SPACE LOGISTICS NETWORK OPTIMIZATION WITH EMBEDDED PROPULSION TECHNOLOGY SELECTION

BY

BINDU BHARGAVI JAGANNATHA

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Aerospace Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2018

Urbana, Illinois

Doctoral Committee:

Assistant Professor Koki Ho, Chair
Professor Victoria L. Coverstone, University of Miami
Professor Bruce A. Conway
Professor Harrison H. M. Kim

Abstract

Establishing long-term human presence beyond low Earth-orbit will require cooperative use of emerging technologies, such as low-thrust solar electric propulsion, along with existing space exploration technologies. Trajectory analysis plays a key role in deciding the costs of using different propulsion technologies. Traditional trajectory design methods are usually confined to analyzing individual missions using high-thrust or low-thrust propulsion options — these techniques also do not usually consider the architectural aspects of the in-space network. On the other hand, determination of network architecture and mission sequence often relies on the expertise of mission designers. Thus, formulating and optimizing the problem as a multi-mission campaign can guide top-level decisions through rigorous mathematical modeling.

However, designing a multi-mission campaign with both propulsion options is generally computationally challenging due to the coupling of logistics network design and space transportation costs. Specifically, conventional space logistics planning methods are unable to account for the use of low-thrust vehicles for transportation due to the inherent nonlinear nature of associated costs. The aim of this work is to develop ways of handling low-thrust trajectory models within space logistics frameworks, so that propulsion technology trade-offs can be conducted internally.

This work develops two new frameworks for optimizing the combined use of low-thrust and high-thrust propulsion options within campaign-level space logistics planning tools. The first framework uses a chromosomal representation of network arc parameters to drive a multiobjective genetic algorithm that explores the tradespace. The second framework combines the generalized multicommodity network flow model with novel event-based time steps for dynamic space logistics optimization in the presence of nonlinear flight times associated with low-thrust transportation.

These methodologies are applied to the case study of Apollo-style crew missions to the lunar surface supported by in-space refueling via predeployed tanks delivered by cargo tugs. Although

the costs for high-thrust trajectories are a part of the inputs to the formulation, the costs for low-thrust trajectories are determined internally because of their dependence on the thrust-to–mass ratio. For an efficient evaluation of low-energy, low-thrust transfers in the Earth–moon system, an approximation method is implemented based on a Lyapunov feedback control law called Q-law, and dynamical systems theory. This preliminary trajectory design technique is validated against literature sources and is then used to closely estimate the costs associated with cargo deliveries using low-thrust tugs.

Numerical results from the new space logistics frameworks reveal campaign profiles where high-thrust and low-thrust propulsion options can be used cooperatively to achieve substantial improvement over the baseline no-refuel cases. These results also present multiple options that involve in-space propellant storage and transfer, deep-space rendezvous, and solar electric propulsion tugs for cargo delivery. By trading mission costs with campaign duration, these methods help quantify the impact of low-thrust solar electric propulsion in logistics supply planning.

To my family, for enduring my stubbornness.

Acknowledgments

Throughout my (long) career as a graduate student, I have been blessed with the most inspiring mentors willing to teach and support me. I am so grateful to have concluded my graduate career with my current advisor, Prof. Koki Ho. He has achieved a fine balance between providing guidance, support and freedom that I needed to mature as a researcher. His vision and focus on the bigger picture has always nudged me in the right direction.

I would also like to thank my committee members Prof. Victoria Coverstone, Prof. Bruce Conway and Prof. Harrison Kim for their interest in my work and for providing additional guidance when I needed it. Prof. Coverstone took me in and gave me opportunities to meander through different (but essential) learning experiences during my initial graduate years. I am also appreciative of the belief that Prof. Gary Swenson, Prof. Rodney Burton and the late Steven D'Urso expressed in me and the opportunities Dr. David Carroll continues to carve out for me. I also deeply appreciate the financial support provided by the Amelia Earhart Fellowship in 2014 by the Zonta International Foundation.

I have also had the good fortune of being surrounded by great colleagues – Bhargav, Danal, Alex & Laura, Erik, Pat, Ryne, Devin, Vicky decided to keep being my friends even after witnessing all the crazy. My current group members JB and Hao have also helped me during the course of my research work. I would also like to thank the staff of the AE department – Staci, Laura, Kendra, Diane – who made grad school paperwork a breeze everytime.

To my family, thank you for molding me into who I am today. And finally to Hari, who has not left my side in spite of all the tumultuous times and continues to give me every bit of love he has to offer. You make me stronger.

Table of Contents

List of Tables

List of Figures

List of Abbreviations

CR3BP	Circular Restricted Three-Body Problem
ED-GMCNF	Event-Driven Generalized Multicommodity Network Flow
EML_1	Earth-Moon Lagrange Point 1
EML_2	Earth-Moon Lagrange Point 2
GA	Genetic Algorithm
GMCNF	Generalized Multicommodity Network Flow
IMLEO	Initial mass in Low Earth Orbit
IMGTO	Initial mass in Geosynchronous Transfer Orbit
ISRU	In-Situ Resource Utilization
LELT	Low-Energy Low-Thrust (trajectories)
LP	Linear Programming
MILP	Mixed-Integer Linear Programming
MINLP	Mixed-Integer Nonlinear Programming
MOGA	Multiobjective Genetic Algorithm
NLP	Nonlinear Programming
PSO	Particle Swarm Optimization
PWL	Piecewise Linear
QLEPSO	Q-law LELT trajectory design with PSO
SEP	Solar Electric Propulsion
TE-GMCNF	Time-Expanded Generalized Multicommodity Network Flow
TLI	Translunar Injection
TOF	Time of flight

Chapter 1

Introduction

1.1 Background

Next-generation space exploration targets are set at sending humans to destinations that are beyond the International Space Station in Low Earth Orbit (LEO), and even past the moon, with the ultimate goal of reaching Mars. A metro map representation of nodes, potentially of interest as transit or target destinations, is shown in Figure 1.1. The road to Mars is expected to be paved

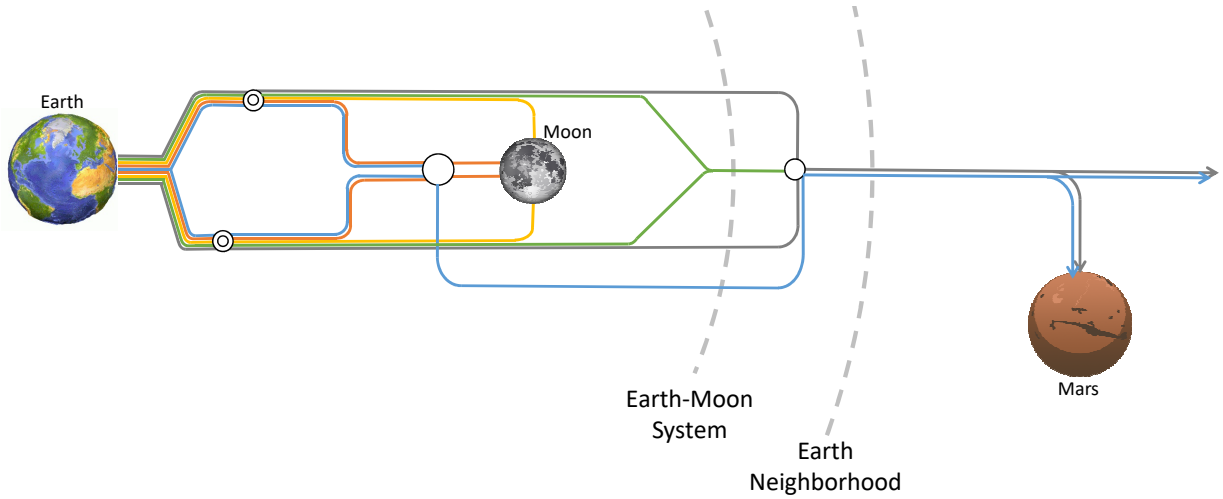

Figure 1.1: Metro-style map of near-term space exploration network illustrating a sample set of potential nodes of interest

with near-term missions to cislunar space meant to fill "strategic knowledge gaps" and progressively build up in-space capabilities. Many suggested architectures for Mars exploration campaigns call for the deployment of new space assets [1], for example, the idea of establishing an orbital outpost in the cislunar space called the Lunar Orbital Platform-Gateway has gained momentum recently. Such initiatives will require multiple long-duration missions; achieving exploration goals in an affordable manner will require a shift in perspective from designing independent missions to finding campaign-

level solutions, where missions can interact with each other in a concerted manner.

In this pursuit of sustaining future human space exploration activities, mathematical modeling can be used to optimize campaign profiles, instead of relying solely on mission designer expertise and intuition. Specifically, space logistics and supply chain planning tools can aid in the design of effective multi-mission campaigns; however, this field has not received adequate attention due to the limited nature of manned missions conducted thus far. The International Space Station (the current largest space facility) is located very close to the Earth and has been maintained by regular resupply of consummables. The only manned ventures beyond LEO were the Apollo missions, which landed humans on the lunar surface. These missions were short duration and required only a modest load of supplies/equipment, and thus were amenable to the carry-along-only strategy, where all required supplies were transported with the crew. However, future space campaigns to expand human activities beyond LEO, to destinations such as near-Earth objects or Mars for example, will be more complex and will need logistics strategies beyond resupply-only or carry-along-only.

Space logistics can also be better managed by taking advantage of new technologies being developed for the purposes of supporting long-term human exploration. For example, chemical propulsion systems are heavily relied upon currently to provide transport, but they possess low cargo capacities when compared to their total system mass. Maturing technologies such as low-thrust solar electric propulsion (SEP) systems can deliver higher mass fractions to their final destinations, albeit at the cost of long flight times. Other advancements such as in-situ resource utilization (ISRU), in-space manufacturing, and on-orbit propellant storage depots can also help expand exploration capabilities in incremental steps. Space logistics planning methods must be able to balance all associated costs of using these new emerging technologies in order to devise robust strategies.

Existing logistics tools for space supply chain management are based on terrestrial analogues. Extensive investigations of terrestrial logistics problems and associated solution methods have been conducted in the past [2]. However, the challenges of space exploration are vastly different from those of its terrestrial analogues. Transportation of commodities is vastly more expensive – using current propulsion technologies, the final deliverable mass is a very small fraction of the total system mass required. Trajectory considerations like launch windows, gravitational assists and low-energy

pathways can further complicate analysis. Additionally, current space logistics and supply chain optimization tools are unable to incorporate the vastly different paradigm of low-thrust (i.e., SEP) transfer costs. Briefly, this is because the associated costs are inherently nonlinear and require the trajectory itself to be designed.

Thus, the current need is for logistics and supply chain tools that can help mission designers make architectural decisions regarding space exploration campaigns by trading off cost, time *and* technology. Such tools will be able to conduct top-level trades in an automated manner and evaluate potential approaches to meeting the long-term strategic objectives. More importantly, by exploring a broad range of logistics strategies and supply chain tradespaces, solutions may emerge that expertise-based designs could miss.

The main purpose of this research is to develop a new way to optimize space logistics planning that also considers the effects of using different propulsion technologies cooperatively in the design of the supply chain network. The subsequent section presents a brief background about existing space logistics methods and their drawbacks specifically with regards to modeling low-thrust propulsion technologies, before explaining the current research objective in further detail.

1.2 Existing Space Logistics Frameworks

The field of campaign-oriented supply chain logistics planning for the purposes of space exploration has its roots in MIT's Space Logistics Project. This program gave rise to the network modeling approach to space logistics, resulting in an open-source discrete-event simulator called *SpaceNet* [3,4]. Use of metaheuristic algorithms to optimize commodities flow was also suggested [5, 6] as a part of the project. The modeling paradigms developed by the program have been continually refined over time and recently produced the generalized multi-commodity network flow (GMCNF) formulation [7, 8], along with its time-expanded variant [9]. Based on mixed-integer linear programming (MILP) formulations, they are both powerful formulations applicable to a broad range of space logistics scenarios.

The GMCNF model uses a graph-theoretic approach to cast the space exploration map into a network, with nodes corresponding to the candidates for surface destinations, orbits and staging locations. Arcs connecting these nodes represent directional flow of materials between nodes and

their associated costs. The mathematical formulation aims to minimize the cost of flow of commodities such as spacecraft, crew, cargo, propellant, etc. between these nodes. It builds on the basic minimum cost flow formulation, where demand/supply constraints are applied at the network nodes to conserve mass flow across the entire network. This basic model is "generalized" to account for flows with gains/losses, e.g., consumption of resources for arc transportation. This generalized model was then further upgraded to allow multiple commodities to flow within the network and include formulation to simulate their interactions, e.g. capacity constraints across arcs or conversion of resources from one type to another. Thus, the generalized multicommodity network flow (GMCNF) model emerged and was applied to campaign-level space logistics network problems with resource generation nodes. It enforces linear relationships in the form of mass balance constraints at the nodes, and flow transformation and concurrency constraints on the arcs. Considering a linear objective function, mixed-integer linear programming (MILP) solvers are then used to derive the optimal logistics strategy.

The major drawback of the conventional GMCNF model was its inability to deal with time-related behavior correctly, which resulted in time paradoxes and infeasible flow-generation loops. These shortcomings related to the static nature of the original GMCNF model were addressed by its time-expanded variant (i.e. the TE-GMCNF model), which assimilated the time dimension by duplicating the entire static network at discrete time steps and connecting them via holdover arcs. The TE-GMCNF model can be used to solve a space logistics problem if its dynamic behavior can be discretized using known time steps. Both models are predicated on the linearity of constraints and the objective function.

However, many of the costs and constraints arising from incorporating emerging technologies into supply chain logistics planning are inherently nonlinear. For example, the cost of traveling on an arc where an SEP spacecraft is the propulsive element depends on the total mass of all commodities being transported by the spacecraft. This means that the time of flight for low-thrust SEP propulsion systems is coupled to the amount of payload that needs to be delivered. Since the flight time between nodes is not fixed in this case, the discrete time steps used to set up the TE-GMCNF model cannot be pre-calculated. Fine or adaptive time-step discretization could be employed to capture low-thrust trajectory costs correctly, but this induces an increase in complexity

and in computational time. Hence, the inherently complex costs of low-thrust transportation cannot yet be captured in existing logistics modeling frameworks in a computationally efficient manner.

In light of this background and motivated by it, the current work develops ways of integrating low-thrust propulsion technologies into space logistics and supply chain planning in a computationally efficient manner. This will help maximize exploration capabilities even in early concept-design phases by aiding technology selection for optimal transportation architecture.

1.3 Research Objective

The aim of this research is to develop a unified space logistics method to design effective space campaigns that make use of both high-thrust and low-thrust propulsion technologies. By doing so, the tradespace of astrodynamics studies will be broadened and campaign designers will be able to explore architectural options that use different propulsion technologies cooperatively. This will also involve concurrent optimization of the network itself, which will be composed from a set of potential nodes.

The specific case study considered here is a campaign for sustaining short-duration crew missions to the lunar surface through in-space refueling support provided by predeployed resupply propellant. Propellant depots, long considered to support sustainable spacefaring activities [10,11], can be delivered to their operational locations before the crew missions commence. While the crew missions themselves will use high-thrust vehicles, the propellant depots can be predeployed to their orbital locations using a fleet of cargo delivery vehicles consisting of high-thrust and low-thrust tugs. By optimizing the problem as a tightly-integrated campaign (rather than single-launch isolated missions), cheaper solutions can be found that still conform to the crew time of flight constraint by relaxing the overall time required.

To illustrate this point, let us consider the Apollo missions. These missions comprised of a fully-fueled vehicle stack launched to lunar orbit in a single launch, where the Translunar Injection (TLI) was provided by the launch vehicle's upper stage and all other maneuvers were conducted by the Command+Service Module or the Lunar Module. The tradespace typically consists of comparing initial mass in LEO (IMLEO) costs with mission flight time, where the maximum limit on the time allowed arises from the limitations on how long the crew can spend in deep space. This

limitation prohibits the use of SEP vehicles alone for crew journeys in the cislunar region, unless fuel efficiencies are pushed beyond the current envelope and/or very large solar arrays can be used to provide megawatts of power.

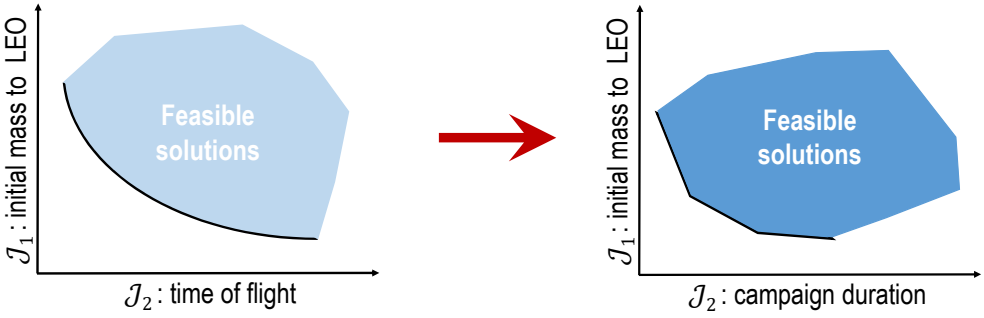

Figure 1.2: Transformation of objectives for typical Pareto front

However, if we move toward "distributed" launch architectures, the fuel for the crewed spacecraft can be delivered in advance to certain way stations along the crew's path, along with cargo not required by the crew during the journey. This strategy also lets mission designers circumvent the necessity of using super heavy-lift launch vehicles for crew launch. These refueling stations can be delivered using propulsion or trajectory options that are not normally considered for crewed spacecrafts, e.g. using cheaper transfers that may be long duration or using solar electric propulsion tugs. The tradespace then expands along the axis that denotes the total campaign duration, as shown in Figure 1.2. An example shape of the expected Pareto-front of IMLEO vs. campaign duration is also sketched out; all prospective solutions will lie above this frontier.

In this problem of refueling the crewed spacecraft during its journey, the final destination and mass delivery requirements for each of the cargo missions are decided by the profile of the crew missions. Additionally, the cargo vehicles used to deliver the refueling tanks must be chosen from an available fleet of vehicles, but the tug choices are unknown *a priori*. Any newly developed logistics and supply chain planning framework must be capable of optimizing the campaign along these dimensions (at minimum).

In further chapters, two frameworks are developed to help mission designers quickly obtain the Pareto front of cost versus time by optimizing propellant supply locations for crew missions concurrently with propulsion technology options (i.e., high-thrust or low-thrust) for corresponding

cargo deliveries. Both models are demonstrated for a return-to-the-Moon campaign, a case study that is derived from the Apollo missions but with added in-space refueling support for the crew. In the first framework, the problem is formulated in a combinatorial manner that allows for the use of multiobjective heuristic optimization algorithms. Its main drawback is the non-deterministic nature of solutions, with no guarantee of producing same results on successive runs or of optimality. Also, the formulation uses some assumptions to bound the problem size. This limiting assumption is eliminated in the second framework, which uses a new event-driven approach to discretizing the time dimension associated with the GMCNF logistics model.

An important aspect of this problem is the computation of setup costs of the in-space supply chain using available propulsion technologies. These mission costs are derived from the transportation of crew and cargo across the network. Many previous researchers have designed both high-thrust and low-thrust trajectories in the cislunar region. However, these studies do not consider the tradeoff between propulsion technologies for the purposes of meeting campaign-level objectives, which the current study aims to elaborate.

Transportation costs associated with the use of high-thrust propulsion systems can be modeled adequately using the impulsive-burn assumption, which means that they will be independent of the total mass. However, the transfer costs of low-thrust trajectories are coupled to the spacecraft's thrust-to-mass ratio (or thrust acceleration) due to the continuous-burn nature of the corresponding propulsive maneuvers. Thus, the costs for employing SEP have to be calculated internally by designing the trajectory itself. As is pertinent to the case study, these low-thrust transfers can further leverage the low-energy pathways found in the cislunar region. These combined low-energy low-thrust (LELT) trajectories are complex – existing design techniques do not align well with the needs of campaign-level analysis. Thus a new preliminary LELT trajectory design method is developed, and applied to the current case study to calculate the cost models required within the space logistics planning frameworks.

Mathematical modeling of space logistics in the manner developed in this work can support campaign-level architectural decisions where different propulsion technologies can be traded-off to predeploy mission elements. This new approach has a higher potential of sustaining long-term human space exploration as local design decision-making can be replaced by campaign-level

optimization.

1.4 Thesis Overview

Figure 1.3 represents the outline of this thesis. Beyond the current introduction, the contents are organized as follows:

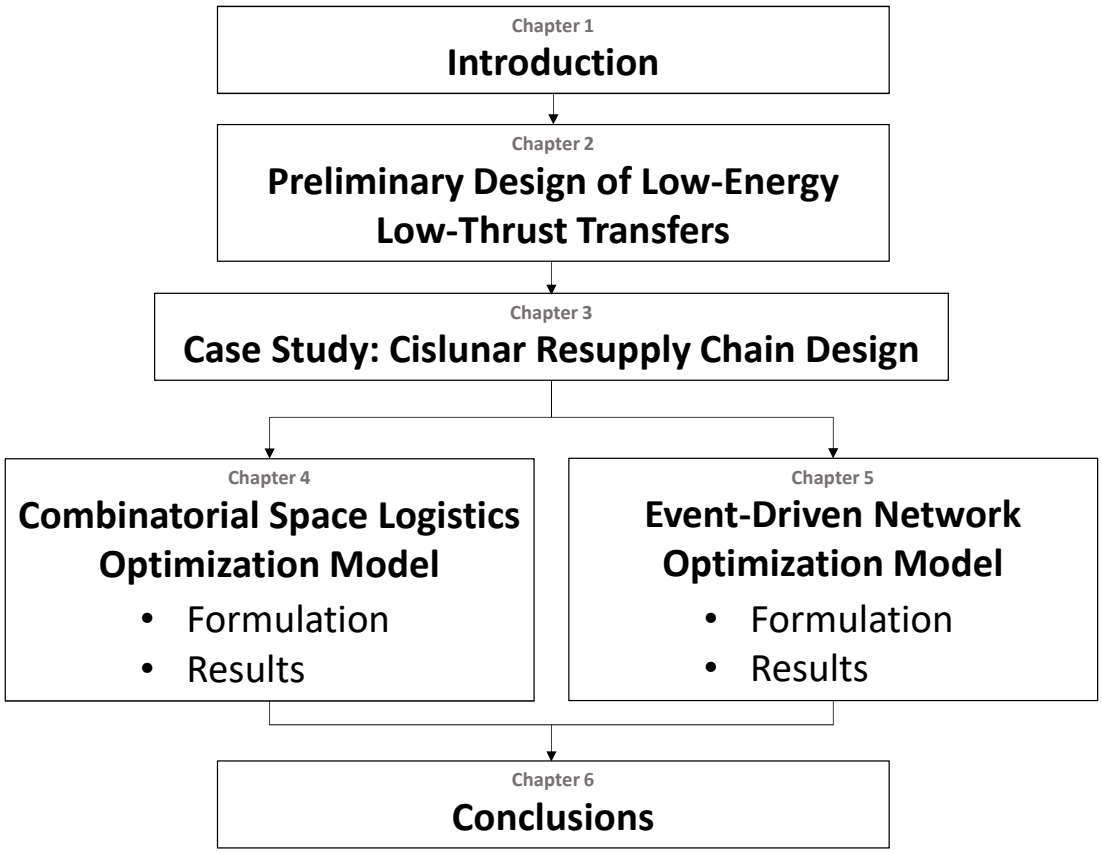

Figure 1.3: Thesis Outline

- Chapter 2 first provides a brief summary of dynamical system theory as applicable to the design of trajectories in the current case study. Then, the method developed for preliminary design of low-energy low-thrust trajectories in a three-body system is described. The efficiency and effectivity of the developed method is quantified by comparing it with methods and results that already exist in the literature.

- Chapter 3 details the problem setting and parameters of the case study, along with the

low-thrust cost model derived for it using the methods developed in the previous chapter.

- Chapter 4 describes the first solution methodology developed to concurrently optimize the space exploration network and the propulsion technology type within a campaign. The tradespace of parameters associated with all network arcs is evaluated using a multiobjective genetic algorithm. Optimal predeployment strategies are calculated for a campaign of one, two, or three human lunar missions, thus quantifying the game-changing impact of low-thrust propulsion for logistics supply planning in the selected case study.

- Chapter 5 presents an improved network logistics optimization framework, based on the GMCNF formulation. A new event-driven discretization is proposed for modeling dynamic behavior, so that the computational efficiency and robustness of using mixed-integer linear programming techniques can be realized. This second framework makes the problem formulation more general and is capable of exploring a wider tradespace of campaign architectures. Higher quality solutions are obtained, some displaying interesting interactions between missions that could be missed in a mission-wise analysis.

- Finally, Chapter 6 concludes this thesis by summarizing the key contributions and providing some ideas for future avenues for continuing this work.

Chapter 2

Preliminary Design of Low-Energy Low-Thrust Transfers in Cislunar Space

2.1 Introduction

Dynamical structures inherent to the circular restricted three-body system have been shown to play a significant role in the design of optimal spacecraft trajectories [12, 13] by providing low-energy pathways, e.g. the "interplanetary superhighway network" for transport within the solar system [14], or the "petit grand tour" of the Jovian moons [15]. These low-energy pathways are of specific interest within the cislunar region for potential applications in near-term space infrastructure deployment. Within this region, the unstable halo orbits around the collinear Earth-moon Lagrange points L_1 and L_2 are being examined to serve as staging nodes for the expansion of human space activities [11]. Trajectories that leverage these unstable periodic orbits and their associated invariant manifolds allow transfers within the cislunar region for lower costs than conventional transfers [16]. Missions such as Hiten [17] and ARTEMIS [18] (among others) have demonstrated the use of these routes to accomplish low-cost transport in the Earth-moon system. Other missions such as Genesis [19] have used similar pathways in the sun-Earth system.

Achieving these same transfers with low-thrust propulsion is currently a growing research field because higher mass fractions can be delivered, albeit at the cost of longer times of flight. A sample low-thrust trajectory between Earth and a halo orbit can be conducted by breaking up the journey into two phases – a spiral thrust arc and a coast arc. The multi-revolution thrust phase connects the given Earth parking orbit to the coasting phase on the invariant manifolds associated with the halo orbit. Previous studies have used different thrusting strategies to design such transfers by casting them into optimal control problems and using gradient-based techniques like nonlinear programming (NLP) to solve them. Senent et al. [20] studied low-thrust transfers to libration point orbits in the sun-Earth/moon system via their manifolds using *variable* specific impulse engines.

A similar problem in the Earth-moon three-body system was solved by Ozimek and Howell [21] by employing iterative continuation to lower the thrust value until problem constraints were satisfied. Other efforts by Mingotti et al. [22–26] used the concept of attainable sets to seed initial guesses for low-energy low-thrust transfers in the Earth-moon system and sun-perturbed Earth-moon system, as well as in the sun-Earth/moon system. Constant specific impulse engines were assumed for these trajectories. Martin and Conway [27, 28] further refined the solution to the Earth-to-halo transfer problem by including the low-thrust escape spiral as a part of the optimization setup. Optimal low-energy low-thrust transfers commencing with ballistic escape from Earth-orbit (i.e., impulsive maneuver provided by launch vehicle) in patched three-body problems have also been studied [29, 30].

Solving these optimal control problems usually requires the mission designer to seed initial guesses using an understanding of the underlying dynamics (e.g., through the use of Poincaré sections) or generate them by solving a homotopically similar problem or continuation techniques. This task can be non-trivial and time consuming. Additionally, solving the fully-formulated optimal problem can be computationally intensive as it can suffer from the curse of dimensionality typically associated with transfers consisting of significant spiraling (on the order of tens or hundreds of revolutions) and the iterative nature of NLP methods. Since these methods usually consume substantial computational resources to generate single point solutions, they do not lend themselves well to rapid trade studies.

To alleviate this issue, simpler shape-based methods or feedback-driven control laws can be used to design the spiral thrust arcs as orbital transfers around the primary bodies. Shape-based low-thrust trajectory design methods usually assume tangential control direction and output the thrust levels needed to achieve the transfer [31,32]. In the past, they have been wrapped inside stochastic solvers to determine the best solutions that satisfy practical engineering constraints [33, 34]. On the other hand, feedback control laws are pertinent for designing trajectories that constant-thrust constant-specific impulse engines can execute. These control laws use Lyapunov control theory [35, 36] or blend optimal control directions [37,38] for each orbital element to calculate the thrust profile using knowledge of only the current and target states. The steering direction provided by blending optimal control directions is adequate for simpler problems, but Lyapunov-based control laws are

usually more robust to the demands of complex transfers. Though generally suboptimal, they can be used to generate fully feasible initial guesses for higher-fidelity gradient-based optimization methods since they use significantly less computational resources and are relatively easy to implement. This also allows them to serve as good approximations for calculating the costs associated with a transfer, namely the propellant mass consumed and the time of flight.

The problem of designing low-energy low-thrust (LELT) trajectories is a challenging one due to the many revolutions required for the transfer as well as coordinating the use of coast arcs with the system dynamics. Existing methods of designing LELT transfers are ineffective for preliminary trajectory analysis because of high computational effort and the need for successive preprocessing. This gap is critical for the planning of support and logistics for crew missions, where early-stage mission analysis tools capable of quick tradespace exploration are required. In response to this literature gap, a preliminary LELT transfer design method was developed using a Lyapunov-feedback control law called Q-law within a particle swarm optimization (PSO) algorithm.

This method, designated QLEPSO, is detailed in the current chapter. The three-body dynamics are developed first in Section 2.2, so that the concept of low-energy pathways can be explained. This introductory section will present the pertinent fundamental concepts; for more detailed expositions, the reader is referred to external literature sources, e.g. [39–41]. Next, the feedback control law, called Q-law, used for the design of low-thrust trajectory segments is introduced in Section 2.3. The rest of the chapter (from Section 2.4 onward) will put these together and deal with the design of low-energy low-thrust (LELT) trajectories for the purpose of estimating transfer costs as required in space logistics modeling (detailed in later chapters).

2.2 Low-Energy Dynamical Model: Circular Restricted Three-Body Problem

The circular restricted three-body (CR3B) model is used throughout this work to represent the Earth-moon system as an adequate first approximation. It provides a simplified system to qualitatively capture the fundamental dynamical behavior of a spacecraft in cislunar space. Dynamical systems theory can be applied to this system to extract useful long-term qualitative behavior from

the representative coupled differential equations. The CR3B model shown in Figure 2.1 describes the motion of an infinitesimal point mass P_3 (of mass $m_3 \approx 0$) under the gravitational influence of two primary bodies P_1 and P_2 (with masses m_1 and m_2 respectively, such that $m_1 > m_2 \gg m_3$). Thus, by virtue of "restricting" the mass of P_3, its influence on the motion of the primaries can be neglected. Further, the frame of reference ($\mathbf{x} - \mathbf{y} - \mathbf{z}$ axes) is centered at the barycenter of P_1 and P_2, and allowed to rotate with a constant angular velocity matching the motion of these primaries around their barycenter. The use of this non-inertial reference frame, also called the synodic frame, reduces the number of first-order differential equations required to describe the system from the 18 required in the inertial frame ($\eta - \xi - \zeta$ axes) to 6.

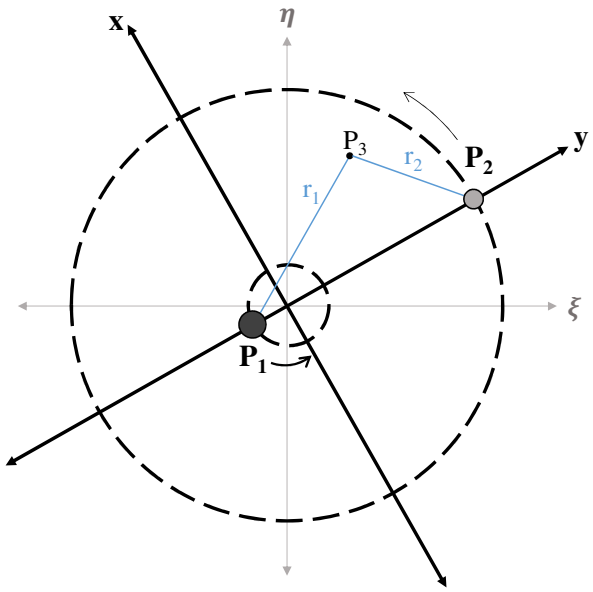

Figure 2.1: Geometry of the circular restricted three-body problem.

This system is further nondimensionalized for efficient numerical scaling – the synodic distance unit ($\mathrm{DU_{syn}}$) corresponds to the distance between the primaries P_1 and P_2, the normalized mass unit takes the combined value of the primary masses ($= m_1 + m_2$), and finally the synodic time unit ($\mathrm{TU_{syn}}$) is chosen to make the orbital period of the primaries about the system barycenter equal to 2π. The value of $\mathrm{TU_{syn}}$ is achieved by setting the universal gravitational constant G to unity.

The system can now be described using a single parameter, the mass ratio μ, which is defined as:

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

By definition, $m_1 \geq m_2$ and so $\mu \in [0, 0.5]$. Thus, the masses of the two primaries can also been normalized as $P_1 = (1 - \mu)$ and $P_2 = \mu$, with their orbital distance measured from the system barycenter as μ and $(1-\mu)$ $\mathrm{DU_{syn}}$ respectively. Using these conventions, the motion of the particle P_3 can be described by the following system of equations:

$$\begin{aligned} \ddot{x} - 2\dot{y} &= \frac{\partial \Psi}{\partial x} \ , \\ \ddot{y} + 2\dot{x} &= \frac{\partial \Psi}{\partial y} \ , \\ \ddot{z} &= \frac{\partial \Psi}{\partial z} \ , \end{aligned} \tag{2.1}$$

where potential function $\Psi = \frac{1}{2}(x^2 + y^2) + \frac{1-\mu}{\sqrt{(x+\mu)^2 + y^2 + z^2}} + \frac{\mu}{\sqrt{(x-1+\mu)^2 + y^2 + z^2}}$.

As formulated in the rotating/synodic reference frame, the above system of equations is time invariant. This allows a possibility of stationary or equilibrium points where the gravitational influences of both primary bodies are balanced against the centripetal forces. Five such points of equilibrium emerge in the circular three-body problem, called the Lagrange points. These stationary points (x^*, y^*, z^*) can be obtained by setting the acceleration and velocity terms in Equation (2.1) to zero:

$$-x^* = -\frac{(1-\mu)(x^* + \mu)}{(r_1)^3} - \frac{\mu\ (x^* - 1 + \mu)}{(r_2)^3} \ , \tag{2.2}$$

$$-y^* = -\frac{(1-\mu)\ y^*}{(r_1)^3} - \frac{\mu\ y^*}{(r_2)^3} \ , \tag{2.3}$$

$$0 = -\frac{(1-\mu)\ z^*}{(r_1)^3} - \frac{\mu\ z^*}{(r_2)^3} \ , \tag{2.4}$$

where $r_1 = \sqrt{(x^* + \mu)^2 + (y^*)^2 + (z^*)^2}$ and $r_2 = \sqrt{(x^* - 1 + \mu)^2 + (y^*)^2 + (z^*)^2}$.

Equation (2.4) can only hold if $z^* = 0$, which then couples equations (2.2) and (2.3). The first

set of equilibrium points arise when $y^* = 0$ also, thereby resulting in:

$$x^* = \frac{(1-\mu)(x^*+\mu)}{|x^*+\mu|^3} + \frac{\mu\ (x^*-1+\mu)}{|x^*-1+\mu|^3}\ . \tag{2.5}$$

This quintic equation in x^* can be solved numerically to obtain three real solutions, corresponding to the location of the three collinear Lagrange points (L_1, L_2 and L_3) – these points lie one each on either side of the primaries and one in between them in the rotating frame. Two other equilibria result when $r_1 = r_2 = 1$ is substituted into Equations (2.2) - (2.4). These equilateral Lagrange points (L_4 and L_5) lead and lag the smaller primary by 60° respectively. All five Lagrange points are marked in Figure 2.2.

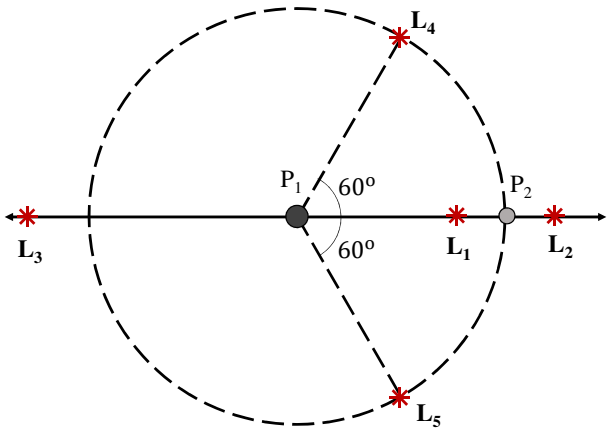

Figure 2.2: Equilibrium points in the CR3BP (Lagrange points).

The CR3B system thus defined admits a single integral of motion, commonly referred to as the Jacobi integral or Jacobi energy:

$$C = 2\Psi - (\dot{x}^2 + \dot{y}^2 + \dot{z}^2) \tag{2.6}$$

As this is the only quantity that is conserved in the system described by 6 differential equations, no analytical solution exists for the circular restricted three-body problem.

2.2.1 Periodic Orbits

Linear stability analysis of the collinear Lagrange points reveals them to be unstable [42], thereby giving rise to many families of periodic and quasi-periodic orbits in their vicinity. Of particular interest in the current work, halo orbits are one such family of non-planar periodic solutions. Continuous families of both Northern and Southern halo orbits exist at each collinear Lagrange point. Sample Northern halo orbits in Earth-moon L_1 and L_2 are shown in Figure 2.3 and colored according to their Jacobi energy – corresponding Southern halo orbits can be obtained by reflecting these across the X-Y plane. Even within the idealized CR3B problem framework, maintaining a spacecraft in a halo orbit will require active station-keeping as they are dynamically unstable.

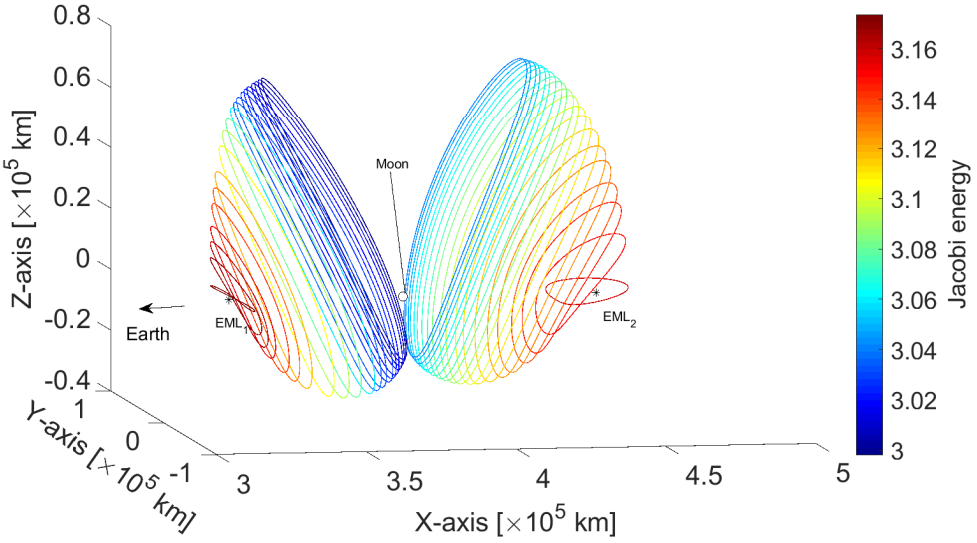

Figure 2.3: Northern halo orbits at EML_1 and EML_2.

These halo orbits (and other periodic orbits) can be generated through numerical means, though analytical approximations can be obtained in certain cases through linearization. In recent work, determining such periodic solutions has been formulated as a nonlinear optimization problem where the objective is to minimize the difference between initial and final states (i.e. position and velocity) at the end of integrating the system equations for some specified time period [27]. In the current work, the above nonlinear optimization problem is solved by direct transcription and fifth-order Gauss-Lobatto collocation [28, 43], with the initial guess solution provided by Richardson's third-order analytical approximation [44]. Once a single halo orbit has been generated, pseudo-arclength

continuation can be used to generate other members of the same family. A detailed description of these techniques as it pertains to the current problem is provided in Appendix A.

2.2.2 Invariant Manifolds

Unstable periodic orbits in the CR3B system have invariant manifolds associated with them. For halo orbits, these manifolds are tube-like structures that allow low-energy transport in the three-body system by virtue of their asymptotic arrival at/departure from the orbit. The manifolds of a sample L_1 halo orbit in the Earth-moon system are shown in Figure 2.4.

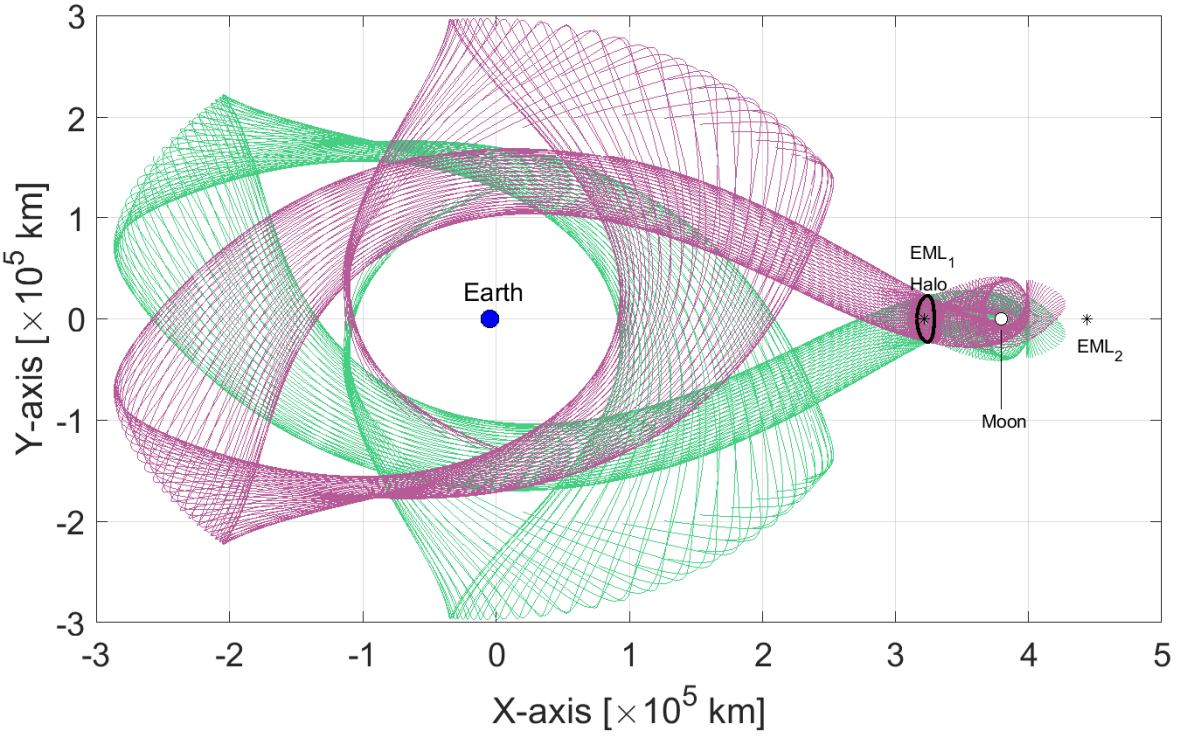

Figure 2.4: Stable (in green) and unstable (in pink) manifolds of an EML_1 halo orbit with out-of-plane amplitude $A_z = 8000$ km.

