Optimal Timing of Control-Law Updates for Unstable Systems with Continuous Control

Eric D. Gustafson*
University of Michigan, Ann Arbor, Michigan 48109-2140
and
Daniel J. Scheeres†
University of Colorado, Boulder, Colorado 80309-0429

The optimal control of a linear system is studied relative to a periodic unstable trajectory using continuous control. Gaussian state uncertainties induce a statistical cost of controlling the state over a long period of time. The length of time between control-law updates directly impacts this cost, and in a hyperbolically unstable system, the time between control updates can take an optimal value. If the amount of uncertainty is fixed, there is an optimal distribution between position and velocity uncertainty. We apply these ideas to study the statistical cost of controlling a spacecraft in the vicinity of a relative equilibrium point and a Halo orbit in the Hill three-body problem.

Nomenclature

$\mathbf{A}(t)$	=	linear system dynamics matrix
a	=	acceleration component
$\mathbf{B}$	=	linear system control matrix
$E[\cdot]$	=	expectation operator
G	=	universal gravitational constant, 6.673×10^{-11} m$^3 \cdot$ kg$^{-1} \cdot$ s^{-2}
J	=	cost function
J^*	=	optimal cost function
l	=	reference length scale for nondimensionalizing
M	=	mass of secondary body in the Hill three-body problem
$\mathbf{m}$	=	mean vector
$\mathbf{P}$	=	covariance matrix
r	=	distance from origin in the Hill three-body problem
$\mathbf{r}$	=	position vector
T	=	halo orbit period
T_u	=	control horizon
T_u^*	=	optimal control horizon
t_f	=	control-law planning horizon
$\mathbf{u}$	=	control vector
$\mathbf{v}$	=	velocity vector
var$[\cdot]$	=	variance operator
$\mathbf{x}$	=	state vector
x, y, z	=	positions of spacecraft in the rotating frame relative to the secondary body in the Hill three-body problem
α	=	Lyapunov characteristic exponent
$\boldsymbol{\lambda}$	=	costate vector
σ^2	=	variance/covariance value
σ_r, λ	=	nondimensional uncertainty parameters
τ	=	reference time scale for nondimensionalizing
$\boldsymbol{\Phi}(t_1, t_0)$	=	state transition matrix mapping linearized state from time t_0 to time t_1
ω	=	angular rotation rate of the secondary body about the primary in the Hill three-body problem

I. Introduction

IN THIS paper, we describe a method to analyze the average (or ensemble) cost of optimal control near a periodic unstable trajectory. Specifically, we focus on control of the time-varying linear system resulting from linearizing the full dynamics about a nominal periodic trajectory. We consider a specific control strategy to take into account the finite horizon of the continuous control and uncertainty in the estimate of the state. This analysis is of direct application to the determination of mission operations for halo orbiters and for the budgeting of statistical fuel costs, especially for spacecraft in the highly unstable Earth–moon system.

Previously, Renault and Scheeres [1] conducted a similar study of optimal statistical control that considered the placement of impulsive control maneuvers near an unstable equilibrium point. The results of this paper serve to reinforce key results by Renault and Scheeres, such as the correlation between optimal control maneuver timing and the characteristic time of the instability of a system. Also, trends derived in Scheeres's previous work [2] concerning the qualitative impact of the update time on control cost using impulsive maneuvers are developed here for the continuous-control case and are shown to be similar.

The control force and system dynamics are assumed to be deterministic, and the state estimates are assumed to have a Gaussian probability distribution. The Gaussian assumption is justified for this analysis, because spacecraft uncertainties are almost always reported according to Gaussian statistics; more detailed information is not typically available. Additionally, this allows for the availability of analytical control laws for linear systems with state uncertainty. Further analysis could be performed under non-Gaussian statistics; however, this would stray from the realm of practicality. All system uncertainty is assumed to be adequately described by uncertainties in the estimation of the state. In addition, note that we almost exclusively focus on optimally updating control laws in the presence of uncertainty, not optimal control per se.

The motivation for these control-law updates is as follows:

1) At some time t_1, we have an estimate of the state with a certain uncertainty level. The uncertainties at this point can be viewed as the steady-state uncertainties of an estimation process.

2) Based on the estimate of the state at t_1, we choose and implement a controller that would nominally cause the state to converge to the target state at time t_2 (in the absence of uncertainty).

3) At time t_2, error exists again due to uncertainties at time t_1.

This process repeats and therefore there is a statistical cost associated with the steady-state control. It is also similar to the actual process used in spacecraft trajectory control [2]. To estimate this statistical cost, we evaluate the expected cost of the control from time interval t_i to t_{i+1} due to propagated uncertainties from interval t_{i-1} to t_i. During each interval, the control force is continuous; however, at the boundary between each interval, a discontinuity results from the choice of a new optimal control for the next control period.

To minimize the cost of regulating the system, we seek to minimize the average cost over time. To achieve this, one must find the time between updates that minimizes the expected cost per segment divided by the time between updates: $E[J]/T_u$, where J is the cost incurred and T_u is the time between updates. That is, the optimal time between updates, T_u^*, is given by

$$T_u^* = \arg \min_{T_u} \frac{E[J]}{T_u} \qquad (1)$$

The time between updates is assumed to be a constant parameter over the analysis period of interest for this study. This choice fits with our desire to develop a steady-state control and is particularly appropriate considering our analysis of control in the vicinity of a periodic system, which naturally lends itself to a repeating control strategy. A topic of future investigation could be to formulate the problem in a slightly more general way: for example, finding the best way to distribute n control-law updates over m periods.

We split the control problem into two pieces:

1) Determine the optimal control to target back to a nominal trajectory/state in a finite time.

2) Determine the effect of state uncertainty on the nominal control and how we can decrease the overall cost in the presence of uncertainty.

In problem 1, the control time T_u is a free parameter, and in the absence of noise, cost is reduced by taking $T_u \to \infty$, even for unstable systems.

For problem 2, in which we do not know the precise initial state, we find that the error can have a catastrophic penalty if our dynamical system is unstable. Hence, this injects a specific structure or natural time scale into our control problem. This optimal update time is nominally related to the characteristic time scale of the instability. Although the control update time cannot be solved in closed form, even for simple linear systems, the natural dynamics still provide more insight than numerical optimization. The simple control law elucidates the practical relevance of the unstable characteristic time when estimating the optimal update time, as will be discussed.

This combined control and measurement strategy is a periodic update procedure in which the optimal control problem is solved using a finite-horizon time span equal to the time between control updates. This can be viewed as an extreme case of receding horizon control (RHC) or model predictive control (MPC), in which the execution horizon is equal to the planning horizon. In RHC, the execution horizon is typically much shorter than the planning horizon [3,4]; however, spacecraft state estimates are made using data from ground-based radar tracking stations, which perform measurements infrequently compared with typical RHC applications. This necessitates a relatively long execution horizon. Extending the planning horizon does allow for a lower expected cost; however, this also increases the steady-state uncertainty, as shown later. For simplicity, we will assume that the control update time equals the planning horizon. One could relax this assumption by modifying the cost function to be a weighted combination of fuel-based cost as well as some uncertainty cost, then solving for optimal values of both the update time and the horizon time. To further clarify, the typical application of RHC is to approximate a feedback control law, which is not the goal here. Instead, we are interested in optimizing the time between control-law updates, which is the key parameter to our overall optimization process.

This paper is split into two main sections. The first section reviews optimal design of statistical correction maneuvers and the second applies these results to the optimal statistical control of a libration point orbiter.

II. Optimal Control-Law Update Timing

A. Deterministic Optimal Control Law

The system resulting from the linearization near a desired trajectory is a linear time-varying system, written as

$$\dot{\mathbf{x}} = \mathbf{A}(t)\mathbf{x} + \mathbf{B}\mathbf{u}$$

Our control objective is to minimize the energy expended over a finite horizon,

$$J^*(t_0, t_f) = \min_{\mathbf{u}} \frac{1}{2} \int_{t_0}^{t_f} \mathbf{u} \cdot \mathbf{u}\, dt \qquad (2)$$

subject to the boundary conditions of a given initial state and a hard terminal constraint:

$$\mathbf{x}(t_0) = \mathbf{x}_0, \quad \text{given } \mathbf{x}(t_f) = \mathbf{0}$$

Our method of choice to integrate this system numerically was the sweep method [5], which is an efficient and robust method for linear systems with terminal constraints. The sweep method for this optimization problem involves the numerical integration of two linear matrix differential equations:

$$\dot{\mathbf{R}}(t; t_0, t_f) = -\mathbf{A}^T \mathbf{R}(t; t_0, t_f)$$

$$\dot{\mathbf{Q}}(t; t_0, t_f) = \mathbf{R}^T(t; t_0, t_f)\mathbf{B}\mathbf{B}^T \mathbf{R}(t; t_0, t_f)$$

subject to the terminal constraints

$$\mathbf{R}(t_f; t_0, t_f) = \mathbf{I} \qquad \mathbf{Q}(t_f; t_0, t_f) = \mathbf{0}$$

Once these equations are integrated backward in time, the costate vector as a function of the initial state is given by

$$\boldsymbol{\lambda}(t; t_0, t_f) = -\mathbf{R}(t; t_0, t_f)\mathbf{Q}^{-1}(t_0; t_0, t_f)\mathbf{R}^T(t_0; t_0, t_f)\mathbf{x}(t_0)$$

