LOW-THRUST TRANSFERS FROM DISTANT RETROGRADE ORBITS TO L₂ HALO ORBITS IN THE EARTH-MOON SYSTEM

Nathan L. Parrish[1], Jeffrey S. Parker[1], Steven P. Hughes[2], Jeannette Heiligers[1, 3]

[1]Colorado Center for Astrodynamics Research, University of Colorado at Boulder;
[2]NASA Goddard Space Flight Center
[3]Delft University of Technology

ABSTRACT

This paper presents a study of transfers between distant retrograde orbits (DROs) and L₂ halo orbits in the Earth-Moon system that could be flown by a spacecraft with solar electric propulsion (SEP). Two collocation-based optimal control methods are used to optimize these highly-nonlinear transfers: Legendre pseudospectral and Hermite-Simpson. Transfers between DROs and halo orbits using low-thrust propulsion have not been studied previously. This paper offers a study of several families of trajectories, parameterized by the number of orbital revolutions in a synodic frame. Even with a poor initial guess, a method is described to reliably generate families of solutions. The circular restricted 3-body problem (CRTBP) is used throughout the paper so that the results are autonomous and simpler to understand.

Index Terms— Electric propulsion, collocation, CRTBP

1. INTRODUCTION

The goal of this paper is to fill a gap in the types of transfers studied, as well as to begin understanding some of the families of transfers which exist for any low-thrust transfer in an N-body force field. Similar types of transfers that have been studied in the literature include: from Earth orbit to Moon orbit using low-thrust [1, 2], from Earth orbit to libration point orbits using low-thrust [3], from Earth to DRO using impulsive maneuvers [4], from Earth to DRO using low-thrust [5], from L₁ halo orbit to L₂ halo orbit, and solar sail transfers between libration point orbits of different Sun-planet systems [6, 7].

For the most part, results in the literature focus on a single example trajectory studied in great detail. However, there are few papers that study families of transfers. Topputo [8] showed that many distinct families of ballistic transfers exist between the Earth and Moon in a four-body model, and others have demonstrated that such variations exist for other types of transfers in Earth-Moon space [9, 10]. By exploring the families of transfers that exist between DROs and L₂ halo orbits, this paper provides deeper insights into the trade space available.

2. BACKGROUND

2.1. Circular restricted three-body problem

The full three-body problem has eluded analytical representation for centuries. Each body has 6 degrees of freedom, for a total of 18. There are only 10 known integrals of motion, so it is impossible to develop an analytical representation. Some common simplifications can be made to make the problem tractable.

The circular restricted three body problem (CRTBP) makes two significant assumptions: the mass of the third body (the spacecraft) is negligible compared to the primary or secondary bodies, and the primary and secondary orbit the system barycenter in perfectly circular orbits [11].

Non-dimensional distance and time units are used such that 1 DU is the distance from Earth to Moon, and 2π TU is the orbital period of Earth and Moon about their barycenter. The non-dimensional mass ratio μ (defined as the mass of the secondary divided by the system's total mass) is used instead of the gravitational parameter of a two-body system. For the Earth-Moon system, μ is approximately 0.012151. A synodic reference frame is used, defined such that the x-axis is positive towards the secondary body. Earth is on the x-axis at $(-\mu)$, and the Moon is on the x-axis at $(1 - \mu)$. The z-axis is defined by the rotation axis of the system, and the y-axis completes the right-handed triad. This reference frame is shown in Figure 1. The differential equations with thrust in the CRTBP in the synodic reference frame are

$$\ddot{x} = -\left(\frac{(1-\mu)}{r_1^3}(x+\mu) + \frac{\mu}{r_2^3}(x-1+\mu)\right) + 2\dot{y} + x + T_x$$

$$\ddot{y} = -\left(\frac{(1-\mu)}{r_1^3}y + \frac{\mu}{r_2^3}y\right) - 2\dot{x} + y + T_y$$

$$\ddot{z} = -\left(\frac{(1-\mu)}{r_1^3}z + \frac{\mu}{r_2^3}z\right) + T_z$$

where μ is the mass ratio of the system, r_1 is the distance from the primary, r_2 is the distance from the secondary, and $T_{[]}$ is the acceleration due to thrust.

Rate of mass change is given by $\dot{m} = -|\vec{T}|/(I_{sp}g_0)$, where I_{sp} is the specific impulse and g_0 is the standard sea-level gravitational acceleration due to Earth.

