquad |\lambda_5| = |\lambda_6| = 1$$

where λ_6^* is the complex conjugate of λ_6. Note, since this is a discrete map, stable/unstable modes lie within/outside the unit circle respectively. Thus, it can readily be seen that halo orbits maintain the same stability characteristics of their associated libration point: center $\times$ center $\times$ saddle. Similar trajectories to those

[12]Note, since $\dot{\Phi}(t,t_0)$ is a 6×6 matrix, solving 2.12 translates to solving thirty-six differential equations simultaneously

near the equilibrium point also appear in the neighborhood of the halo orbits. For a given solution, hyperbolic trajectories depart along its unstable eigenvector direction, while those along the stable eigenvector asymptotically approach the orbit. All the unstable (stable) trajectories emanating from the periodic orbit make up a dense invariant manifold surface called the unstable (stable) manifold respectively. The solutions corresponding to the four remaining center modes together make up the invariant center manifold. For the purposes of this thesis, the following terms are defined:

> *manifold* — a collection of orbits (solutions) that together form a surface of dimension $(n-1)$ that locally has the structure of Euclidean space
>
> *invariant manifold* — a surface of solutions such that a solution starting on the invariant manifold surface remains on its surface for all time

A trajectory on an invariant manifold can be described by equation 2.13:

$$\forall \mathbf{x}(t) \in \bar{W}^I \text{ if } \mathbf{x}_j(t_0) \in \bar{W}^I \rightarrow \mathbf{x}_j(t) \in \bar{W}^I;\ \forall t \geq t_0 \tag{2.13}$$

where $\mathbf{x}_j(t)$ is a trajectory, and $\bar{W}^I$ is an invariant manifold whose surface is composed of trajectories satisfying the above property. From this perspective a qualitative picture of the natural transport of mass in three-body celestial systems can be grasped. The six-dimensional phase space near the collinear Lagrange points can be envisioned as being composed of subspaces of stable and unstable invariant manifolds linked by the center manifold of bounded solutions. Therefore, theoretically fuel-free trajectories can be selected that once entered, will asymptotically transport spacecraft toward/away-from periodic orbits. The difficulty lies in teasing out the specific trajectories of interest in the often times opaque dynamical structure.

From this understanding, Howell, Barden, and Lo [47] developed a methodology to design trajectories for a mission orbiting an unstable libration point that exploits these natural solutions. The steps are summarized as follows:

1. Select a halo orbit with characteristics (e.g., amplitude) that achieves the mission objective. Numerically solve for the solution.

2. Compute the stable and unstable invariant manifold directions through the monodromy matrix evaluated at a set of points along the orbit.

3. Globalize the manifolds to obtain the transfer trajectory solution space.

4. Select solutions on the manifold that allow for efficient Earth-to-halo-orbit transfer and halo-orbit-to-Earth return (if necessary).

For more information on how to numerically solve for the initial conditions of a halo orbit using a differential corrector scheme, refer to section 3.1. To globalize the manifold, simply select a point on the orbit, perturb it along the local, linearly-approximated direction of the (un)stable manifold and propagate the nonlinear time evolution of that point backwards (forwards) in time. Repeating this for several points along the orbit yields an approximation for the surface of the invariant manifold. An equivalent mathematical expression is provided as follows:

$$\mathbf{x}^s(\mathbf{x}_0) = \mathbf{x}_0 + \epsilon \mathbf{Y}^s(\mathbf{x}_0) \rightarrow \mathbf{x}^s \in \bar{W}^s \tag{2.14}$$

$$\mathbf{x}^u(\mathbf{x}_0) = \mathbf{x}_0 + \epsilon \mathbf{Y}^u(\mathbf{x}_0) \rightarrow \mathbf{x}^u \in \bar{W}^u \tag{2.15}$$

where $\mathbf{x}_0$ is a point on the orbit, ϵ is a small perturbation constant, and $\mathbf{Y}^s$ ($\mathbf{Y}^u$) is the stable (unstable) eigenvector of the monodromy matrix respectively.

The Genesis Mission

Howell, Barden, and Lo's [47] DST approach to trajectory design is most clearly displayed in its successful application to the NASA Genesis mission [67]. Launched in 2001, Genesis was a probe sent to Sun-Earth L1 to collect solar wind particles

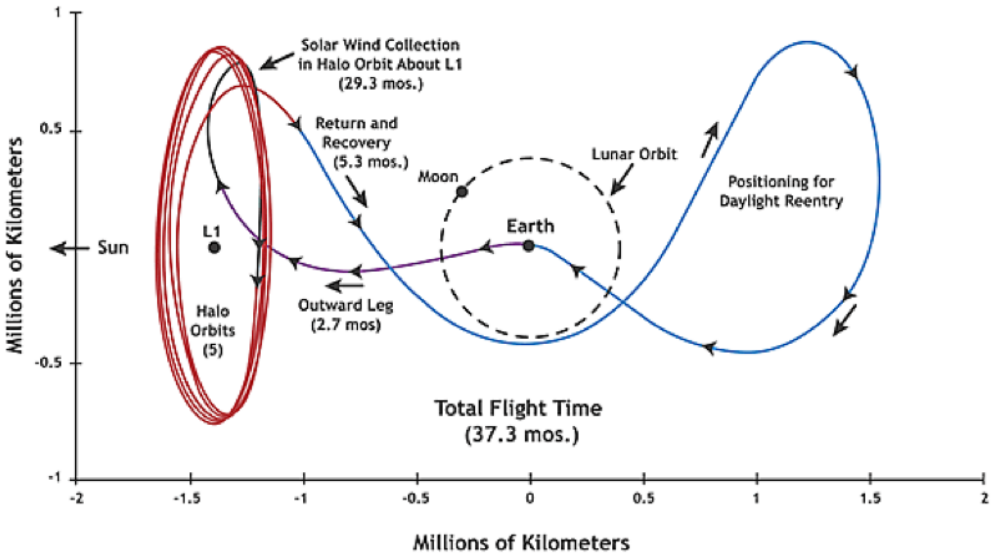

Figure 2-6: The Genesis Mission trajectory (as published by Lo et al. [67])

and return them to Earth. Beyond being the first robotic sample-return mission, the Genesis trajectory was the first to be designed using a Dynamical Systems Theory based approach — specifically by exploiting the operating halo orbit's stable and unstable manifold to achieve low-energy return and transfer trajectories to and from Earth. According to legend, Martin Lo of the NASA Jet Propulsion Laboratory contacted Professor Kathleen Howell of Purdue University and requested a low-energy mission trajectory within a week that satisfied the following high-level requirements: efficient travel to Sun-Earth L1, followed by entrance into a periodic orbit about the Lagrange point, maintenance for two years, and finally a daytime return to earth over Utah. The trajectory shown in Figure 2-6 was developed and delivered on time. By maximizing the use of low-fuel paths found in nature, the entire mission trajectory was designed to require only one deterministic fuel burn ($\Delta V \approx 6$ to $36 m/s$)[13], demonstrating not only tremendous fuel savings, but a leap in understanding of the interplay between three-body problem phase space and trajectory design.

[13]Note: this does not include stochastic corrective maneuvers.

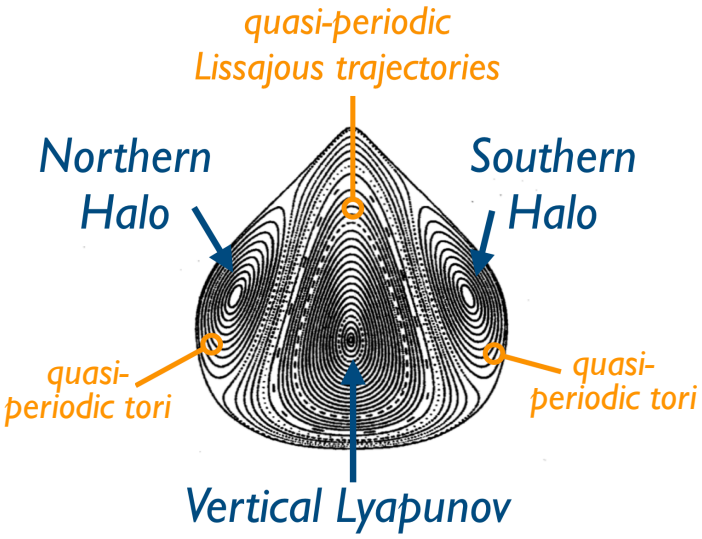

Figure 2-7: Poincaré surface of center manifold (modified from Goméz et al. [35])

2.3 The Center Manifold

As discussed in section 2.2.2, the six-dimensional CR3BP phase space — near the collinear libration points — is composed of solutions on the surface of stable and unstable invariant manifold subspaces, linked by the center manifold surface of bounded solutions. This is a consequence of the center × center × saddle stability structure shared by the equilibrium points and their periodic trajectories. In an investigation aimed at furthering qualitative intuition of their utility for trajectory design, Barden and Howell [9] surveyed the families of solutions that compose the center manifold and the natural phenomena yielding transition between them. Revealed is a rich, yet approachable trajectory design space ideal for use in a formation flying mission.

The constant energy Poincaré surface of the center manifold shown in Figure 2-7 (originally published by Goméz et al. [35]) summarizes their findings. Six classes of motion in three distinct regions of the phase space are qualitatively captured. Each region contains a periodic solution (single point on the surface). These are the vertical Lyapunov orbit in the center, northern halo orbit to the left, and south-

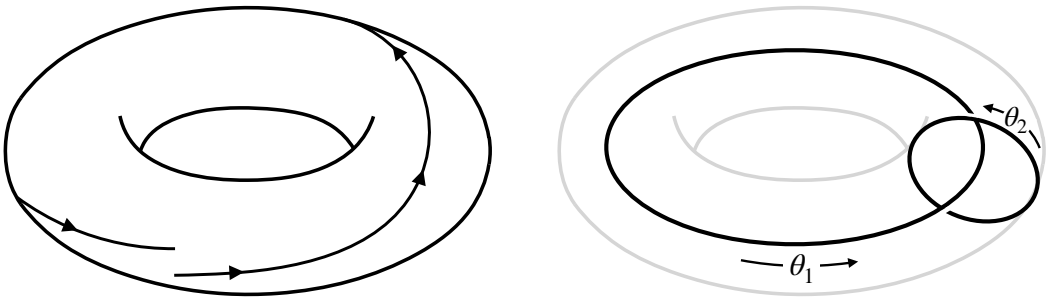

Figure 2-8: Parameterization of motion along a torus

ern halo orbit to the right[14] respectively. Additionally, quasi-periodic solutions — represented as rings surrounding the periodic orbits — exist in each region.

Similar to a coupled oscillator, quasi-periodic orbits (QPO) can be characterized by (at least) two fundamental frequencies of oscillation (ω_1 and ω_2)[15]. By definition, if $\omega_1/\omega_2 \in \mathbb{Q}$ (i.e., the ratio of the frequencies is rational) then the motion is periodic (returns to the same point after some time T). However, if $\omega_1/\omega_2 \notin \mathbb{Q}$ (i.e., the ratio is irrational), then the orbit will continuously wrap and wind — densely covering the surface of a torus — while never repeating or closing. This torus — parameterized by two angles (θ_1 and θ_2) — is the natural topological representation for these solutions: ($\mathbb{T}^2 : \theta_1 \times \theta_2$). As illustrated in Figure 2-8, while θ_1 — the longitudinal angle — wraps around its circumference, θ_2 — the latitudinal angle — winds inward, through it. Any point on the surface of the torus — and thus along the quasi-periodic trajectory — can be represented by these two angles. The solution of the dynamical flow under this topology is expressed simply as follows:

$$\mathbb{T}^2 : \begin{cases} \theta_1(t) = \theta_{1,0} + \omega_1(t - t_0) \\ \theta_2(t) = \theta_{2,0} + \omega_2(t - t_0) \end{cases} \tag{2.16}$$

[14]Note: the northern and southern halo orbit are mirror images of each other in the $x-y$ plane.
[15]The following discussion will be restricted to the case of QPO parameterized by two frequencies (and thus angles). As will be seen, the orbits/tori of interest are completely described as such.

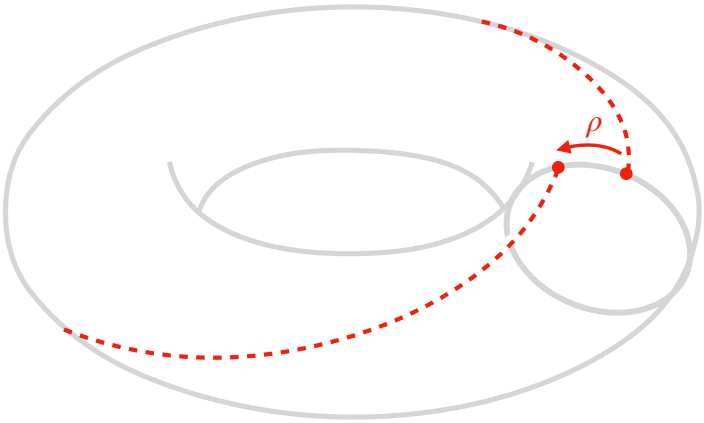

Figure 2-9: Illustration of invariant circle and torus number ρ

where $\theta_{1,0}$ and $\theta_{2,0}$ are the initial latitudinal and longitudinal phase of the solutions and t_0, the initial time. It is assumed that a diffeomorphism exists between the torus manifold and six-dimensional Euclidean space: $u(\theta_1, \theta_2) : \mathbb{T}^2 \rightarrow \mathbb{R}^6$.

Thus, the torus itself is an invariant surface of constant energy, yielding the condition that any solution starting on the torus remains on its surface for all future time. Since both ω_1 and ω_2 are fixed, it follows that the circular cross section of the torus — the set of all points for a fixed θ_1 — is invariant well. Consider the invariant circle illustrated in Figure 2-9. Each solution starting at $\theta_{1,0}$ and any initial phase $\theta_{2,0} \in [0, 2\pi)$, when propagated for the longitudinal period of revolution $T = 2\pi/\omega_1$, returns to the same invariant circle — i.e. $\theta_1(T) = \theta_{1,0} + \omega_1 T = \theta_{1,0} + 2\pi$ — but rotated in latitude by angle $\rho = \theta_2(T) - \theta_{2,0} = 2\pi(\omega_2/\omega_1)$. The latter quantity is termed the torus number.

2.4 Formation Flying Concept of Operations

This section sketches a Concept of Operations that exploits the natural solutions of the CR3BP center manifold subspace to solve the sequence of retargeting maneuvers for an observatory/external occulter formation flying mission. Both an extended

database — of periodic and quasi-periodic solutions — and optimal methods of transfer between these solutions are leveraged in this methodology. This CONOPS frames the approach of this thesis and guides the remaining chapters.

Exploiting the Center Manifold

As discussed in section 2.3, the center manifold subspace contains the families of bounded motion near the collinear Lagrange points. Since it is assumed that the spacecraft operate in the Sun-Earth L2 region, these solutions make ideal candidates to exploit for formation retargeting. Recall that each of the three-dimensional periodic halo orbits is surrounding by dense families of quasi-periodic invariant tori (QPT). Thus, it can be conceived that by constraining the external occulter motion to quasi-periodic solutions on the surface of a torus relative to the observatory — operating from the inner halo orbit — the formation line-of-sight (LOS), and therefore its trace on the celestial sphere, can be commanded. Ultimately, this is the mission-wide objective for the retargeting maneuvers: to trace a path[16] on the celestial sphere that intersects all of the target stars in the DRM. Figure 2-10 provides an illustration of the formation exploiting the center manifold solutions.

Extending the Design Database

Quasi-periodic tori exist within dense families of varying amplitude from their parent halo orbit. By strategically planning efficient transfers between solutions on a given torus, the external occulter can effectively manipulate the formation LOS vector to reach multiple regions of the celestial sphere, and therefore target select stars[17]. In 2018, Restrepo and Russell [88] publicly released a database of thousands of planar periodic orbits created from a meticulous and efficient search through several three-body systems. Two-dimensional orbits do not possess the characteristic

[16]That is, from the formation LOS. For a more detailed description, see section 2.5.1.

[17]It is assumed that the observatory is evolving kinematically in this scenario — refer to section 4.1

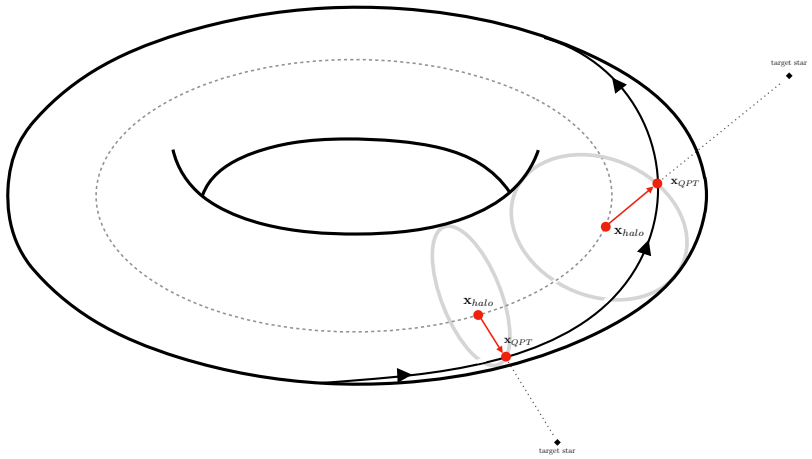

Figure 2-10: Illustration of the formation exploiting the center manifold solutions.

motion sought for this application. However, it is well documented how to take planar periodic orbits and find the bifurcation condition from which three-dimensional orbital families are born[18]. A contribution of this work[19] will be to develop a three-dimensional periodic orbit and quasi-periodic torus database. This catalog will provide quick and efficient search for pathways of the formation LOS across the celestial sphere, as well multitudes of additional natural structures that can be explored in future work[20].

Transferring Between Solutions

Efficient means of transfer by the external occulter between solutions on the torus (as provided by the design database) needs to be determined if low-fuel tracings through the target stars of a DRM is to be achieved. It is shown in section 4.1 that this tasks corresponds to manipulating the latitudinal angle θ_2 of the torus topology. It is expected that a future, fully explored set of retargeting maneuvers that optimally solve a DRM, will require a combination of both manipulating this latitudinal phase on a single torus' surface, and transfer of the formation between

[18]For a further discussion of these topics, refer to section 3.1.3
[19]Refer to Chapter 3 for a full description.
[20]For suggestions, refer to section 6.2.

different halo/QPT combinations to open different pathways[21].

2.5 Simulating the Formation

The following section demonstrates how to convert the formation LOS to angle coordinates in the ICRS J2000 frame. For any time t, the distance from the observatory operating on the halo orbit and the external occulter on the quasi-periodic torus is defined as follows

$$\mathbf{r}(t) = \mathbf{x}_{QPT}(t) - \mathbf{x}_{halo}(t) \tag{2.17}$$

where $\mathbf{x}_{QPT}$ and $\mathbf{x}_{halo}$ are the position of the occulter and the observatory respectively, expressed in the three-body system rotating frame. The formation LOS vector in the rotating frame is simply the unit vector along the distance vector.

$$\hat{\mathbf{r}}(t)_{LOS} = \frac{\mathbf{r}(t)}{|\mathbf{r}(t)|} \tag{2.18}$$

The ICRS J2000 frame is aligned with the Earth equatorial coordinate system. Thus, two rotations are required to convert the formation LOS vector to the inertial frame. The first converts the vector from the rotating frame at time t to the inertial frame with origin at the three-body system barycenter and aligned with the ecliptic plane. The second rotates the resulting frame about its x-axis 23.43 degrees, to align with the Earth equatorial frame. It is assumed that the target stars are at a great enough distance that their location on the celestial sphere is unchanged from the perspective of a reference frame centered at the three-body system barycenter and one at Earth. These rotation matrices and their subsequent application to the formation LOS vector are expressed below:

[21]Note, searching across different halo orbits and tori is beyond the scope of this thesis, and is left for future work. Refer to section 4.1 and 6.2.

$$R_{eq} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos(23.43^o) & -\sin(23.43^o) \\ 0 & \sin(23.43^o) & \cos(23.43^o) \end{bmatrix}, \qquad R = \begin{bmatrix} \cos(t) & -\sin(t) & 0 \\ \sin(t) & \cos(t) & 0 \\ 0 & 0 & 1 \end{bmatrix} \tag{2.19}$$

$$\hat{\mathbf{r}}^I_{LOS} = R_{eq} R \hat{\mathbf{r}}_{LOS} \tag{2.20}$$

where R and R_{eq} are the first and second rotation matrices respectively and $\hat{\mathbf{r}}^I_{LOS}$ is the formation LOS vector expressed in the ICRS J2000 inertial frame. This vector can be converted to angle coordinates — longitude (θ_1) and latitude (θ_2) — via the following formulae:

$$\theta_1 = \theta_{long} = \tan^{-1}\left(\frac{\hat{r}^I_y}{\hat{r}^I_x}\right) \tag{2.21}$$

$$\theta_2 = \theta_{lat} = \tan^{-1}\left(\frac{\hat{r}^I_z}{\sqrt{(\hat{r}^I_x)^2 + (\hat{r}^I_y)^2}}\right) \tag{2.22}$$

where $\hat{r}^I_x$, $\hat{r}^I_y$, and $\hat{r}^I_z$ are the x, y, and z components of the formation LOS vector expressed in the ICRS J2000 inertial frame, respectively. Since the location of all target stars are also expressed in this thesis by a longitude-latitude pair in the ICRS J2000 inertial frame, the result of equations 2.21-2.22 can be considered the projection of the formation LOS vector onto the celestial sphere.

2.5.1 Interpreting Formation Flying Plots

The main visualization tool used to interpret the results of this thesis is the plot of the formation's line-of-sight tracing across the celestial sphere. Figure 2-11 is provided as an example. A trace is defined as the time evolution of the formation LOS vector, rotated to the J2000 inertial frame, and projected to angle space (lat-

itude, longitude) on the celestial sphere. In general, multiple traces are shown per plot (expressed as the thin colored lines in the figure), each representing a simulation starting from a different initial condition by the external occulter, on the torus. The gold line specifically represents the projection of the initial invariant circle — from the perspective of the formation — onto the celestial sphere. All QPT solutions start on this invariant circle, separated by phase angle, thus all traces start from this line[22]. These concepts are further developed in section 4.1. In the example plot provided, the trajectories are simulated for a quarter-period of revolution and the set of target stars represented by red dots.

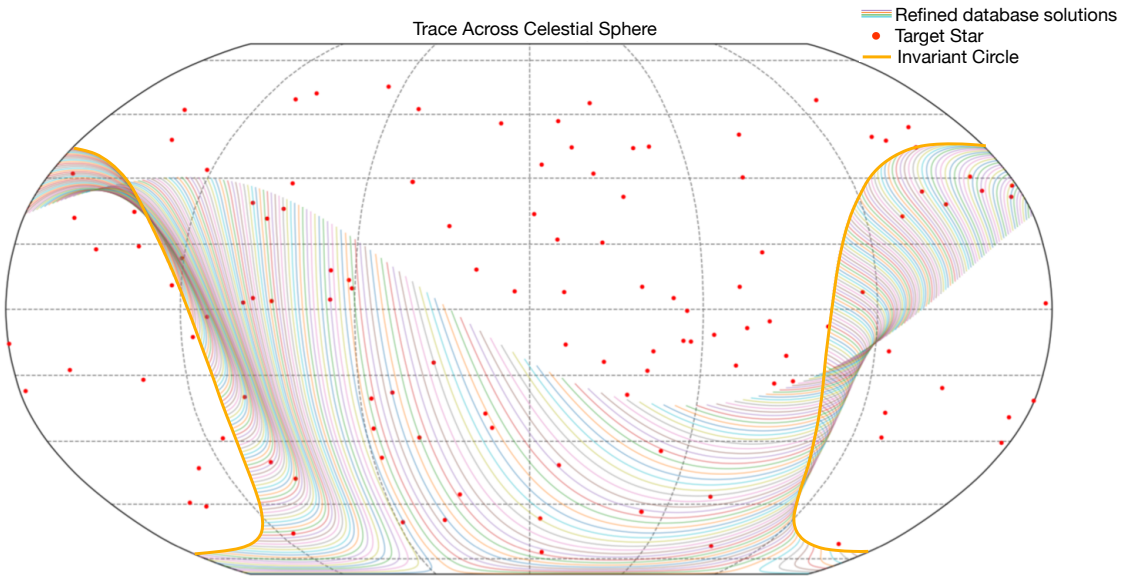

Figure 2-11: Example plot illustrating multiple LOS tracings on the celestial sphere.

[22]It should be emphasized that the stars are static in this visualization (since it is based on an inertial frame of reference). However, the same is not true for the torus surfaces and halo orbits, and thus, their simulated formation LOS tracings on the celestial sphere. They are derived from a rotating frame that can be related to the inertial frame through a rotation matrix that is a function of time. Thus, the exact shapes and geometry in which they manifest on the plot is dependent on the initial time t_0 of the simulation, as well as the relative position of the spacecraft within that rotating frame. This t_0 anchors the initial relative position of the rotating frame with respect to the inertial frame.

Chapter 3

Database Development

This chapter details the design and development of a three-dimensional periodic orbit and quasi-periodic torus database. Recall that in the CR3BP phase space, the center manifold contains the invariant set of all bounded solutions near the collinear libration points[1]. In Chapter 2, a Concept of Operations was introduced that identified the natural periodic and quasi periodic motion within this subspace — specifically halo orbits and their surrounding quasi-periodic tori — as ideal candidates to be exploited by a formation flying observatory and external occulter spacecraft operating in the SEL2 environment, respectively.

Solutions on the manifold exist within dense families [8], highlighting the need for infrastructure to efficiently manage and explore the trajectory design space. In response, a database of three-dimensional periodic and quasi-periodic orbits was created, equipped with a software suite that allows mission designers to rapidly retrieve and analyze candidate solutions that meet a mission objective. The toolkit, developed in the Julia language, incorporates a robust and diverse set of differential corrector algorithms (to converge solutions to the precision needed for numerical propagation in the chaotic phase space), quasi-periodic orbit solvers, adaptive continuation algorithms (to extend orbital searches along families), robust stability

[1]Refer to section 2.3 for an in-depth discussion.

index tracking for phase space bifurcation detection, period-doubling bifurcation solvers, and a manifold visualizer (projected onto configuration space) with meshing algorithm, among others. The database was designed with general mission constraints in mind and includes over 23,000 periodic orbits — including both northern and southern families of halo orbits, near rectilinear halo orbits, butterfly orbits, and planar Lyapunov orbits within the Sun-Earth L2 & L1, and Earth-Moon L2 & L1 regions — as well as over fifty million quasi-periodic solutions existing on the tori surrounding the SEL2 halo orbits. Details on the development of the periodic orbit and quasi-periodic torus database are presented in sections 3.1 and 3.2 respectively.

This work was inspired by Restrepo and Russell [88] who performed a global search for planar periodic orbits along the CR3BP phase space of multiple celestial systems, uncovering and categorizing several types of families. This study does not perform the same exhaustive search, but instead focuses on specific three-dimensional families of interest. However, as the authors' point out, three-dimensional periodic orbits are birthed from planar periodic orbits — which can be identified via known bifurcation conditions (refer to section 3.1.3). Thus, one can consider this contribution as a minor extension of their work.

3.1 Generating Periodic Solutions

This section documents the theory and methods employed to develop the periodic orbit database. Solutions are stored by their initial conditions — converged to thirteenth-order accuracy[2] — by the method of differential corrections (section 3.1.1). Section 3.1.2, describes the method in which solutions are continued across families, while section 3.1.3 details how families of interest are targeted by their bifurcation conditions, signaled by the tracked stability indices of the solutions across their families. Finally, section 3.1.4 provides important details regarding software

[2]In this chapter, "order of accuracy" will refer to the number of digits of precision converged for a given solution.

implementation, demonstrating the toolkit created for interacting with the database.

3.1.1 Method of Differential Correction

Recall from section 2.1, that any trajectory in the phase space can be uniquely determined through knowledge of its state $\mathbf{x}(\tau) = [x,\ y,\ z,\ \dot{x},\ \dot{y},\ \dot{z}]$ at any time τ along the flow. This allows for efficient storage of periodic solutions in the database, which are saved by their converged initial conditions $\mathbf{x}_0$[3] and period of revolution T. The orbits themselves possess center $\times$ center $\times$ saddle stability. Therefore, small variations in the state, including unavoidable numerical inaccuracies, will cause the periodic solutions to diverge exponentially in time — when propagated — along its unstable manifold. Thus, it is necessary for the initial states be computed to a high level of accuracy in order to delay noticeable deviations long enough for the propagated trajectory to be of general utility. As a rule of thumb, it has been found that convergence to at least tenth-order is required, which usually allows for multiple periods of revolution. All entries in the periodic orbit database have been solved to $\mathcal{O}(10^{-13})$.

This was accomplished by the method of differential correction (also called shooting method). As an overview, the algorithm — seeded with a guess of the initial state — targets a desired end condition[4] of the trajectory at a future (potentially unknown) point in time. The initial state is then incrementally corrected at each iteration — via predicted variations of the flow downstream — until the desired end condition is met within some error tolerance. Several comprehensive resources exist on differential correction applied to the CR3BP phase space [39, 48, 61]. The reader is referred to these texts for details that may be omitted from this brief overview.

[3]Note, solving a trajectory (or solution) will refer to computing its initial conditions from this point forward. Additionally, it should be noted that since the CR3BP phase space is time invariant, the initial time τ_0 can be set equal to zero without loss of generality, and hence, neglected.

[4]For example, periodicity.

Approximating variations along the trajectory

The differential corrector method uses predicted variations of the flow along a trajectory to inform the algorithm at each iteration how to adjust its initial conditions to achieve the desired terminal state. A linear approximation of this relationship, following the outline provided by Grebow [39][5], is derived below.

The objective is to estimate the variation of the state of the trajectory at time τ_{i+1} — where $\mathbf{x}_{i+1} = \mathbf{x}(\tau_{i+1})$ — given the variation applied at a previous time $\mathbf{x}_i = \mathbf{x}(\tau_i)$. It is assumed that $\Delta\tau = \tau_{i+1} - \tau_i > 0$. Thus, $\mathbf{x}_{i+1}$ represents the solution of the flow with initial condition $\mathbf{x}_i$ evaluated for $\Delta\tau$. This relationship is expressed in equation 3.1.

$$\mathbf{x}_{i+1} = \phi(\Delta\tau; \mathbf{x}_i) \tag{3.1}$$

It follows that the variation of the terminal state can be expressed via equation 3.2.

$$\delta\mathbf{x}_{i+1} = \phi(\Delta\tau + \delta\tau_{i+1}; \mathbf{x}_i + \delta\mathbf{x}_i) - \mathbf{x}_{i+1} \tag{3.2}$$

where $\delta\mathbf{x}_i$ represents the variation in the initial conditions and $\delta\tau_{i+1}$, the variation in the final time required to reach the new final condition. A Taylor series expansion performed on equation 3.2 yields the following:

$$\mathbf{x}_{i+1} + \delta\mathbf{x}_{i+1} = \phi(\Delta\tau; \mathbf{x}_i) + \left.\frac{\partial\phi}{\partial\mathbf{x}_i}\right|_{\mathbf{x}_i,\Delta\tau} \delta\mathbf{x}_i + \left.\frac{\partial\phi}{\partial\tau}\right|_{\mathbf{x}_i,\Delta\tau} \delta\tau_{i+1} + \text{H.O.T.} \tag{3.3}$$

where by definition, $\left.\frac{\partial\phi}{\partial\mathbf{x}_i}\right|_{\mathbf{x}_i,\Delta\tau} = \Phi(\tau_{i+1}, \tau_i)$ — i.e., the state transition matrix eval-

[5]For further details, including how to handle contemporaneous and non-contemporaneous variations in the state, refer to this work.

uated from τ_i to τ_{i+1} — and $\left.\frac{\partial \phi}{\partial \tau}\right|_{\mathbf{x}_i, \Delta\tau} = f(\mathbf{x}_{i+1}) = \dot{\mathbf{x}}_{i+1}$ — the vector flow at τ_{i+1}. By neglecting higher order terms and substituting equation 3.1 into 3.3, the linear approximation to the variations of the flow is obtained (equation 3.4).

$$\delta \mathbf{x}_{i+1} = \Phi(\tau_{i+1}, \tau_i)\delta \mathbf{x}_i + \dot{\mathbf{x}}_{i+1}\delta\tau_{i+1} \tag{3.4}$$

Differential Correction Strategy

The differential correction strategy manipulates the initial conditions at each iteration, guided by the variational relationship provided in equation 3.4, to achieve a desired terminal state. In the space about the libration points, Roy and Ovenden's Periodicity Theorem [48] informs that three-dimensional periodic solutions are symmetric about the $x - z$ plane [48]. Exploiting this knowledge reduces the degrees of freedom of the problem, greatly simplifying its method of solution. This is demonstrated below.

All initial conditions of the periodic solutions are constrained to lie on the plane of symmetry, fixing both $y_0 = 0$ and its variation, $\delta y_0 = 0$. Furthermore, the mirror condition ensures that the orbit passes through the symmetric plane perpendicularly, fixing both $\dot{x}_0 = \dot{z}_0 = 0$ as well as their variations, $\delta\dot{x}_0 = \delta\dot{z}_0 = 0$. Thus, the initial conditions take the following form:

$$\mathbf{x}_0 = [x_0,\ 0,\ z_0,\ 0,\ \dot{y}_0,\ 0] \tag{3.5}$$

In the formulation, $\dot{y}_0$ is free to change (i.e., be corrected), as well as one of the two remaining states. The other, — either x_0 or z_0 — must be fixed (its variation set equal to zero). An added benefit of exploiting symmetry is that only half the orbit needs to be computed. Thus the initial condition is propagated until the $x-z$ plane is intercepted (for the first time): this state is $\mathbf{x}_f$. Fixing x_0, substituting the symmetry conditions into equation 3.4, and taking the inverse of the resultant

matrix yields the updated variational relationship expressed in equation 3.6.

$$\begin{bmatrix} \delta x_0 \\ \delta \dot{y}_0 \\ \delta \tau \end{bmatrix} = \begin{bmatrix} \frac{\partial y_f}{\partial x_0} & \frac{\partial y_f}{\partial \dot{y}_0} & \dot{y}_f \\ \frac{\partial \dot{x_f}}{\partial x_0} & \frac{\partial \dot{x}_f}{\partial \dot{y}_0} & \ddot{x}_f \\ \frac{\partial \dot{z_f}}{\partial x_0} & \frac{\partial \dot{x}_f}{\partial \dot{y}_0} & \ddot{z}_f \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ \delta \dot{x}_f \\ \delta \dot{z}_f \end{bmatrix} \tag{3.6}$$

Periodicity is achieved when the terminal state crosses the $x-z$ plane perpendicularly — i.e., $y_f = \dot{x}_f = \dot{z}_f = 0$. Note that the condition on y_f is enforced by the propagation event handler. In general, the initial conditions guessed at the onset of the algorithm will yield non-zero $\dot{x}_f,\ \dot{z}_f$. If this is the case, the results from equation 3.6 are used to update the initial conditions on x_o and $\dot{y}_o$, as shown in equation 3.7

$$\begin{aligned} x_{0,c} &= x_0 - \delta x_0 \\ \dot{y}_{0,c} &= \dot{y}_0 - \delta \dot{y}_0 \end{aligned} \tag{3.7}$$

This process is repeated, starting by propagating the corrected initial conditions until the $x-z$ plane is intercepted, until both $|\delta \dot{x}_0| < \epsilon$ and $|\delta \dot{z}_0| < \epsilon$. Where ϵ is some user-defined error tolerance.

Alternatively, if z_0 is fixed instead of x_0, the equations 3.6 and 3.7 are updated to equations 3.8 and 3.9 respectively.

$$\begin{bmatrix} \delta z_0 \\ \delta \dot{y}_0 \\ \delta \tau \end{bmatrix} = \begin{bmatrix} \frac{\partial y}{\partial z_0} & \frac{\partial y}{\partial \dot{y}_0} & \dot{y} \\ \frac{\partial \dot{x}}{\partial z_0} & \frac{\partial \dot{x}}{\partial \dot{y}_0} & \ddot{x} \\ \frac{\partial \dot{z}}{\partial z_0} & \frac{\partial \dot{x}}{\partial \dot{y}_0} & \ddot{z} \end{bmatrix}^{-1} \begin{bmatrix} 0 \\ \delta \dot{x}_f \\ \delta \dot{z}_f \end{bmatrix} \tag{3.8}$$

$$\begin{aligned} z_{0,c} &= z_0 - \delta z_0 \\ \dot{y}_{0,c} &= \dot{y}_0 - \delta \dot{y}_0 \end{aligned} \tag{3.9}$$

It should be noted that studies applying multiple shooting (two-level) differential

correction schemes can often times be found in the literature. This alternative formulation is more robust than the single-shooting method and can help solutions converge in particularly sensitive regions of the phase space, such as near singularities. However, this was not considered for this application as the simpler single-shooting method met the database objective without any issues.