In order to generate the invariant manifolds associated with unstable periodic orbits in the CR3B system, one must begin with the state transition matrix (STM) corresponding to the orbit. An STM, given as $\mathbf{\Phi}(t, t_0)$, simply maps the final state of a given dynamical system at time t to its initial state at time t_0 as:

$$\mathbf{X}(t) = \mathbf{\Phi}(t, t_0)\mathbf{X}(t_0) \ . \tag{2.7}$$

The STM can be useful in gauging the sensitivity of the system's final state to the specified initial conditions, and is propagated using:

$$\dot{\mathbf{\Phi}}(t, t_0) = \mathbf{A}(t)\mathbf{\Phi}(t, t_0) \ , \tag{2.8}$$

where $\mathbf{\Phi}(t_0) = \mathbb{I}$. The state propagation matrix $\mathbf{A}$ for the CR3B problem is given by:

$$\mathbf{A} = \frac{\partial \dot{\mathbf{X}}}{\partial \mathbf{X}} = \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ \overline{\Psi}_{xx} & \overline{\Psi}_{xy} & \overline{\Psi}_{xz} & 0 & 2 & 0 \\ \overline{\Psi}_{yx} & \overline{\Psi}_{yy} & \overline{\Psi}_{yz} & -2 & 0 & 0 \\ \overline{\Psi}_{zx} & \overline{\Psi}_{zy} & \overline{\Psi}_{zz} & 0 & 0 & 0 \end{bmatrix} \ , \tag{2.9}$$

where $\overline{\Psi} = \frac{1}{2}(x^2 + y^2) + \frac{1-\mu}{\sqrt{(x+\mu)^2 + y^2 + z^2}} + \frac{\mu}{\sqrt{(x-1+\mu)^2 + y^2 + z^2}} + \frac{1}{2}\ \mu\ (1-\mu)$, and terms like $\overline{\Psi}_{xx}$ represent the second derivative of $\overline{\Psi}$ with respect to the corresponding state vector component/s. By combining Equations (2.7)-(2.9) and propagating them together with Equation (2.1), the system state and its STM can be obtained at any time.

The first step in generating the manifolds associated with any unstable periodic orbit is to obtain the orbit's STM at the end of a full revolution. If T is the orbital period, this STM $\mathbf{\Phi}(t_0 + T, t_0)$ (also called the monodromy matrix) can be used to characterize the stability of the periodic orbit under consideration [39,41,45]. Furthermore, its eigenvectors are useful in identifying the directions of the stable and unstable manifolds. The monodromy matrix of periodic solutions in the CR3B system has six eigenvalues λ_i (i = 1,2,3..6) corresponding to the eigenvectors $\mathbf{V}_i$. Due to the symplectic nature of the monodromy matrix, its eigenvalues occur in reciprocal pairs. Additionally, the existence of the Jacobi integral in the CR3B system renders one pair equal to unity. The relationships between the three reciprocal pairs of eigenvalues of the monodromy matrix

that specifically apply to unstable halo orbits are listed below:

$$\begin{aligned} &\lambda_1 = \frac{1}{\lambda_2} \quad (\lambda_1, \lambda_2 \text{ are real with } \lambda_1 > 1) \ , \\ &\lambda_3 = \lambda_4^* \quad (\lambda_3, \lambda_4 \text{ are complex with } |\lambda_3| = |\lambda_4| = 1) \ , \\ &\lambda_5 = \lambda_6 = 1 \ . \end{aligned} \tag{2.10}$$

The invariant manifolds of a halo orbit are locally linearly approximated by perturbing the eigenvectors that correspond to λ_1 and λ_2, labeled $\mathbf{V}^U$ and $\mathbf{V}^S$ respectively. The stable manifold (W^S) is defined as the set of all trajectories that asymptotically approach the orbit along its stable eigenvector as time goes to infinity. Similarly, the set of trajectories departing the orbit along its unstable eigenvector comprise its unstable manifold (W^U), i.e. they approach the orbit as time goes to negative infinity. The instantaneous eigenvectors can either be calculated at each point along the discretized orbit, or they can be obtained by simply propagating them by using the STM. Thus, the stable and unstable eigenvectors at time t along the orbit, denoted by $\mathbf{V}_t^S$ and $\mathbf{V}_t^U$, can be determined by propagating the eigenvectors of the monodromy matrix (i.e., $\mathbf{V}^S$ and $\mathbf{V}^U$) using the state transition matrix as:

$$\begin{aligned} \mathbf{V}_t^S &= \mathbf{\Phi}(t, t_0)\mathbf{V}^S \ , \\ \mathbf{V}_t^U &= \mathbf{\Phi}(t, t_0)\mathbf{V}^U \ . \end{aligned} \tag{2.11}$$

These instantaneous eigenvectors are normalized before use in manifold generation. Note that the eigenvalues themselves remain constant along a given periodic orbit in the CR3B system; it is only their corresponding eigenvectors that vary along the orbit.

At any point along the orbit, the invariant manifolds are mapped by perturbing the state along the corresponding eigenvector. A small perturbation ϵ is applied to the state $\mathbf{X}(t)$ to obtain the perturbed state at time t along the orbit:

$$\begin{aligned} \mathbf{X}_S &= \mathbf{X}(t) \pm \epsilon \mathbf{V}^S \ , \\ \mathbf{X}_U &= \mathbf{X}(t) \pm \epsilon \mathbf{V}^U \ . \end{aligned} \tag{2.12}$$

These perturbed states are propagated using Equation (2.1) to obtain the complete set of globalized manifolds associated with the three-body periodic orbit. The propagation is conducted backward

in time to generate the stable manifold and forward in time for the unstable manifold. A body of negligible mass placed on the stable manifold approaches the three-body periodic orbit asymptotically, while it can depart the same orbit on its unstable manifold. For certain orbital energies (characterized by the Jacobi integral), these manifolds structurally resemble tubes; however, this tube-like structure is lost in regions of dynamic instabilities such as close encounters with either of the primary bodies. Each type of manifold consists of two branches, one that initially moves towards the larger primary in the system and the other that initially moves away from it.

The magnitude of the disturbance ϵ used to map the manifolds must be chosen judiciously – it must be small enough to retain the validity of the local linear approximation used to generate the manifolds, but also large enough to move away from the periodic orbit within a practical timespan. The choice of this ϵ (within a reasonable range) only alters the time taken by the manifold to approach/depart the three-body orbit, leaving its characteristics far away from the orbit almost unchanged.

2.3 Low-Thrust Orbital Transfers: Q-law

A complete LELT trajectory consists of a manifold segment (where the spacecraft coasts), and a spiral thrust segment that connects the spacecraft's parking orbit to the manifold. The many-revolution thrusting phase of the transfer is approximated using a Lyapunov feedback control law called Q-law, developed by Petropoulos [36, 46] for designing low-thrust orbital transfers. The native version of this algorithm is presented in this section – modifications to it for use within QLEPSO are described later in Section 2.4.

Q-law decides the thrust direction at any point during the trajectory based on the "remaining distance to go" and can target defined changes in all orbital elements except true anomaly. It has been previously used successfully to provide good initial guesses for optimal transfers [47, 48] using constant specific impulse, constant thrust engines in a two-body system model. This feedback control law provides a non-iterative alternative to calculating a low-thrust transfer between arbitrary orbits, thereby providing a computationally efficient control law.

The basis of calculating these transfers is the following proximity quotient Q, a candidate

Lyapunov function that quantifies the proximity of the current orbit to the target orbit [36]:

$$Q = (1 + W_p P) \sum_{\text{œ}} W_{\text{œ}} S_{\text{œ}} \left[\frac{d(\text{œ}, \text{œ}_T)}{\dot{\text{œ}}_{\text{xx}}} \right]^2 \quad , \tag{2.13}$$

where œ denotes the instantaneous values of the orbital elements being targeted and the subscript T denotes their target values. The orbital elements that can be targeted are semi-major axis (a), eccentricity (e), inclination (i), argument of periapse (ω), and longitude of ascending node (Ω). W_p and $W_{\text{œ}}$ are scalar weights such that $W_{\text{œ}}$ is non-zero for the elements being targeted and W_p is non-zero when enforcing a minimum-periapse-radius condition using the penalty function P:

$$P = \exp \left[k \left(1 - \frac{r_p}{r_{p_{\min}}} \right) \right] \quad , \tag{2.14}$$

where k is a constant, $r_p = a(1 - e)$ and $r_{p_{\min}}$ is the desired lower limit for the osculating periapse radius. $S_{\text{œ}}$ is the scaling function and is suggested to have the following form in order to aid convergence of the semi-major axis to its target value:

$$S_{\text{œ}} = \begin{cases} \left[1 + \left(\dfrac{a - a_T}{m \cdot a_T} \right)^n \right]^{\frac{1}{r}} & \text{for } \text{œ} = a \ , \\ 1 & \text{for } \text{œ} = e, i, \omega, \Omega \ , \end{cases} \tag{2.15}$$

where m, n and r are pre-selected scalar constants. The distance function $d(\text{œ}, \text{œ}_T)$ denotes the difference between current and target orbital elements as:

$$d(\text{œ}, \text{œ}_T) = \begin{cases} \text{œ} - \text{œ}_T & \text{for } \text{œ} = a, e, i \ , \\ \cos^{-1} \left[\cos(\text{œ} - \text{œ}_T) \right] & \text{for } \text{œ} = \omega, \Omega \ , \end{cases} \tag{2.16}$$

where the value chosen for the arc cosine lies between 0 and π. The term $\dot{\text{œ}}_{\text{xx}}$ in Equation (2.13) denotes the maximum rate of change of orbital element due to thrust over thrust direction and over true anomaly. Analytic expressions for it are derived by using the Gauss variational equations.

If the perturbing acceleration $\mathbf{F}$ is identified by its components (F_r, F_θ, F_h) in the radial, circumferential and angular momentum directions respectively, then the Gauss variational equations

are given as:

$$\begin{aligned}
\frac{da}{dt} &= \frac{2a^2}{h}\left(e\sin\theta F_r + \frac{p}{r}F_\theta\right) \ , \\
\frac{de}{dt} &= \frac{1}{h}\left\{p\sin\theta F_r + [(p+r)\cos\theta + re]F_\theta\right\} \ , \\
\frac{di}{dt} &= \frac{r\cos(\theta+\omega)}{h}F_h \ , \\
\frac{d\Omega}{dt} &= \frac{r\sin(\theta+\omega)}{h\sin i}F_h \ , \\
\frac{d\omega}{dt} &= \frac{1}{eh}\left\{-p\cos\theta F_r + (p+r)\sin\theta F_\theta\right\} - \frac{r\sin(\theta+\omega)\cos i}{h\sin i}F_h \ , \\
\frac{d\theta}{dt} &= \frac{h}{r^2} + \frac{1}{eh}\left\{p\cos\theta F_r - (p+r)\sin\theta F_\theta\right\} \ ,
\end{aligned} \tag{2.17}$$

where t represents time, p is the semilatus rectum, h is the angular momentum of the orbit, and r is radius from the central body given by $r = p/(1 + e\cos\theta)$. By representing the vector of targetable orbital elements as $\mathbf{X}$, the above equations can be reduced to the form:

$$\dot{\mathbf{X}} = \mathbf{\Xi}(\mathbf{X}, \theta)\mathbf{F} \ . \tag{2.18}$$

In order to derive $\dot{\alpha}_{\mathrm{xx}}$ from Equation (2.17), the perturbing acceleration $\mathbf{F}$ must be first defined. If T = thrust provided by the engines, m = mass of the spacecraft, α = in-plane thrusting angle and β = out-of-plane thrusting angle, then the vector $\mathbf{\Gamma}$ representing the control acceleration vector provided by the low-thrust engines, specified by its components $(\Gamma_r, \Gamma_\theta, \Gamma_h)$, is:

$$\begin{aligned}
\Gamma_r &= \Gamma\ \cos\beta\sin\alpha \ , \\
\Gamma_\theta &= \Gamma\ \cos\beta\cos\alpha \ , \\
\Gamma_h &= \Gamma\ \sin\beta \ , \\
\text{where } \Gamma &= |\mathbf{\Gamma}| = \frac{T}{m} \ .
\end{aligned} \tag{2.19}$$

Only this acceleration provided by the low-thrust engine is considered to constitute the perturbing acceleration $\mathbf{F}$ for the purposes of deriving $\dot{\alpha}_{\mathrm{xx}}$ (and thus the Lyapunov-optimal control), i.e. $\mathbf{F} = \mathbf{\Gamma}$. By setting $\dot{\alpha}_{\mathrm{xx}} = \max_{\alpha,\beta,\theta} \dot{\alpha}$, the corresponding terms in the expression for Q can be obtained from

Equation (2.17) as:

$$
\begin{aligned}
\dot{a}_{\mathrm{xx}} &= 2\Gamma\sqrt{\frac{a^3(1+e)}{\mu(1-e)}} \ , \\
\dot{e}_{\mathrm{xx}} &= \frac{2p\Gamma}{h} \ , \\
\dot{i}_{\mathrm{xx}} &= \frac{p\Gamma}{h\left(\sqrt{1-e^2\sin^2\omega}-e|\cos\omega|\right)} \ , \\
\dot{\Omega}_{\mathrm{xx}} &= \frac{p\Gamma}{h\sin i\left(\sqrt{1-e^2\cos^2\omega}-e|\sin\omega|\right)} \ , \\
\dot{\omega}_{\mathrm{xx}} &= \frac{\dot{\omega}_{\mathrm{xx}i}+b\ \dot{\omega}_{\mathrm{xx}o}}{1+b} \ ,
\end{aligned}
\tag{2.20}
$$

where $b = 0.01$ nominally, and the other terms in $\dot{\omega}_{\mathrm{xx}}$ are defined as:

$$
\begin{aligned}
&\dot{\omega}_{\mathrm{xx}o} = \max_{\theta}\left(\dot{\varpi}|_{\beta=\pi/2}\right) = \dot{\Omega}_{\mathrm{xx}}|\cos i| \ , \\
&\dot{\omega}_{\mathrm{xx}i} = \max_{\alpha,\theta}\left(\dot{\varpi}|_{\beta=0}\right) = \frac{\Gamma}{eh}\sqrt{p^2\cos^2\theta_{\mathrm{xx}}+(p+r_{\mathrm{xx}})^2\sin^2\theta_{\mathrm{xx}}} \ , \\
&\quad \text{with } \cos\theta_{\mathrm{xx}} = \left[\frac{1-e^2}{2e^3}+\sqrt{\frac{1}{4}\left(\frac{1-e^2}{e^3}\right)^2+\frac{1}{27}}\right]^{\frac{1}{3}} - \left[-\frac{1-e^2}{2e^3}+\sqrt{\frac{1}{4}\left(\frac{1-e^2}{e^3}\right)^2+\frac{1}{27}}\right]^{\frac{1}{3}} - \frac{1}{e} \ , \\
&\quad \text{and } r_{\mathrm{xx}} = \frac{p}{1+e\cos\theta_{\mathrm{xx}}} \ .
\end{aligned}
$$

The proximity quotient Q in Equation (2.13) is now fully defined. Its value is zero at the target orbit and positive along the transfer trajectory. The thrust direction at any instant along the transfer is calculated such that the proximity quotient is driven to zero as quickly as possible, i.e., the direction in which $\dot{Q}$ attains its most negative value. By applying Lyapunov control theory, the instantaneous optimal control vector can be derived using the notation in Equation (2.18) as:

$$
\mathbf{\Gamma} = -\mathbf{\Xi}^T\left(\frac{\partial Q}{\partial \mathbf{X}}\right)^T \ . \tag{2.21}
$$

Symbolic differentiation is employed at this juncture to determine the Q-law control vector.

In addition to providing the Lyapunov-optimal control, the Q-law algorithm has two mechanisms for positioning the coast arcs based on the effectivity of thrusting along the osculating orbit – absolute effectivity cut-off and relative effectivity cut-off. The absolute effectivity (η_a) at the

given true anomaly along the orbital transfer is given as

$$\eta_a = \frac{\dot{Q}_n}{\dot{Q}_{nn}} \ , \tag{2.22}$$

and relative effectivity (η_r) as

$$\eta_r = \frac{\dot{Q}_n - \dot{Q}_{nx}}{\dot{Q}_{nn} - \dot{Q}_{nx}} \ , \tag{2.23}$$

where

$$\dot{Q}_n = \min_{\alpha,\beta} \dot{Q} \ ,$$

$$\dot{Q}_{nn} = \min_{\theta} \dot{Q}_n \ ,$$

$$\dot{Q}_{nx} = \max_{\theta} \dot{Q}_n \ .$$

The values of $\dot{Q}_{nn}$ and $\dot{Q}_{nx}$ are obtained numerically by discretizing the osculating orbit. The algorithm can be set to thrust only when the effectivities are above some cut-off value, i.e. if $\eta_a \geq \eta_{a_{\text{cut}}}$ and/or $\eta_r \geq \eta_{r_{\text{cut}}}$. Generally, higher values for effectivity cut-offs result in longer time-of-flight solutions, but can deliver greater payload mass ratios. These parameters (if used) manage the extent of coasting during the transfer phase governed by this control law.

The candidate Lyapunov function Q assumed in Equation (2.13) represents a "quadratic-time-to-go" with its $\left[\frac{d(\text{œ}, \text{œ}_T)}{\dot{\text{œ}}_{\text{xx}}}\right]^2$ term. Due to the existence of this term, Q is not strictly stable in a Lyapunov sense – it approaches zero not only when osculating orbital elements are close to their target values, but also as the semi-major axis approaches infinity (i.e., as $a \to \infty$). Thus the term $S_{\text{œ}}$ as defined in Equation (2.15) is appended to Q to return it to the region of convergence around the target orbit. Although Q-law's convergence properties have not been mathematically proven, it has been empirically shown to converge to the target orbit through numerical experiments on a broad range of orbital transfers [46, 47, 49]. It has also been shown to provide rapid and acceptable estimates of the trajectory behavior and the performance costs of orbital transfers within a two-body perspective [50]. This feedback control algorithm will be used in the design of LELT transfers as described in the following section.

2.4 Low-Energy Low-Thrust Transfer Design: Q-law with Particle Swarm Optimization (QLEPSO)

To design a complete LELT transfer, the low-thrust spiral arc around the primary body must be patched together with a coast arc on the invariant manifold associated with the desired halo orbit. The method developed during the course of this research for this purpose, called QLEPSO, is specifically demonstrated on a transfer that originates at an Earth orbit and delivers the spacecraft to a halo orbit located at the Earth-moon L_1 Lagrange point (EML_1). The problem being considered here is that of rapid preliminary design of minimum-fuel time-constrained LELT trajectories in a three-body system. The complete trajectory consists of two phases – the powered escape spiral to raise the spacecraft's orbit until it is delivered to the manifold, followed by a ballistic coast on the manifold, which is the low-energy pathway that leads to the desired halo orbit. The low-energy dynamics introduced in Section 2.2 are combined with the low-thrust trajectory design method outlined in Section 2.3, with the two phases patched together at a point on the manifold surface. A stochastic optimization algorithm is then utilized to determine the ideal values of the design parameters that control the two phases, as well as the patch point itself. This algorithm uses cooperative swarm techniques to find solutions that minimize the fuel consumption while satisfying the time-of-flight constraint.

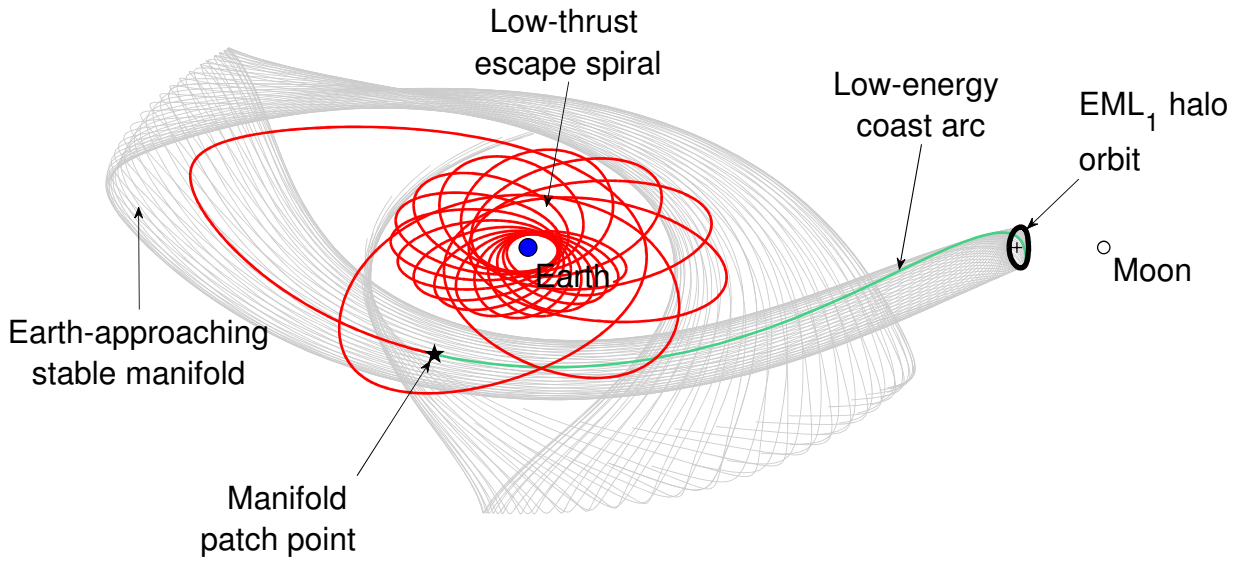

Figure 2.5: Overview of LELT trajectory phases.

The sample trajectory being designed is thematically sketched in Figure 2.5. The process of designing the low-energy phase and the low-thrust phase of the transfer is described first in this section, followed by the procedure to patch them together to yield the desired solutions.

2.4.1 Low-Energy Phase Design with QLEPSO

As introduced in Section 2.2, the invariant manifolds associated with unstable three-body orbits can be used as low-energy transportation pathways. The dynamical structures of significance in the current sample trajectory design are the EML_1 halo orbit and the Earth-approaching branch of its stable manifold. Once a manifold branch is selected, every point on the manifold structure can be parametrized using two values – point of origin of the corresponding manifold trajectory on the halo orbit, τ_h, and time to reach the halo once on that manifold, τ_m. These two quantities are marked in Figure 2.6 for an EML_1 halo orbit with out-of-plane amplitude $A_z = 8000$ km.

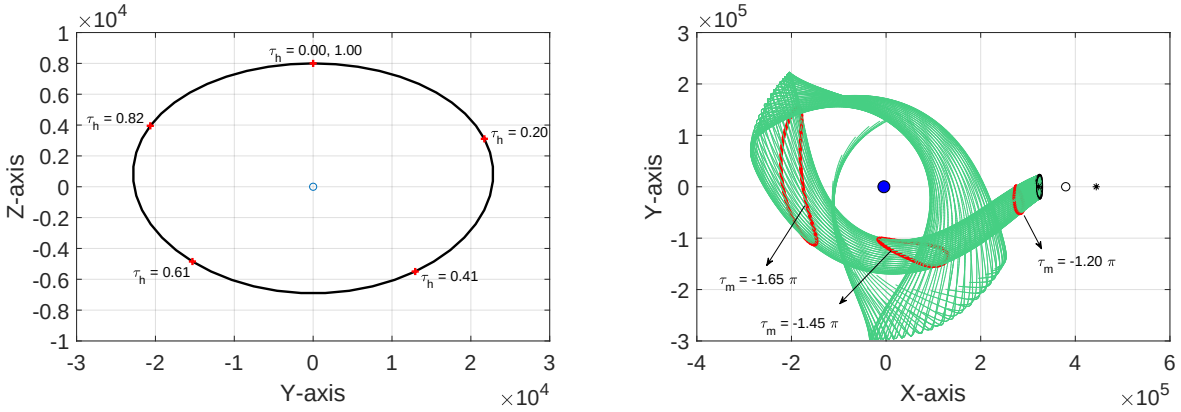

(a) Halo orbit with various τ_h values marked. (b) Stable manifold with various τ_m values marked.

Figure 2.6: Parametrization of the invariant manifold.

The halo orbit insertion parameter (τ_h) represents the normalized position around the periodic orbit with respect to the point where the maximum out-of-plane amplitude is achieved. The manifold insertion parameter (τ_m), on the other hand, is time propagated to arrive at the given point on the manifold. The direction of propagation is forward in time for the case of the unstable manifold, and backward for generating the stable manifold (as noted earlier).

2.4.2 Low-Thrust Phase Design within QLEPSO

By breaking down the desired LELT trajectories as orbital transfers around the primary body (Earth in this case), with the stable manifold providing "free" transport to the destination three-body periodic orbit, Q-law can be applied here to design the escape spiral. The algorithm has been modified to include the effect of third-body gravitational perturbations during system propagation. Gravitational perturbations due to the effect of the moon are dominant in the sample trajectory design problem, especially in the vicinity of the manifolds. The angular position of the moon in the Earth-fixed frame, λ, is appended to the system equations for state propagation as:

$$\frac{d\lambda}{dt} = n_{\,\mathbb{D}} \ , \tag{2.24}$$

where $n_{\,\mathbb{D}}$ is the mean motion of the moon as it orbits the Earth. The perturbing acceleration in Equation (2.17) becomes:

$$\mathbf{F} = \mathbf{\Gamma} + \mathbf{F}_{\,\mathbb{D}} \ , \tag{2.25}$$
$$\text{where } \mathbf{F}_{\,\mathbb{D}} = \mu_{\,\mathbb{D}} \left(\frac{\mathbf{r}_{\,\mathbb{D}} - \mathbf{r}}{|\mathbf{r}_{\,\mathbb{D}} - \mathbf{r}|^3} - \frac{\mathbf{r}_{\,\mathbb{D}}}{|\mathbf{r}_{\,\mathbb{D}}|^3} \right) \ .$$

The vectors $\mathbf{r}$ (spacecraft position) and $\mathbf{r}_{\,\mathbb{D}}$ (lunar position) in terms of their (r, θ, h) components are as follows and are meant for substitution into Equation (2.25):

$$\mathbf{r}_{\,\mathbb{D}} = a_{\,\mathbb{D}} \begin{bmatrix} \cos\theta\cos(\lambda-\Omega) + \cos i \sin\theta \sin(\lambda-\Omega) \\ -\sin\theta\cos(\lambda-\Omega) + \cos i \cos\theta \sin(\lambda-\Omega) \\ -\sin i \sin(\lambda-\Omega) \end{bmatrix} \text{ and } \mathbf{r} = \begin{bmatrix} r \\ 0 \\ 0 \end{bmatrix} \ . \tag{2.26}$$

The mass of the spacecraft is also tracked along the trajectory, through the following equation:

$$\frac{dm}{dt} = \frac{-T}{gI_{sp}} \ , \tag{2.27}$$

where g = standard acceleration due to gravity and I_{sp} = specific impulse imparted by the low-thrust engine.

Although system propagation includes third-body effects in this work, other Q-law settings are

derived from the original references. Since the Lyapunov-optimal control vector is not re-derived to include these additional gravitational effects, the caveat is that the thrust profile may then fight these perturbations rather than exploit them. The Q-law algorithm within QLEPSO implements fixed-step integration as prescribed in the original work [46]. The absolute effectivity (η_a) cutoff is used as the sole coasting mechanism within the escape spiral. The effect of eclipse periods could be included as unplanned shut-offs, but is not implemented here.

2.4.3 Complete LELT Trajectory Design

As an example to demonstrate the QLEPSO trajectory design strategy, consider again the trajectory from an initial Earth parking orbit to an EML_1 halo orbit. The spacecraft must be delivered to the stable manifold for a ballistic transfer to the corresponding halo orbit. Once the manifold is generated, all the points on this manifold are considered viable targets for the low-thrust arc. The equivalent classical orbital representation of these target points is found by converting the coordinates from the Earth-moon rotating frame to an Earth-centered inertial frame (refer Appendix B). These candidate patch points have fixed true anomalies and so, Q-law has to be propagated backwards in time starting from the final manifold location (i.e., the patch point) and ending at initial parking orbit. A complete LELT trajectory can be designed in this manner for each point on the manifold. This trajectory is refined by use of advantageous values for decision variables, selected by a stochastic particle swarm optimization (PSO) algorithm. This complete procedure is summarized in Figure 2.7.

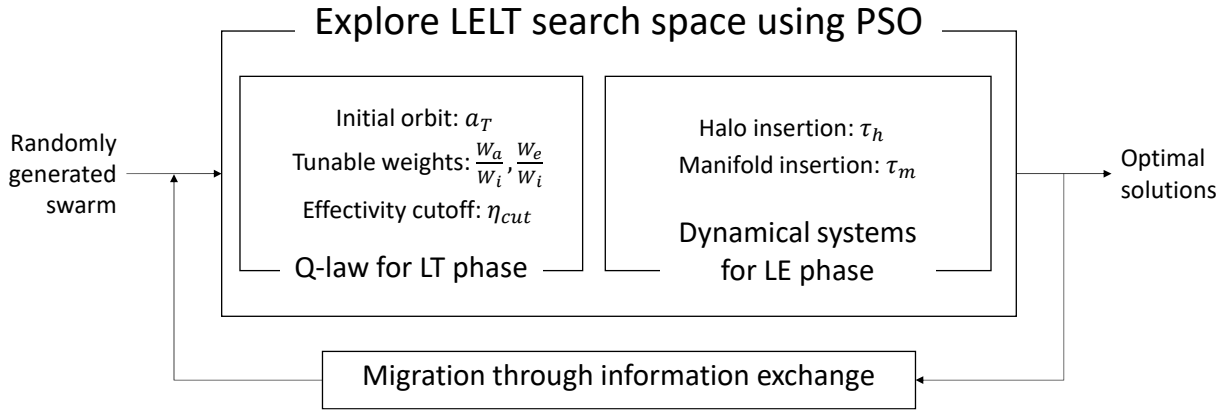

Figure 2.7: Overview of the QLEPSO method.

Cost Function

The LELT trajectories being designed here will consist of two phases connected to each other at the manifold patch point – the orbit raising phase consisting of many revolutions around the Earth, and the final coast arc on the manifold that asymptotically brings the spacecraft to the halo orbit. The objective for this complete LELT trajectory is to minimize the propellant mass consumed. Reduction in consumed propellant mass is typically obtained at the cost of higher time of flight (TOF) in low-thrust trajectory solutions. Thus, the transfers in this study are sought for fixed flight times, where the flight time includes the intermittently-powered escape spiral as well as the ballistic coast on the manifold. Mathematically, this optimization problem can be stated as:

$$\begin{aligned} \underset{\boldsymbol{\chi}}{\text{minimize}} \quad & \mathcal{J}(\boldsymbol{\chi}) = m_{\text{propellant}} \ , \\ \text{subject to} \quad & \mathbf{f}(\boldsymbol{\chi}) \leq \mathbf{b_C}, \\ & \mathbf{b_L} \leq \boldsymbol{\chi} \leq \mathbf{b_U}, \\ \text{where} \quad & \boldsymbol{\chi} = \left[\ a_T, \ \frac{W_a}{W_i}, \ \frac{W_e}{W_i}, \ \eta_{\text{cut}}, \ \tau_h, \ \tau_m \right]^T. \end{aligned} \tag{2.28}$$

This problem is restricted by two constraints, one regarding the initial parking orbit and the other relating to the allowable maximum time of flight. Additionally, the decision variables are also subject to bounds. These constraints ($\mathbf{f}(\boldsymbol{\chi})$ and $\mathbf{b_C}$) and the problem bounds ($\mathbf{b_U}$ and $\mathbf{b_L}$) are described in detail after introducing the decision variables.

Decision Variables

The decision vector $\boldsymbol{\chi}$ is comprised of variables that determine the cost of the transfer. Within the described LELT trajectory design strategy, values of the following variables can be judiciously selected to construct cheaper transfers:

1. Parking orbit specification $\rightarrow a_T$
2. Variables in Q-law computations $\rightarrow \dfrac{W_a}{W_i}, \dfrac{W_e}{W_i}$ and η_{cut}
3. Manifold patch point selection $\rightarrow \tau_h$ and τ_m

Following the trajectory in its temporal order, the first parameter that can be optimized is

the parking orbit from which the spacecraft begins its journey. The two common launch orbits are the Low Earth Orbit (LEO) and the Geosynchronous Transfer Orbit (GTO), which is used to deliver satellites to their operational Geosynchronous Earth Orbit (GEO). Although some previous LELT trajectory design studies have selected the GTO as the initial parking orbit [22, 43], some benefit could be realized by optimizing this orbit itself. The search space for this initial parking orbit is limited in the current study by the Δv that can be imparted by the launch vehicle, which differentiates it from the analysis in previous works by Mingotti et. al. [24, 25]. This available Δv is calculated by fixing the perigee at LEO altitude and then calculating the single impulsive maneuver required to raise the orbit from LEO to GTO (with no inclination change):

$$\Delta \mathrm{v}_{\mathrm{budget}} = v_p^{\mathrm{GTO}} - v_p^{\mathrm{LEO}} \quad . \tag{2.29}$$

The subscript p denotes perigee in the above equation. Given this Δv budget after launch, two parameters can be changed: the apogee and the inclination of the initial parking orbit for the LELT journey. This Δv can be distributed optimally between increasing the semi-major axis and changing the inclination of the parking orbit, such that the residual changes required to reach the manifold can be executed by the low-thrust phase of the trajectory. However, since the semi-major axis and inclination of the parking orbit that can be achieved are constrained by the Δv budget considered in Equation (2.29), only one of them needs to be a decision variable. Assuming that a *single* maneuver conducted at perigee places the spacecraft into its parking orbit, this constraint can be expressed as:

$$\begin{gathered} \sqrt{\left(v_p^{\mathrm{park}}\right)^2 + \left(v_p^{\mathrm{LEO}}\right)^2 - 2 v_p^{\mathrm{park}} v_p^{\mathrm{LEO}} \cos \Delta i} \;= \Delta \mathrm{v}_{\mathrm{budget}} \;, \\ \text{where} \;\; v_{\mathrm{p}}^{\mathrm{park}} = \sqrt{\mu_{\oplus}\left(\tfrac{2}{r_p} - \tfrac{1}{a_T}\right)} \;. \end{gathered} \tag{2.30}$$

Thus, the selection of the parking orbit can be completely parametrized by its semi-major axis a_T. This is possible due to the assumption that the parking orbit is achieved by a single impulse after launch. If multiple impulses are allowed, additional parameters may be included in the decision vector to match the new problem structure.

Consequent to determination of the Earth parking orbit, the low-thrust phase of the transfer

is calculated using Q-law. Only three orbital elements of the parking orbit are targeted – semi-major axis (a), eccentricity (e) and inclination (i) – and so only three weights (W_a, W_e, W_i) need to be specified. The remaining secular elements of the initial parking orbit (i.e., ω and Ω) are only incidentally chosen by the final solutions and can be matched by adjusting the time frame of the spacecraft launch or its insertion into the parking orbit. The low-thrust phase can be optimized by tuning the weights of Q-law by using the relative values of W_a, W_e and W_i in the form of W_a/W_i and W_e/W_i, as well as by adjusting the value of η_{cut}. The coefficients in the formulation of the scaling function S_{oe} in Equation (2.15) (refer [36, 46]) could also be included in the decision vector; however, sensitivity of the low-thrust trajectory to these parameters has been shown to be negligible [49].

The final optimization variables select the manifold patch point that separates the powered section of the trajectory from the purely ballistic portion. It is parametrized by the halo orbit insertion parameter τ_h and the manifold insertion parameter τ_m. As explained before, τ_h is the normalized position; the value of zero (and one due to periodicity) commonly indicates the point at which the halo orbit achieves its maximum value along the positive Z-axis. Similarly, τ_m is typically specified in normalized time units (in the Earth-moon rotating / synodic frame). In order to clearly indicate the direction of propagation, τ_m is negative for the stable manifold and positive for the unstable manifold. The low-thrust escape spiral targets the manifold patch point specified by τ_m and τ_h for spacecraft insertion, beyond which ballistic travel to the desired halo obit completes the transfer. This patch point is converted from its (τ_m, τ_h) representation to equivalent Keplerian orbital elements before input to the modified Q-law algorithm; this is done by using the transformation procedure described in Appendix B.

Constraints

Given this solution approach, the optimization problem in Equation (2.28) is subject to certain constraints and bounds. The first constraint regards the relationship between the decision variable a_T and other parameters that definitively specify the initial parking orbit through the Δv budget.

Combining Equations (2.29)-(2.30), the constraint translates to:

$$\begin{gathered} f_1(\boldsymbol{\chi}) = b_{C_1} \ , \\ \text{where} \ \ f_1(\boldsymbol{\chi}) = f(a_T) = \sqrt{(v_p^{\text{park}})^2 + (v_p^{\text{LEO}})^2 - 2v_p^{\text{park}} v_p^{\text{LEO}} \cos \Delta i} \ , \\ \text{and} \ \ b_{C_1} = \Delta \text{v}_{\text{budget}} = v_p^{\text{GTO}} - v_p^{\text{LEO}}. \end{gathered} \tag{2.31}$$

The underlying assumption here is that the entire budget available to the launch vehicle is exhausted in achieving the parking orbit, prior to commencing the low-thrust phase. This allows for a more equitable calculation of propellant mass consumed in different candidate solutions. The second constraint considers the time of flight restriction and is expressed as:

$$\begin{gathered} f_2(\boldsymbol{\chi}) \leq b_{C_2} \ , \\ \text{where} \ \ f_2(\boldsymbol{\chi}) = f_{\text{Q-law}} \left(a_T, \frac{W_a}{W_i}, \frac{W_e}{W_i}, \eta_{\text{cut}} \right) + |\tau_m| \ , \\ \text{and} \ \ b_{C_2} = \text{TOF}_{\text{max}}. \end{gathered} \tag{2.32}$$

The function $f_{\text{Q-law}}$ produces the time of flight required to complete the orbit raising spiral and is calculated simultaneously as the objective function (i.e., $\Delta m_{\text{propellant}}$) by using Q-law. The time spent coasting on the manifold in order to complete the second phase of the transfer is simply equal to the absolute value of τ_m. This τ_m takes a negative value when used to parameterize the stable manifold because the stable manifold is obtained by propagating the CR3BP system equations backward in time. In addition to these constraints, the search space of the decision variables is restricted by means of upper and lower bounds. These bounds are problem specific and their selection for the current sample case is outlined in Section 2.5.

Selection of Optimization Algorithm for LELT Trajectory Design

Within the current problem setup, the search space is composed of six continuous dimensions and thus attempting all configurations is computationally challenging. Stochastic optimization algorithms are well suited to explore the solution space for the current problem in an automated manner without much guidance from the user. In previous works that investigated the applications

of different stochastic optimization techniques, the general consensus is that the choice must align with the problem structure for effectiveness. In the current work, we have chosen the particle swarm algorithm due to its compatibility with the current problem structure where the search space associated with each of the decision variables is continuous, as well as its higher speed of convergence* [51]. Although the PSO is chosen in this work, it is expected that other algorithms such as differential evolution can also be applied successfully.

The PSO algorithm, originally introduced by Kennedy and Eberhart [52], is driven by cooperation between particles in a swarm by exchange of information about the best-found-solutions. This guides the motion of the initial population towards the optimal point. Beginning with a randomly generated initial random swarm of N particles, $\boldsymbol{\chi}_j^0$ where $j = 1, 2, ...N$, the new position of the particle j at the $(k+1)^{\text{th}}$ iteration is calculated as:

$$\begin{aligned} \boldsymbol{\chi}_j^{k+1} &= \boldsymbol{\chi}_j^k + \boldsymbol{\upsilon}_j^{k+1} \ , \\ \text{where } \boldsymbol{\upsilon}_j^{k+1} &= C_i \boldsymbol{\upsilon}_j^k + C_c R_1 (\boldsymbol{\mathcal{P}}_j - \boldsymbol{\chi}_j^k) + C_s R_2 (\boldsymbol{\mathcal{G}}^k - \boldsymbol{\chi}_j^k) \ . \end{aligned} \tag{2.33}$$

The constants C_i, C_c and C_s are the inertial, cognitive and social weighting coefficients respectively. The fitness of all particles is evaluated at each iteration to determine the best position of each particle over time, $\boldsymbol{\mathcal{P}}_j$, and the best global value in the current swarm, $\boldsymbol{\mathcal{G}}^k$. R_1 and R_2 are randomly generated numbers that lie between 0 and 1, selected from a uniform distribution. If a particle strays outside the specified boundary, the algorithm handles it by setting those components equal to the corresponding bound. The algorithm terminates either when maximum number of iterations N_{IT} is reached, or when the solution has not improved (i.e., has stalled) for N_{ST} iterations.

Method of Treating Constraints

Two constraints are present in the current problem, as laid out in Equations (2.28)-(2.32). Constraints can be treated within the PSO algorithm using the "penalty" method [53, 54], where penalties commensurate with the constraint violations are added to the candidate's fitness function evaluation. However, finding ways to reduce complexity of implementation by exploiting problem structure can benefit the computational resource consumption.

*vs. differential evolution and simulated annealing, when applied to a general class of benchmark problems

In the problem under consideration, the first constraint is a linear equality that deals with controlling the initial parking orbit where the low-thrust journey begins. Due to the structure of the problem, the linear equality constraint in Equation (2.31) can be interpreted as placing bounds on the value of a_T instead. Its lower bound is set at $a_T = a_{\text{LEO}}$, which corresponds to the limiting case where the launch vehicle expends its entire budget to make only an inclination change. The upper bound is calculated in an opposite manner by setting $\Delta i = 0$ in Equation (2.30), resulting in $a_T = a_{\text{GTO}}$. This implementation allows the use of the simpler unconstrained (but bounded) versions of optimization algorithms in the absence of any other constraints.

On the other hand, the inequality constraint in Equation (2.32) restricts the time of flight. Two methods were considered here to handle this constraint, (1) the previously mentioned penalty method, (2) a simple strategy of assigning a fictitious high cost to infeasible candidate solutions (similar to the strategy used in [55] to treat inequality constraints). In trials, the constrained optimization method by employing penalties to feasibilize candidates (i.e., method #1) converged to final solutions in less number of iterations than the second strategy. However, the overall computational runtime was longer because the optimizer needs to evaluate very long Q-law flight time solutions (which are likely infeasible) during the optimization. The simpler approach (method #2) worked well during the course of this work due to low computational runtimes from bounding the Q-law flight times. Thus, candidate solutions that exceed the time of flight constraint are simply marked within the fitness function to consume all of the available mass.

2.5 Problem Setup for QLEPSO Demo

The LELT trajectory design process described in the previous section is demonstrated on a transfer that starts at a geocentric parking orbit and delivers the spacecraft to an EML_1 halo orbit. The PSO algorithm used in this work is the native `particleswarm` available in MATLAB [56]. Progression of this algorithm can be controlled by setting different values for C_c, C_i and C_s, as used in Equation (2.33) for updating velocity of swarm particles at each iteration. The values used in this study are charted out in Table 2.1, along with the conditions for terminating the PSO run. All requisite integration is conducted in corresponding nondimensionalized units within the cost function, with the intent of properly scaling and conditioning the numerical integration process. Three-body

(CR3BP) nondimensional units are used for halo and manifold generation (where 1DU = distance between Earth and moon), and two-body nondimensional units used within Q-law implementation (where 1DU = Earth radius).

Table 2.1: PSO algorithm parameters.

Parameter	Value
Inertial weight C_i	[0.1, 1.1]
Cognitive weight C_c	1.49
Social weight C_s	1.49
Number of particles N	200
Maximum number of iterations allowed N_{IT}	200
Number of stall iterations allowed N_{ST}	20

Table 2.2: Q-law algorithm parameters.

Parameter	Value
Fixed time step	1000 s
Integration scheme	RK6 [57]
Target convergence radius for a	10 km
Target convergence radius for e	0.005
Target convergence radius for i	0.005 rad
Weight on inclination change W_i	1

The LELT trajectory is assembled using the values of decision variables selected by the PSO. The low-thrust engine used to affect this trajectory is assumed to provide a constant thrust of 0.7 N and a fixed specific impulse of 3000 s. The final mass delivered to the halo orbit is fixed at 1000 kg. This means that the spacecraft's *final* thrust acceleration is $7\times10^{-4}\text{m/s}^2$. The low-thrust trajectory is designed using the Q-law algorithm with the parameters specified in Table 2.2. No minimum periapse constraint is enforced and so W_p is set to zero.

Table 2.3: Bounds on decision variables.

Parameter	Lower bound	Upper bound
a_T	6563 km	24510 km
W_a/W_i	0.001	1000
W_e/W_i	0.001	1000
η_{cut}	0	0.5
τ_h	0	1
τ_m	-3π TU	-1π TU

The PSO algorithm also requires definition of the search space with respect to each component of the decision vector ($\boldsymbol{\chi}$ in Equation (2.28)). In order to do so, suitable bounds are defined on each decision variable as specified in Table 2.3. These bounds are either derived from the problem structure or assumed so as to confine the search space to a finite region. The first decision variable is the semi-major axis of the initial Earth parking orbit (a_T). Its perigee is fixed at a LEO-like altitude of 400 km (h_p), while its apogee and inclination are free to vary. These two parameters are however constrained by a Δv budget equal to that required for a transfer between LEO and a GTO-sized orbit at perigee with an impulsive maneuver:

$$\Delta \mathrm{v}_{\mathrm{budget}} = v_p^{\mathrm{GTO}} - v_p^{\mathrm{LEO}} = 2.3988 \text{ km/s} \ . \tag{2.34}$$

Once the optimizer (i.e. the PSO) selects the semi-major axis a_T of the parking orbit, then its inclination can be calculated by substituting the above $\Delta \mathrm{v}_{\mathrm{budget}}$ in Equation (2.30). In the current implementation, this constraint is instead expressed as a bound on a_T as in Table 2.3 – the lower bound is set at the orbital size that can be achieved if the entire Δv budget is used to change the inclination (i.e., lower bound set at LEO), while the upper bound is set where the available budget is used completely for apogee raising only (i.e., GTO). As described earlier, this calculation is based on the assumption that this Δv budget is completely consumed to reach the parking orbit.

