PROPERTIES OF RECEDING HORIZON CONTROLLERS FOR SOME HYBRID SYSTEMS WITH EVENT UNCERTAINTIES

Christos G. Cassandras [*,1]
Reetabrata Mookherjee [*,1]

[*] *Dept. of Manufacturing Engineering and Center for Information and Systems Engineering, Boston University, Brookline, MA 02446*

Abstract: We consider optimal control problems for a class of hybrid systems with switches dependent on an external event process. In the case where all event times in this process are fully known, the solution to such problems was obtained in prior work. When event times are uncertain or unknown, we have proposed in prior work a Receding Horizon (RH) control scheme in which only some future event information is available within a time window of length T and have obtained several properties of this scheme. In this paper, we derive additional properties, including the fact that the error due to lack of future event information is monotonically decreasing under certain conditions and may be zero for segments of the sample path, depending on the window length T.

Keywords: Receding Horizon, Optimal Control, Hybrid Systems, Discrete Event Systems.

1. INTRODUCTION

Hybrid Systems are characterized by the combination of *time-driven* and *event-driven* dynamics. A simple way to think of a hybrid system is as one characterized by a set of operating "modes", each one evolving according to time-driven dynamics described by differential (or difference) equations. The system switches between modes through discrete events which may be controlled or uncontrolled. Controlling the switching times, when possible, and choosing among several feasible modes, whenever such choices are available, gives rise to a rich class of optimal control problems (Branicky *et al.*, 1998),(Piccoli, Dec. 1998),(Sussmann, Dec. 1999),(Hedlund and Rantzer, 1999),(Xu and Antsaklis, Dec. 2000*a*). Unfortunately, not only does one have to deal with the well-known "curse of dimensionality" in such problems, but there are at least two additional sources of complexity to deal with, i.e., the presence of switching events causing transitions from one mode to another (which introduces a combinatorial element into the control), and the presence of event-driven dynamics for the switching times (which introduce nondifferentiabilities). To overcome these difficulties, it is often natural to hierarchically decompose a system into a lower-level component representing physical processes characterized by time-driven dynamics and a higher-level component controlling discrete events related to these physical processes. In this spirit, a hierarchical decomposition method was introduced in (Gokbayrak and Cassandras, March 2000) and,

[1] Supported in part by the National Science Foundation under Grants ACI-98-73339 and EEC-00-88073, by AFOSR under contracts F49620-01-0056 and F49620-01-1-0348, and by ARO under grant DAAD19-01-0610.

independently, in (Xu and Antsaklis, 2000*b*). The explicit solution of the lower and higher-level problems depends on the specifics of the time-driven and event-driven dynamics involved.

To complicate matters, some of the most interesting problems one encounters in dealing with hybrid systems involve some form of uncertainty, which generally calls for stochastic modeling and solution techniques. In this paper, we attempt to deal with this issue for a class of optimal control problems where the event-driven dynamics may be captured through 'max-plus' equations. In such cases, switching times are controllable, but they are dependent upon an external event process in which the associated event times $\{a_1, \ldots, a_N\}$ are generally unknown. If this sequence is fully deterministic, optimization problems formulated as in (Pepyne and Cassandras, 1998),(Cassandras *et al.*, 2001) can be efficiently solved through the "Forward Algorithm" presented in (Cho *et al.*, 2001). If it is not, then one approach is to model it as a stochastic process as in (Gokbayrak and Cassandras, Dec. 1999) where the structure of an optimal policy can be determined, but explicit calculations are difficult.

