Adaptive Optimal Timing Control of Hybrid Systems

Dennis Ding, Yorai Wardi, and Magnus Egerstedt
{ding,ywardi,magnus}@ece.gatech.edu
Georgia Institute of Technology
School of Electrical and Computer Engineering
Atlanta, Georgia, 30332, USA

Extended Abstract

1 Introduction

In this extended abstract we consider the problem of optimizing over the switching times that dictate the transition moments in a multi-modal dynamic system. Moreover, the problem class we are considering is one in which the complete cost-to-go is unavailable. Rather, we only have access to the instantaneous cost when making the control decisions.

This problem falls under the category of optimal control of hybrid systems. This is an area of research that has been extensively investigated during the past ten years. (See [6] and [7] for a representative sample.) This line of inquiry has been motivated by a number of application domains, such as robotics ([2]), sensor scheduling ([3]), and power electronics ([1]).

The novelty of the work in this paper comes from the fact that we explicitly focus on the on-line aspect of the optimization problem. In particular, we will design a process that lets a switch-time control parameter evolve over time in such a way that it always remains locally optimal given the current information. To achieve this, it is assumed that a measure of the "rate of the change" of the cost function can be made in run time. As an example, consider the problem of tracking a reference point whose time evolution is not known in advance. This problem arises for example in the field of power electronics.

Previously, the problem of optimal timing control based on partial state information was considered (see [4]). The research presented in this paper is the extension of this work. The results of this paper can be used to solve a class of problems, in which the problem considered in [4] is one of them. Furthermore, this problem can also be considered as a reinforcement learning problem. However, this paper utilizes special structures in the hybrid system that can not be easily exploited by traditional reinforcement learning techniques.

To summarize, the goal of this extended abstract is to present an approach that allows us to adaptively update the switch-time parameters in such a way that they are locally optimal at each time instant, given the available information up to that point. The approach consists of two phases. In the first phase, an optimal solution is computed based on the initial information of the cost function. This calculation is computationally expensive, and it must be performed once "off-line", before the system starts evolving. In the second phase, as the system runs, the derivative of the optimal solution is computed based on the "rate of change" measured for the cost function. The optimal solution is then updated based on its derivative. The rate of change of the cost function is measured in the form of partial derivatives of the cost function.

2 Instantaneous and Total Costs

The systems we consider are of the form

$$\dot{x}(t) \ \in \ \{f_\alpha(x), \ \alpha \in \mathcal{A}\}, \tag{1}$$

with initial condition x_0 given, where $x \in \mathbb{R}^n$, $\mathcal{A}$ is a finite index set, and for every $\alpha \in \mathcal{A}$, $f_\alpha : \mathbb{R}^n \rightarrow \mathbb{R}^n$ is a continuously differentiable function. The functions f_α, $\alpha \in \mathcal{A}$, correspond to the various modes of the systems, and are referred to as modal functions. The time-horizon is the interval $[0, T]$ for a fixed $T > 0$. A finite sequence $\bar{\alpha} = \{\alpha_1, \ldots, \alpha_N\} \in 2^{\mathcal{A}}$ is used to describe a schedule of the modes corresponding to a sequence of modal functions $\{f_{\alpha_1}, \ldots, f_{\alpha_N}\}$ in the RHS of (1). For the scope of this paper the switching sequence of the system is assumed to be given *a priori*, hence $\bar{\alpha}$ is known and is not part of the control variable of the problem. We denote the switching times between the ith mode and the $(i+1)st$ mode by τ_i, and the vector $[\tau_1, \tau_2, ..., \tau_N]^T$ as the switching time vector $\bar{\tau}$. For convenience, we also define $\tau_0 := 0$ and $\tau_{N+1} := T$, denote $f_i := f_{\alpha_i}$, and $F(x, t) \ := \ f_i(x)$ for all $t \in [\tau_{i-1}, \tau_i), \ i = 1, \ldots, N+1$. Equation (1) then can be written in the following form with initial condition x_0:

$$\dot{x}(t) \ = \ F(x, t). \tag{2}$$

In a general optimal timing control problem, a cost function is defined as follows. Let $L : \mathbb{R}^n \rightarrow \mathbb{R}$ be a continuously differentiable function. The cost function J is defined by $J(\bar{\tau}) = \int_0^T L(x)dt$, where L is referred to as the instantaneous cost. In the setting of this paper, an adaptive version of the instantaneous cost is formulated since the time-varying information required to evaluate the instantaneous cost is only available at real time. Therefore, the instantaneous cost L formulated in this paper is a function of two time indices, t and s. t refers to the *real time*, which is amount of elapsed since the system starts evolving. s refers to the *simulated time*, which is time projected into the future. At each time instant t, the trajectory of the state is simulated on the future time axis indexed by s. L is evaluated over this future state trajectory, denoted by $\tilde{x}(s)$, which is defined as the follows.

$$\dot{\tilde{x}}(s) \ = \ F(\tilde{x}, s), s \in [t, T], \tag{3}$$

with the initial condition $\tilde{x}(t) = x(t)$ and given $\bar{\tau}$.