The inclination of the parking orbit is expressed with respect to the Earth-moon plane. When the parking orbit is allowed to vary freely, the angular separation from the nominally-inclined LEO imparted by the launch vehicle depends on how its fuel budget was consumed for raising the apogee. In certain comparative cases of later sections, the parking orbit is fixed at geosynchronous transfer orbit (GTO) (Table 2.4) in order to match conditions of the referenced studies. Assuming that the GTO lies in the Earth's equatorial plane, its specified inclination reflects the mean angular separation between the Earth's equatorial plane and the plane of the moon's orbit. The angular separation of these two planes varies between 18.28° and 28.58° [58]. Similarly, the GEO specifications are defined in Table 2.4 for later use. However, for the sample case used for demonstration of the QLEPSO method, the inclination is allowed to freely vary within the $\Delta \mathrm{v}_{\mathrm{budget}}$ constraint as explained earlier.

The next two variables in the decision vector are related to the performance of Q-law and

Table 2.4: GTO and GEO specifications.

Parameter	GTO	GEO
Semi-major axis a	24510 km	42164 km
Eccentricity e	0.7234	0.005 †
Inclination i	23.43°	23.43°
Argument of periapse ω	Free	Free
Longitude of ascending node Ω	Free	Free

† Nominal non-zero value specified to avoid singularity in Gauss variational equations

Table 2.5: EML_1 halo orbit specifications.

Parameter	Value
Amplitude A_z	+8000 km
Type	Northern
Period	11.924 days
Initial position	[0.823383, 0, 0.020812] DU
Initial velocity	[0, 0.133228, 0] DU/TU

control the rates at which the semi-major axis and eccentricity converge to their target values with respect to the convergence speed of inclination to its target value. Instead of focusing on the absolute values of these variables, it is sufficient to find their relative values, thereby controlling the relative magnitudes of the components of the instantaneous control vector. Reasonable bounds are chosen for these variables. The fourth decision variable (η_{cut}) decides the placement and duration of coast arcs within the escape spiral — its lower bound is set to zero, which denotes a continuous thrust escape spiral with no coasting. While the natural upper bound on this effectivity cutoff approaches unity (which involves a coast-only trajectory with flight time approaching infinity), the upper bound is set to 0.5 instead. This educated estimate improves the convergence properties of the PSO by guiding it to a feasible region faster.

The EML_1 halo orbit used in this work follows the specifications presented in Table 2.5. A perturbation ϵ of 5 km is used in Equation (2.12) to generate the Earth-approaching branch of the stable manifold. It is propagated backwards in time for 3π normalized time units (synodic TU), a time span that converts to little more than 40 days. The manifold surface exposed for selection of insertion points covers at least two closest approaches to Earth; the solution space is thus deemed large enough for practical engineering purposes. This setup decides the bounds on the

last two decision variables in $\boldsymbol{\chi}$. The penultimate variable τ_h (halo orbit insertion parameter) is already normalized and bound between 0 and 1. Since the values of the final decision variable, τ_m, are negative for the stable manifold, the lower bound on τ_m is the time at which the (backward) integration of the manifold is terminated (-3π here). The upper bound is placed where all the manifolds have approximately departed the lunar sphere of influence (SOI), which is approximately -1π in this case. The radius of this SOI, centered at the moon, is calculated as [39]:

$$r_{\text{SOI}} \approx a_{\,\text{☽}} \left(\frac{m_{\,\text{☽}}}{m_{\oplus}} \right)^{\frac{2}{5}} = 66183.85 \text{ km} \ .$$

The part of the manifold surface of interest in this sample problem is the section that lies beyond the lunar SOI; this represents the low-energy search space for the current case study. These bounds on τ_m represent the region where Q-law is effective, because it is derived using the two-body equations of motion (lunar gravity is added as third-body gravitational perturbation during propagation only).

Although the current sample case handles LELT transfer to the EML_1 halo, the use of this lunar SOI also aids design of transfers to EML_2 halo orbits, whose Earth-approaching stable manifold branch crosses the moon first. In the absence of this bound on τ_m, using the QLEPSO method to design such transfers may otherwise involve using Q-law derived in the Earth-centered frame too close to the moon. Additionally, such transfers would require stricter bounds to be derived and used on the variable τ_h because certain sections of the manifold surface originating at EML_2 halo orbits do not approach Earth after their lunar flyby.

2.6 Results for QLEPSO Demo

The PSO algorithm is run for five trials with 200 particles and the optimization results are organized in Table 2.6. For these runs, the total TOF is constrained to be less than or equal to 90 days. This value of maximum TOF is chosen so as to simplify the process of comparison with results from independent past studies in later sections. The stochastic nature of particle swarm optimization gives rise to variance in the results presented by each run. Table 2.6 summarizes the results from five identical trial runs of the PSO, beginning with randomly distributed particle positions and

velocities. This table proves that the variance of results is within a narrow and acceptable range.

Table 2.6: QLEPSO results for total TOF $\leq$ 90 days.

Parameter	Run #1	Run #2	Run #3	Run #4	Run #5
Raw values of decision variables					
a_T, km	24294.2907	24282.4559	24290.4436	24375.4808	24311.8243
W_a/W_i	588.4616	636.6509	830.1153	89.4069	212.4609
W_e/W_i	814.2662	836.2408	946.7234	122.6418	296.5848
η_{cut}	0.0771	0.0743	0.0668	0.0764	0.0773
τ_h	0.7343	0.7245	0.7377	0.7441	0.7335
τ_m, synodic TU	-1.000π	-1.002π	-1.0043π	-1.0017π	-1.000π
Qualitative properties of solutions					
Orbit raising TOF, days	76.3542	76.3542	76.2847	76.3310	76.3426
Manifold coast TOF, days	13.6427	13.6456	13.7014	13.6657	13.6432
Mass fraction consumed, %	7.7708	7.7769	7.7911	7.7607	7.7688

Some useful insights can be gathered from the results in Table 2.6. Tracing the trajectory again in the temporal order of phases, the results show that almost the entire Δv budget available to the launch vehicle is used to raise its initial parking orbit to the size of GTO, keeping its inclination (with respect to the Earth-moon plane) close to zero. This can be explained by two intuitive causes: (1) the inclination of all the points on the manifold surface lies within 3° of the Earth-moon plane, and (2) the low-thrust trajectory may be able to better satisfy the constraint on the flight time by starting at the parking orbit with the highest permitted apogee. Thus the optimal parking orbit chosen represents the upper limit placed on the a_T variable; this observation is consistent with that made by Kluever [59] when designing combined chemical-electric transfers to GEO.

The other noticeable feature is the range of values of the Q-law weights exhibited in the five runs; however, the ratio of W_a/W_i to W_e/W_i remains approximately constant ($\approx 0.7 - 0.8$). This can be explained by the Q-law formulation where the three components of the thrust vector are decided by the relative (not absolute) values of the weights. The trajectory corresponding to fourth trial run is presented in Figure 2.8. The regions that the Q-law decides to coast within the escape spiral are noticeable and occur at the apogees.

Another prominent feature of the QLEPSO solutions is the consistent selection of τ_m, which identifies the portion of the journey spent coasting on the manifold, to take a value of $\approx -1\pi$. This implies that the length of the manifold coast arc is minimal and the spiral escape arc delivers the spacecraft as close to the halo orbit as allowed by the problem setup (refer Table 2.3). By

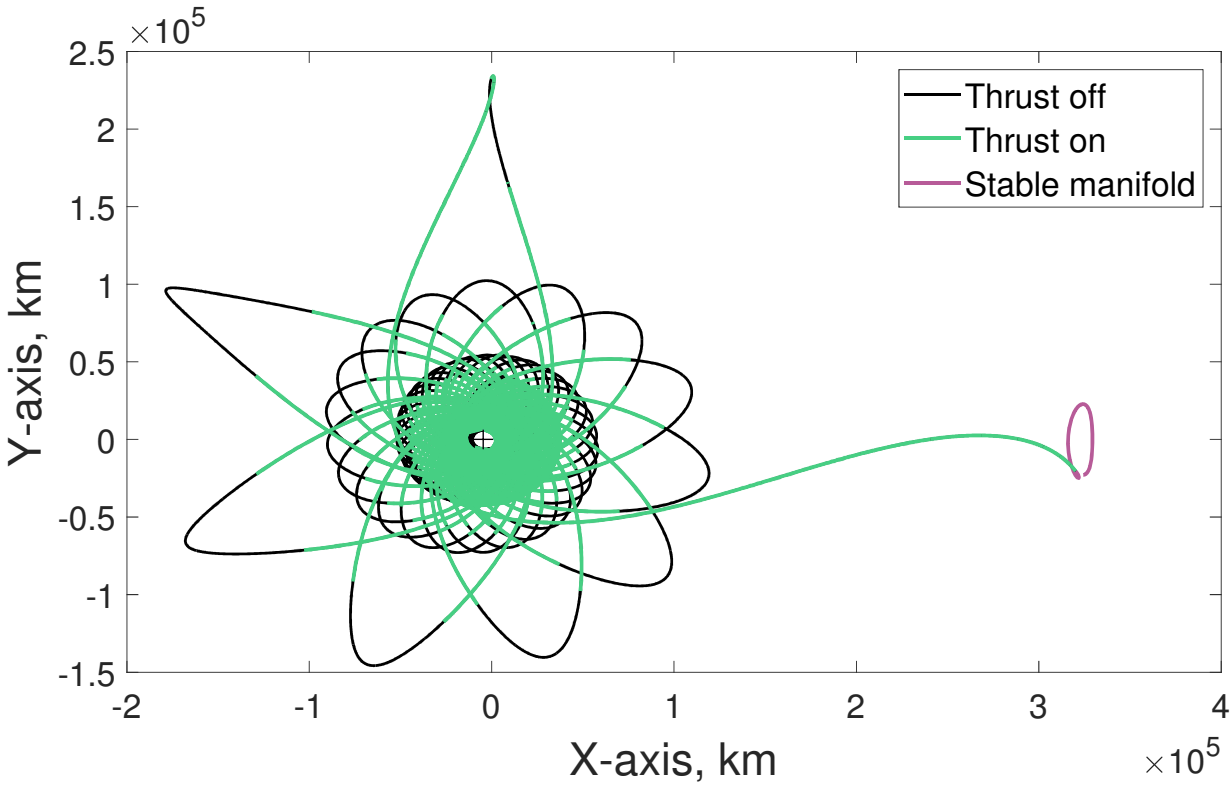

Figure 2.8: Complete QLEPSO trajectory.

minimizing the length of this manifold coast arc, the optimizer seeks to maximize the amount of coasting allowed within the orbit-raising portion of the transfer. These results are in contrast with a separate study that conducted LELT trajectory design using Q-law [60], where an attempt to enable large-scale tradespace exploration was made by heuristically deriving the optimal manifold patch point. This metric, referred to as Q_p, consisted of the weighted sum of relative "distance" between a point on the halo orbit's manifold and the Earth parking orbit:

$$\begin{aligned} Q_p &= \sqrt{\sum_{\overline{\alpha}} \overline{W}_{\alpha} \left[\frac{d(\overline{\alpha}, \overline{\alpha}_T)}{\overline{\alpha}}\right]^2} \quad \text{with } \overline{\alpha} = a, e, i \ , \\ \overline{W}_i &= \ \text{User defined (default} = 100) \ , \\ \overline{W}_e &= \frac{\overline{W}_i}{100} \ , \\ \overline{W}_a &= 0.2\overline{W}_i + \frac{20\overline{W}_i}{\sqrt{|\tau_m|}} \ . \end{aligned} \tag{2.35}$$

Selecting patch points with low values of Q_p reflects the fundamental directive of Q-law to reduce

the distance between the starting and destination orbits. As a general trend, this Q_p reduces in value as the stable manifold is propagated farther back in time; thus, the points on the section of the manifold surface far away from the halo orbit were selected as patch points in [60]. In the referenced study, the physical neighborhood of the heuristically chosen manifold patch point was then iteratively explored in order to find solutions that satisfied a given time-of-flight constraint, with the Q-law weights. However, the results from the current analysis imply that delivering the spacecraft as close to the halo orbit as allowable is the optimal strategy when constraining the total flight time. This allows the orbit-raising spiral arc to consume the majority of the coasting budget, thereby naturally reducing the propellant consumed for the transfer by leveraging the orbital dynamics instead. Though the analysis in [60] was conducted for transfer to an EML_2 halo orbit and the results here pertain to an EML_1 halo orbit, similar reasoning is expected to apply. The patch point that separates the low-thrust phase and the low-energy phase must be chosen with care and the current method can effectively do so, as shown by the trend within the optimal QLEPSO constrained-TOF solutions to pick manifold patch points close to the target halo orbit.

2.7 Comparison with Reference Methods/Solutions

In this section, the results from the QLEPSO method are compared with reference methods/solutions to quantify its accuracy for designing low-energy low-thrust transfers in the three-body system. Available methods in the literature have used different techniques to design LELT trajectories. Here, QLEPSO is first compared with other heuristic methods available in the literature (Section 2.7.1) and then with solutions obtained using deterministic gradient-based methods (Section 2.7.2). Favorable comparisons demonstrate that the proposed method significantly improves computational efficiency while achieving near-optimal solutions. This method can be used to support mission design analysis, for example to seed initial guesses for more accurate trajectory analysis tools, conduct large scale trade studies, or for approximating mission costs in space logistics tools [61].

2.7.1 Use of Tangential Thrust for Orbit Raising Arc

The first reference study chosen is provided by Abraham et.al. [62] where a method for preliminary mission analysis was developed and used to design a transfer from a GEO-energy orbit to a large-amplitude halo orbit. In this reference, the LELT trajectory is similarly constructed using a PSO, with the exception of the control law used to design the low-thrust phase. The escape spiral in [62] instead employs tangential control, where thrust acceleration is applied in the opposite direction as the spacecraft's velocity in the CR3BP frame. The three-body system equations for a powered spacecraft can be modified from Equation (2.1) as:

$$\begin{aligned} \ddot{x} - 2\dot{y} &= \frac{\partial \Psi}{\partial x} + \Gamma_x \ , \\ \ddot{y} + 2\dot{x} &= \frac{\partial \Psi}{\partial y} + \Gamma_y \ , \\ \ddot{z} &= \frac{\partial \Psi}{\partial z} + \Gamma_z \ , \end{aligned} \tag{2.36}$$

where $(\Gamma_x, \Gamma_y, \Gamma_z)$ are the components of the thrust control vector. Beginning at the manifold patch point, this escape spiral is propagated backward using Equation (2.36) by setting the control vector in the opposite direction as the velocity vector until an orbit matching the energy of GEO is achieved. The following cost function was used in [62] within a local PSO to determine solutions where the terminating GEO-energy orbit is also circular:

$$\mathcal{J}_{\text{abraham}} = k_1(e_{\text{terminate}} - e_{\text{desired}}) + k_2\Delta t + k_3\Delta m \ , \tag{2.37}$$

where Δt is the total time of flight and Δm is the propellant consumed. $e_{\text{terminate}}$ is the osculating eccentricity at the termination of the backward-propagated thrust arc, while its desired value, e_{desired}, is zero (i.e., eccentricity of GEO). A process for selecting values of the constants k_1, k_2, k_3 depending on the expected values of the three terms in the cost function is detailed in the original work. The values used here are $k_1 = 1$, $k_2 = 0.002$ and $k_3 = 0.001$.

For this analysis, the QLEPSO method outlined here and the tangential thrust method are applied first to the halo orbit chosen in [62], which is an EML_1 halo of $A_z \approx 72000$ km. To begin, the method in the reference is replicated and tangential control law is used for the orbit raising spiral. Next, the same transfer is calculated using QLEPSO, where the GEO specified in Table 2.4

is used as the desired parking orbit. The decision variables in this case are reduced to only those identifying the manifold patch point (τ_h and τ_m) and η_{cut} is set to zero to simulate continuous thrust trajectories. The parking orbit (a_T) is also fixed. The value of GEO-energy, unreported in the original reference, is assumed to be 9.6 [63]. The comparison between the two cases is shown in Table 2.7. For each case calculated here, the best of ten QLEPSO runs is reported – five for stable plus manifold, five for stable minus manifold, both of which approach Earth for the case of this large-amplitude halo orbit.

Table 2.7: Transfer to EML_1 halo orbit with $A_z = 72000$ km.

	Optimal cost (Equation 2.37)	$e_{\text{terminate}}$	Total TOF	Mass fraction consumed
Solution reported in [62]	0.1201	0.0019	≈116.077 days (26.73 TU)	6.0733%
Method used in [62]	0.1194	∼0	116.767 days	6.1619%
QLEPSO	0.1286	0.0051	113.106 days	6.6692%

The optimal cost reported in the first row of Table 2.7 is taken from [62] and is included here to corroborate the current implementation of the tangential thrust control law. Differences between the first two rows of Table 2.7 may arise due to inconsistencies in the value of GEO-energy or perturbation ϵ used for manifold generation. Another contributing factor could be the removal of the condition from [62] that τ_h take only certain discrete values. The QLEPSO solution also consumes more fuel due to the substantial change in inclination commanded in the problem setup, instead of the incidental change in inclination that can be achieved using tangential control.

Despite accounting for these factors, the results in the last two rows of Table 2.7 seem to indicate that tangential control law is more suitable for designing the LELT journey to the selected halo orbit. In order to demonstrate the real advantage of using Q-law over tangential thrust, this analysis is then repeated for a smaller EML_1 halo orbit of $A_z = 8000$ km that is defined in Table 2.5. The results of this exercise are presented in Table 2.8.

Table 2.8: Transfer to EML_1 halo orbit with $A_z = 8000$ km.

	Optimal cost (Equation 2.37)	$e_{\text{terminate}}$	Total TOF	Mass fraction consumed
Method used in [62]	0.4019	0.3128	60.4443 days	5.7770%
QLEPSO	0.1565	0.0068	93.6938 days	9.6227%

When designing transfers to the smaller halo orbit, the use of tangential control is unable to circularize the orbit within the time required to achieve GEO-energy. The desired outcome of this trajectory design is to bring the eccentricity close to zero; the tangential control method is only able to achieve an eccentricity of 0.3128 using the objective function in Equation (2.37), whereas QLEPSO was able to bring it all the way to 0.0068, which is an almost circular orbit. This is in addition to accomplishing the desired inclination change. Though the time of flight and fuel consumed are smaller in the case where tangential thrusting is used, it failed to produce feasible solutions where the essential condition regarding the parking orbit's eccentricity was satisfied. These results underscore the practicality of using QLEPSO when designing more general LELT transfers. Furthermore, using the method developed in this research allows for targeting of arbitrary parking orbits instead of only low-eccentricity, low-inclination orbits.

2.7.2 Use of Gradient-Based Methods

The accuracy of the LELT design method developed in the current work is quantified in this section by comparing our results with those obtained by using gradient-based methods. Apart from conveying the accuracy of solutions, this exercise will also demonstrate the advantage of QLEPSO in terms of computational efficiency. The first point of comparison is provided by previous studies that searched for locally-minimum LELT point solutions while minimizing a single objective function. This discussion is then concluded by showing comparable performance with a hybrid optimal control method that is capable of minimizing competing objectives.

Local Search: Minimum-Time Solutions

Mingotti et al. [22] conducted one of the earliest studies that applied optimal control techniques to design (fixed I_{sp}) low-thrust transfers to halo orbits in the Earth-moon CR3B system via their manifolds. This early study ignored the inclination between GTO and the Earth-moon plane. Tangential thrusting was used for a majority of the thrust arc under the assumption that it allows for maximal increase in instantaneous energy. The simplifying assumptions used in [22] were removed by Martin et al. [28, 43], where the control history along the entire thrust arc was optimized and the inclination of initial parking orbit (i.e., GTO) was assumed to lie between 18.45° and

28.45°. Throughout these comparisons, the applicable caveat is that the reported value of thrust acceleration used in the current simulations is its *final* value (i.e., at delivery to the halo orbit) due to its inherent nature of backward propagation, while the referenced studies use the same value as their *initial* thrust acceleration. Table 2.9 catalogs the comparison of our results with these past analyses where a similar trajectory was designed from GTO to EML_1 halo orbit of $A_z = 8000$ km. These QLEPSO runs are conducted by eliminating a_T as a design variable and instead fixing the initial parking orbit at GTO. The maximum allowable total time-of-flight is set at 90 days.

The method developed during the course of this research generates solutions with mass delivery fractions within 3% of the optimal solutions found through deterministic gradient-based methods. The LELT trajectories in [22] and [28, 43] consist of the same phases as in the QLEPSO method, where the spacecraft's orbit is raised using low-thrust engines until it is delivered to the stable manifold. The difference, however, is that while the referenced studies use continuous thrust to reach the manifold, our method allows coast arcs within the orbit raising phase. Due to the intrinsic nature of low-thrust trajectories, allowing coast arcs within a transfer reduces fuel consumption at the cost of longer flight times. In this case however, a spacecraft being delivered to the manifold is bound to spend time coasting on it anyway; this coast time is reallocated by the QLEPSO solution to the orbit raising phase with the purpose of reducing fuel consumption. This reassignment of coast locations is the root cause of the low propellant mass consumption in the transfer designed using QLEPSO.

Table 2.9: Comparison with solutions from literature.

	Orbit raising TOF (days)	Coast TOF on manifold (days)	Total TOF (days)	Mass fraction consumed
Mingotti et al. [22]	Unreported	Unreported	91.5	8.92%
Martin et al. [28]	47.1023	41.9463	89.0486	9.68%
QLEPSO	76.343	13.649	89.992	8.35%‡

Results from the work of Zhang et al. [64] are also used here, in which minimum-time low-thrust transfers for 17 different thrust acceleration values were designed to *directly* transfer to a selected point on the halo orbit without traversing its stable manifold. An attempt was made to design similar transfers using QLEPSO, by fixing $\tau_m = 0$ and $\eta_{\text{cut}} = 0$ within the solution process but

‡As a clarifying note, this value is higher than the result in Table 2.6 due to the high inclination of the parking orbit used here.

allowing $\frac{W_a}{W_i}$, $\frac{W_e}{W_i}$ and τ_h to vary. These results are mapped in Figure 2.9. An additional data set tagged as 'QLE' is plotted in the figure – it represents the design of the same minimum-time LELT transfers as [64] with the Q-law weights nominally set to unity and the halo insertion point fixed (i.e., no particle swarm optimization involved). This QLE data illustrates the effect of tuning the problem parameters within QLEPSO.

In the referenced study, the boundaries of the transfer are fixed, i.e. the transfer commences at the perigee of a fixed GTO and concludes at the point where the halo orbit achieves its maximum out-of-plane amplitude. On the other hand, QLEPSO transfers can be initiated from any point on a minimally inclined GTO and the halo insertion parameter is also left free. The selected EML_1 halo orbit ($A_z = 8000$ km) lies entirely within the lunar SOI — QLEPSO will still produce feasible solutions within the boundaries of the lunar SOI, although it is expected that the control directions dictated by Q-law are less effective.

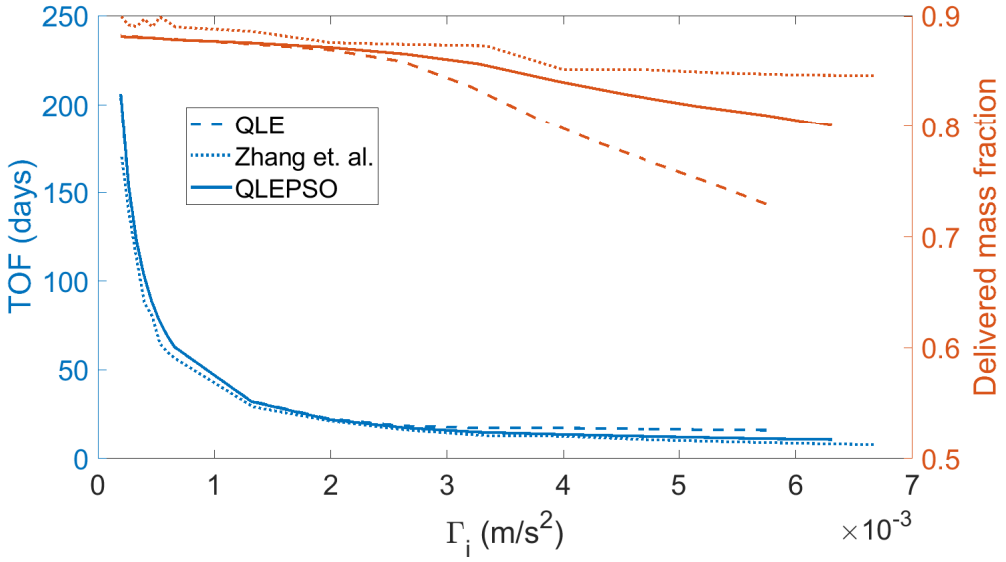

Figure 2.9: Comparison with minimum-time Zhang et. al. solutions [64].

Figure 2.9 shows that results obtained from both QLE and QLEPSO are close in terms of TOF to the optimal solutions in [64]. The deliverable mass fractions calculated, however, deviate from those of the optimal solutions at high values of initial thrust acceleration magnitudes (Γ_i). This will lead to more conservative estimates for propellent consumption in the regime of high Γ_i values.

One of the primary goals of this research project is to ensure that the method developed here can achieve equivalent performance as traditional preliminary mission design methods with a much

lower computational budget. A single QLEPSO run takes less than an hour to complete on a desktop computer with 4 cores running at 3.40GHz. On the other hand, the solutions generated by Martin et al. [28, 43] report a runtime on the "order of days on a desktop computer". Though runtimes are not reported in the other works referenced in Table 2.9, it is expected to be on the order of tens of hours for every final solution generated [§]. Also noteworthy is that using gradient-based methods usually involves generating and providing successive initial guesses through homotopy, like the process implemented in [43] and [64]. However, the comparison in Table 2.9 demonstrates that reasonable approximations to optimal trajectories can be obtained by using QLEPSO without much manual effort.

Local Search: Minimum-Fuel Solutions

In addition to applying gradient-based methods to finding minimum-*time* solutions for direct Earth-to-halo transfer, Zhang et al. [64] also continued the process to obtain minimum-*fuel* solutions for different thrust levels with specified flight times. The results of comparing QLEPSO results in this case are provided in Table 2.10, where the time-of-flight specified is obtained from [64] and serves as the constraint for our method. Specifically, this differs from the QLEPSO results in Table 2.9, where only transfers with 90 days flight time were calculated.

Table 2.10: Comparison with minimum-fuel solutions from Zhang et al. [64].

Thrust level	Total TOF	Mass fraction consumed in reference solutions [64]	Mass fraction consumed in QLEPSO solutions
0.6 N	140.2678 days	8.50%	8.13%
1 N	84.3688 days	8.69%	8.33%
2 N	37.8965 days	8.54%	8.69%
3 N	27.4773 days	9.12%	8.35%
5 N	16.8524 days	9.47%	8.95%

The propellant consumed in direct Earth-to-halo low-thrust trajectories designed using the QLEPSO method is lower than those found in [64]. This is due to the fact that the authors of the referenced work fixed the initial and final points of the transfer, i.e., the transfer commences at a GTO whose orientation (in terms of argument of periapse ω and longitude of ascending node Ω) is completely fixed, and ends at the point on the halo orbit where it achieves maximum

[§]estimated based on computational time of 28h reported by Mingotti et al. [24] for a comparable problem

separation from the Earth-moon plane. This is equivalent to fixing τ_h and τ_m in our formulation of the problem, as well as aiming for all orbital elements while designing the low-thrust orbit raising spiral. However, our Q-law implementation aims to definitively match only the semi-major axis, eccentricity and the inclination of the Earth parking orbit since the other elements can be targeted through the launch itself. Thus, the freedom of not targeting other orbital elements allows our method to produce better results. Regardless of matching the exact boundary conditions, the QLEPSO solution provides a close approximation of the transfer cost. As a point of note, QLEPSO failed to converge when constraining the total TOF to the values reported in [64] for higher values of thrust; acceptable solutions were found in these higher thrust cases when the time of flight constraint was relaxed.

Global Search Solutions

The hybrid optimal control framework developed by Shah et al. [65] provides a fully automated tool, now named DyLAN, for design of three-body transfer problems. It consists of a nested loop architecture to conduct global exploration of the search space associated with the problem. The outer loop uses a multiobjective genetic algorithm to pick the design variables that define the problem structure. This is complemented by the inner loop that solves the optimal control problem defined by the outer loop. The problem is defined to begin at the Earth parking orbit and thrust tangentially until a certain Jacobi energy is achieved. The spacecraft is then allowed one coast arc and a second thrust arc to reach the manifold patch point. The thrust arc beyond the main orbit-raising spiral is solved within the inner loop using the finite-burn low-thrust (FBLT) transcription. In this way, the boundary conditions of the optimal control problem supplied to the inner loop are decided by the outer loop. This hybrid optimal control framework is thus capable of generating the Pareto-front of time-of-flight versus propellant mass consumed.

The final comparison of the QLEPSO method is conducted against the results obtained by the DyLAN framework in [66]. The cost of these transfers, designed between a GTO and an EML_1 halo orbit of $A_z = 5000$ km, is charted in Figure 2.10. These reference trajectories correspond to a thrust acceleration of $1{\times}10^{-3}$ m/s^2. As before, due to the difference in direction of trajectory propagation, this thrust acceleration is the *final* value in the case of QLEPSO, and the *initial*

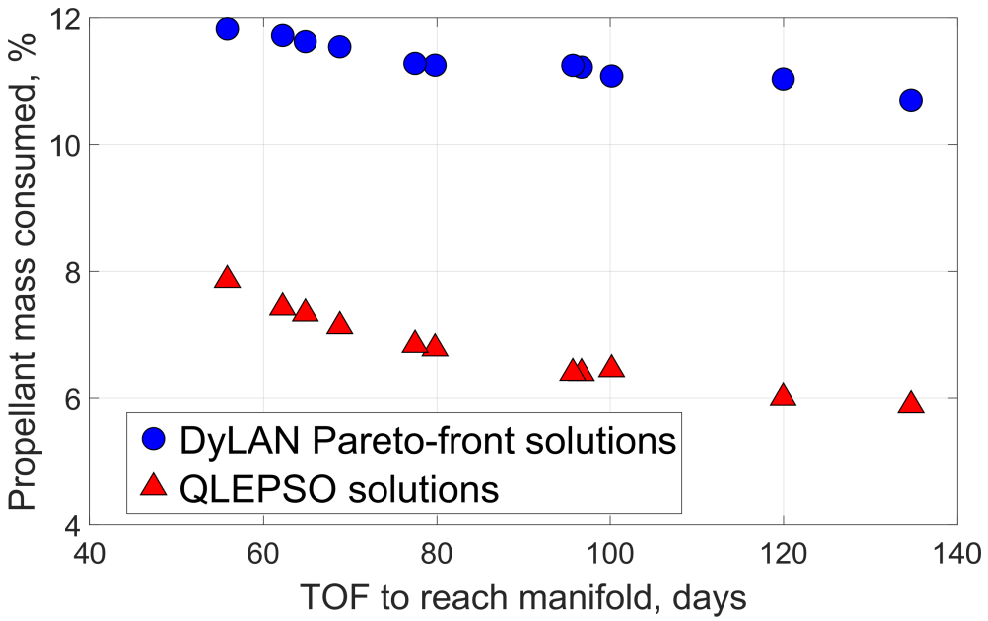

Figure 2.10: Comparison with DyLAN solutions from [66].

value for DyLAN. The Pareto-front of DyLAN solutions in Figure 2.10 was produced by using multiobjective optimization techniques. The QLEPSO data points, however, were obtained one-by-one by setting the TOF constraint to match the referenced results. An important distinction from the previous comparison cases is that only the time to reach the stable manifold is counted for flight time calculation. This is done in order to match the settings of the reported results in [66] and is achieved by excluding the $|\tau_m|$ term from the TOF constraint in Equation (2.32).

As explained earlier, the DyLAN solutions consist of a continuous tangential thrust arc starting at the parking orbit, followed by a coast arc and another thrust arc as decided by the outer loop solver. At the end of these phases, the spacecraft is delivered to the manifold. Thus the FBLT transcription is used only to optimize the second thrust arc. In Figure 2.10, the same trend of reallocation of coast arcs by QLEPSO is observed – the coast arc is removed from the pre-FBLT section and reassigned to the orbit-raising phase. Due to this, all QLEPSO-designed transfers consume less fuel than the DyLAN solutions.

2.8 Chapter Summary

This chapter describes a computationally-efficient design method for quantifying low-energy, low-thrust transfers in a three-body system in a computationally efficient manner. It uses the previously proposed Q-law, a feedback control law based on Lyapunov control theory, to design the low-thrust trajectory, combining it with low-energy dynamical systems theory. This low-energy low-thrust (LELT) trajectory design problem is parametrized by the initial parking orbit, Q-law weights and effectivity cutoff, and the manifold patch point. These decision variables are optimized using a particle swarm algorithm. A simple and effective strategy, comprising of marking infeasible solutions with fictitious high fitness values, is implemented here to deal with the time-of-flight constraint. The complete technique is demonstrated on the design of a transfer that starts at an Earth parking orbit and uses low-thrust propulsion to transport the spacecraft to the stable manifold of an Earth-moon L_1 halo orbit within a specified flight time. The solutions show that the insertion on the invariant manifold occurs at the closest allowed points to the halo orbit in order to spread the coasting time all along the low-thrust section of the trajectory. The parking orbit, constrained by the Δv to transfer from LEO to GTO, was also chosen to have the highest achievable apogee, with minimal inclination change.

Close agreement between the results obtained here and the optimal solutions obtained through gradient-based deterministic methods serves the purpose of validating and assessing the applicability of the developed preliminary Q-law LELT trajectory design with PSO (or simply QLEPSO) method. In addition to its accuracy, QLEPSO also provides significant savings in computational effort due to the use of the non-iterative Q-law for low-thrust arc design. By itself, the estimates provided by this method for percentage mass propellant consumed are suitable for large-scale trade space studies or logistics planning. This method is also well suited for use in the preparatory stages for gradient-based solving techniques (i.e., for seeding initial guesses); the use of QLEPSO for achieving higher-fidelity solutions in this manner is left open as an avenue of future work.

Material in this chapter has appeared in a paper authored by Bindu Jagannatha, Jean-Baptiste Bouvier, and Koki Ho in *Journal of Guidance, Control, and Dynamics* (2018) [67].

Chapter 3

Case Study: Cislunar Resupply Chain Design

Space exploration priorities in the US usually change with the political climate and in recent times have vacillated between Mars-first strategies, Moon-first [68, 69], and the in-between flexible path approach [70]. Each of these approaches requires substantial investment in development of new technologies, as well as deploying and maintaining space infrastructure. For this reason, growing the space economy in a sustainable manner will call for appropriate supply-chain and logistics management.

No matter which space exploration strategy is ultimately followed, NASA's journey to Mars is expected to go through near-term missions to the cislunar space [70, 71], which will allow human space programs to progressively fill knowledge gaps and prove newly developed technologies. Thus, the case study chosen to demonstrate the space logistics frameworks developed in this thesis is to optimally design a cislunar propellant resupply chain for supporting crewed Apollo-like missions to the lunar surface.

3.1 Problem Setting

The campaign considered as the case study in this work consists of repeating Apollo-style crewed missions for lunar surface exploration. For each mission, the crew stack (comprised of the Crew & Service Module, or CSM, and the Lunar Module, or LM) departs low-Earth orbit via a translunar injection maneuver provided by the Saturn V's upper stage. The CSM inserts the entire stack into lunar orbit once in the vicinity of the moon. Lunar ascent and descent operations are then conducted by the LM, which is discarded once the crew is returned to the CSM idling in low lunar orbit. Finally, the crew is returned to Earth by the trans-Earth injection maneuver provided by the CSM. In this scenario, the crew carries its entire fuel supply.

Instead, consider a scenario where these crewed missions can be refueled along their journey by in-space propellant depots. By moving towards distributed launch architectures, the fuel for the crewed spacecraft can be delivered in advance to certain way stations along the crew's path, along with cargo not required by the crew during the journey. The campaign will be composed of two stages – the cargo delivery phase, which then facilitates the crew mission phase. A fleet of cargo tugs, consisting of both high-thrust chemical propulsion (CP) and low-thrust solar electric propulsion (SEP) variants, deliver the crew resupply propellant to way stations in the form of disposable "droptanks". Once the first phase is complete, the crew mission/s can utilize these in-space propellant stockpiles.

This logistics problem can be optimized along multiple dimensions, three of which are: (a) locations of the way stations where the crewed spacecraft can be refueled, (b) amount of propellant supplied at each possible way station, and (c) the tugs selected for delivery of these propellant resupply tanks (assuming the crew vehicles are fixed). Figure 3.1 shows a sample result from such an optimization, where the expected campaign profile in terms of the mission sequence and the nodes to be used as pitstops are all part of the decision variables.

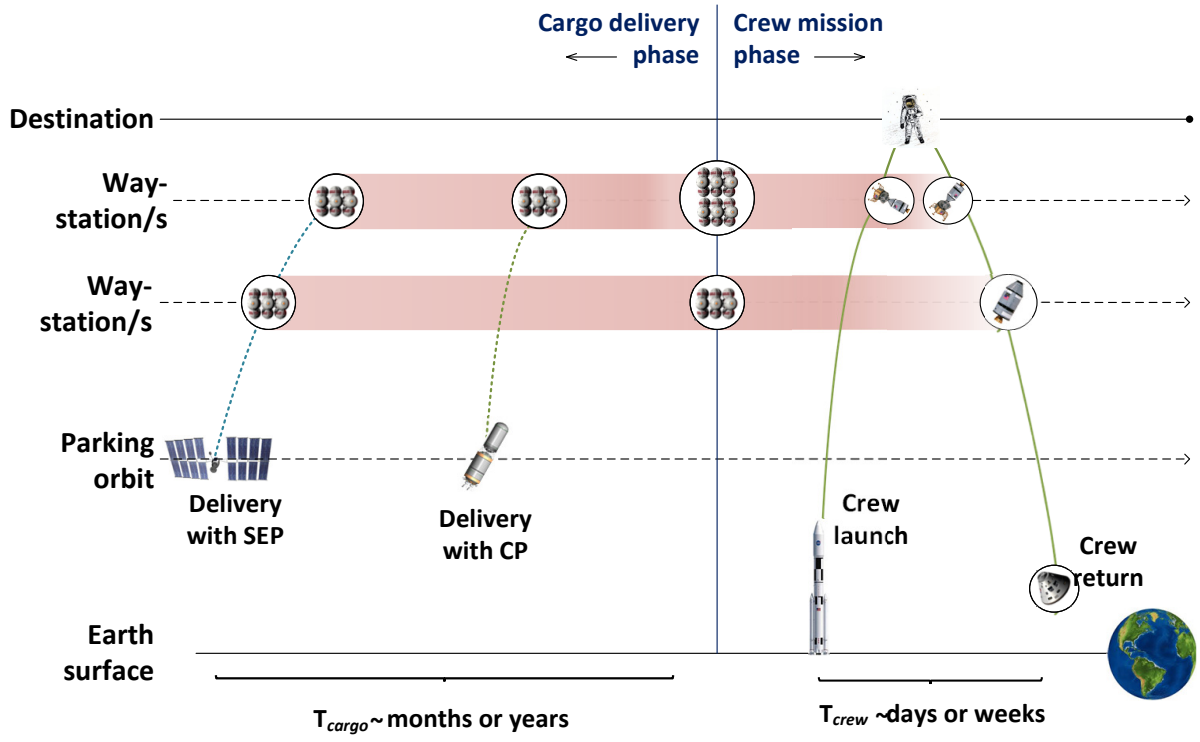

Figure 3.1: Expected campaign profile obtainable from logistics network optimization.

3.2 Network Nodes

The space exploration map is modeled as a network, with nodes corresponding to the candidates for surface destinations, orbits and staging locations. Arcs connecting these nodes represent flow of materials between them. The reduced network chosen in the current application case is sketched in Figure 3.2. Although the EML_2 Lagrange point lies on the farther side of the moon, its location in the figure reflects its use as a way station as opposed to a destination.

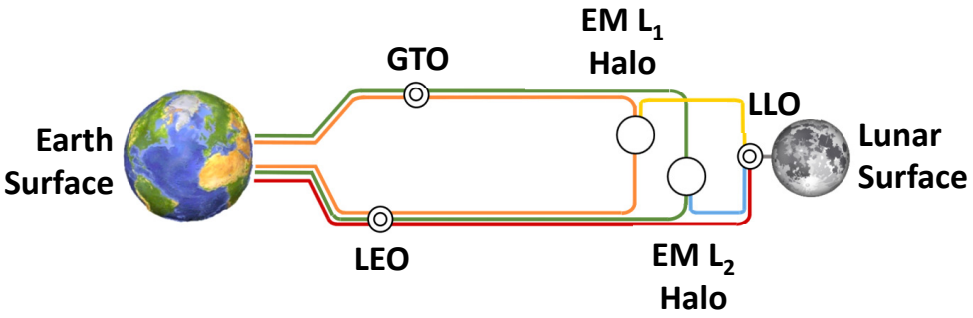

Figure 3.2: Network considered for cislunar resupply chain design.

Commodities flowing along this network include crew, vehicles and propellant. All commodities are launched into parking orbits around the Earth, where they begin their journey towards the moon as required. The parking orbits considered here are the two common launch orbits – the low Earth orbit (LEO) and the geosynchronous transfer orbit (GTO). Particularly within the cislunar space, propellant depots at locations such as the collinear Lagrange points have been suggested in previous architectures [10, 11, 72–75] to provide economic benefits for campaign design. Thus, nodes at Lagrange point halo orbits are considered as possible way stations for refueling the crew stack. These halo orbits have associated invariant manifolds that can be used to provide cheaper low-energy pathways (i.e., lower Δv, but longer flight time) for the tug journeys. The nodes marked in Figure 3.2 are:

1. Earth surface (ES).
2. Low Earth orbit (LEO) of radial altitude 200 km.
3. Geosynchronous Transfer Orbit (GTO) with perigee at 200 km altitude.
4. Halo orbit at EML_1 of 5000 km out-of-plane amplitude.
5. Halo orbit at EML_2 of 5000 km out-of-plane amplitude.

6. Low Lunar Orbit (LLO) of radial altitude ∼100 km.

The inclination of the body-centered orbital nodes (LEO, GTO and LLO) with respect to the Earth-moon plane is assumed to be negligible.

3.3 Vehicles

3.3.1 Crew Vehicles

As mentioned previously, high-thrust propulsion is preferable for in-space transport of crewed spacecraft. The Apollo missions are used as baseline in the current case study and thus use the Crew & Service Module (CSM) in conjunction with the Lunar Module (LM) for each crew sent to the lunar surface. The launch vehicle's upper stage (US) is used to impart the trans-lunar injection (TLI) required by the crew stack. While the crew vehicles that provide transport in space have fixed designs in terms of dry mass and fuel capacity, the upper stage of the crew launch vehicle is sized according to the fuel needed to provide the impulse demanded by the crew stack, i.e. commensurate to the magnitude of the TLI maneuver. This TLI will vary depending on whether the campaign-level architectural solution decided by the logistics framework dictates the first stop to be at LLO, at the EML_1 halo orbit or at the EML_2 halo orbit. Once this first impulse is completed by the upper stage, all succeeding maneuvers are conducted by the CSM or the LM (for lunar landing/ascent operations only). Hence, the crew only needs to be resupplied with either CSM fuel or with LM fuel at any in-space pitstop. The details regarding these crew vehicles are derived from the Apollo program itself and are listed in Table 3.1.

Table 3.1: Specifications of crew vehicles used in analysis.