In (Cassandras and Mookherjee, 2003*a*), we take a different approach in the way we view uncertainties in $\{a_1, \ldots, a_N\}$. In past work, it was assumed that all future events following any time t when a control decision needs to be made are known. Thus, if we associate with t an "information window" $[t, t+T]$, it was assumed that $T = \infty$. A natural next step is to consider $T < \infty$ and within a window $[t, t+T]$ assume that event times (if any are present) are deterministically known. Event times outside $[t, t+T]$ can only be described probabilistically, possibly estimated, or there may be no information at all regarding events beyond time $t+T$. The optimal control problem we now tackle is defined over a "receding horizon" determined by the length T of the window at the controller's disposal. The nature of this scheme lends itself to what we term a *Receding Horizon* (RH) controller. By varying T, we can then analyze the performance of the resulting controller, which becomes an approximation to the optimal one (and recovers it for $T \to \infty$). The idea of a "receding horizon" is one commonly associated with optimal control problems for which feedback solutions are extremely hard or impossible to obtain (e.g., (Mayne and Michalska, 1990)) and it is usually encountered in model predictive control. Intuitively, it is intended to trade off anticipating future unknown events against simpler, tractable control over a shorter, but containing known events, planning horizon. Our work explores the extent to which this way of dealing with uncertainty is effective by analyzing the relationship between the optimal and the RH controller's cost/performance and the role that T plays. In this paper, we use the Receding Horizon (RH) control scheme introduced in (Cassandras and Mookherjee, 2003*a*) and compare sample paths obtained through this scheme to an optimal sample path. An optimal sample path is decomposed into segments which have been shown to be decoupled (Cassandras *et al.*, 2001). Each such segment may include certain events termed "critical" with special properties. Several properties of the RH controller in the absence of critical events in the optimal path were derived in (Cassandras and Mookherjee, 2003*a*). In particular, we have shown that the *error* introduced by the RH controller relative to the optimal one is, under certain conditions, monotonically decreasing and may become zero for parts of the sample path. In this paper, we carry out the analysis of sample paths under the RH controller when the optimal sample path contains such critical events. We are able to show that this "error-reducing" property still holds under certain conditions and fully characterize these conditions.

2. OPTIMAL CONTROL PROBLEM AND SOLUTION

In hybrid systems, the state consists of temporal and physical components. The class of problems we will concentrate on involves event-driven switching time dynamics described by

$$x_i = \max(x_{i-1}, a_i) + s_i(u_i) \qquad (1)$$

where $\{a_i\}$, $i = 1, \ldots, N$, is a given sequence of event times corresponding to an asynchronous event process operating independently of the physical processes $\{z_i(t), t \in [x_{i-1}, x_i)\}$ (we define the physical state of the system after the ith event by $z_i(t)$). This "max-plus" recursive relationship is the well known Lindley equation in queueing theory (Cassandras and Lafortune, 1999). It is worhwhile to mention that the presence of the max function introduces a nondifferentiable component into the solution of the overall problem. In the context of a manufacturing system workstation, the mode switches correspond to jobs that we index by $i = 1, \ldots, N$. A job is associated with a *temporal* state evolving according to (1) where x_i is the departure time of the ith job when the server processes one job at a time on a first-come first-served nonpreemptive basis. Jobs arriving when the server is busy wait in an infinite-capacity queue and $\{a_1, \ldots, a_N\}$ is a sequence of job arrival times. The processing time of the ith job (which we will denote by C_i) is $s_i(u_i)$. A job is also associated with a *physical* state evolving according to $\dot{z}_i = g_i(z_i, u_i, t)$, $z_i(x_{i-1}) = z_i^0, t \in [x_{i-1}, x_i)$ and describing changes the ith job undergoes while in process in such quantities as the temperature, size, weight or some other measure of the "quality" of the job. Thus, the ith "mode" of a workcenter in this context corresponds to the processing of the ith job. The interaction of the time-driven and event-driven dynamics leads to a natural trade-off between temporal requirements on job completion times and physical requirements on the quality of the completed jobs. Let us briefly review the optimization problem introduced in (Cassandras *et al.*, 2001) and solved through the Forward Algorithm developed in (Cho *et al.*, 2001). When the processing of each job stops as soon as a given "quality level" in its physical state z_i is reached and the control is the amount of processing time, i.e., $s_i(u_i) = u_i$ in (1), the problem has been shown to become

$$\min_{u_1,\ldots,u_N} \sum_{i=1}^{N} [\theta_i(u_i) + \psi_i(x_i)] \qquad (2)$$

subject to (1), with control variables u_i assumed to be scalar and not time-dependent. Thus, u_i is the processing time of the ith job, chosen at the beginning of its processing cycle, i.e., at time $\max(x_{i-1}, a_i)$. The cost function $\theta_i(u_i)$ penalizes poor physical quality, in the sense that less processing time monotonically decreases quality and, hence, increases a cost $\theta_i(u_i)$, while $\psi_i(x_i)$ imposes a cost on the departure time x_i. As in (Cassandras *et al.*, 2001), we make the following assumptions:

Assumption A1. For each $i = 1, \ldots, N$, $\theta_i(\cdot)$ is strictly convex, twice continuously differentiable and monotonically decreasing with $\lim_{u_i \to 0^+} \theta_i(u_i) = -\lim_{u_i \to 0^+} \frac{d\theta_i}{du_i} = \infty$ and $\lim_{u_i \to \infty} \theta_i(u_i) = \lim_{u_i \to \infty} \frac{d\theta_i}{du_i} = 0$.

