

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

Eric D. Gustafson and Daniel J. Scheeres

Abstract—The optimal control of a linear system is studied relative to a periodic unstable trajectory using continuous control. Gaussian state uncertainties are included, which induces a statistical cost of controlling the state over a long period of time. The length of time between control law updates directly impacts this statistical cost. When uncertainties are present in a hyperbolically unstable system, the time between control updates can take an optimal value. 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.

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. The control force and system dynamics are assumed to be deterministic, and the state estimates are assumed to have a Gaussian probability distribution.

Previously, Renault and Scheeres [1] conducted a similar study of optimal statistical control which considered the placement of impulsive control maneuvers near an unstable equilibrium point. The results of this paper serve to reinforce key results in 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 [2] concerning the qualitative impact of the update time on control cost using impulsive maneuvers are developed here for the continuous thrust case, and shown to be similar.

This paper does not explicitly consider stochastic accelerations in the system dynamics. All uncertainty is assumed to be adequately described by uncertainties in the estimation of the state. In addition, it is worthwhile to note that we almost exclusively focus on optimally updating control laws in the presence of uncertainty, not optimal control per se.

The procedure is outlined in its general form in the following sequence.

- 1) At some time t_1 , there is uncertainty in the state due to earlier errors. The uncertainties at this point can be

viewed as the steady-state uncertainties of the estimation process.

- 2) Design and implement a controller which would nominally cause the state to converge to the target state at time t_2 based on estimates of the state at t_1 (which are uncertain).
- 3) At time t_2 , error exists again due to uncertainties at time t_1 .
- 4) Design and implement a controller which would nominally cause the state to converge at time t_3 based on estimates of the state at t_2 (assume same covariance as 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.

In order 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; $E[J]/T$, where J is the cost incurred and T is the time-between-updates. That is, the optimal time-between-updates, T^* , is given by

$$T^* = \operatorname{argmin}_T \frac{E[J]}{T}. \quad (1)$$

The time-between-updates is assumed to be a constant parameter over the analysis period of interest for this study, although that is not a requirement for this method.

We split the control problem into two pieces:

- 1) Optimal control to target back to a nominal trajectory/state in a finite time.
- 2) Effect of state uncertainty on the nominal control, and how we can decrease the overall cost in the presence of uncertainty.

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

For problem (2), where we do not know what the initial state is precisely, 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 time is nominally

Graduate Student, Department of Aerospace Engineering, University of Michigan, 1320 Beal Ave, Ann Arbor, MI 48109-2140, email: edgus@umich.edu

A. Richard Seebass Professor, Aerospace Engineering Sciences, Colorado Center for Astroynamics Research, University of Colorado, Engineering Center, ECOT 611, 429 UCB, Boulder, CO 80309-0429, email: scheeres@colorado.edu, Member AAS, Associate Fellow AIAA

related to the characteristic time scale of the instability. This relationship will be discussed in more detail later in the paper.

This combined control and measurement strategy is a periodic update procedure where 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) where 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 to typical RHC applications. This necessitates a relatively long execution horizon. Additionally, extending the planning horizon does not lead to significant differences in the statistical cost of control, and taking the execution horizon equal to the planning horizon is not a limitation. This is mainly due to the insensitivity of the optimal open-loop control law for an unstable system to the length of the planning horizon (i.e. the optimal open-loop control law acts quickly to control the instability of the system, so increasing the planning horizon doesn't significantly change the control). 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 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

In our analysis of the timing of the control law updates, the state is assumed to be a Gaussian random vector (GRV), with the mean and covariance taken as outputs of an independent estimation process. The multivariate Gaussian probability distribution function with mean, \bar{m} , and covariance matrix, P is defined as [5]

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

The expected value of a function is

$$E[\vec{f}(\vec{x})] = \int_{-\infty}^{\infty} \vec{f}(\vec{\xi}) p(\vec{\xi}) d\vec{\xi}$$

and the mean, \bar{m} , and covariance, P , of a random variable are then

$$\begin{aligned} \bar{m} &= E[\vec{x}] = \int_{-\infty}^{\infty} \vec{\xi} p(\vec{\xi}) d\vec{\xi} \\ P &= E[\vec{x}\vec{x}^T] - \bar{m}\bar{m}^T. \end{aligned}$$