Name	Dry mass, t	Propellant capacity, t	I_{sp} of fuel used, s	Structural coefficient of fuel used (ε)
Upper stage	Sized according to TLI	-	421	0.1138*
CSM	12.2	31	314	0.08 **
LM	5.8	12	311	0.08 **

* Calculated based on ∼14t dry mass of Saturn V upper stage and its 123t propellant capacity [76].
** Used to size the required droptanks.

3.3.2 Cargo Vehicles

In the case-study campaign, the droptanks are predeployed to their operational locations using a fleet of cargo tugs. The available cargo fleet consists of vehicles of discrete sizes, with propulsion systems of high-thrust (chemical propulsion) or low-thrust (solar-electric propulsion) type. This means that the dry mass of each tug and its fuel capacity is fixed. The amount of payload that can be carried by a tug across the network is limited only by its fuel capacity. The vehicle choices are specified in Tables 3.2–3.3. Additionally, multiple tugs of the same kind are allowed in the analysis. Each tug will also be permitted to be reused two times, i.e., a total of three uses per tug. A tug must return to its corresponding Earth parking orbit to be refueled/refurbished in order to be considered 'reused'. The advantage of reusing a tug comes from the reduction in launch cost, but the drawback is that the tug needs to carry its return fuel.

Table 3.2: Specifications of high-thrust chemical propulsion (CP) tugs used in analysis.

Type	Dry mass, t	Propellant capacity, t	I_{sp} of fuel used, s	Number of units
1 [77]	2.3	11.5	450	2
2 [78]	5.5	41	450	2
3 [79]	6.5	68	450	3

Table 3.3: Specifications of low-thrust solar-electric propulsion (SEP) tugs used in analysis.

Type	Dry mass, t	Propellant capacity, t	Power, kW	I_{sp} of fuel used, s	Number of units
4 [80]	3.5	11	40	3000	2
5 [81]	7.68	16	150	3000	2
6 [82]	10.7	39	300	2000	1

At this point, it should be noted that the in-space crew vehicles (CSM and LM) use high-density storable propellants and hence in-space storage of refuel propellant can be modeled using simple droptanks. Next generation crew transportation vehicles, like the Orion Multi-Purpose Crew Vehicle, are also expected to use similar hypergolic storable propellants [83]. On the other hand, the chemical propulsion tugs and the Saturn V's upper stage use high-I_{sp} cryogenic propellants; the use of this class of propellants is limited practically to applications where they need to be stored only briefly. Thus, deploying depots for CP tug propellant is not considered in the current problem

setting. Separate in-space storage of propellant for SEP tugs is also not considered.

3.4 Mission Cost Calculations

The total cost of the campaign is measured in terms of the initial mass to low Earth orbit (IMLEO). In the current problem context, an additional cost arises in the form of duration of each mission. This indirect cost regarding the campaign duration has two components – total time required to setup the in-space refuel supply chain using the cargo delivery tugs and the total crew flight time across all considered human missions. These two time-related pseudo-costs are also included in the analysis either as constraints or as objectives, so that the tradeoff between IMLEO and campaign duration can be explored. The campaign costs thus have three aspects – IMLEO, time required to complete predeployment of resupply propellant, and duration of crew missions.

The propellant resupply chain for the crew spacecraft is modeled in the form of disposable droptanks that are sized by their fuel storage requirements through the use of the fuel's structural coefficient (ε). In the current analysis, the costs for rendezvous between the crew spacecraft and the droptanks at way stations is ignored. Stationkeeping costs for maintaining the droptanks at their operational locations (i.e. at the halo orbits or at LLO) are also disregarded. Though not modeled here, the crew can be resupplied with other necessary commodities such as life-support supplies at these pitstops; thus, the extended crew flight times due to the need to transit through way stations for the purposes of refueling is assumed to not have any appreciable effect on the system-level design of the crew spacecraft.

All trajectories in this research are considered within the three-body model*, and some of them make use of the dynamical structures available in the Earth-moon system (refer Chapter 2). Crew transport is conducted using high-thrust chemical propulsion (i.e., Saturn V upper stage, CSM, and LM), but the delivery of the disposable droptanks containing the resupply propellant can be carried out by solar electric propulsion tugs or chemical propulsion tugs. All such cargo deliveries use the low-energy pathways associated with the chosen halo orbits. This means that the tugs cannot access LLO directly from their Earth parking orbits; instead they must transfer through either

*either directly through the circular restricted three-body (CR3B) model or through the two-body model with third-body gravitational perturbation

the EML_1 or EML_2 halo orbit. Another advantage of using the low-energy invariant manifolds is that this permits the same tug to drop off payloads at multiple locations on the same journey, if required. For example, a tug starting at its Earth parking orbit (either LEO or GTO) and going to LLO via the L_1 halo can drop off some of its cargo at the halo orbit and then continue its journey to lunar orbit.

The cost of transporting commodities across a network arc has two components: propellant used (or Δv) and time of flight (TOF). These costs depend on the propulsive element used on the arc. The transportation costs of using high-thrust propulsion can be practically modeled to be independent of the total mass because maneuvers associated with these "ballistic" trajectories are modeled as impulsive burns. However, the transfer costs of low-thrust trajectories are coupled to the spacecraft's thrust-to-mass ratio (or thrust acceleration) due to the continuous-burn nature of the associated propulsive maneuvers. Thus, the costs for employing solar electric propulsion have to be calculated internally; the method chosen for doing so is based on QLEPSO and is described later in this section.

3.4.1 Transfer Costs for Crew Transfers

Past studies have designed minimum Δv trajectories suitable for human missions from LEO to LLO, as well as to periodic orbits around the EML_1 and EML_2 [84–87]. In order to restrict the crew flight time to under safe limits, these transfers do not exploit the low-energy pathways, instead choosing to directly transfer onto the target orbits or to use lunar flybys. These crew-specific spaceflight routes also generally leverage Apollo-like free-return trajectories in order to enable safe aborts. The Δv costs for the different arcs that can be traversed by the crew vehicles are gathered from available literature and summarized in Table 3.4.

3.4.2 Transfer Costs for High-Thrust Tugs

Apart from the crew transportation costs, the other contribution to the total IMLEO cost arises from the launch of tugs and their corresponding cargo. Any resupply propellant (i.e. CSM or LM fuel) stored in space is assumed to be held in disposable droptanks, which is sized linearly with the amount of propellant it stores. The structural mass of these in-space droptanks is calculated using

Table 3.4: Crew transportation costs.

Arc	Δv, km/s	TOF, days	Impulse provided by
*Forward direction arcs***			
LEO to TLI*	3.306	-	Launch vehicle upper stage
TLI to LLO*	0.976	4	CSM
LEO to TL1I*	3.200	-	Launch vehicle upper stage
TL1I to L1 [86]	0.570	5	CSM
L1 to LLO [88]	0.750	3	CSM
LEO to TL2I*	3.400	-	Launch vehicle upper stage
TL2I to L2 [86]	0.275	8.5	CSM
L2 to LLO [88]	0.750	3.5	CSM
*Return direction arcs***			
LLO to ES*	1.091	3	CSM
LLO to L1 [88]	0.750	3	CSM
L1 to ES [86]	0.499	10	CSM
LLO to L2 [88]	0.750	3.5	CSM
L2 to ES [86]	0.275	8.5	CSM

TLI = Translunar injection, TL1I = Translunar injection to go to L_1 halo, TL2I = Translunar injection to go to L_2 halo.
* Approximated/ calculated from Apollo 17 values [76].
** Lunar descent Δv = 2.123 km/s and lunar ascent Δv = 2.239 km/s.

Table 3.5: High-thrust chemical propulsion tug transportation costs.

Arc	Δv, km/s	TOF, days
LEO to/from L1 [89]	3.375	21
L1 to/from LLO [90]	0.259	28
LEO to/from L2 [89]	3.336	17
L2 to/from LLO [90]	0.375	27

the respective fuel's (CSM or LM's) structural coefficient. When these droptanks are delivered to their in-space storage locations, the high-thrust chemical propulsion tugs use the cheaper (but slower) routes through the Earth-moon system. Unlike the crew transfers, the cargo transport using both chemical tugs and solar electric tugs take advantage of the halo orbits' invariant manifolds, instead of inserting onto them directly. The Δv data for cargo deliveries conducted using chemical propulsion tugs is provided in Table 3.5. These high-thrust tugs are assumed to begin their journey in LEO. All tugs can be reused during the campaign if their own propellant tanks are refilled; however, they must return to their Earth parking orbit (which is LEO for the case of high-thrust chemical tugs) to do so. In such cases, the tug will have to carry additional fuel for facilitating its return, apart from the cargo it was required to deliver on its forward journey.

3.4.3 Transfer Costs for Low-Thrust Tugs

The propellant mass fraction consumed by a low-thrust solar electric propulsion (SEP) tug and, more importantly, its flight time are heavily dependent on the ratio of the thrust it can provide and its total initial mass. Thus, the costs of low-thrust transfers cannot be simply be collated from existing literature, and must be calculated within the logistics framework.

The task of delivering cargo (i.e. droptanks) to various cislunar nodes in the network is optimally distributed among the tugs by the space logistics framework. During the course of this process, the thrust-to-mass ratio (i.e., the thrust acceleration) of any given SEP tug takes a range of values, even for the same transportation arc. As mentioned before, obtaining these low-thrust trajectories is a computationally intensive process involving design of the many-revolution thrust arc required to depart/arrive at the Earth/moon [21–25, 43]. Existing methods do not lend themselves well to quick parametric trade studies, as is the application case required here. Thus, a method based on our previous studies [66, 67] is used, where Q-law is combined with dynamical systems theory to design low-energy low-thrust (LELT) trajectories.

The method used in this section for LELT trajectory design is a simpler than the QLEPSO methodology explained in the previous chapter. This is due to one key difference – the low-thrust trajectories pertinent to the ultimate space logistics problem are assumed to be of continuous thrust in nature, i.e. no coast phases are allowed within the orbit raising/lowering spiral. The propellant consumption of low-thrust trajectories inherently decreases as more coast segments are allowed, and tends to zero when the total time of flight approaches infinity. Thus assuming continuous-burn trajectory provides an upper limit on the transfer costs. Additionally, the use of low-energy pathways (i.e. the invariant manifolds) reduces propellant consumption. Thus, the choice is made to model desired low-thrust cislunar transport as continuous burn spiral phase patched with a manifold coast phase.

The initial Earth parking orbit, where the low-thrust tugs begin their journeys, is chosen to be the GTO in compliance with the selection made by other previous studies which designed similar trajectories [22, 43]. The multi-revolution spiral orbital transfer phase is then designed using Q-law assuming constant maximum thrust until the boundary of the coast phase. Feedback control laws are known to be suboptimal, but they simplify the solution process and provide good estimates

for low-thrust transfer costs, thus making the method favorable to large-scale trade studies or for quick approximations such as the application here. In order to use Q-law with dynamical systems theory (i.e., the QLE method) to obtain the low-thrust arc costs, favorable points on the manifold must be first chosen as targets for the spiral section of the trajectory. These patch points on the manifold, which separate the continuous-thrust spiral arc from the coast arc, are chosen via a grid search [66] across three different values of thrust accelerations.

Manifold Patch Point Selection

Journeys ending at the halo orbit will require the *stable* manifold to be generated; both branches of this stable manifold are used for transfers to the EML_1 halo orbit, but only its Earth-approaching branch is used in the case of the EML_2 halo orbit. The same requirement applies to the unstable manifolds used for transfers away from the halo orbits and towards the primary bodies. The generated manifold is discretized in terms of τ_h and τ_m, and is fed into the Q-law algorithm as a viable patch point. The low-thrust trajectory can only be propagated starting from the patch point as the true anomaly is only known at this end of the powered phase; thus the Q-law controlled transfer has to be propagated either backward or forward in time as needed.

This grid search across three values of spacecraft thrust-acceleration (Γ) yields results as shown in Figures 3.4–3.10. In all cases, thruster specific impulse (I_{sp}) of 3000 s is used and the same payload of 1000 kg is being delivered to the manifold and hence to the halo. To search for trends in the grid search results, the evolution of the orbital elements of the manifold trajectories with respect to the primary bodies is also examined. Let us first look at the Earth-approaching branches of both $EML_1$and EML_2 halo orbit manifolds (Figure 3.4 and Figure 3.8). The behavior of the search space related to the Earth-bound stable manifolds seems to correlate with the evolution of the orbital elements of the manifold trajectories – Figures 3.5 and 3.9 at the end of this chapter present this data. Semi-major axis has been normalized with respect to Earth radius and the angles (inclination and true anomaly) are shown in degrees. Comparison of Figures 3.4(b) and 3.4(c) with Figure 3.5 indicates that the minimum fuel for delivery to the stable manifold at high thrust acceleration levels is achieved at apogee points along the manifold, with the minimum point occurring along the manifold trajectories that have the lowest semi-major axis. However, as the

thrust acceleration is decreased, the effect of the inclination starts dominating and the choice of the optimal manifold trajectory moves towards those with higher inclination (due to the high desired inclination of the parking orbit itself). This results in certain manifold trajectories becoming more optimal for insertion, rather than particular regions of the manifold, as evident from Figure 3.4(d). Similar conclusions can be drawn from comparing Figure 3.8 to Figure 3.9 (placed at the end of this chapter).

Though there exists some correlation between orbital specifications of the manifold patch points and the costs associated with transfers to/from these points on the Earth-bound manifold branches, there are no obvious trends in the grid search results for the moon-approaching manifold branches that remain valid across different thrust acceleration magnitudes.

The manifold patch points for the LELT journeys to and from the specified EML_1 and EML_2 halo orbits are thus chosen from these grid searches to reflect the points where minimum costs are incurred across the three different thrust acceleration values considered here. The effectivity of Q-law in designing LELT trajectories *without* a stochastic optimization algorithm serving as an outer loop has been shown briefly in the previous chapter (refer the QLE dataset in Figure 2.9). The manifold patch points chosen for each low-thrust arc permitted in the cislunar logistics network are marked in Figures 3.4(a), 3.6(a), 3.8(a) and 3.10(a).

Low-Thrust Cislunar Transfer Cost Model

Once the manifold patch points are chosen, the trajectories for various combinations of low-thrust tugs (from Table 3.3) and their payloads are designed for all permissible arcs. Both costs – final mass, which is indicative of propellant used, and TOF – are calculated for each tug, by starting at the tug's dry mass as the initial mass along the arc and increasing its payload mass in discrete steps up to the maximum value permissible by the tug's fuel capacity. The thrust magnitude that can be provided by each low-thrust tug is calculated by assuming a 60% efficiency in power transfer to the propulsion system. The results of this analysis are plotted in Figure 3.3. Both costs show an approximately-linear trend in this figure.

This linear behavior of flight time (TOF) with respect to initial mass is similar to the simplified case of very low continuous thrust transfer between two coplanar circular orbits of radii a_i and

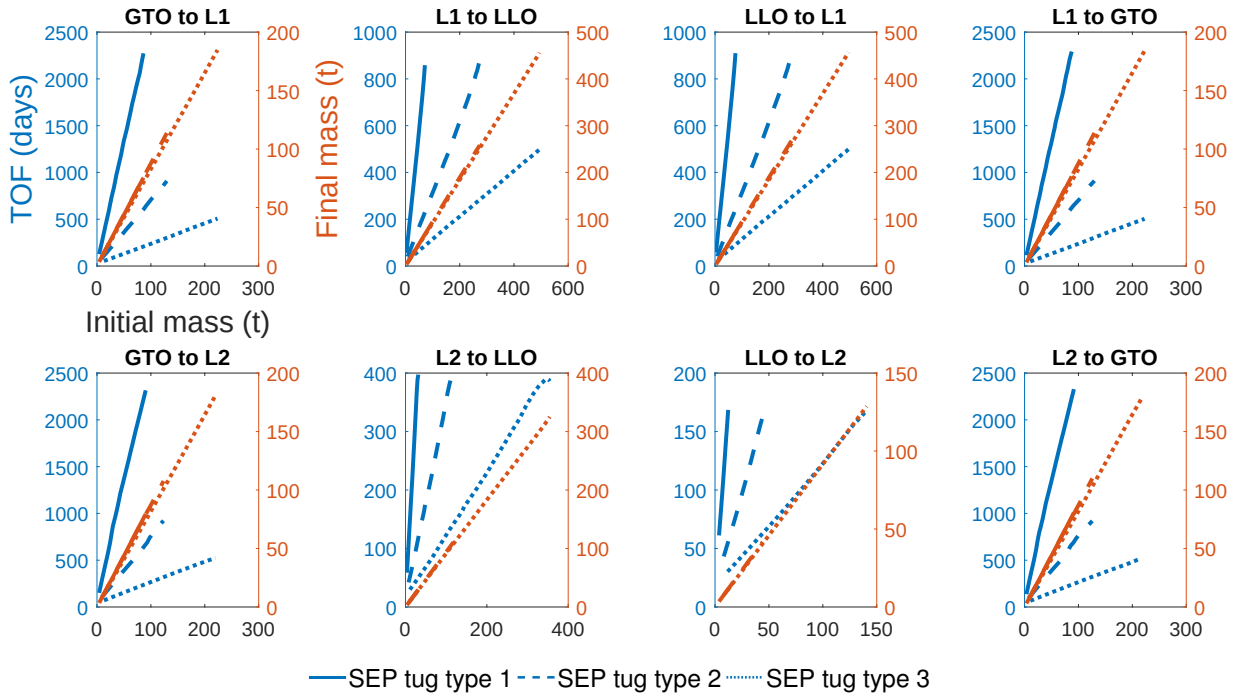

Figure 3.3: Final mass (in t) and TOF (in days) vs. initial mass (in t) for each SEP tug in Table 3.3.

a_f ($a_f > a_i$) using constant tangential thrust (i.e. thrust acceleration vector $\mathbf{\Gamma}$ aligned with spacecraft's velocity vector). In this case, the time of flight is inversely proportional to the thrust-acceleration [91]:

$$t_f = \frac{1}{\Gamma}\left(\sqrt{\frac{\mu}{a_i}} - \sqrt{\frac{\mu}{a_f}}\right) \quad ,$$

where thrust acceleration $\Gamma = \dfrac{\text{thrust}}{\text{mass}}$. This above relationship is derived assuming that the thrust and the mass are unchanged during the course of the transfer. The LELT transfers designed in the current analysis through the use of Q-law are similar transfers, where the orbital transfer is conducted between two points of fixed energies, i.e., each low-thrust transportation arc connects the nodes at GTO and the pre-selected manifold patch point. Thus the thrust-acceleration decreases as mass being carried is increased, which then drives up the time of flight in a nearly linear manner.

Returning to Figure 3.3, a linear curve is fit through the available data and the coefficients used later for modeling the transportation costs:

$$\mathrm{m_{final}} = {}^{1}p \cdot \mathrm{m_{initial}} + {}^{0}p \ , \tag{3.1a}$$

$$\mathrm{TOF} = {}^{1}q \cdot \mathrm{m_{initial}} + {}^{0}q \ , \tag{3.1b}$$

where p and q are constant coefficients. The initial mass on the arc includes the payload being carried by the spacecraft, its dry mass and the fuel onboard at the beginning of the transfer. The coefficients derived for each arc permitted for low-thrust transfers are listed in Tables 3.6 - 3.7.

Table 3.6: Linear fit coefficients for final mass (in t) vs. initial mass (in t).

Arc	SEP tug type 1		SEP tug type 2		SEP tug type 3	
	${}^{1}p$	${}^{0}p$	${}^{1}p$	${}^{0}p$	${}^{1}p$	${}^{0}p$
	To/ from EML_1 halo orbit					
GTO to L1	0.8757	-0.0038	0.8772	0.1202	0.8251	0.2401
L1 to LLO	0.9446	0.0446	0.9447	0.1586	0.9214	-0.0391
LLO to L1	0.9446	0.0429	0.9449	0.1398	0.9214	-0.0407
L1 to GTO	0.8758	-0.0082	0.8764	0.1453	0.8253	0.1808
	To/ from EML_2 halo orbit					
GTO to L2	0.8801	-0.2208	0.8709	0.1763	0.8219	-0.2458
L2 to LLO	0.9397	0.0338	0.9396	0.1337	0.9112	0.5247
LLO to L2	0.9394	0.0276	0.9402	0.0841	0.9128	0.2931
L2 to GTO	0.8818	-0.2571	0.8718	0.1172	0.8232	-0.3195

Table 3.7: Linear fit coefficients for time of flight (in days) vs. initial mass (in t).

Arc	SEP tug type 1		SEP tug type 2		SEP tug type 3	
	${}^{1}q$ ($\times 10$)	${}^{0}q$	${}^{1}q$	${}^{0}q$	${}^{1}q$	${}^{0}q$
	To/ from EML_1 halo orbit					
GTO to L1	2.598	26.631	6.832	19.146	2.164	22.865
L1 to LLO	1.156	8.567	3.074	8.403	0.973	17.465
LLO to L1	1.157	8.015	3.062	9.199	0.973	17.482
L1 to GTO	2.595	27.177	6.872	18.159	2.162	23.674
	To/ from EML_2 halo orbit					
GTO to L2	2.505	88.126	7.180	32.179	2.204	45.036
L2 to LLO	1.257	13.027	3.359	12.346	1.098	13.092
LLO to L2	1.266	13.364	3.329	14.458	1.079	15.511
L2 to GTO	2.476	91.578	7.128	35.899	2.188	45.804

Even though high-thrust propulsion TOF costs are fixed, similar coefficients can be listed for

corresponding arcs by simply setting ${}^1q$ to zero. These coefficients are used later within the mathematical model to relate the transfer times on arcs to the campaign duration, and thereby constrain it.

As a final point of note, one can notice that the initial Earth parking orbit for the low-thrust solar-electric propulsion tugs is GTO, while the cost concerns IMLEO mass. To account for this discrepancy, any mass launched to GTO is penalized to be about 1.74 times the initial mass in LEO; this penalty is derived from the ratio of initial mass to final mass if a rocket with fuel $I_{sp} = 450$ s was used to place a spacecraft in GTO when starting at a LEO of altitude 200 km, which requires $\Delta v \sim 2.4554$ km/s.

3.5 Chapter Summary

This chapter provides a comprehensive description of the case study setup as used to demonstrate the space logistics frameworks developed as the core of this research. The set of network nodes and vehicles that can be potentially utilized are presented. The cost models for both high-thrust and low-thrust transportation arcs relevant to the case study are also derived.

Parts of this chapter appear in two papers: (1) presented by B. Jagannatha, V. Shah, R. Beeson and K. Ho at the 2017 *AIAA/AAS Astrodynamics Specialist Conference* [66], and (2) presented by B. Jagannatha and K. Ho at the 2018 *AAS/AIAA Astrodynamics Specialist Conference* [92].

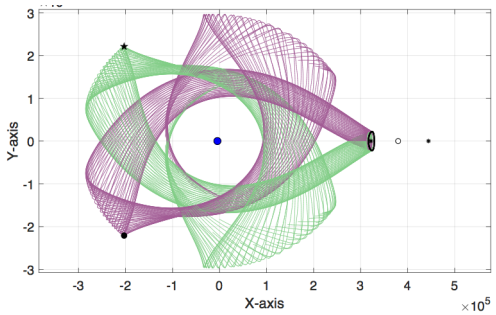

(a) Chosen manifold patch points

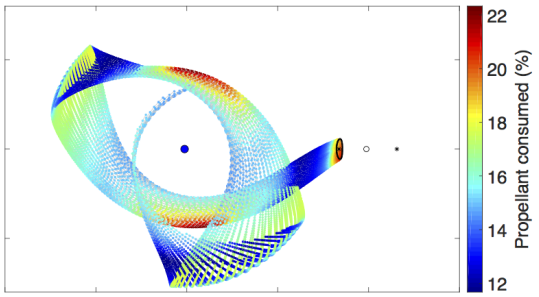

(b) $\Gamma_f = 5 \times 10^{-3}$ m/s^2

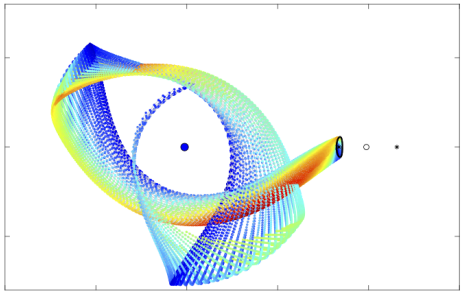

(c) $\Gamma_f = 5 \times 10^{-4}$ m/s^2

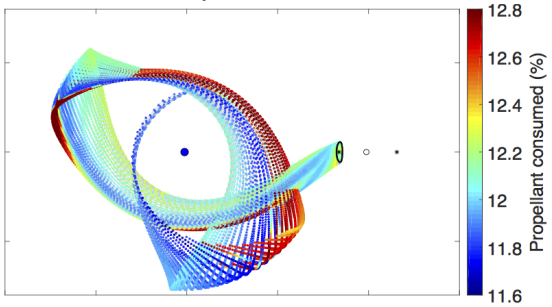

(d) $\Gamma_f = 5 \times 10^{-5}$ m/s^2

Figure 3.4: Propellant mass fraction consumption for low-thrust transfer from GTO to EML_1 halo's discretized stable manifold.

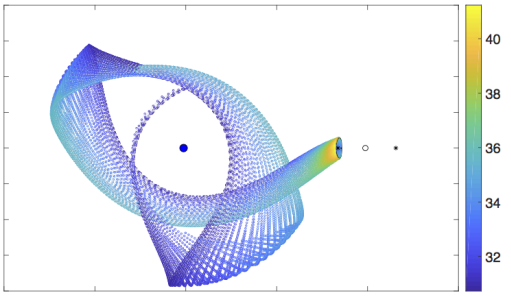

(a) Normalized semi-major axis ($R_\oplus$)

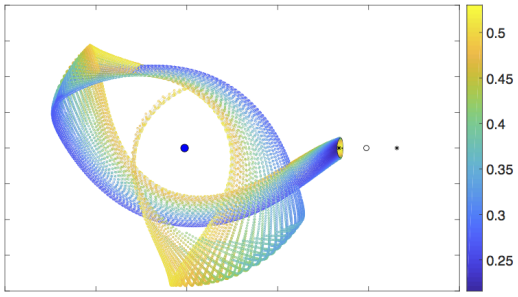

(b) Eccentricity

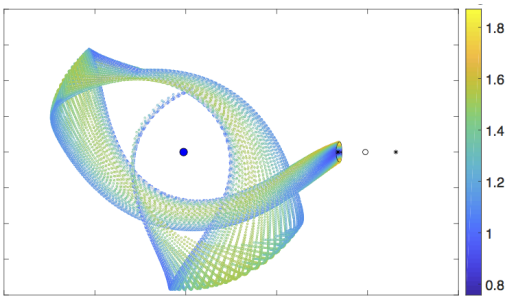

(c) Inclination (deg)

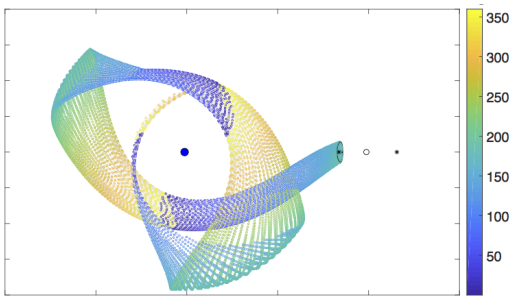

(d) True anomaly (deg)

Figure 3.5: Orbital elements of patch points on the discretized Earth-approaching branch of EML_1 halo's stable manifold.

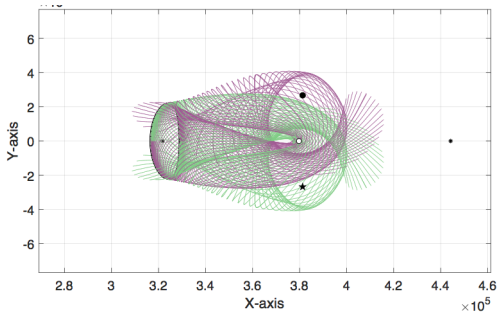

(a) Chosen manifold patch point

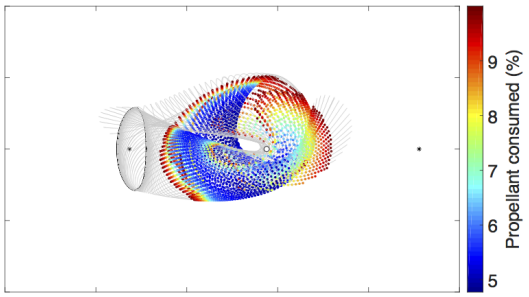

(b) $\Gamma_f = 5 \times 10^{-3}$ m/s^2

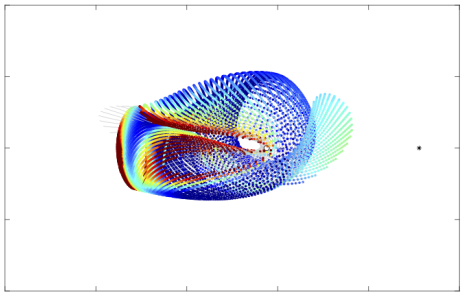

(c) $\Gamma_f = 5 \times 10^{-4}$ m/s^2

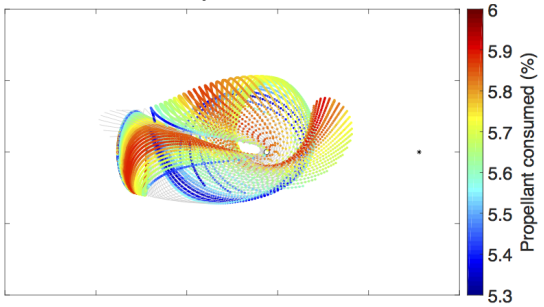

(d) $\Gamma_f = 5 \times 10^{-5}$ m/s^2

Figure 3.6: Propellant mass fraction consumption for low-thrust transfer from EML_1 halo's discretized unstable manifold to LLO.

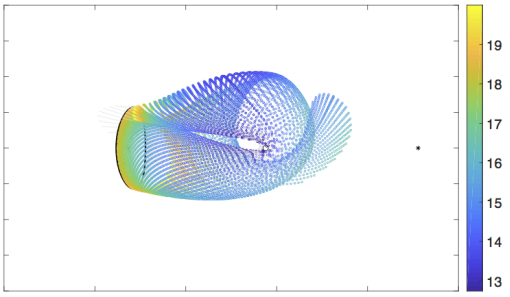

(a) Normalized semi-major axis ($R_☽$)

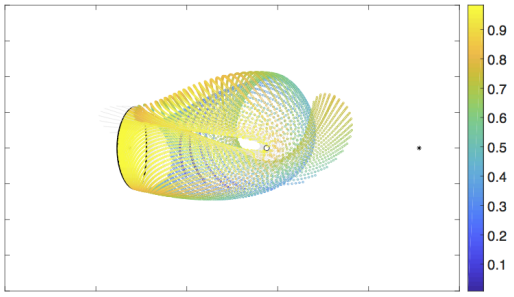

(b) Eccentricity

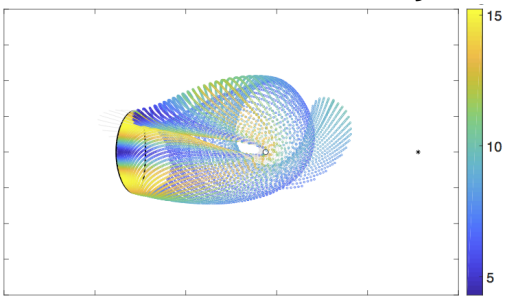

(c) Inclination (deg)

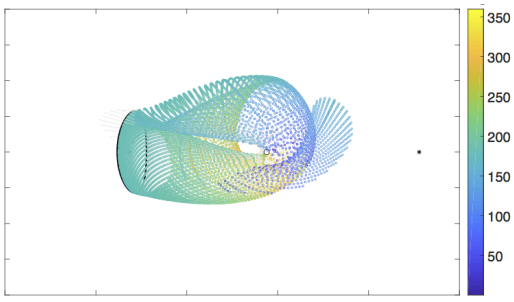

(d) True anomaly (deg)

Figure 3.7: Orbital elements of patch points on the discretized moon-approaching branch of EML_1 halo's unstable manifold.

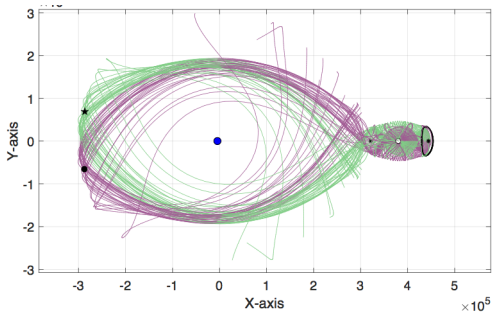

(a) Chosen manifold patch point

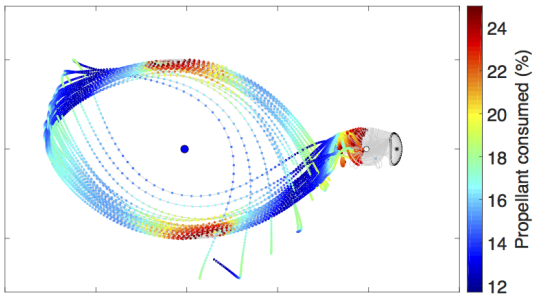

(b) $\Gamma_f = 5 \times 10^{-3}$ m/s^2

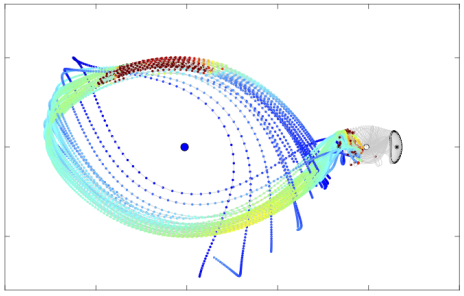

(c) $\Gamma_f = 5 \times 10^{-4}$ m/s^2

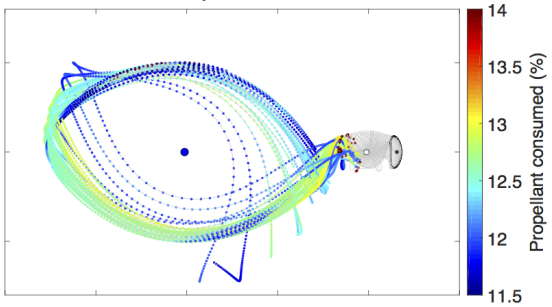

(d) $\Gamma_f = 5 \times 10^{-5}$ m/s^2

Figure 3.8: Propellant mass fraction consumption for low-thrust transfer from GTO to EML_2 halo's discretized stable manifold.

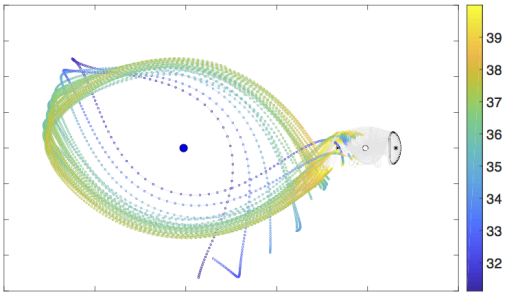

(a) Normalized semi-major axis ($R_\oplus$)

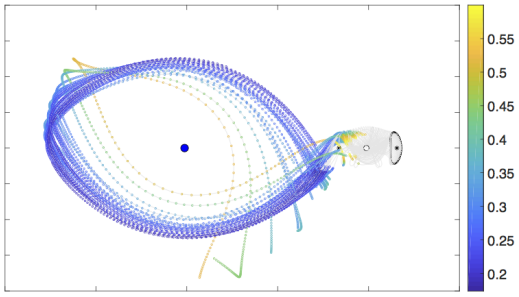

(b) Eccentricity

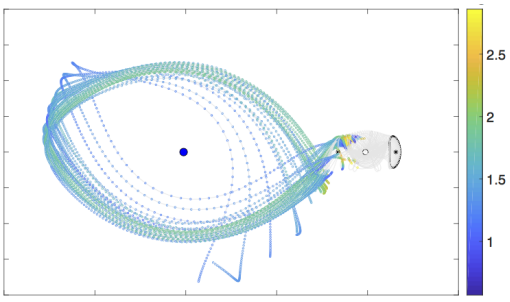

(c) Inclination (deg)

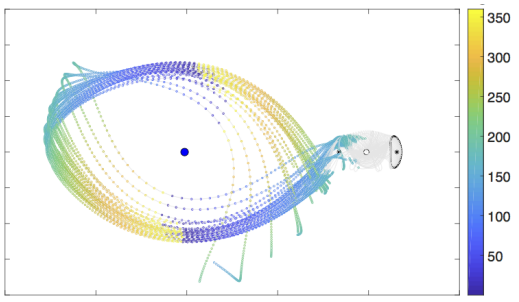

(d) True anomaly (deg)

Figure 3.9: Orbital elements of patch points on the discretized Earth-approaching branch of EML_2 halo's stable manifold.

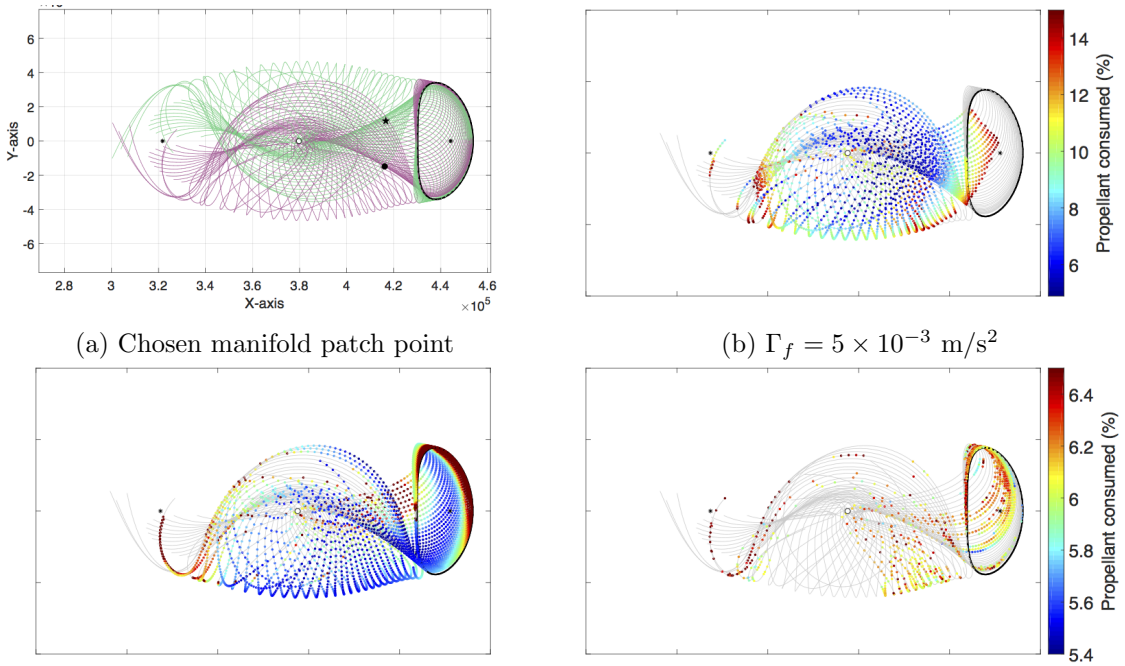

(a) Chosen manifold patch point

(b) $\Gamma_f = 5 \times 10^{-3}$ m/s^2

(c) $\Gamma_f = 5 \times 10^{-4}$ m/s^2

(d) $\Gamma_f = 5 \times 10^{-5}$ m/s^2

Figure 3.10: Propellant mass fraction consumption for low-thrust transfer from EML_2 halo's discretized unstable manifold to LLO.

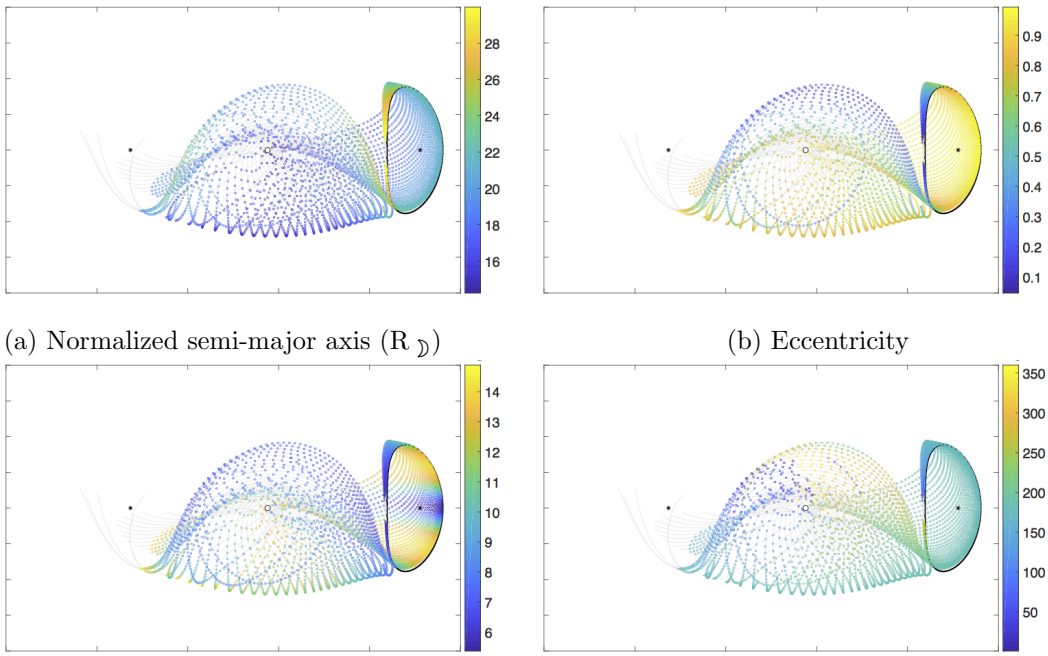

(a) Normalized semi-major axis ($R_☽$)

(b) Eccentricity

(c) Inclination (deg)

(d) True anomaly (deg)

Figure 3.11: Orbital elements of patch points on the discretized moon-approaching branch of EML_2 halo's unstable manifold.

Chapter 4

Combinatorial Optimization Model for In-Space Supply Chain Design

4.1 Introduction

We approach the problem of optimizing space logistics design by first modeling the space exploration map as a network, with nodes representing destinations or parking orbits, and arcs connecting these nodes representing flow of material between them. Each arc in a space exploration map (e.g. Figure 3.2) has discrete and continuous parameters associated with it, such as the payload/vehicle combination, technology used on the arc, time required to traverse the arc, etc. The selection of these parameters influences the campaign-level architecture and must be traded off against each other to affect the end goal efficiently, which is usually to minimize costs. Presence of arcs that use low-thrust spacecraft to provide propulsion complicates mission analysis due to the inherent nonlinear nature of such transfers, where the time of flight and the amount of propellant consumed are dependent on the size of the payload traversing that arc. Thus, optimal logistics design of space exploration campaigns that include low-thrust missions is a mixed-integer nonlinear programming (MINLP) problem.

The modeling paradigm developed in this chapter formulates this MINLP problem such that heuristic search techniques can be used for optimization, with trajectory considerations integrated into this top-level analysis. The case study considered here is that of re-establishing short-term human presence on the lunar surface with one or more crew missions, by supporting Apollo-style missions with in-space refueling capabilities through predeployed disposable droptanks. These droptanks can be delivered using a variety of vehicles, which include both high-thrust chemical and low-thrust solar electric propulsion variants. With respect to this case study, the new framework developed here optimizes the locations of crew mission refueling simultaneously with the mode of prepositioning these propellant resupply tanks; this allows for the cost of all involved missions to

be optimized together as a campaign.

The first space logistics optimization methodology developed during the course of this work and presented in the current chapter frames the design of the cislunar propellant resupply chain as a combinatorial optimization problem. A multiobjective genetic algorithm is applied at the level of campaign planning to find the optimal mission profiles while also calculating the optimal combination of propulsion systems to be used. The three costs of IMLEO, cargo delivery duration and crew flight time are simultaneously optimized through the use of multiple objectives. Optimal predeployment strategies are calculated for a campaign of one, two, or three crewed lunar missions, thus quantifying the game changing impact of low-thrust propulsion for logistics supply planning in the case study.

