(Preprint) AAS 21-374

LOW-THRUST TRAJECTORY OPTIMIZATION USING THE KUSTAANHEIMO-STIEFEL TRANSFORMATION

Kevin Tracy*, Zachary Manchester*

We present a novel trajectory optimization formulation for the low-thrust orbital transfer problem utilizing the Kustaanheimo-Stiefel transformation. For unperturbed two-body motion, this transformation maps the nonlinear Cartesian dynamics to a linear four-dimensional simple-harmonic oscillator. When perturbing accelerations are added, the dynamics become nonlinear, but are significantly better approximated by linearization than alternative state representations. Therefore, the Kustaanheimo-Stiefel dynamics have strong advantages in gradient or Newton-based trajectory optimization algorithms. We formulate a low-thrust trajectory optimization problem with these dynamics and demonstrate empirically that thrust profiles can be found without providing an initial guess to the solver, and with fewer knot points than alternative state representations.

INTRODUCTION

Low-thrust trajectory optimization has been extensively studied over the past 30 years, motivated in part by the increasing popularity of electric propulsion systems.[1] Many orbital state representations have been used for this problem, including classical orbital elements, modified equinoctial orbital elements, and Cartesian position and velocity.[2,3] Unfortunately, all of these representations result in highly nonlinear dynamics with singularities, and are susceptible to severe numerical scaling issues. Another aspect of this problem that makes optimization challenging is the fact that low-thrust maneuvers can take months to complete. A common maneuver that this paper focuses on is the transfer from Geostationary Transfer Orbit (GTO) to Geostationary Orbit (GEO). This maneuver can take anywhere from one to six months, with the resulting trajectory requiring thousands of knot points to represent numerically.

Due to the problem size and nonlinearity of the dynamics, a sophisticated initial guess is often required to navigate a cost landscape that is riddled with local minima. The resulting formulations are large, poorly numerically conditioned, and require significant computational effort. These difficulties have motivated Lyapunov-based online feedback methods, such as those described by Petropoulos and later modified by Varga.[4,5] These control laws are simple to run on-board, benefit from closed-form solutions, and are unaffected by the nonlinear/underactuated nature of the spacecraft. Unfortunately, they are generally sub-optimal, and offline tuning of the controller gains can require substantial effort.

In contrast, our formulation of the low-thrust trajectory optimization problem leverages the Kustaanheimo-Stiefel (KS) transformed orbital dynamics. Trajectory optimization with the KS transformation has been investigated by Thorne[6] and Hernandez,[7] but both cases are limited to

*The Robotics Institute, Carnegie Mellon University, 5000 Forbes Avenue, Pittsburgh, PA 15213.

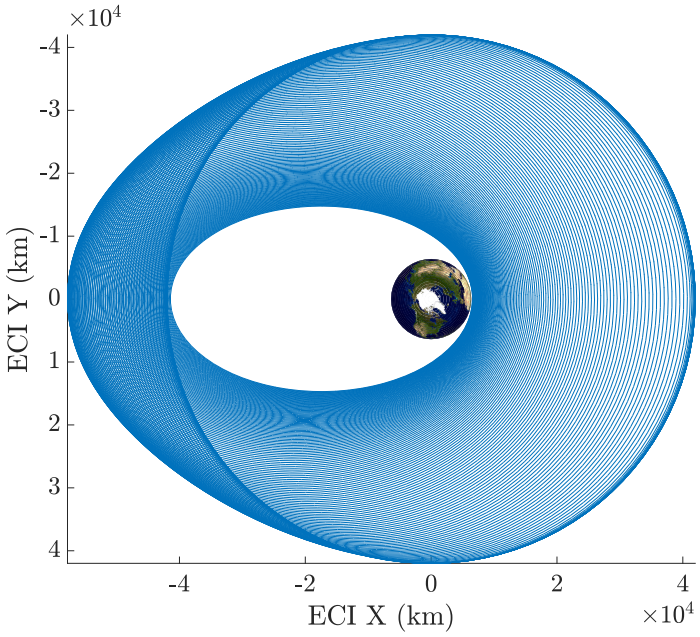

Figure 1. Top-down view for optimized low-thrust 100-day GTO to GEO trajectory

short-duration in-plane transfers of two-dimensional orbits. This paper treats the full three-dimensional case. Our algorithm combines the ease of use of Lyapunov-based methods with the performance and (local) optimality of trajectory optimization for use in three-dimensional orbit transfers involving plane changes. When compared to alternative state representations, the KS-transformed orbital mechanics result in better conditioning and fewer knot points, all without the requirement for a sophisticated initial guess.