From this, the control is defined as

$$\mathbf{u}(t; t_0, t_f) = -\mathbf{B}^T \boldsymbol{\lambda}(t; t_0, t_f)$$
$$= -\mathbf{B}^T \mathbf{R}(t; t_0, t_f)\mathbf{Q}^{-1}(t_0; t_0, t_f)\mathbf{R}^T(t_0; t_0, t_f)\mathbf{x}(t_0)$$

If we define the matrix L as

$$\mathbf{L}(t; t_0, t_f) = -\mathbf{B}^T \mathbf{R}(t; t_0, t_f)\mathbf{Q}^{-1}(t_0; t_0, t_f)\mathbf{R}^T(t_0; t_0, t_f)$$

and $\mathbf{G}$ as

$$\mathbf{G}(t_0, t_f) = \int_{t_0}^{t_f} \mathbf{L}^T(\tau; t_0, t_f)\mathbf{L}(\tau; t_0, t_f)\, d\tau$$

then the form of the optimal control and the optimal cost function are

$$\mathbf{u}^*(t; t_0, t_f) = \mathbf{L}(t; t_0, t_f)\mathbf{x}(t_0) \qquad (3)$$

$$J^*(t_0, t_f) = \frac{1}{2}\mathbf{x}_0^T \mathbf{G}(t_0, t_f)\mathbf{x}_0 = \frac{1}{2}\mathrm{tr}[\mathbf{G}(t_0, t_f)\mathbf{x}_0\mathbf{x}_0^T] \qquad (4)$$

This form of J^* provides for straightforward computation of the expected value of the cost function, because the statistics of $\mathbf{x}_0$ are assumed to be given. This linear method can also be extended to a nonlinear method by using the generating function approach for optimal control [6].

B. Control-Law Updates

We start by studying three simple linear systems to demonstrate the control and measurement strategy as well as the derivation of the expected cost and optimal time between measurements. In our analysis of the timing of the control-law updates, the state is assumed to be a Gaussian random vector, with the mean and covariance taken

as outputs of an estimation process. The multivariate Gaussian probability distribution function for a vector $\mathbf{x} \in \mathbb{R}^n$ with mean $\mathbf{m} \in \mathbb{R}^n$ and covariance matrix $\mathbf{P} \in \mathbb{R}^{n \times n}$ is defined as [7]

$$p(\mathbf{x}) = \frac{1}{\sqrt{(2\pi)^n \det \mathbf{P}}} \exp\left(-\frac{1}{2}(\mathbf{x} - \mathbf{m})^T \mathbf{P}^{-1}(\mathbf{x} - \mathbf{m})\right)$$

The expected value of a function is

$$E[f(\mathbf{x})] = \int_\infty f(\boldsymbol{\xi}) p(\boldsymbol{\xi}) \mathrm{d}\boldsymbol{\xi}$$

and the mean $\mathbf{m}$ and covariance $\mathbf{P}$ of a random variable are then

$$\mathbf{m} = E[\mathbf{x}] = \int_\infty \boldsymbol{\xi} p(\boldsymbol{\xi}) \mathrm{d}\boldsymbol{\xi} \qquad \mathbf{P} = E[\mathbf{x}\mathbf{x}^T] - \mathbf{m}\mathbf{m}^T$$

The concepts involved in this analysis will be explained in the context of three illustrative systems. The choice of the three simple systems is motivated by previous results [2] for impulsive control. Consider the following 2-dimensional system with one scalar input:

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}u$$

The three different cases result from three different $\mathbf{A}$ matrices.
Oscillatory case:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 \\ -\beta^2 & 0 \end{bmatrix}$$

Double-integrator case:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix}$$

Hyperbolically unstable case:

$$\mathbf{A} = \begin{bmatrix} 0 & 1 \\ +\beta^2 & 0 \end{bmatrix}$$

For all three systems, $\mathbf{B} = [0 \quad 1]^T$. The cost function for the optimal control is

$$J = \frac{1}{2}\int_{t_0}^{t_f} u(t)^2 \mathrm{d}t$$

and the optimization problem to be solved is to find $u(t)$ such that $u(t)$ minimizes $E[J]$ subject to given initial and final states.

When the linear-Gaussian assumptions are followed, we can conclude that if optimal control is applied from time t_0 to time t_1, with the following initial conditions,

$$E[\mathbf{x}_0] = \mathbf{0} \tag{5}$$

$$\mathrm{var}\,[\mathbf{x}_0] = \mathbf{P}_0 \tag{6}$$

the expectation and covariance describing the state at time t_1 are given by

$$E[\mathbf{x}_1] = \mathbf{0} \tag{7}$$

$$\mathrm{var}\,[\mathbf{x}_1] = \mathbf{P}_1 = \mathbf{\Phi}(t_1 - t_0)\mathbf{P}_0\mathbf{\Phi}^T(t_1 - t_0) \tag{8}$$

For the preceding three systems, the state transition matrices and covariance matrices are as follows.
Oscillatory case:

$$\mathbf{\Phi}(t, 0) = \begin{bmatrix} \cos(\beta t) & \frac{1}{\beta}\sin(\beta t) \\ -\beta \sin(\beta t) & \cos(\beta t) \end{bmatrix}$$

$P_{11}(t)$
$$= \frac{\beta^2(1 + \cos(2\beta t))P_{11} + 2\beta \sin(2\beta t)P_{12} + (1 - \cos(2\beta t))P_{22}}{2\beta^2}$$

$P_{12}(t) = P_{21}(t)$
$$= \frac{-\beta^2 \sin(2\beta t)P_{11} + 2\beta \cos(2\beta t)P_{12} + \sin(2\beta t)P_{22}}{2\beta}$$

$P_{22}(t)$
$$= -\frac{\beta^2(1 + \cos(2\beta t))P_{11} + 2\beta \sin(2\beta t)P_{12} + (1 - \cos(2\beta t))P_{22}}{2}$$

Double-integrator case:

$$\mathbf{\Phi}(t, 0) = \begin{bmatrix} 1 & t \\ 0 & 1 \end{bmatrix}$$

$$P_{11}(t) = P_{11} + 2P_{12}t + P_{22}t^2 \qquad P_{12}(t) = P_{21}(t) = P_{12} + P_{22}t$$

$$P_{22}(t) = P_{22}$$

Hyperbolically unstable case:

$$\mathbf{\Phi}(t, 0) = \begin{bmatrix} \cosh(\beta t) & \frac{1}{\beta}\sinh(\beta t) \\ \beta \sinh(\beta t) & \cosh(\beta t) \end{bmatrix}$$

$P_{11}(t)$
$$= \frac{\beta^2(1 + \cosh(2\beta t))P_{11} + 2\beta \sinh(2\beta t)P_{12} + (\cosh(2\beta t) - 1)P_{22}}{2\beta^2}$$

$P_{12}(t) = P_{21}(t)$
$$= \frac{-\beta^2 \sinh(2\beta t)P_{11} + 2\beta \cosh(2\beta t)P_{12} + \sinh(2\beta t)P_{22}}{2\beta}$$

$P_{22}(t)$
$$= \frac{\beta^2(\cosh(2\beta t) - 1)P_{11} + 2\beta \sinh(2\beta t)P_{12} + (1 + \cosh(2\beta t))P_{22}}{2}$$

C. Statistical Cost

Let the state to be the random vector $\mathbf{X}(t)$. At time t $(t_0 < t < t_f)$, the state is given by

$$\mathbf{X}(t) = \mathbf{\Phi}(t, t_0)\mathbf{X}(t_0) + \mathbf{F}(t; t_0, t_f)\mathbf{X}(t_0)$$

where

$$\mathbf{F}(t; t_0, t_f) = \int_{t_0}^{t} \mathbf{\Phi}(t, \tau)\mathbf{B}\mathbf{L}(\tau; t_0, t_f)\mathrm{d}\tau$$

Suppose that after a long period of time, we are at the end of an update period just before we take a measurement. Our random state vector should have some steady-state distribution with zero mean and some covariance $\mathbf{P}_{ss}$:

$$\mathbf{P}_{ss} = E[\mathbf{X}(T_u)\mathbf{X}^T(T_u)]$$

We can simulate taking a new measurement by adding a vector taken from this steady-state distribution with a vector that is due to measurement error having covariance $\mathbf{P}_m$. So at the beginning of the next control period, we have

$$\mathbf{X}(t_{0_+}) = \mathbf{M} + \mathbf{V}$$

where $\mathbf{M} \sim \mathcal{N}(\mathbf{0}, \mathbf{P}_{ss})$, $\mathbf{V} \sim \mathcal{N}(\mathbf{0}, \mathbf{P}_m)$, $\mathbf{M}$ and $\mathbf{V}$ are independent, and the $+$ subscript on t_0 indicates an instant just after a