Figure 1. Illustration of the Earth-Moon system in the CRTBP, viewed in the synodic reference frame. The Earth and the Moon are plotted, not to scale. Libration points L_1 through L_5 are also shown, as are some zero-velocity curves of equal Jacobi constant.

2.2. Halo orbits

Halo orbits are so-named because when viewed in a synodic reference frame, they trace a "halo" in space. These orbit the libration points such as L_1 or L_2 [12]. The Orion/Moonrise mission concept would use Earth-Moon L_2 as a low-fuel-cost rendezvous location for the manned Orion capsule and the proposed Moonrise vehicle carrying lunar samples [13].

2.3. Distant retrograde orbits

DROs are a type of orbit that have received increased attention in the past few years because of the unique characteristics they exhibit. DROs are a type of repeating orbit that exists only in the 3-body problem [11]. When viewed in a synodic reference frame, a DRO is retrograde about the secondary body, at a relatively high altitude such that the orbit is significantly perturbed by both the primary and secondary bodies. DROs are unique in that they sit between two-body orbits and libration point orbits in terms of stability. These orbits are often dynamically stable, though it has been shown that perturbations in a high-fidelity model of the solar system may cause a spacecraft to depart an otherwise stable DRO [14]. Parker, Bezrouk, and Davis demonstrated several trajectories that transfer from Earth to a DRO, requiring no maneuvers to capture at the DRO and remaining on the DRO for thousands of years [4].

Mission concepts that have examined DROs include the proposed NASA/JPL Asteroid Redirect Mission [15] and the Orion/MoonRise concept [13], [16]. Both of these mission concepts could benefit from the capability to transit between an asteroid captured in the DRO and a potential space vehicle in the halo orbit. Ongoing research by Davis and Parker is finding that impulsive transfers between those orbits do exist, but they are costly on the order of 150 m/s and require transfer times on the order of weeks to months. The present work finds that spacecraft with SEP have the potential to greatly reduce the propellant mass required to make such transfers, without much increase in time of flight.

2.4. Collocation

The basic principle of collocation is to represent an ordinary differential equation with some continuous function which obeys the differential equations of motion at a set of nodes. Collocation is a direct method that transcribes an optimal control problem to a non-linear programming (NLP) problem which can be solved by any industry-standard NLP software [17]. The IPOPT NLP solver is used here [18]. A variety of collocation-based methods exist, distinguished by the node spacing and the choice of basis functions. Pseudospectral collocation is generally defined on one of three choices of meshes: Legendre-Gauss (LG), which does not have a control node on either endpoint; Legendre-Gauss-Radau (LGR), which has a control node on just one endpoint but not the other, and Legendre-Gauss-Lobatto (LGL), which has control nodes on both endpoints [19].

A helpful way to think of collocation is through a comparison to implicit numerical integration schemes. When propagating a system with known forces, information about the current state and, possibly, the state at previous integration steps is used to calculate the state at some time in the future. In collocation, rather than propagating a known initial state through known forces, the states and controls are optimization parameters subject to constraints. In order to find a solution which obeys the differential equations of motion, a defect is calculated at or between each node. Reference [20] has an excellent description of collocation.

Two collocation methods are used for this research: Legendre pseudospectral, and Hermite-Simpson. These are used as implemented in the open source, optimal control package PSOPT (PseudoSpectral OPTimal control) [21]. In both cases, the NLP solver attempts to minimize the differential defect constraints while minimizing the cost function and minimizing any other constraint defects.

2.4.1. Legendre pseudospectral approximation
The Legendre pseudospectral method approximates each element of the state and control as an N^{th} order Lagrange polynomial at the N quadrature nodes. Time is transformed to

be in the interval $[-1, +1]$. The state x at node τ is approximated by [19, 20, 21]:

$$x(\tau) \approx \sum_{k=0}^{N} x(\tau_k) \mathcal{L}_k(\tau)$$

where $\mathcal{L}_k$ are the Lagrange basis polynomials, and τ is the transformed time. The Lagrange basis polynomials $\mathcal{L}_k(\tau)$ can be expressed as follows:

$$\mathcal{L}_k(\tau) = \frac{1}{N(N+1)L_N(\tau_k)} \frac{(\tau^2 - 1)\dot{L}_N(\tau)}{\tau - \tau_k}$$