Our contributions in this paper are:

1. Using the KS transformation for three-dimensional low-thrust trajectory optimization.

2. A cost function that leverages the mapping between KS state and Cartesian position to perform GTO to GEO transfers.

3. Demonstration of reliable convergence to locally optimal GTO to GEO transfer solutions without an initial guess.

The paper proceeds as follows: The Background section discusses the historical development of the KS transformation, fictitious time, and the differential equations describing unperturbed and perturbed orbital motion. The Trajectory Optimization section describes the formulation of our optimization problem and the cost function used to drive the GTO to GEO transfer solution. The Numerical Experiments section examines successfully optimized low-thrust transfers lasting 30, 60, and 100 days. Finally, the Discussion and Conclusions section elaborates on the impacts of this research for mission operations, and the importance of converging without an initial guess.

BACKGROUND

In 1765 Leonard Euler introduced a linear representation of pure Keplerian motion in one dimension,[8] achieved by storing the square root of distance from the central body and introducing a fictitious time that varies with this distance. The fictitious time slows down near perigee and speeds up near apogee, resulting in a linear simple-harmonic oscillator. The idea was later extended to two dimensions with the Levi-Civita transformation, representing the spacecraft's two-dimensional position as a scaled rotation of the unit-length x-axis.[9] This rotation is expressed using a complex number with a magnitude equal to that of the square root of the radius. Finally, two hundred years after Euler's discovery, Kustaanheimo and Stiefel introduced the three dimensional generalization by expressing the position of the spacecraft again as a scaled rotation from the unit-length x-axis. Where this rotation was expressed using an imaginary number for the two-dimensional case, a quaternion is used in the full three-dimensional case.[9] The KS transformation maps the full two-body Kepler dynamics to a four-dimensional simple-harmonic oscillator.[8,10] By relying on the same fictitious time established by Euler, uniform step sizes can be used for numerical integration, as the dynamics evolve at a constant rate throughout the orbit. This is a powerful feature, as it allows for accurate orbital propagation with fixed-step integrators. The uniformity and linearity in the KS dynamics present significant advantages for optimizing low-thrust trajectories.

Position Mapping

The KS Transform maps Cartesian inertial positions, $x \in \mathbb{R}^3$, to a quaternion, $p \in \mathbb{R}^4$. This quaternion describes the rotation and scaling from the unit-length inertial x-axis to the spacecraft position vector. Since this rotation is not unique, there are an infinite number of p's corresponding to each x, so for the initial conversion, p_1 will be arbitrarily set to 0.1. We will use the following mapping from $x \in \mathbb{R}^3$ to $p \in \mathbb{R}^4$:

$$p_1 = .1, \tag{1}$$

$$p_4 = \sqrt{.5(\|x\| + x_1) - p_1^2}, \tag{2}$$

$$p_2 = \frac{x_2 p_1 + x_3 p_4}{\|x\| + x_1}, \tag{3}$$

$$p_3 = \frac{x_3 p_1 - x_2 p_4}{\|x\| + x_1}. \tag{4}$$

Even though there are four parameters, there remain only three degrees of freedom. Our p parameter has a norm constraint that relates to the corresponding x:

$$\|p\| = \sqrt{\|x\|}. \tag{5}$$

The corresponding inverse transformation is,

$$x_1 = p_1^2 - p_2^2 - p_3^2 + p_4^2, \tag{6}$$

$$x_2 = 2(p_1 p_2 - p_3 p_4), \tag{7}$$

$$x_3 = 2(p_1 p_3 + p_2 p_4), \tag{8}$$

which can be recognized as the vector $[0, 0, 1]^T$ after being rotated by a quaternion p.

Fictitious Time

Before addressing the velocity mapping, the relationship between true time (t) and fictitious time (s) must be established. Two different derivatives will be used in the following sections, those taken with respect to true time and those taken with respect to fictitious time. The former are denoted with a dot,

$$\dot{z} = \frac{dz}{dt}, \tag{9}$$

while derivatives taken with respect to fictitious time (s), are denoted with a prime:

$$z' = \frac{dz}{ds}. \tag{10}$$

The relationship between these two is dependent on the distance from the orbiting body to the central body:

$$dt = \|x\| ds, \tag{11}$$

$$\frac{dt}{ds} = \|x\|. \tag{12}$$

Using the relationship established in equation (5), we can write the following:

$$t' = \|p\|^2. \tag{13}$$

This relationship allows the eccentric anomaly of a spacecraft in Keplerian motion to progress linearly in fictitious time. The resulting orbit can be discretized with equal fictitious time steps, and the resulting points will also be uniformly discretized according to arc length. In order to keep track of the true time (t), we append it to the state and propagate it numerically with the rest of the dynamics.