4.2 Multiobjective Genetic Algorithm for Selecting Campaign Parameters

The selection of parameters associated with each network arc influences the campaign-level architecture and their combinations must be explored in a regular manner to minimize the associated costs. The framework proposed here uses a multiobjective genetic algorithm (henceforth referred to simply as 'MOGA') to compute the ideal combination of technologies for transport and location of refuel tanks for campaign-level logistics design. The MOGA chooses the parameters such as target locations and propulsion technology associated with each cargo mission arc; these pieces of information are fed into the fitness function that calculates the cost of traveling on the chosen arcs in terms of IMLEO and time of campaign. Different combinations of carry-along and refueling architectures can be evaluated in this way to determine the optimal locations of depots where the fuel supply for later crewed mission segments may be stored. The important features of the framework are presented in Figure 4.1, with the details of each block provided in the subsequent sections.

An integer-valued genetic algorithm (GA) is chosen to conduct this combinatorial search for optimal architectures; further, the use of a multiobjective GA (instead of a single-objective GA) allows mission designers to explore the tradespace of IMLEO vs. time thoroughly in a single run.

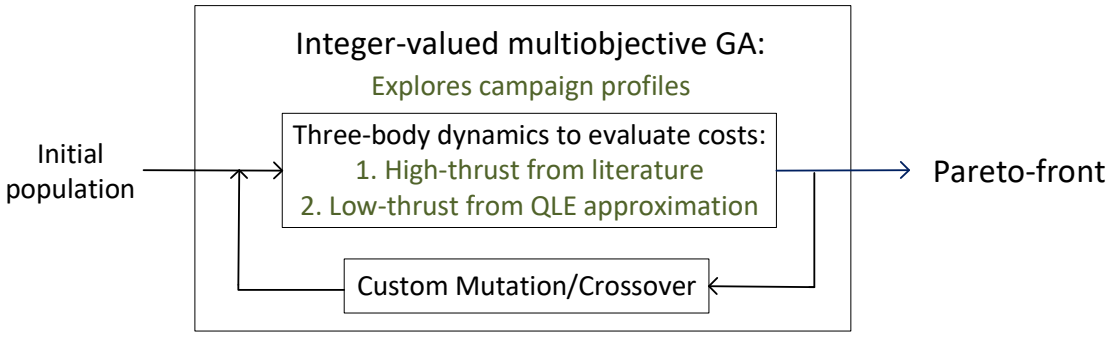

Figure 4.1: Overview of the multi-objective GA framework for space logistics optimization.

The objective is expressed as:

$$\underset{\boldsymbol{\chi}}{\text{minimize}} \quad \mathcal{J}(\boldsymbol{\chi}) = \left\{ \text{IMLEO}(\boldsymbol{\chi}),\ \text{T}_{\text{crew}}(\boldsymbol{\chi}),\ \text{T}_{\text{cargo}}(\boldsymbol{\chi}) \right\} \ , \tag{4.1}$$

where the variable $\boldsymbol{\chi}$ denotes the campaign architecture, T_{crew} represents the total crew flight time, and T_{cargo} is the total duration of the cargo delivery phase. Because the campaign encompasses missions to establish the propellant supply trail as well as the final crewed missions that use the fuel available at these way stations, the components of the cost function can be expressed as:

$$\text{IMLEO} = \text{IMLEO}_{\text{crew}} + \sum \text{IMLEO}_{\text{cargo}} \ , \tag{4.2}$$

$$\text{T}_{\text{crew}} = \sum \text{TOF}_{\text{crew}} \ , \tag{4.3}$$

$$\text{T}_{\text{cargo}} = \max\left(\text{TOF}_{\text{cargo}}\right) \ . \tag{4.4}$$

4.2.1 Chromosomal Representation

Table 4.1 shows the chromosomal representation used for a campaign involving only one crew mission and the cargo missions required to support it. The table also shows the corresponding bounds on the genes, which represent the parameters pertaining to the current problem. The "pitstops" where the crew spacecraft may be refueled are (A) EML_1 halo orbit, (B) EML_2 halo orbit, and (C) LLO. These network nodes were detailed in Chapter 3. Note that Table 4.1 uses **one** crew mission as an example, but it can be naturally extended to a campaign to support refueling of multiple crew missions.

Table 4.1: Chromosomal structure for **one** crew mission campaign.

Gene number	Lower bound	Upper bound	Description
1	0	1	Refuel status at pitstop A on forward journey
2	1	$n_{\max}$	Vehicle ID for delivering above refuel mass
3	0	1	Refuel status at pitstop B on forward journey
4	1	$n_{\max}$	Vehicle ID for delivering above refuel mass
5	0	1	Refuel status at pitstop C on forward journey
6	1	$n_{\max}$	Vehicle ID for delivering above refuel mass
7	0	1	Refuel status at pitstop C on return journey
8	1	$n_{\max}$	Vehicle ID for delivering above refuel mass
9	0	1	Refuel status at pitstop B on return journey
10	1	$n_{\max}$	Vehicle ID for delivering above refuel mass
11	0	1	Refuel status at pitstop A on return journey
12	1	$n_{\max}$	Vehicle ID for delivering above refuel mass

Every odd-numbered gene is valued either 0 or 1 to denote the absence/presence of refueling stations at each considered node along the path of the crew, while the genes immediately succeeding them (i.e., the even-numbered genes) pick integers that map to vehicles types/IDs used to deliver the fuel tanks to that node. A fixed-length chromosomal representation has been used here, with even-numbered genes not contributing to the fitness of the candidate if their preceding genes are not "expressed". For example, gene #2 does not express itself when gene #1 is "off"; this means that the vehicles chosen are ignored if the corresponding node is not involved in the refueling chain. This is similar to the concept of "hidden genes" developed in previous studies for spacecraft trajectory applications [93]. A chromosome representing the no-refuel case is shown in Table 4.4a — the even-numbered genes are greyed out to imply that their values do not influence the fitness of the candidate as the odd-numbered genes are all zero.

The vehicle IDs contained in the chromosome are mapped to a discrete set of tug specifications. The term "vehicle" is used in this work to denote the choice of propulsion system specification, that is, a vehicle can be used on any network arc to deliver any size of payload that is permitted by the size of its fuel tank. Chemical propulsion (CP) tug options particular to this case study are listed in Table 3.2, while the mappings from each vehicle ID to their propulsive system specifications are laid out in Table 4.2. Another example of the chromosome is shown in Table 4.4b where a single tug with ID = 3 is selected to deliver droptanks at both LLO and EML_2 halo.

Table 4.2: CP tugs: vehicle IDs used in chromosomal representation.

Vehicle IDs	Description
1–2	CP tug type 1 vehicle 1–2
3–4	CP tug type 2 vehicle 1–2
5–7	CP tug type 3 vehicle 1–3

Table 4.3: SEP tugs: vehicle IDs used in chromosomal representation.

Vehicle IDs	Description
8–9	SEP tug type 1 vehicle 1–2
10–11	SEP tug type 2 vehicle 1–2
12	SEP tug type 3 vehicle 1

Multiple vehicles of the same kind are allowed in the analysis by assigning them unique identifiers. For example, two CP tugs of type 2 (with dry mass of 5.5 t and propellant capacity of 41 t) are allowed in the analysis by assigning them vehicle identifiers of 3 and 4. The MOGA may then choose to use zero, one, or both of this particular tug type to deliver the propellant droptanks. The possibility of using both of these tugs is shown in Table 4.4c. Similar to the selection of CP tugs, various SEP tug options are encoded in the chromosomal structure using the mapping in Table 4.3, with each engine sized according to Table 3.3.

The upper bound on the vehicle identifier genes ($n_{\max}$ in Table 4.1) comes from how many vehicles are permitted in the analysis as well as how many times they can be reused. A vehicle is considered to be reused when the vehicle identifier is a multiple greater than the number of vehicles available as mapped from Tables 4.2 and 4.3 (= 12 in this case). For example, vehicle ID = 15 (which is greater than the 12 vehicles allowed in the analysis) implies the first reuse of tug with vehicle ID = 3 from Table 4.2, whereas vehicle ID = 27 denotes its second reuse. This chromosome example is encoded in Table 4.4d.

The entire length of 12 "genes" is repeated for every human mission included in the campaign, e.g., the chromosome has 36 genes for a campaign to support refueling needs of three crew missions. A large number of possibilities are expressed by this chromosome due to the permutations involving presence/absence of depots at node locations, multiple tugs of each kind allowed to make cargo deliveries, the ability of each tug to deliver cargo to multiple locations (discussed in more detail in Chapter 3), and the reusability of tugs. Thus, the solution space cannot be evaluated through

Table 4.4: Example chromosomes encoding some possible campaign architectures.

(a) No-refuel

Gene #	Value
1	0
2	x_1
3	0
4	x_2
5	0
6	x_3
7	0
8	x_4
9	0
10	x_5
11	0
12	x_6

(b) Single tug

Gene #	Value
1	0
2	x_1
3	0
4	x_2
5	1
6	3
7	0
8	x_4
9	1
10	3
11	0
12	x_6

(c) Multiple tugs

Gene #	Value
1	0
2	x_1
3	1
4	4
5	1
6	3
7	0
8	x_4
9	1
10	3
11	0
12	x_6

(d) Tug reuse

Gene #	Value
1	0
2	x_1
3	1
4	15
5	1
6	3
7	1
8	27
9	1
10	3
11	0
12	x_6

exhaustive enumeration, and an adaptive heuristic search method (i.e., the MOGA) is adopted instead. Although other evolutionary heuristic formulations may also be applied effectively, as an example, our current work uses the controlled elitist variant of NSGA-II, which is a multiobjective genetic algorithm provided with MATLAB [94]. This specific MOGA ranks individuals in the population based on their relative proximity to the population's nondominated fronts. Each front consists of individuals that cannot outperform other candidates also on the same front. This nondominated ranking system, together with a distance measure that calculates how "far" individuals are from each other on each front, allows the NSGA-II to use strategies during the selection process to encourage diversity in the population as it evolves. Specific implementation details are available in [95].

Although the built-in selection function is used as-is in this research, simple custom creation, crossover, and mutation functions are developed to maintain the integer values of the genes. The creation function generates an initial population of feasible candidates that use each permissible vehicle ID at least once; the crossover function randomly selects one gene pair (i.e., refuel location + tug to deliver droptank to that location) in the first parent to cross over with the second parent; and finally, the mutation function picks a random number of genes in the chromosome and changes their values, while keeping them within their corresponding bounds.

4.3 Cost Evaluation

The cost function parses the incoming chromosome by first extracting the information about the locations of refueling stations. Each refueling event is then assumed to provide the entire propellant required by the crew vehicle to reach the next refueling station or to complete the journey. This includes lunar descent and ascent operations in case refueling is conducted before landing. The IMLEO costs are then calculated in a reverse order, i.e., the costs for the crew mission(s) are calculated first. Knowing the location of the refuel stations, the propellant required by the crew spacecraft for each leg of its mission can be used to compute the masses of the droptanks to be delivered to each node. The costs for high-thrust transfers (used for crew vehicles and chemical propulsion tugs) are derived from existing literature. On the other hand, the costs for low-thrust transfers (used for solar electric propulsion tugs) are calculated internally with a computationally efficient interpolation method based on the low-energy low-thrust transfer cost model derived earlier in Section 3.4.3.

4.3.1 Crew Mission Costs

High-thrust chemical propulsion vehicle and trajectory choices for crew missions have already been presented in Chapter 3. High-thrust rocket burns can be modeled as instantaneous events; this allows for a reasonable estimate of mass propellant consumed during a transfer without heed to the payload mass itself. Given the magnitude of the change in velocity (Δv) required for a maneuver, Tsiolkovsky's rocket equation can be used to determine the propellant consumed for the transfer. For the purposes of the case study demonstrated here, the Δv values for crew trajectories are gathered from available literature (Table 3.4). Once this value is known, the theoretical mass of the propellant required to execute the maneuver can be derived from the rocket equation as:

$$m_{\text{propellant}} = \frac{\phi}{1-\phi}\, m_{\text{payload}} \ , \tag{4.5}$$

where $\phi = 1 - \exp\left(\frac{-\Delta \text{v}}{g I_{\text{sp}}}\right)$. If the maneuver is conducted by the crew vehicle, the "payload mass" would be the mass of the crew stack. The mass to be delivered for refueling is modeled as a disposable droptank, where the propellant is transferred from these tanks to the crewed spacecraft

at a later time. The mass of these droptanks, which then becomes the payload mass for cargo delivery tugs, is found by using the structural coefficient. Assuming that the structural mass of these tanks scales linearly with the propellant mass they contain, a structural coefficient ε can be used to size the droptanks as:

$$m_{\substack{\text{droptank}\\\text{structure}}} = \frac{\varepsilon}{1-\varepsilon} m_{\text{propellant}} \quad , \tag{4.6}$$

with Equation (4.5) providing the value of $m_{\text{propellant}}$. The crew spacecraft's dry mass is assumed to include tanks to which the refuel propellant can be transferred.

The first impulse for the crew transfer (i.e., the TLI) is provided by the upper stage of the launch vehicle. If the crew spacecraft is scheduled to be refueled at some later point along its journey, then less mass needs to be launched and placed beyond TLI with the crew itself, thus reducing the size of the upper stage. Unlike the droptanks, which use the tug propellant to be delivered, the upper stage has to carry fuel to propel itself. This means that the payload mass modeled in the typical rocket equation for the purpose of calculating the size of the upper stage includes its own tank structure as well as the mass of the crew stack. The rocket equation then becomes:

$$\frac{m_{\substack{\text{crew}\\\text{stack}}} + m_{\substack{\text{upper stage}\\\text{tank}}} + m_{\substack{\text{upper stage}\\\text{propellant}}}}{m_{\substack{\text{crew}\\\text{stack}}} + m_{\substack{\text{upper stage}\\\text{tank}}}} = \frac{1}{1-\phi_{\text{TLI}}} \quad , \tag{4.7}$$

where $\phi_{\text{TLI}} = 1 - \exp\left(\frac{-\Delta v_{\text{TLI}}}{g I_{\text{sp}}}\right)$. Using $\hat{\varepsilon} = \frac{\varepsilon}{1-\varepsilon}$, the upper stage empty mass can be expressed as $m_{\substack{\text{upper stage}\\\text{tank}}} = \hat{\varepsilon}\, m_{\substack{\text{upper stage}\\\text{propellant}}}$. Substituting in Equation (4.7) gives:

$$m_{\substack{\text{upper stage}\\\text{propellant}}} = \frac{\phi_{\text{TLI}}}{1-\phi_{\text{TLI}} - \hat{\varepsilon}\phi_{\text{TLI}}} m_{\substack{\text{crew}\\\text{stack}}} \quad , \tag{4.8}$$

Given these preceeding relationships, the part of the IMLEO cost that derives from the crew launch is:

$$\text{IMLEO}_{\text{crew}} = m_{\substack{\text{upper stage}\\\text{propellant}}} + m_{\substack{\text{upper stage}\\\text{tank}}} + m_{\substack{\text{crew}\\\text{stack}}} \quad . \tag{4.9}$$

In the current case study, the crew stack composed of the CSM mated with the LM is transported

between the space network nodes (as shown in Figure 3.2) on the arcs that connect them. The contribution of each crew mission to the total campaign cost is calculated separately.

4.3.2 Cargo Delivery Costs Using Chemical Propulsion Tugs

The second contribution to the IMLEO cost of the campaign comes from the delivery of droptanks to their respective locations in the cislunar space (Equation (4.2)). The amount of propellant to be resupplied to the crew by each droptank is obtained in the previous step; this information helps size the droptanks, which in turn become the payload masses for the cargo delivery tugs. The Δv data for the cargo deliveries using chemical propulsion tugs has provided previously in Table 3.5. These tugs begin their journey at LEO, and the propellant required to conduct each of their maneuvers is given by the rocket equation. Thus, the cost of predeploying refuel tanks using chemical propulsion tugs (if any), including reuse scenarios, is:

$$\underset{\text{chemical tug}}{\text{IMLEO}_{\text{delivery using}}} = f_1 \cdot \underset{\text{mass}}{m_{\text{tug dry}}} + m_{\text{cargo}} + \underset{\text{cargo delivery}}{m_{\text{propellant for}}} \quad , \tag{4.10}$$

where $m_{\text{cargo}} = m_{\text{droptank}} + f_2 \cdot \underset{\text{return to LEO}}{m_{\text{propellant for}}}$. In these equations, f_1 and f_2 are flags that account for reusability. $f_1 = 1$ for the first flight of a cargo vehicle, and 0 for later flights, whereas $f_2 = 1$ if the same tug has to return to LEO for a later reuse, and 0 otherwise.

Similarly, the times for delivery flights are cascaded and summed for tug reuse missions; this is true for SEP tugs as well (described in succeeding section). For example, the time of flight for the second cargo mission using a particular tug is a sum of the flight time of its first delivery, time for returning to its parking orbit around the Earth, and the flight time for its second delivery.

4.3.3 Cargo Delivery Costs Using Solar Electric Propulsion Tugs

The propellant mass consumed by a solar electric propulsion spacecraft and its time of flight are heavily dependent on its initial mass (i.e., the sum of payload mass, dry mass of the vehicle, and onboard fuel mass). This is because the Δv required to traverse an arc using low-thrust propulsion varies with its thrust-to-mass ratio, which in turn is decided by the thrust level of the solar electric propulsion tug, how much payload it is carrying to different nodes, if it is dropping parts of its cargo at multiple locations, and if it is reused later for other cargo missions. Consequently, the Δv

look-up scheme implemented for cargo delivered by the CP tugs cannot be used here, and the cost of traveling on these arcs must be calculated internal to the framework.

Calculating optimal low-energy low-thrust trajectories is a complex and computationally intensive process. In the current analysis, the need is for an estimate of cargo delivery costs (in terms of time and propellant consumed), and not the exact trajectory itself. To achieve this in a computationally efficient manner, the Q-law low-energy method developed in Chapter 3 was used to derive the linear cost model for the low-thrust network arcs (refer Figure 3.3 and Tables 3.6–3.7). For every campaign architecture solution explored by the MOGA, Equations (3.1a)-(3.1b) are used to calculate the associated costs. This incorporation of SEP transfers into a space logistics framework by implementing a method that can estimate the cost of low-thrust trajectories in the cislunar region irrespective of the initial thrust-to-payload-mass ratio is one of the main contributions of this work.

In candidate genetic sequences that use SEP tugs to transport commodities on the network, the GTO is used as the initial Earth parking orbit. This node at GTO serves only as the launch orbit of any deliverable using SEP tugs, in accordance with prior studies dealing with optimal low-thrust transfers [22, 28]. Any SEP-powered transfer would realistically begin at GTO because it would take very long for the tug to exit Earth's gravity well if launched from LEO and would consume excessive fuel due to time spent thrusting continuously in regions of nonnegligible atmospheric drag. This implies that the final launch cost for SEP arcs is obtainable only in terms of initial mass to GTO (IMGTO), instead of in terms of IMLEO. This discrepancy is remedied by using the following relationship:

$$\mathrm{IMLEO}_{\substack{\text{delivery using}\\ \text{SEP tug}}} = \mathrm{P}_{\mathrm{GTO}} \times \mathrm{IMGTO}_{\substack{\text{delivery using}\\ \text{SEP tug}}} \ , \tag{4.11}$$

where the penalty of $\mathrm{P}_{\mathrm{GTO}} = 1.74$ is used (refer Chapter 3).

4.4 Results

The cost of architectures encoded by the MOGA's chromosome are calculated by summing the IMLEO of each mission comprising the campaign using Equations (4.9)–(4.11). The time of cargo

flight is similarly obtained by taking the maximum of the times of flight for cargo missions, including the reuse of tugs. To ensure reasonable crew flight times, the manned spacecraft is not allowed to be refueled at both halo orbits on the same leg of their journey (from LEO to lunar surface is the first leg of the journey, whereas the return journey is the second leg). Similarly, chromosome sequences dictating direct heteroclinic transfers between the EML_1 and EML_2 halos for the cargo delivery tugs are discarded. Table 4.5 details the nondefault settings for the MOGA used in this study. All other settings are left at default values used by MATLAB's `gamultiobj` [94].

Table 4.5: Parameters for the multi-objective GA for space logistics optimization.

Parameter	Value
Population size	150
Maximum number of generations	200
Maximum number of stall generations	50

Given these assumptions and constraints, this framework is implemented to find the optimal architectures for in-space refueling of one, two, and three crew missions. The results obtained in each case are presented in this section against a baseline case consisting of no-refuel Apollo mission/s. This **baseline IMLEO cost is 124.224 t per crew mission** – all optimal solutions obtained are expressed as percentage improvement over this value.

4.4.1 Campaign to support one crew mission

Using the no-refuel Apollo baseline, the Pareto front of the three aforementioned objectives is generated to explore new space architectures where refueling architectures can be compared against a carry-along-only logistics strategy. For this current subcase, no reuse of tugs is allowed, which aligns well with providing support for a single crew mission only. Figure 4.2 shows the Pareto front of percentage IMLEO improvement versus T_{cargo}, with the third objective (T_{crew}) shown using color. Because of the reduction of the three-dimensional Pareto front to a two-dimensional figure, it may seem that certain points such as C1 and C2 are dominated, but they are in fact nondominated due to the presence of the third objective. The color bar represents the third objective, which is the total number of days spent by the crew in space. This Pareto front is examined in some detail in this section to understand the efficacy of the framework presented in this analysis, before studying

other campaigns that examine refueling support for more than one crew mission.

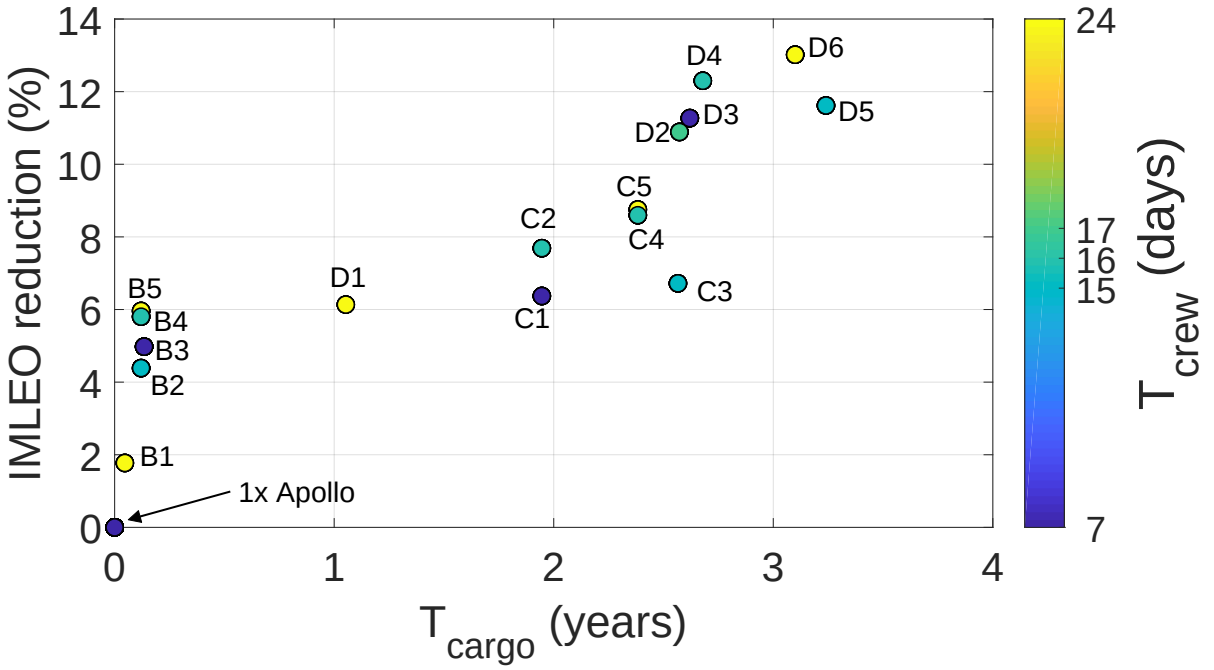

Family	Description
B1-B5	1× CP type 2
C1-C5	1× CP type 1, 1× SEP type 1

Family	Description
D1	1× SEP type 2
D2-D6	1× SEP type 1

Figure 4.2: Pareto front of different refueling architectures for one Apollo-type mission.

This front consists of approximately four families of solutions, with each family segregated by the type of tugs used within the campaign. The three intuitive types of solutions are (a) no refueling, (b) all refueling tanks delivered by chemical propulsion tugs and (c) all refueling tanks delivered by solar electric propulsion tugs. The first family is indeed composed of only the Apollo solution, where no refueling is allowed. The second family (B) represents setting up the propellant supply chain using only CP tugs, while the third idea of using only SEP tugs is reflected in the D family. The details of chosen point solutions from these families are provided in Table 4.6 and Table 4.7.

The points within each family are separated by their architectures (i.e., locations of refuel stations) or by the vehicles used to set up the architecture. For example, the points B3 and B5 use the same CP tug to deliver the required propellant droptanks as shown in Table 4.6, but solution B3 refuels the crew spacecraft only in lunar orbit, whereas solution B5 involves refueling at the

Table 4.6: Chosen solutions from Family B in Figure 4.2 (CP tugs only).

	Point B3		Point B5	
	Forward journey	Return journey	Forward journey	Return journey
Droptank at EML_1 halo	✗	✗	✗	✗
Tug used for delivery	–	–	–	–
Droptank at EML_2 halo	✗	✗	✓	✓
Tug used for delivery	–	–	CP type 1, vehicle 1	CP type 1, vehicle 1
Droptank at LLO	✓	✓	✓	✓
Tug used for delivery	CP type 1, vehicle 1	CP type 1, vehicle 1	CP type 1, vehicle 1	CP type 1, vehicle 1

Table 4.7: Chosen solutions from Family D in Figure 4.2 (SEP tugs only).

	Point D1		Point D4	
	Forward journey	Return journey	Forward journey	Return journey
Droptank at EML_1 halo	✗	✗	✗	✗
Tug used for delivery	–	–	–	–
Droptank at EML_2 halo	✓	✓	✗	✓
Tug used for delivery	SEP type 2, vehicle 1	SEP type 2, vehicle 1	–	SEP type 1, vehicle 1
Droptank at LLO	✓	✓	✓	✓
Tug used for delivery	SEP type 2, vehicle 1	SEP type 2, vehicle 1	SEP type 1, vehicle 1	SEP type 1, vehicle 1

EML_2 halo on both legs of the crew journey. On the other hand, points D1 and D2 use different SEP tugs and different refueling profiles (see Table 4.7).

Family C contains the solutions of special interest, where cargo deliveries are shared between CP tugs and SEP tugs in an optimal manner. The solutions distribute the cargo between one chemical propulsion tug and one solar electric propulsion tug. Both these chosen tugs represent the smallest engines in their group. Although multiple vehicles of each type are available for use in the analysis, almost all the optimal solutions use at most one of each tug type due to the IMLEO cost associated with launching multiple vehicles. The mission profiles as well as the details about the tugs used for cargo deliveries for select family C solutions are presented in Table 4.8. Additionally, point C3 is represented diagrammatically in Figure 4.3.

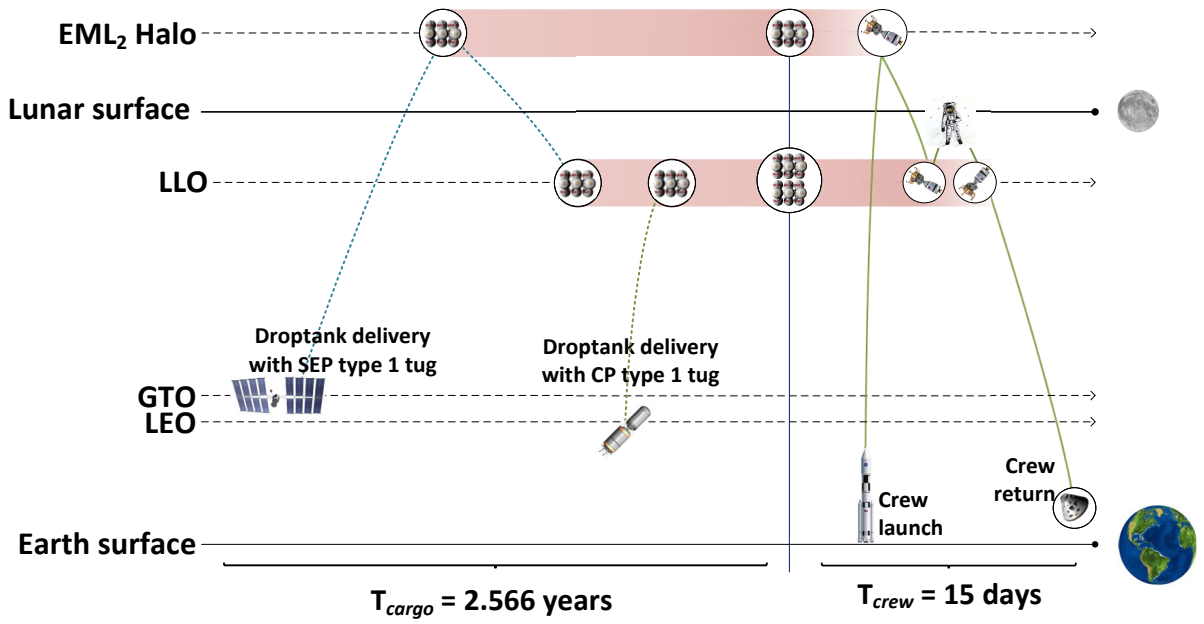

Figure 4.3: Illustration of point solution C3 from Figure 4.2.

Table 4.8: Chosen solutions from Family C in Figure 4.2 (combination of CP and SEP tugs).

	Point C3		Point C5	
	Forward journey	Return journey	Forward journey	Return journey
Droptank at EML_1 halo	×	×	×	×
Tug used for delivery	–	–	–	–
Droptank at EML_2 halo	✓	×	✓	✓
Tug used for delivery	SEP type 1, vehicle 1	–	CP type 1, vehicle 1	CP type 1, vehicle 1
Droptank at LLO	✓	✓	✓	✓
Tug used for delivery	SEP type 1, vehicle 1	CP type 1, vehicle 1	SEP type 1, vehicle 1	SEP type 1, vehicle 1

4.4.2 Campaign to support two crew missions

Figure 4.4 shows the Pareto front for the case of setting up the supply chain for **two** Apollo-style crew missions. The previous analysis is extended to allow for reuse of tug vehicles; up to two reuses of each tug are allowed (i.e., total three uses of the same tug). In this figure, the percentage IMLEO improvement on the Y-axis now is now calculated over the base case of launching two no-refuel Apollo-style missions.

The longest campaign that lies on this Pareto front would require almost 6 years to setup in-space refueling support for just two crew missions. But there exist other solutions on this front

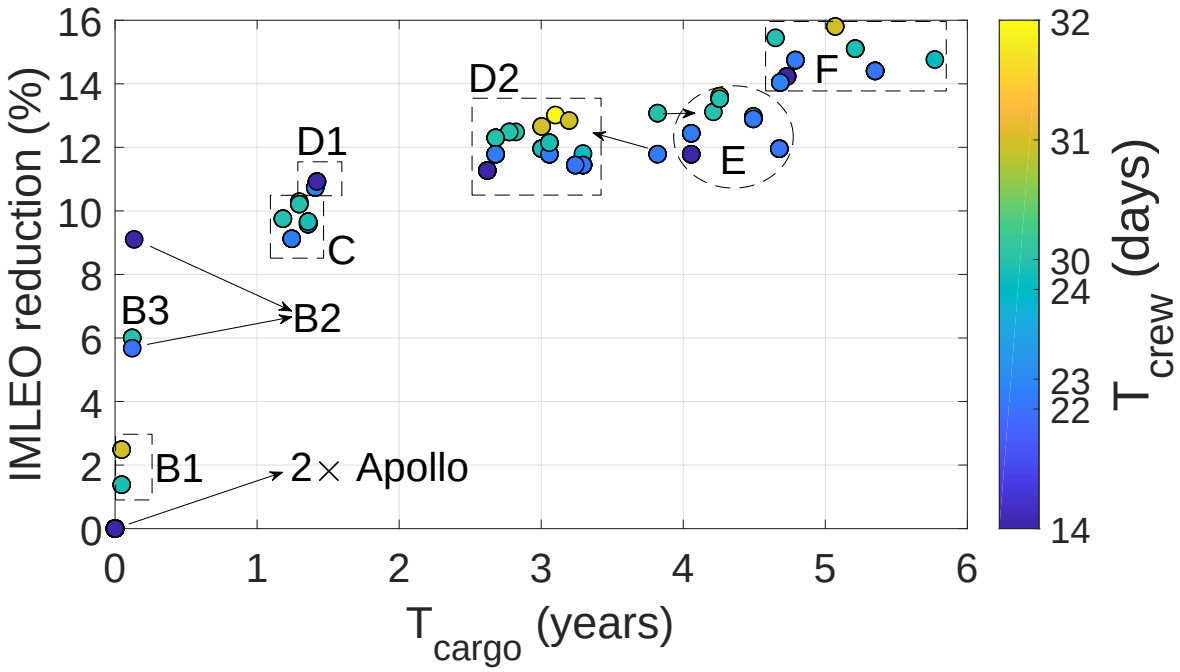

Family	Description
B1	1× CP type 2
B2	1× CP type 3
B3	2× CP type 2
C	1× CP type 1, 1× SEP type 2

Family	Description
D1	1× SEP type 2
D2	2× SEP type 1
E	1× CP type 1, 1× SEP type 1
F	1× SEP type 1

Figure 4.4: Pareto front of different refueling architectures for two Apollo-type missions.

in the region of 1-2 years and ∼3 years that also present attractive options for achieving the same goal. Families of solutions along this front are again classified based on the number and type of tugs used in the campaign.

4.4.3 Campaign to support three crew missions

Finally, the case of supporting the in-space refueling needs for **three** Apollo-style crew missions is studied and the results presented in Figure 4.5. The solutions from different "families" are now intermingled and the Pareto front is seen to plateau after the $T_{cargo} \approx 2$ years mark, with more practical solutions contained within the 1–2 year region. As with the results in Figure 4.2 and Figure 4.4, the solutions that use either CP-only or SEP-only lie at the extremes of the Pareto front. It is the solutions in the middle of the Pareto front that contain the optimal campaign architectures that combine the use of chemical propulsion tugs and solar electric propulsion tugs.

Overall, higher percentage savings are realized in this case than the previous case of two crew missions.

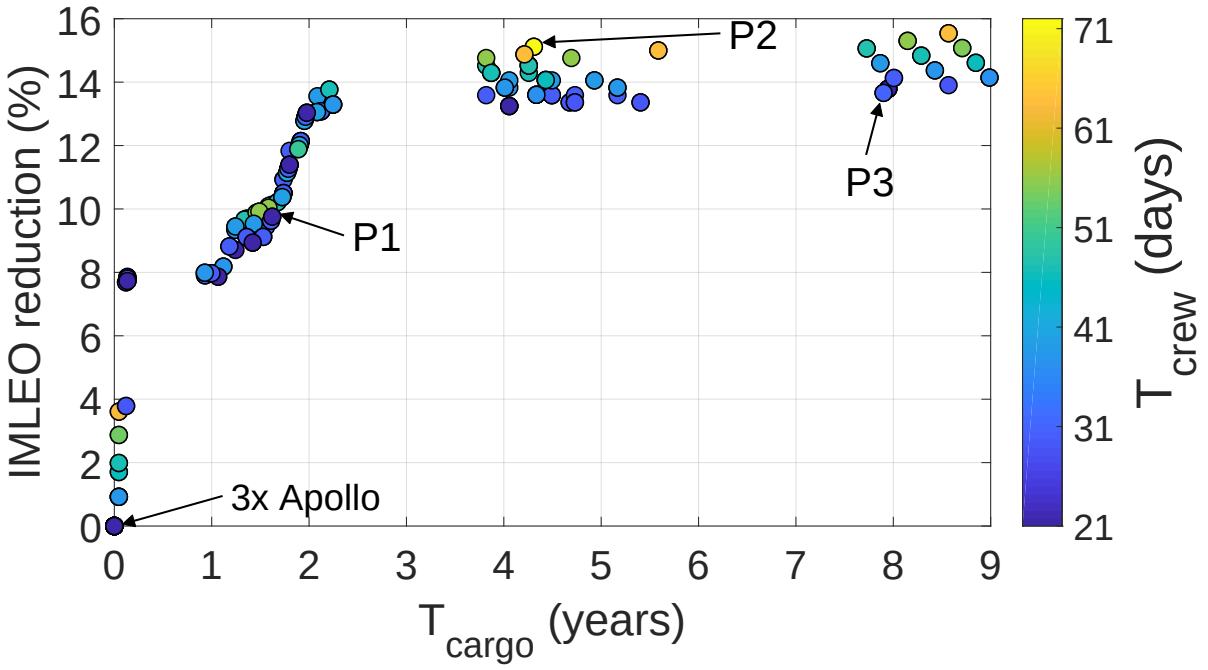

Figure 4.5: Pareto front of different refueling architectures for three Apollo-type missions.

Certain point solutions on the Pareto front that emphasize the unique effectiveness of the methodology developed here are labeled in Figure 4.5. Solution P1 uses a combination of three tugs (2× CP type 1, plus 1× SEP type 2) to predeploy resupply propellant to LLO. On the other hand, solution P2 utilizes both the SEP type 1 tugs available in the fleet to set up the droptanks at the EML_1 halo as well as the LLO. This design also display how the tugs leverage the low-energy pathways to deliver payloads at the halo orbit and LLO on the same journey. Though this campaign profile yields ∼15% IMLEO savings over the baseline no-refuel case, it takes a little more than 4 years for the cargo delivery phase and extends the crew mission phase to 72 days. Both these solutions are detailed in Table 4.9; reading this table in a downward direction yields a mapping to the corresponding chromosomal representation.

Finally, the solution point P3 is also marked as it demonstrates the compromises of reusing tugs. None of the solutions on the Pareto front reuse a CP tug, however, P3 reuses the smallest SEP tug once. This will involve refurbishing the tug when it returns to its parking orbit (GTO)

between its first and second journeys out into the cislunar space.

Table 4.9: Chosen solutions from Figure 4.5.

		Point P1		Point P2	
		Forward journey	Return journey	Forward journey	Return journey
Crew #1	Droptank at EML_1 halo	✗	✗	✗	✗
	Tug used for delivery	–	–	–	–
	Droptank at EML_2 halo	✗	✗	✓	✓
	Tug used for delivery	–	–	SEP type 1, vehicle 1	SEP type 1, vehicle 2
	Droptank at LLO	✓	✓	✓	✓
	Tug used for delivery	SEP type 2, vehicle 1	CP type 1, vehicle 1	SEP type 1, vehicle 2	SEP type 1, vehicle 1
Crew #2	Droptank at EML_1 halo	✗	✗	✗	✗
	Tug used for delivery	–	–	–	–
	Droptank at EML_2 halo	✗	✗	✓	✓
	Tug used for delivery	–	–	SEP type 1, vehicle 1	SEP type 1, vehicle 2
	Droptank at LLO	✓	✓	✓	✓
	Tug used for delivery	SEP type 2, vehicle 1	CP type 1, vehicle 2	SEP type 1, vehicle 1	SEP type 1, vehicle 1
Crew #3	Droptank at EML_1 halo	✗	✗	✗	✗
	Tug used for delivery	–	–	–	–
	Droptank at EML_2 halo	✗	✗	✓	✓
	Tug used for delivery	–	–	SEP type 1, vehicle 1	SEP type 1, vehicle 2
	Droptank at LLO	✓	✓	✓	✓
	Tug used for delivery	SEP type 2, vehicle 1	SEP type 2, vehicle 1	SEP type 1, vehicle 2	SEP type 1, vehicle 2

4.4.4 Discussion

The main point of note here is that these results are specific to the modeling parameters used, as is true with any simulation. For example, the node at the EML_1 halo is largely absent from the optimal solutions due to the slightly higher Δv cost to the crew missions transiting through that node. Another example is that the shape of the Pareto front in all three of the cases is considerably dependent on the value of the IMGTO penalty used in Equation (4.11). The value of this penalty for the current analysis is derived from the ideal cost to transfer netween LEO to GTO; however, this may not be the case in reality. Longer T_{cargo} solutions survive and make it to the Pareto front when a lower penalty is enforced, while the opposite is true when using a higher value.

The results also provide some other useful insights regarding fleet use. In the cases of two and three crew mission campaigns, both the CP and SEP tugs are allowed to be reused up to two times (that is, a total of three uses) each. However, only the smallest solar electric propulsion tug is reused in Pareto front solutions (point P3 from Figure 4.5). This is most likely due to the additional cost of carrying fuel to make the return trip, as well as the longer trip times that result from reusing tugs. Thus, the absence of other solutions that reuse tugs is likely due to the inclusion of T_{cargo} as one of the objectives. Additionally, the largest SEP tug allowed in the analysis is not chosen in the optimal solutions, due to its higher dry mass as well as lower fuel efficiency (refer Table 3.3).

Another useful insight gained through this analysis relates to the ideal locations of refueling nodes. A majority of the optimal solutions contain architectures that include refueling at LLO on both legs of the crew journey. This is because most of the fuel expenditure is around this node (i.e., landing, ascent, and departure). Apart from the refuel node at LLO, the architectures in most of the Pareto front solutions also include refueling at the EML_2 halo, which becomes advantageous due to the use of the low-energy pathways that encourage deliveries to multiple locations using the same tug. This work has thus exploited indirect relationships between cargo delivery pathways and the crew mission goals.

Note that, in the considered analysis, the crew missions are assumed to be launched towards the end of the campaign, after the completion of all cargo deliveries. However, the crew missions can also be potentially launched once their corresponding refueling points have been confirmed to be functional. This may lead to an irregularity in crew launch frequency. This does not preclude the possibility of launching regular crew missions; it just means that the infrastructure established by the cargo missions can only support these crew missions, while other missions can be supported in other ways such as the carry-along strategy.

Finally, the different groupings among the Pareto front solutions obtained through the framework developed in this chapter supply ample options to the mission designer in terms of the number and type of tugs as well as crew mission durations. These Pareto front solutions can be used to support different campaign goals, e.g., a campaign to prove in-space rendezvous technologies can make use of solutions that require multiple docking scenarios with droptanks (i.e., solutions with

longer crew missions), while the goal of demonstrating cargo deliveries using SEP vehicles can utilize solutions with longer cargo delivery times (i.e., solutions with longer T_{cargo}). The optimal solutions should be evaluated for other costs, like vehicle development, station-keeping, rendezvous, etc. which have not been quantified in this analysis.

4.5 Chapter Summary

This chapter proposes a new paradigm for optimizing the logistics network concerning in-space propellant resupply for campaign-level planning, where the locations of the refueling stations are simultaneously optimized with the propulsion technologies that can be used to deliver the refuel propellant. The tradespace of using chemical propulsion tugs vs. solar-electric propulsion tugs for refuel tank deliveries is automatically explored and evaluated by a multiobjective genetic algorithm. The objectives included for populating the Pareto front are IMLEO, time duration of campaign's cargo delivery phase, and time duration of crew flights. This overall methodology was applied to optimally establish a propellant supply trail for a campaign consisting of one to three manned missions to the Moon (Figure 4.2–Figure 4.5).

This framework uses a multiobjective genetic algorithm which utilizes stochastic search techniques to explore the tradespace. This causes the results to be sensitive to the initial population, thereby requiring multiple runs or increase in allowable number of iterations to approach the global minima. Even then, global optimality of the generated Pareto front is not guaranteed. Additionally, the current representation of campaign profiles by genetic sequences requires a simplifying assumption regarding the amount of crew propellant supplied at each pitstop. To address these drawbacks, a new mixed-integer linear programming based formulation for space logistics optimization is developed in the next chapter.

Material in this chapter has appeared in a paper authored by Bindu Jagannatha and Koki Ho in *Journal of Spacecraft and Rockets* (2018) [90]. Although the same case study was explored in the paper, the problem specifics have been changed slightly in this document and the updated results are presented here.

Chapter 5

Event-Driven Network Optimization Model for In-space Supply Chain Design

The second space logistics framework developed in this work for conducting propulsion technology trades is based on the generalized multi-commodity network flow (GMCNF) model introduced in the first chapter. In this chapter, the mathematical formulation of the GMCNF model is presented first, before summarizing its full time-expanded variant (i.e., the TE-GMCNF). Finally, a new dynamic variant is developed based on event-based time steps, rather than discrete time steps. This event-driven GMCNF (ED-GMCNF) formulation is then applied to the top-level design of the cislunar campaign that uses both high-thrust and low-thrust vehicles.

For the specific case of cislunar logistics design, the previous chapter has already demonstrated cost savings that can be realized through the use of an in-space propellant resupply chain. In the current chapter, we aim to improve the rigor of the problem formulation and eliminate some limiting assumptions.