Table 1 Summary of coefficients

	Double-integrator	Oscillatory	Hyperbolically Unstable
J_Δ	$1/T$	$\frac{\beta^2}{\cos(2\beta T)+2\beta^2T^2-1}$	$\frac{\beta^2}{\cosh(2\beta T)-2\beta^2T^2-1}$
J_{rr}	$6/T^2$	$\beta\sin(2\beta T)+2\beta^2T$	$\beta\sinh(2\beta T)+2\beta^2T$
J_{rv}	$3/T$	$1-\cos(2\beta T)$	$\cosh(2\beta T)-1$
J_{vv}	2	$2T-\frac{1}{\beta}\sin(2\beta T)$	$\frac{1}{\beta}\sinh(2\beta T)-2T$

measurement. At the update time,

$$\mathbf{X}(T_u) = \mathbf{\Phi}(T_u, t_0)\mathbf{X}(t_{0_+}) + \mathbf{F}(T_u; t_0, t_f)\mathbf{M}$$
$$= (\mathbf{\Phi}(T_u, t_0) + \mathbf{F}(T_u; t_0, t_f))\mathbf{M} + \mathbf{\Phi}(T_u, t_0)\mathbf{V}$$

For convenience, let $\mathbf{\Phi} = \mathbf{\Phi}(T_u, t_0)$ and $\mathbf{F} = \mathbf{F}(T_u; t_0, t_f)$. The covariance is

$$E[\mathbf{X}(T_u)\mathbf{X}^T(T_u)] = (\mathbf{\Phi} + \mathbf{F})E[\mathbf{M}\mathbf{M}^T](\mathbf{\Phi} + \mathbf{F})^T$$
$$+ \mathbf{\Phi}E[\mathbf{V}\mathbf{V}^T]\mathbf{\Phi}^T\mathbf{P}_{ss} = (\mathbf{\Phi} + \mathbf{F})\mathbf{P}_{ss}(\mathbf{\Phi} + \mathbf{F})^T + \mathbf{\Phi}\mathbf{P}_m\mathbf{\Phi}^T$$

A straightforward solution for $\mathbf{P}_{ss}$ maybe be written with the vectorization (or stacking) operator and the Kronecker product:

$$\text{vec}(\mathbf{P}_{ss}) = \text{vec}((\mathbf{\Phi} + \mathbf{F})\mathbf{P}_{ss}(\mathbf{\Phi} + \mathbf{F})^T) + \text{vec}(\mathbf{\Phi}\mathbf{P}_m\mathbf{\Phi}^T)$$
$$= [(\mathbf{\Phi} + \mathbf{F}) \otimes (\mathbf{\Phi} + \mathbf{F})]\text{vec}(\mathbf{P}_{ss}) + \text{vec}(\mathbf{\Phi}\mathbf{P}_m\mathbf{\Phi}^T)$$
$$\text{vec}(\mathbf{P}_{ss}) = [\mathbf{I} - (\mathbf{\Phi} + \mathbf{F}) \otimes (\mathbf{\Phi} + \mathbf{F})]^{-1}\text{vec}(\mathbf{\Phi}\mathbf{P}_m\mathbf{\Phi}^T)$$

$\mathbf{P}_{ss}$ is obtained by unstacking $\text{vec}(\mathbf{P}_{ss})$. If $T_u = t_f$, then $\mathbf{F} = -\mathbf{\Phi}$ and the preceding equation simplifies to

$$\mathbf{P}_{ss} = \mathbf{\Phi}\mathbf{P}_m\mathbf{\Phi}^T$$

The relevant covariance matrix for calculating the expected cost is the covariance just after taking a measurement $\mathbf{P}_+$:

$$\mathbf{P}_+ = E[\mathbf{X}(t_{0_+})\mathbf{X}^T(t_{0_+})] = \mathbf{P}_{ss} + \mathbf{P}_m$$

Continuing with these examples, we may partition the initial state for convenience. Let r_0 be the initial position and v_0 be the initial velocity. Then $\mathbf{x}_0 = [r_0 \quad v_0]^T$. Similarly, partition $\mathbf{G}$ from Eq. (4) as

$$\mathbf{G} = 2J_\Delta \begin{bmatrix} J_r & J_{rv} \\ J_{rv} & J_v \end{bmatrix}$$

where the factor J_Δ is pulled out for convenience. This yields the following expression for the deterministic cost:

$$J = J_\Delta(J_r r_0^2 + 2J_{rv}r_0v_0 + J_v v_0^2) \tag{9}$$

where J_Δ, J_r, J_{rv}, and J_v are functions of the linear dynamics and the update time, given in Table 1.

Once the deterministic cost is known, the expected value and variance of the cost can be computed. Taking the expectation $E[\cdot]$ of Eq. (4) yields

$$E[J] = \tfrac{1}{2}\text{tr}\{\mathbf{G}(\mathbf{P}_+ + E[\mathbf{x}_0]E[\mathbf{x}_0^T])\} \tag{10}$$

or, in the form of Equation (9),

$$E[J] = J_\Delta(J_r E[r_0^2] + 2J_{rv}E[r_0v_0] + J_v E[v_0^2]) \tag{11}$$

$$E[J] = J_\Delta(J_r(\sigma_r^2(t_0) + E[r_0]^2) + 2J_{rv}(\sigma_{rv}^2(t_0) + E[r_0]^T E[v_0])$$
$$+ J_v(\sigma_v^2(t_0) + E[v_0]^2)) \tag{12}$$

where σ denotes the (co)variance of the given quantity at the beginning of the update interval, just after a measurement.

Because the control law was chosen so that the expected value of the state is the zero vector ($E[\mathbf{r}_0] = E[\mathbf{v}_0] = \mathbf{0}$), Eqs. (10) and (12) simplify to

$$E[J] = \tfrac{1}{2}\text{tr}\{\mathbf{G}\mathbf{P}_+\} \tag{13}$$

$$E[J] = J_\Delta(J_r\sigma_r^2(t_0) + 2J_{rv}\sigma_{rv}^2(t_0) + J_v\sigma_v^2(t_0)) \tag{14}$$

Note that this expected cost is only a function of the time between updates, T_u, and the initial covariances $\mathbf{P}_m$, because for a given system, $\mathbf{G}$ is determined completely by T_u and $\mathbf{P}_+$ is determined by $\mathbf{P}_m$.

The variance of J, var$[J]$, is given by

$$\text{var}[J] = E[J^2] - (E[J])^2 \tag{15}$$

Computing $E[J^2]$ using the Gaussian joint characteristic function yields

$$E[J^2] = \tfrac{1}{2}\text{tr}\{(\mathbf{G}\mathbf{P}_+)^2\} + \tfrac{1}{4}\text{tr}^2(\mathbf{G}\mathbf{P}_+) \tag{16}$$

$$E[J^2] = J_\Delta^2\{3J_r^2(\sigma_r^2)^2 + 4J_{rv}^2(2(\sigma_{rv}^2)^2 + \sigma_r^2\sigma_v^2) + 3J_v^2(\sigma_v^2)^2$$
$$+ 8J_r J_{rv}\sigma_r^2\sigma_{rv}^2 + 2J_r J_v(\sigma_r^2\sigma_v^2 + 2(\sigma_{rv}^2)^2) + 12J_{rv}J_v\sigma_{rv}^2\sigma_v^2\} \tag{17}$$

Substituting Eq. (14) and (17) into Eq. (15) yields the variance of J:

$$\text{var}[J] = \tfrac{1}{2}\text{tr}\{(\mathbf{G}\mathbf{P}_+)^2\} \tag{18}$$

$$\text{var}[J] = 2J_\Delta^2\{J_r^2(\sigma_r^2)^2 + 2J_{rv}^2((\sigma_{rv}^2)^2 + \sigma_v^2\sigma_r^2) + J_v^2(\sigma_v^2)^2$$
$$+ 4J_r J_{rv}\sigma_r^2\sigma_{rv}^2 + 2J_r J_v(\sigma_{rv}^2)^2 + 4J_{rv}J_v\sigma_{rv}^2\sigma_v^2\} \tag{19}$$

As with the expected value of J, the variance of J is only a function of T_u and $\mathbf{P}_m$.

D. Steady-State Minimum Expected Cost

Because of the complicated form of the expression for $E[J]/T_u$, even for simple time-invariant systems, Eq. (1) cannot typically be solved for T_u^* in closed form. Some statements can be made, however, about the behavior of T_u^*, depending on the dynamics of the linear systems under study. For any double-integrator and oscillatory-type dynamics, it can be shown that $E[J]/T_u$ achieves its minimum by letting $T_u \to \infty$ and that the actual value of $E[J]/T_u$ approaches zero. This is analogous to the impulsive control result previously obtained [2] and implies that maneuver execution errors dominate the uncertainty. The behavior is different for a hyperbolically unstable system. As T_u increases, the hyperbolic instability of the system causes $E[J]$ to grow exponentially and drives $E[J]/T_u$ toward ∞. For all three cases, $E[J]/T_u$ goes toward ∞ as T_u

Table 2 Summary of results for control about the H3BP equilibrium point for various systems

	$E[J]/T_u$, nondimensional	ΔV/period, km/s/period	Optimal update time, s
Sun–Earth	1.79×10^{-7}	6.21×10^{-4}	2.71×10^6
Earth–moon	2.88×10^{-5}	4.22×10^{-3}	2.01×10^5
Jupiter–Europa	5.72×10^{-4}	3.22×10^{-2}	2.61×10^4
Jupiter–Io	9.56×10^{-4}	6.46×10^{-2}	1.30×10^4
Saturn–Titan	3.90×10^{-5}	7.19×10^{-3}	1.17×10^5
Saturn–Enceladus	1.18×10^{-1}	8.34×10^{-2}	1.01×10^4

approaches zero. Thus, for the oscillatory and double-integrator cases, it is optimal to let the time between measurements go to infinity; however, for hyperbolically unstable dynamics, there exists an optimal time between measurements.

