AIAA/AAS Astrodynamics Specialist Conference and Exhibit
18 - 21 August 2008, Honolulu, Hawaii

AIAA 2008-6955

Low-Thrust Transfers using Primer Vector Theory and a Second-Order Penalty Method

Anastassios E. Petropoulos*
Jet Propulsion Laboratory, California Institute of Technology, Pasadena, California 91109

Ryan P. Russell†
Georgia Institute of Technology, Atlanta, Georgia 30332

Abstract

The low-thrust trajectory problem is formulated using techniques from calculus of variations, parameter optimization, and differential dynamic programming. Primer vector theory provides the thrust profile as a feedback law, parameterized by the initial co-states. The first and second derivatives of various penalty functions are derived from principles of static-dynamic control and dynamic programming. Both bang-bang/constant ejection velocity and variable specific impulse cases are considered, the former being complicated by the mapping of derivatives across switching times. A trust-region optimizer is then developed, where these derivatives are used to solve transfers between periodic orbits in the circular, restricted, three-body problem.

I. Introduction

The low-thrust trajectory problem has been studied by a multitude of researchers over the past several decades, yielding a variety of techniques applicable to space missions. In this paper, several existing proven methods are combined and extended to build a low-thrust transfer algorithm in the circular restricted three-body problem (CR3BP), enjoying benefits from calculus of variations, parameter optimization, and differential dynamic programming (DDP).

The thrust direction is chosen based on the well-known primer vector control law[1] that arises from an application of calculus of variations to the low-thrust trajectory problem. While the control at each time along the trajectory is replaced by a feedback law, the application of primer vector theory doubles the length of the state vector with the addition of co-state variables that are governed by the Euler-Lagrange equations. For constant ejection velocity (CEV) engines, the calculus of variations further dictates a bang-bang thrusting structure where the thrust is either on or off depending on the sign of a switching function (ignoring the indeterminate case where the function is zero). Numeric implementation of the bang-bang thrusting structure is most easily accommodated through predefined switching times which are iteratively corrected during the solution process. However, in many problems the number of potential switching times is unknown a priori, and it may not be practical to constantly adjust the structure of the problem each time the switching function suggests a new coast arc should appear. Reference 2 outlines a method for implementing the bang-bang structure where the switching times are implicitly calculated during a trajectory propagation and the necessary derivatives are mapped across each switching boundary. Presently, the implicit bang-bang switching structure is extended to second order. The variable specific impulse (VSI) case is also considered, where the derivatives are greatly simplified by the absence of a switching function.

*Senior Member of the Engineering Staff, Guidance, Navigation and Control Section, Mail-Stop 301-121, 4800 Oak Grove Dr. Member AIAA. Member AAS. Email: Anastassios.Petropoulos@jpl.nasa.gov

†Assistant Professor, Guggenheim School of Aerospace Engineering, 270 Ferst Drive. Email: ryan.russell@gatech.edu. Formerly at the Jet Propulsion Laboratory.

A second order trust region algorithm[3] based on analytic derivatives is sought to emulate the robustness and large radius of convergence observed in the Static Dynamic Control (SDC) algorithm that is implemented in Mystic,[4] the high fidelity trajectory optimization tool under development at the Jet Propulsion Laboratory.

The problem is posed as a classic optimal control problem with a scalar performance index that is comprised of initial, final, and integrated cost functions. Constraints are handled through weighted penalty functions. The gradient and Hessian of the performance index with respect to the optimization parameters are calculated via a backward sweep method similar to that found in SDC[4] and DDP.[5] Note that the method does not require numerical integration of state transition matrices or state transition tensors. The optimization variables consist of the initial co-states (they dictate the thrust direction profile), and other variables such as initial position, initial velocity, initial time, flight time, or parameters associated with a particular initial or final orbit. The optimal control problem is reduced to a low-dimensioned second-order parameter optimization problem. An elliptical trust region method[3] is implemented to directly iterate on the unknown variables until the necessary and sufficient conditions of optimality are achieved. The second-order trust region method is favorable for robust convergence especially in the highly non-linear problems such as those typically associated with multi-body low-thrust trajectories.

The method is specifically applied to low-thrust transfers between orbits of the halo family in the Sun-Earth CR3BP. It is also applied to transfers between two distant retrograde orbits (DROs) in the Europa-Jupiter CR3BP. For highly sensitive trajectories with longer flight times, the formulation is amenable to multiple shooting.