Let the control vector be $\vec{u}(t)$. The cost function during each interval between state updates is

$$J = \frac{1}{2} \int_{t_0}^{t_f} \vec{u}^T(t) \vec{u}(t) dt,$$

and the optimization problem to be solved during this interval is to find $\vec{u}(t)$ such that $\vec{u}(t)$ 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,

$$E[\delta\vec{x}_0] = \vec{0} \quad (2)$$

$$\text{Var}[\delta\vec{x}_0] = P_0. \quad (3)$$

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

$$E[\delta\vec{x}_1] = \vec{0} \quad (4)$$

$$\text{Var}[\delta\vec{x}_1] = P_1 = \Phi(t_1 - t_0) P_0 \Phi^T(t_1 - t_0). \quad (5)$$

The optimal control is computed using the sweep method as outlined in [6]. This produces an open-loop control scheme as in common in spacecraft control. This is also practical for the purpose of simulation, because the open-loop control law avoids the singularity present in time-varying gains of feedback control at the end of the execution horizon.

A. Statistical Cost

For convenience, we partition the initial state, \vec{x}_0 , into $[\vec{r}_0^T \ \vec{v}_0^T]^T$, where \vec{r}_0 is the initial position vector and \vec{v}_0 is the initial velocity vector. This yields the following expression for the deterministic cost:

$$J = J_{\Delta}(J_r |\vec{r}_0|^2 + 2J_{rv} \vec{r}_0 \cdot \vec{v}_0 + J_v |\vec{v}_0|^2), \quad (6)$$

where J_{Δ} , J_r , J_{rv} , and J_v are functions of the linear dynamics and the update time.

Once the deterministic cost is known, the expected value and variance of the cost can be computed. Taking the expectation, $E[\cdot]$, of (6) yields

$$\begin{aligned} E[J] &= J_{\Delta} (J_r E[|\vec{r}_0|^2] + 2J_{rv} E[\vec{r}_0 \cdot \vec{v}_0] + J_v E[|\vec{v}_0|^2]) \\ &= J_{\Delta} \left(J_r \left(\sigma_r^2(t_0) + E[|\vec{r}_0|^2] \right) \right. \\ &\quad \left. + 2J_{rv} \left(\sigma_{rv}^2(t_0) + E[\vec{r}_0]^T E[\vec{v}_0] \right) \right. \\ &\quad \left. + J_v \left(\sigma_v^2(t_0) + E[|\vec{v}_0|^2] \right) \right), \quad (7) \end{aligned}$$

where σ denotes the (co)variance of the given quantity at the beginning of the update interval.

Now, since the control law was chosen so that the expected values of the state is the zero vector ($E[\vec{r}_0] = E[\vec{v}_0] = \vec{0}$), (7) simplifies to

$$E[J] = J_{\Delta} (J_r \sigma_r^2(t_0) + 2J_{rv} \sigma_{rv}^2(t_0) + J_v \sigma_v^2(t_0)). \quad (8)$$

It is worthwhile to note that this expected cost is only a function of the time-between-updates, T , and the initial covariances, P_0 , because the coefficients J_* and the covariances σ_*^2 are determined completely by T and P_0 .

The variance of J , $\text{Var}[J]$, is given by

$$\text{Var}[J] = E[J^2] - (E[J])^2. \quad (9)$$

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

$$\begin{aligned} E[J^2] = & J_{\Delta}^2 \left\{ 3J_r^2 (\sigma_r^2)^2 + 4J_{rv}^2 (2(\sigma_{rv}^2)^2 + \sigma_{rv}^2 \sigma_v^2) \right. \\ & + 3J_v^2 (\sigma_v^2)^2 + 4J_r J_{rv} (2\sigma_r^2 \sigma_{rv}^2 + \sigma_{rv}^2 \sigma_r^2) \\ & \left. + 2J_r J_v (\sigma_r^2 \sigma_v^2 + 2(\sigma_{rv}^2)^2) + 12J_{rv} J_v \sigma_{rv}^2 \sigma_v^2 \right\}. \quad (10) \end{aligned}$$

Substituting (8) and (10) into (9) yields the variance of J :

$$\begin{aligned} \text{Var}[J] = & J_{\Delta}^2 \left\{ 2J_r^2 (\sigma_r^2)^2 + 4J_{rv}^2 ((\sigma_{rv}^2)^2 + \sigma_{rv}^2 \sigma_v^2) \right. \\ & + 2J_v^2 (\sigma_v^2)^2 + 4J_r J_{rv} (\sigma_r^2 \sigma_{rv}^2 + \sigma_{rv}^2 \sigma_r^2) \\ & \left. + 4J_r J_v (\sigma_{rv}^2)^2 + 8J_{rv} J_v \sigma_{rv}^2 \sigma_v^2 \right\}. \quad (11) \end{aligned}$$

As with the expected value of J , the variance of J is only a function of T and P_0 .