This can be shown more formally with the limits of $E[J]/T_u$ as $T_u \to 0$ and as $T_u \to \infty$. For all three example systems, $\lim_{T_u \to 0} E[J]/T_u$ is undefined; that is, $E[J]/T_u$ grows without bound as T_u approaches zero. For the oscillatory and double-integrator dynamics,

$$\lim_{T_u \to \infty} \frac{E[J]}{T_u} = 0$$

but for the hyperbolically unstable case, $\lim_{T_u \to \infty} E[J]/T_u$ is undefined ($E[J]/T_u$ also grows without bound as T_u tends toward infinity in this case). Because $E[J]/T_u$ is still finite for any finite value of T_u, this implies the existence of a minimum by continuity. Numerical studies of a wide range of systems show that the value of T_u^* is closely tied to the characteristic time of the unstable mode ($1/\alpha$), although it also depends on the initial values of the covariance matrix. This is also analogous to previously obtained results for impulsive control [1]. For an ideal 1-degree-of-freedom unstable system, it can be shown that the optimal update time for impulsive control equals the characteristic time [2]. For our continuous-control time-varying systems, the relationship is not exact, but numerical simulations support the extension as a rule of thumb. This relationship breaks down when applied to periodic trajectories that are too far from their initial origin, as shown in the example implementation. An example of the cost as a function of T_u for the three linear cases described previously is shown in Fig. 1. For all three cases, $P_{0_{11}} = P_{0_{22}} = 1$ and $P_{0_{12}} = 0$. For the hyperbolically unstable and oscillatory cases, $\beta = 0.1$.

The overall trends of the cost are very predictable as β changes. Figure 2 shows the expected cost per unit time in the unstable system with three different values of β, normalized on each axis so that the update time is scaled by β, and the cost is scaled such that the minimum value equals unity. Although β varies by 2 orders of magnitude, the normalized optimal values of T_u only change by about 10%. Because $E[J]$ scales linearly with $\mathbf{P}_m$, scaling $\mathbf{P}_m$ does not change the location of the optimal update time.

The effect of t_f on $E[J]/T_u$ in the unstable example is shown in Fig. 3. As the horizon is extended, a cost reduction can be obtained. However, as a practical matter, the reduction in cost must be weighed against the growth in $\|\mathbf{P}_{ss}\|$, as shown in Fig. 4. For the examples in this paper, we will assume $T_u = t_f$, as this is common for spacecraft control, due to the limited ability for communication of new control laws.
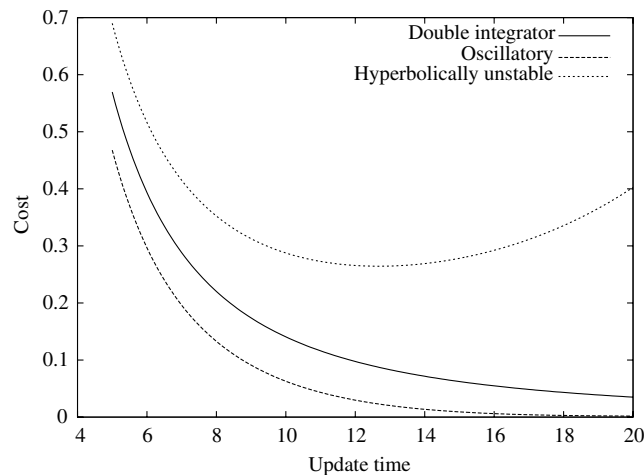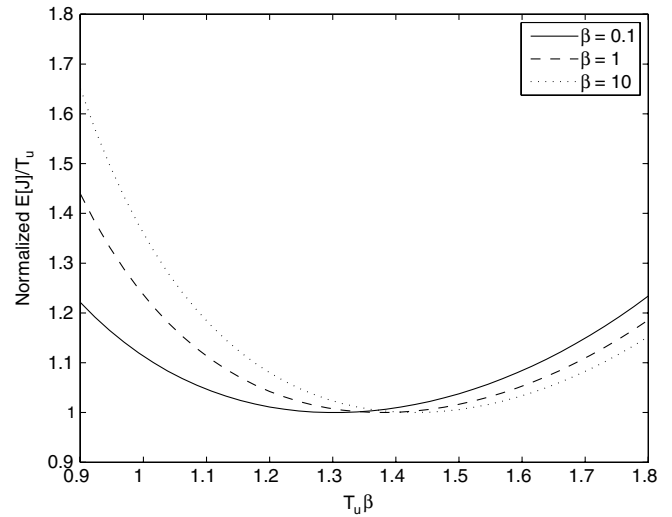
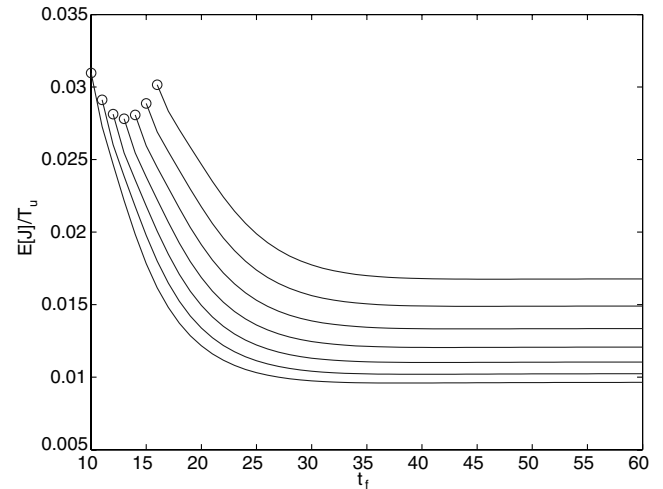Fig. 2 Example cost as a function of update time for three unstable example systems.

Fig. 3 $E[J]/T_u$ as t_f increases for a family of T_u values in the unstable example with $\beta = 1$. The circled points are where $T_u = t_f$ (t_f must be greater than T_u).

Fig. 1 Example cost as a function of update time for three example systems: double-integrator, oscillatory, and unstable.
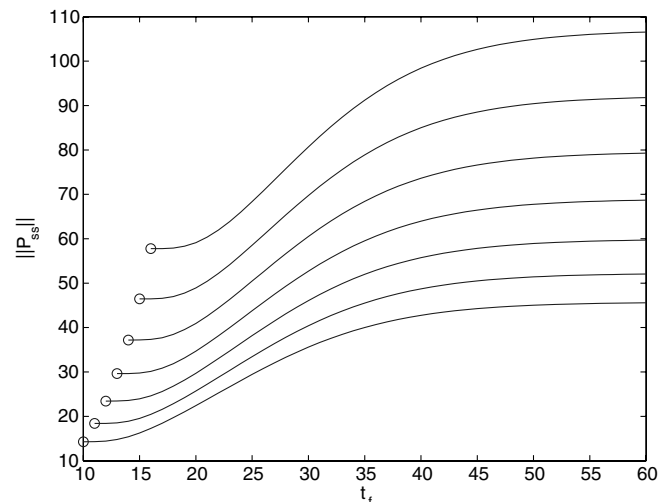
Fig. 4 $\|\mathbf{P}_{ss}\|$ as t_f increases for a family of T_u values in the unstable example with $\beta = 1$. The circled points are where $T_u = t_f$.

E. Estimates for Control Costs

In spacecraft control, one is often concerned with minimizing ΔV: the total velocity change required by the propulsion system to follow a given trajectory. For our continuous-control case, the ΔV between two time periods, t_0 and $t_0 + T_u$, is given by

$$\Delta V = \int_{t_0}^{t_0+T_u} \|\mathbf{u}(t)\| dt \tag{20}$$

$$\Delta V = \int_{t_0}^{t_0+T_u} \sqrt{\mathbf{u}(t) \cdot \mathbf{u}(t)} dt \tag{21}$$

The analysis methods in this paper use an energy cost function as in Eq. (2), although we may obtain an upper bound on the ΔV spent per unit time from our energy cost per unit time. A statement of the Cauchy–Schwarz inequality for two real functions, f and g, is as follows [8]:

$$\left[\int_a^b f(t)g(t) dt \right]^2 \leq \int_a^b [f(t)]^2 dt \int_a^b [g(t)]^2 dt \tag{22}$$

If we let $f(t) = \sqrt{\mathbf{u}(t) \cdot \mathbf{u}(t)}$, $g(t) = 1$, $a = t_0$, and $b = t_0 + T_u$, then we obtain

$$\int_{t_0}^{t_0+T_u} \sqrt{\mathbf{u}(t) \cdot \mathbf{u}(t)} dt \leq \sqrt{\int_{t_0}^{t_0+T_u} (\sqrt{\mathbf{u}(t) \cdot \mathbf{u}(t)})^2 dt \cdot \int_{t_0}^{t_0+T_u} dt} \tag{23}$$

$$\int_{t_0}^{t_0+T_u} \sqrt{\mathbf{u}(t) \cdot \mathbf{u}(t)} dt \leq \sqrt{\int_{t_0}^{t_0+T_u} \mathbf{u}(t) \cdot \mathbf{u}(t) dt \cdot T_u} \tag{24}$$

Therefore, a bound on the ΔV used in each control segment is

$$\Delta V \leq \sqrt{2T_u} \sqrt{J} \tag{25}$$

Taking the expectation of this yields

$$E[\Delta V] \leq \sqrt{2T_u} E[\sqrt{J}] \tag{26}$$

$$E[\Delta V] \leq \sqrt{2T_u} \sqrt{E[J]} \tag{27}$$

where the last inequality again comes from the Cauchy–Schwarz inequality. The expected ΔV spent per unit time is then bounded by

$$\frac{E[\Delta V]}{T_u} \leq \sqrt{2 \frac{E[J]}{T_u}} \tag{28}$$

III. Example Implementation

In this section, we will study two cases of spacecraft control in the Hill three-body problem (H3BP) using continuous thrust. In the first case, we limit ourselves to the planar motion of a spacecraft in the vicinity of one of the relative equilibrium points, and in the second, we study a spacecraft perturbed from a nominal halo orbit. A previous study of the equilibrium-point control problem [1] considered control using impulsive maneuvers. In addition, we show that the results obtained for the linear time-invariant case can be extended to linear time-varying systems.