Assumption A2. For each $i = 1 = 1, \ldots, N$, $\psi_i(\cdot)$ is strictly convex, twice continuously differentiable, and its minimum is obtained at a finite point δ_i.

A typical sample path of this system can be partitioned into "busy" and "idle" periods. In this paper, we make use of the following definitions (see also (Cassandras *et al.*, 2001)). An *idle* period is a time interval (x_k, a_{k+1}) such that $x_k < a_{k+1}$ for any $k = 1, ..., N-1$. A *Busy Period* (BP) is a set of contiguous jobs $\{k, ..., n\}$, $1 \leq k \leq n \leq N$ such that the following conditions are satisfied: (i) $x_{k-1} < a_k$, (ii) $x_n < a_{n+1}$, and (iii) $x_i \geq a_{i+1}$ for every $i = k, ..., n-1$. A *busy-period structure* is a partition of the jobs $1, ..., N$ into busy periods. A job C_i is *critical* if it departs at the arrival time of the next job C_{i+1}, i.e. $x_i = a_{i+1}$. Finally, consider a contiguous job subset $\{C_k, \ldots, C_n\}$, $1 \leq k \leq n \leq N$. This subset is said to be a *block* if (i) $x_{k-1} \leq a_k$ and $x_n \leq a_{n+1}$, and (i) the subset contains no critical jobs.

Obtaining an explicit solution to problem (2) is tantamount to identifying the BP structure of the optimal state trajectory and then solving a nonlinear optimization problem within each BP. Let us denote a BP that starts at a_k and ends at x_n by the job indices (k, n). Then, we define the problem $Q(k, n)$:

$$Q(k,n) : \min_{u_k,\ldots,u_n} \left\{ \begin{array}{c} \sum_{i=k}^{n} \{\theta_i(u_i) + \\ \psi_i(a_k + \sum_{j=k}^{i} u_j)\} : u_i \geq 0 \end{array} \right\} \qquad (3)$$

$$\text{s.t.} \quad a_k + \sum_{j=k}^{i} u_j \geq a_{i+1}, \quad i = k, \ldots, n-1$$

Note that we have set $\psi_i(x_i) = \psi_i(a_k + \sum_{j=k}^{i} u_j)$ since, within a BP, $x_{j+1} = x_j + u_{j+1}$ for all $i = k, \ldots, n-1$. The constraint represents the requirement $x_i \geq a_{i+1}$ for any job $i = k, \ldots, n-1$ belonging to the BP. Since the cost functional is continuously differentiable and strictly convex, the problem $Q(k, n)$ is also a convex optimization problem with linear constraints and has a unique solution at a finite point. The solution of $Q(k, n)$ is denoted by $u_j^*(k, n)$ for $j = k, \ldots, n$, and the corresponding departure times are $x_j^*(k, n)$.

The optimal solution in (2) is denoted by u_i^*, $i = 1, \ldots, N$ and the corresponding departure times are x_i^*. It was shown in (Cassandras *et al.*, 2001) that the optimal solution is unique. The Forward Algorithm for obtaining this solution is based on the fact that x_i^* is given by $x_i^*(k, n)$ for some k, n as follows:

Theorem 1. (Cho *et al.*, 2001) Jobs $k, \ldots, n$ constitute a single busy period on the optimal sample path if and only if the following conditions are satisfied: (i) $a_k > x_{k-1}^*$, (ii) $x_i^*(k, i) \geq a_{i+1}$ for all $i = k, \ldots, n-1$, (iii) $x_n^*(k, n) < a_{n+1}$.

It follows from this theorem that $x_i^* = x_i^*(k, n)$ for all $i = k, \ldots, n$. In particular, the Forward Algorithm proceeds forward in time without the need for multiple forward-backward sweeps that are typically required to solve a two-point-boundary-value problem. Letting $k = 1$, $n = 1$, we first solve the linearly constrained convex optimization problem $Q(k, n)$ and obtain the control $u_j^*(k, n)$, $j = k, \ldots, n$ and departure times $x_j^*(k, n)$, $j = k, \ldots, n$. Then the structure of BPs is identified by checking if $x_n^*(k, n) < a_{n+1}$. If $C_k, \cdots, C_n$ are identified as a single BP, the optimal control is given by $u_j^* = u_j^*(k, n)$, $j = k, \ldots, n$. Then, the process is repeated for a new BP starting at a_{n+1}. This algorithm requires N steps.