B. System Stability

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]:

$$\alpha = \frac{\ln[\max \text{eigenvalue}(\Phi(T, 0))]}{T}.$$

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, i.e., the system is unstable if any of the eigenvalues have a real part greater than zero.

C. Steady-State Minimum Expected Cost

Due to the complicated form of the expression for $E[J]/T$ (even for simple time-invariant systems), (1) cannot typically be solved for T^* in closed-form. Some statements can be made, however, about the behavior of T^* , depending on the dynamics of linear systems under study. For any double-integrator and oscillatory type dynamics, it can be shown that $E[J]/T$ achieves its minimum by letting $T \rightarrow \infty$ and that the actual value of $E[J]/T$ approaches zero. This is analogous to the impulsive control result obtained in [2] and implies that maneuver execution errors dominate the uncertainty. The behavior is different for a hyperbolically unstable system. As $T \rightarrow \infty$, the hyperbolic instability of the system causes $E[J]$ to grow exponentially and drives $E[J]/T \rightarrow \infty$. Thus, in general, there exists an optimum finite value, T^* , which minimizes $E[J]/T$. The value of T^* is strongly dependent on the characteristic time of the unstable mode ($1/\alpha$). In fact, the value of T^* is on the order of the characteristic time, although it also depends more weakly on the initial values of the covariance matrix. This is also analogous to impulsive control results obtain in [1]. For an ideal one 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.

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 in [1] considered control using impulsive maneuvers. In addition, we show that the results obtained for the linear time-invariant case (both in this study, and in the previous study) can be extended to linear time-varying systems.

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

$$\ddot{x} - 2\omega\dot{y} = -\frac{\mu}{r^3}x + 3\omega^2x + a_x \quad (12)$$

$$\ddot{y} + 2\omega\dot{x} = -\frac{\mu}{r^3}y + a_y \quad (13)$$

$$\ddot{z} = -\frac{\mu}{r^3}y - \omega^2z + a_z, \quad (14)$$

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, $\mu = 3.986 \times 10^5 \text{ km}^3/\text{s}^2$, $\omega = 1.991 \times 10^{-7} \text{ rad/s}$, $l = 2.159 \times 10^6 \text{ km}$, and $\tau = 5.023 \times 10^6 \text{ s}$.

The dimensional covariance matrix associated with the state estimates is assumed to be a 4×4 diagonal matrix (typical of spacecraft state estimation) with entries P_r and P_v ,

$$P_d = \begin{bmatrix} P_r \cdot I_2 & 0_{2 \times 2} \\ 0_{2 \times 2} & P_v \cdot I_2 \end{bmatrix}.$$

This covariance matrix may be nondimensionalized to obtain

$$\begin{aligned} P &= \begin{bmatrix} P_r/l^2 \cdot I_2 & 0_{2 \times 2} \\ 0_{2 \times 2} & P_v/(\omega l)^2 \cdot I_2 \end{bmatrix} \\ &= P_r/l^2 \begin{bmatrix} I_2 & 0_{2 \times 2} \\ 0_{2 \times 2} & P_v/(P_r \omega^2) \cdot I_2 \end{bmatrix}. \end{aligned}$$

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 yields the following form for P :

$$P = \sigma_r^2 \begin{bmatrix} I_2 & 0_{2 \times 2} \\ 0_{2 \times 2} & 1/\lambda^2 \cdot I_2 \end{bmatrix}. \quad (15)$$

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 $\lambda = 1.991$.