The equations of motion for a spacecraft's position in the H3BP are [1]

$$\ddot{x} - 2\omega\dot{y} = -\frac{\mu}{r^3}x + 3\omega^2 x + a_x \tag{29}$$

$$\ddot{y} + 2\omega\dot{x} = -\frac{\mu}{r^3}y + a_y \tag{30}$$

$$\ddot{z} = -\frac{\mu}{r^3}y - \omega^2 z + a_z \tag{31}$$

where x, y, and z are the positions of the spacecraft in the rotating frame relative to the secondary body; a_x, a_y, and a_z are the spacecraft control accelerations; ω is the angular velocity of the secondary body about the primary; $\mu = GM$ (M is the mass of the secondary body); and r is the radius ($r = \sqrt{x^2 + y^2 + z^2}$). These equations may be nondimensionalized using the length scale $l = (\mu/\omega^2)^{1/3}$ and time scale $\tau = 1/\omega$. For the Earth–sun system, $\mu = 3.986 \times 10^5$ km^3/s^2, $\omega = 1.991 \times 10^{-7}$ rad/s, $l = 2.158 \times 10^6$ km, and $\tau = 5.023 \times 10^6$ s.

The dimensional covariance matrix associated with the state estimates is assumed to be a 6×6 diagonal matrix (typical of spacecraft state estimation) with entries P_r and P_v:

$$\mathbf{P_d} = \begin{bmatrix} P_r \cdot I_3 & 0_{3\times3} \\ 0_{3\times3} & P_v \cdot I_3 \end{bmatrix}$$

This covariance matrix may be nondimensionalized to obtain

$$\mathbf{P} = \begin{bmatrix} P_r/l^2 \cdot I_3 & 0_{3\times3} \\ 0_{3\times3} & P_v/(\omega l)^2 \cdot I_3 \end{bmatrix} = \frac{P_r}{l^2} \begin{bmatrix} I_3 & 0_{3\times3} \\ 0_{3\times3} & P_v/(P_r\omega^2) \cdot I_3 \end{bmatrix}$$

This may be parameterized to yield further insight into how the uncertainties affect the optimal update time and cost using the parameters $\sigma_r = \sqrt{P_r}/l$ and $\lambda = \omega\sqrt{P_r/P_v}$. This nondimensionalization and parameterization yield the following form for $\mathbf{P}_m$:

$$\mathbf{P}_m = \sigma_r^2 \begin{bmatrix} I_3 & 0_{3\times3} \\ 0_{3\times3} & 1/\lambda^2 \cdot I_3 \end{bmatrix} \tag{32}$$

Typical values of $P_r = (10 \text{ km})^2$ and $P_v = (10^{-6} \text{ km/s})^2$ relating to usual spacecraft uncertainties are used for the simulations, resulting in the nondimensional parameters $\sigma_r = 4.633 \times 10^{-6}$ and $\lambda = 1.991$ in the sun–Earth system and $\sigma_r = 1.13 \times 10^{-4}$ and $\lambda = 26.6$ in the Earth–moon system.

Given a nondimensional expected cost per unit time, $E[J]/T_u$, the dimensional cost per unit time is given as

$$\left(\frac{E[J]}{T_u} \right) \frac{l^2}{\tau^4}$$

To convert nondimensional $\Delta V/T_u$ values to their dimensional values, scale by l/τ^2 instead of l^2/τ^4. The results may also be reported as a cost per period of the secondary body, $2\pi/\omega$. The dimensional cost per secondary-body period is $(E[J]/T_u)2\pi l^2/\tau^3$ and the dimensional ΔV per secondary-body period is $(E[\Delta V]/T_u)2\pi l/\tau$.

For reference, a nondimensional expected value of J/T_u equal to 1×10^{-7} corresponds to an upper bound on the dimensional ΔV per period of 1.21×10^{-3} km/(s · period) in the sun–Earth system or 6.62×10^{-4} km/(s · period) in the Earth–moon system.

A. Planar Equilibrium-Point Control

When the system is nondimensionalized by setting $\mu = \omega = 1$ in Eqs. (29) and (30), the system has two equilibrium points using no control at $x = \pm 3^{-1/3}$ and $y = 0$. Linearizing about either of these points and defining the perturbed state $\delta\mathbf{x} = [\delta x \quad \delta y \quad \delta\dot{x} \quad \delta\dot{y}]^T$ yields the linear system

$$\delta\dot{\mathbf{x}} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 9 & 0 & 0 & 2 \\ 0 & -3 & -2 & 0 \end{bmatrix} \delta\mathbf{x} + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} a_x \\ a_y \end{bmatrix}$$

This system has an unstable mode, a stable mode, and an oscillatory mode, associated with the eigenvalues $+\sqrt{1 + 2\sqrt{7}} \approx 2.5$,

$-\sqrt{1 + 2\sqrt{7}} \approx -2.5$, and $\pm j\sqrt{2\sqrt{7} - 1} \approx \pm 2.1j$, respectively. The unstable mode's characteristic time is then $1/\sqrt{1 + 2\sqrt{7}} \approx 0.4$, leading us to expect the optimal update time to be approximately 0.4 time units.

In this example, the cost function, J, being minimized during each update interval, is the energy used:

$$J = \frac{1}{2}\int_{t_0}^{t_0 + T_u}(a_x^2 + a_y^2)\mathrm{d}t$$

The optimal control law, and hence the trajectories themselves, depend on the final time and are plotted in Fig. 5 for three different final times and various initial conditions. We will show that using an update time of 0.5 time units, corresponding to Fig. 5b, is optimal. Note this optimal update time of 0.5 is near the characteristic time of the unstable mode (0.4).

A plot of the expected cost as a function of update time is shown in Fig. 6, using the uncertainty parameters given previously. Because of the hyperbolically unstable dynamics, an optimal value of T_u clearly exists, which minimizes the expected cost.

These cost values may also be compared with similar studies. For the impulsive control strategy by Renault and Scheeres [1], the
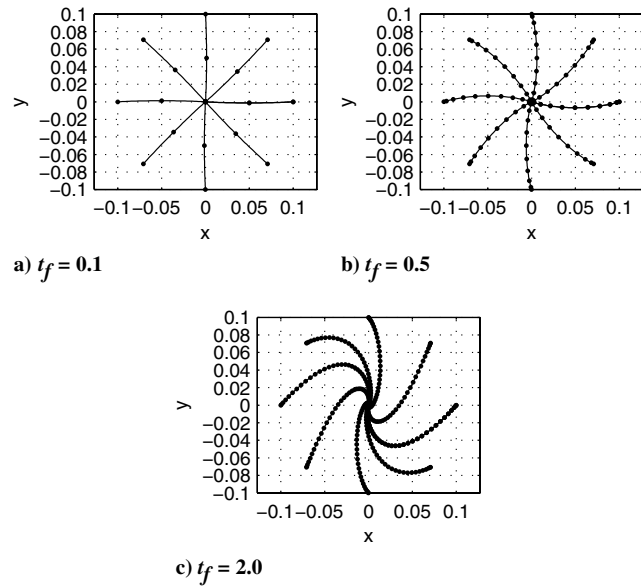
estimate for ΔV per period is 4.70×10^{-4} km/(s · period) in the sun–Earth system, although that assumes slightly larger uncertainties than those used in this study. When using their uncertainties, we obtain a ΔV per period of 6.21×10^{-4} km/(s · period) with our continuous-thrust method. The method has a wide variety of applications in the solar system, some of which are listed in Table 2 along with the corresponding optimal update times and estimated control costs.

Figure 7 shows the effect of the nondimensional parameter λ on the optimal update time for the H3BP using the parameters described previously. The variation in the optimal update time over the range of λ shown is about 1.75 days for the Earth–sun system. Note that the optimal update time does not depend on σ_r itself, only the ratio λ. For reference, a nondimensional time value of 0.5 corresponds to about 29 days in the Earth–sun H3BP and 2.2 days in the Earth–moon HR3BP. Interestingly, even though the formulations for the previous impulsive studies [1] are quite different from this continuous-control derivation, the results are consistent in that the ΔV computed from the continuous-thrust method is reasonably close to the impulsive ΔV and that the optimal update times for each method are close to the characteristic time of the unstable mode.