3. RECEDING HORIZON CONTROL

Throughout the discussion above we have assumed the sequence $\{a_1, \ldots, a_N\}$ to be deterministic, i.e., a schedule of all job arrivals is known in advance. In what follows, we shall assume that knowledge of the future at time t is limited to a "window" $[t, t+T]$ for some given T. It is then natural to solve a sequence of problems of the form (2) replacing $(u_1, \ldots, u_N)$ by some $(u_i, \ldots, u_{i+r})$ where C_i is the next job whose processing time needs to be assigned and C_{i+r} is the last job known to arrive at some time $a_{i+r} \leq t+T$. The window is updated at every decision instant, i.e., upon departure of a job.

We distinguish the optimal controls u_i^* and corresponding departure times x_i^*, $i = 1, \ldots, N$, from those obtained through a control scheme limited in future knowledge by denoting the latter by $\tilde{u}_i$ and $\tilde{x}_i$ respectively. We shall also use the index t to represent the last job processed under such a controller, so that the current information window is $[\tilde{x}_t, \tilde{x}_t + T]$ and the *Receding Horizon* (RH) for the overall problem is $\tilde{x}_t + T$. We will also assume that any arrival time information provided at $\tilde{x}_t$ is "perfect" in the sense that both the optimal and the RH controller make decisions based on the same $\{a_i\}$ such that $\tilde{x}_t < a_i \leq \tilde{x}_t + T$. Assuming the current decision time to be $\tilde{x}_t$, the index of the last job contained within $[\tilde{x}_t, \tilde{x}_t + T]$ is given by $l = \arg\max_{r \geq t}\{a_r : a_r \leq \tilde{x}_t + T\}$ where we note that $l = t$ indicates that there are no arrival events in $[\tilde{x}_t, \tilde{x}_t + T]$. Similar to $Q(k, n)$ in (3), let us now consider a problem $\tilde{Q}(t+1, n)$ defined for a sample path generated by the RH controller when a decision time for job C_{t+1} comes up:

$$\tilde{Q}(t+1, n) : \min_{\tilde{u}_{t+1}, \ldots, \tilde{u}_n} \left\{ \begin{array}{c} \displaystyle\sum_{i=t+1}^{n} \{\theta_i(\tilde{u}_i) + \\ \psi_i[\max(\tilde{x}_t, a_{t+1}) + \\ \displaystyle\sum_{j=t+1}^{i} \tilde{u}_j]\} : \tilde{u}_i \geq 0 \end{array} \right\} \tag{4}$$

$$\text{s.t.} \quad \max(\tilde{x}_t, a_{t+1}) + \sum_{j=t+1}^{i} \tilde{u}_j \geq a_{i+1}, \quad i = t+1, \ldots, n-1$$

Setting $k = t+1$ in (3), the only difference between $Q(t+1, n)$ and $\tilde{Q}(t+1, n)$ lies in replacing a_{t+1} by $\max(\tilde{x}_t, a_{t+1})$. This is due to the fact that $Q(t+1, n)$ is *always* solved with the knowledge that C_{t+1} starts a BP. However, in the RH control scheme C_{t+1} does not necessarily start a BP. In particular, at time $\tilde{x}_t$ the controller can determine whether $a_{t+1} \leq \tilde{x}_t$, in which case C_{t+1} cannot start a BP. When this is true, $\tilde{Q}(t+1, n)$ above is solved with $\max(\tilde{x}_t, a_{t+1}) = \tilde{x}_t$ for $n = t+1, \ldots, l$ as if $\tilde{x}_t$ were initiating a BP and there are only $l - t$ jobs left to process. Since in this scheme all controls $\tilde{u}_j$, $j \leq t$ are already fixed at the time that $\tilde{u}_{t+1}$ is evaluated, our goal is to determine the "optimal" controls for the remainder of a BP or $\tilde{x}_t + T$, whichever comes first. The solution of $\tilde{Q}(t+1, n)$ is denoted by $\tilde{u}_i(t+1, n)$, $i = t+1, \ldots, n-1$. The detailed RH controller operation is described in (Cassandras and Mookherjee, 2003*a*).