II. Dynamical model and primer vector control law

The well-known CR3BP dynamics are used, wherein two primary bodies (e.g. Jupiter and Europa), taken as point masses, move in circular orbits around their centre of mass, while a third body of negligible mass (e.g. a spacecraft) moves under the gravitational influence of the primaries. In addition, the spacecraft is subject to the thrust acceleration of a low-thrust propulsion system. The equations of motion are expressed in the customary dimensionless, rotating coordinate frame where the primaries are fixed on the x-axis at -1 and 0 (origin is at the lighter primary), the z-axis is along their inertial orbital angular momentum vector, and the y-axis is given by $\hat{z} \times \hat{x}$:

$$\begin{aligned} \ddot{x} - 2\dot{y} &= \frac{\partial \Omega}{\partial x} + f_x \\ \ddot{y} + 2\dot{x} &= \frac{\partial \Omega}{\partial y} + f_y \\ \ddot{z} &= \frac{\partial \Omega}{\partial z} + f_z \end{aligned} \tag{1}$$

where

$$\Omega = \frac{1}{2}[(x + 1 - \mu)^2 + \frac{1-\mu}{r_1} + \frac{\mu}{r_2} \tag{2}$$

where μ is the customary mass ratio of the system and corresponds to the nondimensional distance of the heavier primary from the centre of mass, r_1, r_2 denote the distances between the spacecraft and the heavier and lighter primaries, respectively: $r_1^2 = (x+1)^2 + y^2 + z^2$ and $r_2^2 = x^2 + y^2 + z^2$, and f_x, f_y, f_z are the components of the thrust acceleration, f.

Two different simple thruster models are used. In both cases

$$\begin{aligned} \dot{m} &= -\frac{fm}{I_{sp} g_0} \\ f &= \frac{2\eta P}{m I_{sp} g_0} \\ 0 \leq P &\leq P_{max} \end{aligned}$$

where I_{sp} is the engine specific impulse, m is the mass of the spacecraft, $g_0 = 9.80665\text{m/s}^2$ is the standard acceleration due to gravity, P is the engine input power, η is the engine efficiency and P_{max} is the maximum

input power. In the constant ejection velocity (CEV) case, I_{sp} is held constant while P is allowed to vary. In the variable specific impulse (VSI) case, not only can P vary, but so can the I_{sp}, without limit as long as it is positive.

The primer vector control law is developed as normal (Ref. 2 provides a more complete and nuanced presentation) using the Euler-Lagrange equations.

$$\begin{aligned} H &\equiv \vec{\lambda} \cdot \vec{X} \\ \dot{\vec{\lambda}} &= -\frac{\partial H}{\partial \vec{X}} \\ \vec{0} &= \frac{\partial H}{\partial \vec{u}} \end{aligned}$$

where H is the Hamiltonian, $\vec{X} = [x, y, z, \dot{x}, \dot{y}, \dot{z}, m]$, $\vec{\lambda}$ is the vector of co-states, and $\vec{u}$ is the vector of control variables. In the CEV case, $\vec{u}$ contains the direction cosines of the thrust vector and the power, P. In the VSI case, $\vec{u}$ additionally contains the I_{sp}. In both cases, the optimal thrust direction must lie along the primer vector, $-\vec{\lambda}_v$ (components 4-6 of $\vec{\lambda}$). Pontryagin's maximum principle dictates that in the CEV case, the power (or equivalently the thrust magnitude) must follow a P_{min}-P_{max} bang-bang profile according to a switching function, while in the VSI case that $P = P_{max}$ and $I_{sp} = P|\vec{\lambda}_v|/(m\lambda_m)$, λ_m being the last component of $\vec{\lambda}$.

The equations of motion and co-state equations are numerically integrated, allowing the thrust vector to be completely determined at each integration step. The initial values of the co-states become part of the optimisation problem below, and may be thought of as parameters for the thrust profile expressed as a feedback control law.