Figures 8 and 9 show the effect of λ on the value of the cost incurred over an update interval divided by the optimal update time: that is,

$$\min_{\mathbf{u},T_u} E[J(\mathbf{u}, T_u)]/T_u$$

a) $t_f = 0.1$ b) $t_f = 0.5$

c) $t_f = 2.0$

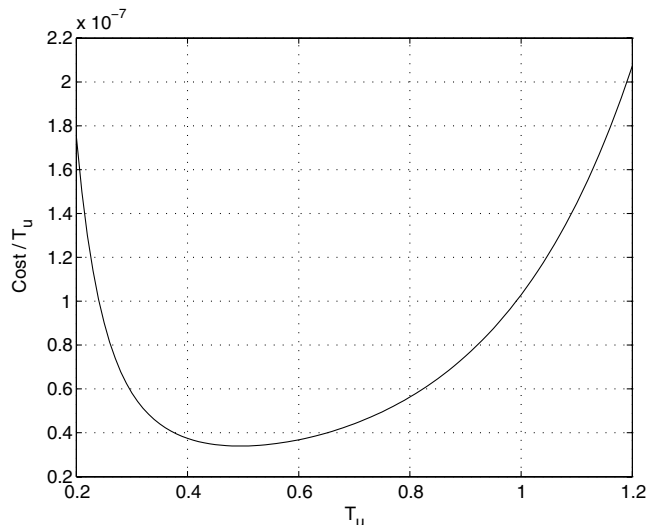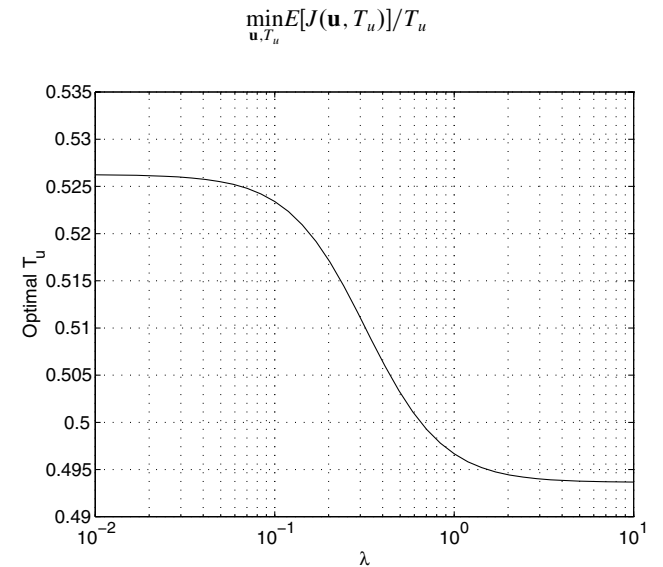**Fig. 5 Example trajectories in the planar H3BP with varying transfer times. Dots are placed every 0.05 time units.**

Fig. 7 Optimal nondimensional update time as a function of λ, for fixed σ_r in the planar H3BP.

Fig. 6 Expected cost divided by T_u as a function of T_u in the planar H3BP.
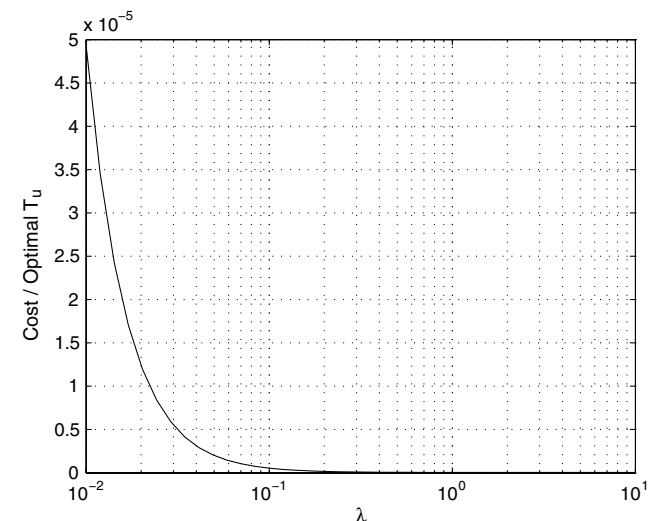

Fig. 8 Optimal nondimensional cost as a function of λ, σ_r fixed in the planar H3BP.

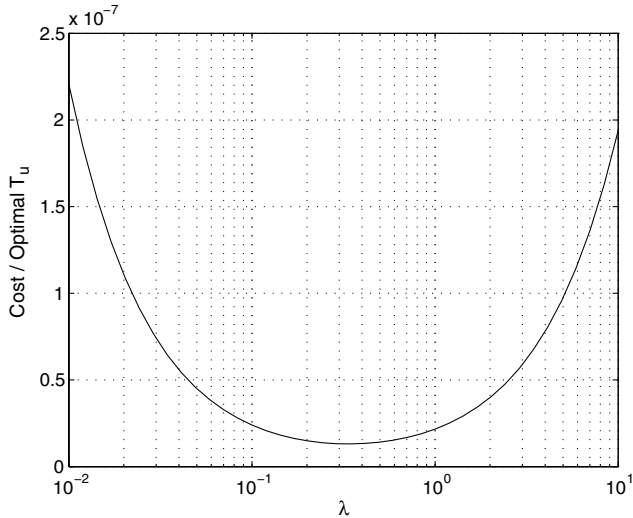Fig. 9 Optimal nondimensional cost as a function of λ, |**P**| fixed in the planar H3BP.

Fig. 10 Nominal halo orbit trajectories.

As can be seen in Fig. 8, if σ_r is fixed, it is optimal to let λ go to infinity, which is equivalent to letting P_v approach zero (i.e., low uncertainty in the velocity components). However, if |**P**| is held constant as λ varies (note the presence of an optimal value of λ in Fig. 9, $\lambda \approx 0.34$), indicating that given a certain amount of uncertainty (measured by a constant |**P**|), there is an optimal way to distribute the position and velocity uncertainties. For the Earth–sun system, $\lambda = 0.34$ corresponds to a ratio between 1-σ uncertainties $\sqrt{P_r/P_v} \approx 1.7 \times 10^6$, which is close to our assumed ratio between these measurement uncertainties. For the Earth–moon system, the optimal ratio of uncertainties is approximately 1.3×10^5. For a position uncertainty of 1 km, the optimal velocity uncertainty is about 0.75 cm/s. This aspect of the problem will be investigated in the future.

The curve in Fig. 9 scales with |**P**|, so that the value of λ yielding the minimum value does not change with |**P**|. From Eq. (13), the expected value of the cost divided by the optimal update time scales linearly as the entries in **P** are scaled. because $|c\mathbf{P}|$ scales in proportion to $c^n|\mathbf{P}|$, where c is a scalar and $\mathbf{P} \in \mathbb{R}^{n \times n}$, we have that

$$\left(\min_{\mathbf{u}, T_u} E[J(\mathbf{u}, T_u)]/T_u \right) \sim |\mathbf{P}|^{1/n}$$

In this planar case, $n = 4$, and so the cost scales with $|\mathbf{P}|^{0.25}$.

B. Halo Orbit Control

From the two oscillatory modes mentioned in the previous section, we see that near the equilibrium point, the linearized system is capable of producing planar periodic orbits. These orbits can also be found in the full nonlinear dynamics by examining the monodromy matrix $\boldsymbol{\Phi}(T, 0)$. The monodromy matrix locally captures the nonlinear periodic dynamics in a discrete-time, linear, time-invariant mapping and thus allows the system to be analyzed by eigenvalue methods. As the amplitude of these periodic orbits is increased, the eigenvalues of the monodromy matrix bifurcate and a new family of periodic orbits is produced. This new family is called the family of halo orbits, which are no longer in the plane and cannot be predicted using the equilibrium-point linearization. The halo orbits used in this paper, denoted as orbits A though E in Fig. 10, may be parameterized by their initial x coordinate, x_0. The values of x_0 for orbits A through E are 0.769, 0.7, 0.6, 0.5, and 0.45, respectively. The other initial conditions for the halo orbits were obtained using a method developed by Howell [9] that takes advantage of symmetry in the system. We developed software to numerical integrate the equations of motion using an eighth-order arbitrary-precision symplectic Runge–Kutta method [10] with 256 bits of precision [11].
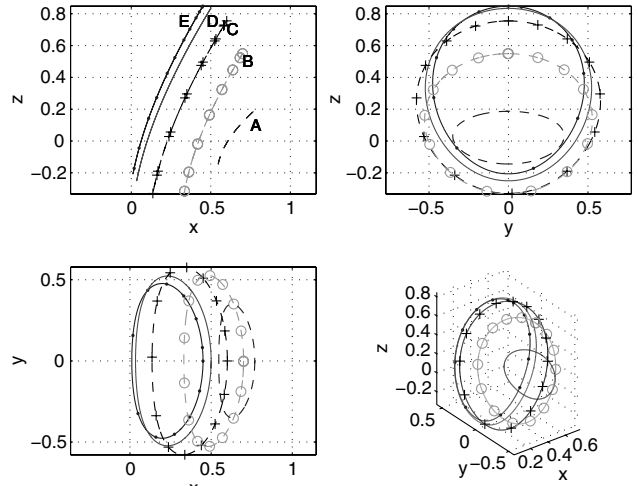
In the previous time-invariant example, each segment of control had the same statistical cost. Therefore, we only needed to consider the cost of one segment of control to draw conclusions about the long-term average cost. However, in this time-varying case, each segment will generally have a different cost. We may still determine the long-term average cost by considering only a finite length of time, due to the periodic nature of our system. We simply need to consider a period of time long enough such that the costs associated with all segments of the nominal trajectory are included. A natural choice for this is to choose two positive integers, n and m, such that the update time is approximated by $T_u \approx (n/m)T$, where T is the period of the system. We then only need to include the cost of segments up to time nT, because any segments after that will have already been included in the average long-term cost. An additional complication is that for each update time, the average cost per unit time will vary with the starting point of the algorithm along the orbit. Therefore, to obtain a statistical result that is independent of an arbitrary starting time, an average is performed with respect to the starting time. The unaveraged expected cost per update time for a halo orbit that is highly sensitive to the starting time is shown in Fig. 11, along with the average value for comparison. Another view of the same data is show in Fig. 12, which displays the data in a similar manner to Fig. 6.