3.1.2 Continuation

The objective of the database is to compute a discrete sampling of solutions across several families of interest. The process of finding and computing nearby family members from a previously converged orbit is called continuation. This is performed for periodic solutions by simply linearly extrapolating from prior two converged initial conditions, by some small perturbation ε along the free states of the differential corrector method[6]. For example, if x_0 was fixed, z_0 and $\dot{y}_0$ are updated as follows:

$$z_{0,new} = \frac{z_{0,2} - z_{0,1}}{x_{0,2} - x_{0,1}}\varepsilon + z_{0,2} \tag{3.10}$$

$$\dot{y}_{0,new} = \frac{\dot{y}_{0,2} - \dot{y}_{0,1}}{x_{0,2} - x_{0,1}}\varepsilon + \dot{y}_{0,2} \tag{3.11}$$

$$\mathbf{x}_{0,new} = [x_{0,2} + \varepsilon,\ 0,\ z_{0,new},\ 0,\ \dot{y}_{0,new},\ 0] \tag{3.12}$$

Since ε is applied directly to the fixed state, that dimension is considered the search direction. The differential corrector is then applied, with $\mathbf{x}_{new}$ as the initial guess, to converge the new family member to the order of accuracy required. When less than two solutions are present in the database, the continuation method is initialized by perturbing the initial solution along the search direction only.

The fixed state is always set as the most rapidly changing of the two dimensions. As consequence, the search direction is adaptive. This check — simply determining

[6]Note, more complicated and robust methodologies could be utilized — such as cubic spline interpolation — as well as techniques to handle reflections (i.e., cases when families fold back on themselves). However, these paths were not pursued in this thesis.

the greater of the two: $|x_{0,2} - x_{0,1}|$ or $|z_{0,2} - z_{0,1}|$ — is performed after converging a solution and before continuation.

3.1.3 Stability of Periodic Solutions and Bifurcations

Section 3.1.2 discussed how to continue solutions to obtain a discrete sampling of orbits across a family. It was shown that the parameter ε ultimately determines the fineness of the database. Lowering the value yields a denser representation of the available solutions within a family, but at the cost of computing additional members. Its selection should reflect this trade-off. On the contrary, loss of resolution caused by raising the value of ε can be aided by the software infrastructure. Section 3.1.4 demonstrates the suite of tools developed to both search among the discrete entries of the database and interpolate solutions within its gaps.

It remains to be seen how specific families of periodic orbits can be targeted within the phase space. Recall from section 2.2, the stability of a periodic orbit is determined from the eigenvalues of the monodromy matrix — which is simply the state transition matrix evaluated after one period $\Phi(T, 0)$ for a point along the orbit. These eigenvalues, six in total, come in three reciprocal pairs — $(\lambda_i, 1/\lambda_i)$, $i = 1, 2, 3$ — which can be summarized by the set of stability indices defined below[7]:

$$b_i = \lambda_i + \frac{1}{\lambda_i} \tag{3.13}$$

Note, since the monodromy matrix is a discrete map for a point along the orbit, stability is determined by where the eigenvalues lie in the complex plane with respect to the unit circle. The stability indices fully capture this information. It follows that $|b_i| > 2$ corresponds to a pair of eigenvalues where one is stable and the other

[7]Many texts exist defining and working with stability indices in the CR3BP phase space. Refer to Hénon [43], Howell & Campbell [49], and Papadakis and Zagouras [82] as a start.

unstable, $|b_i| \leq 2$ to a pair of center eigenvalues lying on the unit circle, and $|b_i| = 2$ specifically to the case of two eigenvalues lying jointly on the real axis at either the point (1,0) or (-1,0). This final case is of critical importance. It will be shown that it is by tracking how the stability indices change and pass through this condition across a family that the birthing of new types of solutions can be predicted.

Bifurcations

In Dynamical Systems Theory, a bifurcation is a structural change in the system's phase space caused by varying a parameter — such as Jacobi constant or orbital amplitude — past some critical value. In general, crossing the bifurcation condition[8] alters some stability characteristic of the flow, such as that of periodic solutions and/or equilibrium points. In some applications this can be destructive, as stable equilibrium points can abruptly switch to become unstable — for example, when increasing the axial load on a beam until it buckles. In other cases, a bifurcation can mean the birth of new, and potentially useful, classes of motion. In the CR3BP phase space, bifurcations can be triggered by varying the orbital amplitude/distance away from the Lagrange point[9] as is naturally done through continuation. At times, new families of periodic orbits are birthed through this process.

It is the stability indices that signal bifurcation conditions as they are tracked across a family [39, 43, 49]. For the classes of solution of interest, this occurs when any of the indices crosses through $|b_i| = 2$. This translates to the collision of two eigenvalues at the point (1,0) or (-1,0) on the complex plane — for $b_i = 2$ and $b_i = -2$ respectively. If $|b_i|$ crosses in the direction of increasing value then the pair previously corresponding to a center mode slide along the unit circle (as complex conjugates) until colliding on the real axis and splitting off as a stable/unstable pair. This case is illustrated in Figure 3-1. The reverse case occurs when crossing $|b_i| = 2$

[8] That is, the critical value of the parameter.

[9] Simply x_0 and z_0

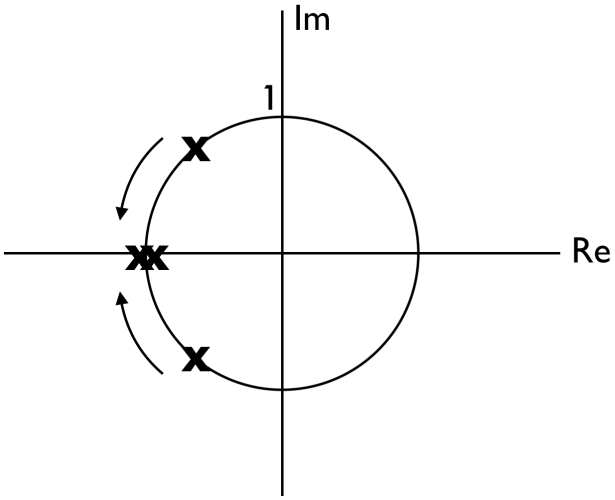

Figure 3-1: Collision of eigenvalues at (1,0) on the complex plane as the stability index crosses $|b_i| = 2$.

from the opposite direction.

It is well known that three-dimensional halo orbits bifurcate from the planar Lyapunov family in this way. As the orbital amplitude of the planar orbits increase one of the stability indices crosses the bifurcation condition. The set of initial conditions at which this occurs can be solved within a specified error tolerance using a bisection search method. Once found, the initial condition is perturbed in the direction of the eigenvector corresponding to the two eigenvalues that collided. This value is seeded as the initial guess for the differential corrector, which then converges to the first member of the halo orbit family[10].

The near rectilinear halo orbits do not explicitly bifurcate from the halo orbits, but instead represent of a subset of the family, classified as those with all stability indices within a small bound of ± 2 [119].

Butterfly orbits are birthed from the near rectilinear halo orbit family of the Earth-Moon system under a period-doubling bifurcation near the six-day-period

[10]Currently this process of switching families is done manually, after the continuation procedure across a family of solutions is completed. Automating this task so is left for future work.

members [40]. Again the bifurcation condition can be found by observing when $|b_i| = 2$ is crossed. A similar procedure to the halo orbits is followed where the guess of the initial conditions for the first orbit is found by perturbing the NRHO solution at the bifurcation condition along the corresponding eigenvector.

On a practical note, the stability indices need to be tracked, and thus labeled, consistently when generating the database (i.e., b_1 is appropriately labeled b_1, b_2 is labeled b_2, etc.) so that there are no discontinuities in the values. This is accomplished in the continuation procedure by comparing both the value and rate of change (via a finite difference) of the indices to the two previous database entries. This prevents false positives if any cross each other in value.

3.1.4 Software Implementation

The software infrastructure developed to both build and interact with the periodic orbit database was written in the Julia language. Julia is a dynamically typed, general-purpose programming language built with high performance in mind [15]. Through its high-level syntax, developers can reap the productivity of popular scripting languages such as Python or MATLAB. However its design, which interplays carefully selected technologies such as multiple dispatch and just-in-time (JIT) compilation, allows the code to run with the performance of a statically-typed, compiled language, such as C or Fortran.

The field of astrodynamics has strong, historical links to traditional scientific computing languages (e.g., Fortran). However, the advantages of Julia are hard to ignore. A study performed by Eichhorn et al. [26] compared the performance of six programming languages (Fortran, C++, Java, MATLAB, Python, and Julia) applied to four classic astrodynamics problems. The authors found that Julia overall offered comparative performance to the compiled languages for the cases analyzed, surpassing interpreted languages on all accounts. These reasons supported

the decision to work within the Julia framework[11]. As a final note, Verner's "Most Efficient" 9/8 Runge-Kutta method — called through the Julia "DifferentialEquations.jl" package [86] (which also provides built-in event handling) — was used for numerical propagation throughout this thesis.

Database Interaction and Visualization

The final periodic orbit database is composed of over twenty-three thousand solutions ($\varepsilon = 10^{-4}$)[12]. Represented are the planar Lyapunov, northern halo orbit, and southern halo orbit families of the SEL1, SEL2, EML1, and EML2 regions, as well as the northern and southern butterfly orbits near EML2. As discussed in section 3.1.3, near rectilinear halo orbits are subsets of the halo orbit families, and are thus included by default. For each solution, its six-dimensional (initial) state on the rotating frame's $x - z$ plane (converged to thirteenth-order accuracy), orbital period, Jacobi constant, and stability indices are recorded. Several elements of the software toolkit developed to interact with the database are demonstrated in the following subsections.

Visualizing the families

The orbital families — projected onto three-dimensional configuration space — can be visualized by creating a triangular mesh between adjacent orbits in the database. This is accomplished by constraining the numerical propagator to solve for an equal number of points along each trajectory, and ordering the resultant states along the paths into a sequence of mesh vertices, as illustrated in Figure 3-2. This result can

[11] It should be noted, that it is the position of the author that continued research and development on the application of the Julia to problems in astrodynamics is needed before any sort of formal conclusions on performance (as compared to compiled and scripting languages) are drawn.

[12] This refers to the step size during continuation, section 3.1.2. To give a sense of the computational cost, several hundreds of converged solutions were achieved across the halo orbit families within minutes. This was performed on a laptop computer equipped with a 2.8 GHz Quad-Core Intel Core i7 processor and 16GB of memory.

then be passed to a 3D mesh plot function in a graphing library of choice. Continuing this method across a family produces a surface plot of the set of solutions. Several examples from the database are provided in Figure 3-3.

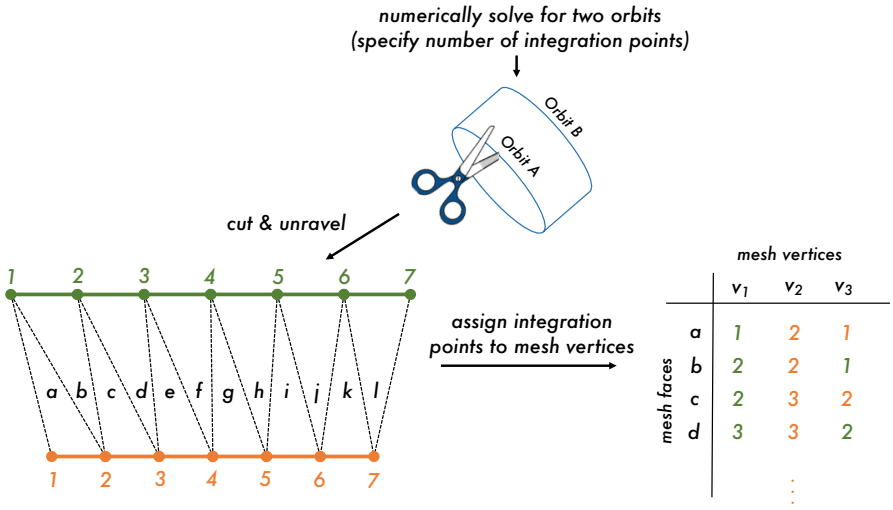

Figure 3-2: Illustration of triangular meshing between orbits methodology

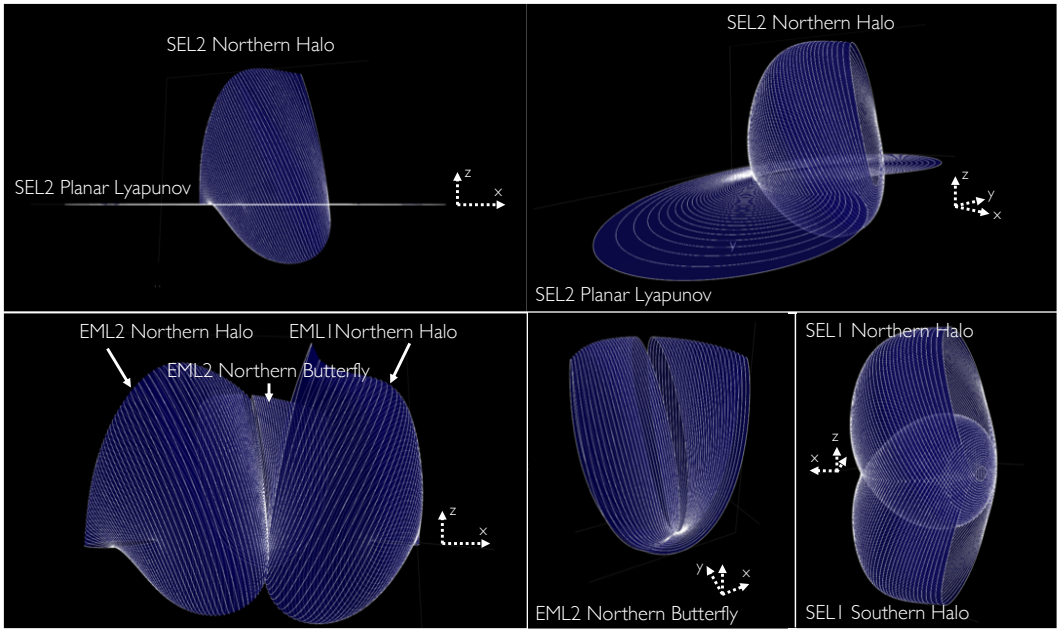

Figure 3-3: Visualization of several orbital families, projected onto the configuration space of the CR3BP rotating frame

Plotting data across a family and searching for solutions

The software structure allows mission designers to plot parameters stored in the database across a family with a single function call. Figure 3-4 displays the plots of x_0 versus b_2 (left) and x_0 versus T (orbital period) (right), for the SEL2 Northern Halo family.

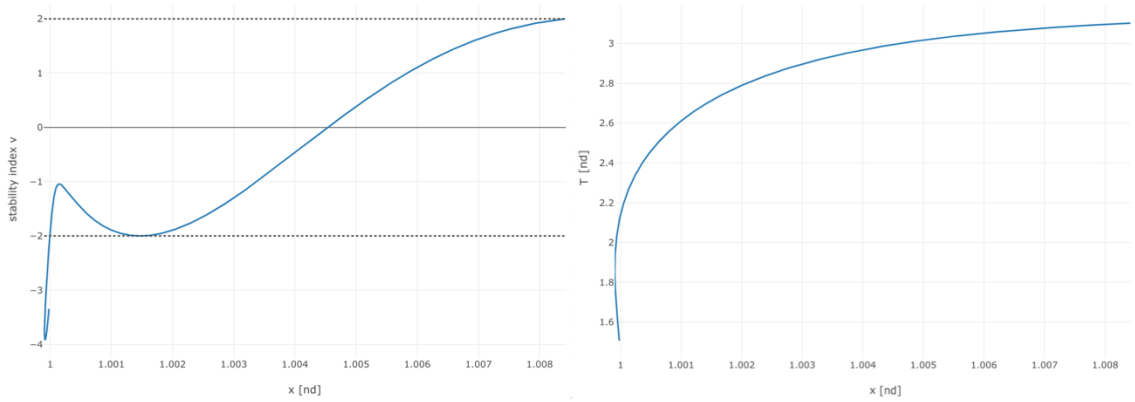

Figure 3-4: Plotting data across the SEL2 Northern Halo family

A "search periodic orbit database" function call allows orbits of interest — selected from the graphs — to be extracted from the database. The function is passed the family name, parameter of choice, value requested, and a series of flags and returns either the closest solution stored in the database matching the value, or solves for the orbit within a specified error tolerance using a bisection search method combined with a differential corrector. With this functionality, the solutions(s) matching the bifurcation condition $|b_i| = 2$ can be extracted.

Finding Stable and Unstable Manifold Solutions

As a final note, the stable and unstable manifold solutions of unstable periodic orbits can be generated with a call to a single black box function. Figure 3-5 displays the unstable and stable manifold generated for an EML2 and EML1 halo

orbit respectively. The intersection of these manifold with the $y-z$ plane (rotating frame) is shown in Figure 3-6. This data can be used to design transfer maneuvers through the Earth-Moon region.

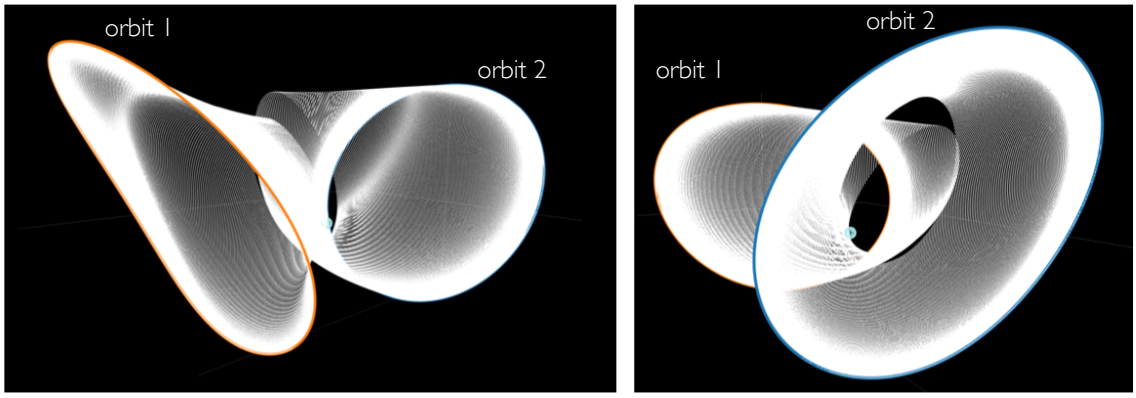

Figure 3-5: Unstable and stable manifolds for halo orbits in the EML2 (orbit 1) and EML1 region (orbit 2) respectively.

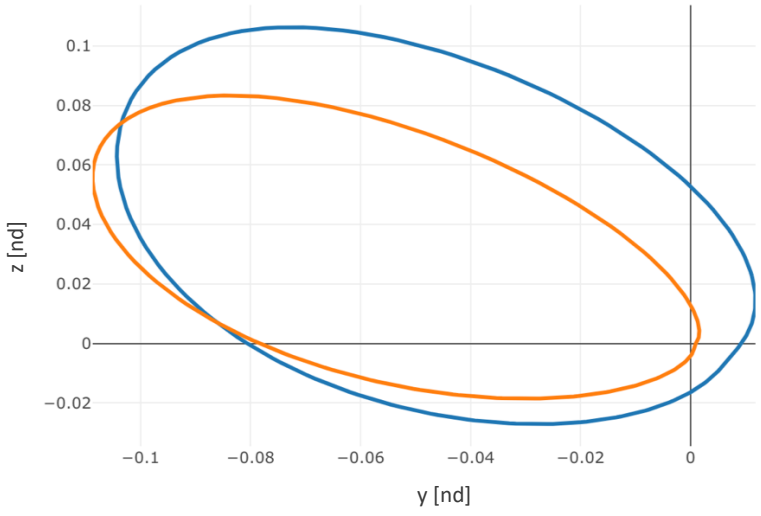

Figure 3-6: Intersection of orbit 1's unstable manifold (blue) and orbit 2's stable manifold (orange) on the rotating frame's $y-z$ plane.

3.2 Generating Quasi-Periodic Invariant Tori

This section describes the development of the quasi-periodic torus (QPT) database. Recall from section 2.3, the center manifold is dense with quasi-periodic solutions — on the surface of tori — that surround[13] the periodic orbits. Of these quasi-periodic tori, numerous types of families can exist. In an attempt to manage this large design space, the database was restricted to only the region of interest for the formation flying mission. Thus, only the QPT solutions surrounding the SEL2 northern halo orbits were generated, and of those, only the constant period — equal to that of their parent halo orbit — family computed[14]. It will be shown that even when narrowing the scope of computation to this subset, the resultant QPT database grows vastly greater than the periodic orbit database. In section 3.2.1 the GMOS method of computing QPT solutions — accomplished by pairing a differential corrector with the stroboscopic mapping condition of the tori's invariant circle — is documented. Finally, section 3.2.2 provides a discussion on the specifics of software implementation, including exploiting high-performance computing and database interaction and visualization.

3.2.1 GMOS: The QPT Solution Algorithm

This section presents the GMOS method implemented to compute the QPT solutions of the quasi-periodic orbit database. As documented by Baresi [13], two notable procedures exist that exploit the method of differential corrections to converge toroidal solutions in the phase space of interest[15]. Named after their originators, these are the KKG [60] and GMOS [12, 81] algorithms. In his study, Baresi performed a detailed

[13]This is specifically for the case of halo orbits. The tori of Lissajous trajectories are related (and near) the planar Lyapunov and vertical period orbits but do not specifically surround them.

[14]For further explanation of why only this family was considered, refer to sections 4.1 and 5.1.

[15]Note, the author also presents two partial differential equation solvers. However, in order to maintain consistency with the periodic orbit database solution methods, these were not explored in this thesis.

comparison of the two methods concluding that the GMOS algorithm is superior in both flexibility and accuracy. It should also be noted how the two algorithms differ in their design. While KKG employs surface of sections to characterize the evolution of solutions on the torus, GMOS relies on stroboscopic mappings, which due to the properties of the invariant circle — as will be shown — is a more natural framework for parameterizing the problem.

Several resources exist in the literature that comprehensively describe the GMOS algorithm [12–14, 75, 81]. The reader is referred to these texts for any details that may be either omitted in this review or require further clarification[16].

Initializing the Solver

Recall from section 2.3 that the torus is invariant, meaning that all solutions starting on its surface remain on the surface for all future time. This behavior is best characterized by the torus' cross section — *the invariant circle*. Under a stroboscopic mapping, solutions return to the same invariant circle from which they began, rotated by the torus number ρ. Since solutions are dense on the quasi-periodic torus[17], the invariant circle can be considered the locus of initial states — that when propagated for at least one period — cover the torus. Thus, one can gain an accurate representation of the solutions available to each QPT by solving for a discrete set of N points on its invariant circle. These points are equally spaced and their latitudinal phase θ_2 (subscripts dropped from this point forward) obtained as the set of $\theta_i = 2\pi i/N$ for $i = 0, 1, 2, ..., N-1$. It should be noted, the invariant circle is not unique for a given QPT. In fact there are an infinite many (one at every cross section of the torus). However, due to the invariance property, only one needs to

[16]For example, how to compute the stable and unstable invariant manifolds associated with a QPT.

[17]In actuality, it is a single solution that makes up the torus when propagated for an infinite time horizon — constantly winding and wrapping about the surface, never intersecting itself. However, for practical, finite time horizon problems, it appears as a continuum of solutions dependent on the starting latitudinal phase on the torus. For further explanation refer to section 4.1.1

be solved to gain the full information of the torus. The constraints added to the differential corrector of the GMOS solver ensure that the computed solutions lie on the the same invariant circle.

The GMOS algorithm pairs a differential corrector with the invariance condition of the torus to compute N solutions on its invariant circle. As with the periodic solutions of section 3.1, the algorithm needs to be seeded with a good initial guess for each of these states. Recall, the ultimate objective is to find the set of tori that surround a given halo orbit. It is known that the eigenvector with corresponding eigenvalues on the unit circle — obtained from the Poincaré map of the periodic orbit — points in the direction of the center subspace containing the quasi-periodic solutions [8]. Thus one can obtain the set of initial guesses by perturbing the halo orbit initial conditions by some ϵ along this eigenvector direction, and rotating the corresponding vector by the set of θ_i. These individual perturbations are defined by equation 3.14,

$$\mathbf{u}_{0,i} = \epsilon(\cos(\theta_i)\mathrm{Re}[\nu] - \sin(\theta_i)\mathrm{Im}[\nu]) \tag{3.14}$$

where $\mathbf{u}_{0,i}$ is the initial guess with respect to the halo orbit condition[18], $\mathrm{Re}[\nu]$ is the real component of the corresponding eigenvector, and $\mathrm{Im}[\nu]$ its imaginary component. The torus number also needs to be estimated, which can be accomplished as follows:

$$\rho_0 = \tan^{-1}\frac{\mathrm{Im}[\lambda]}{\mathrm{Re}[\lambda]} \tag{3.15}$$

where $\mathrm{Re}[\lambda]$ and $\mathrm{Im}[\lambda]$ are the real and imaginary components of the eigenvalue corresponding to ν respectively. This initialization procedure is only performed for the first torus solution found per halo orbit. Continuation methods are to then used

[18]such that $\mathbf{x}_{o,i} = \mathbf{x}_H + \mathbf{u}_{0,i}$ is the state in the rotating frame corresponding to the initial guess on the invariant circle.

to converge to all subsequent solutions.

Exploiting the Invariance Condition

The GMOS algorithm computes a discrete set of solutions on the torus by incrementally correcting the vector of free variables $\mathbf{Z}$, until a targeted set of terminal states — defined relative to $\mathbf{Z}$ by the error vector $\mathbf{F}(\mathbf{Z})$ — are achieved within some tolerance. The vector of free variables uniquely defines the torus and is constructed as follows:

$$\mathbf{Z} = \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_{N-1} \\ \mathbf{u}_N \\ T \\ \rho \end{bmatrix} \tag{3.16}$$

where $\mathbf{u}_i$ is the initial state of the i^{th} solution[19] on the discretized invariant circle (N total), T is the period (i.e., stroboscopic mapping time) of the torus, and ρ is the torus rotation number. When satisfied, the error vector $\mathbf{F}(\mathbf{Z})$ — built from a set of carefully selected conditions — constrains the set of N solutions to a uniquely identified invariant circle, and thus, quasi-periodic torus. These conditions are outlined in the remainder of this section.

The GMOS algorithm is anchored by the the invariance condition of the quasi-periodic torus. This characteristic of the flow can be expressed as follows:

$$\mathbf{u}_i(\theta_i + \rho) = \phi(T; \mathbf{u}_i(\theta_i)) \tag{3.17}$$

[19]That is, converged six-dimensional state vector defined relative to the halo orbit initial conditions.

where $\phi(T; \mathbf{u}_i(\theta_i))$ is the stroboscopic map (i.e., mapping time of one period, T) applied to the initial state $\mathbf{u}_i(\theta_i)$ defined by its phase on the invariant circle θ_i. The combined longitudinal and latitudinal flow along the torus causes the mapped solution to return to the invariant circle, but rotated in phase from the starting location by the torus number ρ. If a rotation operator can be defined that effectively removes the torus rotation from the imaged state, then the invariance condition could be expressed as a mathematical constraint:

$$R_{-\rho}[\phi(T; \mathbf{u}_i(\theta_i))] - \mathbf{u}_i(\theta_i) = 0 \tag{3.18}$$

where $R_{-\rho}$ is this operator. Equation 3.18 is a function of $\mathbf{u}_i$, ρ, and T, that when satisfied, ensures that $\mathbf{u}_i$ lies on the invariant circle of the torus.

Note that the states $\mathbf{u}_i$ are equally distributed in θ along the invariant circle. This outcome of the stroboscopic map formulation allows for the discretized invariant circle to be naturally represented by a series of Fourier coefficients — which can be obtained by multiplying the states by the Discrete Fourier Transform matrix show below.

$$D = \begin{bmatrix} e^{jk_1\theta_1} & e^{jk_1\theta_2} & \dots & e^{jk_1\theta_N} \\ e^{jk_2\theta_1} & e^{jk_2\theta_2} & \dots & e^{jk_2\theta_N} \\ \vdots & \vdots & \ddots & \vdots \\ e^{jk_N\theta_1} & e^{jk_N\theta_2} & \dots & e^{jk_N\theta_N} \end{bmatrix} \tag{3.19}$$

where $j = \sqrt{-1}$, $\mathbf{k} = [-\frac{1}{2}(N-1)..., -1, 0, 1, ..., \frac{1}{2}(N-1)]$ is the vector of indices (note that this assumes that N is odd), and $\theta_i = 2\pi i/N$ for $i = 0, 1, 2, ..., N-1$. It follows that a matrix realization of $R_{-\rho}$ can be constructed that rotates the resulting coefficients by the desired angle and converts them back to Euclidean space — with torus angle removed — by the inverse DFT. This is expressed as follows:

$$R_{-\rho} = D^{-1}Q(-\rho)D \tag{3.20}$$

where $Q(-\rho)$ is the diagonal matrix that rotates the Fourier coefficients of the state by $-\rho$ (equation 3.21). The inverse DFT matrix is defined in equation 3.22.

$$Q(-\rho) = \begin{bmatrix} e^{-jk_1\rho} & 0 & \dots & 0 \\ 0 & e^{-jk_2\rho} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & e^{-jk_N\rho} \end{bmatrix} \tag{3.21}$$

$$D^{-1} = \frac{1}{N}\begin{bmatrix} e^{jk_1\theta_1} & e^{jk_2\theta_1} & \dots & e^{jk_N\theta_1} \\ e^{jk_1\theta_2} & e^{jk_2\theta_2} & \dots & e^{jk_N\theta_2} \\ \vdots & \vdots & \ddots & \vdots \\ e^{jk_1\theta_N} & e^{jk_2\theta_N} & \dots & e^{jk_N\theta_N} \end{bmatrix} \tag{3.22}$$

The invariance condition, applied to all states on the invariant circle, is the first constraint of the error vector $\mathbf{F}(\mathbf{Z})$, and is defined in vector-matrix form below

$$c_1(\mathbf{X}_1, \mathbf{X}_0) = [R_{-\rho} \otimes I_6]\mathbf{X}_1 - \mathbf{X}_0 \tag{3.23}$$

where $R_{-\rho} \otimes I_6$ is the Kronecker product between rotation operator and the 6×6 identity matrix, $\mathbf{X}_1$ is the $6N \times 1$ vector of the mapped initial states, and $\mathbf{X}_0$ is the $6N \times 1$ vector of initial states (shown below).

$$\mathbf{X}_1 = \begin{bmatrix} \phi(T; \mathbf{u}_1) \\ \phi(T; \mathbf{u}_2) \\ \vdots \\ \phi(T; \mathbf{u}_N) \end{bmatrix}, \qquad \mathbf{X}_0 = \begin{bmatrix} \mathbf{u}_1 \\ \mathbf{u}_2 \\ \vdots \\ \mathbf{u}_N \end{bmatrix}$$

Additional Constraints

The invariance constraint on its own cannot uniquely define a quasi-periodic torus[20]. A second condition (e.g., on energy, ρ, etc.) — added to the error vector $\mathbf{F}(\mathbf{Z})$ — is required to achieve this. The following constraint was selected to generate the constant period family of QPT solutions:

$$c_2(T, T_{fixed}) = T - T_{fixed} \tag{3.24}$$

where T_{fixed} is the desired period of the torus. Recall, the invariant circle also needs to be uniquely identified on the torus[21]. This is accomplished by adding two phase constraints [81] to the error vector that exploit a previously converged member of the database, $\tilde{\mathbf{Z}}$[22]. The first is as follows:

$$c_3\big(\mathbf{X}_0, \tilde{\mathbf{X}}_0, \frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_1}\big) = \frac{1}{N} < \mathbf{X}_0 - \tilde{\mathbf{X}}_0, \frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_1} > \tag{3.25}$$

where $\tilde{\mathbf{X}}_\mathbf{0}$ is the vector of converged states on the invariant circle of the previously solve QPT and $< \cdot, \cdot >$ represents the vector inner product operation. The partial derivative of $\tilde{\mathbf{X}}_\mathbf{0}$ with respect to longitudinal angle is provided below:

$$\frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_1} = \frac{2\pi}{T}\left(f(\tilde{\mathbf{X}}_0) - \frac{\rho}{T}\frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_2}\right) \tag{3.26}$$

where $f(\tilde{\mathbf{X}}_0) = \dot{\tilde{X}}_0$ is the vector flow of the previously converged states. The partial derivative of $\tilde{\mathbf{X}}_\mathbf{0}$ with respect to latitudinal angle is provided below:

[20]Note that equation 3.23 provides $6N$ equations, and $\mathbf{Z}$ is of dimension $6N + 2$.

[21]It can both shift in longitudinal space and rotate in latitude while satisfying the first two constraints. The importance of this is especially true during the continuation procedure, where the algorithm can converge to a different invariant circle of the same torus if not properly constrained [75]

[22]Note, to compute the first element of the database, one can perturb two sets of initial guesses along the halo orbit's center subspace eigenvector. One that simulates a previously converged solution and the other (perturbed from the previous), which will serve as the true first initial guess that will be passed to the GMOS algorithm. A special thanks to Dr. Nicola Baresi for providing this tip.

$$\frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_2} = [D^{-1}WD \otimes I_6]\tilde{\mathbf{X}}_0 \tag{3.27}$$

where,

$$W = \begin{bmatrix} jk_1 & 0 & \dots & 0 \\ 0 & jk_2 & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & jk_N \end{bmatrix} \tag{3.28}$$

The second phase constraint can be computed readily thereafter, as shown in equation 3.29.

$$c_4\big(\mathbf{X}_0, \frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_2}\big) = \frac{1}{N} < \mathbf{X}_0, \frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_2} > \tag{3.29}$$

Psuedo-arclength continuation was employed to help compute solutions across the QPT family. Its condition is added to the error vector as a fifth constraint as follows:

$$c_5(\mathbf{X}_0, \tilde{\mathbf{X}}_0, T, \tilde{T}, T', \rho, \tilde{\rho}, \rho', \delta s) = \frac{1}{N} < \mathbf{X}_0 - \tilde{\mathbf{X}}_0, \mathbf{X}' > +(T - \tilde{T})T' + (\rho - \tilde{\rho})\rho' - \delta s \tag{3.30}$$

where δs is the step length and $\mathbf{Z}'$ is the family tangent, computed as follows:

$$\mathbf{Z}' = \begin{bmatrix} \mathbf{X}' \\ T' \\ \rho' \end{bmatrix} = \frac{1}{\alpha}\left(\begin{bmatrix} \mathbf{X}_0 \\ T \\ \rho \end{bmatrix} - \begin{bmatrix} \tilde{\mathbf{X}}_0 \\ \tilde{T} \\ \tilde{\rho} \end{bmatrix} \right) \tag{3.31}$$

where[23],

[23]Again, a special thanks to Dr. Nicola Baresi who provided insight into this normalizing constant.

$$\alpha = \sqrt{\frac{1}{N} < \mathbf{X}', \mathbf{X}' > + (T')^2 + (\rho')^2}$$

Putting it Together

The error vector, consisting of the invariance condition (equation 3.23), period constraint (3.24), phase conditions (equations 3.25 and 3.29), and the pseudo-arclength contuation condition 3.30, is constructed as shown in equation 3.32.

$$\mathbf{F}(\mathbf{Z}) = \begin{bmatrix} c_1(\mathbf{X}_1, \mathbf{X}_0) \\ c_2(T, T_{fixed}) \\ c_3\big(\mathbf{X}_0, \tilde{\mathbf{X}}_0, \frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_1}\big) \\ c_4\big(\mathbf{X}_0, \frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_2}\big) \\ c_5(\mathbf{X}_0, \tilde{\mathbf{X}}_0, T, \tilde{T}, T', \rho, \tilde{\rho}, \rho', \delta s) \end{bmatrix} \tag{3.32}$$

As with all differential corrector methodologies, the Jacobian of the error vector $\mathbf{F}(\mathbf{Z})$ with respect the set free variables $\mathbf{Z}$ is required (provided in equation 3.33).

$$\mathcal{D}F = \begin{bmatrix} \frac{\partial c_1}{\partial \mathbf{X}_0} & \frac{\partial c_1}{\partial T} & \frac{\partial c_1}{\partial \rho} \\ \big(\frac{\partial c_3}{\partial \mathbf{X}_0}\big)^T & 0 & 0 \\ \big(\frac{\partial c_4}{\partial \mathbf{X}_0}\big)^T & 0 & 0 \\ \mathbf{0} & 1 & 0 \\ \big(\frac{\partial c_5}{\partial \mathbf{X}_0}\big)^T & \frac{\partial c_5}{\partial T} & \frac{\partial c_5}{\partial \rho} \end{bmatrix} \tag{3.33}$$

The nontrivial partial derivatives of the Jacobian matrix are provided below. Starting with the invariance constraint,

$$\frac{\partial c_1}{\partial \mathbf{X}_0} = [R_{-\rho} \otimes I_6]\mathbf{\Phi} - I_{6N} \tag{3.34}$$

where $\mathbf{\Phi}$ is the block diagonal matrix of state transition matrices (computed from $t_0 = 0$ to $t_f = T$) for each state on the invariant circle and I_{6N} is the $6N \times 6N$

identity matrix. The partial derivatives with respect to the period and torus angle are computed as follows:

$$\frac{\partial c_1}{\partial T} = [R_{-\rho} \otimes I_6]\dot{\mathbf{X}}_1 \tag{3.35}$$

$$\frac{\partial c_1}{\partial \rho} = -[(D^{-1}WD) \otimes I_6][R_{-\rho} \otimes I_6]\mathbf{X}_1 \tag{3.36}$$

where $\dot{\mathbf{X}}_1 = f(\mathbf{X}_1)$ is the vector flow of the mapped initial states. The partial derivatives of the phase constraints and pseudo-arclength continuation condition readily follow from the formulation, as shown in equations 3.37-3.39.

$$\frac{\partial c_3}{\partial \mathbf{X}_0} = \frac{1}{N}\frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_1} \tag{3.37}$$

$$\frac{\partial c_4}{\partial \mathbf{X}_0} = \frac{1}{N}\frac{\partial \tilde{\mathbf{X}}_0}{\partial \theta_2} \tag{3.38}$$

$$\frac{\partial c_5}{\partial \mathbf{X}_0} = \frac{1}{N}\mathbf{X}' \qquad \frac{\partial c_5}{\partial T} = T' \qquad \frac{\partial c_5}{\partial \rho} = \rho' \tag{3.39}$$

Finally, the correction step can be applied to the vector of free variables $\mathbf{Z}$. The procedure is iterated until the norm of the error vector falls below a user defined tolerance.

$$\mathbf{Z}_c = \mathbf{Z} - [(\mathcal{D}F)^T(\mathcal{D}F)]^{-1}(\mathcal{D}F)^T\mathbf{F} \tag{3.40}$$

Once converged, the initial guess of the subsequent family member is provided as follows: $\mathbf{Z}_0 = \mathbf{Z}_{0,c} + \delta s \mathbf{Z}'$.