where L_N are the Legendre polynomials of order N of the form

$$L_N(\tau) = \frac{1}{2^N N!} \frac{d^N}{d\tau^N} (\tau^2 - 1)^N$$

The derivative of the state vector is analytically approximated as

$$\dot{x}(\tau_k) \approx \sum_{i=0}^{N} D_{ki} x^N(\tau_i)$$

where D is the differentiation matrix with size $(N+1) \times (N+1)$. The elements of D are given by

$$D_{ki} = \begin{cases} -\dfrac{L_N(\tau_k)}{L_N(\tau_i)} \dfrac{1}{\tau_k - \tau_i}, & k \neq i \\ \dfrac{N(N+1)}{4}, & k = i = 0 \\ -\dfrac{N(N+1)}{4}, & k = i = N \\ 0, & else \end{cases}$$

Differential defect constraints are calculated by taking the difference of the analytical derivative of the approximate state vector with the actual differential equations describing the dynamics.

A limitation of PSOPT is that the number of phases must be defined a priori. For the Legendre pseudospectral approximation, we will use the term "phase" to mean a duration of time in the mission that is defined using the collocation nodes. Liu, Hager, and Rao have developed [23] a method for automatically refining the number of phases in addition to refining the number of nodes in each phase, using LGR nodes. Automatically refining the number of phases would permit greater accuracy near times of quickly-changing dynamics, such as a lunar flyby, or discontinuous dynamics, such as thrust turning on or off. As it is, PSOPT suffers decreased accuracy in these situations because it is impractical to increase the number of nodes in the single phase to be high enough to have more than a few nodes near flybys.

2.4.2. Hermite-Simpson approximation
The Hermite-Simpson method defines a vector of differential defect constraints ζ at node τ_k as follows [20, 22]:

$$\zeta(\tau_k) = x(\tau_{k+1}) - x(\tau_k) - \frac{h_k}{6}\left(f_k + 4\bar{f}_{k+1} + f_{k+1}\right)$$

where

$$\bar{f}_{k+1} = f\left[\bar{x}_{k+1}, \bar{u}_{k+1}, p, \tau_k + \frac{h_k}{2}\right]$$

$$\bar{x}_{k+1} = \frac{1}{2}\left(x(\tau_k) + x(\tau_{k+1})\right) + \frac{h_k}{8}\left(f_k - f_{k+1}\right)$$

When using the Hermite-Simpson method, PSOPT is able to perform automatic mesh refinement – placing more nodes near times of quickly-changing dynamics.

3. METHODS AND APPROACH

Throughout the analysis presented here, the initial spacecraft mass is 1500 kg, and the I_{sp} is 3,000 seconds. The DRO used is completely in the Earth-Moon plane. It crosses the *y*-axis with positive *y*-velocity at $x = 0.9$ DU. The orbital period is approximately 5.55 days. The L_2 halo orbit used has a maximum *z*-amplitude of approximately 53,000 km and an orbital period of approximately 14.0 days. These initial and final orbits are shown in several of the figures below, such as Figure 4.

3.1. Optimal control problem implementation in PSOPT

The endpoints of the trajectory were constrained to lie on the DRO (for the initial point) and on the halo orbit (for the final point), but the optimizer could choose where to depart the DRO and where to arrive on the halo orbit: a set of 100 states on the DRO and on the halo orbit are hard-coded into the software, and then linear interpolation is used to find the state at an arbitrary location as requested by the optimizer through a static optimization parameter. The initial and final time are also free parameters. However, upper and lower bounds for time were implemented somewhat above and below the expected time of flight in order to scale the problem better.

The spacecraft state at each node is a 7-element vector: 3 components for position, 3 for velocity, and 1 for mass. The controls were represented as a 3-element vector at each node, for the *x*-, *y*-, and *z*-components of the thrust. Three path constraints were used: the upper limit of thrust magnitude, the lower limit for distance to the Moon, and the lower limit for distance to the Earth.

Three objective functions were used separately: zero cost, minimum time of flight, and maximum final mass. The zero cost objective function was used to quickly find feasible, non-optimal transfers. These transfers were used only as intermediaries in order to accelerate the optimization of one of the other two objective functions. Resulting families of trajectories from each are presented later.