Sensitivity calculations across a switching time

Most gradient based parameter optimization algorithms require continuity in the objective and its derivatives in order to forecast favorable search directions. The bang-bang control problem is historically difficult to optimize because the number of switches in the control is generally unknown a priori. If the bang-bang control law is strictly adhered to based on the sign of the derived switching function, then there is no need to specify the number of thrust arcs during a particular trajectory. In this regard, the search space is dramatically reduced because switching times can be removed from the unknown vector and the requirement that the switching function be zero at the switch can further be removed from the constraint vector. In order to maintain continuity at a fine resolution in the performance index, the implicit bang-bang control law must interpolate the switch times to exactly (machine precision) a switching function of zero. While this small requirement addresses continuity in the objective, a discontinuity exists in the derivatives at the switch points. Updating the derivatives at each switch time is not trivial to derive, but simple to implement. Reference[2] gives a detailed derivation for the first order derivative discontinuities at an implicit switch. Here, we extend the work to second order (i.e., we need Jacobian and Hessian of each state variable with respect to the state vector across a switch time). The equations deal effectively with first and second order sensitivities across a switch. Successive application of Eqs. 3 and 4 will handle trajectories with multiple switches. There is no fundamental limit on the number of switches; however, the practical limit depends on the sensitivity characteristics of a particular problem. In this paper we provide an example involving 42 switches.

Also, we should mention that discontinuities still arise in this problem when the number of thrust arcs changes from iteration to iteration. This is inherent to the method and often leads to chatter between iterations. In such cases, we find it is best to start over in a different local basin of attraction.

As might be expected, we will find useful the following first- and second-order derivatives, which show how changing the state and co-state just before a switching time affects the state just after the switch. The state and co-state are lumped into a single vector, $\mathbf{X}$, for convenience.

$$\frac{\partial \mathbf{X}_+}{\partial \mathbf{X}_-} = \mathbf{I} + (\mathbf{f}_+ - \mathbf{f}_-) \frac{1}{\dot{S}_-} \frac{\partial S_-}{\partial \mathbf{X}} \tag{3}$$

$$\frac{\partial^2 \mathbf{X}_+}{\partial \mathbf{X}_-^2} = \mathbf{A} + \mathbf{B} + \mathbf{C} \tag{4}$$

where

$$\mathbf{A} = (\mathbf{f}_+ - \mathbf{f}_-)\left\{-\frac{1}{\dot{S}_-^2}\left[\frac{\partial S_-}{\partial \mathbf{X}}^T \frac{\partial \dot{S}_-}{\partial \mathbf{X}} + \left(\frac{\partial \dot{S}_-}{\partial \mathbf{X}}\right)^T \frac{\partial S_-}{\partial \mathbf{X}}\right] + \frac{1}{\dot{S}_-}\frac{\partial^2 S_-}{\partial \mathbf{X}^2} + \frac{\ddot{S}_-}{\dot{S}_-^3}\left(\frac{\partial S_-}{\partial \mathbf{X}}\right)^T \frac{\partial S_-}{\partial \mathbf{X}}\right\}$$

$$\mathbf{B} = -\left(\dot{\mathbf{f}}_+ - \dot{\mathbf{f}}_-\right)\frac{1}{\dot{S}_-^2}\left(\frac{\partial S_-}{\partial \mathbf{X}}\right)^T \frac{\partial S_-}{\partial \mathbf{X}}$$

$$\mathbf{C} = -\frac{1}{\dot{S}_-}\left(\frac{\partial S_-}{\partial \mathbf{X}}^T \frac{\partial \mathbf{f}_-}{\partial \mathbf{X}} + \left(\frac{\partial \mathbf{f}_-}{\partial \mathbf{X}}\right)^T \frac{\partial S_-}{\partial \mathbf{X}}\right)$$

where $\mathbf{I}$ is the identity matrix, subscripts $_-$ and $_+$ denote quantities just before or just after the switch time, respectively, $\mathbf{f} \equiv \dot{\mathbf{X}}$, S is the switching function, and superscript ^{T} denotes the matrix transpose.

III. The optimisation problem and approach

To allow some flexibility in the types of problems that can be addressed, the cost function is assumed to be of the following form:

$$J = E(\vec{X}_0, t_0, \tau_0,) + \int_{t_0}^{t_f} F(\vec{X}, \vec{u}))\mathrm{d}t + G(\vec{X}_f, t_f, \tau_f) \tag{5}$$

where E, G and F are user-specified functions, t_0 and t_f are the initial and final times of the transfer, and τ_0 and τ_f are the initial and final values of τ, an angular phase variable which is applicable when transfer is sought to or from a specific point on a final or initial orbit. Initial and terminal constraints are thus handled by the penalty functions E and G, while quanitites commonly optimised in mission design, such as final mass, flight time or total impulse, can be handled either through the function G or through the integral term in the cost function.