3.2.2 Software Implementation

The QPT database software infrastructure was developed in the Julia language[24]. Approximately 2,000 quasi-periodic tori of increasing size[25] were computed for 1,200 different halo orbits of the northern SEL2 family. For each of these tori, their invariant circle was discretized by $N = 21$ states. Thus, approximately fifty million total solutions were computed and stored, making the QPT database three orders of magnitude greater in size than the periodic orbit database.

The computational cost of running the GMOS algorithm over the total number of solutions far exceeded the demands of the periodic orbit database. However, in the lexicon of high-performance computing, this is an embarrassingly parallel problem. The family of tori surrounding each halo orbit can be solved independently from each other[26], opening the door for the problem to be approached as a series of parallel batch jobs. Through the MIT SuperCloud and Lincoln Laboratory Supercomputing Center [89], the work was split across sixty-four Intel Xeon Gold 6248 processors. This resource reduced the time to compute the total number of solutions in the database from the order of years[27], to less than a month.

3.2.3 Database Interaction and Visualization

Each QPT solution stored in the database contains $N = 21$ converged initial conditions on its invariant circle as well as its orbital period T, torus number ρ, and average separation distance from its parent halo orbit[28]. Specific tori are found readily within the database by first calling the search functionality of the periodic

[24]For further discussion on the selection of Julia as the development language, refer to section 3.1.4.

[25]That is, of increasing radius of the invariant circle.

[26]However, tori within a family cannot be solved independently since the method of continuation depends on previously computed members.

[27]In serial on a laptop computer equipped with a 2.8 GHz Quad-Core Intel Core i7 processor and 16GB of memory.

[28]Computed for 3,000 points along the natural dynamical evolution of the trajectories.

orbit database to find a halo orbit of interest. This halo orbit is then passed to a "search QPT Database" function along with a search criteria (such as ρ, or average separation distance from halo orbit) and search value. The closest QPT solution in the database to the desired value is then returned[29]. Various mechanisms for plotting combinations of the halo orbit, invariant circle, and QPT solutions were developed, an example with all of these overlaid is shown in Figure 3-7.

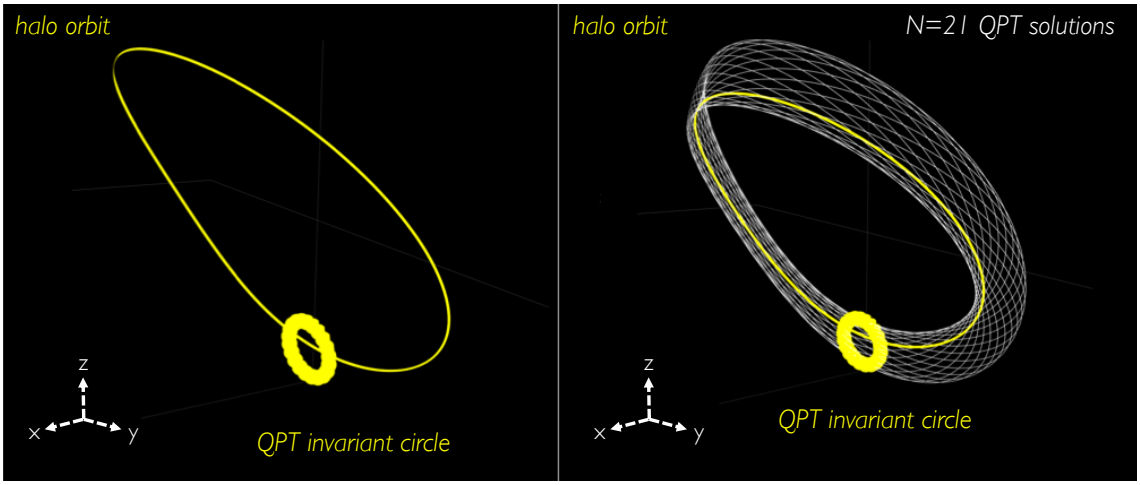

Figure 3-7: Halo orbit with invariant circle overlaid (left) and QPT solutions (right)

[29]Note, a search by bisection method for finding solutions within the gaps of the database as left for future work.

Chapter 4

Trajectory Design Methodology

The Trajectory Design Methodology (TDM) presented in this chapter satisfies the thesis objective introduced in section 1.3. It is rooted in Dynamical Systems Theory in the spirit that maximizing the exploitation of natural solutions correlates to lower fuel costs. This is accomplished by leveraging the periodic and quasi-periodic torus database of Chapter 3, which provides the design space of natural, bounded motion for the observatory and external occulter spacecraft respectively. Additionally, transitioning between elements of the database, preferentially in a dynamically informed manner, is necessary for successful operation. This process is described in section 4.3. It should be noted that ultimately, the TDM is a composition of design decisions, and the one presented is just a single manifestation of these possibilities. Thus, the objective of this chapter is demonstrate a foundation — built from the onset in DST — that is both adaptable to competing mission objectives and a platform from which future work can develop.

The Trajectory Design Methodology divides the spacecraft operating life cycle into a series of more manageable time windows (TW), solving for a sequence of trajectories constrained to the quasi-periodic torus from which the external occulter can transition. The occulter is fixed to the toroidal surface at the formation alignment condition for each target star, only departing to later re-enter during transfers.

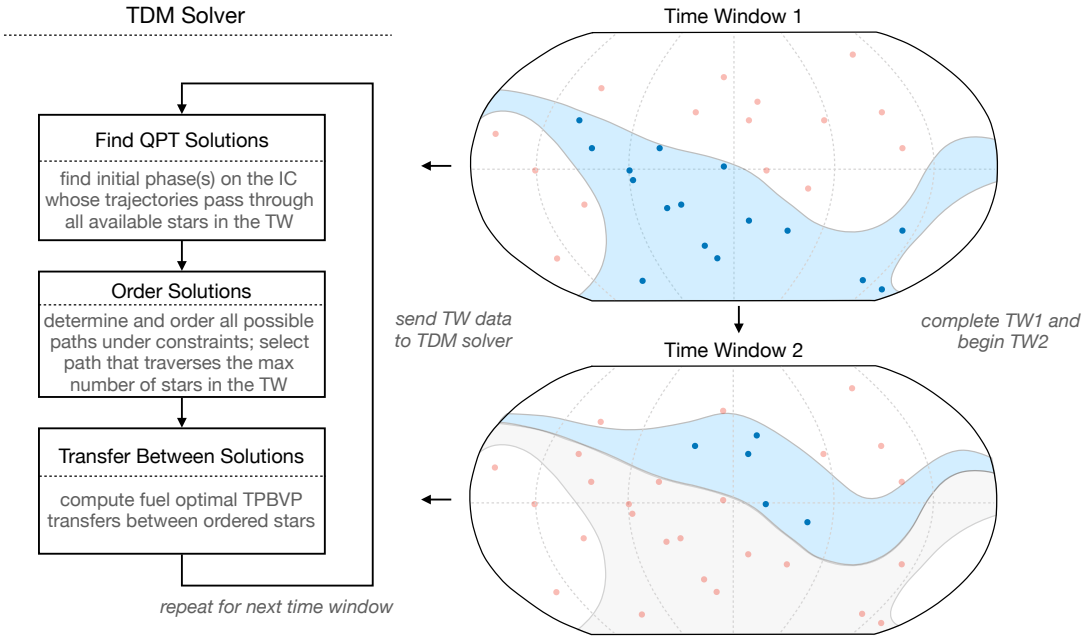

Figure 4-1: The Trajectory Design Methodology overview

The observatory remains in its operating halo orbit, undisturbed[1], throughout the mission. Only a subset of stars are reachable by the formation line-of-sight within a given time window. For each of these time windows, the TDM process can be described in three steps: 1) Find the set of trajectories on the torus that pass through all stars reachable in the time window, 2) Determine and order all possible paths under a given set of constraints and extract the path that traverses the maximal number of stars within the time window, and 3) Transit between the optimally ordered solutions in a dynamically informed, fuel-optimal way. This process is then repeated for the subsequent time window. A graphical representation of this process is provided in Figure 4-1.

In this chapter, sections 4.1, 4.2, and 4.3 describe the three steps of the TDM

[1]In reality, stationkeeping is required to maintain the orbit. However, this is neglected in the analysis

respectively. Chapter 5 demonstrates the Trajectory Design Methodology through its application to a sample mission.

4.1 Finding Solutions on the Torus

The objective of the first step of the Trajectory Design Methodology is to find combinations of naturally existing, ballistic trajectories on the surface of quasi-periodic tori (QPT) and periodic solutions (i.e. halo orbits) that pass the formation line-of-sight (LOS) through the target stars of a design reference mission[2]. Even by restricting the motion of the spacecraft to these subspaces of the Sun-Earth L2 region's center manifold, limitless combinations of QPT and halo orbits are still available to the mission designer. This can be observed in the database developed in Chapter 3, which is a discretization of the continuous families of halo orbits and their surrounding quasi-periodic tori[3]. Thus, careful consideration needs to be made to reduce the search space to a manageable subset.

Two options immediately present themselves. The first is to consider utilizing multiple halo orbits and/or QPT combinations for a single mission, selected from a subset of the database bounded by a range on some parameter such as amplitude or energy. The second is to utilize a single halo orbit and QPT for the duration of the mission. This option has several advantages. First, restricting the analysis to a single halo orbit frees the observatory from having to transfer between solutions, thus simplifying operations and reducing fuel consumption for the spacecraft. This also provides adaptability in the likely case that the operating halo orbit is determined by design constraints outside of formation flight. Second, quasi-periodic solutions are dense on the torus (refer to section 2.3). Since QPT nominally surround their

[2]Recall from the CONOPS of section 2.4, the observatory operates in the halo orbit while the occulter is bounded to the surface of the QPT.

[3]Recall from section 3.2.1 that even QPT come in many different varieties including constant energy, constant frequency ratio, and constant mapping time families. Refer to McCarthy [75] for how to solve for all three.

parent halo orbit, any point on the celestial sphere is in theory reachable by a single halo/QPT combination, providing ample design space for mission planning while eliminating the complexity of handling multiple surfaces of solutions. Finally, the principles developed for this simplified case are general frameworks and can be scaled up to more complex solutions in the future. For these reasons, the later option was selected for this thesis[4]. Section 5.1 details the criteria made in selecting the QPT for the analysis in this thesis.

It should be noted that these simplifying assumptions restrict certain aspects the analysis. For instance, though any latitude and longitude on the celestial sphere can theoretically be reached by a single torus, the precise separation distance between the observatory and external occulter at any point on its surface is determined by the structure of the specific QPT (which if chosen wisely — again see section 5.1 — can get the spacecraft within a known bound) and dynamical flow. Direct exoplanet imaging will require a specific and precise separation distance. It seems promising that by expanding the design space to include multiple tori of varying amplitudes, periods, energies, etc, natural solutions that meet these additional criteria can be included. However, these considerations are beyond the scope of this thesis, where only retargeting between stars is considered. For more discussion on the scope of future work that can stem from this body of research, refer to section 6.2.

4.1.1 Parameterizing the Search Space

By fixing the operating halo orbit of the observatory, the problem reduces to finding the set of solutions on the toroidal surface that pass the formation line-of-sight vector through each of the target stars. Accomplishing this task requires an understanding of how these tracings manifest on the celestial sphere so that they can be

[4]It is recognized that global fuel-optimal solutions are likely not found by constraining the design space to a single halo orbit and QPT. The aim of this thesis is to provide a new path for future work that will continue to maximize exploitation of the naturally available, dense trajectory design space. For further discussion see sections 5.4 and 6.2

differentiated, evaluated, and subsequently selected or discarded. This process is called parameterizing the search space. Figure 4-2 (top) illustrates the importance of this step. Graphed are the formation LOS tracings initialized at the twenty-one invariant circle solutions provided by the database for a single QPT. In this figure, they are propagated for a full period[5]. A first glance shows promise of good coverage of the celestial sphere as expected, but it is impossible by visual inspection to distinguish the curves from each other, and thus extract any meaningful characteristics. However, upon restricting the propagation time to a smaller window, such

[5]Recall, the database QPT share the same period as their parent halo orbit by design and are discretized by twenty-one points along the invariant circle.

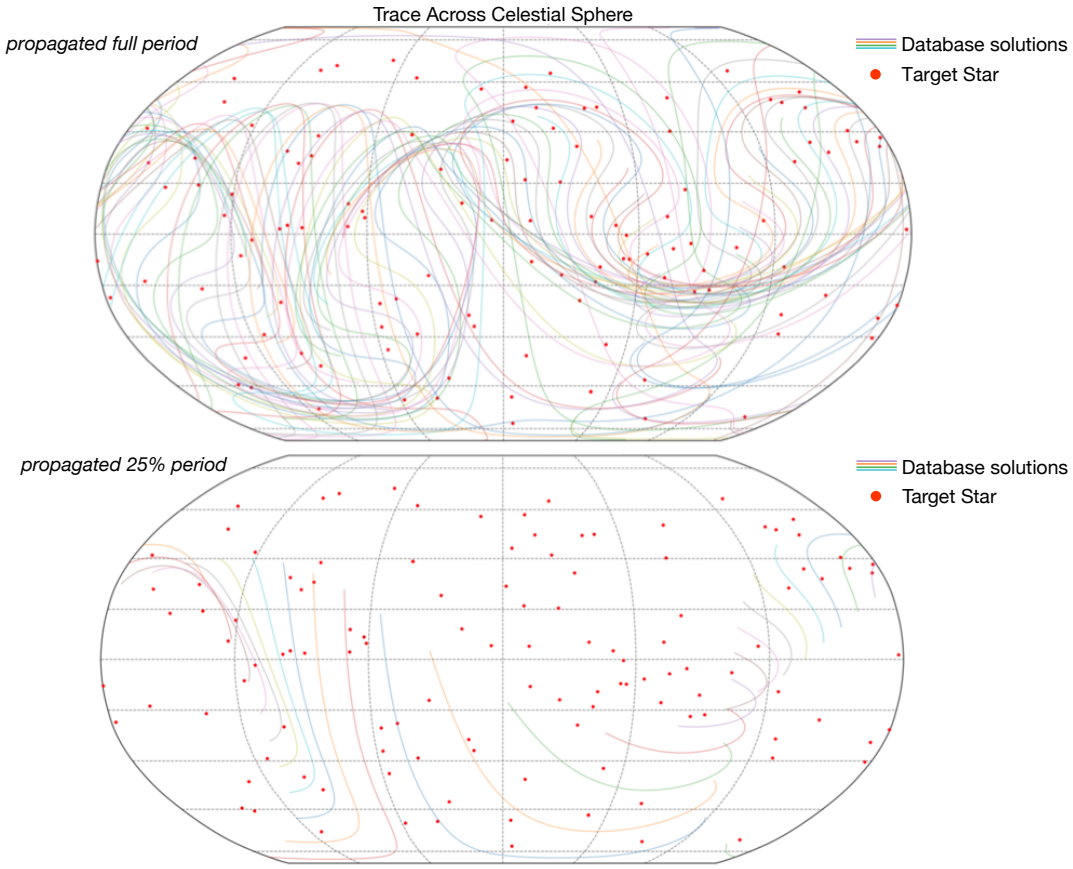

Figure 4-2: Database solutions for a single torus propagated for a full period (top) and single time window — quarter period — (bottom).

as a quarter of the period, individual features become much more distinguishable. This case is shown at the bottom of Figure 4-2.

Most notable are their starting location and direction of flow[6]. All solutions start from the same invariant circle — the cross-section of the torus — and are initially equally separated in phase angle. The projected formation LOS from all points on the circle is marked by the bold gold line in the graph[7]. From this figure, it is clear that the individual solutions are distinguishable by their starting location on the invariant circle. Thus, it is the natural phase space representation of the surface of solution — the torus — that provides the parameterization we seek. These are the latitudinal and longitudinal angles transformable back into to Euclidean space via the diffeomorphism $u(\theta_1, \theta_2)$ (refer to section 2.3).

Expanding further, recall that the torus is made dense by a single quasi-periodic solution that is continuously-winding and never-intersecting. Thus, the invariant circle can be considered a locus of solution points. In the theoretical case of infinite time, a solution starting from any initial phase angle would eventually pass through any and every point on the torus, completely covering the manifold, yielding $\theta_{2,0}$ — the starting latitudinal angle (phase) on the initial invariant circle[8] — an irrelevant parameter. However, for the applicable finite-time-horizon problem, $\theta_{2,0}$ becomes a crucial dial, directly shifting solutions of interest along the projected invariant circle curve on the celestial sphere. Furthermore, by restricting the subset of tori to families with equal period to the underlying halo orbit, as has been done in this thesis, the longitudinal angle (θ_1) becomes equivalent to (and can be replaced by) time.

To summarize, finding solutions on the torus corresponds to finding the set of

[6]Recall from section 2.5.1 that this also inherently depends on t_0, the initial time of the simulation, since it anchors the rotating frame to the J2000 inertial frame. From this point forward, this dependency is assumed and no longer explicitly stated.

[7]This is better illustrated in Figure 4-4.

[8]From this point forward, initial invariant circle corresponds to the invariant circle solved for in the database by the GMOS algorithm. For the purposes of this thesis, this invariant circle represents the locus of initial conditions by their initial phase angle.

M points — defined by initial latitudinal phase angle $\theta_{2,0,i=1:M} \in [0, 2\pi)$ — on the invariant circle, that when propagated for time $t_i < T_w$, yield alignment of the formation LOS with the M available target stars in the time window, where T_w is the length of the time window. Of course, the external occulter cannot exist on all M solutions simultaneously. Instead the TDM must determine an optimal ordering (and transfers) between a subset of the available solutions that maximizes total number of intercepts and minimizes total fuel consumption. These steps are covered in sections 4.2 and 4.3 respectively.

4.1.2 Refining Database QPT Solutions

Every QPT database entry is discretized by $N = 21$ points around its initial invariant circle. As discussed in Chapter 3, N was chosen to balance computational cost and the resolution of solutions readily available per torus. In consequence, N initially constrains the fineness of the search space parameter $\theta_{2,0}$. Figure 4-2 (bottom) illustrates the resulting LOS tracings — made from the N database solutions of a QPT propagated along a quarter-period time window — projected onto the celestial sphere. As can be seen, the curves are sparse, leaving large angular gaps in the celestial sphere untraversed, with several target stars nominally lying in between solutions. The objective of this step of the TDM is to uniquely determine the solutions, by their initial phase, that pass the LOS through each target star available in a time window. To accomplish this, a finer discretization of the initial invariant circle is needed.

QPT database entries are members of continuous families, where each invariant circle of converged solutions builds upon the previously solved member of smaller size. This is done by expanding the same N points by a small δ in radius from the parent halo orbit and providing these new points to the QPT solver algorithm[9].

[9]This is a simplified explanation, refer to 3.2.1 for a more in-depth discussion on pseudo-arclength continuation.

Algorithm 1 Refine Invariant Circle

1: **procedure** RefineInvariantCircle($\mathbf{X_i}, N, \alpha, \rho$)
2: `push!`($\mathbf{X_{ref}}$, $\mathbf{X_i}$) ▷ insert in dequeue
3: **for** $j = 1 : \alpha - 1$ **do**
4: $\mathbf{X} = (R_{-(2\pi j/N\alpha)} \otimes I_6)\mathbf{X_i}$ ▷ I_6 is the 6×6 identity matrix
5: $\mathbf{X_{jN+i}}$ = `solveQPTGMOSFixed`$\rho(\mathbf{X}, \rho, ...)$
6: `push!`($\mathbf{X_{ref}}$, $\mathbf{X_{jN+i}}$)
7: **end for**
8: **return** $\mathbf{X_{ref}}$
9: **end procedure**

Recall from section 3.2.1 that the GMOS QPT solution method works by enforcing the invariance condition of the torus (reproduced below):

$$R_{-\rho}[\phi(T; \mathbf{u}_i(\theta_i))] - \mathbf{u}_i(\theta_i) = 0 \tag{4.1}$$

where $\phi(\cdot)$ is the stroboscopic map, which is applied to some initial condition $\mathbf{u}_i$ defined at an initial phase on an invariant circle θ_i. Recall, after being propagated for the longitudinal period T, the solution returns to the same invariant circle from which it started, but rotated in phase by ρ. The rotation operator $R_{-\rho}[\cdot]$ defined in equation 4.1 can be employed to increase the density of points on the invariant circle by a scalar multiple of N. This process is outlined in Algorithm 4.1[10].

The objective is to increase the total number of solutions on the invariant circle by a factor α. Since all points on the invariant circle should be equally spaced, their final phase separation should become $\gamma = 2\pi/N\alpha$. This is accomplished through the rotation operator (line 3), which transforms the original N converged invariant circle solutions into their Fourier coefficients and rotates them by γ. Figure 4-3 provides an illustration of this step. Since the original N points are converged to sub-tenth-order accuracy, the new points should nearly lie on the same invariant circle as the original. To ensure that they do, the GMOS solver is run on the new

[10]Note, since the database torus solutions are already converged to a high level of precision, this refinement process may be skipped in favor of a simpler two-dimensional interpolation based approach. However, this methodology was not tested in this work.

set of points, fixing the torus angle ρ to the value saved in the database as opposed to the longitudinal period[11]. This change corresponds to swapping two digits in the error Jacobian. This procedure is repeated $\alpha - 1$ times.

Figure 4-4 displays the results of an invariant circle refined by a factor of $\alpha = 5$ for a total of 105 initial phases. As can be seen, a much denser picture is obtained for possibilities of tracing the formation LOS across the celestial sphere. Furthermore, target stars can much more readily be associated with a nearest neighbor solution, and thus, an initial latitudinal phase angle: $\theta_{2,0}$.

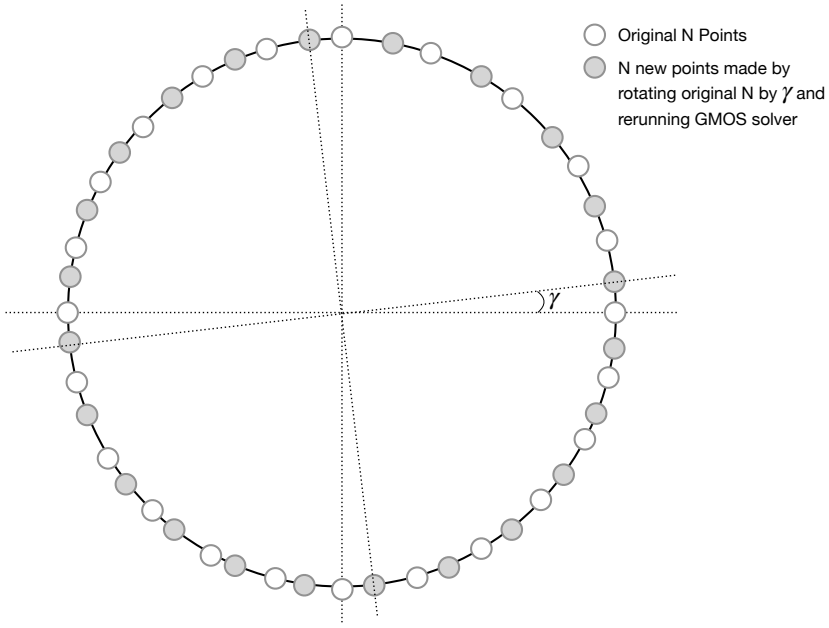

Figure 4-3: Doubling the discretization refinement of the database QPT invariant circle ($\alpha = 2$)

[11]Note: the GMOS phase conditions are enforced by refining the previous and subsequent database entries *without calling GMOS*, and applying the conditions accordingly. Although mathematical guarantees are not provided that the invariant circle does not slip in phase, the precision of the original points combined with the added constraints allow us to considered it negligible for the purposes of this thesis. The alternative of solving all the points together as a batch can quickly become unwieldy depending on the size of α

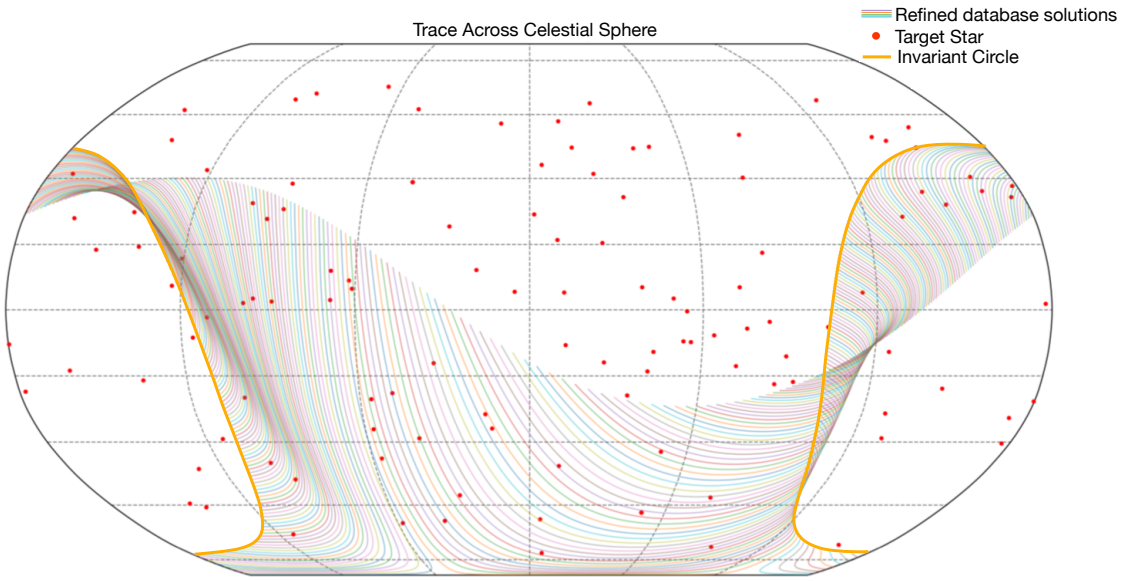

Figure 4-4: Refined invariant circle solutions propagated for a quarter period

4.1.3 First Pass Search: Search by Sorting

A two-step process is presented to find the QPT trajectories that pass the formation line-of-sight through all available target stars in a given time window: a coarse, first pass followed by a second, fine pass. The first pass, "Search by Sorting" is performed by searching through a sorted data structure to find the set of discrete initial phase angles from the refined invariant circle whose corresponding solutions yield closest approach to each target star available in the time window.

The data structure is a matrix of tuples, where each column represents a target star in the DRM and each row, an initial phase angle on the refined invariant circle. To fill the entries of the matrix, each initial condition on the refined torus — denoted by its phase — is numerically propagated for the duration of the TDM time window. The error (equation 4.2) from the solution to each target star is the computed as the Euclidean distance in latitude-longitude space between the star and the point of closest approach by the solution.

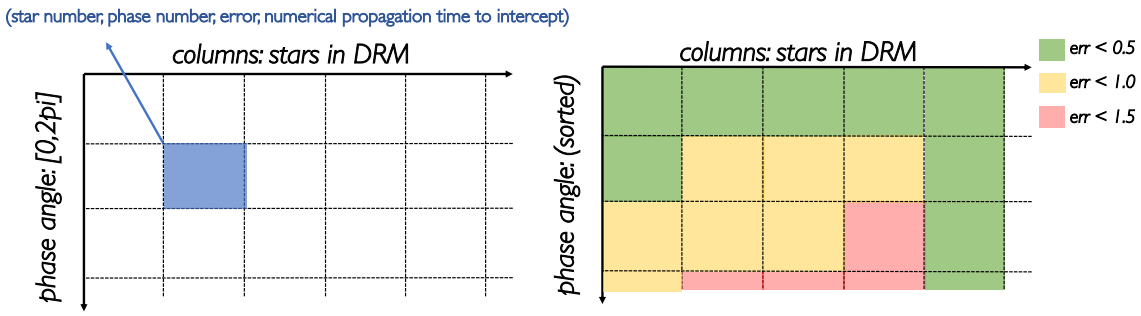

Figure 4-5: Data structure created from database entries: a matrix of tuples (left) that when sorted by column (right) reveal the QPT trajectories that pass the formation LOS closest to each star in the DRM

$$e_{ji} = \min_{k} ||(\theta^{j}_{k,1}, \theta^{j}_{k,2}) - (\theta^{i}_{ts,1}, \theta^{i}_{ts,2})||_2 \tag{4.2}$$

where e_{ji} is the error between the j^{th} trajectory and i^{th} target star, $(\theta^{i}_{ts,1}, \theta^{i}_{ts,2})$ is the latitude and longitude of the i^{th} target star, and $(\theta^{j}_{k,1}, \theta^{j}_{k,2})$ is the latitude and longitude of the j^{th} trajectory evaluated at the k^{th} integration point. Each matrix element is then filled with a four-tuple consisting of the target star's identifying star number, the trajectory's initial phase angle, the error e_{ji}, and the trajectory's time to intercept star i: t_i where $t_i < T_w$. A pictorial representation of the data structure is provided in Figure 4-5 (left).

The rows of each column of the data structure are then sorted by error so that the first row of the newly sorted matrix — see Figure 4-5 (right) — provides the initial phase and error for the closest approaching trajectory to each star in the DRM. A threshold is defined, and any star whose closest approaching trajectory's error lies above this threshold is deemed unavailable in the time window and discarded. Similarly, trajectories that do not rank first in the available stars sorted data structure columns are discarded. The final result is a collection of the initial phases on the refined invariant circle whose corresponding solutions pass the formation LOS tracings closest to each available target star in the time window. These trajectories

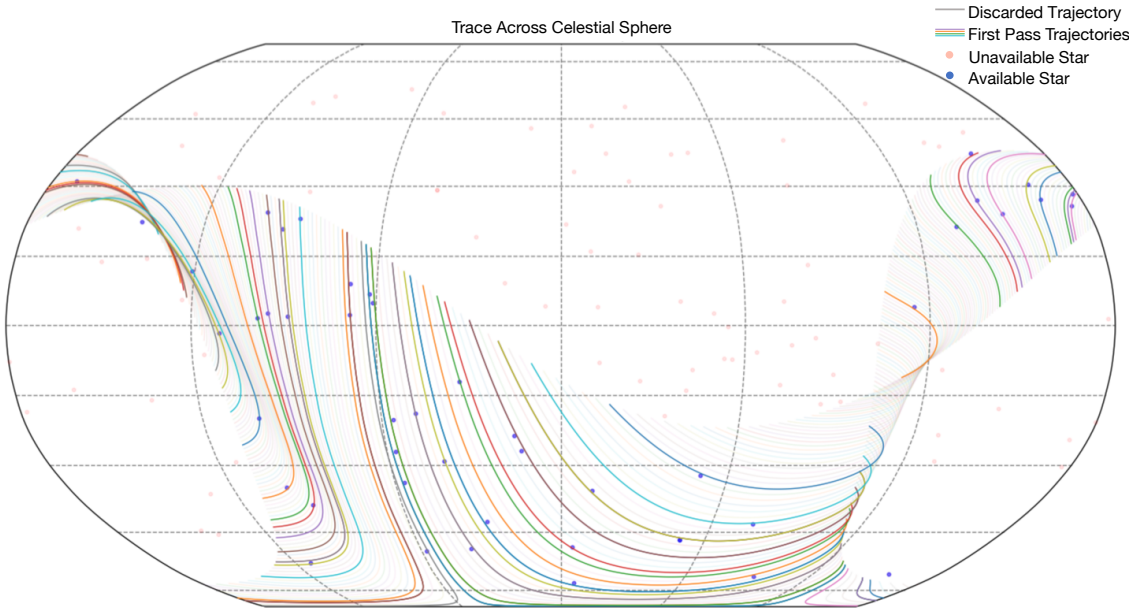

Figure 4-6: First pass search for trajectories for the first time window

(colored) and selected stars (blue) are shown in Figure 4-6 against a backdrop of gray discarded stars and solutions.

4.1.4 Second Pass Search: Search by Shooting

The second phase: "Search by Shooting" is a fine pass search that drives the error down further by converging to solutions in between the discrete initial phases of the refined invariant circle. This is accomplished through a bisection search algorithm, which is sketched as Algorithm 4.2. For simplicity it is assumed here that the phase of interest lies between $\theta_{2,0,i+1}$ and $\theta_{2,0,i}$ and that $\theta_{2,0,i+1} > \theta_{2,0,i}$.

Note that in lines six and seven of the algorithm, the original N invariant circle points from the database are rotated by the newly computed phase angle $\theta_{2,new}$. Its corresponding initial condition in Euclidean space is solved through the GMOS algorithm with fixed torus number ρ (this is similar to the invariant circle refinement procedure of section 4.1.2). The desired state at $\theta_{2,new}$ is saved, the rest are

Algorithm 2 Bisection Search Algorithm Sketch

1: **while** err>tol && $\delta\theta$>threshold **do**
2: $\quad \delta\theta = \frac{1}{2}|\theta_{2,0,i+1} - \theta_{2,0,i}|$
3: $\quad \theta_{2,new} = \theta_{2,0,i} + \delta\theta$
4: $\quad \mathbf{X} = (R_{-\theta_{2,new}} \otimes I_6)\mathbf{X_i}$
5: $\quad \mathbf{X_{jN+i}} = \texttt{solveQPTGMOSFixed}\rho(\mathbf{X}, \rho, ...)$
6: $\quad \mathbf{X_{new}} = \mathbf{X_{jN+i}}[\mathbf{1:6}]$ ▷ save new IC state @ $\theta_{2,new}$, discard the rest
7: $\quad (\theta_{k,1}, \theta_{k,2}) = \texttt{formationFlightSimulation}(\mathbf{X_{new}}, \mathbf{X_{halo}})$
8: $\quad err = \min_k ||(\theta_{k,1}, \theta_{k,2}) - (\theta_{ts,1}, \theta_{ts,2})||_2$
9: $\quad \theta_{2,0,i} = \text{argmin}_{\theta_2}(\text{err}(\theta_{2,0,i+1}),\text{err}(\theta_{2,0,i}))$
10: $\quad \theta_{2,0,i+1} = \theta_{2,new}$
11: **end while**

discarded. In lines eight and nine, the formation flight simulation is run for the newly solved initial condition and its error with respect to the desired target star computed. The phase angle — from the following: $(\theta_{2,0,i+1}, \theta_{2,0,i}, \theta_{2,new})$ — corresponding to maximum error is discarded. The process repeats with the remaining two phase angles until some limit on error or tolerance is reached.

The final QPT solutions found by the second-pass search method are projected onto the celestial sphere and plotted with the target stars they intercept in Figure 4-7. Thus obtained for a given time window are the subset of available targets stars from the DRM[12], the set of trajectories on the torus that naturally pass the formation LOS through these stars[13], and the time to intercept each star starting from each initial condition. This result satisfies the objective of the first step of the Trajectory Design Methodology.

4.2 Ordering Solutions

Section 4.1 demonstrated how to find the set of trajectories on the quasi-periodic torus that naturally lead to alignment between the formation LOS and the set

[12]Whic are identified by their star number.
[13]Which are identified by their initial latitudinal phase on the invariant circle.