Velocity Mapping

Cartesian velocities expressed in the inertial frame, $\dot{x} \in \mathbb{R}^3$, can be mapped to fictitious velocities of our parameter p:

$$p_1' = \frac{1}{2}(p_1\dot{x}_1 + p_2\dot{x}_2 + p_3\dot{x}_3), \tag{14}$$

$$p_2' = \frac{1}{2}(-p_2\dot{x}_1 + p_1\dot{x}_2 + p_4\dot{x}_3), \tag{15}$$

$$p_3' = \frac{1}{2}(-p_3\dot{x}_1 - p_4\dot{x}_2 + p_1\dot{x}_3), \tag{16}$$

$$p_4' = \frac{1}{2}(p_4\dot{x}_1 - p_3\dot{x}_2 + p_2\dot{x}_3), \tag{17}$$

where the inverse of this operation is:

$$\dot{x}_1 = \frac{2}{(p^T p)}(p_1 p_1' - p_2 p_2' - p_3 p_3' + p_4 p_4'), \tag{18}$$

$$\dot{x}_2 = \frac{2}{(p^T p)}(p_2 p_1' + p_1 p_2' - p_4 p_3' - p_3 p_4'), \tag{19}$$

$$\dot{x}_3 = \frac{2}{(p^T p)}(p_3 p_1' + p_4 p_2' + p_1 p_3' + p_2 p_4'). \tag{20}$$

These transformations are unique, and are derived from the relationship described in equation (5).

Unperturbed Dynamics

For the unperturbed case, the dynamics are conventional Keplerian motion. We refer to the following state vector, $\mathbf{X}_{KS} \in \mathbb{R}^{10}$, as the ***KS state***:

$$\mathbf{X}_{KS} = \begin{bmatrix} p \\ p' \\ h \\ t \end{bmatrix}, \tag{21}$$

where h is the negative specific orbital energy:

$$h = -\epsilon = \frac{\mu}{\|x\|} - \frac{\|\dot{x}\|^2}{2}. \tag{22}$$

For the unperturbed case, this energy is constant, allowing us to simplify the equations of motion to:

$$p'' = -\frac{h}{2}p, \tag{23}$$

$$h' = 0, \tag{24}$$

$$t' = \|p\|^2. \tag{25}$$

Upon visual inspection, the dynamics for the quaternion p can be identified as a linear simple-harmonic oscillator. This oscillator is undamped with an effective stiffness-to-mass ratio equal to one half of the specific orbital energy. The solutions to this second-order differential equation are sinusoidal with constant frequencies.

Perturbed Dynamics

In order to incorporate perturbations like non-spherical gravity, atmospheric drag, and spacecraft propulsion, equations of motion that include additional forces must be formulated. For ease of notation, the matrix function $L(p)$ is defined,

$$L(p) = \begin{bmatrix} p_1 & -p_2 & -p_3 & p_4 \\ p_2 & p_1 & -p_4 & -p_3 \\ p_3 & p_4 & p_1 & p_2 \\ p_4 & -p_3 & p_2 & -p_1 \end{bmatrix}. \tag{26}$$

The perturbing accelerations will be collected and combined into one inertial acceleration vector, a_p. With this, the perturbed dynamics can be written as follows:

$$p'' = -\frac{h}{2}p + \frac{\|p\|^2}{2}[L(p)]^T \begin{bmatrix} a_p \\ 0 \end{bmatrix}, \tag{27}$$

$$h' = -2[p']^T[L(p)]^T \begin{bmatrix} a_p \\ 0 \end{bmatrix}, \tag{28}$$

$$t' = \|p\|^2. \tag{29}$$

The resulting equations of motion are nonlinear, but are in a relatively benign class referred to as "perturbed linear".[10] They behave far better under linearization than other state representations.