All states and controls were scaled to be approximately of $O(1)$. Non-dimensional units for distance and time were used as defined in section 2.1. Mass was scaled by a factor of 1,000 so that the initial mass was 1.5 mass units rather than 1,500 kg. Thrust was used in the physical units of Newtons.

In general, we know that the optimal (minimum fuel) transfer will have the form "bang-off-bang" – maximum thrust, coast, maximum thrust. Since collocation methods represent the states and controls as continuous polynomials, it is impossible to find a perfectly sharp thrust cutoff using a single phase. This limitation can be avoided by using a multiphase problem formulation, where the trajectory is continuous only over a phase. Phase endpoints can then be constrained to match in position and velocity, but not thrust. PSOPT has the capability of solving multi-phase problems; however, for this work, a single phase was used. It is expected that using three phases (for thrust-coast-thrust) or more will allow PSOPT to find slightly more accurate transfers, as the thrust cut-off can then be perfectly sharp. Examples of thrust profiles found are presented in Figure 2 and Figure 3.

For all results shown here, the NLP tolerance used was 10^{-4}, and the maximum number of iterations allowed was 3,000. A smaller tolerance will lead to solutions that match the actual dynamics more accurately, but more iterations within the NLP solver will be necessary. Several transfers were solved at a tolerance of 10^{-5}, but it was deemed unnecessary to use such a small tolerance when performing a broad search.


Figure 2. Example of a thrust profile for a single-revolution transfer. There are two clear coast arcs visible, with fairly sharp thrust cutoffs. Use of multiple phases could achieve a very small accuracy improvement.


Figure 3. Example of a thrust profile for a two-revolution transfer. Here, there are three coast arcs, but the thrust

cutoffs are not as clean. Using multiple phases would clean up the thrust profile substantially and improve accuracy.

3.2. Initial guess generation

At the start of this research, attempts were made to find a "close" initial guess by choosing various control laws to propagate forward from the DRO and backward from the halo orbit, searching for intersection points. However, no close initial guess could be found in this way. Instead, a method was developed that allowed the problem to converge even when given a poor initial guess.

Initial guesses were formed by a very simple means:
1) Propagate an initial state on the DRO forward in time.
2) Jump to an arbitrary point on the halo orbit and propagate an initial state on the halo orbit forward in time.
3) Concatenate the states in the DRO with the states in the halo orbit.

Now, we have a list of states and times that obey the force model at all points except the middle, where there is an instantaneous jump from the DRO to the halo orbit. The control was initialized to zero.
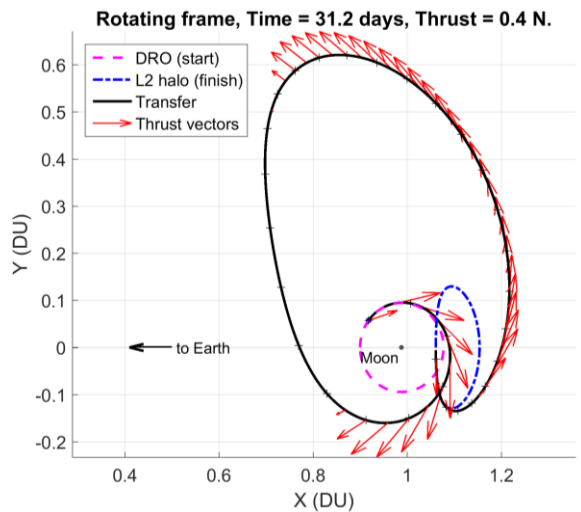

Figure 4. A 1-revolution transfer, using 0.4 N max thrust (for an initial mass of 1500 kg), viewed in the synodic reference frame. The spacecraft begins in the DRO and ends in the L_2 halo orbit. Tick marks appear at ½ day intervals, and thrust vectors appear at ¼ day intervals.

It was found that the optimized trajectories produced by PSOPT generally consisted of the same number of "revolutions" about the Moon as the initial guess. This was found to be true independently for the DRO and for the halo orbit. The converged trajectory will generally have the same number of revolutions about the Moon as the initial guess,

and the same number of revolutions about L_2 as the initial guess. Although collocation methods have been found to have a sufficiently wide basin of attraction to solve N-body transfers such as these, the solutions found with such methods cannot claim to be globally optimal.

Therefore, families of transfers could be selected to some extent by adjusting the number of revolutions about the DRO and the halo orbit in the initial guess. Three representative examples of transfers from 1-revolution, 2-revolution, and 4-revolution families are shown in Figure 4, Figure 5, and Figure 6, respectively.