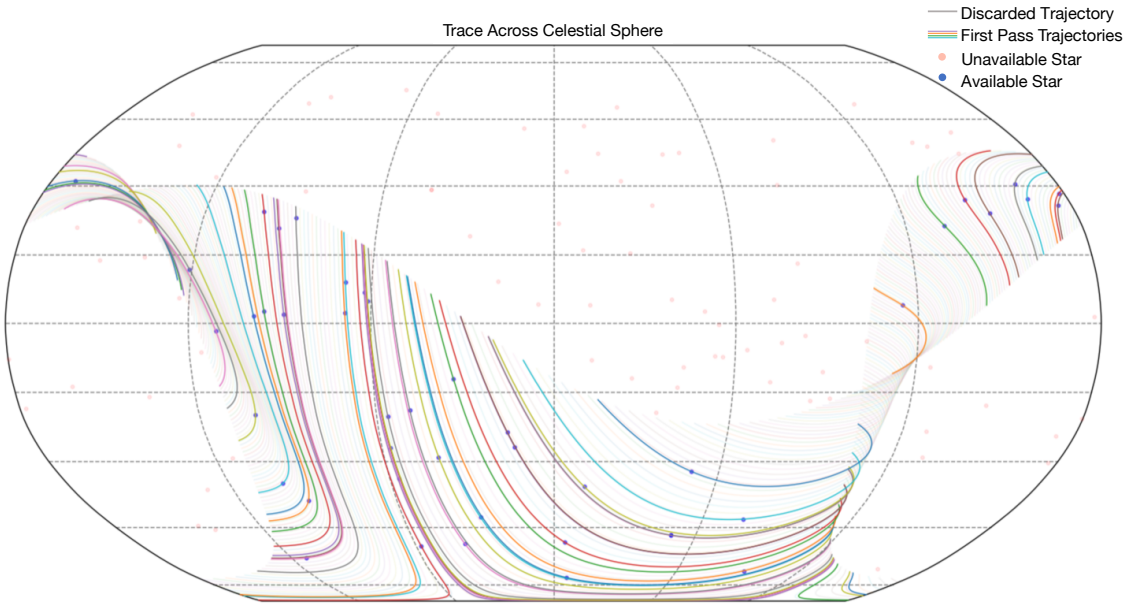

Figure 4-7: Second pass search for trajectories for the first time window

of available target stars within a given time window. In general, each trajectory provides a path to intercept a single star (this can be observed in Figure 4-7). Clearly, the external occulter cannot exist on all paths simultaneously and must transition between the solutions in a manner that maximizes the total number of stars imaged. Since the quasi-periodic torus is an invariant surface, transitions connecting natural trajectories must, in general, be forced solutions. However, due to constraints — such as time and thruster saturation — not all transfers are feasible. Thus, starting from any star, the occulting spacecraft has a potentially limited choice of subsequent stars to which it can next transition[14].

The objective of the second step of the Trajectory Design Methodology[15] is

[14]Note: transition here means transfer the spacecraft via active control to the subsequent star's QPT solution in time for formation/star alignment

[15]As previously discussed, as well as in section 6.2.2, the TDM is a composition of design decisions. The methodology outlined in this chapter reflects the objective of balancing fuel costs and science efficiency, as outlined in the thesis statement, however, it can and should be adapted to future competing mission objectives. This should be a subject of future work.

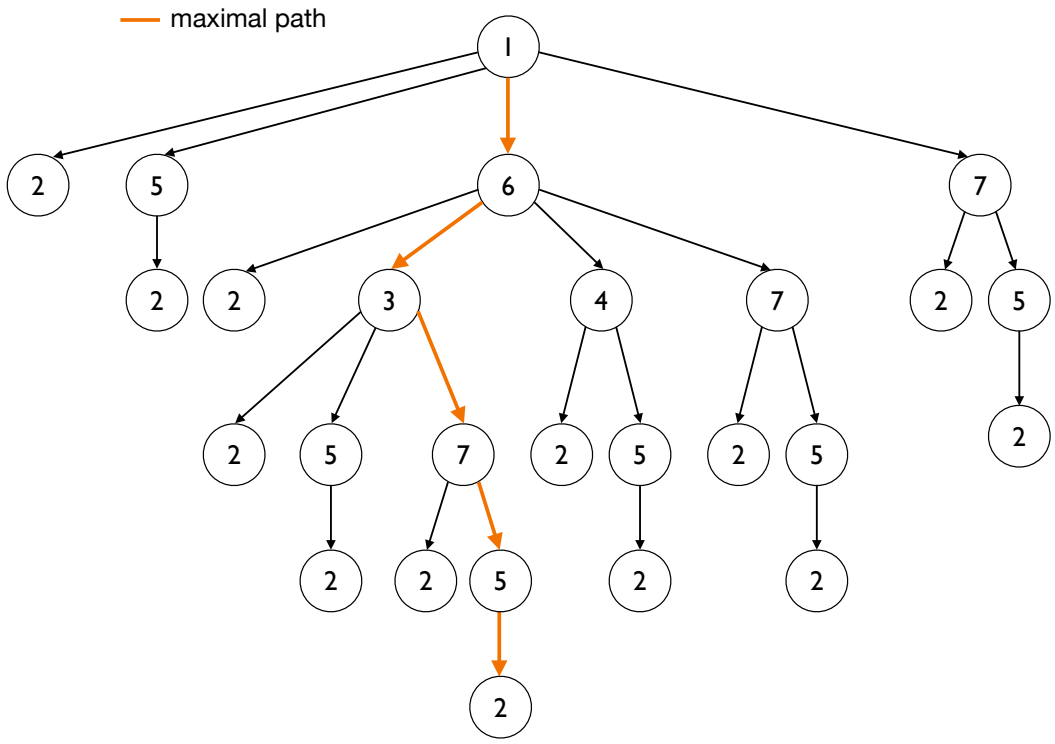

Figure 4-8: Illustration of the tree-based structure of the star ordering and maximal path search problem.

to select and order the targets stars available in the time window into the set of all feasible paths, and extract the path that intercepts the maximum number of stars[16].As illustrated in Figure 4-8, this can be formulated as a search problem over a singly-connected graph[17].

4.2.1 Graph Search Method

Figure 4-8 illustrates the tree-based structure of an example star ordering and maximal path search problem. The graph is directed, with nodes corresponding to available stars. Each node is labeled with its identifying star number, and edges

[16]Note, solution of the boundary value problem that frames the set of transfers between solutions is covered in section 4.3.1.

[17]A singly-connected graph has a tree-based structure with no loops.

indicate feasible transitions between stars and their direction[18]. For instance, since star 1 is a parent of star 5, a transition from star 1 to star 5 is possible given the underlying constraints, however, the reverse is not true. Additionally, since star 2 is a child of every star (and a parent of none), it can be concluded from the graph that it is a terminal star; that is, no other star's alignment condition can be reached from star 2 within the time window. Thus, the longest path from root to leaf on the graph is the ordered path that yields the maximum number of intercepts. This is called the maximal path, and in the example given, is shown in gold. It should be interpreted as the order in which the external occulter should transfer between the QPT solutions corresponding to the stars (e.g., stars 1, 6, 3, 7, 5, and 2 respectively). Section 4.3 will cover how to compute these transfer maneuvers.

Feasible transitions are determined through a set of star transit rules that govern the search process. These can be set by the mission designer to meet a specific criteria but at minimum must include that transitions can only occur forward in time. Recall that the first step of the TDM outputs not only QPT solution for a each available star, but also the time to intercept (τ). Thus, a transition from star i to star j can only occur for time windows less than the period if $\tau_i > \tau_j$. It is recommended that constraints are also made that correlate to thruster saturation levels[19].

4.2.2 Solving the Search Problem

The search problem is solved recursively[20], taking inspiration from the field of Dynamic Programming. The method is outlined below:

1. Determine the moves available (feasible transitions) to each star from the star

[18]Note, the cost of a transfer is not considered in this step, only that it is feasible.

[19]To show how the star tranist rules can be applied in a sample mission, see section 5.1.1.

[20]It is acknowledged that there are many ways to approach a tree-based search. A variety of methods should be explored and compared in future work.

transit rules[21] — this is equivalent to determining the children to each star in the graph (Figure 4-8).

2. Recursively score each star based on the number of remaining moves available. Start with terminal stars — those with no children, and hence no available available moves — these receive a score of zero.

3. Each subsequent star is scored as the maximum score of its children plus one.

4. Progressively work backwards, scoring the stars from those with smallest number of children to those with the largest. This ordering process will minimize the number of cases where a star's child has not been scored before the parent is attempted. If this occurs simply push it further down the stack for a later attempt.

5. The star with the maximum score yields the maximal path. To extract the path, simply start at that star, select its maximum scoring child, followed by that child's maximum scoring child, and so forth until a terminal star is reached.

6. If multiple paths share a maximum score, select the one with minimum latitudinal phase change between the ordered solutions.

4.2.3 Simple Ordering Example

A simple example of a star ordering problem solved via the graph search method is provided in this section. For illustrative purposes, the following simplified star transit rules are applied: 1) transitions can only occur forward in time, and 2) the maximum latitudinal angle that can change between QPT solutions is limited to 0.8

[21] Again, refer to section 5.1.1 for the set of star transit rules applied to a sample mission.

radians: $(\Delta\theta_{2,max} = 0.8)$[22].

Figures 4-9-4-11 illustrate the steps of the graph search method applied to the example problem. Note that though similar, this problem does not produce the same graph show in Figure 4-8. Figure 4-9 demonstrates determining the moves available for the first four stars (all others omitted). Figure 4-10 demonstrates the stars being recursively scored. Finally, Figure 4-11 demonstrates extraction of the maximal path. Note: two stars shared a maximum score, thus the one with minimum latitudinal phase change was selected.

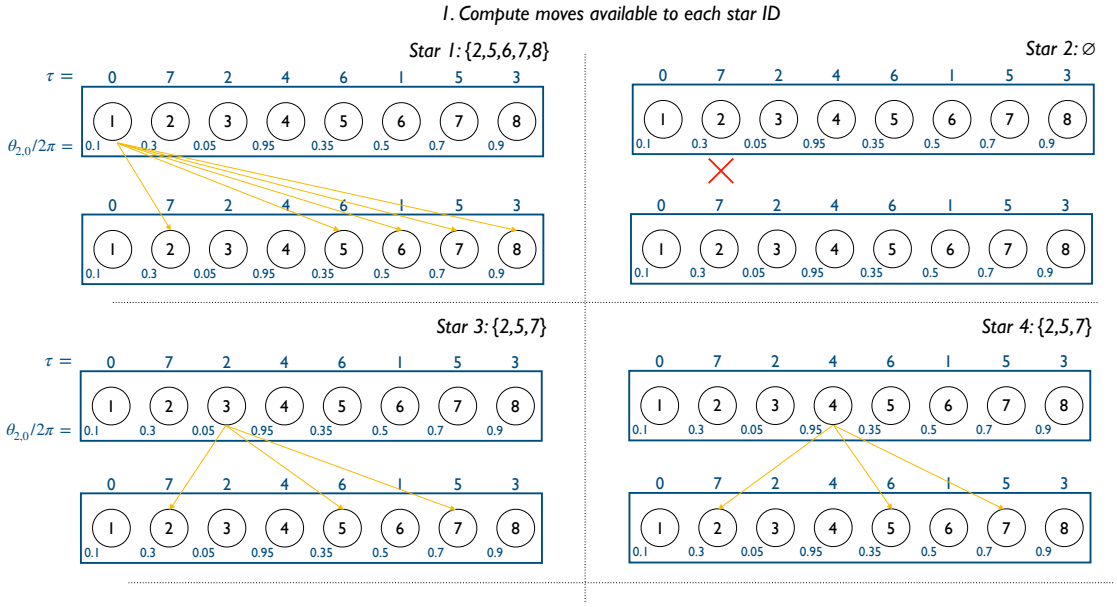

Figure 4-9: Step 1 of the graph search method illustrated.

[22]Note, this value is chosen for illustrative purposes only and does not represent any physical significance. In the sample mission analysis, this method of constraining $\Delta\theta_{2,max}$ is eschewed altogether in favor of a heuristic method that relates transfer conditions to thruster saturation.

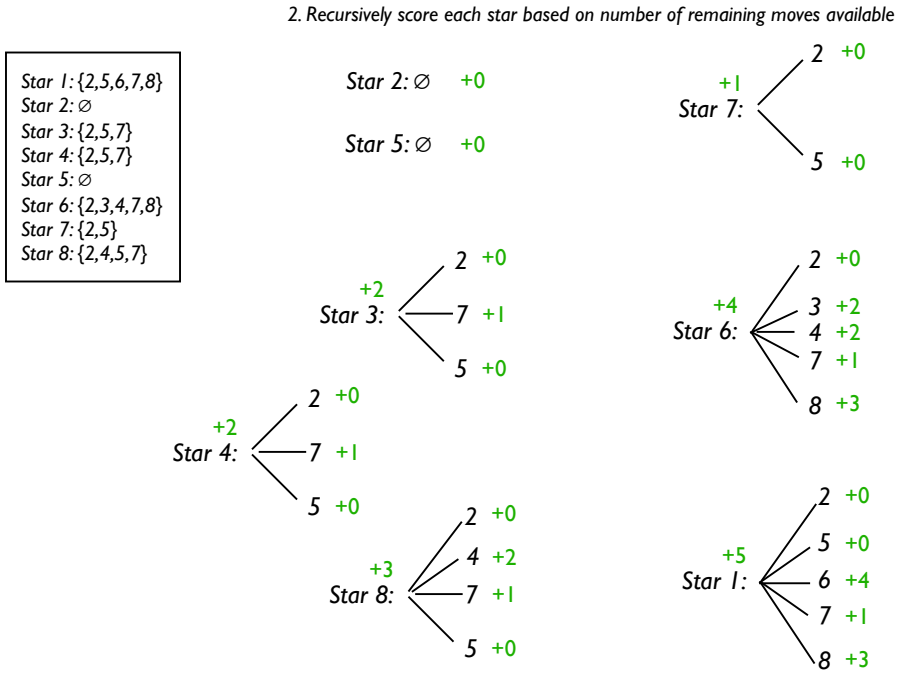

Figure 4-10: Steps 2-4 of the graph search method illustrated.

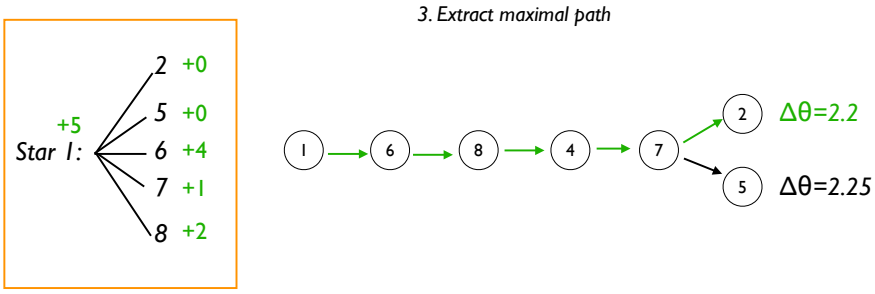

Figure 4-11: steps 5-6 of the graph search method illustrated.

4.3 Transferring Between Solutions

The final step of the Trajectory Design Methodology is to connect the solutions identified on the quasi-periodic torus[23] that naturally yield alignment between the formation and the path of stars sequenced and selected in step two[24]. This is accomplished through a series of transfer maneuvers, that ultimately produce a continuous trajectory through the time window, intercepting all stars in the maximal path. This section demonstrates a method for computing these transfers[25] in a manner that satisfies the objectives of this thesis.

4.3.1 Minimum Fuel-Optimal Control Problem Formulation and Solution

Depending on the boundary conditions of the problem, the design space of feasible transfers between two QPT trajectories can be dense with solutions. Thus, special care needs to be taken to ensure that those computed meet both the constraints of the mission and the overarching thesis objective. Recall from section 1.3, the TDM should maximize the opportunity to exploit natural solutions — in an effort to reduce fuel consumption — while minimizing sacrifices to science efficiency. These are oftentimes competing objectives. However, they can be balanced by framing each transfer as a minimum-fuel optimal control problem. As discussed in this section, solutions to this class of problems are expected to maximally exploit natural phenomena. Combining this approach with the star ordering method of section 4.2 ensures that the maximum number of stars is reached under the thruster saturation limits.

It should again be noted that the TDM is a composition of design choices, and

[23]Recall, this was accomplished in step one of the TDM: refer to section 4.1 for a detailed discussion.

[24]Refer to section 4.2.

[25]Note, this includes the control and state history for each maneuver.

many directions could be taken that still satisfy its guiding philosophy. This is especially true for the transfer problem. The TDM presented in this thesis is just one set of possible decisions. There is still significant room for innovation in this domain, it is believed and encouraged that improvements can still be made through community engagement. See section 6.2 for a discussion on other directions that were explored to solve the transfer problem as well as recommendations for future work.

Transfers as Two-Point Boundary Value Problems

The ordered star list produced by step two of the TDM naturally frames each individual transfer as a two-point boundary value problem with fixed final time. To demonstrate, consider a transfer between two subsequent stars from the list: star i and star $i+1$. Recall, the maximal path describes the set of stars sorted by their time of intercept, τ. Thus, it is known that $\tau_{i+1} > \tau_i$. It is assumed that the formation is aligned with star i at initial time $t_0 = \tau_i$. Since the observatory is assumed to operate ballistically in the halo orbit, it will arrive at the formation alignment condition for star $i+1$ at final time $t_f = \tau_{i+1}$. Thus, the external occulter must do the same — fixing the final time of the transfer maneuver. If $\mathbf{x}_{0,i}$ and $\mathbf{x}_{0,i+1}$ are the initial conditions on the invariant circle leading to formation alignment for stars i and $i+1$ at their respective intercept times[26], then the initial and final state conditions of the occulter are fixed to $\mathbf{x}(t_0) = \phi(\tau_i; \mathbf{x}_{0,i})$ and $\mathbf{x}(t_f) = \phi(\tau_{i+1}; \mathbf{x}_{0,i+1})$ respectively, where $\phi(\tau; \mathbf{x}_0)$ is the solution of the flow at time τ starting from initial condition $\mathbf{x}_0$.

It is important to note that not all two-point boundary value problems with fixed final time have a solution. For the purposes of this thesis, these cases can be correlated to the maximum thrust producible by the onboard propulsion and built in to the star transit rules (refer to section 4.2.1), which evaluate if a transfer is

[26]These are the QPT solutions found in step one of the TDM.

feasible during star ordering. Section 5.1.1 demonstrates how this is done for the sample mission analyzed.

Problem Formulation and Solution

Feasible transfers have the potential to have multiple solutions. The objective of the TDM is to solve each two-point boundary value problem in the most fuel efficient way possible. This is accomplished by solving the constrained minimum-fuel optimal control problem formulated in equations 4.3 - 4.6.

$$\min_{\mathbf{u}} J(\mathbf{u}) = \min_{\mathbf{u}} \int_{t0}^{tf} \left[\sum_{i=1}^{6} u_i(t) \right] dt \tag{4.3}$$

subject to:

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}, \mathbf{u}) \tag{4.4}$$

$$0 \leq u_i \leq u_{max} \tag{4.5}$$

$$t_0 = \tau_i, \quad t_f = \tau_{i+1}, \quad \mathbf{x}(t_0) = \phi(\tau_i; \mathbf{x}_{0,i}), \quad \mathbf{x}(t_f) = \phi(\tau_{i+1}; \mathbf{x}_{0,i+1}) \tag{4.6}$$

where $\mathbf{u}$ is the control input and $\mathbf{x}$ is the state. Equation 4.4, where $f(\mathbf{x}, \mathbf{u})$ represents the vector flow, constrains the problem to the CR3BP phase space. Equation 4.5 enforces thruster saturation limits[27]. Finally, equation 4.6 fixes the two-point boundary value conditions described at the start of this section. Both the optimal control history and resulting trajectory as a function of time are sought as a function of time.

Certain classes of minimum-fuel optimal control problems with linear constraints can be solved analytically via Pontryagin's Minimum Principle. A characteristic control history (in the absence of singular arcs) typically emerges termed "bang-

[27]To avoid adding absolute values to the integrand, the number of thrusters was doubled. That is, instead of having three thrusters for control of motion in the x, y, and z direction. Two thrusters are added per axis, one in the positive direction and one in the negative. This appears in the equations of motion as such: $\dot{x} = f(x) + u_{x,p} - u_{x,n}$. By enforcing that $u_i \geq 0$ and solving, this becomes an equivalent (and better behaved) problem for a min-fuel cost function

off-bang," which refers to its pattern of saturating the thrusters until a natural trajectory connected to the final state in configuration space is reached, coasting along the natural solution, and saturating the thrusters in the terminal phase until the end conditions are met. For a more comprehensive discussion on minimum-fuel optimal control problems and their classes of solutions refer to Kirk [56].

However, the problem defined above is a nonlinear program, and cannot be solved analytically. Instead, a dedicated trajectory optimization software package is required to solve this problem numerically. Though the transfer solutions may not be as simple as a "bang-off-bang" pattern, it is anticipated that by enforcing the minimum-fuel objective in the problem formulation, that the solver will exploit natural trajectories whenever they are available. The commercially available MATLAB software program GPOPS-II, which uses hp-adaptive Gaussian quadrature collocation methods to solve multi-phase nonlinear optimal control problems, was selected for its stability, ease of programming, and wide-spread use in both industry and academia. For more information on the GPOPS-II solution methods, refer to Patterson and Rao [84].

Interpreting Transfer Solutions via DST

The discussion on minimum-fuel optimal control solvers, and their potential exploitation of natural trajectories, leads to the question if these solutions are interpretable via the tools of Dynamical Systems Theory. This was the subject of a series of papers by Lo and Anderson [3–5,68] who investigated low-thrust trajectory design for Jovian moon tours. The authors concluded, through the use of Poincaré sections, that invariant manifolds play a significant role in traversing the unstable regions of the CR3BP phase space and that their optimization, given no prior knowledge of these natural solutions, converged to these paths. It is believed that a similar study of the solutions produced in the fuel-efficient transfers between QPT solutions would be a fruitful endeavor. However, this is left for future work.

Chapter 5

Sample Mission Analysis

This chapter demonstrates the Trajectory Design Methodology described in Chapter 4 through its application to a sample mission with target stars taken from the HabEx Interim Design Report [32]. These stars are tabulated with their Hipparcos number and location on the celestial sphere (right ascension and declination — ICRS J2000 frame) in Appendix B. Recall, the TDM was designed to be a foundation — built from the onset in DST — that is both adaptable to competing mission objectives and a platform from which future work can develop. Quantitative results on fuel expenditure and ΔV are provided, however, these largely reflect the TDM objective selected[1] and the assumptions made in the analysis (section 5.1). Thus, in the spirit of DST, emphasis is placed instead on the qualitative insights gained through applying this framework to solve the sample mission. In general, this can be summarized by two major steps. The first is a segmented analysis — documented in section 5.2 — that produces the local maximal path and trajectory for several subsequent time windows. The second — demonstrated in section 5.3 — combines the results of the local analyses to produce the global maximal path across an

[1]Recall, the TDM was designed to find the maximal path through a given time window. That is, the path that passes the formation LOS through the maximum number of stars, made continuous via the set of minimum-fuel, optimal, low-thrust transfers. Refer to sections 5.4 and 6.2.2 for more realistic mission scenarios that should be considered in the future.

extended time horizon. In addition, a star targeting analysis — documented in section 5.4 — is demonstrated, which showcases the adaptability of the TDM and its ability to contextualize alternative paths against competing objectives. Finally, section 5.5 considers the capabilities of the TDM framework within the context of the state-of-the-art.

5.1 Assumptions and Considerations

Several assumptions were made to scope this thesis to a tractable, yet representative, problem that is both insightful to the greater exoplanet imaging mission and scalable to increasing complexity. These assumptions are outlined in this section. First, the trajectory design space for the analysis is currently restricted to a single halo orbit and quasi-periodic torus for the observatory and external occulter respectively. Recall, this is inherent to the current TDM formulation (section 4.1). For further discussion on the implications of this assumption refer to section 6.2.3.

The observatory is assumed to operate in a Sun-Earth L2 northern halo orbit with an out-of-plane amplitude of 513,166 km and a period of nearly six months. The projection — in non-dimensional units — of the orbit onto the planes formed by the axes of the Sun-Earth rotating frame is shown in Figure 5-1. Table 5.1 lists key parameters characterizing the halo orbit, such as the out-of-plane amplitude,

Z amplitude	0.003430 [nd] = 513,166 km
T_p	3.090597 [nd]=0.490854 yrs
J_c	3.000758 [nd]
ν_1	1252.234 [nd]
ν_2	1.848699 [nd]
ν_3	2.000000 [nd]

Table 5.1: Observatory SEL2 halo orbit parameters

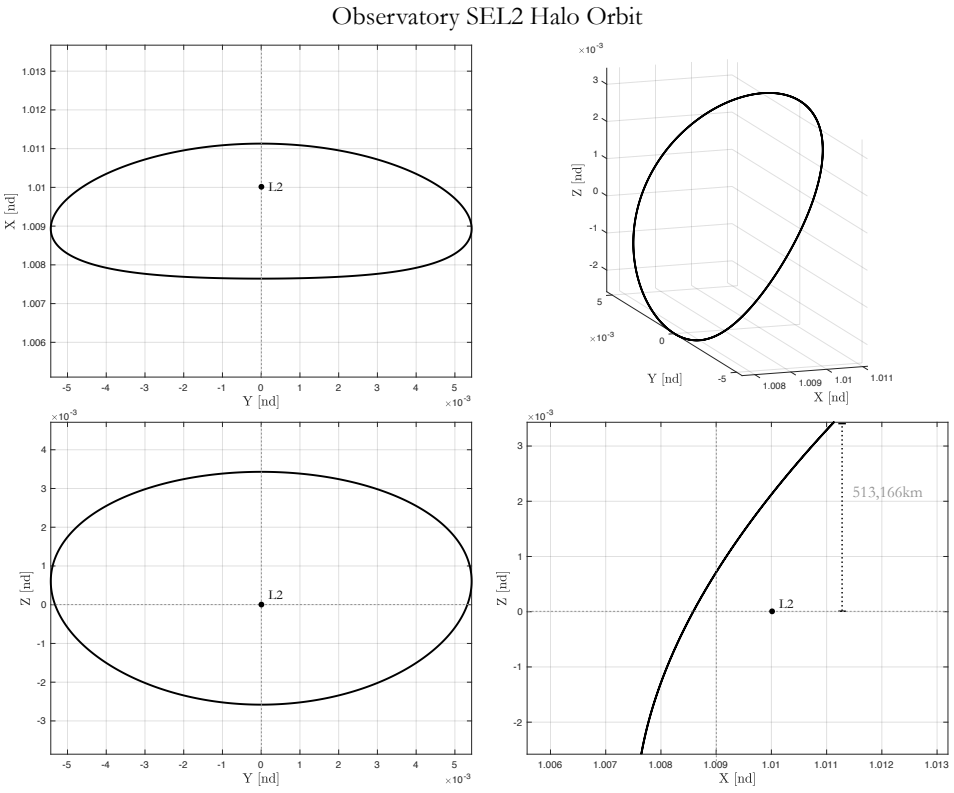

Figure 5-1: The observatory SEL2 northern halo orbit, plotted in the Sun-Earth system's rotating reference frame – non-dimensional units.

period T_p, Jacobi constant J_c, and its stability indices ${\nu_i}^2$.

As previously noted, the external occulter is assumed to exploit the set of solutions available from a single quasi-periodic torus throughout the duration of the mission. The QPT was chosen from the constant-period family surrounding the observatory halo orbit detailed in Table 5.1. The period[3] of the family was constrained to be equivalent to that of the halo orbit. Of those present in the family, the torus with average separation distance from its parent halo orbit[4] equal to 100,000 km — which is representative of the required separation distance for the formation during

[2]Refer to section 3.1.3 for a discussion on how to interpret the stability indices.

[3]That is, the time in which a trajectory on the torus maps back to the same invariant circle from which it started.

[4]Evaluated across one period, at three thousand equally space points in time, for all twenty-one trajectories saved in the database for the torus.

imaging — was selected. Table 5.1 lists several parameters of the QPT including the diameter of the invariant circle and torus number, ρ.

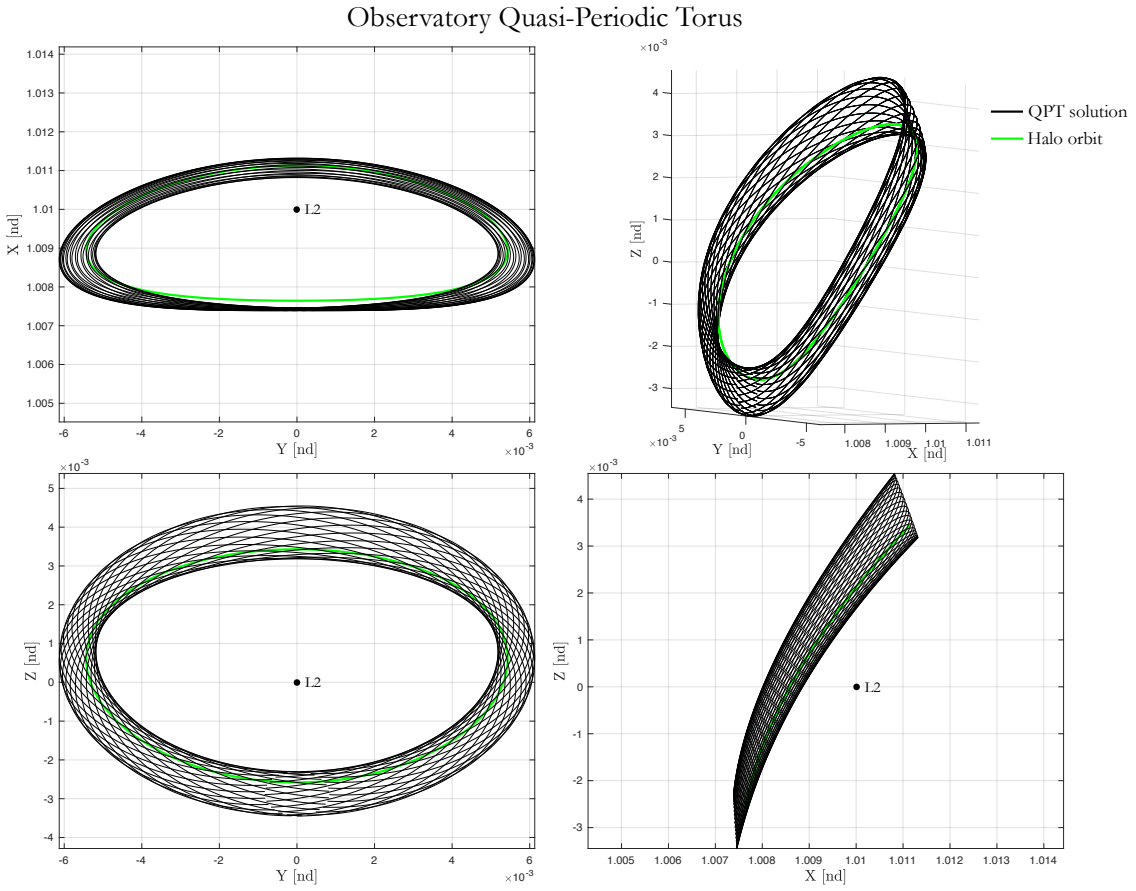

Figure 5-2: The external occulter operating QPT, plotted in the Sun-Earth system's rotating reference frame – non-dimensional units.

IC diameter	100,000 km
T_p	3.090597 [nd]=0.490854 yrs
J_c	3.000738 [nd]
ρ	0.455290 [rad]

Table 5.2: External Occulter QPT parameters

The TDM demonstrated in this thesis is currently restricted to the problem of slewing the external occulter to transfer the formation LOS between target stars. An operational direct exoplanet imaging mission will require the spacecraft to inertially fix their line-of-sight — maintaining precise separation and lateral displacement conditions — at each star for the required instrument integration time. This problem is considered beyond the scope of this thesis and is not included in the sample mission analysis. Instead, retargeting of the formation LOS to star $i+1$ in the maximal path (determined in step two of the TDM) begins immediately upon alignment with star i. Since only transfer is considered, the external occulter is assumed to be a particle of mass m. Where $m = 1200$ kg, in accordance with the NASA S5 Starshade Report [30]. The observatory is also considered to be a particle, though this is inconsequential since its state evolves kinematically in its orbit. Actuation is restricted to the use of low-thrust, ion propulsion engines to minimize fuel mass. Finally, the analysis is restricted to the Circular Restricted Three Body Problem phase space; no additional disturbances are included. Despite these assumptions, it is should be noted that this sample mission analysis is meant to be demonstrative of a general framework that can, and should be scaled up to more complex solutions in the future. As the framework evolves, so should the complexity of cases added to the analysis. For a more detailed discussion on recommendations for future work that relates to these assumptions, refer to section 6.2.

5.1.1 Star Transit Rules

Recall, the set of feasible retargeting maneuvers available to the external occulter after each formation/star alignment is determined by the star transit rules of section 4.2.1[5]. Ultimately, the objective is to determine if the two-point boundary value

[5]Here, a feasible retargeting maneuver refers to the ability — or lack thereof — of the external occulter to transfer from its current state on the torus to a subsequent QPT trajectory that will yield formation alignment with a target star at a known future time. From section 4.3, it was demonstrated that these transfers are the solution of a two-point boundary value problem with

problem constraining the transfer maneuver is solvable. For the purposes of this thesis, these cases can be correlated to the maximum thrust producible by the onboard propulsion.

While ordering the trajectory optimization software to attempt every possible transfer within the time window would guarantee an answer, this brute force method would be an inefficient use of computational resources. Instead a two-layer heuristic method was employed. The premise is as follows: when analyzing the feasibility of transfers between QPT solutions, thruster saturation will manifest itself as a curve in $(\Delta\theta_{0,2}, \Delta t)$ space, where $\Delta\theta_{0,2}$ is the change in initial latitudinal angle between the two QPT solutions defining the initial and final conditions of the transfer, and Δt is the fixed time of transfer. Any transfer between solutions on the torus can be defined by these two parameters. For a given propulsion system, every transfer of $\Delta\theta_{0,2}$ will have a corresponding Δt_{sat} at which the thrusters will saturate; yielding the condition that all transfers commanded within $\Delta t < \Delta t_{sat}$ are not solvable, while those constrained to $\Delta t > \Delta t_{sat}$ are solvable. After evaluating and plotting several $\Delta\theta_{0,2}$, these saturation points fit to a curve. In general, solvability is state dependent, and a given $\Delta\theta_{0,2}$ will have multiple Δt_{sat}; though these values are found to be tightly grouped. To capture this uncertainty, upper and lower bounds are determined empirically and placed around the saturation curve. The analysis then proceeds as follows:

- For a given transfer, determine $\Delta\theta_{0,2}$ and Δt. If the point ($\Delta\theta_{0,2}$,Δt) is outside of the bounds and to the left of the saturation curve, then mark the transfer as unsolvable and discard.

- If the point ($\Delta\theta_{0,2}$,Δt) is outside the bounds and to the right of the saturation curve, then mark the transfer as feasible (solvable).

- If the point ($\Delta\theta_{0,2}$,Δt) is anywhere within the bounds of the saturation curve,

fixed final time, which often fail to have a solution.

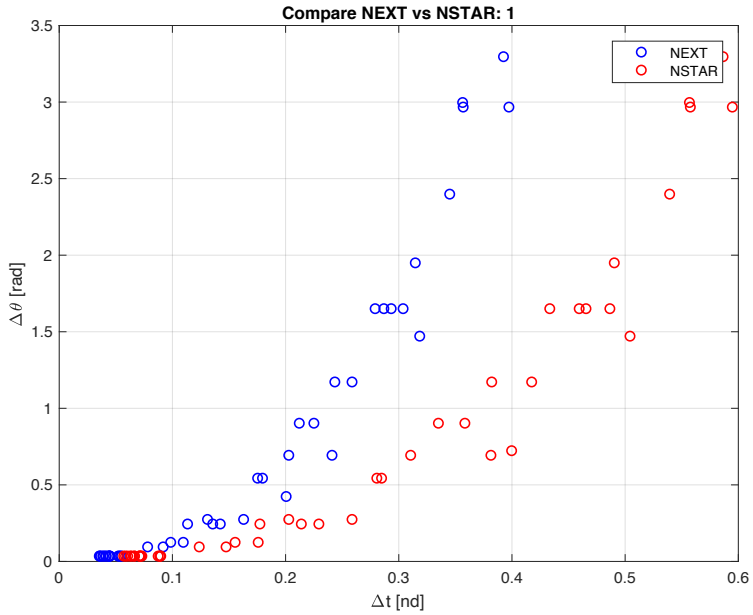

Figure 5-3: Saturation curves for NSTAR and NEXT ion propulsion engines.

attempt to solve the two-point boundary value problem and mark the transfer according to its success or failure.