TRAJECTORY OPTIMIZATION

Optimal control problems are often discretized and formulated as numerical optimization problems with direct transcription,[11] where both the states and the controls are treated as decision variables:

$$\begin{aligned} \underset{x_{1:N},\, u_{0:N-1}}{\text{minimize}} \quad & \ell_N(x_N) + \sum_{k=0}^{N-1} \ell_k(x_k, u_k) \\ \text{subject to} \quad & x_{k+1} = f(x_k, u_k), \qquad k = 0, \ldots, N-1, \\ & g_k(x_k, u_k) \leq 0 \qquad \forall k, \\ & c_k(x_k, u_k) = 0 \qquad \forall k. \end{aligned} \tag{30}$$

Here the states and controls must minimize the cost functions $\ell(x_k, u_k)$ and $\ell_N(x_n)$, subject to the nonlinear dynamics, $f(x_k, u_k)$, and the the constraint functions $g_k(x_k, u_k)$ and $c_k(x_k, u_k)$. Many software tools exist for converting continuous trajectory optimization problems into this form, including GPOPS-II,[12] and general purpose nonlinear programming (NLP) solvers like SNOPT[13] and IPOPT[14] can then be used to solve them. While these methods have significant heritage, general purpose NLP solvers fail to exploit many of the unique structural features present in trajectory optimization problems.

Instead, the trajectories in this paper are optimized using the Augmented Lagrangian Trajectory Optimizer (ALTRO),[15] which exploits the Markovian structure of (30) for faster convergence. ALTRO is based on Differential Dynamic Programming (DDP), with an augmented Lagrangian framework for handling both conic and non-convex constraints. In particular, ALTRO is able to natively handle convex second-order cone (SOC) constraints without having to rely on linear approximations.[16] This feature enables ALTRO to handle norm constraints on thrust vectors efficiently and without suffering from numerical ill-conditioning that can occur with alternative methods. ALTRO can often be orders of magnitude faster than traditional direct transcription approaches.[15]

Cost Function

For the GTO-to-GEO problem, the spacecraft starts in an inclined elliptical orbit and must raise the semi-major axis while also eliminating both eccentricity and inclination. While easy to express in terms of orbital elements, some care must be taken to formulate this cost in terms of the KS state. The first part of the cost function penalizes deviations from the geostationary orbit radius:

$$\min_x \; [(\|x\| - r_{geo})^2]. \tag{31}$$

Driving this cost term to zero achieves both the desired semi-major axis and zero eccentricity. Inclination is driven to zero by penalizing the z-component of the spacecraft's Cartesian position:

$$\min_x [x_3^2]. \tag{32}$$

These cost terms are convex in terms of x, but the mapping from p to x is nonlinear. To address this, the state vector will be augmented to include these Cartesian terms, and they will be propagated

through the integrator. This simply shifts non-convexity from the constraints into the dynamics:

$$\mathbf{X} = \underset{12\times1}{\begin{bmatrix} p \\ p' \\ h \\ t \\ \|x\| \\ x_3 \end{bmatrix}} = \begin{bmatrix} p \\ p' \\ h \\ t \\ \|p\|^2 \\ 2(p_1 p_3 + p_2 p_4) \end{bmatrix}. \tag{33}$$

The cost function can now be computed easily using only the elements in the state,

$$\ell(\mathbf{X}) = \|u\|^2 + \alpha(\|x\| - r_{geo})^2 + \beta x_3^2, \tag{34}$$

where α and β are tuning parameters that can be adjusted for each mission scenario.

Finally, because the ALTRO solver only natively supports zero-order hold control input discretization, we further augment the state vector with an integrator on the controls to better approximate the continuously moving thrust vector. The controls solved for in ALTRO are now the derivative of the thrust vector with respect to the fictitious time, allowing the resulting integrated thrust vector to linearly interpolate between knot points. The full set of ODEs as they appear in the solver are as follows:

$$\mathbf{X}'_{ALTRO} = \begin{bmatrix} p \\ p' \\ h \\ t \\ \|x\| \\ x_3 \\ u \end{bmatrix}' = \begin{bmatrix} p' \\ -\frac{h}{2}p + \frac{(p^T p)}{2} L(p)]^T \begin{bmatrix} a_p \\ 0 \end{bmatrix} \\ -2[p']^T [L(p)]^T \begin{bmatrix} a_p \\ 0 \end{bmatrix} \\ p^T p \\ 2p^T p' \\ 2(p'_1 p_3 + p_1 p'_3 + p'_2 p_4 + p_2 p'_4) \\ u' \end{bmatrix}. \tag{35}$$

With these dynamics, the final trajectory optimization problem can be formulated as,

$$\begin{aligned} \underset{u'_{0:N-1}}{\text{minimize}} \quad & \sum_{k=1}^{N} \|u_k\|^2 + \alpha(\|x_k\| - r_{geo})^2 + \beta x_{3k}^2 \\ \text{subject to} \quad & x_0 = x_{init}, \\ & x_{k+1} = f(x_k, u'_k), \qquad k = 0, \ldots, N-1, \\ & \|u_k\| \leq u_{max}, \qquad k = 0, \ldots, N-1, \end{aligned} \tag{36}$$

where $f(x_k, u'_k)$ describes the discretized dynamics presented in (35), integrated with a fixed-step 4th-order Runge-Kutta scheme.

