To Measure or to Control: Optimal Control with Scheduled Measurements and Controls

Orhan C. Imer and Tamer Başar

***Abstract*— We consider a linear discrete-time optimal control problem where the controller has to choose between measurement and control. The controller is restricted in such a way that at a given time instance, it may either send a control packet to the actuator, or it can request a measurement packet from the sensor. We show that under a quadratic cost structure that does not penalize the control directly, the optimal control is a linear threshold policy, where the controller decides to measure or control by comparing its last measurement of the plant state against a pre-determined threshold.**

I. Introduction

Optimal quadratic control of discrete-time linear systems has been extensively studied in the literature [1]. In the standard formulation of the linear quadratic regulator (LQR) problem, the controller is assumed to be capable of measuring the state and controlling the plant *simultaneously*. In this paper, we consider the case where the controller has to choose between measurement and control. More specifically, the controller is restricted in such a way that at a given time instance, it may either send a control packet to the actuator, or it can request a measurement packet from the sensor.

Constraints of this type arise especially in networked control system applications. For instance, in controller area networks (CAN), the controller, sensor, and actuator nodes are connected over a serial bus, and only a single node can access it at a given time [2], [3], [4]; also see Figure 1. In these types of networks, the question is how to best *schedule* the measurement and control packets to achieve a desired system performance [5], [6]. In most industrial applications where CANs are used, the communication network is allocated before the runtime, and the resulting schedule is not deviated from. This type of scheduling is called *static*, and in the case of fixed intervals between control system events, the stability may be guaranteed by an appropriate choice of a *communication sequence* [7]. Communication sequences quantify the amount of "attention" that the decision maker pays to each component of a control system. In [8], the idea of communication sequencing is used to address the problem of stabilizing an LTI plant under limited access constraints. The discrete-time optimal control problem, where the network allows the transmission of one control signal at a time is considered in [9]. Earlier work in this field include [10], [11], where the focus is on picking the best measurement schedule, when we are constrained in looking at only one of the data signals available from the sensors at a time.

In this paper, we determine the optimal measurement and control policy of the controller under a quadratic cost structure that does not penalize the control directly. We show that the optimal policy is a linear threshold policy, where the controller decides to measure or control by comparing its last measurement of the plant state against a pre-determined threshold.

The rest of the paper is organized as follows. In Section II, we formally define the problem, and discuss the difference between open-loop and closed-loop scheduling policies. The solution is derived in Section III using a dynamic programming type argument. We present some numerical solutions in Section IV, and the paper ends with the concluding remarks of Section V where we also discuss some directions for future research.

II. Problem Statement

A. Problem definition

Consider the linear scalar plant described by the discrete-time dynamics

$$x_{k+1} = Ax_k + a_k u_k + w_k, \; k = 0, 1, \ldots, N-1 \quad (1)$$

where $x_k \in \mathbf{R}$ is the state, $u_k \in \mathbf{R}$ is the control, $a_k \in \{0, 1\}$ is the decision that represents whether we decide to apply a control or not, and $w_k \in \mathbf{R}$ is a zero-mean *i.i.d.* Gaussian process with finite variance, σ_w^2, describing the plant noise. In (1), N denotes the decision horizon, and the initial state x_0 is of zero mean and is characterized by its probability distribution P_{x_0}.

The controller is restricted in such a way that at a given time k, it may either send a control packet to the actuator, or it can request a measurement packet from the sensor. For example, the sensor and the actuator may be connected to the controller via a bus which does not allow simultaneous transmission of measurement and control packets; see Figure 1.

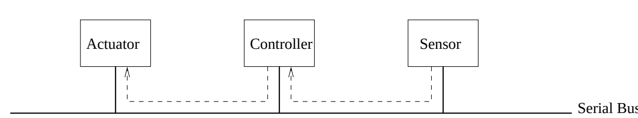

Fig. 1. Illustration of a Simple Control Bus Configuration.