Rotating frame, Time = 27.4 days, Thrust = 0.4 N.

Figure 5. A 2-revolution transfer, using 0.4 N thrust.

**Rotating frame, Time = 32.6 days, Thrust = 0.4 N.**

Figure 6. A 4-revolution transfer, using 0.4 N thrust.

3.3. Method for finding families

It was found that for the low-thrust DRO to L_2 halo orbit transfer problem, the pseudospectral method was more likely to converge than the Hermite-Simpson method when given a poor initial guess. However, there is a danger that a close approach to the Moon will not be represented well by the pseudospectral method due to the fixed node spacing. For instance, if a transfer involved multiple revolutions about the Moon and there are too few nodes, the Legendre pseudospectral approximation can break down and have extremely poor accuracy. Although the optimizer may converge, the solution does not have physical meaning. The constraints on the Legendre pseudospectral approximation are met, but the nodes are spaced too far apart for the approximation to be accurate. An analogous limitation that many astrodynamicists are familiar with is the necessity of keeping time steps appropriately small when numerically propagating an orbit with a Runge-Kutta integration scheme. An example of having too few nodes is shown in Figure 7.

Having too few nodes can be resolved by mesh refinement techniques. Pseudospectral methods use global polynomials by definition, and since PSOPT uses only the phases defined a priori, the pseudospectral collocation method used here requires adding more nodes throughout the entire transfer in order to add nodes near a close approach to the Moon. The Hermite-Simpson method permits local mesh refinement, so more nodes could be added near the close approach only, without the need to add nodes elsewhere in the trajectory.

Figure 7. An incomplete trajectory using too few nodes. This is from a converged, minimum propellant solution. With only four nodes representing the inner revolution about the Moon, the solution has a poor accuracy and may change significantly when more nodes are added.

The Hermite-Simpson method provides a workaround to this challenge via automatic mesh refinement – placing more nodes near times when the dynamics change quickly. In PSOPT, the pseudospectral method becomes very slow as the

number of nodes grows. The greatest number of nodes used with the pseudospectral method was 160, which required a few hours to converge. The same transfer with the Hermite-Simpson method required less than one hour to converge. With these considerations in mind, the following algorithm was developed to reliably generate families of transfers:

1) Generate an initial guess with the appropriate number of revolutions about the Moon, as described in section 3.2.
2) Using the pseudospectral method, run the problem with zero cost function. This allows the optimizer to quickly find a feasible (but not optimal) transfer.
3) Using the Hermite-Simpson method, set the objective function to maximize the final mass, and run the optimizer.
4) Decrease the maximum thrust limit slightly. Using the Hermite-Simpson method again and the solution from step (3) as the initial guess, run the optimizer.
5) Repeat step (4) until the problem no longer converges.

By following the above algorithm, the families of transfers described in the figures below were found.

In some cases, the optimizer would not converge. Even when the problem has not been fully solved, PSOPT will print solution files. Re-running the optimizer, with the failed solution as the new initial guess would sometimes result in a successful solution. This is due to PSOPT automatically recalculating the Jacobian and re-weighting the problem.

One difficulty that was encountered many times was that the optimizer would (correctly) find that thrusting during a lunar flyby would improve the cost function. However, unless a large number of nodes were used initially, the lunar flyby would be represented by only one or two nodes. The accuracy of the entire trajectory would then deteriorate to a point such that it was impossible to interpolate the existing solution accurately enough to add more nodes.


Figure 8. Multiple close approaches to the Moon represented by too few nodes. 80 nodes are used here with the pseudospectral method. The close approaches result in a lower propellant mass (only 15 kg, as opposed to 23-28 kg for most of the transfers represented in this paper). We can see that the closest approach is represented by only a single node, which is not enough.


Figure 9. An example of a successfully-converged solution which uses a powered, close lunar flyby. The powered flyby results in lower propellant costs, but comes at the expense of a more sensitive trajectory.

In order to make the trajectory in Figure 8 more meaningful, either a larger number of nodes should be used in the pseudospectral method, or the Hermite-Simpson method should be used with automatic mesh refinement. In this case, the trajectory did not converge when the solution shown in Figure 8 was used as an initial guess for the Hermite-Simpson method.