4. RECEDING HORIZON (RH) CONTROLLER PROPERTIES

Our analysis of the RH controller properties, compared to the behavior obtained under the optimal control, is organized in two parts. First, we consider BPs (k, n) on the optimal sample path such that *no critical job is present* and establish properties of the RH controlled system for jobs indexed by t such that $k-1 \leq t < n$. As mentioned earlier, we have already investigated this case in (Cassandras and Mookherjee, 2003*a*). In the remainder of this paper, we target BPs on the optimal sample path that *do include critical jobs*. We establish the fact that the same properties still apply under additional conditions.

In the case where an optimal path BP (k, n) contains no critical jobs, we have shown in (Cassandras and Mookherjee, 2003*a*) that the RH controlled system will also not have any critical jobs over the range $k-1, \ldots, n-1$; regarding the nth job, it is possible that it becomes critical, i.e., $\tilde{x}_n = a_{n+1}$. Based on these facts, we have established the following two key properties of the RH controller: (i) $\tilde{x}_i \geq x_i^*$, $k \leq i \leq n$, and (ii) The error of job C_i, defined as $\varepsilon_i = \tilde{x}_i - x_i^*$, is monotonically decreasing when the condition $\tilde{x}_t + T \geq a_n$ holds. The implication of these properties is that the RH controller often incurs no error upon completion of some BPs and this error remains bounded. Formally, these properties are stated and proved in (Cassandras and Mookherjee, 2003*b*).

In the case where an optimal path BP (k, n) contains critical jobs, we can obtain several simple properties similar to the ones in (Cassandras and Mookherjee, 2003*a*). In addition, concentrating on the last block of an optimal path BP, assuming the presence of at least one critical job in this BP, we obtain the following two theorems.

Theorem 2. Let (k, n) define the **last block** of an optimal path BP. Let $\tilde{x}_t$, $k-1 \leq t < n$, be the current decision time on some BP of the RH sample path that ends with $C_{\tilde{n}}$. Then, $\tilde{x}_i \geq x_i^*$ for all $i = t+1, \ldots, \min(n, \tilde{n})$.

Theorem 3. Let (k, n) define the **last block** of an optimal path BP. Let $\tilde{x}_t$, $k-1 \leq t < n$, be the current decision time on some BP of the RH sample path and assume that $\tilde{x}_t + T \geq a_n$. Then, (i) If $t < k$ and C_k also starts a block in the RH path: $\varepsilon_i = 0$ for all $i = k, \ldots, n$, (ii) If $k \leq t < n$, $\varepsilon_{i+1} \leq \varepsilon_i$ for all $i = t+1, \ldots, n$ and if $\varepsilon_j = 0$ for some $j \in \{t+1, \ldots, n\}$, then $\varepsilon_i = 0$ for all $i \in \{j+1, \ldots, n\}$.

Note that the results in Theorems 2 and 3 apply only to the last block of an optimal path BP. In our earlier work, we were able to establish the same results for a complete optimal path BP containing no critical jobs. An obvious question then is whether similar results apply for blocks other than the last one in an optimal path BP. Generally, this is not the case, which implies that the error ε_i may in fact increase even if $\tilde{x}_t + T \geq a_n$ in a block (k, n); it is only when the last block is encountered that we can rely on error reduction. However, it turns out that modifying the nature of the cost function in (2) allows us to establish similar properties for all blocks. Thus, let us proceed by modifying Assumption **A2** to consider functions $\psi_i(\cdot)$ that are monotonically increasing. This is in fact a reasonable condition to impose from a practical standpoint, as discussed in (Zhang and Cassandras, Dec. 2001). Thus, we will replace Assumption **A2** by the following:

Assumption A3. For each $i = 1 = 1, \ldots, N$, $\psi_i(\cdot)$ is strictly convex, twice continuously differentiable, and monotonically increasing.

Clearly, all properties derived so far still hold under Assumption **A3** as well. In addition, the next Theorem is significant because it allows us to establish the fact that $\tilde{x}_i$ upper bounds x_i^* under certain conditions.