5.1 Static GMCNF Model

The static GMCNF formulation approaches the problem using graph-theoretic modeling techniques, where linear programming (LP) and mixed-integer linear programming (MILP) is used to minimize costs while satisfying the given set of linear constraints. Consider a **static** graph consisting of a set of nodes $\mathbb{N}$, connected by a set of directed arcs $\mathbb{A}$ which allow transportation of multiple commodities across it as represented in Figure 5.1a. For each arc from node i to node j, the multi-commodity flow can be split into outflow $\mathbf{x}_{ij}^{+}$ (i.e., flowing out of node i towards node j) and inflow $\mathbf{x}_{ij}^{-}$. Within this static graph, alternative trajectory/propulsion options along each arc can be represented by the multi-graph (Figure 5.1b), where multiple arcs connect the same pair of end nodes. For now, only *transportation* arcs are considered as a part of the static network (i.e. $i \neq j$); the concept of

holdover arcs is introduced later to connect nodes to themselves across time steps. Cost coefficients $\mathbf{c}^{+}_{ijv}$ are assigned to the outflow on each arc.

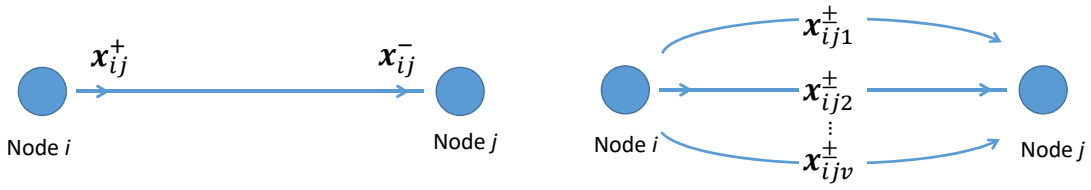

(a) Transportation arc (b) Multi-graph

Figure 5.1: Arcs used in the static network graph formulation.

Using the notation developed above, the GMCNF formulation can be expressed as:
Minimize

$$\mathcal{J} = \sum_{(i,j,v)\in\mathbb{A}} \mathbf{c}^{+}_{ijv}\mathbf{x}^{+}_{ijv} \ , \tag{5.1}$$

subject to

$$\sum_{\substack{(j,v):(i,j,v)\in\mathbb{A}\\ i\neq j}} \mathbf{x}^{+}_{ijv} - \sum_{\substack{(j,v):(j,i,v)\in\mathbb{A}\\ i\neq j}} \mathbf{x}^{-}_{jiv} \leq \mathbf{d}_i \qquad \forall\ i \in \mathbb{N} \ , \tag{5.2a}$$

$$\mathbf{B}_{ijv}\mathbf{x}^{+}_{ijv} = \mathbf{x}^{-}_{ijv} \qquad \forall\ (i,j,v) \in \mathbb{A} \ , \tag{5.2b}$$

$$\mathbf{C}^{+}_{ijv}\mathbf{x}^{+}_{ijv} \leq \mathbf{s}^{+}_{ijv} \qquad \forall\ (i,j,v) \in \mathbb{A} \ , \tag{5.2c}$$

$$\mathbf{x}^{\pm}_{ijv} \geq \mathbf{0}_{k\times 1} \qquad \forall\ (i,j,v) \in \mathbb{A} \ . \tag{5.2d}$$

The above equations represent the mathematical formulation of the static GMCNF model as proposed by Ishimatsu et. al. [7, 8], minus the equations that correspond to holdover arcs. Equation (5.1) represents the objective function as the sum of all transportation costs across the network.

The first constraint, Equation (5.2a), deals with mass balance by ensuring that the demand at every node is met under the supply conditions. Equation (5.2b) uses the k-by-k mass transformation matrix $\mathbf{B}_{ijv}$ to show gains/losses and to account for instantaneous commodity transformations. This constraint is used in space logistics applications to model the propellant consumption by

instantaneous propulsive maneuvers required for arc transport. Concurrency constraints over all arcs are managed by Equation (5.2c). If n_c such constraints have to be enforced, then the matrix $\mathbf{C}^+_{ijv}$ will be of the size n_c-by-k and $\mathbf{s}^+_{ijv}$ a vector of size n_c-by-1. Finally, Equation (5.2d) represents the nonnegative nature of flows across all arcs.

This GMCNF model can also be used to consider continuous mass transformations, e.g. resource generation at an ISRU plant. If an arc (i, j) has only continuous mass transformations across it and $\frac{d\mathbf{x}}{dt} = \mathbf{B}^c\mathbf{x}$ represents this conversion, then the $\mathbf{B}_{ijv}$ term in Equation ((5.2b) can be expressed as:

$$\mathbf{B}_{ijv} = \exp\left(\mathbf{B}^c_{ijv}\Delta t_{ijv}\right) \ , \tag{5.3}$$

where Δt_{ijv} is the "arc length" or the time required to complete the continuous process. If both types of transformations (instantaneous and continuous) occur over an arc, the corresponding constraint over it can be aggregated by serially multiplying the $\mathbf{B}$ matrices in an order matching the sequence of transformation events/processes across that arc.

It may be argued that this above concept could be used to model continuous propellant consumption over a low-thrust trajectory; however, the fuel consumption in such a case is not only continuous, but also a function of the ratio of thrust to initial mass on the arc. The easiest way to express this idea is that Δt in Equation (5.3) is a function of $|\mathbf{x}^+|$, or $\Delta t_{ijv} = f(|\mathbf{x}^+_{ijv}|)$, thus turning the mass transformation constraint in Equation (5.2b) nonlinear.

5.1.1 Full Time-Expanded GMCNF Model

Now, consider the **dynamic** network model where the time dimension is introduced by duplicating the static network at each time step. Figure 5.2 illustrates this idea regarding representation of dynamic flow using a notional static network. While transportation arcs still facilitate mass flow between different nodes, the concept of *stock* can be explored using the *holdover* arcs that connect the same node across time steps. As always, these arcs only represent the possibility of flow, and need not be all utilized in the final solution.

Though a powerful formulation, the time-expanded GMCNF is based on the assumption of fixed arc times, i.e., the time steps separating the different copies of the static network must be known beforehand for defining the full time-expanded network. If any of the transportation arcs use low-

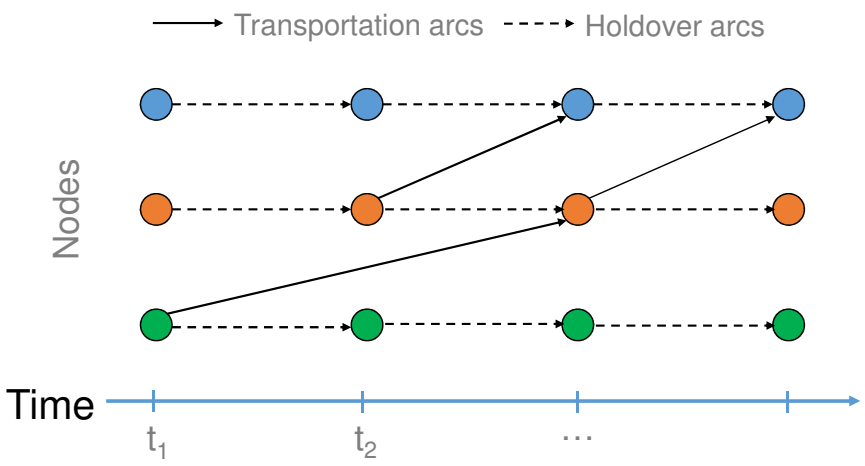

Figure 5.2: Full time-expanded network.

thrust propulsion, using TE-GMNCF will require fine/adaptive time-step discretization to model the corresponding costs correctly. However, this is not computationally efficient as it increases the problem size and causes scaling issues. Thus a new framework is developed to capture the time costs of low-thrust arcs efficiently.

5.2 Event-Driven Network Model

5.2.1 General MILP Formulation

Instead of expanding the dynamic network model along discrete time steps, the time dimension can be introduced through the use of **event-based time steps**, with holdover arcs facilitating the flow between these static "layers". These holdover arcs connect the copies of the same node i across event layers e and $e+1$, where each event e belongs to the set of pre-determined events $\mathbb{E}$. These arcs are always directed forward in time, which means that a holdover arc originating in event layer e can cannot connect to its node copy in event layer $e-1$. The network arcs can thus be separated into the set of intra-layer transportation arcs $\mathbb{A}_T$ and the set of inter-layer holdover arcs $\mathbb{A}_H$, such that $\mathbb{A} = \mathbb{A}_T \cup \mathbb{A}_H$.

Due to this expansion of the static model into the time-dimension, each transportation arc now has an additional index to denote its event layer, i.e., an arc (i, j, v, e) allows commodities to flow from node i to node j $(i \neq j)$ by using propulsion provided by vehicle v in an event layer e.

The concept of multi-graph is used to denote the choice of discrete alternatives for propulsion, i.e., vehicle v is chosen from the set $\mathbb{V}$ denoting the complete available fleet. Holdover arcs, on the other hand, are represented by the index (i, i, e) to denote the arc that connects node i to itself across event steps e and $e+1$. Both categories of arcs are sketched in Figure 5.3.

If k commodities are allowed to flow across the network graph, then $\mathbf{x}^{+}_{ijve}$, $\mathbf{x}^{-}_{ijve}$ and $\mathbf{c}^{+}_{ijve}$ are all k-by-1 vectors. All flows $\mathbf{x}^{\pm}_{ijve}$ are nonnegative. Finally, each node i also has an associated demand vector $\mathbf{d}_{ie}$, which represents the demand or supply of each commodity at node i at event-based time step e. Negative components of $\mathbf{d}_{ie}$ indicate demand, whereas supply is denoted by positive components.

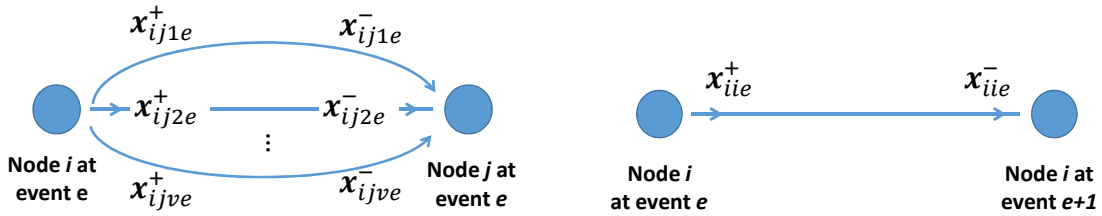

(a) Multi-graph transportation arc (b) Holdover arc

Figure 5.3: Arcs used in the event-driven dynamic network graph formulation.

Using the above developed notation, the event-driven network model can be expressed by the following formulation:

$$\text{Minimize} \quad \mathcal{J} = \sum_{e\in\mathbb{E}} \sum_{(i,j,v,e)\in\mathbb{A}_T} \mathbf{c}^{+}_{ijve}\mathbf{x}^{+}_{ijve} + \sum_{e\in\mathbb{E}} \sum_{(i,i,e)\in\mathbb{A}_H} \cancelto{0}{\mathbf{c}^{+}_{iie}}\mathbf{x}^{+}_{iie} \ , \tag{5.4}$$

subject to the following constraints:

$$\sum_{\substack{(j,v):\\(i,j,v)\in\mathbb{A}_T}} \mathbf{x}^{+}_{ijve} - \sum_{\substack{(j,v):\\(j,i,v)\in\mathbb{A}_T}} \mathbf{x}^{-}_{jive} + \sum_{\substack{i:\\(i,i,e)\in\mathbb{A}_H}} \mathbf{x}^{+}_{iive} - \sum_{\substack{i:\\(i,i,e-1)\in\mathbb{A}_H}} \mathbf{x}^{-}_{ii(e-1)} \le \mathbf{d}_{ie} \quad \forall\ i\in\mathbb{N} \quad \forall\ e\in\mathbb{E} \ , \tag{5.5}$$

$$\mathbf{B}_{ijv}\mathbf{x}^{+}_{ijve} = \mathbf{x}^{-}_{ijve} \qquad \forall\ (i,j,v,e)\in\mathbb{A}_T \quad \forall\ e\in\mathbb{E} \ , \tag{5.6a}$$

$$\cancelto{\mathbb{I}_{k\times k}}{\mathbf{B}_{iie}}\mathbf{x}^{+}_{iie} = \mathbf{x}^{-}_{iie} \qquad \forall\ (i,i,e)\in\mathbb{A}_H \quad \forall\ e\in\mathbb{E} \ , \tag{5.6b}$$

$$\mathbf{C}^{+}_{ijv}\mathbf{x}^{+}_{ijve} \leq \mathbf{s}^{+}_{ijv} \qquad \forall\ (i,j,v,e) \in \mathbb{A}_T \quad \forall\ e \in \mathbb{E}\ , \tag{5.7a}$$

$$\mathbf{C}^{+}_{iie}\mathbf{x}^{+}_{iie} \leq \mathbf{s}^{+}_{iie} \qquad \forall\ (i,i,e) \in \mathbb{A}_H \quad \forall\ e \in \mathbb{E}\ , \tag{5.7b}$$

$$\mathbf{x}^{\pm}_{ijve} \geq \mathbf{0}_{k\times 1} \qquad \forall\ (i,j,v,e) \in \mathbb{A}_T \quad \forall\ e \in \mathbb{E}\ , \tag{5.8a}$$

$$\mathbf{x}^{\pm}_{iie} \geq \mathbf{0}_{k\times 1} \qquad \forall\ (i,i,e) \in \mathbb{A}_H \quad \forall\ e \in \mathbb{E}\ , \tag{5.8b}$$

$$y^{\pm}_{ijve} = |\mathbf{x}^{\pm}_{ijve}| \qquad \forall\ (i,j,v,e) \in \mathbb{A}_T \quad \forall\ e \in \mathbb{E}\ , \tag{5.9}$$

$$f(\mathbf{x}^{+}_{ijve}) \leq T \qquad \forall\ v \in \mathbb{V}' \subseteq \mathbb{V} \quad \forall\ e \in \mathbb{E}' \subseteq \mathbb{E}\ . \tag{5.10}$$

Equation (5.4) is the cost of transporting commodities across the network, plus the cost of holding stock at nodes across layers (which is zero in this case). These commodities can be integer-valued, binary-valued or continuous in nature, for example, the commodities representing the tugs, CSM, or LM take only integer values (binary specifically) as they are all discrete-sized vehicles. A diagonal matrix $\mathbf{M}$ can be used to convert these integer-valued commodity variables into corresponding masses when required.

As before, the first constraint in Equation (5.5) deals with mass balance by ensuring that the demand at every node is met under the supply conditions. Equation (5.6a) uses the k-by-k mass transformation matrix $\mathbf{B}_{ijv}$ to show gains/losses and account for commodity transformation. This constraint is used to model the propellant consumption by propulsive maneuvers required for arc transport. No mass transformation occurs over holdover arcs in the considered case, which is reflected by setting $\mathbf{B}_{iie} = \mathbb{I}_{k\times k}$ in Equation (5.6b), thereby enforcing $\mathbf{x}^{+}_{iie} = \mathbf{x}^{-}_{iie}$.

Concurrency constraints over transportation arcs are handled by Equation (5.7a) and over holdover arcs by Equation (5.7b). If n_c such constraints have to be enforced, then the matrix $\mathbf{C}^{+}_{ijv}$ (and $\mathbf{C}^{+}_{ii}$) will be of the size n_c-by-k and $\mathbf{s}^{+}_{ijv}$ (and $\mathbf{s}^{+}_{ii}$) a vector of size n_c-by-1. These constraints vary across arcs depending on the layer that they belong to, and will be listed in detail after the event-based layers are introduced. Similarly, while Equations (5.8a)-(5.8b) represent the nonnegative nature of flows across all arcs, they are set to be equal to zero (instead of greater-than-or-equal to zero) for certain commodities depending on the event layer, for example, tugs are not allowed to flow in event layers corresponding to the crew missions.

The last constraint in Equation (5.10), which can be expanded to include multiple constraints, represents specific time-related bounds that need to be respected by the campaign. The time-related

constraints are dependent on the total incoming mass on individual arcs, which is calculated by adding the mass of all incoming commodities flowing on any given arc as shown in Equation (5.9). These constraints are appended to the formulation as a means of exploring the cost-versus-time tradespace for the campaign and represent an improvement over the static GMCNF as well as the time-expanded GMCNF models. The form that these constraints assume for the current case study is explained after the event layers considered for it are defined in the following section.

5.2.2 Event-Based Time Steps for Dynamic Network Modeling of Case Study

The static network model lacks any mechanism for dealing with dynamic behaviors and can result in erroneous flow generation loops. As opposed to the existing dynamic variant [9] where discrete time steps are used to expand the static model, event-based time steps are used in the current work to allow easy integration of low-thrust vehicles flight times which are not known *a priori* and thus preclude the use of discrete time step sizes. Figure 5.4 shows the breakdown of event-based time steps for the current case study, where four event steps are used for each use of the cargo tugs and two additional event steps for every crew mission. Each of these layers represents a copy of the complete static network – they span a variable time period that is calculated and adjusted internally to meet the overall time constraints.

The first four static layers or event steps (E^{F1}_{cargo}, E^{F2}_{cargo}, E^{R1}_{cargo}, and E^{R2}_{cargo}) permit only cargo deliveries by tugs chosen from the available fleet. Layers E^{F1-F2}_{cargo} allow only those arcs that move toward the moon, while E^{R1-R2}_{cargo} only allow returning arcs. Restricting transport directions in this way avoids flow generation loops. A demarcation between E^{F1}_{cargo} and E^{F2}_{cargo} layers is further made such that the **longest** path between the source (Earth surface) and the destination (lunar surface) can be simulated between them. Parallel arcs (e.g., LEO $\rightarrow$ EML_1 halo and LEO $\rightarrow$ EML_2 halo) that do not interact with each other are allowed within the same layer. By allowing a single layer to encompass all flow until the end of an event (e.g. forward motion towards the halo orbits, in layer E^{F1}_{cargo}), commodity interaction at nodes can be simulated. For example, if two tugs finish at the same node at the end of an event layer, then they could exchange cargo, and more importantly, they could even exchange propellant if both tugs employ the same type of propulsion systems (high-thrust or low-thrust).

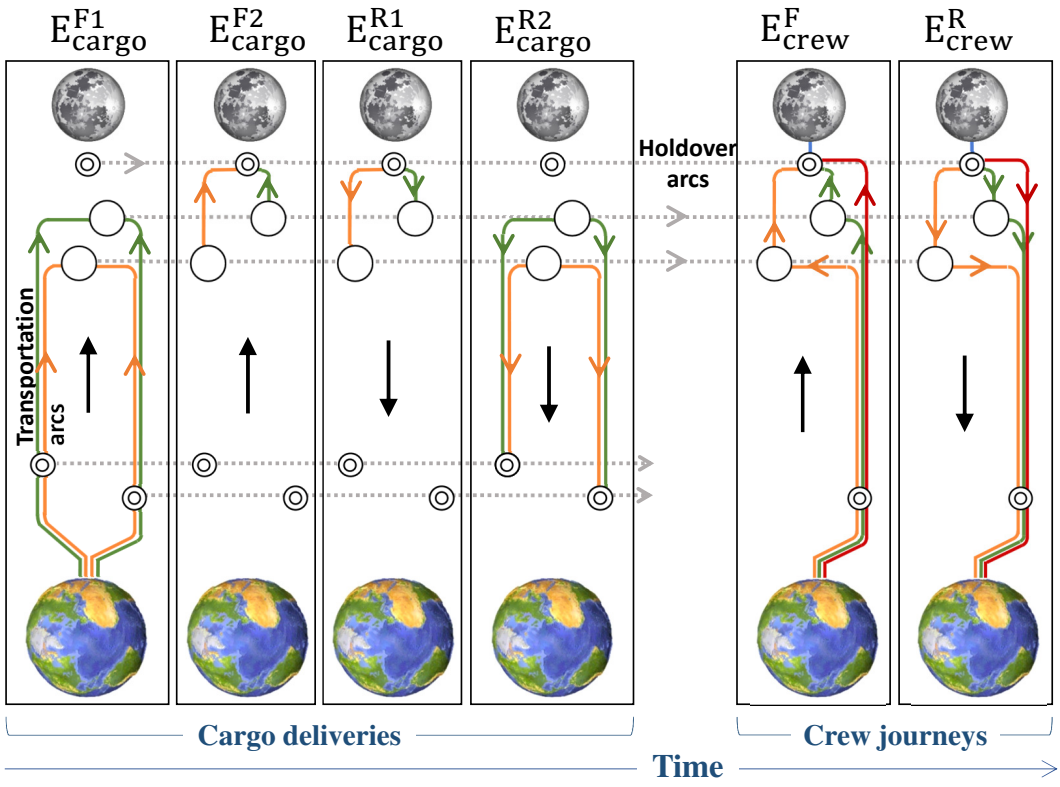

Figure 5.4: Event-driven network for the case study campaign.

Similarly, the crew journeys are also split into layers E^{F}_{crew} and E^{R}_{crew}. The four layers required for the tug flight are collapsed to two layers for crew travel because all the commodities that the crew can interact with (i.e., the droptanks and their contents) have definitively been delivered in the previous layers and the crew vehicle (the CSM) is not allowed to interact with any other commodities.

The layers E^{F1-F2}_{cargo} and E^{R1-R2}_{cargo} are repeated three times in the current problem to correspond to the three uses allowed for each tug. Likewise, the layers E^{F}_{crew} and E^{R}_{crew} are repeated three times to signify the three crew missions that the current campaign is set up for. If $\mathbb{E}_{cargo}$ consolidates all the cargo-related layers such that $\mathbb{E}_{cargo} = E^{F1}_{cargo} \cup E^{F2}_{cargo} \cup E^{R1}_{cargo} \cup E^{R2}_{cargo}$, then the only propulsive elements allowed in the event set $\mathbb{E}_{cargo}$ are the cargo tugs. Along the same lines, only crew vehicles are permitted to provide propulsion along the arcs in the event set $\mathbb{E}_{crew} = E^{F}_{crew} \cup E^{R}_{crew}$. The cargo (i.e. refuel droptanks) delivery-related events are assumed to occur before the crew-related events and these two sets of event layers are connected via the holdover arcs. Together these

two sets ($\mathbb{E}_{\text{cargo}}$ and $\mathbb{E}_{\text{crew}}$) form the set of all events used to discretize the time dimension, i.e., $\mathbb{E} = \mathbb{E}_{\text{cargo}} \bigcup \mathbb{E}_{\text{crew}}$.

Similarly, the vehicles are also divided into two sets – the set of cargo tugs $\mathbb{V}_{\text{cargo}}$ and the set of crew vehicles $\mathbb{V}_{\text{crew}}$ together form the set representing the entire vehicle fleet $\mathbb{V} = \mathbb{V}_{\text{cargo}} \bigcup \mathbb{V}_{\text{crew}}$. The multi-graph is an important modeling feature used in the cargo layers ($\text{E}_{\text{cargo}}^{\text{F1}-\text{F2}}$ and $\text{E}_{\text{cargo}}^{\text{R1}-\text{R2}}$) that aids the correct consolidation of flight times of different tugs across event layers. The exact process of incorporating the time of flight constraints across these layers is detailed in the next section.

5.2.3 Treating Time Constraints in Case Study

Equation (5.10) can be used to place bounds on the total time taken to set up the in-space propellant resupply chain (T_{cargo}) and to constrain the total crew flight time across all missions (T_{crew}) by splitting it into two conditions. The nonlinearity of including low-thrust transfer times into space logistics frameworks arises at this juncture from these two conditions. However, if this system model can be reduced to a set of linear relationships, then mixed-integer linear programming (MILP) techniques can be applied to yield optimal solutions. This section describes in detail the manner of dealing with each time-related constraint.

Time of flight across transportation arcs

The length Δt_{ijve} of a transportation arc (i, j, v, e) is defined as the time taken to traverse that arc by the propulsive vehicle v. This arc length, or the time of flight across the arc, can be expressed as:

$$\Delta t_{ijve} = g(\mathbf{x}_{ijve}^{+}) \ . \tag{5.11}$$

where g is a function that is derived next on a case-by-case basis.

If no vehicle transports any commodity across a particular pair of nodes in the event layer, then the flight time over that arc is simply zero, i.e., $g(\mathbf{x}_{ijve}^{+}) = 0$. In the case that an arc uses one of the available high-thrust vehicles to transport commodities, the arc length is fixed, but can be

modeled by the following linear relationship:

$$g(\mathbf{x}^{+}_{ijve}) = \cancelto{0}{{}^{1}q_{ijve}} \cdot y^{+}_{ijve} + {}^{0}q_{ijve} \cdot \mathrm{x}^{v}_{ijve} \ , \tag{5.12}$$

where $y^{+}_{ijve} = |\mathbf{x}^{+}_{ijve}|$ as shown in Equation (5.9) and x^{v}_{ijve} is the component of the commodities vector $\mathbf{x}^{+}_{ijve}$ specifying the flow of vehicle v over the arc (i, j, v, e). The coefficient ${}^{1}q_{ijve}$ takes a value of zero because high-thrust arc costs are independent of the total mass. The term x^{v}_{ijve} is assigned to be binary-valued as mentioned earlier, and hence the coefficient ${}^{0}q_{ijve}$ simply assumes the value of the high-thrust flight time that is derived from literature.

On the other hand, this arc length is dependent on the total mass of commodities flowing across the arc in the case of low-thrust vehicle. For a general case, a cost model such as the one shown by the dashed blue line in Figure 5.5 can be represented by piecewise linear (PWL) approximations

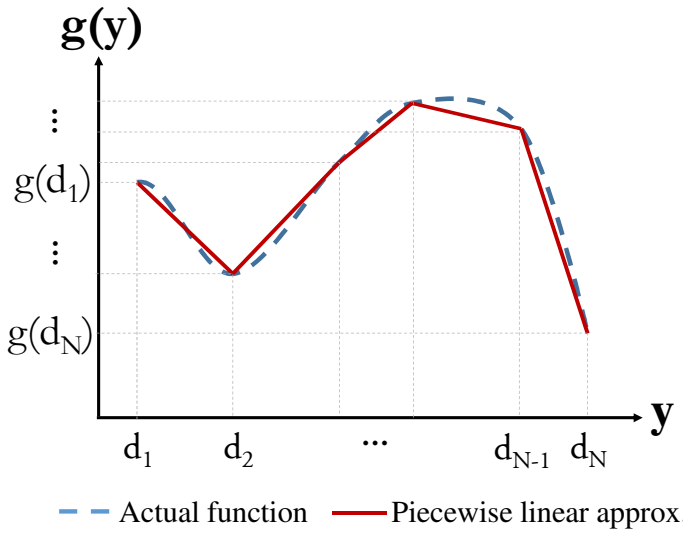

Figure 5.5: Piecewise linear approximation of a nonlinear curve.

by using the breakpoints $\{d_1, d_2, \ ... \ d_N\}$. If we introduce a set of additional continuous variables $\{\lambda_1, \lambda_2, \ ... \ \lambda_N\}$ that belong to an SOS2 set*, then the value of $g(\mathbf{x}^{+}_{ijve})$ in Equation (5.11) can be found as [96, 97]:

$$\lambda_1 + ... + \lambda_N = 1 \ , \tag{5.13a}$$

*An SOS2 (special ordered sets of type 2) constraint basically says that at most two of the λ can be nonzero and these two nonzero elements of the set must be consecutive.

$$\lambda_1 d_1 + ... + \lambda_N d_N = y^+_{ijve} \ , \tag{5.13b}$$

$$\lambda_1 g(d_1) + ... + \lambda_N g(d_N) = g(y^+_{ijve}) \ , \tag{5.13c}$$

$$g(\mathbf{x}^+_{ijve}) = g(y^+_{ijve}) \ . \tag{5.13d}$$

This above set of equations can be used for each arc where a nonlinear cost model is encountered. In the case where Equation (5.13) is used to represent a cost model with PWL functions, the decision vector available to the MILP solver expands from $\boldsymbol{\chi} = \left[\mathbf{x}^{\pm}_{ijve},\ y^+_{ijve},\ \mathbf{x}^{\pm}_{iiv}\right]^T$ to $\boldsymbol{\chi} = \left[\mathbf{x}^{\pm}_{ijve},\ y^+_{ijve},\ \mathbf{x}^{\pm}_{iiv},\ \boldsymbol{\lambda}_{ijve}\right]^T$ where the vector $\boldsymbol{\lambda}_{ijve}$ contains the variables needed for the PWL representation of arc costs. Additionally, the constraints in Equation (5.13) will be appended to the network formulation in Equations (5.4)–(5.10) in such a case.

Although this above PWL formulation can be used in more general cases, the propellant mass and flight time costs in the current case study can be related to the initial mass on the arc using the linear relationships derived in Equations (3.1a)–(3.1b) (in Chapter 3). The time of flight of vehicle v across a given arc (i, j, v, e) can thus be modeled by modifying Equation (3.1b) as:

$$g(\mathbf{x}^+_{ijve}) = {}^1q_{ijve} \cdot y^+_{ijve} + {}^0q_{ijve} \cdot \mathrm{x}^v_{ijve} \ , \tag{5.14}$$

where the coefficients ${}^1q$ and ${}^0q$ for each low-thrust arc (i, j, v, e) were calculated and organized in Table 3.7.

With this, all the relationships required to calculate the arc length according to Equation (5.11) are now available – Equation (5.12) for high-thrust, Equation (5.13) for a generic nonlinear model, and Equation (5.14) for the specific cost model obtained for low-thrust arcs in the current case study campaign.

Cargo delivery duration

The concept essential to incorporating time of flight in the event-driven network formulation is the multi-graph. Consider a single pair of nodes within any layer in the set $\mathbb{E}_{\text{cargo}}$ – if multiple arcs connect them in the form of a multi-graph, then each arc in the multi-graph set is assigned to a single tug unit. The implication is that the propulsion required to traverse a single arc among the

multi-graph arcs between the same pair of nodes can only be provided by a particular tug unit.

If the set of arcs that belong to event layer e is $\mathbb{A}_e$ (such that $\mathbb{A}_e \subset \mathbb{A}_T$), then the total time of flight t_{ve} of vehicle v over all the arcs within the layer e is denoted as:

$$t_{ve} = \sum_{(i,j):(i,j,v,e)\in\mathbb{A}_e} \Delta t_{ijve} \quad \forall v \in \mathbb{V}_{\text{cargo}} \quad \forall e \in \mathbb{E}_{\text{cargo}}. \tag{5.15}$$

where Δt_{ijve} is the arc length derived earlier in Equation (5.11). From here, the total duration t_e of a single static layer e within the dynamic network can be expressed as the maximum of sum of flight times of all tugs flowing in that layer:

$$\max_{v} \left(t_{ve}\right) = t_e \ . \tag{5.16}$$

The above equality constraint is converted into a linear inequality through the introduction of additional continuous variables $\widetilde{t}_e$ (one per event layer) and then constraining these variables to be greater than each tug's total time of flight within the corresponding event layer $e \in \mathbb{E}_{\text{cargo}}$ as:

$$t_{ve} \leq \widetilde{t}_e \ . \tag{5.17}$$

The desired bound on the duration of the campaign's cargo delivery phase can now be implemented as:

$$\sum_{e\in\mathbb{E}_{\text{cargo}}} \widetilde{t}_e \leq T_{\text{cargo}} \ . \tag{5.18}$$

Thus, by deriving the relationship between $\widetilde{t}_e$ and $\mathbf{x}^{\pm}_{ijve}$, this above Equation (5.18) can be correlated to the form specified in Equation (5.10). For the current case study campaign, this is done by condensing Equation (5.12) and Equation (5.14) into Equations (5.15)-(5.18) as:

$$\left\{ \sum_{(i,j):(i,j,v)\in\mathbb{A}_e} \left({}^{1}q_{ijve} \cdot y^{+}_{ijve} + {}^{0}q_{ijve} \cdot \mathrm{x}^{v}_{ijve} \right) \right\} \leq \widetilde{t}_e \quad \forall\ v \in \mathbb{V}_{\text{cargo}} \quad \forall\ e \in \mathbb{E}_{\text{cargo}} \ , \tag{5.19a}$$

$$\sum_{e\in\mathbb{E}_{\text{cargo}}} \widetilde{t}_e \leq T_{\text{cargo}} \ . \tag{5.19b}$$

This provides the equations that account for the duration of the first phase of the case study

campaign by replacing Equation (5.10) in the ED-GMNCF formulation.

Crew time of flight

Deriving a constraint on the total allowable crew time of flight across all missions is done in a similar manner and begins by simply calculating the total time of flight of the CSMs across all three missions. Every crew layer arc uses high-thrust propulsion, and thus corresponding arc lengths are given by:

$$\Delta t_{ijve} = \cancel{{}^{1}q_{ijve}}^{0} \cdot y^{+}_{ijve} + {}^{0}q_{ijve} \cdot \mathrm{x}^{v}_{ijve} \ , \tag{5.20}$$

where v = CSM only and $e \in \mathbb{E}_{\text{crew}}$. Due to the lack of alternative in-space crew vehicles in the analysis, the $\max_{v}(t_{ve})$ relationship used in Equation (5.16) can be discarded and the time constraint is formulated directly as:

$$\sum_{e\in\mathbb{E}} \sum_{(i,j):(i,j,v)\in\mathbb{A}_e} \left({}^{0}q_{ijve} \cdot \mathrm{x}^{v}_{ijve}\right) \leq T_{\text{crew}} \quad \forall\ v \in \mathbb{V}_{\text{crew}} \quad \forall\ e \in \mathbb{E}_{\text{crew}} \ . \tag{5.21}$$

5.2.4 Complete Event-Driven Dynamic Network Formulation

At this point, the derivation of the mathematical model is complete and can be supplied to a MILP optimizer. The formulation for the event-driven dynamic network model is given below by collating Equations (5.4)-(5.21), where the time-related constraints in the last three equations are specific to the case study:

$$\min_{\boldsymbol{x}} \quad \mathcal{J} = \sum_{e\in\mathbb{E}} \sum_{(i,j,v,e)\in\mathbb{A}_T} \mathbf{c}^{+}_{ijve}\mathbf{x}^{+}_{ijve} \quad \text{subject to :} \tag{5.22}$$

$$\sum_{\substack{(j,v):\\(i,j,v)\in\mathbb{A}_T}} \mathbf{x}^{+}_{ijve} - \sum_{\substack{(j,v):\\(j,i,v)\in\mathbb{A}_T}} \mathbf{x}^{-}_{jive} + \sum_{\substack{i:\\(i,i,e)\in\mathbb{A}_H}} \mathbf{x}^{+}_{iive} - \sum_{\substack{i:\\(i,i,e-1)\in\mathbb{A}_H}} \mathbf{x}^{-}_{ii(e-1)} \leq \mathbf{d}_{ie} \quad \forall\ i \in \mathbb{N} \quad \forall\ e \in \mathbb{E} \ , \tag{5.23a}$$

$$\mathbf{B}_{ijv}\mathbf{x}^{+}_{ijve} = \mathbf{x}^{-}_{ijve} \qquad \forall\ (i,j,v,e) \in \mathbb{A}_T \quad \forall\ e \in \mathbb{E} \ , \tag{5.23b}$$

$$\mathbf{x}^{+}_{iie} = \mathbf{x}^{-}_{iie} \qquad \forall\ (i,i,e) \in \mathbb{A}_H \quad \forall\ e \in \mathbb{E} \ , \tag{5.23c}$$

$$\mathbf{C}^{+}_{ijv}\mathbf{x}^{+}_{ijve} \leq \mathbf{s}^{+}_{ijv} \qquad \forall\ (i,j,v,e) \in \mathbb{A}_T \quad \forall\ e \in \mathbb{E} \ , \tag{5.23d}$$

$$\mathbf{C}^{+}_{iie}\mathbf{x}^{+}_{iie} \leq \mathbf{s}^{+}_{iie} \qquad \forall\ (i,i,e) \in \mathbb{A}_H \quad \forall\ e \in \mathbb{E} \ , \tag{5.23e}$$

$$\mathbf{x}^{\pm}_{ijve} \geq \mathbf{0}_{k\times 1} \qquad \forall\ (i,j,v,e) \in \mathbb{A}_T \quad \forall\ e \in \mathbb{E}\ , \tag{5.23f}$$

$$\mathbf{x}^{\pm}_{iie} \geq \mathbf{0}_{k\times 1} \qquad \forall\ (i,i,e) \in \mathbb{A}_H \quad \forall\ e \in \mathbb{E}\ , \tag{5.23g}$$

$$y^{\pm}_{ijve} = |\mathbf{x}^{\pm}_{ijve}| \qquad \forall\ (i,j,v,e) \in \mathbb{A}_T \quad \forall\ e \in \mathbb{E}\ , \tag{5.23h}$$

$$\left\{ \sum_{(i,j):(i,j,v)\in\mathbb{A}_e} \left({}^{1}q_{ijve} \cdot y^{+}_{ijve} + {}^{0}q_{ijve} \cdot \mathrm{x}^{v}_{ijve} \right) \right\} \leq \widetilde{t}_e \qquad \forall\ v \in \mathbb{V}_{\text{cargo}} \quad \forall\ e \in \mathbb{E}_{\text{cargo}}\ , \tag{5.23i}$$

$$\sum_{e\in\mathbb{E}_{\text{cargo}}} \widetilde{t}_e \leq T_{\text{cargo}}\ , \tag{5.23j}$$

$$\sum_{e\in\mathbb{E}} \sum_{(i,j):(i,j,v)\in\mathbb{A}_e} \left({}^{0}q_{ijve} \cdot \mathrm{x}^{v}_{ijve} \right) \leq T_{\text{crew}} \qquad \forall\ v \in \mathbb{V}_{\text{crew}} \quad \forall\ e \in \mathbb{E}_{\text{crew}}\ . \tag{5.23k}$$

The decision vector available to the MILP-solver to optimize the above model in Equations (5.22) – (5.23) is given by:

$$\boldsymbol{\chi} = \begin{bmatrix} \mathbf{x}^{\pm}_{ijve} \\ y^{+}_{ijve} \\ \mathbf{x}^{\pm}_{iie} \\ \widetilde{t}_e \end{bmatrix}\ , \tag{5.24}$$

where $\mathbf{x}^{\pm}_{ijve}$ represents the commodity flow over transportation arcs, y^{+}_{ijve} gives the total initial mass across transportation arcs, $\mathbf{x}^{\pm}_{iie}$ is the flow over holdover arcs, and $\widetilde{t}_e$ are the variables introduced to denote the length of each cargo-related event layer. This decision vector will contain $\boldsymbol{\lambda}_{\text{ijve}}$ only if PWL approximations were adopted for deriving the arc costs.

5.3 Parameters for Case Study Campaign

Commodities List

Table 5.1 lists all commodities allowed to flow on the network for the case study campaign. As catalogued in this table, some of the commodities can only take binary values (such as the vehicles used for propulsion over an arc), while others (such as tug fuels and crew vehicles' fuels) can assume the entire continuous range of real nonnegative values. A diagonal mass matrix $\mathbf{M}$ is used to convert the binary-valued commodities to their corresponding masses.

Table 5.1: List of commodities

Commodity name	Abbreviated name	Variable type	Notes
Upper stage structure	strUS	Continuous	-
Upper stage fuel	fUS	Continuous	-
CSM	CSM	Binary	-
CSM fuel	fCSM	Continuous	-
LM	LM	Binary	-
LM fuel	fLM	Continuous	-
Droptank structure	strDtank	Continuous	-
Chemical tug fuel	fHIGH	Continuous	-
SEP tug fuel	fLOW	Continuous	-
Tug #1	tug1	Binary	Chemical tug type 1
Tug #2	tug2	Binary	Chemical tug type 1
Tug #3	tug3	Binary	Chemical tug type 2
Tug #4	tug4	Binary	Chemical tug type 2
Tug #5	tug5	Binary	Chemical tug type 3
Tug #6	tug6	Binary	Chemical tug type 3
Tug #7	tug7	Binary	Chemical tug type 3
Tug #8	tug8	Binary	SEP tug type 1
Tug #9	tug9	Binary	SEP tug type 1
Tug #10	tug10	Binary	SEP tug type 2
Tug #11	tug11	Binary	SEP tug type 2
Tug #12	tug12	Binary	SEP tug type 3

Demand and Supply

The entire vehicle fleet starts at the Earth surface and units are launched according to need, hence this node provides supply of one unit of each tug, and the requisite number of crew vehicles (i.e. one CSM and one LM for each manned mission). The Earth surface also supplies practically infinite amount of all other commodities that are continuous variables (such as the droptank structural mass, upper stage structural mass, and all propellants). Cost coefficients are thus only applied to those arcs that launch commodities from the Earth surface to either LEO or GTO. In order to simulate the disposal of the LM after the surface crew has been transferred back to the CSM, the LLO node is modeled to have a demand for one LM every crew mission. Also, since the LM propellant amount required is fixed and is only consumed during the lunar descent/ascent operations, a corresponding demand for LM fuel is assigned at LLO for each crew mission. This information represents the demand/supply at nodes for use in Equation (5.23a).

Flow Transformation

Commodities undergo transformation as they flow through the network, i.e. propellant is consumed in order to provide transportation across arcs. Since there are two kinds of propulsion available, the transformation matrix in Equation (5.23b) can take the two different forms listed below.

- For high-thrust chemical propulsion, propellant consumption is modeled as an impulsive maneuver. Thus the Δv values gathered from the literature can be used with the Tsiolkovsky rocket equation to write the transformation constraint:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ -\phi & -\phi & 1-\phi \end{bmatrix}_{ijv} \times \mathbf{M} \times \begin{bmatrix} \text{payload} \\ \text{vehicle} \\ \text{propellant} \end{bmatrix}^{+}_{ijve} = \mathbf{M} \times \begin{bmatrix} \text{payload} \\ \text{vehicle} \\ \text{propellant} \end{bmatrix}^{-}_{ijve} , \tag{5.25}$$

  where $\phi = 1 - \exp(-\Delta v / g_0 I_{sp})$, Δv is the change in vehicle's velocity provided by its high-thrust propulsion system, I_{sp} is the corresponding fuel's specific impulse and g_0 is standard gravity.

- For low-thrust solar electric propulsion, Equation (3.1a) is recalled to model the propellant consumption:

$$\left\{ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ {}^{1}p_{ijv}-1 & {}^{1}p_{ijv}-1 & {}^{1}p_{ijv} \end{bmatrix} \times \mathbf{M} - \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & {}^{0}p_{ijv} & 0 \end{bmatrix} \right\} \begin{bmatrix} \text{payload} \\ \text{vehicle} \\ \text{propellant} \end{bmatrix}^{+}_{ijve} = \mathbf{M} \times \begin{bmatrix} \text{payload} \\ \text{vehicle} \\ \text{propellant} \end{bmatrix}^{-}_{ijve} . \tag{5.26}$$

  The above transformation constraints can be modified as follows if using piecewise linear approximations to model the relationship between initial and final mass on a given arc, instead of the linear fit derived in Equation (3.1a). If doing so, the payload and vehicle that comprise the outflow and inflow will be equal to each other, but the following equations can be used to correlate the total outflow and inflow masses:

$$\lambda_1 + ... + \lambda_N = 1 \ , \tag{5.27a}$$

$$\lambda_1 d_1 + ... + \lambda_N d_N = y^{+}_{ijve} \ , \tag{5.27b}$$

$$\lambda_1 h(d_1) + ... + \lambda_N h(d_N) = y^{-}_{ijve} \ , \tag{5.27c}$$

  where the function h represents the relationship between the initial and final mass on the

arc. The breakpoints $\{d_1, ...d_n\}$ and the SOS2 variables $\{\lambda_1, ...\lambda_N\}$ are derived similarly as Equation (5.13). The use of PWL approximations is also demonstrated using a simple toy problem in Section 5.4.3.

- There is no commodity transformation associated with flow over holdover arcs as no propulsion is required to traverse these arcs.

Flow Concurrency

As mentioned earlier, the concurrency constraints in Equations (5.23d)-(5.23e) vary across the event layers, and their general form is listed below.

- Propellant being carried by a spacecraft (tug or crew vehicle) for its own transportation should not be exceed its maximum fuel capacity:

$$\Big[\text{vehicle fuel}\Big]^+_{ijve} \leq M_v^{\text{fcap}} \Big[\text{vehicle}\Big]^+_{ijve} \ , \tag{5.28}$$

  where v = vehicle and M_v^{fcap} = vehicle fuel capacity. When applying this constraint for an arc that transports the crew to beyond TLI, this changes to:

$$\hat{\varepsilon}_{US} \Big[\text{vehicle fuel}\Big]^+_{ijve} \leq \Big[\text{vehicle}\Big]^+_{ijve} \ , \tag{5.29}$$

  where "vehicle" is the upper stage's structural mass, v = crew launch vehicle's upper stage and $\hat{\varepsilon}_{US} = \varepsilon_{US}/(1 - \varepsilon_{US})$. This change is made to accommodate the continuous nature of the upper stage fuel as well as its structural mass. It is modeled in this way to simulate the savings in crew IMLEO achieved from a distributed launch, as opposed to using a single heavy-lift launch vehicle.