In a few rare circumstances (such as shown in Figure 9), the optimizer succeeded in adding subsequent nodes such that the trajectory became physically meaningful. However, this was not reliable.

Additionally, a powered lunar flyby would be dangerous from an operations standpoint. Due to the chaotic nature of the dynamics, slight errors in the state estimate or in maneuver execution could have drastic effects on the orbit after the flyby. Looking at Figure 9, we can see that the duration of a close approach is roughly 0.5-1 days, depending on the definition of "close approach". Ground control during the close approach would be difficult at the least. In order to avoid these complications, a "keep-out" zone was enforced so that the spacecraft could never get closer than 0.04 DU, or about 9 lunar radii.

4. ANALYSIS OF RESULTS

Using the methods described above, four families of transfers were examined: 1-, 2-, and 3-revolutions, minimizing propellant; and 1-revolution, minimizing time.

Within a family of transfers, lower thrust generally requires higher time of flight. Within each family, there are also branches which connect the families. For instance, in Figure 10, the 1-revolution, minimum time case has three points which appear to lie on the 1-revolution, min propellant curve.

The propellant mass required for these transfers generally lies between 23-28 kg. There were no observed trends between propellant mass and any other parameters used to describe the families of transfers. Figure 11 shows the propellant mass as a function of the thrust limit.

Within a family of transfers, decreasing the thrust requires departing further from the Moon, as seen in Figure 12. This could have implications for navigation and/or science objectives. At some point, each family hit a limit for thrust, below which they would not converge. In general, the trajectory would take on an extra loop far from the Moon and violate the differential constraints from the dynamics. An example of this is shown in Figure 13. Although a similar feature may exist in some real solutions, whenever it appeared in this work, the optimizer could not converge.


Figure 10. Time of flight as a function of the thrust limit. The 2-revolution and 3-revolution cases were only examined at thrust levels of 0.4 N and below.


Figure 11. Propellant mass as a function of thrust. There is no clear relationship, and the propellant mass for almost every case lies between 23-28 kg. It is clear in the 1-revolution, min time case, especially, that there are multiple branches within each family, some of which are more favorable than others.

**Figure 12. Maximum distance from the Moon asd a function of the thrust limit.**

Figure 13. A non-physical solution illustrating a common failure mode for the continuation method described. The loop which appears near x=0.8, y=0.45 came in for many cases in which the optimizer could not converge.

5. CONCLUSION & FUTURE WORK

Trajectories in N-body force fields are among the most difficult in astrodynamics to understand and optimize. Whenever possible, the problem has been effectively simplified by using the CRTBP. Even with those simplifications, however, there are only a handful of special types of orbits which are well-defined, such as the DRO and the L_2 halo orbit used in this paper. Predicting a transfer trajectory between any two N-body orbits remains a great challenge.

This work provides a modest advancement in generating such transfers and exploring the design space of all possible transfers. The greatest challenge faced so far is the extreme sensitivity of the dynamics – even the error in evaluating a collocation approximation can be great enough to change a solution.

Further study is warranted to explore a wider variety of initial guesses and their impact on the optimized trajectories. Although it is possible to find solutions even from poor initial guesses, the structure of the initial guess has a strong impact on the structure of converged solutions. Thus, different initial guesses should reveal even more distinct families of solutions. Exploration of low-thrust "manifolds" connecting the DRO and halo orbit should help identify more families of transfers – including some that travel to the opposite side of the Earth-Moon system and back again. Similarly, the size and out-of-plane motion of the DRO and L_2 halo orbit could be varied.

PSOPT is an effective tool, but it has limitations. The most significant limitation encountered in this research is that it cannot use adaptive mesh refinement for the pseudospectral method. Other implementations of pseudospectral optimal control permit automatically adding more phases as necessary.

In further study, more accurate dynamics should be used. This would involve using a full ephemeris model (rather than the CRTBP), and a shadowing model to cycle thrusting when the spacecraft is not in view of the Sun.

Other future research areas include examination of other Earth-Moon system transfers and reducing the maximum thrust level to the capability of existing EP systems. Doing so will require new initial guesses and possibly new methodology.

6. ACKNOWLEDGEMENTS

7. REFERENCES

[1] J. T. Betts and S. O. Erb, "Computing optimal low thrust trajectories to the moon," *European Space Agency, (Special Publication) ESA SP*, vol. 2, no. 516, pp. 143–146, 2003.