Therefore, the instantaneous cost function is denoted by $L(t, \tilde{x}(s), s)$ and is assumed to be continuously differentiable. A picture of the system at time t is shown in figure 1 to illustrate the usage of the two time indices.

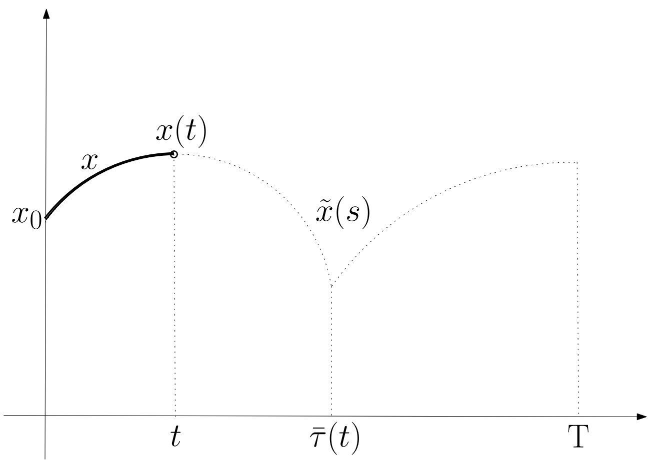

Figure 1: Current State and the State Simulated into the future.

Similarly, in this paper we formulate a cost-to-go function defined on the adaptive instantaneous cost mentioned above.

$$J(t, x(t), \bar{\tau}) = \int_t^T L(t, \tilde{x}(s), s)ds. \tag{4}$$

Observe that the cost-to-go function J depends on the current time t, the initial condition of the simulated state trajectory $x(t)$, and the switching time vector $\bar{\tau}$.

In this extended abstract henceforth we denote the problem of optimizing the cost-to-go function at time t by Π_t.

3 Switch-time Dynamics

The approach of this paper allows the means to compute the trajectory of the switching time vector, such that it is a local minima to the cost-to-go function at each time instant t. This optimal switching time vector is therefore a function of time t as well, and is denoted as $\bar{\tau}(t)$.

Observe that if we obtain the optimal switching time vector $\bar{\tau}(t)$ and substitute it in (4), then $J(t, x(t), \bar{\tau}(t))$ is locally optimal and

$$\frac{\partial J}{\partial \tau}(t, x(t), \bar{\tau}(t)) = 0, t \in [0, T]. \tag{5}$$

The aim of this paper is to provide the time derivative of the optimal switching time vector $\bar{\tau}(t)$ at time t, namely $\dot{\bar{\tau}}(t)$, so that optimality is conserved. In another word, we aim to compute $\dot{\bar{\tau}}(t)$, so that if $\bar{\tau}(t)$ is a solution point for Π_t, then $\bar{\tau}(t + dt)$ computed by this continuous process

$$\bar{\tau}(t + dt) = \bar{\tau}(t) + \dot{\bar{\tau}}(t)dt \tag{6}$$

is a solution point for Π_{t+dt}, and $\frac{\partial J}{\partial \tau}(t + dt, x(t + dt), \bar{\tau}(t + dt)) = 0$. Therefore, if we start at initial condition $\bar{\tau}(0)$ that optimizes the problem Π_0, then we can obtain the entire trajectory of the optimal switching time vector.

As mentioned earlier, to obtain the optimal switching time trajectory, the initial condition $\bar{\tau}(0)$ needs to be computed. This computation is referred to as the "off-line" phase of the approach. There are multiple approaches to achieve this computation. The authors use a gradient descent method proposed earlier in their research (see [5]). In each iteration of the gradient descent approach, a gradient of cost as a function of switching time vector, $\frac{\partial J}{\partial \tau}$, is computed. Then the switching time is updated by $\bar{\tau} = \bar{\tau} - \gamma \frac{\partial J}{\partial \tau}$, where γ is a suitable step-size. The iterations terminate when $\frac{\partial J}{\partial \tau} = 0$.

This paper first shows that the general formulation to compute $\bar{\tau}(t)$ is through the following proposition.

Proposition 2.1

Assume that $\bar{\tau}(\xi)$ is a local minimum to Π_ξ for $\xi \in [0, t]$. As a consequence, the second derivative of J with respect to $\bar{\tau}$ is strictly positive. Then the time derivative of $\bar{\tau}(t)$ can be determined by:

$$\dot{\bar{\tau}}(t) = -\left(\frac{\partial^2 J}{\partial \tau^2}(t, x(t), \bar{\tau}(t))\right)^{-1} \left(\frac{\partial^2 J}{\partial t \partial \tau}(t, x(t), \bar{\tau}(t)) + \frac{\partial^2 J}{\partial x \partial \tau}(t, x(t), \bar{\tau}(t))\dot{x}(t)\right). \tag{7}$$

The proof of this proposition follows from (5), and is omitted from this extended abstract.