- Amount of LM/CSM fuel flowing across any arc (transportation or holdover) is not to exceed the sum of capacity of LM/CSM vehicles and the droptank structural mass:

$$\begin{bmatrix} \hat{\varepsilon}_{\text{LM}} & \hat{\varepsilon}_{\text{CSM}} \end{bmatrix} \begin{bmatrix} \text{LM fuel} \\ \text{CSM fuel} \end{bmatrix}^+_{ijve} \leq \begin{bmatrix} 1 & \hat{\varepsilon}_{\text{LM}} \cdot M_{\text{LM}}^{\text{fcap}} & \hat{\varepsilon}_{\text{CSM}} \cdot M_{\text{CSM}}^{\text{fcap}} \end{bmatrix} \begin{bmatrix} \text{dtank struct} \\ \text{LM} \\ \text{CSM} \end{bmatrix}^+_{ijve} \ , \tag{5.30}$$

where $\hat{\varepsilon}_{\text{LM}} = \varepsilon_{\text{LM}}/(1-\varepsilon_{\text{LM}})$, $\hat{\varepsilon}_{\text{CSM}} = \varepsilon_{\text{CSM}}/(1-\varepsilon_{\text{CSM}})$, $M_{\text{LM}}^{\text{fcap}}$ = LM fuel capacity, $M_{\text{CSM}}^{\text{fcap}}$ = CSM fuel capacity. ε_{CSM} and ε_{LM} represent the structural coefficients for CSM fuel and LM fuel respectively. This above constraint is repeated for holdover arcs also.

- Any commodity can flow freely across a holdover arc. Thus additional constraints are required to ensure that the tug fuel is contained within tugs across events on holdover arcs and is within their fuel capacity:

$$\Big[\text{high-thrust tug fuel}\Big]^{+}_{iie} \leq \sum_{v} M_{v}^{\text{fcap}} \Big[\text{vehicle}\Big]^{+}_{iie} \quad \text{where vehicle} \in \text{set of high-thrust tugs} \ , \tag{5.31a}$$

$$\Big[\text{low-thrust tug fuel}\Big]^{+}_{iie} \leq \sum_{v} M_{v}^{\text{fcap}} \Big[\text{vehicle}\Big]^{+}_{iie} \quad \text{where vehicle} \in \text{set of low-thrust tugs} \ . \tag{5.31b}$$

Flow Bounds

In addition to the mass balance, mass transformation, and concurrency constraints, Equations (5.23f)-(5.23g) have to be set to specify non-negative values for all commodity flow. These bounds are set to reflect the following restrictions:

- Cargo delivery layers can only be served by tugs. This means that crew vehicles (LM, CSM and upper stage) are not allowed in the cargo-delivery event layers. The opposite is also true where tugs and tug propellants are not allowed in the crew layers.
- The only commodities that can act as "payloads" in the cargo delivery layers are the droptank structural mass, LM fuel and CSM fuel. This ensures that tugs are not piggybacking on each other.
- In the cargo delivery events steps, each arc in a multi-graph set allows only one tug and its corresponding fuel type, along with corresponding payload commodities.
- The launch vehicle upper stage (and its corresponding fuel) is only allowed on the arcs that take the crew from LEO to TLI.

When a commodity k is not allowed on a given arc, its flow is set as $x_{ijve}^{k} = 0$ within the vector $\mathbf{x}^{+} \pm_{ijve}$.

5.4 Results

This section presents the results of implementing the event-driven network model for the cislunar propellant resupply chain design. All permitted arcs are recorded in Table 3.7 and Tables 3.4-3.5. It is noteworthy that direct transfers from L_1 to L_2 are not permitted. The MILP problem resulting from Equation (5.22)-(5.23) is solved in MATLAB using the Gurobi 6.5 solver on an Intel Core i7-4770 3.4 GHz platform. The results are compared against a baseline case consisting of no-refuel 3× Apollo missions by setting T_{cargo} to be zero and $T_{\text{crew}} = 21$ days (7 days for each crew mission) within the ED-GMNCF framework. This **baseline cost is IMLEO = 372.671 t** – all optimal solutions obtained are expressed as percentage improvement over this value.

5.4.1 MILP-based Results

In order to demonstrate the true tradeoffs between cost and time, a Pareto front is manually generated for discrete values of maximum time bounds on crew flight times, and in steps of 4 months for maximum allowable time for cargo deliveries. This triple-objective front is displayed in Figure 5.6 – the colorbar represents the total crew flight time across all three human missions. New options exploring in-space refueling architectures are available along this front. As expected, cost and time display a competing nature where more savings can be realized as the time constraint is relaxed.

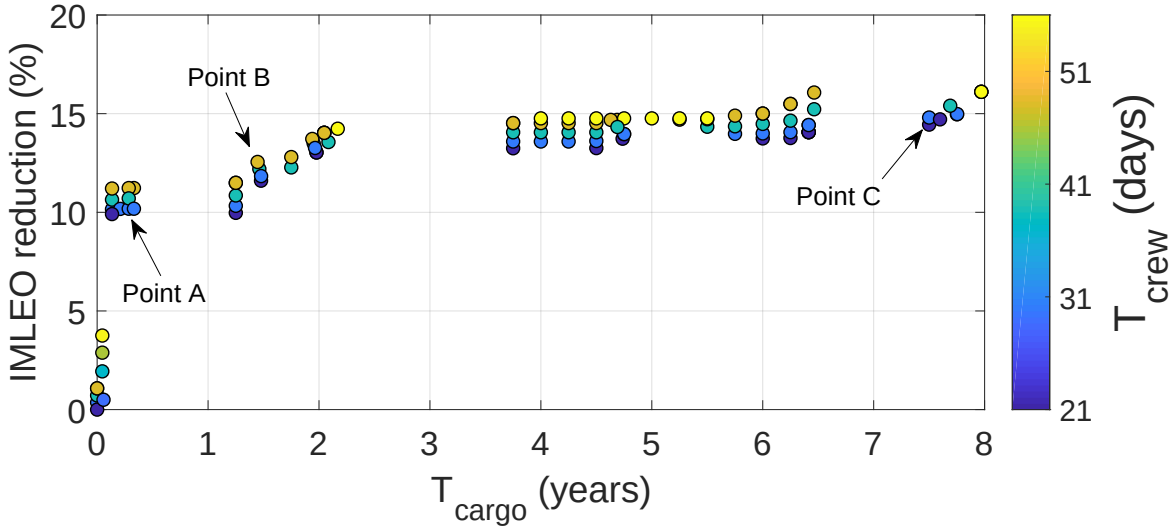

Figure 5.6: Pareto front obtained from the event-driven GMCNF model.

This Pareto front of optimal solutions is examined in detail by focusing on campaign profiles depicting characteristics obtainable only by using the MILP-based ED-GMCNF formulation developed in this work. These points are marked out in Figure 5.6 and the breakdown of commodities transport across two of these network solutions is recorded in Tables 5.2-5.3. The general trend is that architectures use chemical propulsion tugs only at the left-end of the Pareto front (Point A), and then add solar electric propulsion tugs as the axis of time duration is broadened (Points B–C). Some salient features of the solution points on the Pareto front are:

1. Cargo relaying: In some solutions (such as Points A–B), two tugs are used for cargo deliveries, but one of them is launched with minimal cargo to a way station in order to relay to a farther node the cargo that the other tug delivers. This behavior is further explored and explained later in the subsequent text.
2. SEP tug reuse: Each tug itself can be used upto a total of three times, by returning to its Earth parking orbit for refueling. The tug would however have to carry additional propellant to facilitate its return, apart from thepropellant required to make its cargo deliveries. The advantage, on the other hand, is that the architecture does not have to launch the dry mass of the tug into Earth orbit again, thereby driving the IMLEO down. This solution is seen in Point C, where the architecture only uses the smallest SEP tug, but reuses it once.

The specific Pareto front is dependent on the vehicle fleet chosen for the problem and the Δv costs. Let us now consider the solutions points marked in Figure 5.6 in depth. The commodities flowing across all arcs for two chosen solutions are shown in Table 5.2 and Table 5.3. The arcs are listed clearly and are tagged with the event-based layers that they belong to; as a reminder, layers 1 through 12 contain arcs that deliver cargo (i.e. the refuel droptanks), while the succeeding layers simulate commodities flow with respect to crew journeys. The reader can also refer to Table 5.1 to understand the variable type used for each commodity; in general, all vehicles except the upper stage have discrete sizes and hence are modeled as integer variables, specifically as binary variables.

Point A

The solution architecture in Point A depicts the advantages of relaying cargo with different tugs in a very prominent manner. Looking at Table 5.2, the reader can notice that two chemical tugs

are used to carry out all the cargo deliveries – Tug #2 (chemical propulsion tug type 1), which is the smallest vehicle of its kind, and Tug #7 (chemical propulsion tug type 3), which is the largest high-thrust tug. The flow of commodities is illustrated in Figure 5.7.

Table 5.2: Solution Point A from Figure 5.6
(integer variables specified in number of units used and continuous variables in kg).

Layer	Arc Name	CSM	LM	strUS	fCSM	fLM	fUS	strDtank	fHIGH	tug2	tug7	TOF
Arcs within cargo delivery-related event layers (E1–E12)												
1	ES to LEO	0	0	0	14875	33140	0	4175	77780	1	1	0
1	LEO to L1	0	0	0	13735	29630	0	3771	68000	0	1	21
1 to 3	L1 to L1	0	0	0	13735	29630	0	3771	2988	0	1	-
1	LEO to L2	0	0	0	1139	3510	0	404	9780	1	0	17
1 to 2	L2 to L2	0	0	0	1139	3510	0	404	694	1	0	-
2	L2 to LLO	0	0	0	0	3510	0	305	694	1	0	27
2 to 3	LLO to LLO	0	0	0	0	3510	0	305	139	1	0	-
3	LLO to L1	0	0	0	0	0	0	0	139	1	0	28
3 to 10	L1 to L1	0	0	0	13735	29630	0	3771	2988	1	1	-
3 to 10	LLO to LLO	0	0	0	0	3510	0	305	0	0	0	-
10	L1 to LLO	0	0	0	13735	29630	0	3771	2988	1	0	28
Arcs within crew flight-related event layers (E13–E18)												
13	ES to LEO	1	1	4622	6711	0	35986	0	0	0	0	0
13	LEO to TLI	1	1	4622	6711	0	35986	0	0	0	0	0
13	TLI to LLO	1	1	0	6711	0	0	0	0	0	0	4
13 to 14	LLO to LLO	1	0	0	13735	22093	0	4076	0	0	0	-
14	LLO to L2	1	0	0	3364	0	0	0	0	0	0	3.5
14	L2 to ES	1	0	0	1139	0	0	0	0	0	0	8.5
14 to 15	LLO to LLO	0	0	0	10372	22093	0	4076	0	0	0	-
15	ES to LEO	1	1	4622	6711	0	35986	0	0	0	0	0
15	LEO to TLI	1	1	4622	6711	0	35986	0	0	0	0	0
15	TLI to LLO	1	1	0	6711	0	0	0	0	0	0	4
15 to 16	LLO to LLO	1	0	0	10372	11047	0	4076	0	0	0	-
16	LLO to ES	1	0	0	5186	0	0	0	0	0	0	3
16 to 17	LLO to LLO	0	0	0	5186	11047	0	4076	0	0	0	-
17	ES to LEO	1	1	4622	6711	0	35986	0	0	0	0	0
17	LEO to TLI	1	1	4622	6711	0	35986	0	0	0	0	0
17	TLI to LLO	1	1	0	6711	0	0	0	0	0	0	4
17 to 18	LLO to LLO	1	0	0	5186	0	0	4076	0	0	0	-
18	LLO to ES	1	0	0	5186	0	0	0	0	0	0	3
Total IMLEO cost = 334.7268 t, T_{cargo} = 104 days, T_{crew} = 30 days												

Tug #2 launches to the EML_2 halo with the crew propellant in droptanks, drops off all of its payload, continues to the EML_1 halo via the LLO to wait for tug #7. Tug #7 then launches to LEO and travels to EML_1 with the rest of the LM fuel and CSM fuel required for crew resupply, along with required droptank structural mass, bundled as its cargo. Upon reaching the L_1 node, this tug #7 transfers all of its cargo to the awaiting tug #2. More interestingly, tug #2 also receives additional tug propellant from the incoming tug #7. Ultimately, this allows tug #2 to

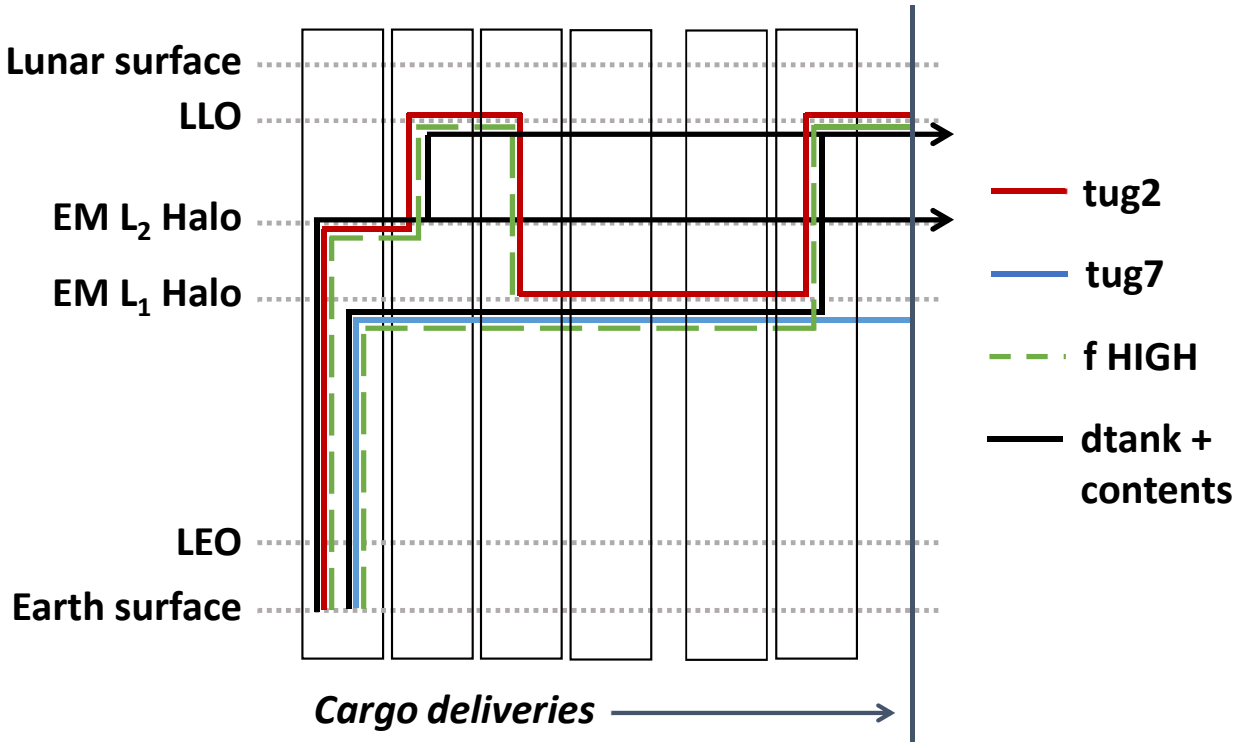

(a) Cargo-related event layers $\mathbb{E}_{\text{cargo}} = \{\text{E1, E2, ... E6}\}$

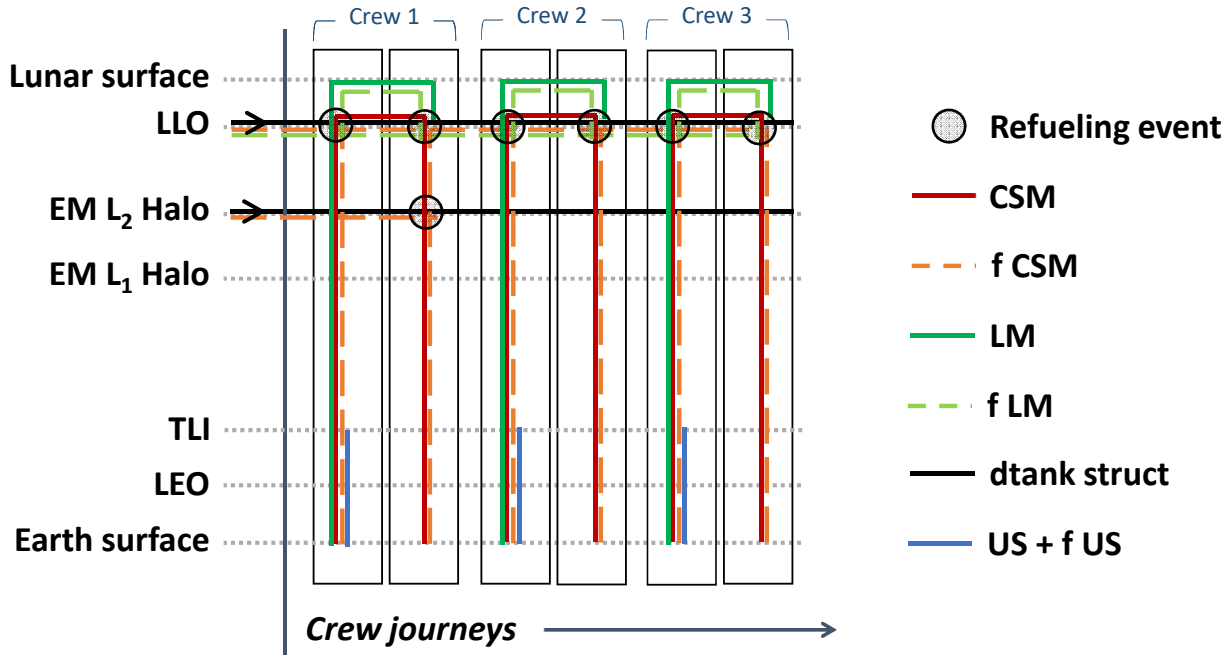

(b) Crew-related event layers $\mathbb{E}_{\text{crew}} = \{\text{E12, E13, ... E17}\}$

Figure 5.7: Commodities flow in point solution A from Figure 5.6.

relay what was initially tug #7's cargo from L_1 to LLO, where the droptank can then flow across holdover arcs to later refuel the incoming crew spacecraft (in layers E13–E18). In this particular point solution, the crew flight time is constrained at 30 days, and hence only one of the three crews will visit EML_2 for refueling on their return journey, while the other two crews will be resupplied at LLO itself and return directly to Earth.

Point B

The minimum crew flight time is 21 days, which accounts for 7 days of flight time for each Apollo-style crew mission. Longer crew flight times are permissible in other solutions, for example in those points marked by say, green and yellow color on the Pareto front in Figure 5.6. Consider Point B as an example of this case, where the crew flight flight is restricted to $\leq$50 days instead of just 21 days. Table 5.3 lists the commodities flow across the network in this solution. This solution has been chosen for showcasing because it exhibits the cooperative use of tugs for cargo relaying again, but with a CP tug and an SEP tug (so tug propellant itself cannot be exchanged at in-space nodes). This relaxation on crew flight time bound allows the crew to return via the L_2 halo repeatedly, because the LLO$\rightarrow L_2 \rightarrow$ Earth surface is cheaper than the direct return. Tug #1 (smallest high-thrust tug variety) launches without any cargo payload, but with a full tank of tug propellant to the L_1 node. Here, it waits for tug #10 (mid-sized SEP tug) to bring it the total crew resupply droptanks, and then ferries these droptanks to LLO and further to L_2. This trajectory route is possible due to the low-energy pathways that are exploited in this analysis.

Point C

Points C lies farther on the right side of the Pareto front, where long campaign durations are allowed. However, its uniqueness lies in the reuse of a tug, which in this case is the smallest SEP tug available within the vehicle fleet. Since the crew flight time is constrained to 21 days in this solution, all crew propellant resupply has to occur at the LLO itself. Thus, the tug #8 delivers this entire droptank cargo to LLO in two journeys. The time duration to set up this resupply chain however is very long because of the lower thrust that can be provided by tug #8, as compared to tug #10 used in Point B.

Table 5.3: Solution Point B from Figure 5.6
(ifnteger variables specified in number of units used and continuous variables in kg).

Layer	Arc Name	CSM	LM	strUS	fCSM	fLM	fUS	strDtank	fHIGH	fLOW	tug1	tug10	TOF
Arcs within cargo delivery-related event layers (E1–E12)													
5	ES to GTO	0	0	0	13509	33140	0	4056	0	8034	0	1	0
5	ES to LEO	0	0	0	0	0	0	0	10737	0	1	0	0
5	LEO to L1	0	0	0	0	0	0	0	10737	0	1	0	21
5	GTO to L1	0	0	0	13509	33140	0	4056	0	8034	0	1	473
5 to 6	L1 to L1	0	0	0	13509	33140	0	4056	3769	0	1	1	-
6	L1 to LLO	0	0	0	13509	33140	0	4056	3769	0	1	0	28
6 to 7	L1 to L1	0	0	0	0	0	0	0	0	0	0	1	-
6 to 7	LLO to LLO	0	0	0	13509	33140	0	4056	533	0	1	0	-
7	LLO to L2	0	0	0	3418	0	0	297	533	0	1	0	27
Arcs within crew flight-related event layers (E13-E18)													
13	ES to LEO	1	1	4622	6711	0	35986	0	0	0	0	0	0
13	LEO to TLI	1	1	4622	6711	0	35986	0	0	0	0	0	0
13	TLI to LLO	1	1	0	6711	0	0	0	0	0	0	0	4
13 to 14	L2 to L2	0	0	0	3418	0	0	297	0	0	0	0	-
13 to 14	LLO to LLO	1	0	0	10091	22093	0	3759	0	0	0	0	-
14	LLO to L2	1	0	0	3364	0	0	0	0	0	0	0	3.5
14	L2 to ES	1	0	0	1139	0	0	0	0	0	0	0	8.5
14 to 15	L2 to L2	0	0	0	2279	0	0	297	0	0	0	0	-
14 to 15	LLO to LLO	0	0	0	6727	22093	0	3759	0	0	0	0	-
15	ES to LEO	1	1	4622	6711	0	35986	0	0	0	0	0	0
15	LEO to TLI	1	1	4622	6711	0	35986	0	0	0	0	0	0
15	TLI to LLO	1	1	0	6711	0	0	0	0	0	0	0	4
15 to 16	L2 to L2	0	0	0	2279	0	0	297	0	0	0	0	-
15 to 16	LLO to LLO	1	0	0	6727	11047	0	3759	0	0	0	0	-
16	LLO to L2	1	0	0	3364	0	0	0	0	0	0	0	3.5
16	L2 to ES	1	0	0	1139	0	0	0	0	0	0	0	8.5
16 to 17	L2 to L2	0	0	0	1139	0	0	297	0	0	0	0	-
16 to 17	LLO to LLO	0	0	0	3364	11047	0	3759	0	0	0	0	-
17	ES to LEO	1	1	4622	6711	0	35986	0	0	0	0	0	0
17	LEO to TLI	1	1	4622	6711	0	35986	0	0	0	0	0	0
17	TLI to LLO	1	1	0	6711	0	0	0	0	0	0	0	4
17 to 18	L2 to L2	0	0	0	1139	0	0	99	0	0	0	0	-
17 to 18	LLO to LLO	1	0	0	3364	0	0	0	0	0	0	0	-
18	LLO to L2	1	0	0	3364	0	0	0	0	0	0	0	3.5
18	L2 to ES	1	0	0	1139	0	0	0	0	0	0	0	8.5
Total IMLEO cost = 325.892 t, T_{cargo} = 528 days, T_{crew} = 48 days													

5.4.2 Comparison with Previous MOGA-based Method

In the previous chapter, the same problem of optimally designing the in-space refueling architectures for Apollo-style missions with cost vs. time tradeoffs was examined. This earlier analysis used a formulation based on multiobjective genetic algorithm (MOGA) where different in-space architectures were represented by genetic sequences. Each gene pair in a fixed-length chromosome was used to indicate if the crew gets refueled at an intermediary node on their way to and back from the lunar surface (i.e. at L_1, L_2 or LLO), and the tug used to deliver the corresponding cargo.

In order to showcase the advantages of using the new framework developed in this chapter, the results are compared with those obtained from the existing MOGA-based formulation. The

problem settings were matched in both methods (i.e. same fleet of vehicles and identical cost models); the Pareto front obtained from the MOGA method is shown in Figure 5.8 overlaid with the one obtained from the MILP-based method. The families of solutions that exist on this MOGA Pareto front have been discussed in detail previously (Chapter 4).

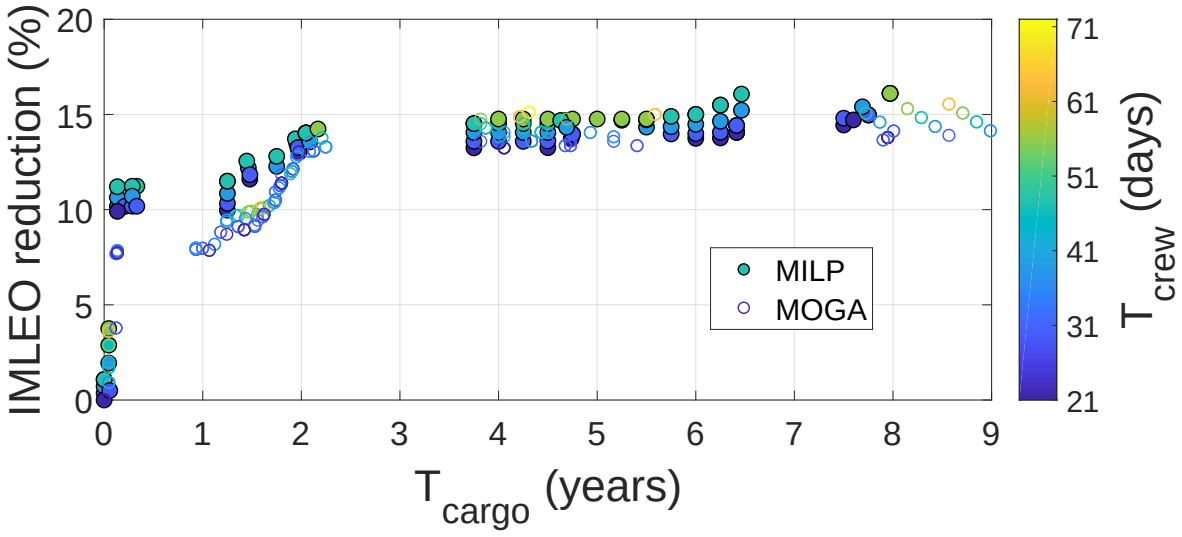

Figure 5.8: Pareto front obtained from GA-based formulation.

The highlight of this comparison is that the solutions obtained from the MILP-based formulation realize higher overall cost savings than the MOGA-based method. One of the core assumptions used in the MOGA-based formulation was that when the crew is refueled along its journey, it is given the exact amount of fuel to reach the next refuel droptank location or to complete the journey. Thus, the flexibility of carrying some of its own fuel was not afforded to the crew missions being resupplied. Furthermore, the MOGA-based formulation is deficient in modeling of cargo/propellant relaying between tugs, and cannot exploit this powerful feature of the optimal solutions as generated by the MILP-based ED-GMCNF formulation. Due to the lack of the rigorous modeling enabled by MILP techniques, the quality of MOGA-based results is lower, i.e., the realizable IMLEO cost savings are decreased due to the limiting assumptions made on propellant refuel amount at way stations. By inherently eliminating restrictive assumptions, the proposed event-driven network approach is capable of exploring a substantially larger tradespace of solutions.

As the time duration allowed for cargo deliveries is increased (i.e., moving towards the right along the Pareto front), more and more cargo is delivered using SEP tugs. Due to this, at some

point, almost all the propellant required to complete the crew journey is provided to it at in-space waystations, instead of carrying it along. In such cases, the SEP tugs are chosen to provide high-efficiency transportation, as long as the allowed campaign duration is long enough. Hence, the solutions obtained from the event-driven network method approach those obtained from the MOGA-based formulation as we move further to the right along the time axis.

Due to the more general problem formulation, the event-driven network method represents a substantial enhancement over the previous MOGA-based formulation. Usually an improvement in fidelity comes at the cost of computational expense. However, in this case, the increase in computational resources required to generate the Pareto front are insignificant. While a single solution with the MILP-based event-driven network method takes only $\sim$3 seconds on a desktop computer, the manual generation of the entire Pareto front takes about 10 minutes. The MOGA, on the other hand, completes its run in little over a minute (200 generations using population size of 300 candidates), but is stochastic in nature and is often sensitive to the initial population of candidate solutions used. The deterministic nature of MILP-techniques is another advantage over metaheuristic optimization methods such as MOGA. Finally, optimality is guaranteed within the specified tolerance when using MILP techniques.

5.4.3 Piecewise Linear (PWL) Approximation of Low-Thrust Arc Costs

The results presented in this section so far have demonstrated the utility of the ED-GMCNF formulation when using the linear fits for low-thrust arc costs as derived in Chapter 3. In order to clarify the applicability of this framework when using more general nonlinear curves to represent cost models, results from a small toy problem are presented here that employs PWL approximations of the cost curves (instead of a linear fit).

Let us consider a simplified problem of setting up the cislunar propellant resupply chain for two Apollo-style crew missions, where only one tug of SEP type 1 is allowed in the analysis (instead of 12 tugs as before). This tug can only be used once and is only allowed to traverse the GTO to L_1 arc. The ED-GMCNF framework is then used to calculate optimal campaign architectures for this toy problem.

The high-thrust crew vehicle costs are calculated as before; no high-thrust tugs are allowed

in this toy problem. The difference arises when modeling the low-thrust costs – the piecewise linear relationships expressed in Equations (5.13) and (5.27) are used to formulate the time and concurrency constraints instead. Table 5.4a is populated from the linear fit coefficients derived in Table 3.6 and 3.7 earlier, while the piecewise linear representation of the same curve uses the values specified in Table 5.4b. The breakpoints for the PWL representation are chosen randomly in the present toy problem; the interested reader is referred to existing literature such as [98] for methods regarding the optimal selection of such breakpoints.

Table 5.4: Linear vs. piecewise linear approximations for low-thrust arc costs.

(a) Linear fit coefficients w.r.t. initial mass (in t)

	Coefficient	Value
Flight time (in days)	${}^1q$	25.98
	${}^0q$	26.631
Final mass (in t)	${}^1p$	0.8757
	${}^0p$	-0.0038

(b) PWL coefficients w.r.t. initial mass (in t)

Breakpoints in initial mass	TOF at breakpoints (in days)	Final mass at breakpoints (in t)
0	0	0
3.9850	0	0
3.9855	127.2829	3.5
47.4639	1272.01669	41.5
86.2539	2272.88475	75.5

The results from both approximations of the low-thrust cost models are presented in Table 5.5 where the T_{cargo} and T_{crew} are restricted to be under 1000 days and 20 days respectively. The campaign architecture from the two methods yield similar results in terms of IMLEO and T_{cargo}.

Table 5.5: Results of linear vs. PWL approximations for low-thrust arc costs.

	IMLEO (t)	T_{cargo} (days)	T_{crew} (days)
Linear fit	242.673	958	18
PWL fit	242.752	968	18

These results are meant to show that PWL approximation for low-thrust costs can easily be implemented within the ED-GMCNF formulation with minor modifications regarding the concurrency and time constraints. Currently, the error in the optimization results cannot be correlated to the error from the SEP cost curve fitting from either of these treatments. This task is recognized as a potential avenue for future work, as is the methodological selection of breakpoints for representing any given nonlinear low-thrust cost curve with piecewise linear functions.

5.5 Chapter Summary

This chapter proposes a new event-driven generalized multi-commodity network flow model to incorporate dynamic behaviors arising from use of low-thrust propulsion in the design of space logistics. Dynamic behavior is added to the static network model through the use of event-based time steps, instead of discrete-sized time-steps as has been done before. These event-based layers are more amenable to the problem of incorporating low-thrust cost models into the network-based formulation. Events such as end-of-forward-journey and completion-of-return-journey are used to demarcate the different static network layers, with only holdover arcs connecting them.

The utility of the developed framework is illustrated by applying it to a campaign to design a cislunar crew propellant resupply chain by using tugs with high-thrust propulsion and low-thrust propulsion cooperatively. The resupply chain is simulated to serve three Apollo-type crew missions through the use of disposable droptanks delivered to their operational locations using the fleet of tugs. Optimal solutions found for a range of time-of-flight constraints for cargo and crew are used to populate the Pareto front. All architectures are compared against the baseline case of no-refuel 3× Apollo missions. The most notable feature of the Pareto front solutions is the optimal relaying of cargo between different sets of tugs to collectively set up the in-space crew propellant resupply chain. Another highlight is tug reuse when long cargo delivery times are allowed; tugs are considered reused when they are return to their Earth parking orbit to be refueled and collect more cargo.

Finally, the solutions are also compared with architectures suggested by in Chapter 4. The current model eliminates certain limiting assumptions from the model in that chapter, thus making the problem formulation more general and exploring a wider tradespace of campaign architectures. The increase in computational effort to populate the Pareto front is insubstantial, and worthwhile because of the higher quality solutions that can be obtained.

Chapter 6

Conclusions

The research in this dissertation is motivated by the need to incorporate dynamic nonlinear behaviors arising from the use of low-thrust propulsion into the design of space logistics and supply chain planning. Low-thrust propulsion is more efficient than current high-thrust technologies in terms of deliverable payload masses, but this advantage is realized at the expense of long flight times. Thus, space logistics tools must be capable of balancing the competing nature of cost, time *and* technology to design the most economically efficient space exploration campaigns.

Chapter 1 presents the necessary background for this problem, along with the objective of this work. Literature review shows that the modeling paradigms underlying current space logistics tools cannot adequately capture low-thrust transportation costs in a computationally efficient manner. This is because these costs are inherently nonlinear, whereas the conventional logistics frameworks are based on linear formulations. This nonlinearity arises from the need to continuously power low-thrust systems for long time periods to impart the Δv necessary to complete any mission. This also couples the time of flight (TOF) and the propellant mass fraction consumed, which are the two primary dimensions of cost, to the total spacecraft mass. High-thrust trajectory design, on the other hand, uses the impulsive-burn model and hence remains independent of the total spacecraft mass.

This work specifically examined the space logistics and supply chain planning problem for the **setup of simple in-space refueling infrastructures to support manned missions to the lunar surface**. This case study campaign, described in detail in Chapter 3, was chosen due to its potential for cooperative use of high-thrust chemical propulsion and low-thrust solar electric propulsion technologies. Costs associated with the use of high-thrust propulsion systems can be modeled adequately by collating values available in the literature. However, the costs (i.e. Δv and TOF) of low-thrust transportation were calculated internal to the logistics framework due to

its dependence on the total spacecraft mass. Cargo delivery in the cislunar space can additionally leverage "low-energy" pathways that result from the circular restricted three-body (CR3B) system model. By applying the CR3B model to the Earth-moon system, the invariant manifolds associated with halo orbits emerge as low-energy routes that were employed in the case study campaign to transport cargo. The halo orbits themselves were treated as potential way stations for storage of refuel tanks.

In this research, two new space logistics frameworks that have the critical capability of trading off cost and time with regards to different propulsion technologies were developed and demonstrated on the cislunar case study. In the course of developing these frameworks, research efforts were also successfully invested in developing computationally efficient methods for preliminary design of low-energy low-thrust transfers as specifically required by the case study considered in this work.

6.1 Contributions to Preliminary Low-Energy Low-Thrust Trajectory Design and Future Work

Existing methods for design of low-energy low-thrust (LELT) trajectories are based on solving the corresponding optimal control problem using iterative nonlinear programming techniques, and hence are computationally expensive. These methods also require initial guesses, which are hard to obtain and need considerable manual involvement for each trajectory designed. Hence, they do not lend themselves well to the current analysis' need for quick cost estimates. Chapter 2 describes the new preliminary LELT trajectory design method developed as a part of the current research effort. The LELT trajectory that connects the geocentric or selenocentric orbits to the halo orbit was designed to consist of a powered spiral around the primary body, patched with a coast arc on the halo orbit's invariant manifold. The computational burden of designing continuous low-thrust spiral trajectories was relieved by the use of a non-iterative Lyapunov control law called Q-law; the native Q-law was modified to include third-body gravitational perturbations during system propagation.

This LELT trajectory design problem was parametrized by the Q-law weights and effectivity cutoff, and the manifold patch point. These decision variables are optimized using a particle

swarm algorithm. The efficiency and effectivity of the developed method, labeled QLEPSO, was quantified by comparing it with results and methods that exist in the literature. Close agreement was observed between the QLEPSO results and the solutions obtained through a variety of other stochastic methods as well as gradient-based deterministic optimal methods. Thus, QLEPSO was shown to provide close estimates of optimal LELT trajectories with a smaller computational budget than existing optimal methods.

Future Work

As with the native Q-law for two-body orbital transfers, QLEPSO results can be used for seeding initial guesses within gradient-based optimal control formulations for LELT trajectory design. This progression in terms of improving the optimality of the solution is left open as an avenue of future work. QLEPSO can also be used to aid initial guess generation in multi-body regimes, where extensive parameter sweeps may be necessary as the spacecraft trajectory in different two-body or three-body frames is patched together. We estimate that QLEPSO-generated initial guesses will complement the use of direct transcription techniques such as multiple shooting or finite-burn low-thrust (FBLT) [99] for solving the entire spiral trajectory as an optimal control problem.

6.2 Contributions to Space Logistics Modeling and Future Work

Space logistics modeling begins by representing the space exploration map as a network, with nodes representing destinations or parking orbits, and arcs connecting these nodes representing flow of material between them. The network details of the case study campaign are provided in Chapter 3, such as the specifications of the discrete-sized vehicles allowed in the analysis and attributes of the nodes considered. The development and validation of the LELT trajectory design method using Q-law in Chapter 2 provided the necessary tools for obtaining the cost models for transportation on different network arcs that exist in the case study campaign. High-thrust transfer costs are borrowed from the available literature, while the cost model is derived for low-thrust transfers using the Q-law LELT trajectory design method. These cost models are tabulated in Chapter 3.

Research efforts were then focused on developing formulations that optimize the mixed-integer nonlinear programming problem of designing logistics for space exploration campaigns that include

low-thrust missions. Chapter 4 presents the first framework, where chromosomal representation of the arc parameters such as payload/vehicle combination and origin/destination nodes was devised so that a multiobjective genetic algorithm (MOGA) could be used to automatically explore the tradespace. This methodology was applied to the problem of establishing a propellant resupply trail for a campaign consisting of one to three Apollo-type manned missions to the moon. Optimal campaign profiles were found along the Pareto front, where the initial mass in LEO (IMLEO) cost was traded against the time required to set up the propellant supply chain and duration of crew mission(s). The solutions of special interest used propulsion technology types cooperatively to deliver refueling-related cargo; the IMLEO improvement in these solutions quantifies the impact of employing low-thrust propulsion technology in the case study campaign.

The main drawback of using heuristic optimization algorithms such as a genetic algorithm lies in the stochastic nature of their search techniques and the lack of guarantee of optimality. To address these concerns and remove limiting assumptions, a second framework was developed in Chapter 5 which uses a deterministic mixed-integer linear programming formulation to optimally plan campaign logistics. This new framework builds on the static generalized multicommodity network flow (GMCNF) model by embedding the time dimension through event-driven steps or "layers", instead of discrete-sized time steps as has been done before. This new event-driven GMCNF (ED-GMCNF) captures low-thrust cost models into network-based formulation in a computationally efficient manner. The basic idea is to duplicate the static network at events such as end-of-forward motion and completion-of-return motion, with only holdover arcs connecting the static layers. Each layer has a time duration associated with it, which is adjusted internally to conform to the time-related constraints.

With this complete model formulated, optimal solutions found for a range of flight time constraints for cargo and crew were used to populate the Pareto front for the case study campaign. The most notable feature of the Pareto front solutions is the optimal relaying of cargo (and tug propellant in some cases) between different sets of tugs to collectively set up the in-space crew propellant resupply chain. Such solutions can only be obtained through rigorous mathematical modeling, where technology and trajectories can be traded to provide top-level architectural decisions.

Although the cost models in the current case study campaign depend linearly on the total

mass flowing on the arc, the ED-GMCNF framework is applicable to problems with more general nonlinear costs or constraints. This was demonstrated with a toy problem where piecewise linear approximations were used to develop a modified set of constraints.

The campaign profiles derived during the course of this investigation not only reduce costs but can also be used for recommending roadmaps for technology development, such as in-space propellant storage and transfer, deep-space rendezvous, and solar electric propulsion tugs for cargo delivery. The crew missions in these campaigns require in-space rendezvous, thereby trading the operational simplicity of using heavy lift launchers with that of exploiting the high efficiency of solar electric propulsion tugs, as well as low-energy pathways in this particular case study. Thus, by shifting the focus away from single mission design, this dissertation has demonstrated the potential of campaign-level selection of propulsion technology to economically sustain future manned space exploration.

Future Work

Although this dissertation considers a specific case study, the techniques developed through the course of this work can be applied to a wide range of problems that contain nonlinear technology costs/constraints. Nonlinear cost models can be generically approximated using piecewise linear functions as demonstrated in the toy problem. An important next step could be the improvement (in terms of computational speed) of these piecewise linear approximations of nonlinear cost models with techniques such as those described by [100, 101]. Additionally, this work could also be advanced by developing a method to optimally select the number and location of breakpoints required for effective piecewise linear approximation of the given function [98]. This would also allow the application of the event-driven GMCNF model to other problems such as developing on-orbit servicing paradigms, or to include in-situ resource utilization (ISRU) strategies for future campaigns. In applications to other problems, the selection of event layers in a rigorous manner will also be an open path.

Potential short-term future work with regards to the current case study would involve developing a better, more optimal low-thrust cost model for the network arcs involving cargo delivery. The Q-law LELT trajectory method involves suboptimality in terms of the control profile dictated

by the Lyapunov control law due to the third-body perturbations, as well as the selection of manifold patch points for each network arc. However, the main advantage of this method with respect to campaign design is its continued relevance across different magnitudes of spacecraft thrust acceleration. Current optimal control formulations for this problem suffer from the curse of dimensionality as the number of variables required to feasibly define the trajectory rise with the flight time, which is further dependent on the thrust acceleration. In the case that abundant computing resources are available, hybrid differential dynamic programming (HDDP) [102–104] could be employed to design the multi-revolution spiral required by the ensuing low-energy low-thrust trajectory.

Appendix A

Computing Halo Orbits

Certain periodic orbits in the circular restricted three-body model can be analytically constructed by linearizing the system equations near the Lagrange points. However, most other periodic orbits in this regime must be constructed numerically. In this work, numerical generation of the desired halo orbit is favored due to their accuracy. The symmetry of halo orbits about the XZ-plane requires $y = v_x = v_z = 0$ at the point where the halo orbit achieves its maximum/ minimum out-of-plane displacement – this property can be exploited to compute the orbit. If the period of the orbit is denoted by T_h, then the states at $t = 0$ and $t = T_h$ must match. These conditions can be mathematically translated as:

$$\text{minimize} \quad \mathcal{J} = \sqrt{\Big(x(T_h) - x(0)\Big)^2 + \Big(z(T_h) - z(0)\Big)^2 + \Big(v_y(T_h) - v_y(0)\Big)^2} \tag{A.1a}$$

$$\begin{aligned} \text{subject to} \quad & y(0) = y(T_h) = 0 \\ & v_x(0) = v_x(T_h) = 0 \\ & v_z(0) = v_z(T_h) = 0 \\ & z(0) = A_z \end{aligned} \tag{A.1b}$$

where A_z is the maximum out-of-plane amplitude of the desired halo orbit. Halo orbits are specified in this work through their A_z, but other popular ways to uniquely identify orbits **within the same family** include specifying their Jacobi energy, or the x value at the XZ-plane crossing.