A. Planar Equilibrium Point Control

When the system is nondimensionalized by setting $\mu = \omega = 1$ in (12) and (13), the system has two equilibrium points using no control at $x = \pm 3^{-1/3}$, $y = 0$. Linearizing about either of these points and defining the perturbed state $\delta\vec{x} = [\delta x \ \delta y \ \delta \dot{x} \ \delta \dot{y}]^T$ yields the linear system

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

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,

$$J = \frac{1}{2} \int_{t_0}^{t_f} (a_x^2 + a_y^2) dt.$$

The optimal control law, and hence the trajectories themselves, depend on the final time and are plotted in Figure 1 for three different final times and various initial conditions. In this deterministic analysis, the cost uniformly decreases as the final time increases, however, in the following section when uncertainties are included in the analysis, this is not the case. We will show that using an update time of 0.5 time units, corresponding to Figure 1b, 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 Figure 2, using the uncertainty parameters given above. Due to the hyperbolically unstable dynamics, an optimal value of T_u clearly exists which minimizes the expected cost.

Figure 3 shows the effect of the nondimensional parameter λ on the optimal update time for the H3BP using the parameters described above. The variation in the optimal update time over the range of λ shown is about 1.75 days for the Earth-Sun system. It is important to 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, this continuous thrust result is consistent with the impulsive results in [1].

Figures 4 and 5 show the effect of λ on the value of the cost incurred over an update interval divided by the optimal update time, i.e.

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

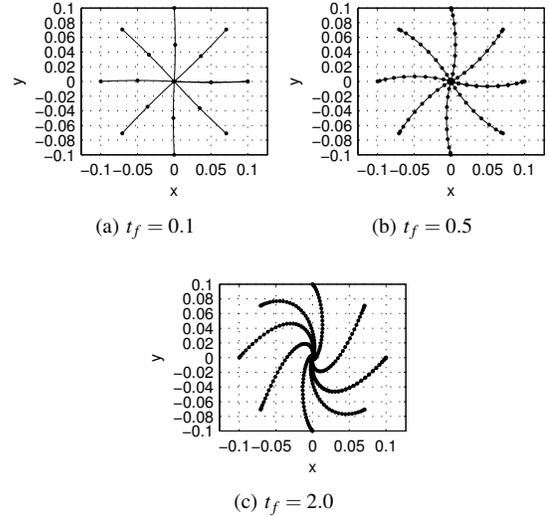


Fig. 1. Example trajectories with varying transfer times. Dots are placed every 0.05 time units.

As can be seen in Figure 4, 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 Figure 5, $\lambda \approx 0.34$, 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 surprisingly close to the actual 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 Figure 5 scales with $|P|$, so that the value of λ yielding the minimum value does not change with $|P|$. If the optimal cost divided by the optimal update time is plotted against $|P|$ on a log-log graph, the curve is a straight line with slope 0.25, indicating that $(\min_{\vec{u}, T_u} E[J(\vec{u}, T_u)]/T_u) \sim |P|^{0.25}$.

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. 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. One of the halo orbit used for this example is shown in Figure 6.