NUMERICAL EXPERIMENTS

To demonstrate the numerical advantage of KS-transformed dynamics, the energy behaviors of three common orbital state representations are compared in Fig 2. A geostationary transfer orbit was used for the comparison, with the only perturbation coming from J2. All three state representations

were integrated with 100 steps per orbit, with equinoctial and Cartesian using true time, and the KS transformed dynamics using fictitious time. This fictitious time allows for efficient sampling of the dynamics function throughout the orbit. The frequency in which the dynamics function is sampled is directly related to the speed of the dynamics, with sampling happening more often at perigee than apogee. The result is consistent energy behavior, even with relatively few knot points. This is noteworthy compared to the equinoctial propagation that loses energy during every orbit, and the Cartesian propagation that sees a spike in energy before going completely unstable.

For our trajectory optimization experiments, a 10,000 kg spacecraft traveling from GTO to GEO was used to explore three possible propulsion scenarios: a 30-day transfer with 1 N of thrust, a 60-day transfer with 0.59 N of thrust, and a 100-day transfer with 0.32 N of thrust. The starting GTO was inclined at 27°, with a perigee at 200 km, and apogee at 35,000 km. Each solve was initiated with zeros for the thrust controls. In all three of these scenarios, ALTRO reliably converged to a solution with only this naive guess. We intentionally did not exploit any *a priori* knowledge about the solution to warm-start the solver. The timing results for the three solver runs are presented in Table 1, with all of the examples run on a laptop computer with a 2.5GHz quad-core Intel i7-4870HQ and 16GB of RAM. The number of knot points and the magnitude of the thrust constraint were the determining factors in time for convergence, with the 100-day transfer taking the longest time at 30.92 minutes.

The solutions for all three scenarios are plotted in Fig. 3. In every case, the semi-major axis, eccentricity, and inclination all arrived to the target values around the same time. There is a tail to each of the curves that can be seen in Fig. 4 as the time where the thrust is being throttled. For all three of the thrust magnitudes, the spacecraft is commanding maximum thrust for almost the entirety of the transfer. Towards the very end of the transfer when the spacecraft is nearly at GEO, there is a minor tail in the errors on the orbital elements. This is an expected result when using quadratic cost functions. An example output of the thrust plan from ALTRO is illustrated in figure 5. When looking at the full thrust plan, it is difficult to identify the true shape of the solution. The two subplots with closer views show the periodicity of the solution, something that would be a reasonable expectation for an optimal thrust solution. To give an idea of the sort of shape that one of these trajectories produces, a top-down view of the 100 day transfer is shown in Fig 1. All of the code to reproduce these examples is available at `github.com/RoboticExplorationLab/KSLowThrust`.

Transfer Time	**Knot Points**	**Iterations**	**Solve Time**
30-day	1501	2,763	2.86 min
60-day	2801	6,158	12.56 min
100-day	3501	10,789	30.92 min

Table 1. Timing results for ALTRO to converge on solutions for the GTO to GEO trajectory optimization problems of varying transfer time.

DISCUSSION AND CONCLUSIONS

The solution for the 100-day transfer was expressed using only 3,501 knot points. This was possible for the KS dynamics due to the fictitious time, but would be intractable with other representations. For the state representations that use true time, this would have resulted in knot points