The first and second derivatives of J with respect to the initial values of the states and co-states (as well as the other relevant problem variables), are calculated through a backward sweep of the adjoint equations as is common in DDP[5] and SDC.[4] The sensitivity equations are valid across a thrusting and coasting arc. In the VSI case, there are no coast arcs and therefore the derivatives of J with respect to the inital state are straightforward to obtain. In the CEV case, however, discontinuities in the derivatives exist due to the implicit following of the bang-bang switching profile. The derivatives therefore must be mapped across the switch times. To further simplify the equations and implementation, we also note that all of the variables (states, controls, co-states, times, and other optimization parameters) are incorporated into a single extended State vector, $\mathbf{Y}$. Evaluating second order Taylor series on the switching function, switch times, and dynamics, and using Eqs. 3 and 4 leads to Eqs. 6 and 7 which map the Jacobian and Hessian of the performance index with the extended State vector. The Jacobian and the Hessian will often be sparse. For brevity only the summarized results are given here.

$$\frac{\partial J}{\partial \mathbf{Y}_-} = \frac{\partial J}{\partial \mathbf{Y}_+}\frac{\partial \mathbf{Y}_+}{\partial \mathbf{Y}_-} \tag{6}$$

$$\frac{\partial^2 J}{\partial \mathbf{Y}_-^2} = \frac{\partial \mathbf{Y}_+}{\partial \mathbf{Y}_-}^T \frac{\partial^2 J}{\partial \mathbf{Y}_+^2}\frac{\partial \mathbf{Y}_+}{\partial \mathbf{Y}_-} + \frac{1}{2}\left(\mathbf{\Phi} + \mathbf{\Phi}^T\right) + \frac{1}{2}\left(\mathbf{\Psi} + \mathbf{\Psi}^T\right) \tag{7}$$

where

$$\mathbf{\Phi} = \frac{2}{\dot{S}_-}\frac{\partial S_-}{\partial \mathbf{Y}}^T \frac{\partial J}{\partial \mathbf{Y}_+}\frac{\partial \mathbf{f}_+}{\partial \mathbf{Y}}\frac{\partial \mathbf{Y}_+}{\partial \mathbf{Y}_-} \tag{8}$$

$$\mathbf{\Psi} = \frac{\partial J}{\partial \mathbf{Y}_+}\frac{\partial^2 \mathbf{Y}_+}{\partial \mathbf{Y}_-^2}^{T*} \tag{9}$$

and ^{T*} is the transpose of a 3-tensor such that if $B = A^{T*}$, then $B^{ijk} = A^{kji}$. The i^{th} row of $\mathbf{\Psi}$ is $\frac{\partial J}{\partial \mathbf{Y}_+}\left[\frac{\partial^2 \mathbf{Y}_+}{\partial \mathbf{Y}_-^2}^{T*}(:,:,i)\right]$.

To faciliate general use of the algorithm, an algebraic manipulator (MAPLE) is used to generate the equations of motion (*viz.* for the co-state) and partial derivatives necessary for a specific problem. The required user input is the general equations of motion (for the state), the performance index, and any constraints. The algebraic manipulator then evaluates the first and second order initial and terminal partial derivatives as well as the first and second order sensitivity equations of motion with respect to the full vector Y. Further, the output from the algebraic manipulator is Fortran code that is easily and directly included for compilation. We find that this system is remarkably accurate and removes many of the burdens typically associated with implementing analytic derivatives.

Finally, an elliptic trust-region optimiser[3] is developed based on these derivatives. This optimisation technique was chosen because of its promise of relatively rapid convergence and large radii of convergence.

IV. Applications

In this work, we consider orbit transfers from one halo orbit to another, and from one DRO to another. The position and velocity on the halo orbits are represented by 50^{th}-order Chebyshev polynomials where the independent variable, τ, is the angular position on the orbit measured around the x- axis. Similar curve fits are used for the DROs. Some of the transfers are phase free, that is they can start or end anywhere on the initial or final orbit (τ_0, τ_f are unconstrained). The transfers take place in the Sun-Earth or Jupiter-Europa systems, idealised by the CR3PB.