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

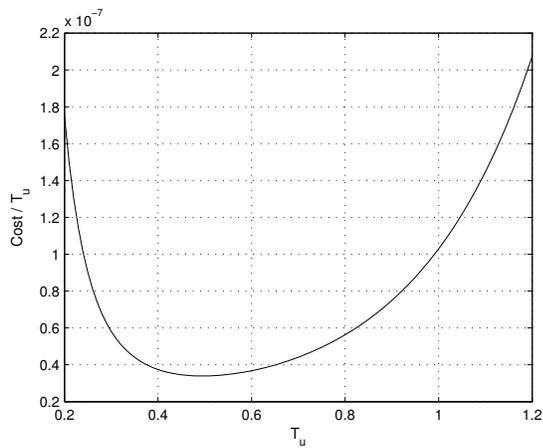


Fig. 2. Expected cost divided by T_u as a function of T_u

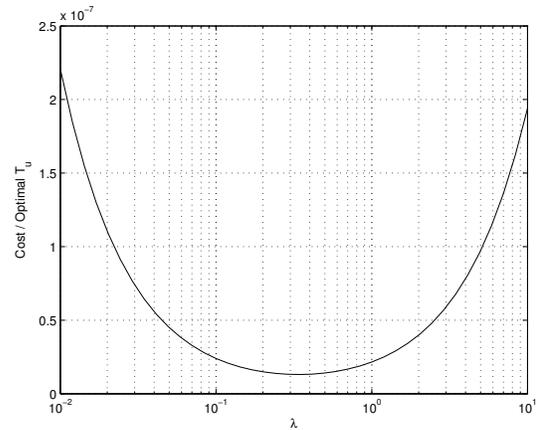


Fig. 5. Optimal nondimensional cost as a function of λ , $|P|$ fixed

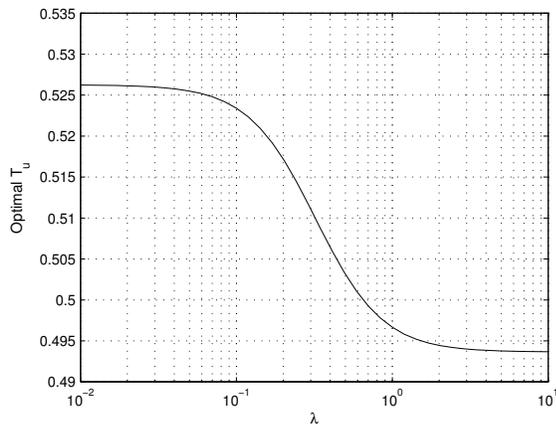


Fig. 3. Optimal nondimensional update time as a function of λ , for fixed σ_r

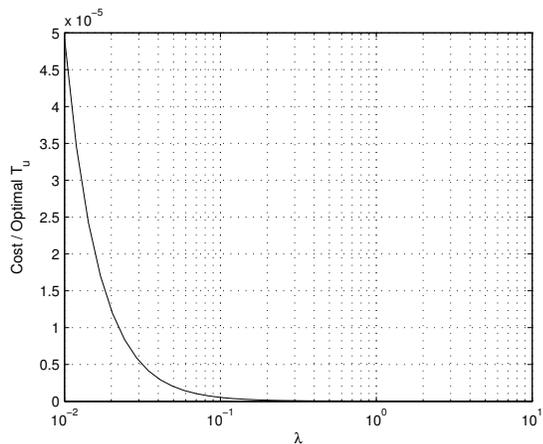


Fig. 4. Optimal nondimensional cost as a function of λ , σ_r fixed

in order 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 cost associated with all segments of the nominal trajectory are included. A natural choice for this type of analysis is to choose two positive integers, n and m , such that the update time is approximated by $T_u \approx \frac{n}{m}T$, 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, in order to obtain a statistical result that is independent of an arbitrary starting time, an average is performed with respect to the starting time.

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 Figure 7. For the example halo orbit with $x_0 = 0.769$ 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, showing a correlation between the characteristic time of the instability and the actual optimal update time.

As seen in Figure 7, the structure of the cost 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 to the rest of the orbit. Combining (12) through (14) into standard first-order form with state $\vec{x} = [x \ y \ z \ \dot{x} \ \dot{y} \ \dot{z}]^T$ and linearizing about the periodic orbit, we find $\delta\dot{\vec{x}} = A(t)\delta\vec{x}$. The induced norm of $A(t)$ gives an indication of how the eigenvalues of $A(t)$ vary along the orbit, which in turn make the trajectory sensitive to uncertainties. The larger the norm, the stronger the sensitivity. Figure 8 shows a plot of