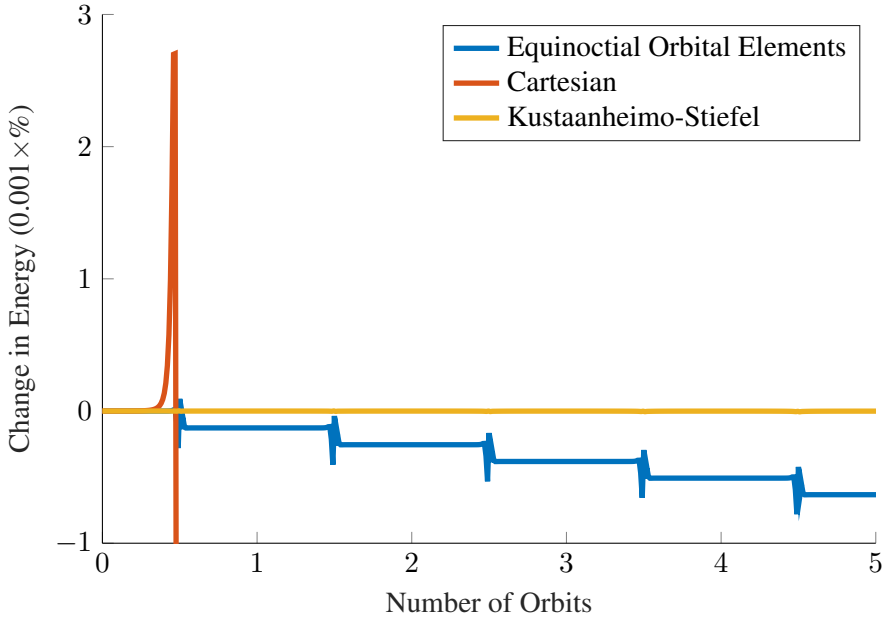

Figure 2. Specific orbital energy behavior for three orbital state representations in a J2 perturbed GTO. Integration was performed with a 4th-order fixed-step Runge-Kutta scheme with 100 knot points per orbit. Both equinoctial and Cartesian state representations see inaccurate energy behavior, with the spacecraft losing energy when it should be constant. The KS state representation maintains the expected consistent energy despite the large time-steps.

being over 45 minutes apart, and the integration errors near perigee would have made the solution too inaccurate to be useful. The KS dynamics enable a significantly reduced problem size, aiding in both the solve speed and numerical conditioning.

A significant contribution of this work is the ability to converge on a low-thrust orbital transfer without an initial guess. Due to the nonlinear dynamics constraints, low-thrust trajectory optimization problems are non-convex. This means that the cost function can have many local minima, with no guarantee of convergence to a global minimum. In order to aid in convergence to a favorable minimum, trajectory optimization solvers are often instantiated with a guess trajectory. This guess normally requires significant insight in to the problem and trial and error to ensure convergence to a desirable local minimum. Because of this, one trajectory optimization problem can have many different solutions depending on the fidelity of the guess trajectory. With the method outlined in this paper, we have introduced a formulation of a low-thrust GTO to GEO transfer that requires no initial guess for reliable convergence. ALTRO is simply instantiated with zeros for the controls, and the initial state trajectory is a forward rollout of the dynamics without any thrust. This feature makes the formulation in this paper attractive for real-time mission planning during low-thrust orbit raising. During a long trajectory, the trajectory optimization problem could be reliably and efficiently resolved given updated initial conditions to account for potential unmodeled effects in space missions.