IV.A. CEV, constant-thrust halo-to-halo transfer

A partly phase-free, CEV transfer is sought between two halos of the same family about the L1 (interior) Lagrange point of the Jupiter-Europa system ($\mu = 25.3 \times 10^{-6}$). The initial state is a fixed point on the larger halo, while the final point is allowed to be anywhere on the smaller. The cost function involves the final mass, m_f, and penalty constraint terms for ensuring that the initial and final states are on the desired halos. That is, maximum final mass is sought. Constant thrust is assumed, that is, the switching function is ignored. Thus, the problem is equivalent to finding the minimum flight time. Although minimum flight time transfers are seldom used in practice, the simpler computations that are needed may be considered a good exercise for the more complex computations for maximum-mass transfers.

The independent variables are $\vec{X}_f, t_f$ and $\vec{\lambda}_0$. The initial guess for these quantities is obtained by propagating a number of randomly chosen initial conditions up to a stopping condition. The cost function is checked at every integration step of every propagation until a value has been found that is sufficiently good. This last propagation then provides the initial guess values for each of the independent variables. This process of initial guess generation is repeatedly applied as an attempt at reasonably sampling the space of local minima. More sophisticated initial guess generation methods are left for future research.

In this example, the initial halo orbit reaches minimum and maximum z-coordinate values of $(-16601, 4231)$ km, while the target halo extends in z from approximately -10959 km to about 6643 km. The spacecraft and engine parameters are $m_0 = 25,000$ kg, $P_{max} = 1.2$ MW, $\eta = 1$ (an ideal engine), $I_{sp} = 7365$ s, yielding a thrust of 33.2 N and an initial thrust acceleration of 1.3 mm/s^2. The thrust acceleration is such that at least approximately two revolutions are needed to complete the transfer. Such high thrust levels are used to avoid having to perform more revolutions, which would necessitate a multiple shooting approach to handle the sensitivity of the trajectory to the initial conditions.

An optimal trajectory is shown in Fig. 1, where the short line segments represent the thrust direction. The final mass was $24,874$ kg. The mass and thrust profiles are shown in Fig. 2.

IV.B. VSI halo-to-halo transfer

A phase-free, VSI transfer is sought between two L1 halos in the Sun-Earth system ($\mu = 3.0 \times 10^{-6}$), from the larger to the smaller. The cost function used is the same as in the CEV constant-thrust case. Other cases exploiting the flexibility of the functional form of J (Eq. 5) were also computed, but not presented

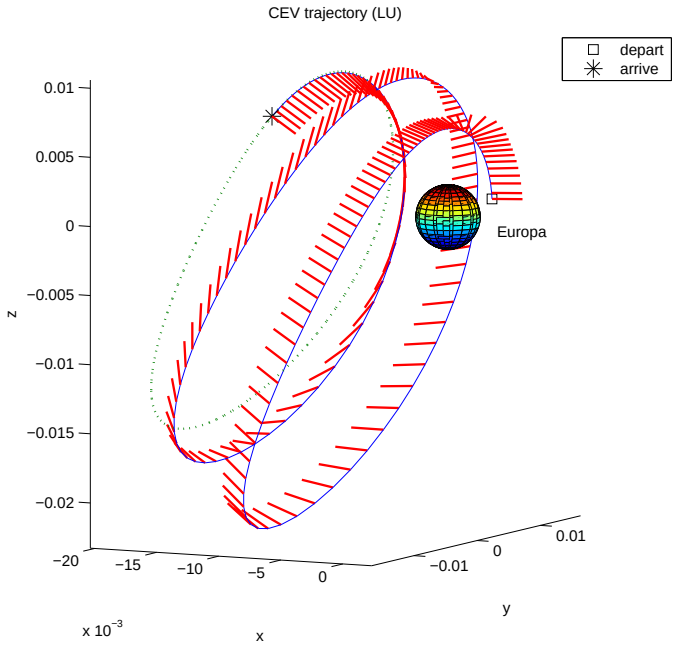

Fig. 1. CEV minimum-flight-time, constant-thrust, halo-to-halo transfer in the Jupiter-Europa CR3BP.