These feasibility constraints are dependent on the actuators selected, and it is expected that a more capable propulsion system will shift the curves favorably. This is precisely what was observed. The analysis was performed for both the NSTAR [103] and NEXT [83] ion propulsion engines. Their computed saturation curves are shown together in Figure 5-3. The curve for the NEXT engine (which saturates at 236 mN of thrust) is shifted to the left of the NSTAR engine (with maximum thrust of 90 mN), corresponding to a greater space of solvable solutions for the former configuration. Figure 5-4 summarizes the performance for the NSTAR (left) and NEXT (right) thrusters. The saturation points are shown as inverted triangles, the quadratic curve fit through those points is shown in black, and the bounds applied around the curve are shown in red. The curves are reflective across $\Delta\theta_{0,2} = \pi$, acknowledging the periodic structure of the invariant circle. Every transfer combi-

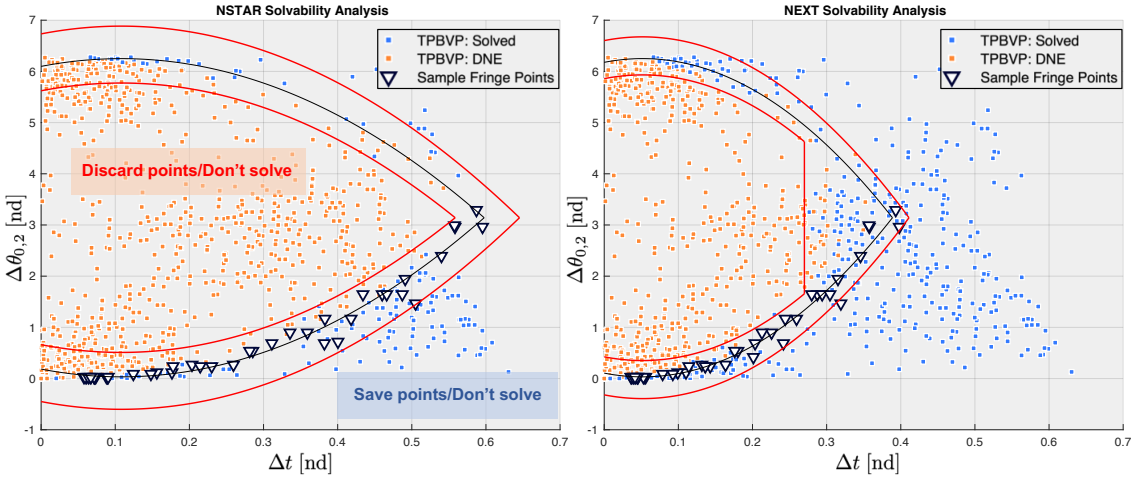

Figure 5-4: TPBVP solutions plotted in $(\Delta\theta_{0,2}, \Delta t)$ space, with saturation curves shown for NSTAR and NEXT ion propulsion engines.

nation within Time Window 1 (refer to section 5.2) was directly computed by the GPOPS-II solver and marked as a red dot if the solution did not exist (DNE) or blue if a solution was found. The discard and save conditions are explicitly marked on the left image. Thus, as compared to the NSTAR engine, the higher performing NEXT thruster configuration opens a larger space of transfers that are both solvable and do not need to be directly computed. All future analysis will assume use of this engine, and will correspondingly use the heuristic regions illustrated in the right figure to frame the star transit rules during the star ordering step of the TDM.

Final Selection of the Maximal Path

It was mentioned in section 4.2.2 that multiple paths[6] can exist through a given time window that yield an equivalent number of star alignments. In general, a trajectory designer can select from these paths in accordance with the mission objective. For the TDM developed, the path expending the minimum amount of fuel[7]

[6]Path here refers to a sequence of stars in a time window to target, and not specifically the external occulter trajectory that will pass the formation LOS through them.

[7]Under the two point boundary value problem with fixed final time constraints.

from those yielding the maximum number of intercepts is desired. This path is selected — when multiple are encountered — as the one that minimizes the total change in initial latitudinal phase between its corresponding QPT solutions: $\Delta\theta^{total}_{2,0,j} = (\sum_{i=1}^{n-1} |\theta^{i+1}_{2,0,j} - \theta^{i}_{2,0,j}|)$ — where n is the total number of star intercepts of path j and $\theta^{i}_{2,0}$ is the latitudinal phase of the QPT solution on the initial invariant circle that passes the formation LOS through target star star i. This heuristic is consistent with the solution parameterization expressed in section 4.1.1, and has predicting the desired path in all time windows analyzed. Figure 5-5 illustrates the set of LOS tracings formed by simulating the trajectories that pass the formation alignment through each of the "maximal paths" available within Time Window 5. The path of minimum fuel expenditure, costing 13.26 kg, is shown in green. That of maximum fuel expenditure, costing 17.95 kg is shown in red. All others in between these extremes are plotted in gray. Thus, by choosing the path that maximizes exploitation of natural solutions, one can save over 4.5 kg for just this one time window.

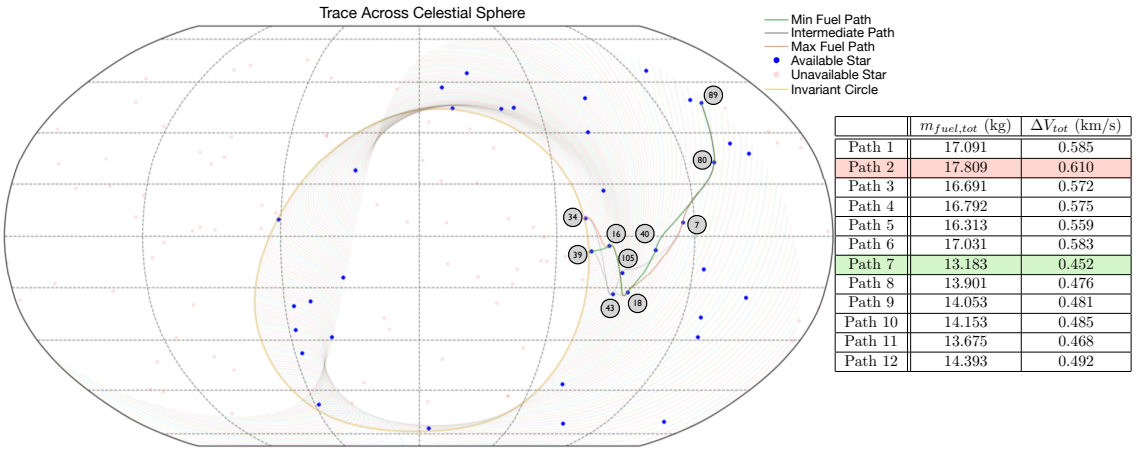

	$m_{fuel,tot}$ (kg)	ΔV_{tot} (km/s)
Path 1	17.091	0.585
Path 2	17.809	0.610
Path 3	16.691	0.572
Path 4	16.792	0.575
Path 5	16.313	0.559
Path 6	17.031	0.583
Path 7	13.183	0.452
Path 8	13.901	0.476
Path 9	14.053	0.481
Path 10	14.153	0.485
Path 11	13.675	0.468
Path 12	14.393	0.492

Figure 5-5: Trace of all paths leading through the maximum number of star intercepts for Time Window 5. The path corresponding to minimum fuel expenditure is marked in green, while the path of maximum fuel expenditure is marked in red.

5.2 Time Window Analyses: Finding Local Maximal Paths

This section demonstrates the Trajectory Design Methodology applied to four subsequent time windows, each a quarter period in length (0.123 years). This is a segmented analysis that individually produces the local maximal path and trajectory for each of these windows. Together, they form a basis for computing the global maximal path across an extended time horizon (e.g., a full mission). This latter task is accomplished via the procedure documented in section 5.3. As a note, the analysis was extended through Time Windows 5 – 10, the results of which are cataloged in Appendix A.

Beyond characterizing their performance, the objective is to gain qualitative insight into how the maximal paths selected by the TDM fit into the greater dynamical landscape. To that end, three figures are provided for each time window analyzed. First is the set of LOS tracings simulated from each initial phase on the refined invariant circle[8] (see section 4.1.2), which provides a dense visualization of the flow across the celestial sphere through each segment of time. Second is a plot of the subset of tracings, made by the converged QPT solutions found by the search by shooting method of section 4.1.4. These two plots are effectively the output of the first step of the TDM, with the latter illustrating the set of pathways passing through each available target star within the time window. The final plot shows the tracing made by the simulated trajectory through the maximal path, made continuous by the transfers computed in step three of the TDM. Tables detailing estimates of the fuel mass and ΔV for each transfer are included.

[8]Here, "starting from each initial phase," refers to the initial condition of the external occulter — since the observatory starts from the same halo orbit initial condition. These tracings are produced by simulation.

Time Window Analyses: Results

Recall from section 4.2 that in order to maximize science efficiency, the TDM first finds the set of paths that pass the formation LOS through the maximum number of stars available in the time window. If there are multiple paths passing through an equal, maximum number of stars, the path of lowest fuel consumption from that subset is selected. This competing objective directly manifests in the results obtained. While in general it is observed that the TDM selects local maximal paths that align well with the dynamical flow, individual transfer maneuvers can be observed that deviate from this trend to greater maximize the total science yield. Ultimately, the decision to pursue these additional, costly stars at the expense of fighting nature (and thus, fuel) will need to be determined by the mission designer[9]. It should be noted though that the TDM provides the dynamical context from which these trade-off decisions can be made. For a deeper discussion and analysis on the incremental cost of targeting individual stars, refer to section 5.4.

This behavior is best observed in the results of the Time Window 1 analysis, summarized by Figures 5-6 and 5-7. The local maximal path intercepts six target stars [10] for a total estimated cost of 10.204 kg of fuel and 0.350 km/s of ΔV (refer to section 5.5 for a discussion on how these estimates were computed). An effective way of contextualizing the results within the dynamical landscape is to observe the angle at which the maximal path crosses the tracing produced by the natural flow. Figure 5-6 (top) illustrates a dense view of these tracings (shown again in Figure 5-7 in grey to highlight the maximal path in green). Transfers that minimize this angle generally align well with natural solutions. On the contrary, segments of the

[9]It should be noted that the trivial case exists such that if the TDM were designed to strictly maximize fuel-efficiency, without considering science yield, only one star would be intercepted for each window for zero cost (since the formation would reach its alignment condition ballistically and no further transfers would be made). Thus, to make any sort of interesting progress for real mission, the objectives need to be balanced. For further discussion, refer to section 6.2.2.

[10]It should be noted that from the perspective of the formation, stars 74 and 75 are close in latitudinal/longitudinal space, and are thus indistinguishable from each other on the celestial sphere plots.

maximal path that pass nearly perpendicular to the flow rely heavily on the onboard thrusters to fight the natural dynamics at the expense of fuel. This phenomena can be observed in the cost of transfer between stars 79 to 55. This maneuver accounts

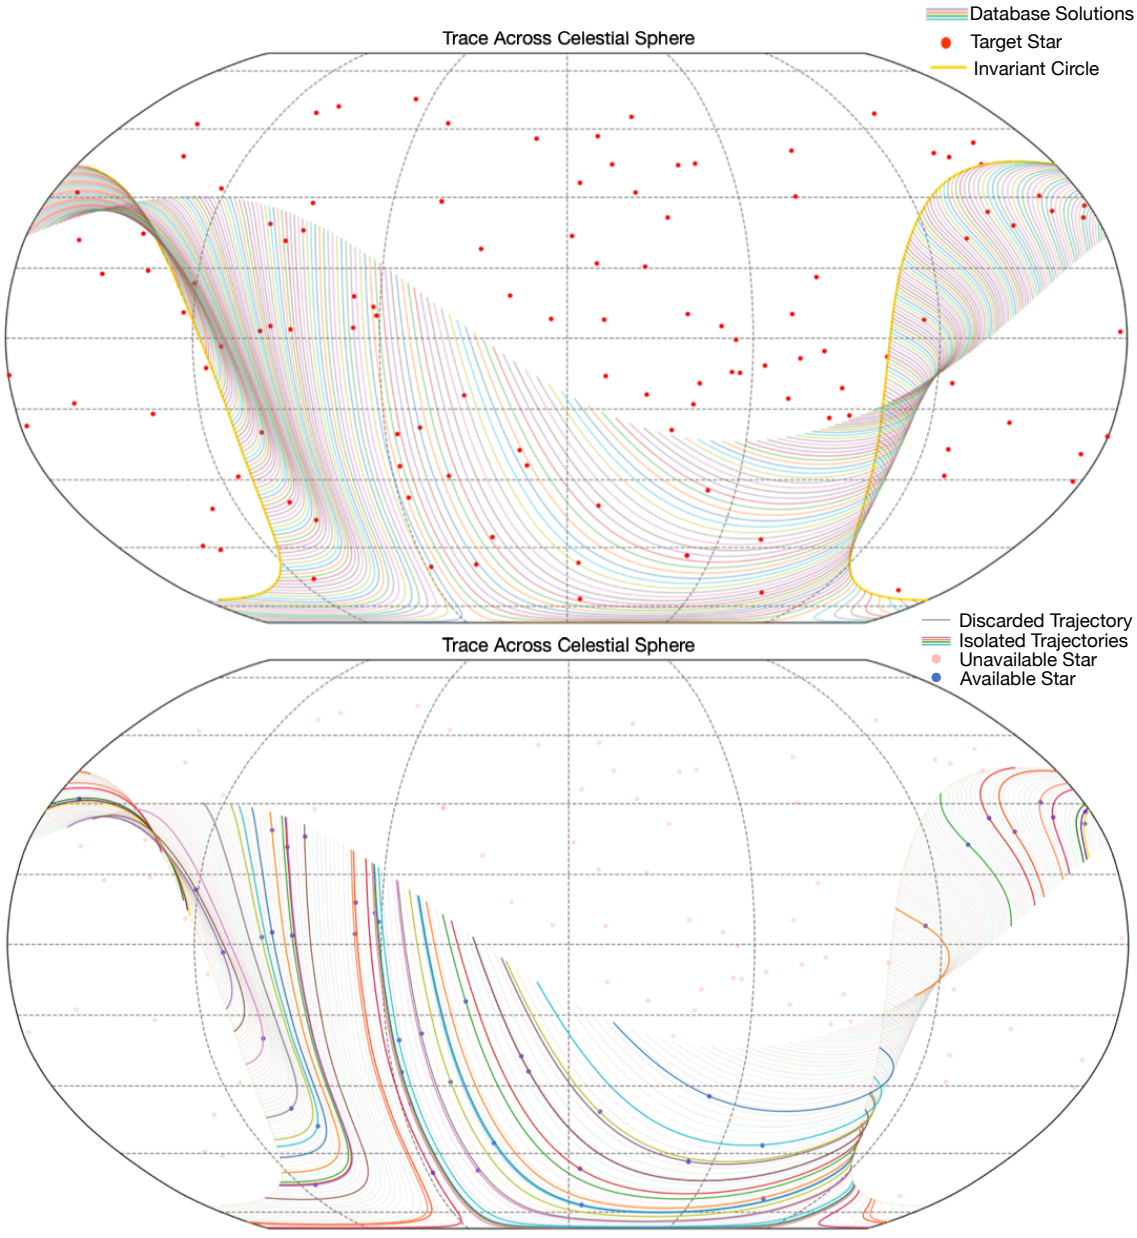

Figure 5-6: Time Window 1: tracings starting from the refined invariant circle (top), and the subset that pass through each available star in the time window (bottom).

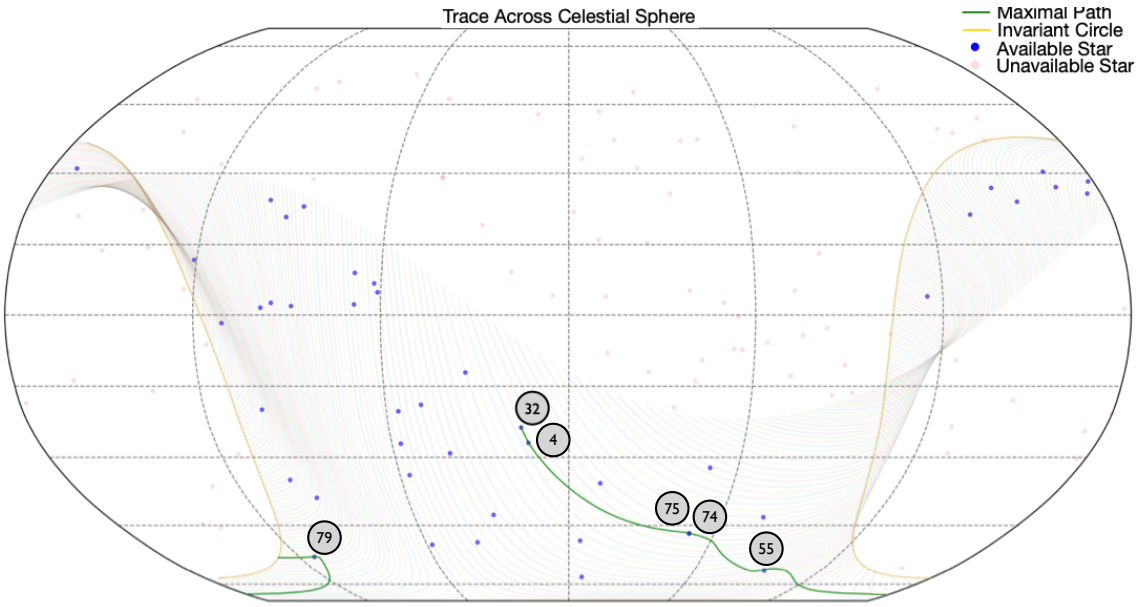

Star ID Transfer	$79 \rightarrow 55$	$55 \rightarrow 74$	$74 \rightarrow 75$	$75 \rightarrow 4$	$4 \rightarrow 32$	*Total*
m_{fuel} (kg)	5.156	2.515	0.022	2.201	0.310	10.204
ΔV (km/s)	0.177	0.086	0.001	0.075	0.011	0.350

Figure 5-7: Time Window 1: path leading to the maximum amount of star intercepts in the minimum amount of fuel (green). Table summarizing each transfer included.

for 50% of the fuel expenditure across the entire time window. From Figure 5-7, it is evident this transfer does is in fact run nearly perpendicular to several natural transit paths. However, following this maneuver, the TDM selects four stars that are relatively "well-aligned" with the flow and together have a total retargeting cost equal to the first. In the future, mission designers will need to either weigh selecting or forgoing these hard-to-reach stars or further tune the TDM objective to handle these case under the mission constraints. For recommendations on how these concepts should be approached in future work, refer to section 6.2.2.

A different scenario can be viewed in the maximal path produced through Time Window 2 (Figure 5-8 and 5-9). In this region, the TDM converged to a local maximal path that, in general, aligns well with the natural flow tracings.

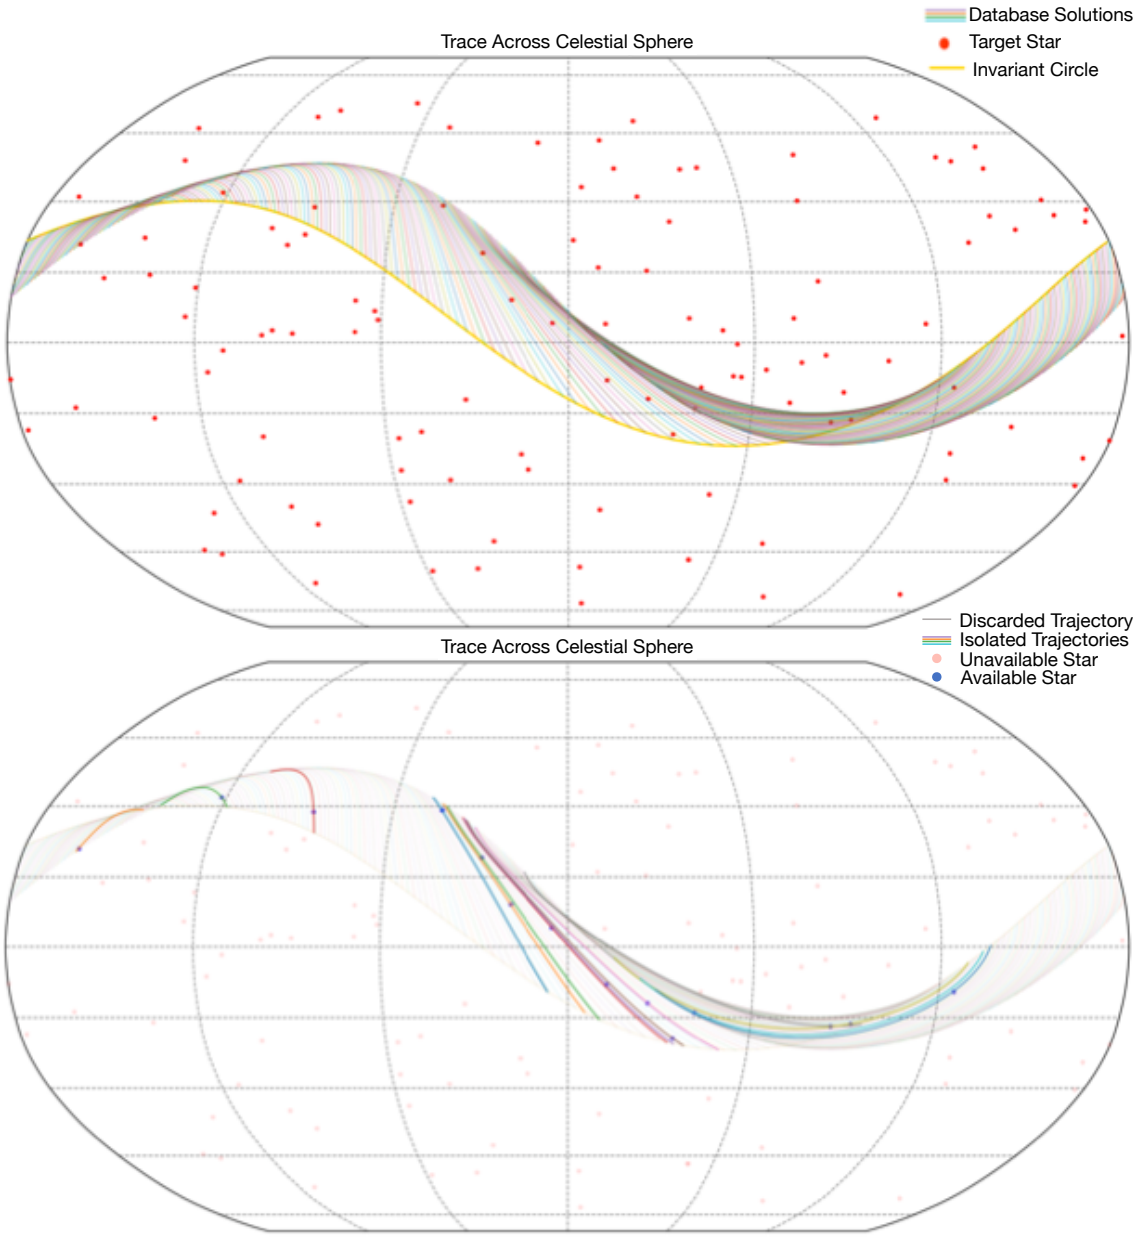

Figure 5-8: Time Window 2: tracings starting from the refined invariant circle (top), and the subset that pass through each available star in the time window (bottom).

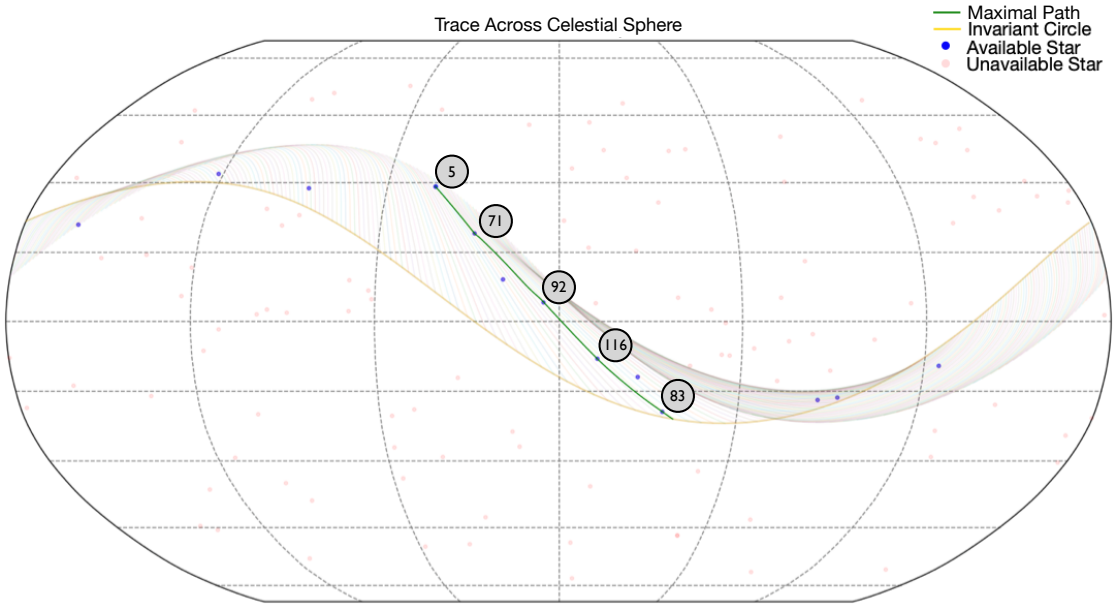

Star ID Transfer	$83 \rightarrow 116$	$116 \rightarrow 92$	$92 \rightarrow 71$	$71 \rightarrow 5$	*Total*
m_{fuel} (kg)	0.811	0.389	2.317	0.749	4.226
ΔV (km/s)	0.028	0.013	0.079	0.026	0.146

Figure 5-9: Time Window 2: path leading to the maximum amount of star intercepts in the minimum amount of fuel (green). Table summarizing each transfer included.

As a result, a total of five stars are intercepted for just over 40% of the fuel cost of the six intercepts made in Time Window 1. These were performed at a cost of 4.226 kg of fuel and 0.146 km/s of ΔV.

The plots corresponding to Time Window 3 are provided in Figures 5-10 and 5-11. Only two stars can be targeted within this segment under the current TDM constraints. The cost is a total of 5.661 kg of fuel and 0.194 km/s of ΔV. It can be observed that solutions at the top left of the plot reach an apex in latitude, allowing for the possibility of retargeting previously intercepted stars. However, it should be noted that the apex does not imply that stars above this latitude are unreachable. The formation LOS tracing is dependent on both the natural flow and position of the rotating frame with respect to the inertial frame. It can be observed in Appendix A that other windows later in the year provide intercept of stars outside the scope of these first four.

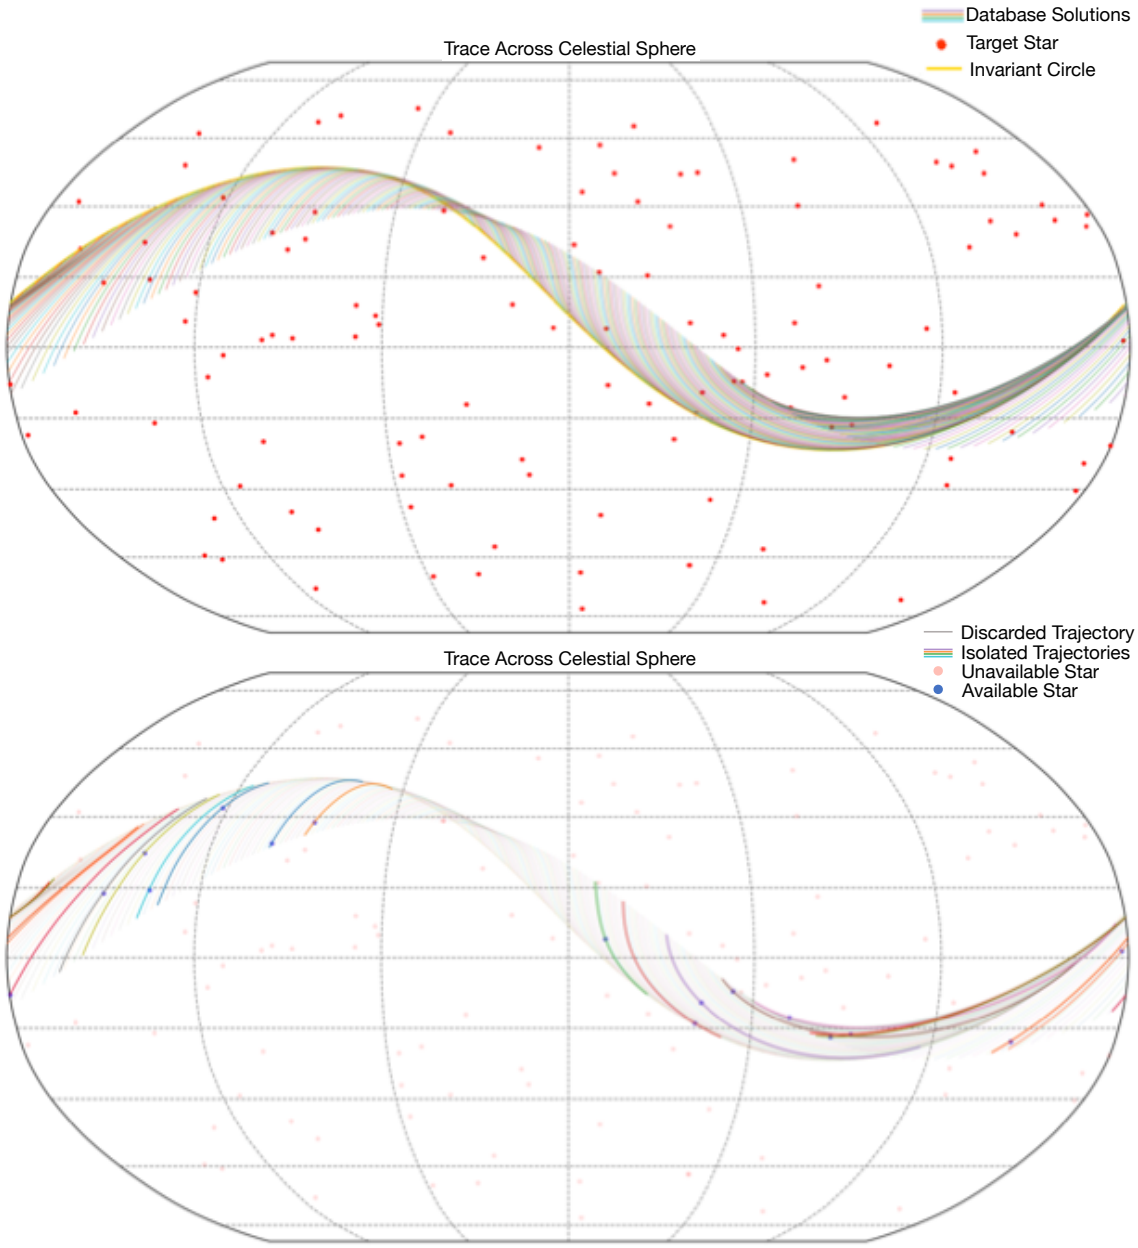

Figure 5-10: Time Window 3: tracings starting from the refined invariant circle (top), and the subset that pass through each available star in the TW (bottom).

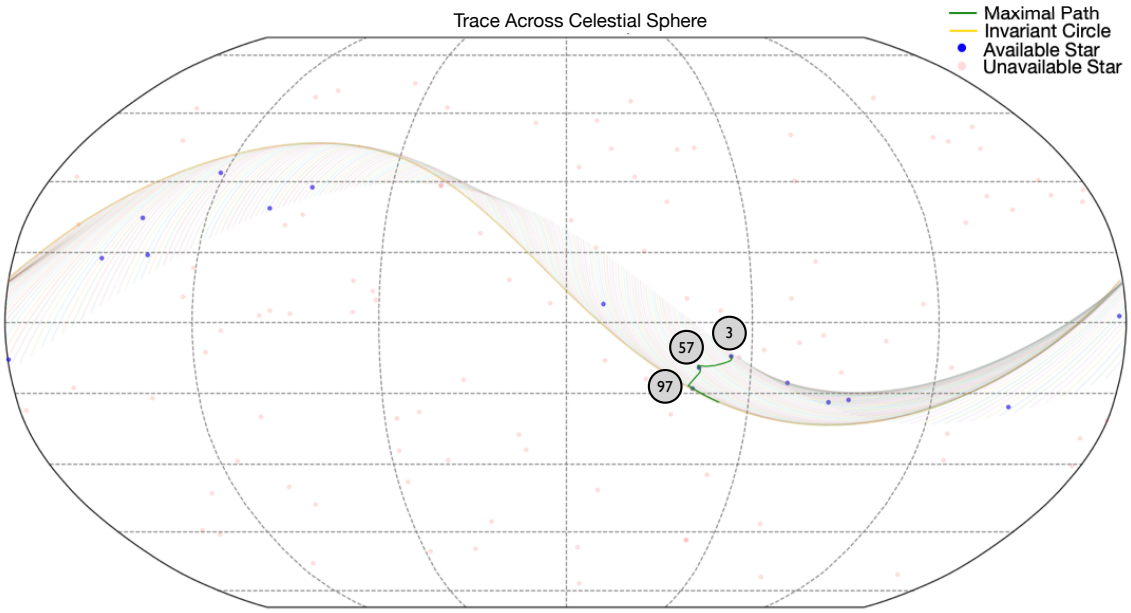

Star ID Transfer	$97 \rightarrow 57$	$57 \rightarrow 3$	*Total*
m_{fuel} (kg)	2.550	3.111	5.661
ΔV (km/s)	0.087	0.107	0.194

Figure 5-11: Time Window 3: path leading to the maximum amount of star intercepts in the min. amount of fuel (green). Table summarizing each transfer included.

Finally, the results of Time Window 4 are summarized in Figures 5-12 and 5-13. Here, the TDM again converges to a local maximal path that aligns well with the natural flow, intercepting five target stars for a total cost of 4.334 kg of fuel and 0.148 km/s of ΔV. Though Time Window 4 completes an orbital period of revolution, tracings do not repeat themselves until the rotating reference frame completes a revolution about the inertial, celestial frame. For the Sun-Earth system this takes nearly a year, and can be observed starting in Time Window 9.

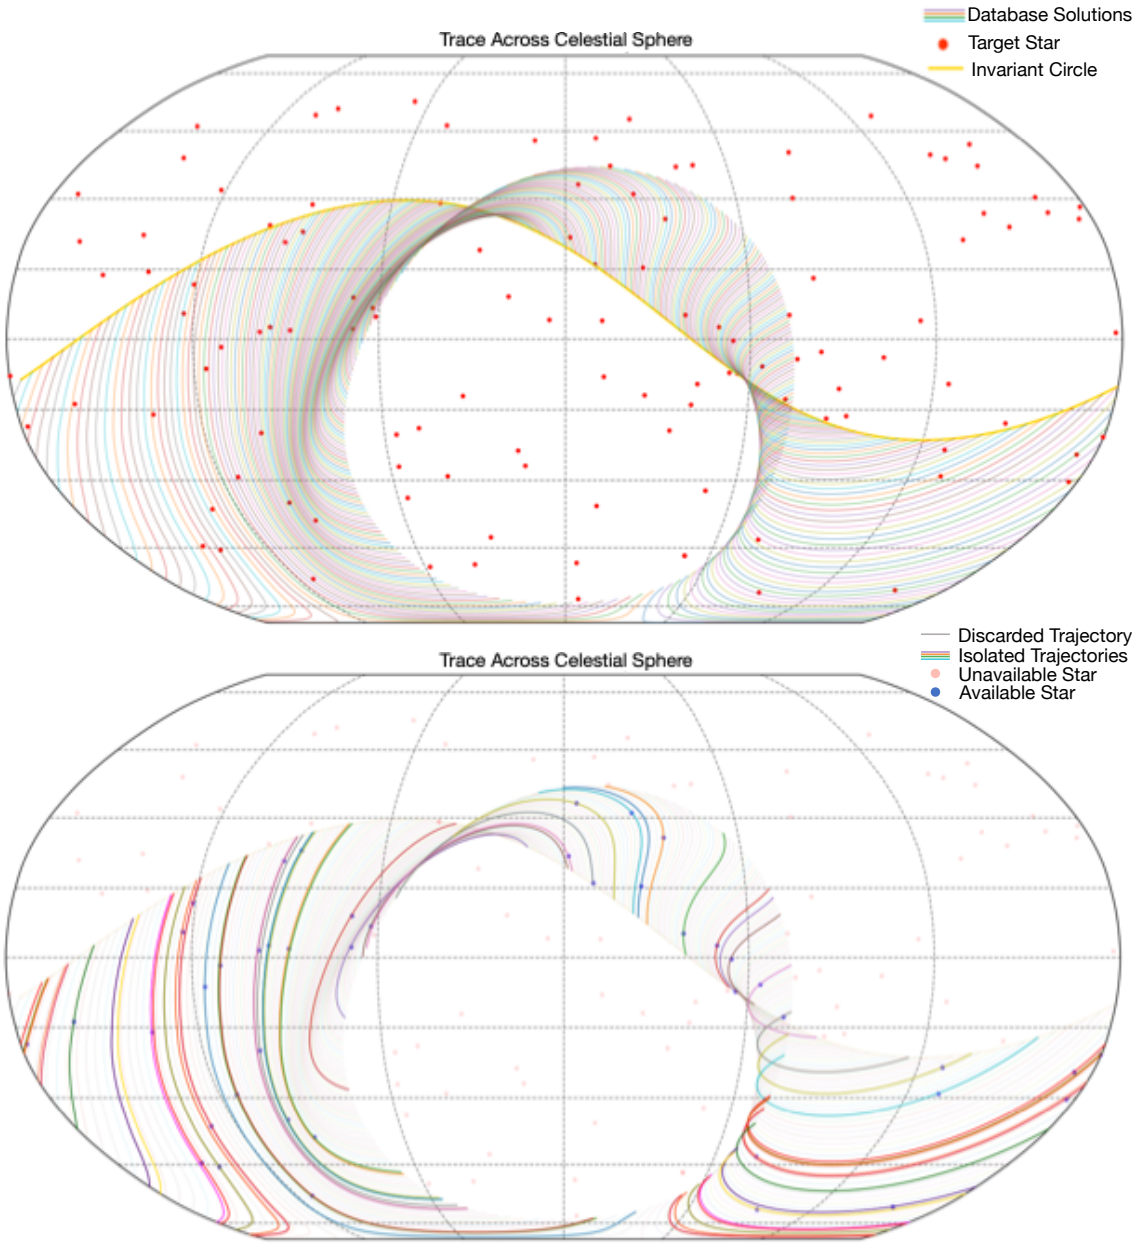

Figure 5-12: Time Window 4: tracings starting from the refined invariant circle (top), and the subset that pass through each available star in the TW (bottom).