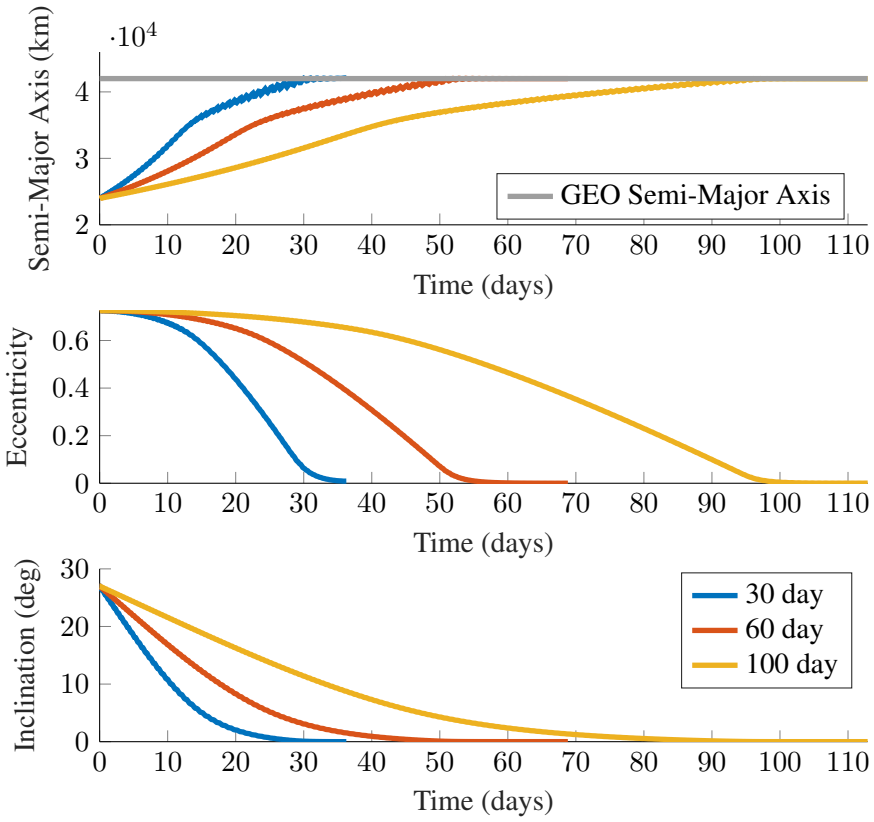

Figure 3. Orbital elements for GTO to GEO transfers. Target semi-major axis was that of GEO, and the target eccentricity and inclination were both zero.

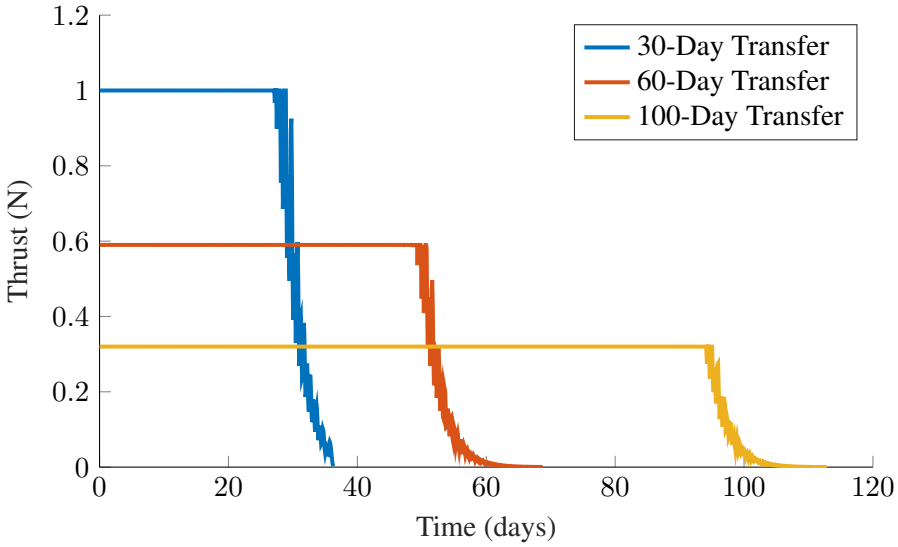

Figure 4. Thrust magnitudes for GTO to GEO transfers. The max thrust constraint is satisfied down to a tolerance of 1e-6, with the spacecraft at 100% throttle for almost the entirety of the transfer.

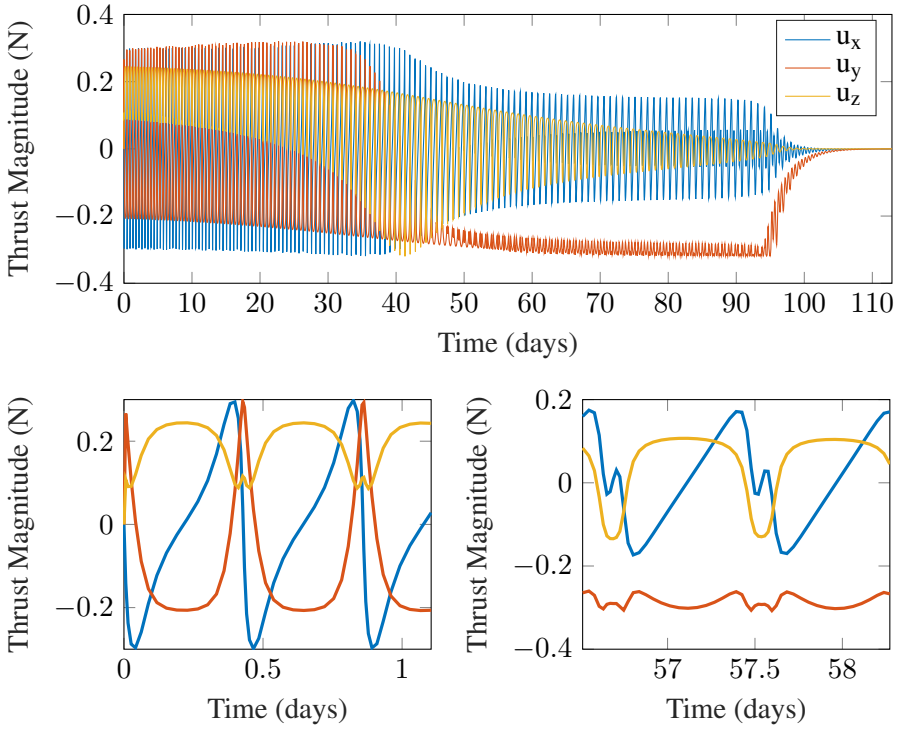

Figure 5. ALTRO thrust solution for the 100-day GTO to GEO transfer. The full solution is on top, with two alternate zoomed in views on the bottom. These close up views show the periodicity of the solution and the impact that a first-order hold had on the shape of the thrust plan.