Significant computational effort can be saved when computing the cost associated with multiple update times by choosing n and m wisely. By setting m to be the number of starting times we wish to
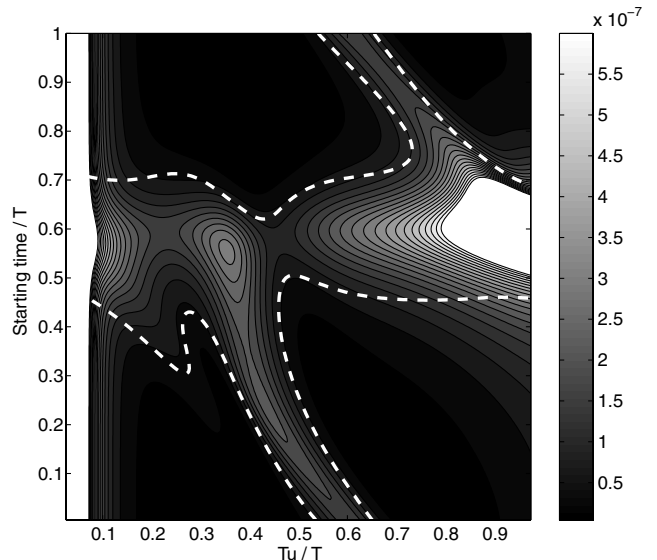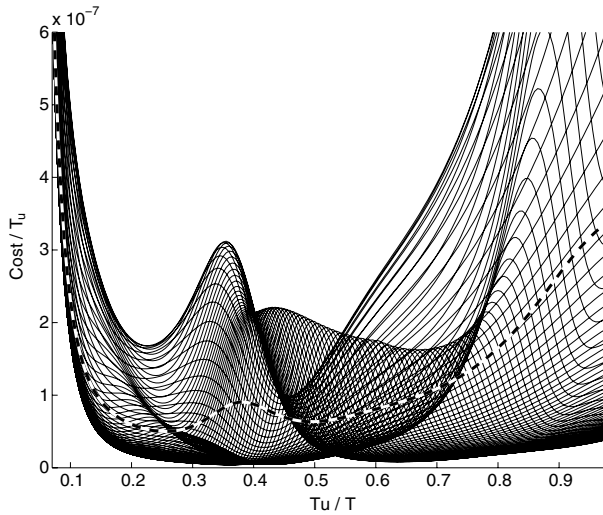

Fig. 11 Contours of the unaveraged expected cost per update time as a function of the update time and starting time for halo orbit E. Each axis is scaled by the orbit period and the average value for a given update time is shown as the dashed white line.

Fig. 12 Unaveraged expected cost per update time as a function of the update time for halo orbit E. Each solid line corresponds to a different starting time and the thick dashed line is the average over all starting times.

average over, the expected cost of control for each segment can be stored for each starting time and each n of interest. An illustration of this is shown in Fig. 13 for $m = 6$ with $n = 1$ and 2. In this analysis, we used $m = 100$ with n ranging from 5 to 95 for orbits A through D and $m = 200$ with n ranging from 10 to 190 for orbit E. This gives a worst-case resolution of T_u equal to 0.031 nondimensional time units, corresponding to orbit A. Once the expected cost for each segment has been calculated, an arithmetic average is taken with respect to each starting time.

To determine whether a given periodic orbit is stable or not, we define the Lyapunov characteristic exponent using the associated monodromy matrix, similar to the definition in [2]:

$$\alpha = \frac{\ell n \ \max_i |\lambda_i|}{T}$$
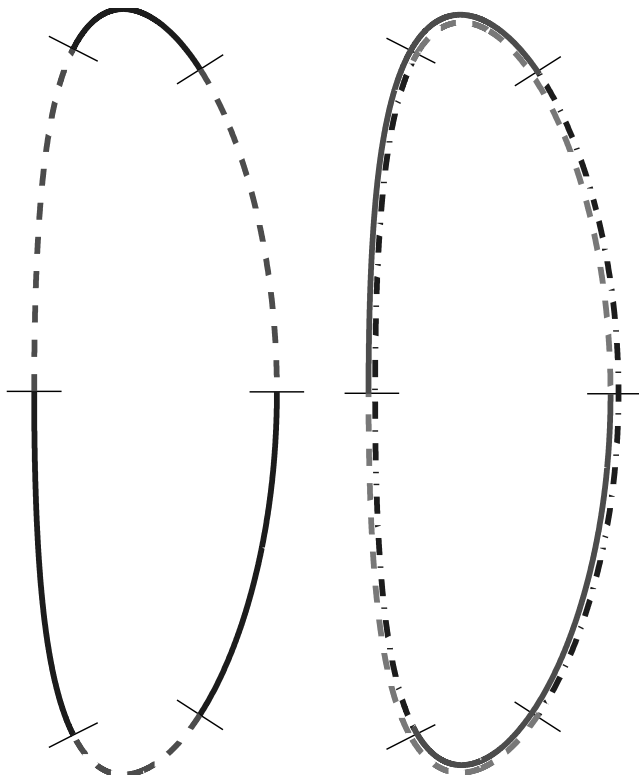
Fig. 13   An example halo orbit and control segments divided into $m = 6$ equal-time segments, with $n = 1$ on the left and $n = 2$ on the right.
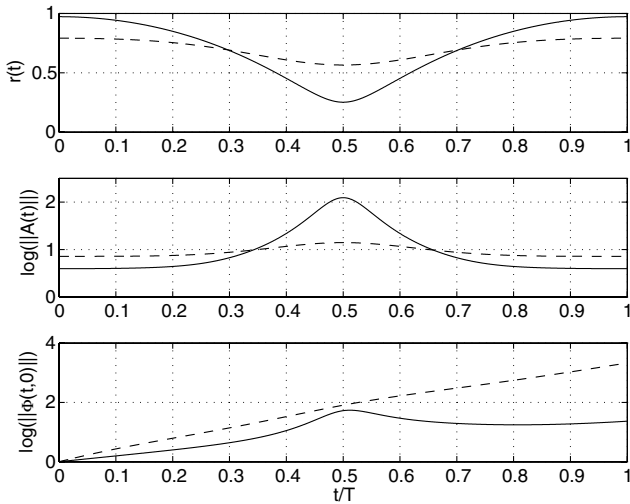
where λ_i are the eigenvalues of $\mathbf{\Phi}(T, 0)$. The characteristic exponent gives an idea of how quickly the state of the system will grow in time (on the order of $e^{\alpha t}$). If $\alpha > 0$, the system is unstable. The characteristic time is then $1/\alpha$, which gives a time scale on which the exponential effects develop. For linear time-invariant systems, this simplifies to the usual condition on the eigenvalues of the dynamics matrix; that is, the system is unstable if any of the eigenvalues have a real part greater than zero. For Hamiltonian systems, the existence of a stable manifold implies the existence of an unstable manifold.