Theorem 4. Let (k, n) define an optimal path BP with at least one critical job. Let $\tilde{x}_t$, $k-1 \leq t < n$, be the current decision time on some BP of the RH sample path and assume that $\tilde{x}_t + T \geq a_n$. Under Assumption **A3**, if there are r critical jobs indexed by $B_1 < \ldots < B_r$ in the optimal path between C_{t+1} and C_n and none of these jobs is critical in the RH path, then $\tilde{x}_i \geq x_i^*$ for all $i = t+1, \ldots, n$, and $\tilde{x}_i > x_i^*$ for all $i = B_1, \ldots, B_r$.

Our final result deals with the *error reducing property* which comes into play under the same condition as that of Theorem 4.

Theorem 5. Let (k, n) define an optimal path BP with at least one critical job. Let $\tilde{x}_t$, $k-1 \leq t < n$, be the current decision time on some BP of the RH sample path and assume that $\tilde{x}_t + T \geq a_n$. Under Assumption **A3**, if there are r critical jobs indexed by $B_1 < \ldots < B_r$ in the optimal path between C_{t+1} and C_n and none of these jobs is critical in the RH path, then $\varepsilon_{i+1} \leq \varepsilon_i$ for all $i = t+2, \ldots, n$.

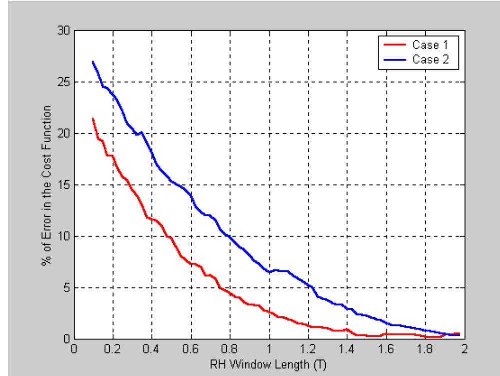

Fig. 1. Ensemble Average of % Error in Cost function vs Receding Horizon Window Length (T) over 25 samples

An interesting observation resulting from the theorem above is that as long as no critical jobs are observed on an unfolding RH sample path, we can conclude that the error reducing property holds.

5. NUMERICAL EXAMPLES

As an illustration of the performance of the RH controller compared to the Forward Algorithm (Cho *et al.*, 2001), known to provide the unique optimal solution of problem (2), we consider a problem with $N = 30$ jobs whose arrival times are uniformly distributed over $[0, 15]$. In this problem formulation we assume two different cost functions: In *Case* 1, $\psi_i(\cdot)$ is strictly convex and monotonically increasing, while in *Case* 2, $\psi_i(\cdot)$ is not monotonically increasing. In particular, in *Case* 1 : $\theta_i(u_i) = 1/u_i$ and $\psi_i(x_i) = (x_i - a_i)^2$ where $(x_i - a_i)$ is the *system time* of job C_i, and in *Case* 2 : $\theta_i(u_i) = 1/u_i$ and $\psi_i(x_i) = (x_i - d_i)^2$ where d_i is a specified deadline of job C_i. In this example, we set $d_i = a_i + 0.5$. Using the Forward Algorithm (Cho *et al.*, 2001) with all 30 jobs we have obtained the optimal sample path (u_i^* and x_i^* for all $i = 1, \ldots, 30$). For different values of the RH window size parameter T, we have also determined the sample path under RH control ($\tilde{u}_i(T)$ and $\tilde{x}_i(T)$ for all $i = 1, \ldots, 30$). We then investigate the effect of T on the error. To do so, we define the *% error in the cost function* for a given T and set of arrivals as the fraction [(RH cost −Optimal cost)/(Optimal cost)] $\times 100$. Averaging over 25 different realizations of the arrival process defined above and varying T, we have obtained Fig. 1. We observe that the error becomes negligible for T of about 1.5.

In order to visualize the error reducing property of ε_i, we have plotted the error ε_i for all $i = 1, \ldots, 100$ for different values of T, as shown in

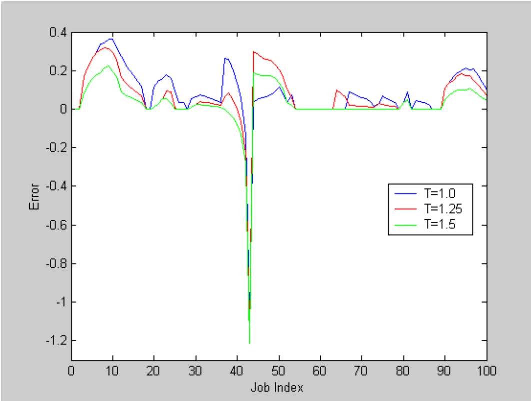

Fig. 2. Error $\varepsilon_i(T)$ over job index i for different T values under Case 1