[2] G. Mingotti, F. Topputo and F. Bernelli-Zazzera, "Numerical Methods to Design Low-Energy, Low-Thrust Sun-Perturbed Transfers to the Moon," *49th Israel Annual Conference on Aerospace Sciences*, no. February 2016, 2009.

[3] M. T. Ozimek and K. C. 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, pp. 533–549, 2010.

[4] J. S. Parker, C. Bezrouk and K. E. Davis, "Low-Energy Transfers to Distant Retrograde Orbits," in *Advances in the Astronautical Sciences Spaceflight Mechanics 2015*, 2015, p. AAS 15–311.

[5] J. F. C. Herman and J. S. Parker, "Low-energy, low-thrust transfers between earth and distant retrograde orbits about the moon," *AAS Guidance, Navigation & Control Conference*, pp. 1–11, 2015.

[6] J. Heiligers, G. Mingotti and C. McInnes, "Optimisation of Solar Sail Interplanetary Heteroclinic Connections," in *2nd Conference on Dynamics and Control of Space Systems*, 2014.

[7] J. Heiligers, G. Mingotti and C. R. McInnes, "Optimal solar sail transfers between Halo orbits of different Sun-planet systems," *Advances in Space Research*, vol. 55, no. 5, pp. 1405–1421, 2015.

[8] F. Topputo, "On optimal two-impulse Earth-Moon transfers in a four-body model," *Celestial Mechanics and Dynamical Astronomy*, vol. 117, no. 3, pp. 279–313, 2013.

[9] J. S. Parker, "Families of low-energy lunar halo transfers," *Advances in the Astronautical Sciences*, vol. 90, no. JANUARY 2006, p. AAS 06–132, 2006.

[10] K. E. Davis and J. S. Parker, "Prograde Lunar Flyby Trajectories from Distant Retrograde Orbits," 2015, pp. 1–12.

[11] V. G. Szebehely, *Theory of Orbits: The Restricted Problem of Three Bodies*. 1967.

[12] R. W. Farquhar, "The utilization of halo orbits in advanced lunar operations," 1971.

[13] L. Alkalai, J. Hopkins and B. L. Jolliff, "Orion/moonrise joint human-robotic lunar sample return mission concept," 2013.

[14] C. Bezrouk and J. Parker, "Long Duration Stability of Distant Retrograde Orbits," no. August, pp. 1–9, 2014.

[15] J. N. Pelton and F. Allahdadi, "Handbook of Cosmic Hazards and Planetary Defense," pp. 535–542, 2015.

[16] J. Hopkins, "OSCAR : Crew-Assisted Lunar Sample Return with Orion Orion Sample Capture and Return," 2014.

[17] G. Elnagar, M. A. Kazemi and M. Razzaghi, "The Pseudospectral Legendre Method for Discretizing Optimal Control Problems," *IEEE Transactions on Automatic Control*, vol. 40, no. 10, pp. 1793–1796, 1995.

[18] A. Wachter and L. T. Biegler, "On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming," *Mathematical Programming*, vol. 106, no. 1, pp. 25–57, 2006.

[19] D. Garg, M. A. Patterson, W. W. Hager, A. V. Rao, D. Benson and G. T. Huntington, "An Overview of Three Pseudospectral Methods for the Numerical Solution of Optimal Control Problems," in *AAS/AIAA Astrodynamics Specialist Conference, August 9-13, 2009, Pittsburgh, Pennsylvania*, 2009.

[20] I. M. Ross and F. Fahroo, "Legendre pseudospectral approximations of optimal control problems," *New Trends in Nonlinear Dynamics and Control and their Applications*, vol. 295, no. January, pp. 327–342, 2003.

[21] V. M. Becerra, "Solving complex optimal control problems at no cost with PSOPT," *Proceedings of the IEEE International Symposium on Computer-Aided Control System Design*, pp. 1391–1396, 2010.

[22] J. T. Betts, *Practical Methods for Optimal Control and Estimation Using Nonlinear Programming*. Society for Industrial and Applied Mathematics, 2010.

[23] F. Liu, W. W. Hager and A. V. Rao, "An hp Mesh Refinement Method for Optimal Control Using Discontinuity Detection and Mesh Size Reduction," in *53rd IEEE Conference on Decision and Control, Los Angeles*.