The primary result of this analysis is that an optimal control-law update time exists for unstable time-varying systems, just as in the time-invariant case, as shown in Fig. 14. For halo orbit A using the same levels of uncertainty, the characteristic time of the instability was 0.42 time units, with the actual value occurring at about 0.61 time units (about 35 days for the Earth–sun system and about 2.7 days for the Earth–moon system). The cost associated with using the characteristic time as the update time is only 10% higher than the true minimum cost for this orbit, supporting a correlation between the characteristic time of the instability and the actual optimal update time. Note that the minimum cost per unit of time occurs very near to $T_u = 1/\alpha$ for orbits A, B, and C.
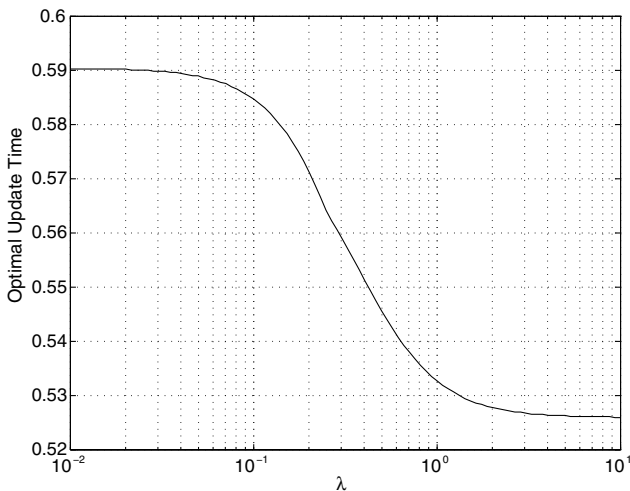
As seen in Fig. 14, the structure of the cost for orbits D and E bifurcates into a double-minimum case. This is due to the interesting dynamics of the halo orbits; as the orbits move farther out of plane, they make a closer approach to the secondary body, resulting in dynamics that are very strong compared with the rest of the orbit. Combining Eqs. (29–31) into standard first-order form with state $\mathbf{x} = [x \quad y \quad z \quad \dot{x} \quad \dot{y} \quad \dot{z}]^T$ and linearizing about the periodic orbit, we find $\delta\dot{\mathbf{x}} = \mathbf{A}(t)\delta\mathbf{x}$. The induced norm of $\mathbf{A}(t)$ gives an indication of how the eigenvalues of $\mathbf{A}(t)$ vary along the orbit, which in turn make the trajectory sensitive to uncertainties. The larger the norm, the stronger the sensitivity. Figure 15 shows a plot of $\log \|\mathbf{A}(t)\|$ and $\log \|\mathbf{\Phi}(t, 0)\|$ for two halo orbits: one highly out-of-plane and the other more in-plane. Note that for the highly-out-of-plane orbit, the sensitivity varies by up to 1.5 orders of magnitude throughout the orbit, whereas in the more in-plane orbit, it varies by less than 0.3. Because of this variation, the cost of control along a halo orbit varies depending on where measurements are taken. For example, consider a control segment in which $\|A\|$ is large initially then decreases quickly. In this, the unstable effect on the probability distribution is greatly enhanced, resulting in a higher control cost for the next segment. For a given update time, if the segments are structured such that $\|A\|$ is large when measurements are made, the cost is much higher than if $\|A\|$ were only large between measurements. This behavior is strong enough to hold even through the orbit average and is clearly visible in Fig. 14, particularly in orbits D and E. Each local maxima occurs just before the halfway point of the corresponding orbit, where $\|A\|$ is large, as in Fig. 15. In practice, orbit
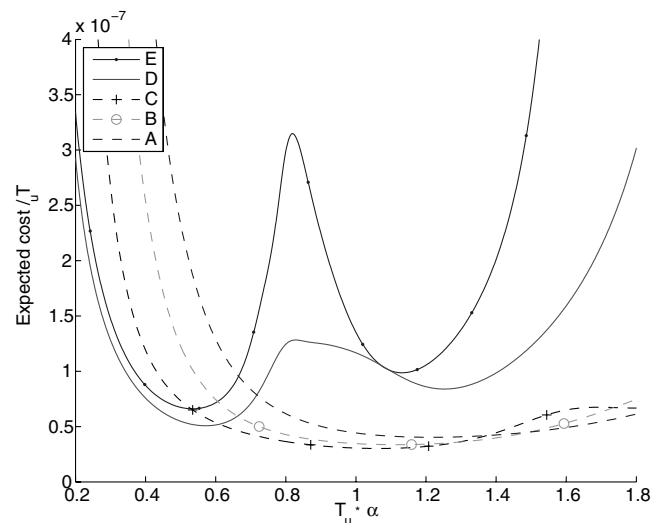
Fig. 14   Expected cost divided by T_u as a function of T_u (scaled by the characteristic time of the unstable mode) near several halo orbits.

Fig. 15   Plots of $r(t)$ (the distance from the secondary body), $\log \|A(t)\|$, and $\log \|\Phi(t,0)\|$ for halo orbits D (solid) and A (dashed), plotted against a fraction of their respective orbit periods, T.

Fig. 16   Halo orbit A optimal nondimensional update time as a function of λ, σ_r fixed.

The expected value of ΔV was calculated by a Monte Carlo simulation with 10,000 trials for each of the 100 starting points along the orbit. Confidence interval analysis shows that $E[\Delta V]/T_u$ is within 1.7% of the reported value, with 99% confidence. More accurate performance estimates could be performed with more information about detailed spacecraft specifications such as engine efficiencies.

Hill et al. [12] conducted a study of a lunar L_2 orbiter with measurement uncertainties of $P_r = 1$ km^2 and $P_v = 0$ km^2/s^2 and obtained a ΔV estimate of approximately 16 cm/s per year, although they budget 1 m/s per year. Using a similar halo orbit and those uncertainties, we obtain an upper bound of 12.3 cm/s per year, an actual value of 11.2 cm/s per year, and an optimal update time of 3.00 days. If we instead use a velocity uncertainty of $P_v = (10^{-5})^2$ km^2/s^2, we obtain an upper bound of 20.5 cm/s per year, an actual value of 19.7 cm/s per year, and an optimal update time of 3.13 days. It should be noted that the dynamics of the Earth–moon system are not well represented by the H3BP; however, the approximation is adequate for relative cost comparison [1].

IV.   Conclusions

This paper describes a method to analyze the average cost of controlling a linear system near an unstable trajectory. Such analyses are needed for mission planning purposes and for budgeting fuel needs due to statistical control. In particular, we show that for unstable systems, there is an optimal control-law update time, which is related to the system characteristic time. Additionally, if the total level of uncertainty is fixed, there is an optimal way to distribute uncertainty between the position and velocity states. These concepts are applied to spacecraft control in the vicinity of a halo orbit in the Hill three-body problem (H3BP) as well as one of the relative equilibrium points.

References

[1] Renault, C. A., and Scheeres, D. J., "Statistical Analysis of Control Maneuvers in Unstable Orbital Environments," *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 5, Sept.–Oct. 2003, pp. 758–769.
doi:10.2514/2.5110

[2] Scheeres, D. J., "Navigation of Spacecraft in Unstable Orbital Environments," *Proceedings of the International Conference on Libration Point Orbits and the Applications*, World Scientific Publishing, Singapore, June 2002, pp. 399–438.

[3] Michalska, H., and Mayne, D. Q., "Robust Receding Horizon Control of Constrained Nonlinear Systems," *IEEE Transactions on Automatic Control*, Vol. 38, No. 11, Nov. 1993, pp. 1623–1633.
doi:10.1109/9.262032

[4] Bellingham, J., Richards, A., and How, J., "Receding Horizon Control of Autonomous Aerial Vehicles," *Proceedings of the American Control Conference*, Vol. 5, Inst. of Electrical and Electronics Engineers, Piscataway, NJ, 2002, pp. 3741–3746.

[5] Bryson, A. E., and Ho, Y.-C., *Applied Optimal Control*, Hemisphere, New York, 1975, pp. 158–167.

[6] Park, C., Guibout, V., and Scheeres, D. J., "Solving Optimal Continuous Thrust Rendezvous Problems with Generating Functions," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 2, 2006, pp. 321–331.
doi:10.2514/1.14580

[7] Gubner, J. A., *Probability and Random Processes for Electrical and Computer Engineers*, Cambridge Univ. Press, New York, 2006.

[8] Bay, J. S., *Fundamentals of Linear State Space Systems*, WCB/McGraw–Hill, New York, 1999, pp. 68–69.

[9] Howell, K., "Three-Dimensional, Periodic, 'Halo' Orbits," *Celestial Mechanics and Dynamical Astronomy*, Vol. 32, No. 1, Jan. 1984, pp. 53–71.

[10] Hairer, E., Lubich, C., and Wanner, G., *Geometric Numerical Integration: Structure Preserving Algorithms for Ordinary Differential Equations*, 2nd ed., Springer, New York, 2004, pp. 24–34.

[11] Granlund, T., *GNU Multiple Precision Arithmetic Library*, Ver. 4.2.3, http://gmplib.org/manual/, July 2008 [retrieved 31 July 2008].

[12] Hill, K., Parker, J., Born, G. H., and Demandantez, N., "A Lunar L_2 Navigation, Communication, and Gravity Mission," AIAA/AAS Astrodynamics Specialist Conference, AIAA Paper 2006-6662, Aug. 2006.

determination may be more accurate near periapsis, which could potentially be weighed against the higher statistical control cost.

The optimal update time for this time-varying system has a dependence on λ that is very similar to the equilibrium-point example. Figure 16 shows the effect of λ on the optimal update time for orbit A, using the same parameters as in the equilibrium-point analysis. The variation in the optimal update time over the range of λ shown is about 3.74 days for the Earth–sun system.

We find a strong correlation between the characteristic time of an unstable trajectory and the optimal update time for the control of this trajectory. For strongly varying trajectories, we also find additional structure in the optimal time update, due to interactions between the trajectory and the gravitating bodies. These interactions raise interesting questions about how the interplay of measurement uncertainty, instability, and statistical costs are interrelated. We plan to further study these issues in the future and to consider the effect of stochastic accelerations on the system.

The actual nondimensional minimum expected value of J/T_u for orbit A in the sun–Earth system is 4.55×10^{-8}, corresponding with $T_u = 0.55$ nondimensional time units. Therefore, the upper bound on the expected ΔV per period from Eq. (28) is 8.15×10^{-4}, whereas the actual value is 5.97×10^{-4}, or 73.2% of the upper bound. As a dimensional value, this equates to

$$E[\Delta V]/\text{period} = 5.97 \times 10^{-4} \text{ km}/(\text{s} \cdot \text{period})$$