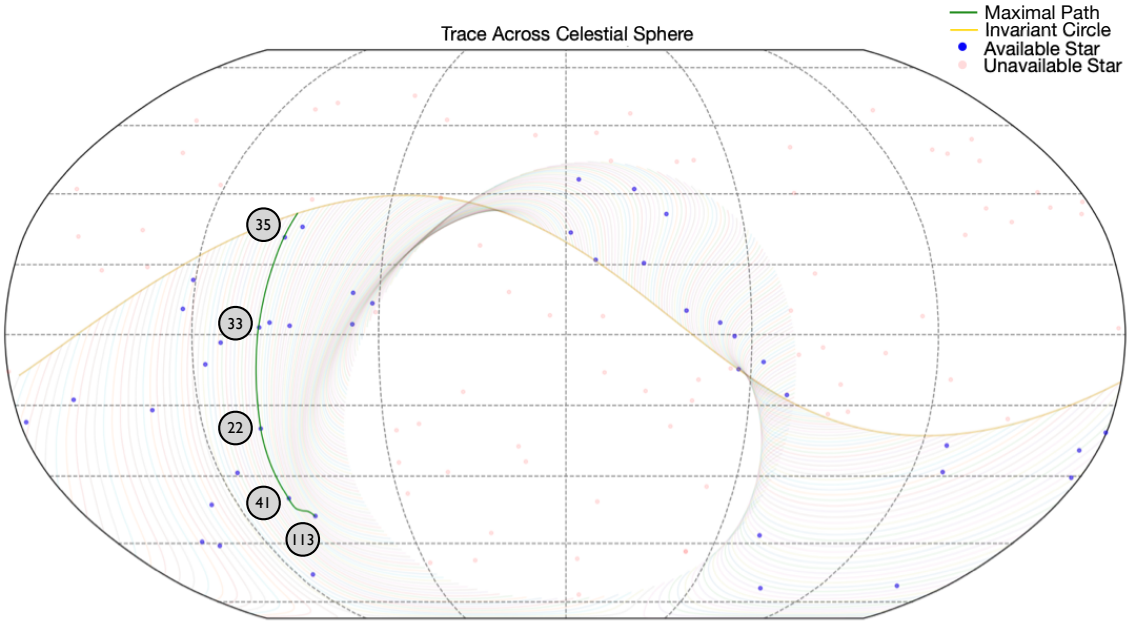

Star ID Transfer	$35 \rightarrow 33$	$33 \rightarrow 22$	$22 \rightarrow 41$	$41 \rightarrow 113$	*Total*
m_{fuel} (kg)	0.103	0.435	1.071	2.725	4.334
ΔV (km/s)	0.149	0.015	0.037	0.093	0.148

Figure 5-13: Time Window 4:path leading to the maximum amount of star intercepts in the min. amount of fuel (green). Table summarizing each transfer included.

5.3 Combining Time Windows: Finding Global Maximal Paths

The ultimate objective of the sample mission analysis is to produce the global maximal path across an extended time horizon — such as a full period, year, or mission. This is accomplished by combining the local results of the segmented time window analyses demonstrated in section 5.2. The procedure outlining this task is documented below[11].

First, it is assumed that all steps of the TDM for N subsequent time windows have been precomputed and their results stored. The objective is to combine the

[11]To note, both the individual time window and combined time window analyses were completed on the order of hours using a laptop computer equipped with a 2.8 GHz Quad-Core Intel Core i7 processor and 16GB of memory.

data from the N segments into a single, continuous time window and produce the global maximal path[12]. This is done in a backwards fashion.

1. Determine the set of feasible transfers between the stars of Time Window $N-1$ and Time Window N. This effectively stitches the time windows together by providing the set of all possible connecting trajectories.

2. Recursively re-score the stars of both time windows via the star transit rules outlined in section 5.1.1, starting from the terminal stars of Time Window N, working down to the root stars of Time Window $N-1$. (At this point the data from these two time windows are effectively combined. The maximal path through Time Windows $N-1$ and N can be extracted or the analysis continued through subsequent windows.)

3. To connect Time Window $N-2$, determine the set of feasible transfers between the available stars of Time Windows $N-2$ and $N-1$ only. It is assumed that Time Window N is far enough in the future that skipping Time Window $N-1$ and transferring directly to the stars of Time Window N is always suboptimal, with respect to the objective of maximizing the total number of star alignments[13].

4. Recursively re-score all stars within the analyzed time windows ($N-2$ through N). This reflects the objective that the maximal path through the extended time horizon is sought.

5. Continue the process until Time Window 1 is connected.

[12]As with the local analyses, maximal path refers to the path that intercepts the formation LOS with the maximum number of stars in the extended time horizon. If multiple exist, the minimum-fuel path is selected.

[13]Part of the computational savings made by breaking an extended time window into smaller segments are realized in this step. Computing the set of feasible transfers between available stars is the most expensive step of the TDM. By restricting these calculations to the stars within a time window and adjacent time windows, needless complexity is discarded.

The constraint is added that the formation cannot retarget the same star consecutively. To clarify, the external occulter is restricted from slewing from alignment with star i to alignment with star i. There has to be at least one other unique star imaged in between revisits. This case can appear if a given star is available to two subsequent time windows. As a final note, it is believed that the methodology described would apply well to a receding time horizon formulation — as inspired by the field of Model Predictive Control. Instead of the batch method described, the time window analysis would be resolved at the end of each star alignment, and the maximal path recomputed to the end of the segment. However, this study is left for future work. For further discussion refer to section 6.2.

Combining Time Windows: Results

To demonstrate the extended time horizon analysis, the global maximal path through Time Windows 1 – 2 and Time Windows 1 – 4 were computed. The results are summarized in Figures 5-14 and 5-15 respectively. Similar to the segmented analysis, fuel consumption and ΔV costs are provided, however, emphasis is placed on the qualitative insight gained into how the TDM selects stars from the local path to produce the global trajectory.

First, it should be noted that the global maximal path is not simply the superposition of the local maximal paths for each window stitched together. Instead, the algorithm reoptimizes for the global solution given all available data. This can be observed in Figure 5-14, which illustrates the global maximal path across the first two time windows. While the algorithm incorporates the first six stars from the maximal path of Time Window 1, the transfer between stars 32 and 83 that would graft-in the local solution of Time Window 2 is not feasible under the given constraints. This limits the total number of stars that can be targeted to eight total intercepts, which is less than what the individual window analyses would suggest.

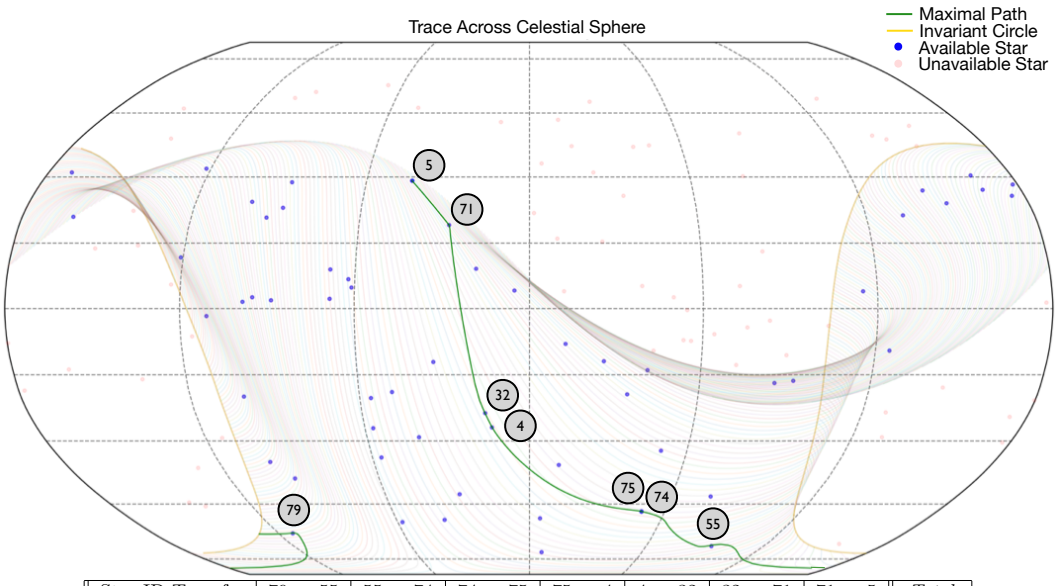

Star ID Transfer	$79 \rightarrow 55$	$55 \rightarrow 74$	$74 \rightarrow 75$	$75 \rightarrow 4$	$4 \rightarrow 32$	$32 \rightarrow 71$	$71 \rightarrow 5$	*Total*
m_{fuel} (kg)	5.155	2.515	0.022	0.309	1.809	1.805	0.749	12.762
ΔV (km/s)	0.177	0.082	0.001	0.011	0.011	0.062	0.026	0.437

Figure 5-14: Time Windows 1 & 2 combined: Path leading to the max amount of star intercepts in the min amount of fuel. Table summarizing each transfer included.

This is performed at a cost of 12.762 kg of fuel and 0.437 km/s of ΔV[14].

Similar behavior can be observed in the global maximal path produced through Time Windows 1 – 4 (refer to Figure 5-15), where thirteen total stars are intercepted. In this extended time horizon, the algorithm incorporated all stars from the local maximal path for both Time Windows 1 and 4. However, only the last two stars of the local maximal path from Time Window 2 were included, none from that of Time Window 3 made the list. Instead, one star that was previously suboptimal (star 118) was selected in their place. This is achieved at a cost of 24.333 kg and 1.480 km/s of ΔV.

As a final note, the combined maximal path through Time Windows 1 – 10 (just over 1.2 years) is provided in tabular form at the end of Appendix A. Thirty-seven total star intercepts were achieved costing a total of 110.277 kg of fuel and 3.777

[14]As discussed in section 5.2, the performance of the maximal path directly reflects the TDM objective. Since it seeks first to maximize science yield, expensive maneuvers that deviate from the natural flow can occur. The objective can and should be adapted to meet future mission objectives.

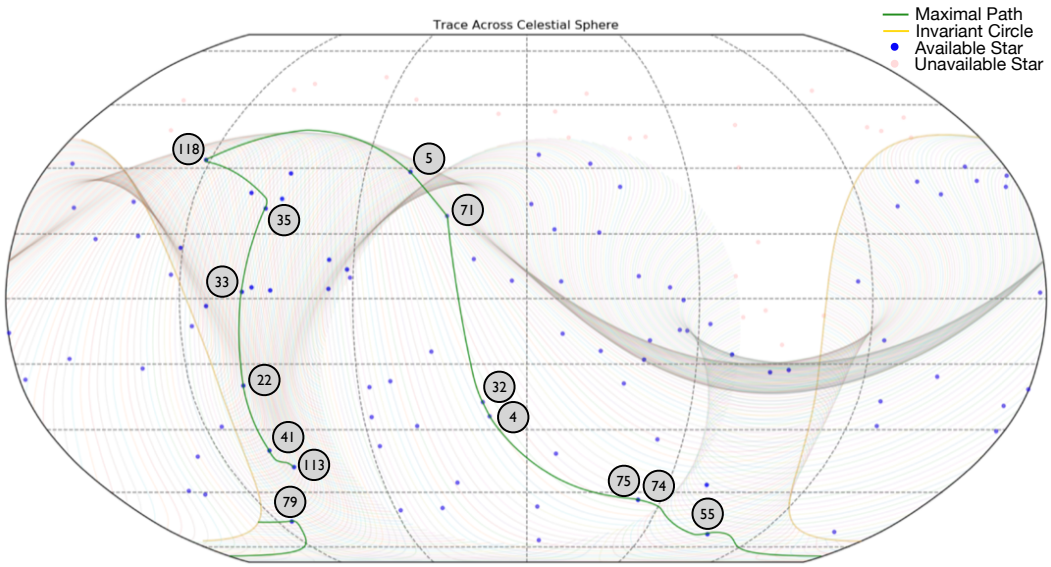

Star ID Transfer	$79 \rightarrow 55$	$55 \rightarrow 74$	$74 \rightarrow 75$	$75 \rightarrow 4$	$4 \rightarrow 32$	$32 \rightarrow 71$	$71 \rightarrow 5$	*Total*
m_{fuel} (kg)	5.180	2.557	0.022	2.205	0.307	1.780	0.692	
ΔV (km/s)	0.177	0.088	0.001	0.076	0.105	0.610	0.024	
Star ID Transfer	$5 \rightarrow 118$	$118 \rightarrow 35$	$35 \rightarrow 33$	$33 \rightarrow 22$	$22 \rightarrow 41$	$41 \rightarrow 113$		
m_{fuel} (kg)	3.376	3.935	0.074	0.424	1.012	2.769		24.333
ΔV (km/s)	0.116	0.135	0.003	0.015	0.035	0.095		0.0833

Figure 5-15: Time Windows 1-4 combined: Path leading to the max amount of star intercepts in the min amount of fuel. Table summarizing each transfer included.

km/s of ΔV. It should be noted that the final maximal path through the ten time windows contains several star revisits. This reflects the objective of finding the trajectory that maximizes the number of star alignments, agnostic to their identity, in the minimum amount of fuel. In reality, science — not engineering — will drive this decision. The complexities of prioritizing certain stars and specific revisits (over fuel or total number of alignments) will need to be integrated into the framework to adapt to the needs of future missions. A cursory examination of this problem is provided in section 5.4.

5.4 Adapting the TDM: Targeting Select Stars

It was demonstrated in section 5.3 how to combine the results from subsequent time window analyses to produce the global maximal path across an extended time

horizon — such as a full period, year, or mission. It should be noted that this result is comprehensive, and effectively solves the observatory/external occulter formation flying retargeting problem as scoped for this thesis[15]. Referring to the problem statement introduced in section 1.3, a Trajectory Design Methodology was sought that maximizes exploitation of natural solutions to retarget the formation LOS, while minimizing sacrifice to science efficiency. It is argued that the framework developed — by constraining formation alignments to occur along the surface of a quasi-periodic torus, extracting the maximum sequence of stars that can be feasibly retargeted within a time window, and connecting the natural solutions through optimized, minimum-fuel transfers — presents a novel approach that satisfies this objective under the aforementioned assumptions.

Despite this, in no way is it claimed that the maximal path produced is the globally optimal solution. A host of directions can be taken to further develop this work, stripping back assumptions to add realism and complexity, pursuing more topological structures of the CR3BP design space, and adapting to competing mission objectives. Most of these considerations are left — and encouraged — for future work (refer to section 6.2). However, one key scenario is highlighted and considered necessary for examination.

The method in which the demonstrated TDM selects and orders stars into paths is agnostic to their identity. Though this feature is consistent with objective of this thesis, a more realistic mission scenario would prioritize targeting certain stars over others, potentially at the cost of total fuel consumed and/or raw number of intercepts[16]. Thus, tallying multiple revisits to the same star — as was observed in the combined maximal path for Time Windows 1 – 10 — would likely be an undesirable outcome in this case. Ultimately, it is advocated that a balanced approach should

[15]Refer to section 5.1 for a discussion on how this scope is defined.

[16]A realistic scenario could also be that if a subset of target stars are prioritized — and a low raw number of intercepts is an acceptable outcome — then fuel could be saved by reducing the total number of retargeting maneuvers overall.

be taken. One way this could be accomplished is by enforcing alignment with the desired subset in step two of the TDM while still adhering to the natural flow of the phase space by retaining steps one and three. The result will be the desired exploitation of natural solutions while capturing as many lower-priority stars as possible in between the targeted subset.

The TDM framework was designed to be adaptable to all of these scenarios. A cursory examination of this problem is provided, demonstrating the trade-off in fuel for targeting specific stars — and hence selecting alternative trajectories — instead of following the maximal path. This is performed first for a single time window, and then extended to multiple.

5.4.1 Targeting Select Stars: Single Time Window

This section analyzes the cost of targeting stars outside the maximal path. To start, only a single time window is considered — Time Window 1 (assumed a representative case). In the spirit of the maximal path formulation, the case of multiple windows — demonstrated in the following subsection — is analyzed by combining the results from individual segments.

The objective is to contextualize the performance of the maximal path — in both number of stars intercepted and fuel cost — with respect to all alternative paths. The first step in accomplishing this task is to compute all possible paths passing through the 44 target stars available in the time window. Recall, step two of the TDM formulates the time windowed star ordering problem as a search over a directed, singly-connected graph 4.2.1. To create the tree-based structure, the set of all feasible retargeting maneuvers from each star is determined and stored. In the graph, stars are represented by nodes, and edges connect parent stars to their children in the direction of feasible transition. Just as the maximal path can be determined recursively, the set of all possible paths can be determined starting at the terminal stars (graph leaves) and working backwards until all roots are reached.

In total, 1267 paths were determined to traverse the time window [17], each starting from the trace of the initial invariant circle and ending at a terminal star. The total fuel mass expended (cost) for the external occulter to complete all available paths was computed via the minimum-fuel optimization formulation formulated in section 4.3. Figure 5-16 presents a plot of each path (numbered 1 through 1267) versus total fuel mass consumed in kilograms. The paths are grouped and colored according to the total number of target stars they intercept.

Only two 6-star paths exist. The maximal path — costing 10.2 kg of fuel as documented in section 5.2 — and path 1267, costing 13.9 kg of fuel. The two paths vary only in their root node. Three star paths make up the majority of the total number. Interestingly, the path of maximum fuel consumption is a three star path

[17] One-star paths are considered trivial and omitted for non-terminal stars, since such a path starts on its corresponding QPT solutions and reaches the alignment condition ballistically.

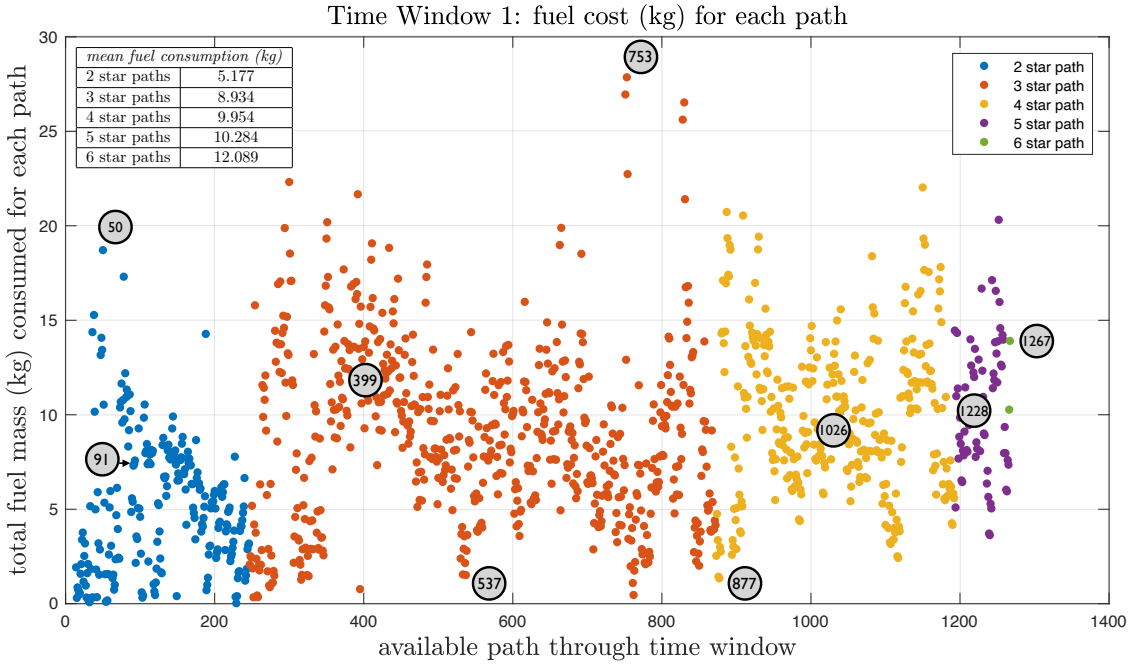

Figure 5-16: Plot of all paths available within Time Window 1 and their corresponding cost in fuel mass (kg).

Path Number	Star ID Transfer	m_{fuel} (kg)
50	$2 \rightarrow 99$	18.713
91	$66 \rightarrow 111$	7.305
399	$36 \rightarrow 14 \rightarrow 35$	12.304
537	$113 \rightarrow 14 \rightarrow 35$	1.515
753	$79 \rightarrow 45 \rightarrow 99$	27.861
877	$25 \rightarrow 26 \rightarrow 23 \rightarrow 91$	1.342
1026	$52 \rightarrow 74 \rightarrow 75 \rightarrow 32$	9.170
1228	$79 \rightarrow 55 \rightarrow 75 \rightarrow 4 \rightarrow 32$	10.257
1267	$41 \rightarrow 55 \rightarrow 74 \rightarrow 75 \rightarrow 4 \rightarrow 32$	13.908

Table 5.3: Data corresponding to sampled paths from Figure 5-16.

(consuming 27.9 kg). This is due to the fact that having less stars to intercept would allow for more time for longer, and more expensive transfers. The mean fuel consumption for each path type is provided in the top left corner of the figure. It can be seen that for the high density paths (i.e., 3, 4, and 5 star paths), the mean fuel consumption varies on the order of a kilogram or less. Even the 6 star path mean (though only two data points exist) is less than two kilograms more than that of the 5 star paths. As suspected, this seems to indicate that fuel consumption is connected more to the specific stars that are ordered and not the total number stars that are ordered in itself. Nine representative paths (close to the mean and outliers) were sampled, their location marked on the figure, and their star IDs and consumed fuel provided in Table 5.4.1.

Figure 5-17 (top) provides a stacked histogram, illustrating the total number of paths passing through each target star, by path type. The bottom figure presents a box plot outlining the mean and range of fuel consumed for the paths passing through each star. From the histogram, it can be deduced which stars are reachable by which type of path. For instance stars 51, 33, 61, and 45 — among others — cannot be reached by anything greater than a 3 star path. Thus, if one of these is prioritized, it would be at the expense of total number of intercepts. As previously discussed, this does not directly translate to low fuel consumption. While the paths

intercepting stars 33 and 61 have a mean below 5kg and maximums below 10kg, the paths passing through star 45 have a mean of approximately 15kg with a range varying from nearly 1 kg to 27 kg.

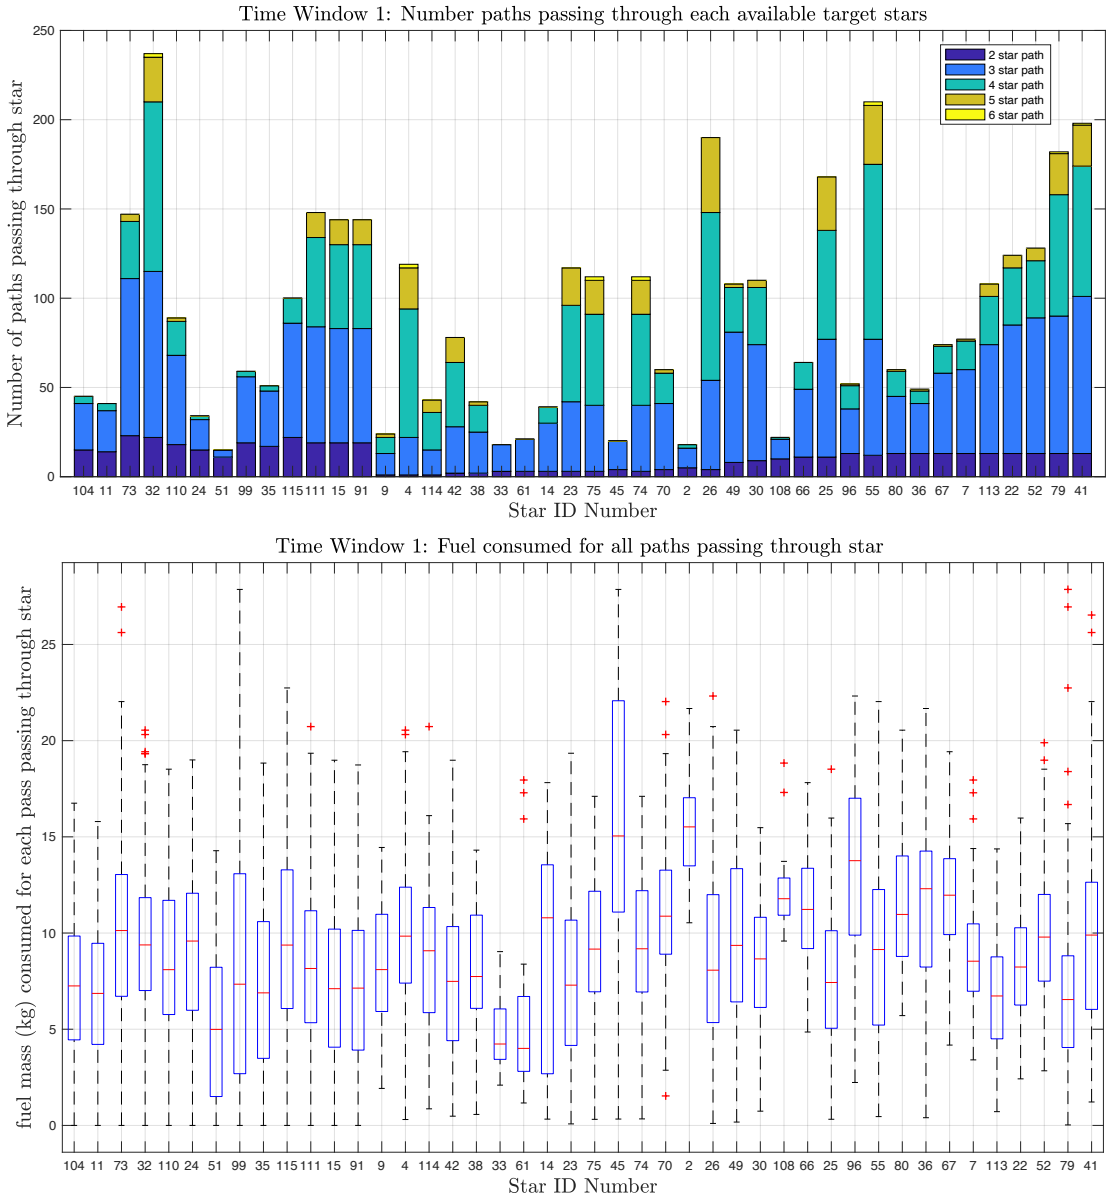

Figure 5-17: Histogram of total number of paths passing through each available star in Time Window 1 (top) and Box plot of the fuel consumed for the paths passing through each available star. (bottom)

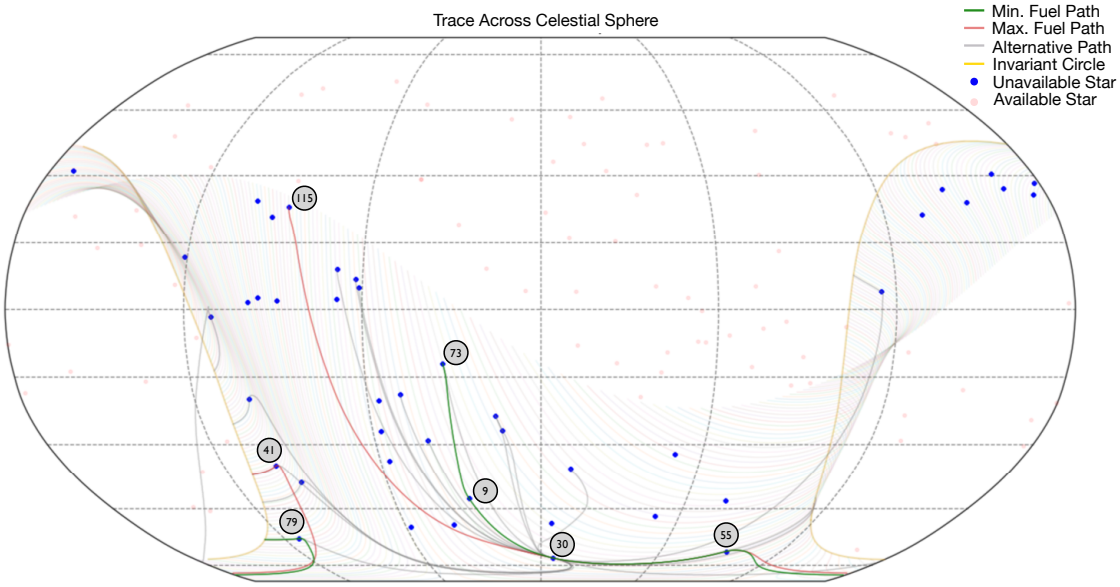

	Star ID Transfer	m_{fuel} (kg)
min. fuel path	$79 \rightarrow 55 \rightarrow 30 \rightarrow 9 \rightarrow 73$	7.769
max. fuel path	$41 \rightarrow 55 \rightarrow 30 \rightarrow 115$	15.478

Figure 5-18: Plot of all paths through TW 1 passing through star 30.

Once all possible paths through the time window are found, it is trivial to search through the set to find those passing through a specific star. Figure 5-18 illustrates the formation LOS tracings on the celestial sphere for all paths passing through star 30[18]. From that subset, the maximal path can be determined. That is, the minimum-fuel trajectory that intercepts the maximum number of stars, including star 30. This is a 5 star path costing 7.769 kg of fuel ($\Delta V = 0.191$ km/s). Thus, the trade-off for choosing to pass through star 30 is one less star intercept than the overall time window maximal path. However, approximately 2.5 kg are saved. The overall most fuel expensive path passing through star 30 (15.478 kg of fuel and 0.387 km/s of ΔV) is also provided.

The paths passing through two or more select stars can be found as the intersection of the subsets of the paths passing through each star. If the intersection is

[18]There are 110 total, many are fractions of others since theoretically, a path can start from any star's QPT solution and transit ballistically to its alignment condition.

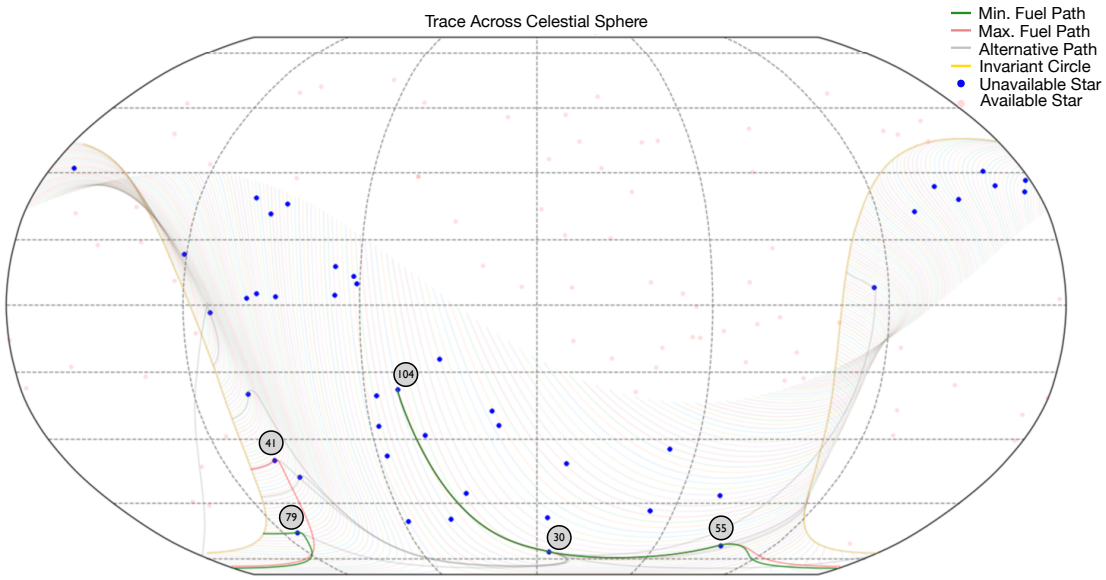

	Star ID Transfer	m_{fuel} (kg)
min. fuel path	$79 \rightarrow 55 \rightarrow 30 \rightarrow 104$	7.245
max. fuel path	$41 \rightarrow 55 \rightarrow 30 \rightarrow 104$	10.882

Figure 5-19: Plot of all paths through TW 1 passing through stars 30 & 104.

empty, an infeasible transfer exists between two or more of the stars. Thus, completion will require multiple time windows. Figure 5-19 demonstrates the set of all paths through Time Window 1 passing through both stars 30 and 104. Ten of these exist, of which the maximal path passes through four stars (including 30 and 104) and costs 7.245 kg of fuel (ΔV=0.179 km/s).

5.4.2 Targeting Select Stars: Multiple Time Windows

The analysis can be extended to include multiple time windows. In the same spirit as finding the combined maximal path across an extended time horizon, the time windows can be solved individually and their results combined. As an example, the maximal path passing through stars 30, 18, 118, and 8 was computed (Figure 5-20). These stars appear in Time Windows 1, 2, 3, and 4 respectively. Each time window was solved independently, as demonstrated in the previous subsection, and their

resulting paths joined via knowledge of feasible transitions between them[19]. A total of 6000 paths were found. Sixteen of these were found to pass through the maximum number of twelve stars. The minimum fuel path of these was calculated to cost 47.6 kg (ΔV=1.03 km/s) of fuel and the maximum 54.14 kg (ΔV=1.30 km/s). Note, this is double the cost of the overall maximal path computed in section 5.3 for Time Windows 1 – 4.

Solving the target selection problem by computing all available paths and finding their intersection quickly becomes intractable when extending the analysis beyond six time windows. A recursive scoring method that constrains the search through the stars of interest should be developed. However, this is left for future work.

[19]Recall, this was obtained in section 5.3

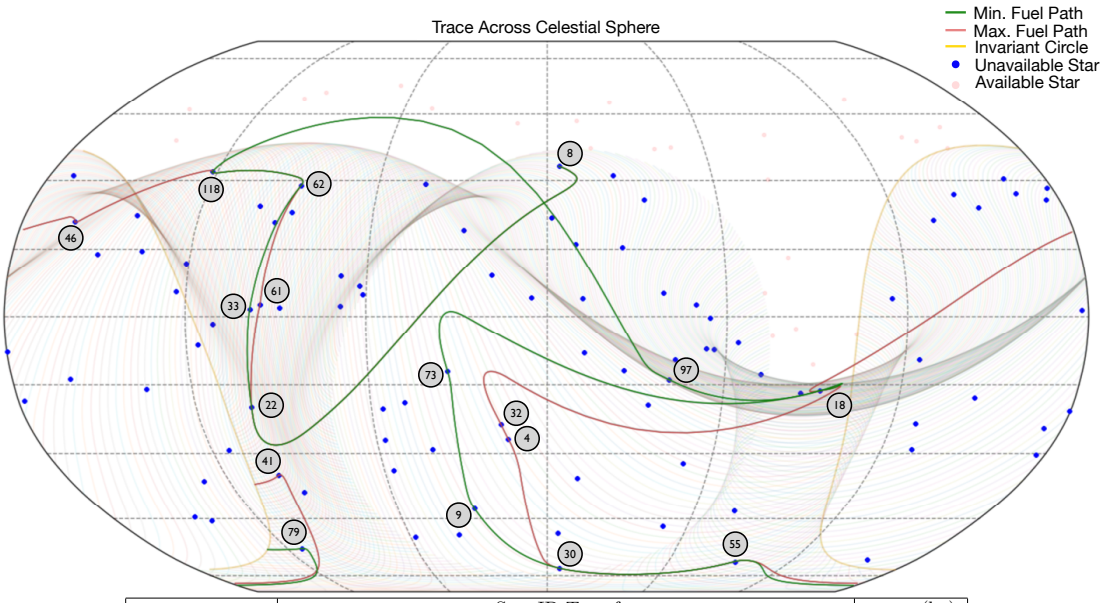

	Star ID Transfer	m_{fuel} (kg)
min. fuel path	79 → 55 → 30 → 9 → 73 → 18 → 97 → 118 → 62 → 33 → 22 → 8	47.577
max. fuel path	41 → 55 → 30 → 4 → 32 → 18 → 46 → 118 → 62 → 61 → 22 → 8	54.142

Figure 5-20: Time Windows 1-4: Plot of minimum fuel and maximum fuel maximal paths, passing through stars 30, 18, 118, and 8.

5.5 Performance Considerations

The objective of this section is to contextualize the performance of the TDM to other observatory/external occulter formation flight methodologies available in the literature. Prior to this, a brief summary of the methods employed to estimate the costs of the transfer maneuvers in both fuel mass and ΔV is provided.

Estimating Expended Fuel Mass and ΔV

The fuel mass consumed by each transfer maneuver is estimated throughout the sample mission analysis via the Tsiolkovsky Rocket Equation (5.1).

$$m_{fuel} = \int_0^{t_f} \dot{m}_{fuel}(t)dt = \int_0^{t_f} \frac{T(t)}{I_{sp}g_0}dt \approx \sum_{k=0}^{N} \frac{T_k}{I_{sp}g_0}\Delta t_k \tag{5.1}$$

where T_k is the commanded thrust at time step k of the transfer maneuver (outputted by the trajectory optimization software package), I_{sp} is the specific impulse of the spacecraft propulsion system, and g_0 is the acceleration due to gravity (on Earth). The transfer ΔV can be computed in a similar manner, as shown in Equation 5.2.

$$\Delta V \approx \sum_{k=0}^{N} \frac{|T_k|}{m_{s/c}}\Delta t_k \tag{5.2}$$

where $m_{s/c}$ is the mass of the spacecraft.