A.1 Direct Transcription

Individual halo orbits are constructed in this work by using direct transcription to convert the continuous optimal control problem in Equation (A.1) into a constrained parameter optimization

problem so that nonlinear programming (NLP) techniques may be employed. This process is illustrated in Figure A.1. The halo orbit is discretized into N segments and the states (position components x, y, z and velocity components v_x, v_y, v_z) at these discrete points form the decision vector, along with the time period of the orbit. The halo orbit is calculated within a purely ballistic system model and hence no controls are included in this decision vector. Thus, the discretized state $\mathbf{X}_i$ $(i = 1, 2, ... N+1)$ in Figure A.1 consists of $[x_i, y_i, z_i, v_{x_i}, v_{y_i}, v_{z_i}]$. The complete decision vector $\boldsymbol{\chi}$ will be composed of all $\mathbf{X}_i$ and T_h.

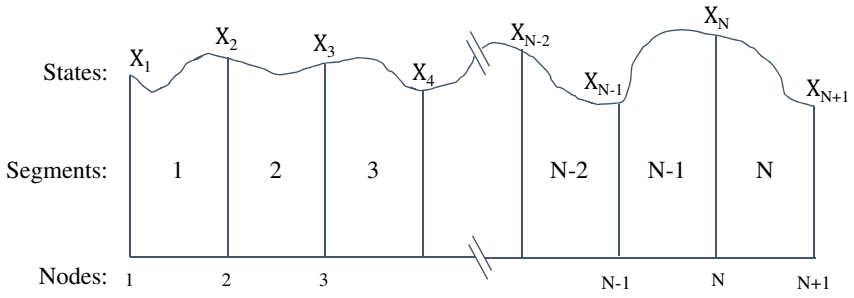

Figure A.1: Direct transcription of continuous optimal control problem

Among available methods, multiple shooting and collocation are the prominently used in similar application as the current halo orbit generation problem. Both these techniques have been implemented and used in this work, though the results from both differ only negligibly.

Multiple Shooting

Within each segment of the discretized trajectory (i.e. the halo orbit), the states at the left boundary X_{i-1} are integrated towards the right boundary using the system differential equations to obtain estimated state $\widetilde{\mathrm{X}}_i$. This mismatch between the estimated state at the end of one segment and the state at the beginning of the next is called the *defect* (Figure A.2) and is added to the constraint vector in Equation (A.1b):

$$\Delta_i = \widetilde{\mathbf{X}}_i - \mathbf{X}_i \tag{A.2}$$

The estimated state at t_i can be calculated as $\widetilde{\mathbf{X}}_i = \boldsymbol{\Phi}(t_{i-1}, t_i)\mathbf{X}_{i-1}$, where $\boldsymbol{\Phi}$ is the state transition matrix as introduced in Chapter 2.

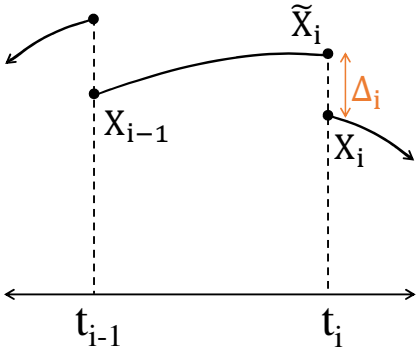

Figure A.2: Multiple shooting constraint formulation

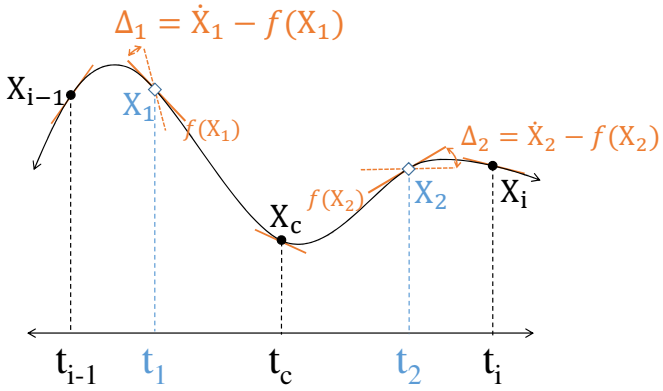

Figure A.3: Fifth-order Gauss-Lobatto constraint formulation

Collocation

Collocation methods use polynomials of selected order to approximate the states within segments and then enforce constraints to match the slope of the polynomial at certain interior points (called collocation points) to the derivative calculated from system equations. These methods are classified according to the number of interior points used and the degree of polynomial chosen. The interior points and the resulting constraints regarding defects for the fifth-order Gauss-Lobatto collocation method [105] are sketched in Figure A.3. In this case, $\mathbf{X}_c$ will also be a part of the decision vector. The interior points in this scheme are given by:

$$\begin{aligned} t_1 &= t_c - \frac{h}{2}\sqrt{\frac{3}{7}} \\ t_1 &= t_c + \frac{h}{2}\sqrt{\frac{3}{7}} \end{aligned} \tag{A.3}$$

and the states at these interior points are expressed as:

$$\begin{aligned} \mathbf{X}_1 = \frac{1}{686}\Big\{&(39\sqrt{21}+231)\mathbf{X}_{i-1} + 224\mathbf{X}_c + (-39\sqrt{21}+231)\mathbf{X}_i + \\ &h\Big((3\sqrt{21}+21)\mathbf{f}_{i-1} - 16\sqrt{21}\mathbf{f}_c + (3\sqrt{21}-21)\mathbf{f}_i\Big)\Big\} \\ \mathbf{X}_2 = \frac{1}{686}\Big\{&(-39\sqrt{21}+231)\mathbf{X}_{i-1} + 224\mathbf{X}_c + (39\sqrt{21}+231)\mathbf{X}_i + \\ &h\Big((-3\sqrt{21}+21)\mathbf{f}_{i-1} + 16\sqrt{21}\mathbf{f}_c + (-3\sqrt{21}-21)\mathbf{f}_i\Big)\Big\} \end{aligned} \tag{A.4}$$

The defects that are set to zero as a part of the constraint vector are:

$$\begin{aligned}\Delta_1 = \frac{1}{360}\Big\{&(32\sqrt{21}+180)\mathbf{X}_{i-1} - 64\sqrt{21}\mathbf{X}_c + (32\sqrt{21}-180)\mathbf{X}_i + \\ &h\Big((9+\sqrt{21})\mathbf{f}_{x-1} + 98\mathbf{f}_1 + 64\mathbf{f}_c + (9-\sqrt{21})\mathbf{f}_i\Big)\Big\} \\ \Delta_2 = \frac{1}{360}\Big\{&(-32\sqrt{21}+180)\mathbf{X}_{i-1} + 64\sqrt{21}\mathbf{X}_c + (-32\sqrt{21}-180)\mathbf{X}_i + \\ &h\Big((9-\sqrt{21})\mathbf{f}_{x-1} + 98\mathbf{f}_2 + 64\mathbf{f}_c + (9+\sqrt{21})\mathbf{f}_i\Big)\Big\}\end{aligned} \tag{A.5}$$

where $\mathbf{f}$ represents the system's first-order differential equations of motion and h is the width of the time segment (i.e. $h = t_i - t_{i-1}$).

A.2 Initial Guess

Once the formulation of the optimal control problem as an NLP problem is complete, the solver needs an initial guess to begin the iterative process of finding the optimal solution. In the current case, the initial estimate for the halo orbit is provided by Richardson's third-order analytical approximation [44]. This approximation holds **valid in close vicinity of the Lagrange points** as it is derived by linearizing the CR3B system equations about L_1 and L_2 and applying the Lindstedt-Poincare method to resulting system. The details of implementing this method for generating the L_1 or L_2 halo orbits are provided below – the interested reader is referred to the original work for the complete derivation.

The 3$^{\text{rd}}$-order solution to the linearized equations that is specific to halo orbits can be expressed as:

$$\begin{aligned}x &= r_1\left\{a_{21}A_x^2 + a_{22}A_z^2 - A_x\cos(\tau) + \Big(a_{23}A_x^2 - a_{24}A_z^2\Big)\cos(2\tau) + \Big(a_{31}A_x^3 - a_{32}A_xA_z^2\Big)\cos(3\tau)\right\}, \\ y &= r_1\left\{kA_x\sin(\tau) + \Big(b_{21}A_x^2 - b_{22}A_z^2\Big)\sin(2\tau) + \Big(b_{31}A_x^3 - b_{32}A_xA_z^2\Big)\sin(3\tau)\right\}, \\ z &= r_1\left\{\delta_m A_z\cos(\tau) + \delta_m d_{21}A_xA_z\Big(\cos(2\tau) - 3\Big) + \delta_m\Big(d_{32}A_zA_x^2 - d_{31}A_z^3\Big)\cos(3\tau)\right\},\end{aligned} \tag{A.6}$$

where $\delta_m = 2 - m$ for $m = 1,\ 3$ *. A_x and A_z represent in-plane and out-of-plane amplitudes of

*$m = 1$ gives the Northern halo orbit family while $m = 3$ gives the Southern halo orbit family

the halo orbit respectively and are constrained by the nonlinear relationship

$$l_1 A_x^2 + l_2 A_z^2 + \Delta = 0 \tag{A.7}$$

where expressions for $l_1,\ l_2$ and Δ will be presented shortly. λ represents both the in-plane and out-of-plane frequency, and ϕ and ψ are phase angles constrained by

$$\psi = \phi + \frac{m\pi}{2}\ . \tag{A.8}$$

The constant k is given by:

$$k = \frac{2\lambda}{\lambda^2 + 1 - c_2}\ . \tag{A.9}$$

The constants $c_2,\ c_3,\ c_4$ are defined by

$$c_n = \frac{1}{\gamma^3}\left[(\pm 1)^n \mu + (-1)^n \frac{(1-\mu)\gamma^{n+1}}{(1 \mp \gamma)^{n+1}}\right] \text{ where n} = 2,\ 3,\ 4 \tag{A.10}$$

where the upper sign applies to the orbits around L_1 and the lower sign for those around L_2. The dimensionless quantity γ is given as:

$$\gamma = \frac{r_1}{a}\ , \tag{A.11}$$

where a is the distance between the two primary bodies and r_1 is the distance between the Lagrange point (L_1 or L_2 only) from the smaller primary, both in normalized synodic units.

Given these terms, the linearized frequency λ can be found as the solution to the following quartic equation:

$$\lambda^4 + (c_2 - 2)\lambda^2 - (c_2 - 1)(2c_2 + 1) = 0\ . \tag{A.12}$$

Finally, the terms in Equation (A.6) and Equation (A.7) are calculated from the relationships below:

$$\Delta = \lambda^2 - c_2$$
$$d_1 = \frac{3\lambda^2}{k}\Big(k(6\lambda^2 - 1) - 2\lambda\Big)$$
$$d_2 = \frac{8\lambda^2}{k}\Big(k(11\lambda^2 - 1) - 2\lambda\Big)$$

$$a_{21} = \frac{3c_3(k^2-2)}{4*(1+2*c_2)}$$
$$a_{22} = \frac{3c_3}{4(1+2c_2)}$$
$$a_{23} = -\frac{3c_3\lambda}{4kd_1}\Big(3k^3\lambda - 6k(k-\lambda) + 4\Big)$$
$$a_{24} = -\frac{3c_3\lambda}{4kd_1}\Big(2 + 3k\lambda\Big)$$
$$a_{31} = -\frac{9\lambda}{4d2}\Big[4c_3(ka_{23} - b_{21}) + kc_4(4+k^2)\Big] + \left(\frac{9\lambda^2+1-c_2}{2d_2}\right)\Big[3c_3(2a_{23} - kb_{21}) + c_4(2+3k^2)\Big]$$
$$a_{32} = -\frac{1}{d_2}\left\{\frac{9\lambda}{4}\Big[4c_3(ka_{24} - b_{22}) + kc_4\Big] + \frac{3}{2}(9\lambda^2+1-c2)\Big[c_3(kb_{22} + d_{21} - 2a_{24}) - c_4\Big]\right\}$$

$$b_{21} = -\frac{3c_3\lambda}{2d_1}\Big(3k\lambda - 4\Big)$$
$$b_{22} = \frac{3c_3\lambda}{d_1}$$
$$b_{31} = \frac{3}{8d_2}\left\{8\lambda\Big[3c_3(kb_{21} - 2a_{23}) - c_4(2+3k^2)\Big] + (9\lambda^2+1+2c_2)\Big[4c_3(ka_{23} - b_{21}) + kc_4(4+k^2)\Big]\right\}$$
$$b_{32} = \frac{1}{d_2}\left\{9\lambda\Big[c_3(kb_{22} + d_{21} - 2*a_{24}) - c_4\Big] + \frac{3}{8}(9\lambda^2+1+2c_2)\Big[4c_3(ka_{24} - b_{22}) + kc_4\Big]\right\}$$

$$d_{21} = -\frac{c_3}{2\lambda^2}$$
$$d_{31} = \frac{3}{64\lambda^2}\Big(4c_3a_{24} + c_4\Big)$$
$$d_{32} = \frac{3}{64\lambda^2}\Big(4c_3(a_{23} - d_{21}) + c_4(4+k^2)\Big)$$

$$s_1 = \frac{1}{2\lambda\big[\lambda(1+k^2) - 2k\big]}\left\{\frac{3c_3}{2}\Big[2a_{21}(k^2-2) - a_{23}(k^2+2) - 2kb_{21}\Big] - \frac{3c_4}{8}\Big(3k^4 - 8k^2 + 8\Big)\right\}$$
$$s_2 = \frac{1}{2\lambda\big[\lambda(1+k^2) - 2k)\big]}\left\{\frac{3c_3}{2}\Big[2a_{22}(k^2-2) + a_{24}(k^2+2) + 2kb_{22} + 5d_{21}\Big] + \frac{3c_4}{8}\Big(12 - k^2\Big)\right\}$$
$$a_1 = -\frac{3c_3}{2}\Big(2a_{21} + a_{23} + 5d_{21}\Big) - \frac{3c_4}{8}\Big(12 - k^2\Big)$$
$$a_2 = \frac{3c_3}{2}\Big(a_{24} - 2a_{22}\Big) + \frac{3c_4}{8}$$
$$l_1 = a_1 + 2\lambda^2 s_1$$

$l_2 = a_2 + 2\lambda^2 s_2$

A.3 Continuation

Once a single solution for a periodic orbit is found, homotopy or continuation techniques can be employed to generate other members of the same family. This is achieved by using the solution for a certain amplitude orbit as an initial guess in calculating the a "nearby" periodic orbit with similar amplitude.

Appendix B

Converting (τ_h, τ_m) to Keplerian orbital elements

The invariant manifold branch associated with a given unstable halo orbit can be parametrized using the halo insertion (or departure) parameter τ_h and the manifold insertion (or departure) parameter τ_m. Any point on the manifold can be reached by perturbing the state along the orbit's eigenvector at point τ_h and propagating it for τ_m synodic time units. This propagation will yield the coordinates of the point in the rotating frame as $\mathbf{r}_r = \begin{bmatrix} x_r, & y_r, & z_r \end{bmatrix}$ and $\mathbf{v}_r = \begin{bmatrix} v_{x_r}, & v_{y_r}, & v_{z_r} \end{bmatrix}$. Before obtaining the set of corresponding Keplerian orbital elements, these coordinates have to be converted to their equivalent inertial state.

Converting from Synodic to Inertial Frame

If the origin of the rotating synodic frame can be expressed in inertial coordinates s $\mathbf{R_i}$, and the transformation matrix $\mathbf{T}$ denotes the rotation of the synodic frame with respect to the inertial frame, then the inertial coordinates of the point under consideration are given as

$$\mathbf{r_i} = \mathbf{R_i} + \mathbf{T} \cdot \mathbf{r_r} \quad \text{and}$$

$$\mathbf{v_i} = \frac{\mathbf{dR_i}}{\mathbf{dt}} + \dot{\mathbf{T}} \cdot \mathbf{r_r} + \mathbf{T} \cdot \mathbf{v_r} \ ,$$

$$\text{where } \mathbf{T} = \begin{bmatrix} \cos\tilde{\theta} & -\sin\tilde{\theta} & 0 \\ \sin\tilde{\theta} & \cos\tilde{\theta} & 0 \\ 0 & 0 & 1 \end{bmatrix} \ .$$

$\tilde{\theta}$ represents the angle between the x-axis of the rotating frame and x-axis of the inertial frame; due to the normalization used, its time rate of change is unity (i.e., $\dot{\tilde{\theta}} = 1$).

Converting $\left[\mathbf{r_i}, \mathbf{v_i}\right]$ to Classical Orbital Elements

Given the inertial state of a spacecraft $\left(\mathbf{r_i}, \mathbf{v_i}\right)$, the classical orbital elements $\left(a,\ e,\ i,\ \omega,\ \Omega,\ \theta\right)$ can be given by:

- Semi-major axis (a)

$$a = \frac{r}{2 - rv^2/\mu_b} \ ,$$

  where $r = |\mathbf{r_i}|$, $v = |\mathbf{v_i}|$ and μ_b is the gravitational parameter of the central primary body.

- Eccentricity (e)

$$\mathbf{e} = \frac{\mathbf{1}}{\mu_\mathbf{b}} \left[\left(\mathbf{v^2} - \frac{\mu_\mathbf{b}}{\mathbf{r}} \right) \mathbf{r} - (\mathbf{r} \cdot \mathbf{v}) \mathbf{v} \right] \ ,$$

  and $e = |\mathbf{e}|$.

- Inclination (i)

  If $\hat{\mathbf{k}}$ is defined as a unit vector in the out-of-plane direction of the considered inertial frame, then a nodal vector can be defined as

$$\mathbf{n} = \hat{\mathbf{k}} \times \frac{\mathbf{h}}{\mathbf{h}} \quad \text{with } \mathbf{n} = |\mathbf{n}| = \sin \mathbf{i} \ ,$$

  where $\mathbf{h}$ is the angular momentum vector that points in the direction normal to the orbital plane. Thus, inclination is given as

$$\cos i = \frac{\mathbf{h} \cdot \hat{\mathbf{k}}}{h} = \frac{h_z}{h} \ ,$$

  where $h = |\mathbf{h}|$. This relationship automatically yields the inclination in the correct range $0 \leq i \leq \pi$.

- Longitude of the ascending node (Ω)

$$\cos \Omega = \frac{\mathbf{n} \cdot \hat{\mathbf{i}}}{n} \ ,$$

  where $0 \leq \Omega \leq \pi$ if $n_y = \mathbf{n} \cdot \hat{\mathbf{j}} \geq \mathbf{0}$.

- Argument of periapse (ω)

$$\cos \omega = \frac{\mathbf{n} \cdot \mathbf{e}}{ne} \ ,$$

where ω lies either in the first- or second-quadrant if $e_z = \mathbf{e} \cdot \hat{\mathbf{k}} \geq \mathbf{0}$

- True anomaly (θ)

$$\cos\theta = \frac{\mathbf{e} \cdot \mathbf{r}}{er} \ ,$$

where θ lies either in the first- or second-quadrant if $\mathbf{r} \cdot \mathbf{v} \geq \mathbf{0}$

Protections to avoid divergence from singularities in eccentricity or any of the angles should be implemented.

References

[1] K. E. Goodliff, B. Mattfeld, C. Stromgren, H. Shyface, and W. Cirillo, "Comparison of Human Exploration Architecture and Campaign Approaches," *AIAA SPACE 2015 Conference and Exposition*, 2015, p. 4413.

[2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin, *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, 1993.

[3] E. Gralla, S. Shull, and O. De Weck, "A Modeling Framework for Interplanetary Supply Chains," *AIAA SPACE Forum*, 2006.

[4] O. L. De Weck, D. Simchi-Levi, R. Shishko, J. Ahn, E. Gralla, D. Klabjan, J. Mellein, A. Shull, A. Siddiqi, B. Bairstow, *et al.*, "Spacenet v1. 3 User's Guide," *NASA/TP*, Vol. 214725, 2007.

[5] C. Taylor, O. d. Weck, and D. Klabjan, "Integrated Transportation System Design for Space Exploration Logistics," *57th International Astronautical Congress*, 2006.

[6] C. Taylor, M. Song, D. Klabjan, O. deWeck, and D. Simchi-Levi, "Modeling Interplanetary Logistics: A Mathematical Model for Mission Planning," *SpaceOps 2006 Conference*, 2006.

[7] T. Ishimatsu, O. L. de Weck, J. A. Hoffman, Y. Ohkami, and R. Shishko, "A Generalized Multi-Commodity Network Flow Model for Space Exploration Logistics," *AIAA SPACE Conference & Exposition*, 2013, 10.2514/6.2013-5414.

[8] T. Ishimatsu, O. L. de Weck, J. A. Hoffman, Y. Ohkami, and R. Shishko, "Generalized Multicommodity Network Flow Model for the Earth-Moon-Mars Logistics System," *Journal of Spacecraft and Rockets*, Vol. 53, No. 1, 2015, pp. 25–38, 10.2514/1.A33235.

[9] K. Ho, O. L. d. Weck, J. A. Hoffman, and R. Shishko, "Dynamic Modeling and Optimization for Space Logistics using Time-Expanded Networks," *Acta Astronautica*, Vol. 105, No. 2, 2014, pp. 428–443, 10.1016/j.actaastro.2014.10.026.

[10] D. Smitherman and G. Woodcock, "Space Transportation Infrastructure Supported by Propellant Depots," *AIAA SPACE Conference & Exposition*, 2011, 10.2514/6.2011-7160.

[11] R. W. Farquhar, "The Utilization of Halo Orbits in Advanced Lunar Operations," tech. rep., NASA Report TN D-6365, 1971.

[12] K. Howell, B. Barden, and M. Lo, "Application of Dynamical Systems Theory to Trajectory Design for a Libration Point Mission," *The Journal of the Astronautical Sciences*, Vol. 45, No. 2, 1997, pp. 161–178.

[13] R. L. Anderson and M. W. Lo, "Role of Invariant Manifolds in Low-Thrust Trajectory Design," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 6, 2009, pp. 1921–1930.

[14] M. W. Lo, "The Interplanetary Superhighway and the Origins Program," *Aerospace Conference Proceedings (2002)*, Vol. 7, IEEE, 2002, p. 3543–3562.

[15] W. S. Koon, M. W. Lo, J. E. Marsden, and S. D. Ross, *Dynamical Systems, the Three-Body Problem and Space Mission Design*. Free online Copy: Marsden Books, 2008.

[16] J. S. Parker and G. H. Born, "Modeling a Low-Energy Ballistic Lunar Transfer Using Dynamical Systems Theory," *Journal of Spacecraft and Rockets*, Vol. 45, 2008, pp. 1269–1281.

[17] E. A. Belbruno and J. K. Miller, "Sun-Perturbed Earth-to-Moon Transfers with Ballistic Capture," *Journal of Guidance, Control, and Dynamics*, Vol. 16, No. 4, 1993, pp. 770–775.

[18] D. C. Folta, M. Woodard, K. Howell, C. Patterson, and W. Schlei, "Applications of Multi-Body Dynamical Environments: the ARTEMIS Transfer Trajectory Design," *Acta Astronautica*, Vol. 73, 2012, pp. 237–249.

[19] M. Lo, B. Williams, W. Bollman, D. Han, Y. Hahn, J. Bell, E. A. Hirst, R. Corwin, P. Hong, K. Howell, B. Barden, and R. Wilson, "Sun-Perturbed Earth-to-Moon Transfers with Ballistic Capture," *The Journal of the Astronautical Sciences*, Vol. 49, No. 1, 2001, pp. 169–184.

[20] J. Senent, C. Ocampo, and A. Capella, "Low-Thrust Variable-Specific-Impulse Transfers and Guidance to Unstable Periodic Orbits," *Journal of Guidance, Control, and Dynamics*, Vol. 28, No. 2, 2005, pp. 280–290.

[21] M. Ozimek and K. Howell, "Low-Thrust Transfers in the Earth-Moon System, including Applications to Libration Point Orbits," *Journal of Guidance, Control, and Dynamics*, Vol. 33, No. 2, 2010, pp. 533–549.

[22] G. Mingotti, F. Topputo, and F. Bernelli-Zazzera, "Combined Optimal Low-Thrust and Stable-Manifold Trajectories to the Earth-Moon Halo Orbits," *AIP Conference Proceedings*, Vol. 886, 2007, pp. 100–112.

[23] G. Mingotti, F. Topputo, and F. Bernelli-Zazzera, "Low-Energy, Low-Thrust Transfers to the Moon," *Celestial Mechanics and Dynamical Astronomy*, Vol. 105, No. 1-3, 2009, p. 61.

[24] G. Mingotti, F. Topputo, and F. Bernelli-Zazzera, "Optimal Low-Thrust Invariant Manifold Trajectories via Attainable Sets," *Journal of Guidance, Control, and Dynamics*, Vol. 34, No. 6, 2011, p. 1644.

[25] G. Mingotti, F. Topputo, and F. Bernelli-Zazzera, "Transfers to Distant Periodic Orbits around the Moon via their Invariant Manifolds," *Acta Astronautica*, Vol. 79, 2012, pp. 20–32.

[26] G. Mingotti, J. Sánchez, and C. McInnes, "Combined Low-Thrust Propulsion and Invariant Manifold Trajectories to Capture NEOs in the Sun–Earth Circular Restricted Three-Body Problem," *Celestial Mechanics and Dynamical Astronomy*, Vol. 120, No. 3, 2014, pp. 309–336.

[27] C. Martin, B. A. Conway, and P. Ibanez, "Optimal Low-Thrust Trajectories to Interior Earth-Moon Lagrange Point," *Space Manifold Dynamics* (E. Perozzi and S. Ferraz-Mello, eds.), ch. 6, Springer, 2010.

[28] C. Martin and B. A. Conway, "Optimal Low-Thrust Trajectories Using Stable Manifolds," *Spacecraft Trajectory Optimization* (B. A. Conway, ed.), ch. 9, Cambridge University Press, 2010.

[29] G. Mingotti and P. Gurfil, "Mixed Low-Thrust Invariant-Manifold Transfers to Distant Prograde Orbits around Mars," *Journal of Guidance, Control, and Dynamics*, Vol. 33, No. 6, 2010, pp. 1753–1764, 10.2514/1.49810.

[30] G. Mingotti, F. Topputo, and F. Bernelli-Zazzera, "Earth–Mars Transfers with Ballistic Escape and Low-Thrust Capture," *Celestial Mechanics and Dynamical Astronomy*, Vol. 110, No. 2, 2011, pp. 169–188, 10.1007/s10569-011-9343-5.

[31] A. E. Petropoulos and J. M. Longuski, "Shape-Based Algorithm for the Automated Design of Low-Thrust, Gravity Assist Trajectories," *Journal of Spacecraft and Rockets*, Vol. 41, No. 5, 2004, pp. 787–796, 10.2514/1.13095.

[32] D. Izzo, "Lambert's Problem for Exponential Sinusoids," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 5, 2006, pp. 1242–1245, 10.2514/1.21796.

[33] B. J. Wall and B. A. Conway, "Shape-Based Approach to Low-Thrust Rendezvous Trajectory Design," *Journal of Guidance, Control, and Dynamics*, Vol. 32, No. 1, 2009, pp. 95–101, 10.2514/1.36848.

[34] E. Vellutini and G. Avanzini, "Shape-Based Design of Low-Thrust Trajectories to Cislunar Lagrangian Point," *Journal of Guidance, Control, and Dynamics*, Vol. 37, No. 4, 2014, pp. 1329–1335, 10.2514/1.G000165.

[35] D. E. Chang, D. F. Chichka, and J. E. Marsden, "Lyapunov-Based Transfer between Elliptic Keplerian Orbits," *Discrete and Continuous Dynamical Systems Series B*, Vol. 2, No. 1, 2002, pp. 57–68, 10.3934/dcdsb.2002.2.57.

[36] A. Petropoulos, "Low-Thrust Orbit Transfers using Candidate Lyapunov Functions with a Mechanism for Coasting," *AIAA/AAS Astrodynamics Specialist Conference*, August 2004, 10.2514/6.2004-5089.

[37] C. A. Kluever, "Simple Guidance Scheme for Low-Thrust Orbit Transfers," *Journal of Guidance, Control, and Dynamics*, Vol. 21, No. 6, 1998, pp. 1015–1017, 10.2514/2.4344.

[38] A. Ruggiero, P. Pergola, S. Marcuccio, and M. Andrenucci, "Low-Thrust Maneuvers for the Efficient Correction of Orbital Elements," *32nd International Electric Propulsion Conference*, Paper IEPC-2011-102, September 2011, pp. 11–15.

[39] J. S. Parker, *Low-Energy Ballistic Lunar Transfers*. PhD thesis, University of Colorado at Boulder, 2007.

[40] T. A. Pavlak, *Trajectory Design and Orbit Maintenance Strategies in Multi-body Dynamical Regimes*. PhD thesis, Purdue University, 2013.

[41] E. M. Zimovan, K. C. Howell, and D. C. Davis, "Near Rectilinear Halo Orbits and Their Application in Cis-Lunar Space," *IAA Conference on Dynamics and Control of Space Systems*, 2017, pp. Paper Number IAA–AAS–DyCoSS3–125.

[42] J. E. Prussing and B. A. Conway, "The Three-Body Problem," *Orbital Mechanics*, ch. 4, Oxford University Press, 2013.

[43] C. Martin, *A New Numerical Optimization Method Based on Taylor Series for Challenging Trajectory Optimization Problems*. PhD thesis, University of Illinois at Urbana-Champaign, 2011.

[44] D. L. Richardson, "Analytic Construction of Periodic Orbits about the Collinear Points," *Celestial Mechanics and Dynamical Astronomy*, Vol. 22, No. 3, 1980, pp. 241–253.

[45] J. V. Breakwell and J. V. Brown, "The 'Halo' Family of 3-Dimensional Periodic Orbits in the Earth-Moon Restricted 3-Body Problem," *Celestial Mechanics*, Vol. 20, No. 4, 1979, pp. 389–404.

[46] A. Petropoulos, "Refinements to the Q-law for Low-Thrust Orbit Transfers," *AAS/AIAA Space Flight Mechanics Conference*, AAS Paper 05-162, January 2005.

[47] A. Petropoulos and S. Lee, "Optimisation of Low-Thrust Orbit Transfers Using the Q-Law for the Initial Guess," *AAS/AIAA Astrodynamics Specialist Conference*, AAS Paper 05–392, August 2005.

[48] G. Whiffen, "Mystic: Implementation of the Static Dynamic Optimal Control Algorithm for High-Fidelity, Low-Thrust Trajectory Design," *AIAA/AAS Astrodynamics Specialists Conference*, August 2006.

[49] S. Lee, P. v. Ailmen, W. Fink, A. Petropoulos, and R. J. Terrile, "Design and Optimization of Low-Thrust Orbit Transfers," *IEEE Aerospace Conference*, March 2005, 10.1109/AERO.2005.1559377.

[50] A. E. Petropoulos, Z. B. Tarzi, G. Lantoine, T. Dargent, and R. Epenoy, "Techniques for Designing Many-Revolution Electric-Propulsion Trajectories," *AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 14-373, January 2014.

[51] J. Vesterstrom and R. Thomsen, "A Comparative Study of Differential Evolution, Particle Swarm Optimization, and Evolutionary Algorithms on Numerical Benchmark Problems," *Congress on Evolutionary Computation*, Vol. 2, IEEE, 2004, pp. 1980–1987, 10.1109/CEC.2004.1331139.

[52] J. Kennedy and R. Eberhart, "Particle Swarm Optimization," *IEEE International Conference on Neural Networks*, Vol. 4, 1995, pp. 1941–1948, 10.1109/ICNN.1995.488968.

[53] K. E. Parsopoulos, M. N. Vrahatis, *et al.*, "Particle Swarm Optimization Method for Constrained Optimization Problems," *Intelligent Technologies–Theory and Application* (V. Kvasnicka, J. Pospíchal, and P. Sincák, eds.), pp. 214–220, IOS Press, 2002.

[54] R. l. Perez and K. Behdinan, "Particle Swarm Approach for Structural Design Optimization," *Computers & Structures*, Vol. 85, No. 19-20, 2007, pp. 1579–1588, 10.1016/j.compstruc.2006.10.013.

[55] M. Pontani and B. A. Conway, "Particle Swarm Optimization Applied to Space Trajectories," *Journal of Guidance, Control, and Dynamics*, Vol. 33, No. 5, 2010, pp. 1429–1441, 10.2514/1.48475.

[56] MATLAB R2016a, "Particle Swarm Optimization Algorithm," https://www.mathworks.com/help/gads/particle-swarm-optimization-algorithm.html. [retrieved 15 May 2018].

[57] J. R. Dormand and P. J. Prince, "Higher Order Embedded Runge—Kutta Formulae," *Journal of Computational and Applied Mathematics*, Vol. 7, No. 1, 1981, pp. 67–75, 10.1016/0771-050X(81)90010-3.

[58] NASA, "Moon Fact Sheet," https://nssdc.gsfc.nasa.gov/planetary/factsheet/moonfact.html. [retrieved 30 Jan 2018].

[59] C. A. Kluever, "Designing Transfers to Geostationary Orbit Using Combined Chemical–Electric Propulsion," *Journal of Spacecraft and Rockets*, Vol. 52, No. 4, 2015, pp. 1144–1151, 10.2514/1.A33259.

[60] N. Arora and N. Strange, "A Low Energy, Low Thrust Unified Solver for Rapid Mission Design," *AAS/AIAA Space Flight Mechanics Meeting*, AAS Paper 13-255, February 2013.

[61] B. B. Jagannatha and K. Ho, "Optimization of In-Space Supply Chain Design Using High-Thrust and Low-Thrust Propulsion Technologies," *Journal of Spacecraft and Rockets*, Vol. 55, No. 3, 2018, pp. 648–659, 10.2514/1.A34042.

[62] A. J. Abraham, D. B. Spencer, and T. J. Hart, "Early Mission Design of Transfers to Halo Orbits via Particle Swarm Optimization," *The Journal of the Astronautical Sciences*, Vol. 63, No. 2, 2016, pp. 81–102, 10.1007/s40295-016-0084-2.

[63] J. S. Parker and R. L. Anderson, "Low-Energy Lunar Trajectory Design," *DESCANSO Series*, JPL, 2014.

[64] C. Zhang, F. Topputo, F. Bernelli-Zazzera, and Y.-S. Zhao, "Low-Thrust Minimum-Fuel Optimization in the Circular Restricted Three-Body Problem," *Journal of Guidance, Control, and Dynamics*, Vol. 38, No. 8, 2015, pp. 1501–1510, 10.2514/1.G001080.

[65] V. Shah, R. Beeson, and V. Coverstone, "Automated Global Optimization of Multi-Phase Trajectories in the Three-Body Problem using a Variable Chromosome Transcription," *AAS/AIAA Space Flight Mechanics Conference*, AAS Paper 17-445, February 2017.

[66] B. B. Jagannatha, V. Shah, R. Beeson, and K. Ho, "Exploration of Low-Thrust Trajectories to Earth-Moon Halo Orbits," *AIAA/AAS Astrodynamics Specialist Conference*, Paper 17-757, August 2017.

[67] B. B. Jagannatha, J.-B. H. Bouvier, and K. Ho, "Preliminary Design of Low-Energy Low-Thrust Transfers to Halo Orbits using Feedback Control," *Journal of Guidance, Navigation, and Control*, Vol. Accepted, 2018.

[68] A. Commission, "A Journey to Inspire, Innovate, and Discover," https://history.nasa.gov/aldridge_commission_report_june2004.pdf, June 2004. Accessed: May 4, 2017.

[69] N. P. Release, "NASA Administrator Statement on Space Policy Directive-3," https://www.nasa.gov/press-release/nasa-administrator-statement-on-space-policy-directive-3, June, 2018. Accessed: Sept 4, 2018.

[70] G. Williams and J. Crusan, "Pioneering Space – the Evolvable Mars Campaign," https://www.nasa.gov/sites/default/files/files/20150408-NAC-Crusan-EMC-v7a_TAGGED.pdf. Accessed: May 4, 2017.

[71] K. E. Goodliff, C. Stromgren, Z. Dickert, M. K. Ewert, J. Hill, and C. Moore, "Logistics Needs for Future Human Exploration Beyond Low Earth Orbit," *AIAA SPACE and Astronautics Forum and Exposition*, Paper AIAA 2017-5122, 2017.

[72] F. Chandler, D. Bienhoff, J. Cronick, and G. Grayson, "Propellant Depots for Earth Orbit and Lunar Exploration," *AIAA SPACE 2007 Conference & Exposition*, 2007.

[73] E. Hurlbert, M. Baine, and G. Grush, "An Open Exploration Architecture Using an L-1 Space Propellant Depot," *AIAA SpaceOps Conference*, Paper 2010-2158, 2010.

[74] Advanced Design Team, "Lunar L1 Gateway Conceptual Design Report V.1.0," tech. rep., NASA Report JSC EX15-01-001, Oct 2001.

[75] P. Troutman, D. D. Mazanek, F. Stillwagen, J. Antol, T. Sarver-Verhey, D. Chato, R. Saucillo, D. Blue, W. Siegfried, D. Carey, *et al.*, "Orbital Aggregation and Space Infrastructure Systems (OASIS)," *International Astronautical Congress*, Paper IAC-02-IAA.13.2.06, Houston, TX, 2002.

[76] R. W. Orloff, "Apollo by the Numbers: A Statistical Reference," tech. rep., NASA/SP-2000-4029, NAS 1.21:4029, LC-00-061677, 1996.

[77] M. Schaffer, "Lunar Surface Access from Earth-Moon L1/L2: A novel lander design and study of alternative solutions," http://www.sei.aero/archive/Schaffer-FISO-SEV-L1L2-Lander-11-28-2012-revA.pdf. Accessed: May 4, 2017.

[78] F. Zegler, B. Kutter, and J. Barr, "A Commercially Based Lunar Architecture," *AIAA SPACE Conference & Exposition*, Paper 2009-6567, 2009.

[79] B. F. Kutter and G. F. Sowers, "Cislunar-1000: Transportation Supporting a Self-Sustaining Space Economy," *AIAA SPACE Conference & Exposition*, Paper 2016-5491, 2016.

[80] J. Brophy and S. Oleson, "Spacecraft Conceptual Design for Returning Entire Near-Earth Asteroids," *AIAA/ASME/SAE/ASEE Joint Propulsion Conference & Exhibit*, 2012-4067, 2012.

[81] T. Percy, M. McGuire, and T. Polsgrove, "Combining Solar Electric Propulsion and Chemical Propulsion for Crewed Missions to Mars," *IEEE Aerospace Conference*, 2015, 10.1109/AERO.2015.7119289.

[82] D. H. Manzella and K. Hack, "High-Power Solar Electric Propulsion for Future NASA Missions," *AIAA/ASME/SAE/ASEE Joint Propulsion Conference*, Paper 2014-3718, 2014, 10.2514/6.2014-3718.

[83] ESA, "Orion," https://www.esa.int/Our_Activities/Human_Spaceflight/Orion/Propulsion. Accessed: Sept 12, 2018.

[84] R. Martinez, G. Condon, and J. Williams, "Time and Energy, Exploring Trajectory Options Between Nodes in Earth-Moon Space," *Global Space Exploration Conference*, GLEX-2102.D1.2.8x12585, 2012.

[85] W. D. Pratt, C. Buxton, S. Hall, J. B. Hopkins, and A. Scott, "Trajectory Design Considerations for Human Missions to Explore the Lunar Farside from the Earth-Moon Lagrange point EM-L2," *AIAA SPACE Conference & Exposition*, Paper 2013-5478, 2013, 10.2514/6.2013-5478.

[86] D. C. Folta, T. A. Pavlak, A. F. Haapala, and K. C. Howell, "Preliminary Design Considerations for Access and Operations in Earth-Moon L1/L2 Orbits," *AAS/AIAA Space Flight Mechanics Meeting*, Paper 13-339, 2013.

[87] J. O. Burns, D. A. Kring, J. B. Hopkins, S. Norris, T. J. W. Lazio, and J. Kasper, "A Lunar L2-Farside Exploration and Science Mission Concept with the Orion Multi-Purpose Crew Vehicle and a Teleoperated Lander/Rover," *Advances in Space Research*, Vol. 52, No. 2, 2013, https://doi.org/10.1016/j.asr.2012.11.016.

[88] R. W. Farquhar, D. W. Dunham, Y. Guo, and J. V. McAdams, "Utilization of Libration Points for Human Exploration in the Sun-Earth-Moon System and Beyond," *Acta Astronautica*, Vol. 55, No. 3-9, 2004, pp. 687–700, 10.1016/j.actaastro.2004.05.021.

[89] L. Mingtao and Z. Jianhua, "Impulsive Lunar Halo Transfers using the Stable Manifolds and Lunar Flybys," *Acta Astronautica*, Vol. 66, No. 9-10, 2010, pp. 1481–1492, 10.1016/j.actaastro.2009.11.014.

[90] B. B. Jagannatha and K. Ho, "Optimization of In-Space Supply Chain Design Using High-Thrust and Low-Thrust Propulsion Technologies," *Journal of Spacecraft and Rockets*, Vol. 55, No. 3, 2018, pp. 648–659, 10.2514/1.A34042.

[91] J. E. Prussing and B. A. Conway, "Continuous-Thrust Orbit Transfer," *Orbital Mechanics*, ch. 8, Oxford University Press, 2013.

[92] B. B. Jagannatha and K. Ho, "Event-Driven Space Logistics Network Optimization for Cislunar Supply Chain Design with High-Thrust and Low-Thrust Propulsion Technologies," *AAS/AIAA Astrodynamics Specialist Conference*, Paper 18-337, August 2018.

[93] A. Gad and O. Abdelkhalik, "Hidden Genes Genetic Algorithm for Multi-Gravity-Assist Trajectories Optimization," *Journal of Spacecraft and Rockets*, Vol. 48, No. 4, 2011, pp. 629–641.

[94] MATLAB R2016a, "gamultiobj Algorithm," `https://www.mathworks.com/help/gads/gamultiobj-algorithm.html`. [retrieved 15 May 2018].

[95] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A Fast and Elitist Multiobjective Genetic Algorithm: NSGA-II," *IEEE Transactions on Evolutionary Computation*, Vol. 6, No. 2, 2002, pp. 182–197.

[96] E. L. M. Beale and J. A. Tomlin, "Special Facilities in a General Mathematical Programming System for Nonconvex Problems using Ordered Sets of Variables," *Proceedings of the 5th International Conference on Operations Research*, Tavistock Publications, London, 1970, p. 447–454.

[97] A. B. Keha, I. R. d. Farias Jr, and G. L. Nemhauser, "A Branch-and-Cut Algorithm Without Binary Variables for Nonconvex Piecewise Linear Optimization," *Operations Research*, Vol. 54, No. 5, 2006, pp. 847–858.

[98] E. Camponogara and L. F. Nazari, "Models and Algorithms for Optimal Piecewise-Linear Function Approximation," *Mathematical Problems in Engineering*, 2015.

[99] D. H. Ellison, *Robust Preliminary Design for Multiple Gravity Assist Spacecraft Trajectories*. PhD thesis, University of Illinois at Urbana-Champaign, 2018.

[100] J. P. Vielma, S. Ahmed, and G. Nemhauser, "Mixed-Integer Models for Nonseparable Piecewise-Linear Optimization: Unifying Framework and Extensions," *Operations research*, Vol. 58, No. 2, 2010, pp. 303–315.

[101] J. P. Vielma and G. L. Nemhauser, "Modeling Disjunctive Constraints with a Logarithmic Number of Binary Variables and Constraints," *Mathematical Programming*, Vol. 128, No. 1-2, 2011, pp. 49–72.

[102] G. Lantoine and R. P. Russell, "A Hybrid Differential Dynamic Programming Algorithm for Constrained Optimal Control Problems. Part 1: theory," *Journal of Optimization Theory and Applications*, Vol. 154, No. 2, 2012, pp. 382–417.

[103] G. Lantoine and R. P. Russell, "A Hybrid Differential Dynamic Programming Algorithm for Constrained Optimal Control Problems. Part 2: Application," *Journal of Optimization Theory and Applications*, Vol. 154, No. 2, 2012, pp. 418–442.

[104] J. D. Aziz, J. S. Parker, D. J. Scheeres, and J. A. Englander, "Low-Thrust Many-Revolution Trajectory Optimization via Differential Dynamic Programming and a Sundman Transformation," *The Journal of the Astronautical Sciences*, Vol. 65, No. 2, 2018, pp. 205–228.

[105] A. L. Herman and B. A. Conway, "Direct Optimization using Collocation Based on High-Order Gauss-Lobatto Quadrature rules," *Journal of Guidance, Control, and Dynamics*, Vol. 19, No. 3, 1996, pp. 592–599.