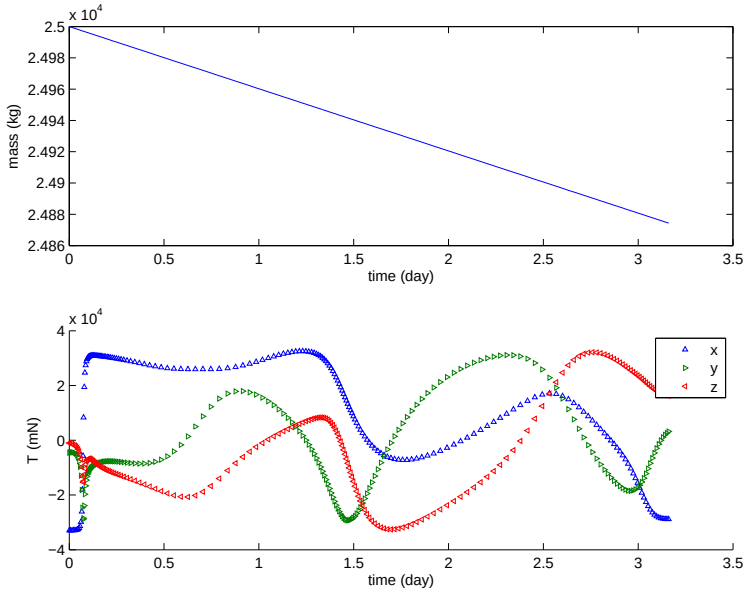

Fig. 2. Mass and thrust profiles for a CEV minimum-flight-time, constant-thrust, halo-to-halo transfer in the Jupiter-Europa CR3BP.

here, wherein terms were added under the integral to encourage the I_{sp} to be close to either of two specific values, thereby modelling approximately an engine with two I_{sp} levels. The initial guesses to the trust region optimiser are obtained analogously to the CEV case.

The initial and target halos have extremal z-coordinate values of approximately $(-1.786, 0.562)$ $\times 10^6$ km and $(-0.808, 0.568)$ $\times 10^6$ km, respectively. The spacecraft and engine parameters are $m_0 = 2000$ kg, $P_{max} = 10$ kW, $\eta = 1$ (an ideal engine). The resulting thrust acceleration is such that approximately two revolutions are needed to complete the transfer.

An optimal trajectory is shown in Fig. 3, where the short, red line segments represent the thrust direction, and the length of the line segments is proportional to the thrust level. The final mass was 1998.16 kg. The very small amount of propellant consumed (under 2 kg) indicates not only the advantages of continuous-thrust, high-I_{sp} propulsion, but also the unstable nature of these neighbouring periodic orbits. The I_{sp}, thrust acceleration and thrust profiles are shown in Fig. 4.

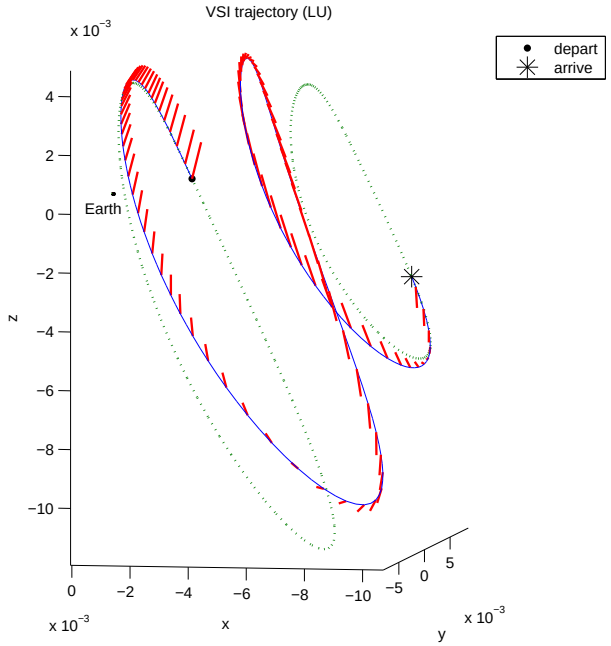

Fig. 3. Trajectory plot for a VSI maximum-final-mass, halo-to-halo transfer in the Sun-Earth CR3BP.