5.5.1 Contextualizing the Methodology

In this section the capabilities of the TDM framework are considered within the context of the state-of-the-art. First, it is important to acknowledge the large breadth of technical challenges that make up a fully scoped observatory/external occulter, exoplanet imaging mission. In consequence, all studies have made a series of underlying assumptions, in order to frame a research problem that is both tractable

for study within a reasonable amount of time, and impactful. Often times these assumptions only partially overlap, and more importantly, focus on different technical challenges. Thus, the objective of this section is not to simply exhibit a competing approach to designing an observatory/external occulter formation flying mission, but to provide context to how pieces — that is, independent contributions such as this thesis — can potentially fit into a greater whole in the future.

As far as the author is aware, Kolemen [58, 59] was the first to consider the retargeting observatory/external occulter formation flying problem from the context of star ordering and mission design. This analysis assumed operation in the Sun-Earth L2 environment, with dynamics governed by the CR3BP equations of motion and the observatory operating from a halo orbit. The occulter motion was not designed to be constrained or influenced by any specific class of motion[20]. The star ordering problem was solved by an exhaustive search method. That is, all combinations of the minimum-fuel optimal control problem to transfer the external occulter between the alignment conditions of all stars in the DRM were solved. The path through these transfers was then ordered and solved as a Traveling Salesman Problem. It should be noted that neither knowledge of dynamical structures, nor the natural flow of the phase space/solutions were explicitly exploited or included in the design. The author achieved 150 observations at a 50,000 km separation distance for a total of 3.6 km/s of ΔV and 80 kg of fuel, using low-thrust propulsion [59]. However, the external occulter was modeled identically to the SMART-1 spacecraft, which has a dry mass of 267 kg. This is almost a fifth of the value used in this analysis. When matching the mass and thruster constraints to those of the SMART-1 analysis, it was found that the estimated fuel consumed to traverse the maximal path for Time Window 1 dropped by a fifth as well (to 2.2 kg), extrapolating this scaling factor to the combined Time Windows 1 – 10 result suggests a comparable 22 kg of fuel consumed in about 1.2 years with 37 stars imaged. It is critical to note,

[20]The author did consider the possibility of operating from a QPT [58] but did not implement.

however that all analysis performed in this thesis exploited a torus with average separation distance of 100,000km. A fair comparison would require this value to be equal among methodologies. This is left for future work.

Soto [100–102] continued Kolemen's work, specifically tackling the complex imaging constraints of a direct-imaging, formation flying mission. The environmental assumptions are the same as Kolemen's, except that solar radiation pressure (SRP) is included in the dynamical model. However, only impulsive maneuvers via bi-propellant thrusters ($I_{sp} = 308$s) were considered. Furthermore, the separation distance was restricted to approximately 37,000 km over a three year mission. While twenty-nine observations were made, a total of 2.094 km/s of ΔV was accumulated during this time, costing 1750 kg of fuel. Again, the contribution made by Soto in handling imaging constraints should be noted (e.g., imgaging keep-out zones due to instrument blinding by celestial bodies/reflectance, integration time, life-time constraints, observation blocks due to shared instrument time, etc.) since they are critical for considering a fully-scoped mission. As discussed in section 6.2.1, integrating these considerations into the TDM should be of high-priority moving forward.

Finally, Folta et al. [31], Leitner [65], and Webster et al. [112] each examined mission scenarios with a fixed DRM — implying that star ordering was constrained. Thus, only point-to-point transferring costs for the external occulter across the sequence was considered, and a design/optimization methodology/framework was not developed. Fuel costs for cases solved using impulsive maneuvers range on the order of thousands of kilograms, and a separation distance of 72,000 km and a 1000 kg external occulter assumed[21]. However, Folta also solved several DRMs using low thrust propulsion at a separation distance of 72,000 km, reducing the total fuel consumed to a level near 200 kg. It is unclear how many stars were targeted in the study, their angular separations, timing constraints, or maximum thrust available

[21]Note: Webster [112] simulated a separation distance of 36,000 km and thus reports lower numbers.

(although an specific impulse of 4100s was used).

This thesis, having focused on designing the retargeting maneuvers of an observatory/external occulter formation flying mission, also builds on the work of Kolemen [59]. By rooting the Trajectory Design Methodology in Dynamical Systems Theory, the design process is informed by the natural flow and structure of the phase space, allowing for the efficient search of solutions aligning with this phenomena, as opposed to exhaustive search methods. As demonstrated, by understanding how these paths pass through the phase space, trajectories can quickly be adapted to included specific stars of interest. This methodology, combined with the database developed, opens an immense design space of natural solutions including continuous families of quasi-periodic tori and periodic orbits that can continue to be developed in potentially limitless ways (refer to section 6.2). However, the subject of this research would benefit greatly from combining the contributions of this thesis with Soto's work [100], since handling the imaging constraints would be the next step to further ground this retargeting methodology to real-life mission constraints. This task is left to future work.

Chapter 6

Conclusion

This thesis investigates the use of natural solutions to frame and solve the formation retargeting maneuvers of an observatory/external occulter exoplanet imaging mission. By illuminating the classes of natural motion that can be exploited, fuel costs can be minimized[1], but more importantly, the set of all available paths contextualized within the dynamical landscape. This provides a baseline from which solutions can be interpreted and mission design trade-offs analyzed. To this end, a Trajectory Design Methodology — introduced in Chapter 4 — was developed that leverages the natural periodic and quasi-periodic motion of the CR3BP phase space's center manifold to guide the spacecrafts' motion. The TDM determines the fuel-optimal[2] path that passes the formation LOS through the maximum number of stars within an extended time window. Since the framework is dynamically informed, the incremental costs of deviating from this maximal path, to achieve a specific science objective, can be readily considered. This was accomplished by rooting the Trajectory Design Method within Dynamical Systems Theory perspective — which provides the lens from which the qualitative structure of the phase space can be analyzed. Its steps

[1] As previously stated, this is ultimately dependent on the mission objective: which should be reflected in the TDM formulation. What the TDM provides is the flexibility to select, plan, contextualize, and evaluate these low fuel paths with respect to their effect on science yield.

[2] That is, under the given two-point boundary value problem with fixed final time constraints.

are summarized below.

Trajectory Design Methodology Summary

1. **Divide**: the problem into a series of finite time horizons — only search for paths up to a specified point in the future (e.g., quarter period).

2. **Find**: trajectories that pass the formation LOS through a maximum number of stars available in this time horizon: Trajectories lie on the same torus surrounding the operational halo orbit of the observatory and are uniquely defined by their initial phase on the invariant circle. The torus is defined by a constant longitudinal period equal to the halo orbit.

   i *Run a first pass search*: via a sorted tuple data structure containing the stars and initial 21 phase solutions provided by the database. Obtain the closest database entry to each star of interest.

   ii *Refine*: the initial phases by bisection and GMOS solver to converge to solutions in-between the discrete database entries.

3. **Order**: the solutions by time of intercept (equivalent to longitudinal angle) and initial latitudinal phase using a tree-based search inspired by dynamic programming. Find the maximal feasible pass through of stars with minimal $\Delta\theta_{2,0}$.

4. **Transfer**: between solutions via a minimum-fuel optimal control solver using low-thrust propulsion. Frame the problem as a sequence of two-point boundary value problems with fixed final time.

5. **Repeat**: for the next time window

The classes of natural periodic orbits and quasi-periodic tori exploited by the TDM were made readily available for search and analysis by the database cataloged

in Chapter 3. Included are descriptions of the differential corrector methods used to converge solutions in the dynamically sensitive Circular Restricted Three-Body Problem Phase space, as well as details on the design of the software infrastructure developed in the Julia language. In Chapter 5, The TDM is demonstrated on a sample mission analysis, with target stars taken from the HabEx study. A summary of the final contributions made by this thesis is provided in section 6.1. Recommendations for future work are provided in section 6.2.

6.1 Contributions

A summary of the contributions made by this thesis is provided in this section.

- A trajectory design methodology for the retargeting maneuvers of an observatory/external occulter formation flying mission operating at Sun-Earth L2, rooted in Dynamical Systems Theory, that exploits the natural motion provided by the center manifold of the Circular Restricted Three-Body Problem phase space. This includes:
  - A search method to find the set of natural trajectories on the quasi-periodic torus (for the external occulter) that naturally pass the formation line-of-sight through alignment with all reachable stars in a given time window
  - A method to determine and order the transfers between natural solutions on the torus that yields the maximum number of intercept between the formation line-of-sight and the set of available target stars within a given time window.
  - A method to connect the natural solutions on the torus — through minimum-fuel, low-thrust transfers — into a continuous trajectory, passing the formation line-of-sight through the maximum number of target

stars in a given time window

  - A method for combining the results of individual time window analyses into a global solution across an extended time horizon.

- A queryable database of three-dimensional periodic and quasi-periodic orbits existing in the CR3BP phase space. This includes a large infrastructure of solvers and tools for mission designers to further develop and interact with the database — including rapid retrieval of the set of orbits that achieve a given mission objective. The database includes halo orbits, surrounding quasi-periodic tori, butterfly orbits, near-rectilinear halo orbits, and natural connections between select family members.

6.2 Future Work and Recommendations

The Trajectory Design Methodology was designed to be a foundation — built from the onset in DST — that is both adaptable to competing mission objectives and a platform from which future work can develop. Demonstrated is a general framework, guided by the philosophy that complex trajectory design should be rooted in understanding of the underlying dynamical landscape and that natural solutions should be exploited whenever possible. However, in no way is this work considered complete. A host of nearly limitless directions can be taken to both search for global fuel-optimal solutions and understand the formation retargeting problem within the context of Dynamical Systems Theory. The recommendations for future work generally fall within three categories. The first is in stripping back assumptions and adding additional realism/complexity to the trajectory design process. The second expands on the design decisions made to construct the TDM. The final category is in expanding the scope of the trajectory design space. Several points of interest are detailed below.

6.2.1 Eliminating Assumptions and Adding Complexity

As indicated in section 5.1, several assumptions were made to scope the thesis to a tractable, yet representative, problem that is both insightful to the greater exoplanet imaging mission and scalable to increasing complexity. It is believed that the inherit adaptability of the Trajectory Design Methodology satisfies this objective. In that regard, the recommendations for future work presented are considered as additions/developments to the foundation developed in this thesis and not in any way negating the progress made.

The first set of assumptions considered pertain to the dynamical model. This thesis considered only the Circular Restricted Three-Body Problem (CR3BP) phase space for the analysis and construction of trajectories. Though qualitatively, the structure of solutions developed are expected to continue to more complex models [8], a full mission analysis will need to account for additional environmental disturbances. Chief of these will be solar radiation pressure (SRP). The external occulter in many ways resembles a solar sail. Thus, in general, additional fuel will need to be consumed to counteract its effects. Adding SRP to the CR3BP model is as straightforward as computing the expected acceleration from a model and adding it to the vector flow (equations 2.8)-2.10 [3]. A separate study on the specific impact of SRP on formation retargeting and fuel costs would be insightful, as well as considerations of how — if at all — it can be leveraged in design. Ultimately, the objective would be to move completely beyond the CR3BP phase space into a full ephemeris model. Solutions generated can be converted from the lower to higher fidelity model with the assistance of a two-level differential correction process — which breaks trajectories into multiple segments and corrects the set along a series of patch points [39,61]. Finally, a full six-degree of freedom model (position and attitude) should to be developed for each of the spacecraft, and the torque requirements for pointing (as opposed to just positional alignment) computed.

[3] This was done by Soto [39] in his observatory/external occulter formation flight study.

The second set of assumptions concern the imaging constraints of the greater mission. Only formation retargeting — i.e., realignment of the spacecraft LOS from one target star to the next — was considered in this thesis. Realistic constraints for a full mission are much more complex. The primary two that should be added to the analysis are keep-out zones (i.e., accounting for sun, earth, or other bright celestial body exclusion angles) and fixing the formation LOS during imaging. Currently, it is assumed that thruster saturation and timing alone determine the set of reachable stars within a time window. In reality, imaging constraints — such as blinding by celestial objects, including the sun, and reflected light from the external occulter — often determine the eligibility of a target. Given a model or set of logical conditions, these constraints can naturally be included during the star ordering step of the TDM. Finally, the spacecraft need to not only pass their LOS through each target star (as is currently assumed), but fix it in inertial space during imaging for the required integration time [55]. Furthermore, this needs to be accomplished while meeting precise separation and lateral displacement conditions between the spacecraft[4]. One avenue that could be considered to assist with this problem is searching for — and potentially aligning imaging at — points in the phase space where the formation have non-zero velocity in the rotating frame yet pass through zero velocity in the inertial frame. As previously noted, Soto [100] focused his analysis on handling these imaging constraints. Integrating the formulation presented in his work into the TDM of this thesis should be of high-priority moving forward.

6.2.2 Exploring Further Design Decisions

Ultimately, the TDM is a composition of design decisions, and the one presented is just a single manifestation of these possibilities. Many directions could be taken to develop a framework that still satisfies the guiding philosophy of exploiting natural solutions when possible, indicating that there is still significant room for innovation

[4]Section 6.2.3 considers dynamical structures that can assist with this problem.

in this domain. This was first considered in section 5.4, where the effects of deviating from the maximal path to target specific stars was examined. In general, competing mission objectives can be handled by augmenting the TDM, such as electing to revisiting stars of high-science-value over lower priority stars yet to be imaged[5]. To assist with these decisions, dynamic ordering and selection constraints could be created that weight multiple objectives and depend on both the present state and history of observations. A means to fully automate this process would need to be developed. A promising direction is to implement a receding time horizon approach — as is performed in the field of model predictive control — that recomputes the optimal (or set of available) path(s) after every star intercept. The method of combining multiple time windows for the star targeting case (section 5.4) should also be further developed to include a recursive star scoring method — as was developed for the maximal path formulation — so that full set of all possible paths do not need to be directly computed.

As a quick note, it was originally desired to develop a continuous controller that could directly manipulate the latitudinal frequency of solutions on the torus to skirt along its surface, as opposed to the method implemented of computing fuel-optimal control solutions via a trajectory solver. However, in general, applying control changes the energy of the spacecraft. Since the torus is a constant energy solution, this causes the spacecraft to depart, piercing through surfaces of solutions until the thrusting ceases. However, as with optimal orbital plane changes, there may be specific conditions in which thrust can be applied that do not change the spacecraft's energy. This is a potential direction that could be explored. Regardless, the design of controllers that are directly informed by the underlying dynamics should be pursued. However, if the fuel-optimal control methods are continued, future work should employ the methods of Lo and Anderson [3–5, 68] to interpret the solutions within the dynamical landscape.

[5]Note, this simply a subcase of the targeting specific star scenario analyzed in section 5.4.

A final interesting investigation to mention would be the pairing of the series of formation retargeting maneuvers with attitude control. In a future full mission scenario, where on-board thrusters are used to point the spacecraft, differential thruster firing could be used pair lateral and rotational motion achieving multiple objectives simultaneously. The effects of the converse scenario — where attitude or retargeting maneuvers disrupt or counteract each other — should also be considered.

6.2.3 Expanding the Scope of the Design Space

The final point considered is in expanding the scope of the trajectory design space. Over one thousand periodic solutions and tens of millions of QPT were computed in the SEL2 region for the orbital database. However, the current analysis was restricted to a single orbit and invariant torus. Developing the TDM to search for paths along the full database would open an immense set of new possibilities of how natural motion can be exploited for the observatory/external occulter formation flying problem. Since an infinite number of solutions are present within each family, a means to manage and guide the search would need to be determined. These directions could be explored incrementally. For example, the effects of changing the halo/QPT combination along members of the halo orbit family[6] could first be inspected. This analysis could then be followed by a study conducted by fixing the halo orbit and operating from tori of varying amplitude. Ultimately, the objective would be to gain intuition into how and when (if ever), the spacecraft should jump between solutions. A scenario where this could be particularly useful is in meeting the spacecraft separation distance required during imaging, since the toroidal families vary continuously in amplitude from their parent halo orbit[7]. Even beyond the

[6]Recall, this corresponds to varying the halo orbit amplitude and distance from the Lagrange point. It is assumed that the average distance of the QPT from the parent halo orbit is held constant.

[7]The alignment condition on one torus family member may meet the separation requirement for one target star, but not another — and vice versa.

database, there are multiple types of QPT families that were not computed (e.g., constant torus angle ρ, constant energy, varying period). It is believed that the two-dimensional torus plots[8] introduced by Baresi and Kawakatsu [10] could potentially be a powerful tool in helping interpret and compare the utility of different tori in such an investigation. In essence, there are near limitless combinations of naturally existing solutions yet to be considered, and all hold the potential for discovery of a global optimal solutions that maximally exploit natural phenomena when paired with an augmented TDM.

[8]These two-dimensional plots represent the "unraveled" torus surface, where the longitude and latitude over a revolution are plotted on the x and y axes respectively. This representation allows an investigator to project physical quantities — such as radial velocity, the location of stars on the celestial sphere, etc. — onto the torus, as opposed to the inverse case of projecting the quasi-periodic solutions onto the inertial, celestial sphere — as was done in this thesis. The advantage is that from this perspective, the dynamical flow on the torus is represented by sets of simple diagonal lines across the plane, with slope equal to ω_2/ω_1 (refer to equation 2.16). This may provide an easy way to visualize and compare the utility of different tori. However, target stars will not appear static in this visualization. Thus, if they are to be included in the plot, a means to capture their motion sweeping across the plane would need to be determined.

Appendix A

Sample Mission: Additional Time Windows

Plots from the analysis for Time Windows 5-10 are provided in this appendix. Additionally, Time Windows 1-10 can be combined via the method described in section 5.3. The maximal path is provided in Table A.1. A total of 37 target stars are intercepted within 1.227 years for an estimated 110.277 kg of fuel and 3.777 (km/s) of ΔV. To quote from section 5.3, "It should be noted that the final maximal path through the ten time windows contains several star revisits. This reflects the objective of finding the trajectory that maximizes the number of star alignments, agnostic to their identity, in the minimum amount of fuel. In reality, science — not engineering — will drive this decision. The complexities of prioritizing certain stars and specific revisits (over fuel or total number of alignments) will need to be integrated into the framework to adapt to the needs of future missions — which will in general mean selecting suboptimal paths with respect to fuel. For a preliminary look at this problem, refer to section 5.5."

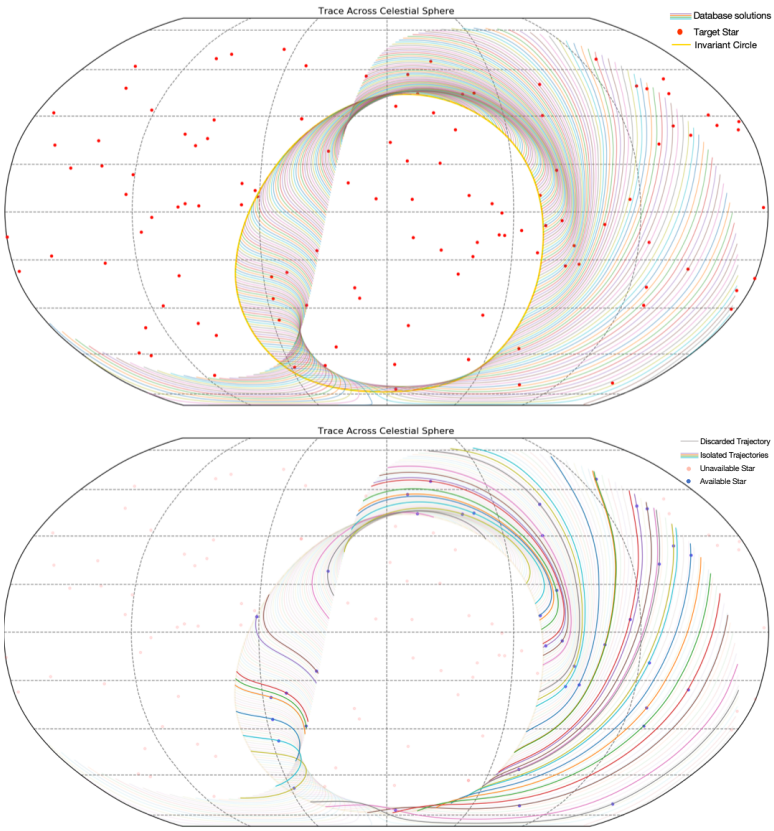

Figure A-1: Time Window 5: solutions available on the refined invariant circle (top), and the subset of converged trajectories that pass the formation LOS through each available star in the time window (bottom).

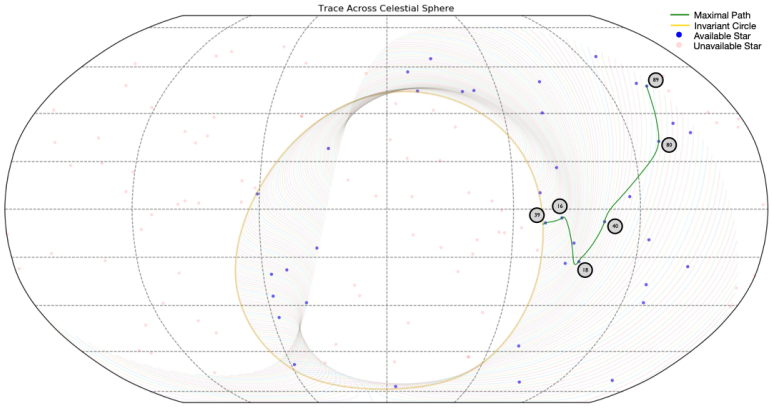

Star ID Transfer	$39 \rightarrow 16$	$16 \rightarrow 18$	$18 \rightarrow 40$	$40 \rightarrow 80$	$80 \rightarrow 89$	*Total*
m_{fuel} (kg)	0.780	5.806	2.031	3.279	1.352	13.247
ΔV (km/s)	0.027	0.199	0.070	0.112	0.046	0.454

Figure A-2: Time Window 5: Path leading to the maximum amount of star intercepts in the minimum amount of fuel.

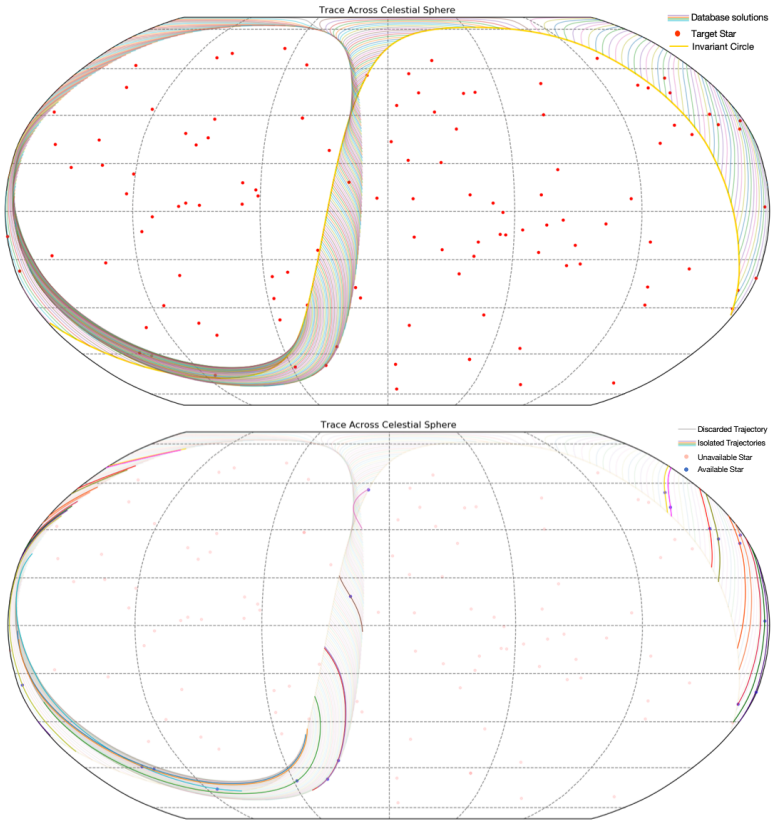

Figure A-3: Time Window 6: solutions available on the refined invariant circle (top), and the subset of converged trajectories that pass the formation LOS through each available star in the time window (bottom).

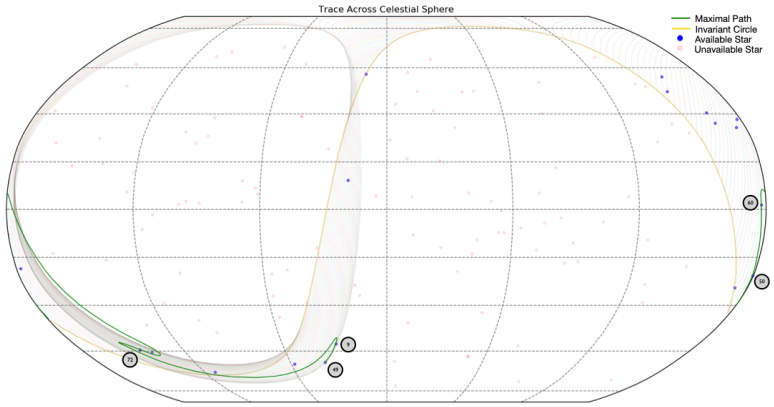

Star ID Transfer	$56 \rightarrow 60$	$60 \rightarrow 72$	$72 \rightarrow 9$	$9 \rightarrow 49$	*Total*
m_{fuel} (kg)	1.515	8.615	6.302	0.454	16.888
ΔV (km/s)	0.052	0.295	0.216	0.016	0.578

Figure A-4: Time Window 6: Path leading to the maximum amount of star intercepts in the minimum amount of fuel.

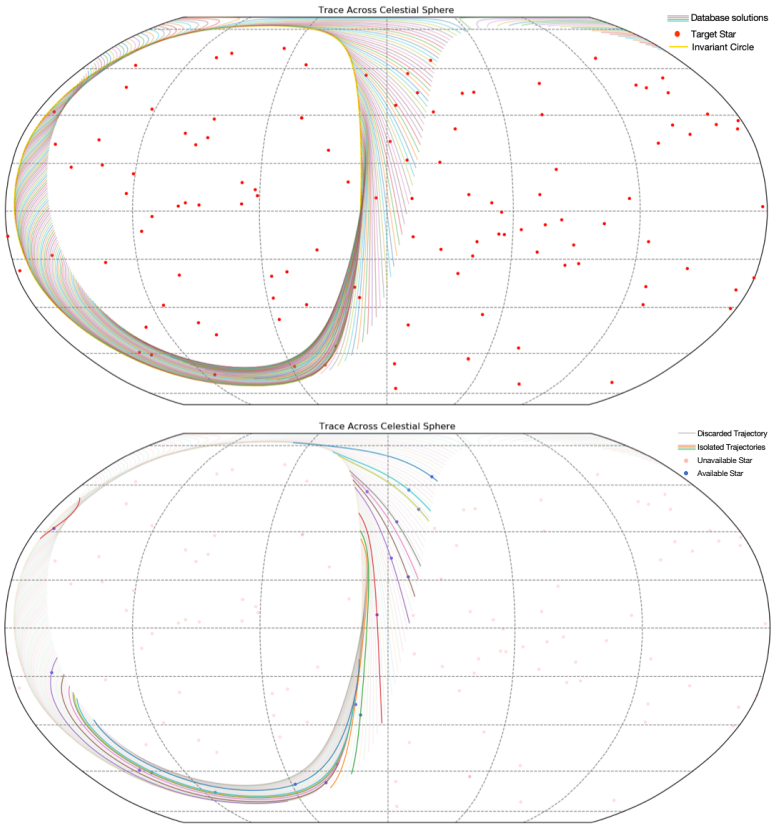

Figure A-5: Time Window 7: solutions available on the refined invariant circle (top), and the subset of converged trajectories that pass the formation LOS through each available star in the time window (bottom).

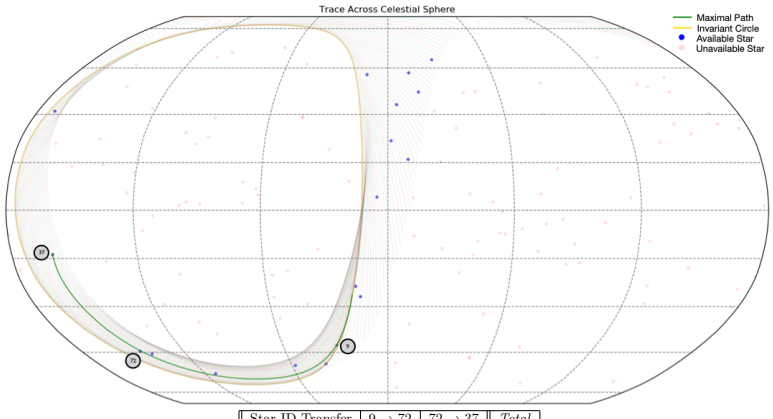

Star ID Transfer	$9 \rightarrow 72$	$72 \rightarrow 37$	*Total*
m_{fuel} (kg)	0.556	2.322	2.879
ΔV (km/s)	0.019	0.080	0.099

Figure A-6: Time Window 7: Path leading to the maximum amount of star intercepts in the minimum amount of fuel.

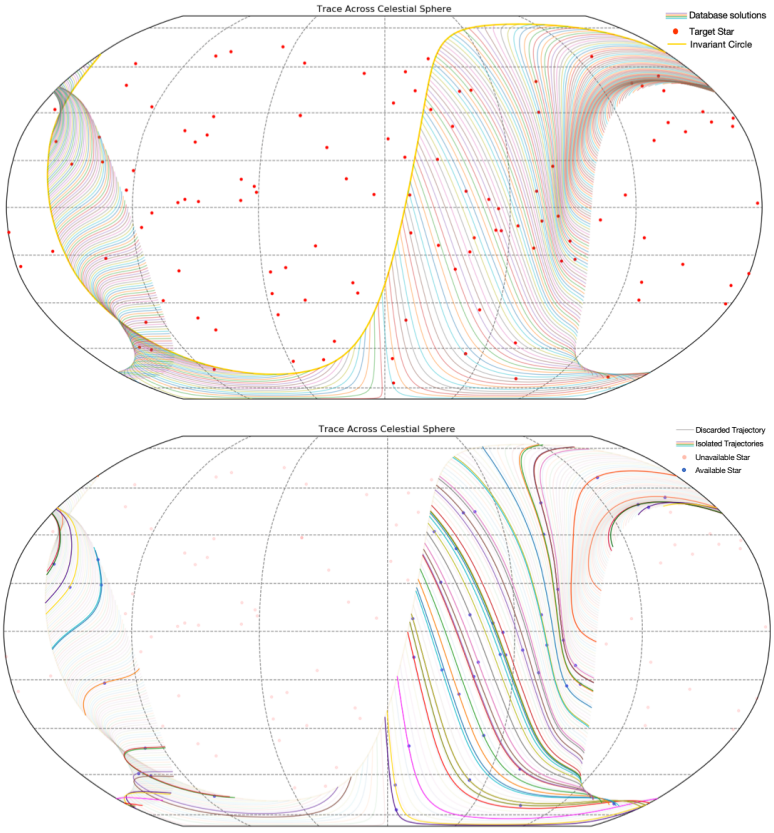

Figure A-7: Time Window 8: solutions available on the refined invariant circle (top), and the subset of converged trajectories that pass the formation LOS through each available star in the time window (bottom).

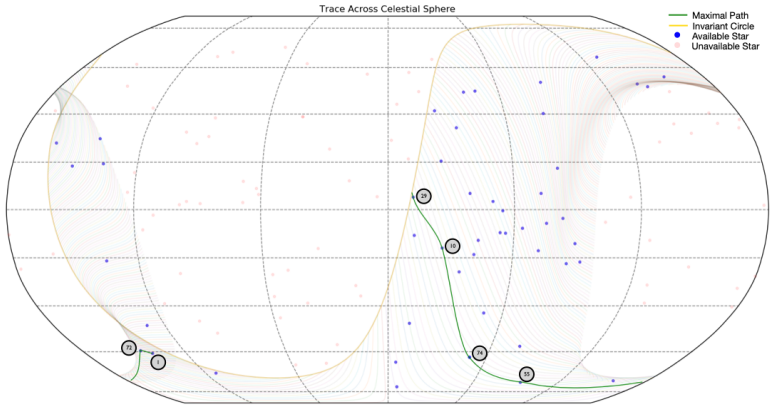

Star ID Transfer	$29 \rightarrow 10$	$10 \rightarrow 74$	$74 \rightarrow 55$	$55 \rightarrow 72$	$72 \rightarrow 1$	*Total*
m_{fuel} (kg)	2.700	1.560	0.462	3.420	0.916	9.058
ΔV (km/s)	0.092	0.053	0.158	0.117	0.031	0.310

Figure A-8: Time Window 8: Path leading to the maximum amount of star intercepts in the minimum amount of fuel.

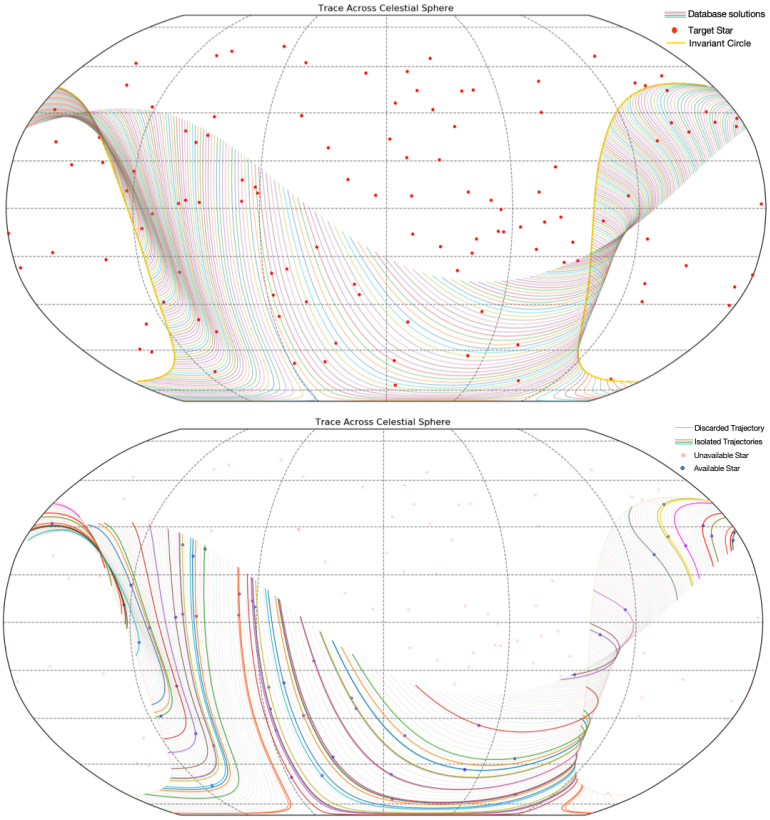

Figure A-9: Time Window 9: solutions available on the refined invariant circle (top), and the subset of converged trajectories that pass the formation LOS through each available star in the time window (bottom).

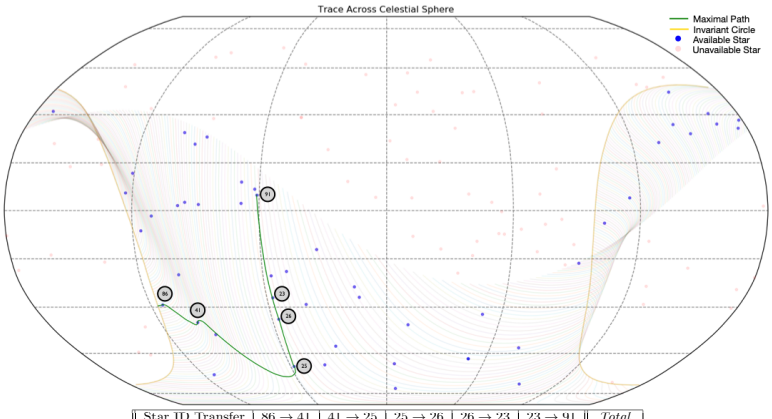

Star ID Transfer	$86 \rightarrow 41$	$41 \rightarrow 25$	$25 \rightarrow 26$	$26 \rightarrow 23$	$23 \rightarrow 91$	*Total*
m_{fuel} (kg)	2.662	6.134	1.086	0.301	0.070	10.252
ΔV (km/s)	0.091	0.210	0.037	0.010	0.002	0.351

Figure A-10: Time Window 9: Path leading to the maximum amount of star intercepts in the minimum amount of fuel.

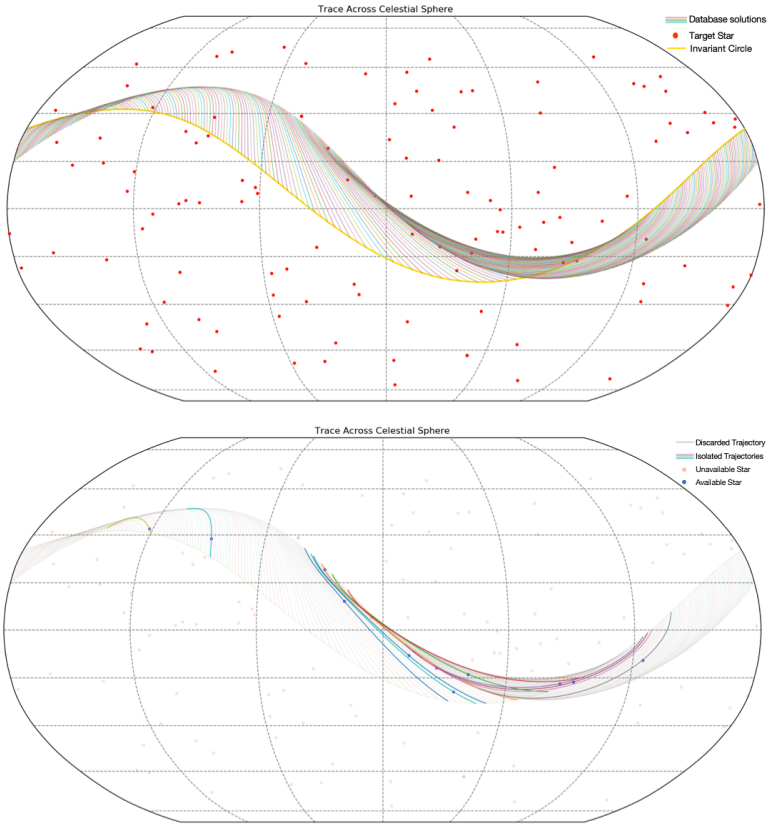

Figure A-11: Time Window 10: solutions available on the refined invariant circle (top), and the subset of converged trajectories that pass the formation LOS through each available star in the time window (bottom).