Furthermore, it is shown that the terms $\frac{\partial^2 J}{\partial t \partial \tau}(t, x(t), \bar{\tau}(t)) + \frac{\partial^2 J}{\partial x \partial \tau}(t, x(t), \bar{\tau}(t))\dot{x}(t)$ can be obtained by the following equation as a result of simple Taylor expansions.

$$\frac{\partial^2 J}{\partial t \partial \tau}(t, x(t), \bar{\tau}(t)) + \frac{\partial^2 J}{\partial x \partial \tau}(t, x(t), \bar{\tau}(t))\dot{x}(t) = \lim_{dt \to 0} \frac{1}{dt} \frac{\partial J}{\partial \tau}(t + dt, x(t + dt), \bar{\tau}(t)) \tag{8}$$

Equation (7) alone does not provide explicit forms of equations to update $\bar{\tau}(t)$. It is needed to carry out the analysis further to reveal the explicit equations to compute the RHS of (8). Through these equations we can see how the rate of change of the cost function measured at time t are needed to compute this time evolution. This results will be included in this extended abstract without derivation.

Let us denote

$$\frac{\partial L}{\partial \tilde{x}}(t, \tilde{x}, s) := L_x(t, \tilde{x}, s). \tag{9}$$

Additionally, denote the state transition matrices corresponding to the systems $\dot{\tilde{x}}(s) = -\frac{\partial f_i}{\partial \tilde{x}}(\tilde{x}(s))\tilde{x}(s)$ as $\Phi_i(t_1, t_2)$. Since the switching sequence is given a priori, these state transition matrices can be calculated a priori as well. The main result of the paper is contained in the following proposition.

Proposition 2.2

Assume that $\bar{\tau}(\xi)$ is a local minimum to Π_ξ for $\xi \in [0, t]$, then the time derivative of $\bar{\tau}(t)$ can be determined by:

$$\dot{\bar{\tau}}(t) \quad = \quad -\left(\frac{\partial^2 J}{\partial \tau^2}(t, x(t), \bar{\tau}(t))\right)^{-1} [M(1), M(2), ..., M(N)]^T \tag{10}$$

where

$$M(i) = \int_{\bar{\tau}_i(t)}^{\bar{\tau}_{i+1}(t)} \Big(\frac{\partial L_x}{\partial t}(t, \tilde{x}, s) + \frac{\partial L_x}{\partial \tilde{x}}(t, \tilde{x}, s)\dot{x}(t)\Big)\Phi_i(s, \bar{\tau}(t))ds\Big(f_i(x(\bar{\tau}_i(t))) - f_{i+1}(x(\bar{\tau}_i(t)))\Big) \tag{11}$$

It should be remarked that the Hessian matrix $\frac{\partial^2 J}{\partial \tau^2}(t, x(t), \bar{\tau}(t))$ can be explicitly obtained when dimension of the switching time vector N is small. As N grows, it is no long feasible to compute it explicitly and numerical approximations will be used.

4 Concluding Remarks

In this extended abstract we proposed a way of computing the evolution of the switch-time vector in such a way that it is locally optimal at each time instant, given the currently available information. In the final paper, we will give the formal proofs to the main propositions in the previous section. We will also provide a simulation example that highlights the operation of the proposed method on a linear as well as a nonlinear system.

References

[1] A. Beccuti, G. Papfotiou and M. Morari. Optimal Control of the Boost dc-dc Converter *IEEE Conf. on Decision and Control*, Seville, Spain, 2005

[2] M. Boccadoro, Y. Wardi, M. Egerstedt, and E. Verriest. Optimal Control of Switching Surfaces in Hybrid Dynamical Systems. *Journal of Discrete Event Dynamic Systems*, Vol. 15, No. 4, pp. 433-448, Dec.2005.

[3] R. Brockett. Stabilization of Motor Networks. *IEEE Conf. on Decision and Control*, pp. 1484–1488, 1995.

[4] M. Egerstedt, S. Azuma, and Y. Wardi. Optimal Timing Control of Switched Linear Systems Based on Partial Information. In *Nonlinear Analysis: Theory, Methods & Applications*, 2006.

[5] M. Egerstedt, Y. Wardi, and H. Axelsson. Transition-Time Optimization for Switched-Mode Dynamical Systems. *IEEE Trans. on Automatic Control*, Vol. AC-51, pp. 110-115, 2006.

[6] M.S. Shaikh and P. Caines. On Trajectory Optimization for Hybrid Systems: Theory and Algorithms for Fixed Schedules. *IEEE Conf. on Decision and Control*, pp. 1997–1998, 2002.

[7] X. Xu and P. Antsaklis. Optimal Control of Switched Autonomous Systems. *IEEE Conf. on Decision and Control*, pp. 4401–4406, 2002.