REFERENCES

[1] T. Polsgrove, L. Kos, R. Hopkins, and T. Crane, "Comparison of Performance Predictions for New Low-Thrust Trajectory Tools," *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Keystone, Colorado, American Institute of Aeronautics and Astronautics, Aug. 2006, 10.2514/6.2006-6742.

[2] J. T. Betts and S. O. Erb, "Optimal Low Thrust Trajectories to the Moon," *SIAM Journal on Applied Dynamical Systems*, Vol. 2, Jan. 2003, pp. 144–170, 10.1137/S1111111102409080.

[3] J. T. Betts, *Practical Methods for Optimal Control Using Nonlinear Programming*, Vol. 3 of *Advances in Design and Control*. Philadelphia, PA: Society for Industrial and Applied Mathematics (SIAM), 2001.

[4] A. Petropoulos, "Low-Thrust Orbit Transfers Using Candidate Lyapunov Functions with a Mechanism for Coasting," *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, Providence, Rhode Island, American Institute of Aeronautics and Astronautics, Aug. 2004, 10.2514/6.2004-5089.

[5] G. I. Varga and J. M. S. Perez, "MANY-REVOLUTION LOW-THRUST ORBIT TRANSFER COMPUTATION USING EQUINOCTIAL Q-LAW INCLUDING J2 AND ECLIPSE EFFECTS," p. 10.

[6] J. D. Thorne and C. D. Hall, "Minimum-Time Continuous-Thrust Orbit Transfers Using the Kustaanheimo-Stiefel Transformation," *Journal of Guidance, Control, and Dynamics*, Vol. 20, May 2012, p. 3.

[7] S. Hernandez and M. R. Akella, "Lyapunov-Based Guidance for Orbit Transfers and Rendezvous in Levi-Civita Coordinates," *Journal of Guidance, Control, and Dynamics*, Vol. 37, July 2014, pp. 1170–1181, 10.2514/1.62305.

[8] T. Bartsch, "The Kustaanheimo-Stiefel Transformation in Geometric Algebra," *Journal of Physics A: Mathematical and General*, Vol. 36, June 2003, pp. 6963–6978, 10.1088/0305-4470/36/25/305.

[9] J. Waldvogel, "Quaternions and the Perturbed Kepler Problem," p. 14.

[10] E. L. Stiefel and G. Scheifele, *Linear and Regular Celestial Mechanics*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1971, 10.1007/978-3-642-65027-7.

[11] C. R. Hargraves and S. W. Paris, "Direct Trajectory Optimization Using Nonlinear Programming and Collocation," *J. Guidance*, Vol. 10, No. 4, 1987, pp. 338–342.

[12] M. A. Patterson and A. V. Rao, "A General-Purpose MATLAB Software for Solving Multiple-Phase Optimal Control Problems Version 2.3," p. 71.

[13] P. E. Gill, W. Murray, and M. A. Saunders, "SNOPT: An SQP Algorithm for Large-Scale Constrained Optimization," *SIAM Review*, Vol. 47, Jan. 2005, pp. 99–131, 10.1137/S0036144504446096.

[14] A. Wächter and L. T. Biegler, "On the Implementation of an Interior-Point Filter Line-Search Algorithm for Large-Scale Nonlinear Programming," *Mathematical Programming*, Vol. 106, Mar. 2006, pp. 25–57, 10.1007/s10107-004-0559-y.

[15] T. A. Howell, B. E. Jackson, and Z. Manchester, "ALTRO: A Fast Solver for Constrained Trajectory Optimization," *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Macau, China, Nov. 2019.

[16] B. E. Jackson, T. Punnoose, D. Neamati, K. Tracy, and R. Jitosho, "ALTRO-C: A Fast Solver for Conic Model-Predictive Control," *International Conference on Robotics and Automation (ICRA)*, Xi'an, China, 2021, p. 8.