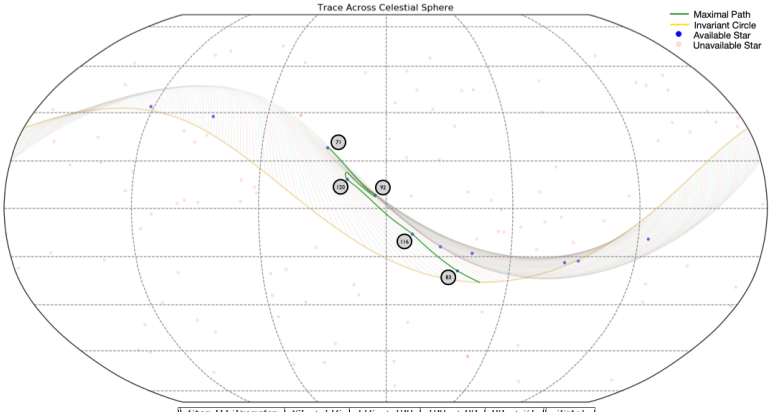

Star ID Transfer	$83 \rightarrow 116$	$116 \rightarrow 120$	$120 \rightarrow 92$	$92 \rightarrow 71$	*Total*
m_{fuel} (kg)	0.865	3.235	4.775	3.921	12.796
ΔV (km/s)	0.030	0.111	0.164	0.134	0.438

Figure A-12: Time Window 10: Path leading to the maximum amount of star intercepts in the minimum amount of fuel.

Star ID Transfer	m_{fuel} (kg)	ΔV (km/s)
$79 \rightarrow 55$	5.179	0.177
$55 \rightarrow 74$	2.557	0.088
$74 \rightarrow 75$	0.022	0.001
$75 \rightarrow 4$	2.205	0.076
$4 \rightarrow 32$	0.307	0.011
$32 \rightarrow 71$	1.780	0.061
$71 \rightarrow 5$	0.692	0.024
$5 \rightarrow 118$	3.376	0.116
$118 \rightarrow 66$	1.076	0.037
$66 \rightarrow 94$	2.694	0.092
$94 \rightarrow 86$	1.475	0.051
$86 \rightarrow 79$	0.122	0.004
$79 \rightarrow 88$	7.891	0.270
$88 \rightarrow 43$	2.096	0.072
$43 \rightarrow 18$	2.616	0.090
$18 \rightarrow 40$	1.958	0.067
$40 \rightarrow 80$	3.304	0.113
$80 \rightarrow 89$	1.330	0.046
$89 \rightarrow 72$	8.637	0.296
$72 \rightarrow 9$	6.193	0.212
$9 \rightarrow 49$	0.444	0.015
$49 \rightarrow 79$	1.938	0.066
$79 \rightarrow 29$	7.023	0.241
$29 \rightarrow 10$	2.694	0.092
$10 \rightarrow 74$	1.563	0.054
$74 \rightarrow 55$	0.438	0.015
$55 \rightarrow 72$	3.352	0.115
$72 \rightarrow 1$	0.969	0.033
$1 \rightarrow 79$	1.602	0.055
$79 \rightarrow 25$	2.847	0.098
$25 \rightarrow 26$	1.097	0.038
$26 \rightarrow 42$	1.312	0.045
$42 \rightarrow 116$	17.537	0.601
$116 \rightarrow 120$	3.257	0.112
$120 \rightarrow 92$	4.767	0.163
$92 \rightarrow 71$	3.927	0.134
Total	110.277	3.777

Table A.1: Table of the global maximal path through Time Windows 1-10 combined. Included are the list of targets stars as well as the estimated fuel mass and ΔV required to complete each transfer.

Appendix B

Table of Target Stars

Presented are the list of stars targeted in the sample mission analysis of Chapter 5. Presented are the target stars published in the HabEx study's Interim Report [32]. The same stars were chosen for analysis in the sample mission of Chapter 5. The right ascension (RA) and declination for each star were collected from the SINBAD astronomical database [114].

HIP	Star Number	RA (hr)	RA (min)	RA (sec)	DEC (deg)	DEC (min)	DEC (sec)
71683	1	14	39	36.494	-60	50	2.3737
54035	2	11	3	20.194	35	58	11.5682
16537	3	3	32	55.8449634	-9	27	29.731165
114046	4	23	5	52.03545455	-35	51	11.05875752
104214	5	21	6	53.93961007	38	44	57.89702436
104217	6	21	6	55.26406519	38	44	31.36214091
37279	7	7	39	18.1195	5	13	29.9552
1475	8	0	18	22.88496679	44	1	22.63727192
108870	9	22	3	21.6542295	-56	47	9.537018193
8102	10	1	44	4.0834226	-15	56	14.926552

HIP	Star Number	RA (hr)	RA (min)	RA (sec)	DEC (deg)	DEC (min)	DEC (sec)
105090	11	21	17	15.26885765	-38	52	2.510022611
49908	12	10	11	22.14002092	49	27	15.24915875
19849	13	4	15	16.319726	-7	39	10.338087
88601	14	18	5	27.28518	2	30	0.3558
97649	15	19	50	46.99855	8	52	5.9563
25878	16	5	31	27.39584753	-3	40	38.02155182
96100	17	19	32	21.590299	69	39	40.234737
29295	18	6	10	34.61525102	-21	51	52.65802119
45343	19	9	14	22.77545197	52	41	11.79284076
73184	20	14	57	28.00089285	-21	24	55.72689756
3821	21	0	49	6.2907161	57	48	54.67574
84478	22	17	16	13.3627375	-26	32	46.13309153
99461	23	20	11	11.93848533	-36	6	4.353559455
15510	24	3	19	55.6509352	-43	4	11.217495
99240	25	20	8	43.6094697	-66	10	55.443275
99701	26	20	13	53.39639035	-45	9	50.47346032
114622	27	23	13	16.9747821	57	10	6.076520993
12114	28	2	36	4.90238448	6	53	12.43158871
3765	29	0	48	22.97634387	5	16	50.209562
2021	30	0	25	45.07036	-77	15	15.286
7981	31	1	42	29.76349327	20	16	6.660242064
113283	32	22	56	24.0532946	-31	33	56.03506461
85295	33	17	25	45.23230437	2	6	41.1236681
22449	34	4	49	50.4109057	6	57	40.588294
86974	35	17	46	27.5266778	27	43	14.437984
61317	36	12	33	44.5448195	41	21	26.924857
64924	37	13	18	24.3142756	-18	18	40.304648
1599	38	0	20	4.2599713	-64	52	29.25476
23311	39	5	0	48.99913461	-5	45	13.22034196
32984	40	6	52	18.0504547	-5	10	25.36616519
84720	41	17	19	3.83574	-46	38	10.4467
99825	42	20	15	17.39165586	-27	1	58.7135846
27072	43	5	44	27.790894	-22	26	54.180763

HIP	Star Number	RA (hr)	RA (min)	RA (sec)	DEC (deg)	DEC (min)	DEC (sec)
17378	44	3	43	14.9008787	-9	45	48.208444
57939	45	11	52	58.76838016	37	43	7.240082865
64394	46	13	11	52.3937856	27	52	41.453553
15457	47	3	19	21.6963205	3	22	12.715139
57443	48	11	46	31.07199197	-40	30	1.279976346
105858	49	21	26	26.6048372	-65	21	58.314484
56452	50	11	34	29.48628408	-32	49	52.81989327
56997	51	11	41	3.015935829	34	12	5.882438337
81300	52	16	36	21.44929977	-2	19	28.51248573
68184	53	13	57	32.0591734	61	29	34.29935805
8362	54	1	47	44.83362505	63	51	9.00733143
29271	55	6	10	14.47411941	-74	45	10.96358713
58345	56	11	57	56.2063624	-27	42	25.36424248
13402	57	2	52	32.12818537	-12	46	10.97064946
14632	58	3	9	4.0198629	49	36	47.799638
10644	59	2	17	3.235387641	34	13	27.24344485
57757	60	11	50	41.718239	1	45	52.991019
86400	61	17	39	16.9163262	3	33	18.875718
88972	62	18	9	37.41628109	38	27	27.99592156
3093	63	0	39	21.8055114	21	15	1.716052732
12777	64	2	44	11.987042	49	13	42.41112
42808	65	8	43	18.03040054	-38	52	56.5700256
78072	66	15	56	27.1826948	15	39	41.8205
47080	67	9	35	39.50180534	35	48	36.48407023
67927	68	13	54	41.07892	18	23	51.7946
72848	69	14	53	23.7667403	19	9	10.081308
23693	70	5	5	30.65617823	-57	28	21.7289318
109176	71	22	7	0.662057274	25	20	42.3761385
69972	72	14	19	4.834136742	-59	22	44.53502748
107556	73	21	47	2.4442395	-16	7	38.233507
15330	74	3	17	46.16326057	-62	34	31.15424748
15371	75	3	18	12.81854126	-62	30	22.91730028
77257	76	15	46	26.61442914	7	21	11.04164744

HIP	Star Number	RA (hr)	RA (min)	RA (sec)	DEC (deg)	DEC (min)	DEC (sec)
41926	77	8	32	51.49583115	-31	30	3.062906467
26779	78	5	41	20.33572837	53	28	51.81062985
80686	79	16	28	28.13964655	-70	5	3.822057572
43587	80	8	52	35.81132821	28	19	50.95690137
40693	81	8	18	23.94696925	-12	37	55.81020257
24813	82	5	19	8.475463075	40	5	56.58964388
10798	83	2	18	58.50469439	-25	56	44.47347686
58576	84	12	0	44.46115932	-10	26	46.05502359
85235	85	17	25	0.098270858	67	18	24.15014156
80337	86	16	24	1.290597026	-39	11	34.73461124
51459	87	10	30	37.58036785	55	58	49.93641669
22263	88	4	47	36.29176076	-16	56	4.041927355
46853	89	9	32	51.4338939	51	40	38.281102
7513	90	1	36	47.84216	41	24	19.6443
98036	91	19	55	18.792563	6	24	24.342501
116771	92	23	39	57.0413764	5	37	34.647529
544	93	0	6	36.78409436	29	1	17.41039011
79672	94	16	15	37.27037212	-8	22	9.981989277
16852	95	3	36	52.14480225	0	23	58.53691774
53721	96	10	59	27.97386449	40	25	48.92238892
12843	97	2	45	6.203440566	-18	34	21.47862997
102422	98	20	45	17.375554	61	50	19.616737
84862	99	17	20	39.56753951	32	28	3.877348066
25278	100	5	24	25.46297733	17	23	0.729080928
42438	101	8	39	11.70432653	65	1	15.26827319
70497	102	14	25	11.7970287	51	51	2.676894
75181	103	15	21	48.15111711	-48	19	3.462483973
102485	104	20	46	5.7326265	-25	16	15.231155
28103	105	5	56	24.2930042	-14	10	3.718884
59199	106	12	8	24.8165241	-24	43	43.950354
47592	107	9	42	14.41652543	-23	54	56.05407907
49081	108	10	1	0.656779927	31	55	25.2168485
5262	109	1	7	13.88427484	49	33	19.66550726
3583	110	0	45	45.59296562	-47	33	7.146506536

HIP	Star Number	RA (hr)	RA (min)	RA (sec)	DEC (deg)	DEC (min)	DEC (sec)
95447	111	19	24	58.20024068	11	56	39.88231262
82860	112	16	56	1.6892483	65	8	5.263139
86796	113	17	44	8.703634228	-51	50	2.591049123
95501	114	19	25	29.90139	3	6	53.2061
88745	115	18	7	1.53971	30	33	43.6896
3909	116	0	50	7.588589216	-10	38	39.58481106
71284	117	14	34	40.81718085	29	44	42.46368598
77760	118	15	52	40.54113721	42	27	5.451105113
50954	119	10	24	23.7059714	-74	1	53.803578
112447	120	22	46	41.5811758	12	10	22.385447

Appendix C

Bibliography

[1] RL Akeson, X Chen, D Ciardi, M Crane, J Good, M Harbut, E Jackson, SR Kane, AC Laity, S Leifer, et al. The nasa exoplanet archive: data and tools for exoplanet research. *Publications of the Astronomical Society of the Pacific*, 125(930):989, 2013.

[2] Kyle T Alfriend, Srinivas R Vadali, and Hanspeter Schaub. Formation flying satellites: Control by an astrodynamicist. *Celestial Mechanics and Dynamical Astronomy*, 81(1-2):57–62, 2001.

[3] Rodney Anderson and Martin Lo. The role of invariant manifolds in low thrust trajectory design (part i). In *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, page 288, 2004.

[4] Rodney Anderson and Martin Lo. The role of invariant manifolds in low thrust trajectory design (part ii). In *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, page 5305, 2004.

[5] Rodney L Anderson and Martin W Lo. Role of invariant manifolds in low-thrust trajectory design. *Journal of guidance, control, and dynamics*, 32(6):1921–1930, 2009.

[6] Brian Barden. *Using Stable Manifolds to Generate Transfers in the Circular Restricted Problem of Three Bodies*. PhD thesis, Purdue University, 1994.

[7] Brian Barden. *Applications of Dynamical Systems Theory in Mission Design and Conceptual Development for Libration Point Missions*. PhD thesis, Purdue University, 2000.

[8] Brian T Barden and Kathleen C Howell. Fundamental motions near collinear libration points and their transitions. *Journal of the Astronautical Sciences*, 46(4):361–378, 1998.

[9] BT Barden. Formation flying in the vicinity of libration point orbits spaceflight mechanics 1998. In *Proceedings of the AAS/AIAA Space Flight Mechanics Meeting, Monterey, 1998*, pages 969–988, 1998.

[10] N Baresi and Y Kawakatsu. Quasi-periodic motion around phobos: Applications to the martian moons exploration (mmx). In *International Symposium on Space Technology and Science (ISTS), Fukui, Japan*, pages 1–8, 2019.

[11] Nicola Baresi. Spacecraft formation flight on quasi-periodic invariant tori. 2017.

[12] Nicola Baresi. *Spacecraft Formation Flight on Quasi-Periodic Invariant Tori*. PhD thesis, University of Colorado, Boulder, 2017.

[13] Nicola Baresi, Zubin P Olikara, and Daniel J Scheeres. Fully numerical methods for continuing families of quasi-periodic invariant tori in astrodynamics. *The Journal of the Astronautical Sciences*, 65(2):157–182, 2018.

[14] Nicola Baresi and Daniel J Scheeres. Bounded relative motion under zonal harmonics perturbations. *Celestial Mechanics and Dynamical Astronomy*, 127(4):527–548, 2017.

[15] Jeff Bezanson, Alan Edelman, Stefan Karpinski, and Viral B Shah. Julia: A fresh approach to numerical computing. *SIAM review*, 59(1):65–98, 2017.

[16] William J Borucki, David Koch, Gibor Basri, Natalie Batalha, Timothy Brown, Douglas Caldwell, John Caldwell, Jørgen Christensen-Dalsgaard, William D Cochran, Edna DeVore, et al. Kepler planet-detection mission: introduction and first results. *Science*, 327(5968):977–980, 2010.

[17] William J Borucki, David G Koch, Gibor Basri, Natalie Batalha, Timothy M Brown, Stephen T Bryson, Douglas Caldwell, Jørgen Christensen-Dalsgaard, William D Cochran, Edna DeVore, et al. Characteristics of planetary candidates observed by kepler. ii. analysis of the first four months of data. *The Astrophysical Journal*, 736(1):19, 2011.

[18] Brendan Bowler. Imaging extrasolar giant planets. *Publications of the Astronomical Society of the Pacific*, 128(968), 2016.

[19] G Chauvin. Two decades of exoplanetary science with adaptive optics. *Adaptive Optics Systems VI*, 10703, 2018.

[20] G. Chauvin, A.M. Lagrange, C. Dumas, B. Zuckerman, D. Mouillet, I. Song, J.L. Beuzit, and P Lowrance. A giant planet candidate near a young brown dwarf-direct vlt/naco observations using ir wavefront sensing. *Astronomy & Astrophysics*, 425(2), 2004.

[21] C.C. Conley. Low energy transit orbits in the restricted three-body problem. *SIAM Journal of Applied Mathematics*, 16(4), 1968.

[22] Bruce A Conway and John E Prussing. *Orbital Mechanics*. Oxford University Press, 1993.

[23] Brendan P Crill and Nicholas Siegler. Space technology for directly imaging and characterizing exo-earths. In *UV/Optical/IR Space Telescopes and Instruments: Innovative Technologies and Concepts VIII*, volume 10398, page 103980H. International Society for Optics and Photonics, 2017.

[24] Camilla Danielski et al. Atmospheric characterization of directly imaged exoplanets with jwst/miri. *The Astronomical Journal*, 156(6):276, 2018.

[25] V Domingo, B Fleck, and Arthur I Poland. The soho mission: an overview. *Solar Physics*, 162(1-2):1–37, 1995.

[26] Helge Eichhorn, Juan Luis Cano, Frazer McLean, and Reiner Anderl. A comparative study of programming languages for next-generation astrodynamics systems. *CEAS Space Journal*, 10(1):115–123, 2018.

[27] AJ Eldorado et al. Numerically optimized coronagraph designs for the habitable exoplanet imaging mission (habex). *In Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, 10968, 2018.

[28] Leonhard Euler. De motu rectilineo trium corporum se mutuo attrahentium. *Novi commentarii academiae scientiarum Petropolitanae*, pages 144–151, 1767.

[29] Robert Farquhar. The control and use of libration-point satellites. Technical report, NASA TR R-346, 1970.

[30] Thibault Flinois, Michael Bottom, Stefan Martin, Daniel Scharf, et al. S5: Starshade technology to trl5 milestone 4 final report: Lateral formation sensing and control. Technical report, Jet Propulsion Laboratory, California Institute of Technology, 2018.

[31] David Folta. Formation flying of a telescope/occulter system with large separations in an l2 libration orbit. 2008.

[32] B Scott Gaudi, Sara Seager, Bertrand Mennesson, Alina Kiessling, Keith Warfield, Gary Kuan, Kerri Cahoy, John T Clarke, Shawn Domagal-Goldman, Lee Feinberg, et al. The habitable exoplanet observatory (habex) mission concept study interim report. *arXiv preprint arXiv:1809.09674*, 2018.

[33] B. Scott Gaudi, Sara Seager, Bertrand Mennesson, Alina Kiessling, Keith R. Warfield, Kerri Cahoy, John T. Clarke, Shawn Domagal-Goldman, Lee Feinberg, Olivier Guyon, N. Jeremy Kasdin, Dimitri Mawet, Peter Plavchan, Tyler Robinson, Leslie Rogers, Paul Scowen, Rachel Somerville, Karl Stapelfeldt, Chris Stark, Daniel Stern, Margaret Turnbull, et al. The habitable exoplanet observatory. *Nature Astronomy*, 2(8):600–604, 2018.

[34] G Gómez and Mondelo J. The dynamics around the collinear equilibrium points of the rtbp. *Physica D: Nonlinear Phenomena*, 157(4):282–321, 2001.

[35] G Gómez, A Jorba, J Masdemont, and C Simó. Final report: Study refinement of semi-analytical halo orbit theory. *ESOC Contract Report, Technical Report*, 8625:89, 1991.

[36] G Gómez, A Jorba, J Masdemont, and C Simó. Study refinement of semi-analytical halo orbit theory. *Final Report, ESOC Contract*, (8625/89), 1991.

[37] G Gómez, M Marcote, JJ Masdemont, and JM Mondelo. Zero relative radial acceleration cones and controlled motions suitable for formation flying. *Journal of the Astronautical Sciences*, 53(4):413–432, 2005.

[38] Gerard Gómez, A Jorba, Jt Masdemont, and C Simó. Study of the transfer from the earth to a halo orbit around the equilibrium pointl 1. *Celestial Mechanics and Dynamical Astronomy*, 56(4):541–562, 1993.

[39] D Grebow. Generating periodic orbits in the circular restricted three-body problem with applications to lunar south pole coverage. *MSAA Thesis, School of Aeronautics and Astronautics, Purdue University*, 2006.

[40] Daniel J Grebow, Martin T Ozimek, Kathleen C Howell, and David C Folta. Multibody orbit architectures for lunar south pole coverage. *Journal of Spacecraft and Rockets*, 45(2):344–358, 2008.

[41] Nicholas Hamilton, David Folta, and Russell Carpenter. Formation flying satellite control around the l2 sun-earth libration point. In *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, page 4528, 2002.

[42] Anthony Harness, Steve Warwick, Ann Shipley, and Webster Cash. Ground-based testing and demonstrations of starshades, 2016.

[43] M Hénon. Vertical stability of periodic orbits in the restricted problem. i. equal masses. *Astronomy and Astrophysics*, 28:415, 1973.

[44] Damennick B Henry and Daniel Scheeres. Generalized spacecraft formation design through exploitation of quasi-periodic tori families. In *AIAA Scitech 2020 Forum*, page 0950, 2020.

[45] George William Hill. On the part of the motion of the lunar perigee which is a function of the mean motions of the sun and moon. *Acta mathematica*, 8(1):1–36, 1886.

[46] Sasha Hinkley. Adaptive optics observations of exoplanets, brown dwarfs, and binary stars. *Proceedings of the International Astronomical Union*, 7(S282):181–188, 2011.

[47] Kathleen Howell, Brian Barden, and Martin Lo. Application of dynamical systems theory to trajectory design for a libration point mission. *Journal of the Astronautical Sciences*, 45(2), 1997.

[48] Kathleen C Howell and Henry J Pernicka. Numerical determination of lissajous trajectories in the restricted three-body problem. *Celestial Mechanics*, 41(1-4):107–124, 1987.

[49] KC Howell and ET Campbell. Three-dimensional periodic solutions that bifurcate from halo families in the circular restricted three-body problem (aas 99-161). *ADVANCES IN THE ASTRONAUTICAL SCIENCES*, 102:891–910, 1999.

[50] FY Hsiao and DJ Scheeres. Design of spacecraft formation orbits relative to a stabilized trajectory. *Journal of Guidance, Control, and Dynamics*, 28(4):782–794, 2005.

[51] Carl GJ Jacobi. Sur le mouvement d'un point et sur un cas particulier du probleme des trois corps. *Compt. Rend*, 3:59–61, 1836.

[52] Angel Jorba and Josep Masdemont. Dynamics in the center manifold of the collinear points of the restricted three body problem. *Physica D: Nonlinear Phenomena*, 132(1-2):189–213, 1999.

[53] Seungyun Jung and Youdan Kim. Formation flying along halo orbit using switching hamiltonian structure-preserving control. In *7th European Conference for Aeronautics and Aerospace Science*, pages 2017–39, 2017.

[54] Paul Kalas, James Graham, Eugene Chiang, Michael Fitzgerald, Mark Clampin, Edwin Kite, Karl Stapelfeldt, Christian Marois, and John Krist. Optical images of an exosolar planet 25 light-years from earth. *Science*, 322(5906):1345–1348, 2008.

[55] Dean R Keithly, Dmitry Savransky, Daniel Garrett, Christian Delacroix, and Gabriel Soto. Optimal scheduling of exoplanet direct imaging single-visit observations of a blind search survey. *Journal of Astronomical Telescopes, Instruments, and Systems*, 6(2):027001, 2020.

[56] D.E. Kirk. *Optimal Control Theory: An Introduction*. Dover Books on Electrical Engineering Series. Dover Publications, 2004.

[57] David G Koch, William J Borucki, Gibor Basri, Natalie M Batalha, Timothy M Brown, Douglas Caldwell, Jørgen Christensen-Dalsgaard, William D Cochran, Edna DeVore, Edward W Dunham, et al. Kepler mission design, realized photometric performance, and early science. *The Astrophysical Journal Letters*, 713(2):L79, 2010.

[58] Egemen Koleman. *Optimal Configuration of a Planet-Finding Mission Consisting of a Telescope and a Constellation of Occulters*. PhD thesis, Princeton University, 2008.

[59] Egemen Kolemen and N Jeremy Kasdin. Optimization of an occulter-based extrasolar-planet-imaging mission. *Journal of guidance, control, and dynamics*, 35(1):172–185, 2012.

[60] Egemen Kolemen, N Jeremy Kasdin, and Pini Gurfil. Multiple poincaré sections method for finding the quasiperiodic orbits of the restricted three body problem. *Celestial Mechanics and Dynamical Astronomy*, 112(1):47–74, 2012.

[61] 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, Pasadena, CA, USA*, 2006.

[62] S Lacour, M Nowak, J Wang, O Pfuhl, F Eisenhauer, R Abuter, A Amorim, N Anugu, M Benisty, JP Berger, et al. First direct detection of an exoplanet by optical interferometry-astrometry and k-band spectroscopy of hr 8799 e. *Astronomy & Astrophysics*, 623:L11, 2019.

[63] Joseph-Louis Lagrange. Essai sur le probleme des trois corps. *Prix de l'académie royale des Sciences de paris*, 9:292, 1772.

[64] Jesse Leitner. Formation flying system design for a planet-finding telescope-occulter system. In *UV/Optical/IR Space Telescopes: Innovative Technologies and Concepts III*, volume 6687, page 66871D. International Society for Optics and Photonics, 2007.

[65] Jesse Leitner. Formation flying system design for a planet-finding telescope-occulter system. volume 6687, pages 6687 – 6687 – 10, 2007.

[66] Jaume Llibre, Regina Martinez, and Carles Simo. Transversality of the invariant manifolds associated to the lyapunov family of periodic orbits near l2 in the restricted three-body problem. *Journal of Differential Equations*, 58:104–156, 1985.

[67] Martin Lo, Bobby Williams, Williard Bollman, Dongsuk Han, Yungsun Hahn, Julia Bell, Edward Hirst, Robert Corwin, Philip Hong, and Kathleen Howell. Genesis mission design. In *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, page 4468, 1998.

[68] Martin W Lo, Rodney L Anderson, Try Lam, and Greg Whiffen. The role of invariant manifolds in lowthrust trajectory design (part iii). page 190, 2006.

[69] Martin W Lo and Shane D Ross. Low-energy interplanetary transfers using lagrangian points. *NASA Tech Brief*, 23, 1999.

[70] MW Lo and KC Howell. Trajectory design using a dynamical systems approach with application to genesis. 1997.

[71] Bernard Lyot. The study of the solar corona and prominences without eclipses (george darwin lecture, 1939). *Monthly Notices of the Royal Astronomical Society*, 99:580, 1939.

[72] Christian Marois, Bruce Macintosh, Travis Barman, B Zuckerman, Inseok Song, Jennifer Patience, David Lafrenière, and René Doyon. Direct imaging of multiple planets orbiting the star hr 8799. *Science*, 322(5906):1348–1352, 2008.

[73] Christian Marois, B Zuckerman, Quinn M Konopacky, Bruce Macintosh, and Travis Barman. Images of a fourth planet orbiting hr 8799. *Nature*, 468(7327):1080, 2010.

[74] Michel Mayor and Didier Queloz. A jupiter-mass companion to a solar-type star. *Nature*, 378(6555):355, 1995.

[75] Brian P McCarthy. *Characterization of Quasi-Periodic Orbits for Applications in the Sun-Earth and Earth-Moon Systems*. PhD thesis, Purdue University Graduate School, 2019.

[76] Lindsay D Millard and Kathleen C Howell. Optimal reconfiguration maneuvers for spacecraft imaging arrays in multi-body regimes. *Acta Astronautica*, 63(11-12):1283–1298, 2008.

[77] Jürgen Moser and Walter T Kyner. *Lectures on Hamiltonian systems*. Number 81. American Mathematical Soc., 1968.

[78] Carl D Murray and Stanley F Dermott. *Solar system dynamics*. Cambridge university press, 1999.

[79] Erc Nielson et al. The gemini planet imager exoplanet survey: Giant planet and brown dwarf demographics from 10-100 au. *arXiv preprint arXiv:1904.05358*, 2019.

[80] KW Ogilvie and MD Desch. The wind spacecraft and its early scientific results. *Advances in Space Research*, 20(4-5):559–568, 1997.

[81] Zubin P Olikara and Daniel J Scheeres. Numerical method for computing quasi-periodic orbits and their stability in the restricted three-body problem. *Advances in the Astronautical Sciences*, 145(911-930), 2012.

[82] KE Papadakis and CG Zagouras. Bifurcation points and intersections of families of periodic orbits in the three-dimensional restricted three-body problem. *Astrophysics and space science*, 199(2):241–256, 1993.

[83] Michael Patterson, John Foster, Thomas Haag, Vincent Rawlin, George Soulas, and Robert Roman. Next: Nasa's evolutionary xenon thruster. In *38th AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, page 3832, 2002.

[84] Michael A Patterson and Anil V Rao. Gpops-ii: A matlab software for solving multiple-phase optimal control problems using hp-adaptive gaussian quadrature collocation methods and sparse nonlinear programming. *ACM Transactions on Mathematical Software (TOMS)*, 41(1):1–37, 2014.

[85] Henri Poincaré and R Magini. Les méthodes nouvelles de la mécanique céleste. *Il Nuovo Cimento (1895-1900)*, 10(1):128–130, 1899.

[86] Christopher Rackauckas and Qing Nie. Differentialequations. jl–a performant and feature-rich ecosystem for solving differential equations in julia. *Journal of Open Research Software*, 5(1), 2017.

[87] David Redding, Keith Coste, Otto Polanco, Claudia Pineda, Kevin Hurd, Howard Tseng, Jose Quezada, Stefan Martin, Joel Nissen, Kevin Schulz, et al. Habex lite: a starshade-only habitable exoplanet imager alternative. In *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, volume 10698, page 106980X. International Society for Optics and Photonics, 2018.

[88] Ricardo L Restrepo and Ryan P Russell. A database of planar axisymmetric periodic orbits for the solar system. *Celestial Mechanics and Dynamical Astronomy*, 130(7):49, 2018.

[89] Albert Reuther, Jeremy Kepner, Chansup Byun, Siddharth Samsi, William Arcand, David Bestor, Bill Bergeron, Vijay Gadepally, Michael Houle,

Matthew Hubbell, et al. Interactive supercomputing on 40,000 cores for machine learning and data analysis. In *2018 IEEE High Performance extreme Computing Conference (HPEC)*, pages 1–6. IEEE, 2018.

[90] David Richardson. Halo orbit formualtion for the isee-3 mission. *Journal of Guidance and Controll*, 3(6), 1980.

[91] George R Ricker, Joshua N Winn, Roland Vanderspek, David W Latham, Gáspár Á Bakos, Jacob L Bean, Zachory K Berta-Thompson, Timothy M Brown, Lars Buchhave, Nathaniel R Butler, et al. Transiting exoplanet survey satellite. *Journal of Astronomical Telescopes, Instruments, and Systems*, 1(1):014003, 2014.

[92] Rodolfo Rosales. *Various Lecture Notes for 18385*. MIT Department of Mathematics, September 2012.

[93] Shane D Ross. The interplanetary transport network: Some mathematical sophisstication allows spacecraft to be maneuvered over large distances using little or no fuel. *American Scientist*, 94(3):230–237, 2006.

[94] Daniel P Scharf, Stefan R Martin, Carl Christian Liebe, Zahidul H Rahman, Carl R Seubert, Martin Charles Noecker, and George H Purcell. Precision formation flying at megameter separations for exoplanet characterization. *Acta Astronautica*, 123:420–434, 2016.

[95] Hanspeter Schaub, Srinivas R Vadali, John L Junkins, and Kyle T Alfriend. Spacecraft formation flying control using mean orbit elements. *Journal of the Astronautical Sciences*, 48(1):69–87, 2000.

[96] Daniel J Scheeres, F-Y Hsiao, and NX Vinh. Stabilizing motion relative to an unstable orbit: applications to spacecraft formation flight. *Journal of Guidance, Control, and Dynamics*, 26(1):62–73, 2003.

[97] Sara Seager, Margaret Turnbull, William Sparks, Mark Thomson, Stuart B Shaklan, Aki Roberge, Marc Kuchner, N Jeremy Kasdin, Shawn Domagal-Goldman, Webster Cash, et al. The exo-s probe class starshade mission. In *Techniques and Instrumentation for Detection of Exoplanets VII*, volume 9605, page 96050W. International Society for Optics and Photonics, 2015.

[98] Andrew J Skemer, Mark S Marley, Philip M Hinz, Katie M Morzinski, Michael F Skrutskie, Jarron M Leisenring, Laird M Close, Didier Saumon, Vanessa P Bailey, Runa Briguglio, et al. Directly imaged lt transition exoplanets in the mid-infrared. *The Astrophysical Journal*, 792(1):17, 2014.

[99] Daniel Smith, Steven Warwick, Tiffany Glassman, Michael Novicki, Megan Richards, Anthony Harness, and Keith Patterson. Measurements of high-contrast starshade performance in the field, 2016.

[100] Gabriel Soto, Dean Keithly, Daniel Garrett, Christian Delacroix, and Dmitry Savransky. Optimal starshade observation scheduling. In *Space Telescopes and Instrumentation 2018: Optical, Infrared, and Millimeter Wave*, volume 10698, page 106984M. International Society for Optics and Photonics, 2018.

[101] Gabriel Soto, Amlan Sinha, Dmitry Savransky, Christian Delacroix, and Daniel Garrett. Starshade orbital maneuver study for wfirst. In *Techniques and Instrumentation for Detection of Exoplanets VIII*, volume 10400, page 104001U. International Society for Optics and Photonics, 2017.

[102] Gabriel J Soto, Dmitry Savransky, Daniel Garrett, and Christian Delacroix. Parameterizing the search space of starshade fuel costs for optimal observation schedules. *Journal of Guidance, Control, and Dynamics*, 42(12):2671–2676, 2019.

[103] James S Sovey, Vincent K Rawlin, and Michael J Patterson. Ion propulsion development projects in us: Space electric rocket test i to deep space 1. *Journal of Propulsion and Power*, 17(3):517–526, 2001.

[104] Lyman Spitzer. The beginnings and future of space astronomy. *American Scientist*, 50(3):473–484, 1962.

[105] Karl Stapelfeldt. Extrasolar planets and star formation: science opportunities for future elts. *Proceedings of the International Astronomical Union*, 1(S232):149–158, 2005.

[106] Karl Stapelfeldt. Exo-c: Imaging nearby worlds coronagraph probe mission study: Final report. 2015.

[107] Edward C Stone, AM Frandsen, RA Mewaldt, ER Christian, D Margolies, JF Ormes, and F Snow. The advanced composition explorer. *Space Science Reviews*, 86(1-4):1–22, 1998.

[108] Steven Strogatz. *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, and Engineering*. Westview Press, 2 edition, 2015.

[109] Victor Szebehely. Theory of orbits: The restricted problem of three bodies. *Science*, 160(3827), 1968.

[110] LUVOIR Team et al. The luvoir mission concept study interim report. *arXiv preprint arXiv:1809.09668*, 2018.

[111] John T Trauger and Wesley A Traub. A laboratory demonstration of the capability to image an earth-like extrasolar planet. *Nature*, 446(7137):771, 2007.

[112] Casandra Webster and David Folta. Understanding the sun-earth libration point orbit formation flying challenges for wfirst and starshade. Technical report, IWSCFF 17-74 NASA Goddard Space Flight Center, 2017.

[113] Cassandra M Webster et al. Wide-field infrared survey telescope and starshade formation flying dynamics at sun-earth l2. *69th International Astronautical Congress*, 2018.

[114] Marc Wenger, François Ochsenbein, Daniel Egret, Pascal Dubois, François Bonnarel, Suzanne Borde, Françoise Genova, Gérard Jasniewicz, Suzanne Laloë, Soizick Lesteven, et al. The simbad astronomical database-the cds reference database for astronomical objects. *Astronomy and Astrophysics Supplement Series*, 143(1):9–22, 2000.

[115] K-P Wenzel. The isee-c mission and the ims. In *The scientific satellite programme during the International Magnetospheric Study*, pages 81–85. Springer, 1976.

[116] Stephen Wiggins. *Introduction to applied nonlinear dynamical systems and chaos*, volume 2. Springer Science & Business Media, 2003.

[117] P Willems. Starshade to trl5 (s5) technology development plan. *Technology Report, Jet Propulsion Laboratory*, 2018.

[118] Joshua N Winn. Exoplanet transits and occultations. *Exoplanets*, 1:55–77, 2010.

[119] Emily M Zimovan, Kathleen C Howell, and Diane C Davis. Near rectilinear halo orbits and their application in cis-lunar space. In *3rd IAA Conference on Dynamics and Control of Space Systems, Moscow, Russia*, page 20, 2017.