IV.C. CEV, bang-bang DRO-to-DRO transfer

A bang-bang transfer is sought between two DROs, starting from the positive x-axis on the larger and ending anywhere on the smaller. Apart from assuming a fixed flight time, the initial guess generation process is similar to that used in the CEV halo-halo case.

The initial and target DROs have extremal x-coordinate values of approximately $(-49982, 50116)$ km and $(-19926, 20060)$ km, respectively. The spacecraft and engine parameters are $m_0 = 25,000$ kg, $I_{sp} = 7365$ s, $P_{max} = 180$ kW, $\eta = 1$ (an ideal engine). The resulting thrust acceleration is such that approximately eleven revolutions are needed to complete the transfer with the allowance of reasonably sized coast arcs.

An optimal trajectory is shown in Fig. 5, where the short, red line segments represent the thrust direction, and the length of the line segments is proportional to the thrust level. The final mass was $24,885$ kg. The thrust acceleration and thrust profiles are shown in Fig. 6.

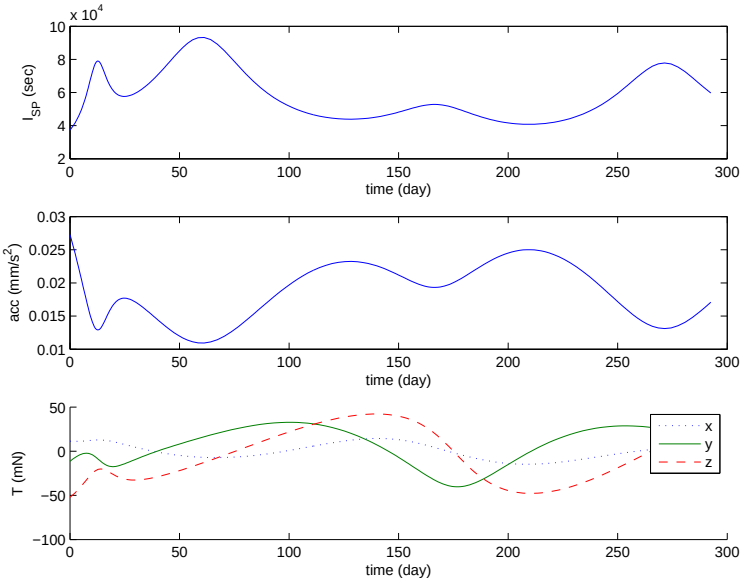

Fig. 4. Engine profile for a VSI maximum-final-mass, halo-to-halo transfer in the Sun-Earth CR3BP.

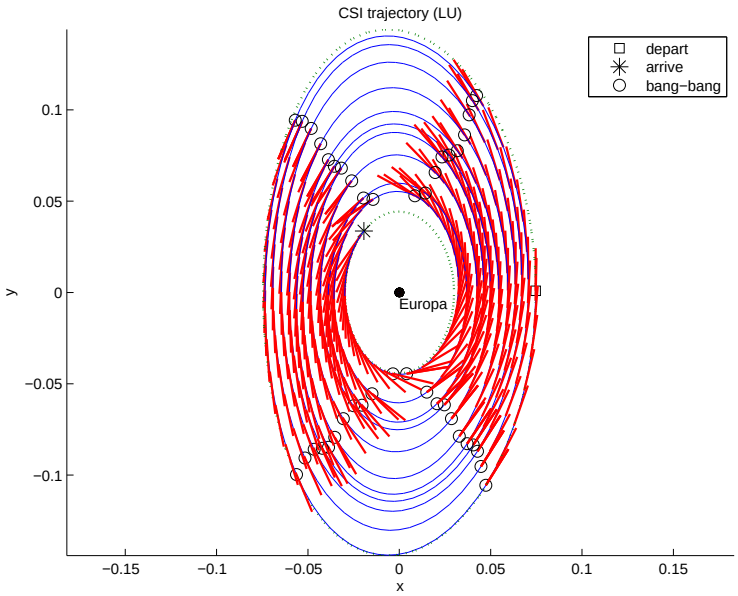

Fig. 5. Trajectory plot for a CEV, bang-bang maximum-final-mass, DRO-to-DRO transfer in the Jupiter-Europa CR3BP.