Fig. 2 under *Case* 1 (similar results are obtained for *Case* 2). An arrival sequence is considered with 100 jobs whose arrival times are uniformly distributed over $[0, 45]$. Note that the error can become highly negative (however, this phenomenon is rare), but the error reducing property immediately takes effect and brings it back to 0. In addition, observe that substantial parts of the sample paths have zero error; in the case of $T = 1.5$, the majority of the sample path has in fact zero error.

6. FUTURE WORK

The receding horizon approach provides one way to deal with event uncertainties, when a controller has some limited look-ahead capability, without resorting to stochastic models and methods. We have assumed so far that the look-ahead event time information is perfect, a condition that may be relaxed if one is willing to estimate such event times. Ongoing research is addressing this possibility, as well as the obvious dependence of the accuracy of the RH controller on the window size T.

REFERENCES

Branicky, M. S., V. S. Borkar and S. K. Mitter (1998). A unified framework for hybrid control: Model and optimal control theory. *IEEE Trans. on Automatic Control* **43**(1), 31–45.

Cassandras, C. G. and R. Mookherjee (2003*a*). Receding horizon control for a class of hybrid systems with event uncertainties. In: *Proc. of 2003 American Control Conf.* To appear.

Cassandras, C. G. and R. Mookherjee (2003*b*). Receding horizon optimal control for some stochastic hybrid systems. *Technical Report, Dept. of Manuf. Eng., Boston University*. See also http://vita.bu.edu/cgc/Published /RHTechRep.pdf.

Cassandras, C. G. and S. Lafortune (1999). *Introduction to Discrete Event Systems*. Kluwer Academic Publishers.

Cassandras, C. G., D. L. Pepyne and Y. Wardi (2001). Optimal control of a class of hybrid systems. *IEEE Trans. on Automatic Control* **AC-46**(3), 398–415.

Cho, Y. C., C. G. Cassandras and D. L. Pepyne (2001). Forward decomposition algorithms for optimal control of a class of hybrid systems. *Intl. Journal of Robust and Nonlinear Control* **11**(5), 497–513.

Gokbayrak, K. and C. G. Cassandras (Dec. 1999). Stochastic optimal control of a hybrid manufacturing system model. In: *Proceedings of 38th IEEE Conf. On Decision and Control*. pp. 919–924.

Gokbayrak, K. and C. G. Cassandras (March 2000). Hybrid controllers for hierarchically decomposed systems. In: *Proceedings of 3rd Intl. Workshop on Hybrid Systems: Computation and Control*. pp. 117–129.

Hedlund, S. and A. Rantzer (1999). Optimal control of hybrid systems. *Proceedings of 38th IEEE Conf. On Decision and Control* pp. 3972–3977.

Mayne, D. and L. Michalska (1990). Receding horizon control of nonlinear systems. *IEEE Trans. on Automatic Control* **AC-35**(7), 814–824.

Pepyne, D. L. and C. G. Cassandras (1998). Modeling, analysis, and optimal control of a class of hybrid systems. *Journal of Discrete Event Dynamic Systems: Theory and Applications* **8**(2), 175–201.

Piccoli, B. (Dec. 1998). Hybrid systems and optimal control. In: *Proceedings of 37th IEEE Conf. On Decision and Control*. pp. 13–18.

Sussmann, H. J. (Dec. 1999). A maximum principle for hybrid optimal control problems. In: *Proceedings of 38th IEEE Conf. On Decision and Control*. pp. 425–430.

Xu, X. and P. J. Antsaklis (2000*b*). Optimal control of switched systems: New results and open problems. *Proceedings of the ACC* pp. 2683–2687.

Xu, X. and P. J. Antsaklis (Dec. 2000*a*). A dynamic programming approach for optimal control of switched systems. In: *Proceedings of 39th IEEE Conf. On Decision and Control*. pp. 1822–1827.

Zhang, P. and C. G. Cassandras (Dec. 2001). An improved forward algorithm for optimal control of a class of hybrid systems. In: *Proceedings of 40th IEEE Conf. On Decision and Control*. pp. 1235–1236.