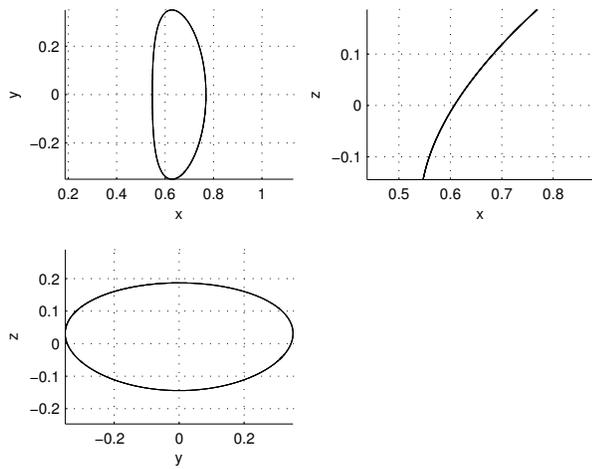


Fig. 6. Nominal halo orbit trajectory

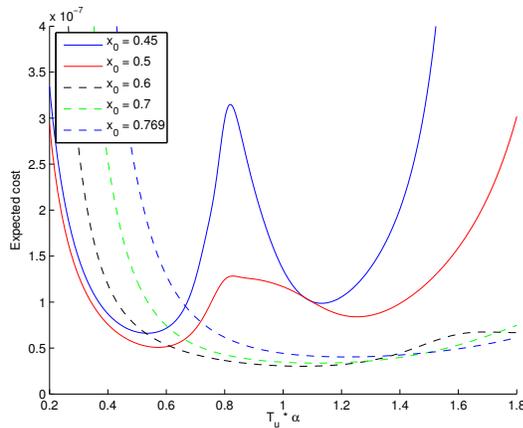


Fig. 7. Expected cost divided by T_u as a function of T_u (scaled by the characteristic time of the uncertainty) near several halo orbits. The halo orbits are parameterized by the initial value of their x coordinate. The lower the value of x_0 , the more out-of-plane the orbit.

$\log \|A(t)\|$ and $\log \|\Phi(t, 0)\|$ for two halo orbits; one highly out-of-plane, 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. Due to this variation, the cost of control along a halo orbit varies depending on where measurements are taken. For example, consider a control segment where $\|A\|$ is large initially, then decreases quickly. In this case, a measurement is taken when $\|A\|$ is large, and 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 Figure 7, particularly the orbits with $x_0 = 0.45$ and $x_0 = 0.5$. Each local maxima occurs just before the halfway point of the corresponding orbit, where $\|A\|$ is large, as in Figure 8.

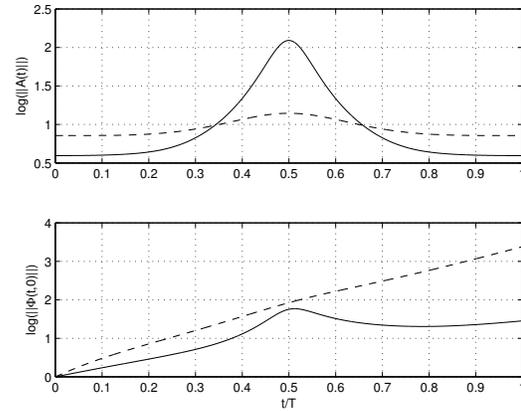


Fig. 8. Plots of $\log \|A(t)\|$ and $\log \|\Phi(t, 0)\|$ for the halo orbit with $x_0 = 0.5$ (solid) and $x_0 = 0.769$ (dashed), plotted against a fraction of their respective orbit periods, T .

IV. CONCLUSIONS AND FUTURE WORK

A. Conclusions

This paper describes a method to analyze the average cost of controlling a linear system near an unstable trajectory. 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.

B. Future Work

An interesting topic for future work is to further study the relation between the optimal update time and the system's characteristic time. Future goals also include extending the analysis to nonlinear stochastic systems, and to explain the relationship between position and velocity uncertainty.

REFERENCES

- [1] C. A. Renault and D. J. Scheeres, "Statistical analysis of control maneuvers in unstable orbital environments," *Journal of Guidance, Control, and Dynamics*, vol. 26, no. 5, September-October 2003.
- [2] D. J. Scheeres, "Navigation of spacecraft in unstable orbital environments," in *International Conference on Libration Point Orbits and the Applications*, June 2002.
- [3] H. Michalska and D. Q. Mayne, "Robust receding horizon control of constrained nonlinear systems," *IEEE Transactions on Automatic Control*, vol. 38, no. 11, November 1993.
- [4] J. Bellingham, A. Richards, and J. How, "Receding horizon control of autonomous aerial vehicles," in *American Control Conference*, vol. 5, 2002, pp. 3741–3746.
- [5] J. A. Gubner, *Probability and Random Processes for Electrical and Computer Engineers*. Cambridge University Press, 2006.
- [6] A. E. Bryson and Y.-C. Ho, *Applied Optimal Control*. Hemisphere Publishing Corporation, 1975.
- [7] C. Park and D. J. Scheeres, "Determination of optimal feedback terminal controllers for general boundary conditions using generating functions," *Automatica*, vol. 42, pp. 869–875, 2006.
- [8] R. F. Stengel, *Optimal Control and Estimation*. Dover Publications, 1994.
- [9] D. T. Greenwood, *Classical Dynamics*. Dover Publications, Inc., 1977.