V. Conclusions

We have presented a novel approach to the design of optimal low-thrust orbit transfers and applied the approach to transfers in the circular restricted three-body problem. The approach combines optimal control theory, differential dynamic programming, penalty methods and second-order trust-region parameter optimisation techniques. Each of these components offers a distinct advantage in the overall approach.

The optimal control theory and the penalty methods allow the problem to be posed as a small-scale parameter optimisation problem (unlike the traditional direct methods, where the problem typically has a much larger scale, and unlike traditional indirect methods where a two-point boundary value problem

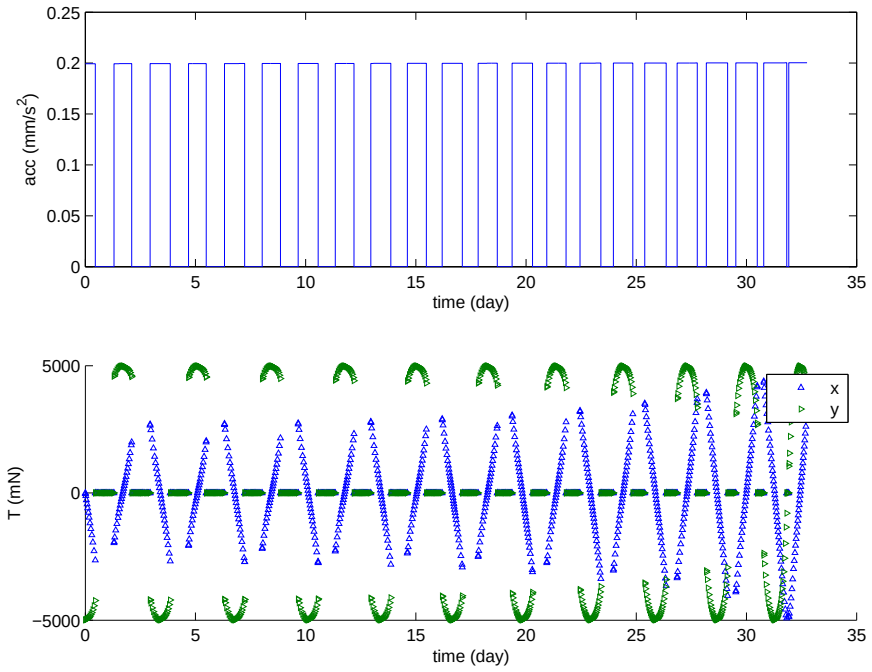

Fig. 6. Thrust profile for a CEV, bang-bang maximum-final-mass, DRO-to-DRO transfer in the Jupiter-Europa CR3BP.

must be solved). Furthermore, by carefully obtaining analytic derivatives across the switching times in the bang-bang optimal control problem, the need for including the numerous switching times in the vector of unknowns is eliminated. (However, there is the drawback of discontinuity when the thrust structure, i.e. number of switch times, changes. This can sometimes prevent the thrust structure from changing, or can lead to chatter between two thrust structures.) The differential dynamic programming allows the number of equations that must be numerically propagated to be drastically reduced (compared to state transition methods). The second-order trust region method offers good and easily tunable convergence characteristics.

Altogether, the resulting small scale of the problem and the speed of convergence make the present approach well-suited to global searches, and hence a valuable tool in mission design of low-thrust trajectories.

Acknowledgements

We are grateful for the support of the Research and Technology Development Program at the Jet Propulsion Laboratory. This research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

References

[1]Lawden, D. F., *Optimal Trajectories for Space Navigation,* Butterworth & Co., London, 1963. pp. 5, 59.

[2]Russell, R.P., "Primer Vector Theory Applied to Global Low-Thrust Trade Studies," *Journal of Guidance, Control, and Dynamics,* Vol. 30, No. 2, 2007, pp. 460-472.

[3]Conn, A.R., Gould, N. I. M., Toint, P. L., *Trust Region Methods,* SIAM, Philadelphia, 2000.

[4]Whiffen, G. J., Sims, J. A., "Application of the SDC optimal control algorithm to low-thrust escape and capture trajectory optimization," Paper AAS 02-208, AAS/AIAA Space Flight Mechanics Meeting, San Antonio, Texas, 2002.

[5]Jacobson, D.H., Mayne, D.Q., *Differential Dynamic Programming,* American Elsevier Publishing Co. Inc., New York, 1970, Ch. 2.