More precisely, at the beginning of each period k the controller makes a decision as to whether to receive an observation y_k from the sensor, or send a control command u_k

First author is with GE Global Research, One Research Circle, Niskayuna, NY 12309. Second author is with Coordinated Science Laboratory, University of Illinois at Urbana-Champaign, 1308 W. Main Street, Urbana, IL 61801. imer@research.ge.com,tbasar@control.csl.uiuc.edu

Research supported by the NSF grant NSF CCR 00-85917

to the actuator. We assume that the observations are noiseless for ease of derivation, but the noisy observation case can be treated in a similar manner. The observation equation has the form:

$$y_k = (1-a_k)x_k, \; k = 0,1,\ldots,N-1$$

The noise process $\{w_k\}$ and the initial state x_0 are assumed to be independent. Let I_k denote the information available to the controller at time k. We have $I_0 = \emptyset$, and for $k = 1,2,\ldots,N-1$

$$I_k = \{(1-a_0)x_0, \ldots, (1-a_{k-1})x_{k-1}, u_0^{k-1}\}$$

Consider the class of policies consisting of a sequence of pairs of functions $\pi = \{(\alpha_0,\mu_0),(\alpha_1,\mu_1),\ldots,(\alpha_{N-1},\mu_{N-1})\}$, where each function α_k maps the information vector into the control subspace $\{0,1\}$, and each function μ_k maps the information vector I_k into the control subspace $\mathbf{R}$ unless $a_k = 0$. In other words, the policy pair $(\alpha_k(I_k),\mu_k(I_k)) = (a_k,u_k)$ is restricted to be mapped to either $(0,0)$ or $(1,u_k)$, where $u_k \in \mathbf{R}$. Such control policies are called *admissible*. We want to find an admissible policy π that minimizes the cost:

$$J_\pi = E\left\{x_N^2 + \sum_{k=0}^{N-1} x_k^2\right\} \tag{2}$$

subject to the system equation (1). Note that the decision policy at time k consists of two components $a_k = \alpha_k(I_k)$, $u_k = \mu_k(I_k)$ with the property that $(a_k,u_k) = (0,0)$, or $(a_k,u_k) = (1,u_k)$, $u_k \in \mathbf{R}$. At time k, $a_k = 0$ or $a_k = 1$, depending on whether we measure or control, respectively.

B. Open-loop versus closed-loop policies

Each given sequence $\{a_k\}_{k=0}^{N-1}$ of 0's and 1's corresponds to a *scheduling sequence* between the measurement and control packets. Now, given a particular scheduling sequence $\{a_k\}_{k=0}^{N-1}$, the optimal control is given by

$$u_k = \begin{cases} 0 & \text{if} \quad a_k = 0 \\ -AE\{x_k|I_k\} & \text{if} \quad a_k = 1 \end{cases} \tag{3}$$

since the cost function is separable. The optimal *open-loop* cost can then be calculated by an exhaustive search over all 2^N possible scheduling sequences. However, there is an easier way to obtain the optimum open-loop scheduling sequence. For this purpose, we first show two properties the optimal scheduling sequence should satisfy. Note that the information vector, I_k, evolves according to

$$I_{k+1} = \begin{cases} \{I_k, x_k, (0,0)\} & \text{if} \quad a_k = 0 \\ I_k & \text{if} \quad a_k = 1 \end{cases}$$

depending on whether we measure or control at time k.

Lemma 1: At time $k = 0$, the optimal scheduling sequence is such that $a_k = 0$.

Proof: Recall that at time 0, $I_0 = \emptyset$. In other words, we only know the *a priori* information about x_0 at time 0. Using (3), we see that $u_0 = 0$, regardless of a_0. Now, looking at the evolution of the information state, I_k, we see that

$$I_1 = \begin{cases} \{x_0, (0,0)\} & \text{if} \quad a_0 = 0 \\ \emptyset & \text{if} \quad a_0 = 1 \end{cases}$$

Since the control, u_0, is the same for both choices of a_0, we choose to measure at time 0 to increase our knowledge about the state at time $k = 1$, i.e. I_1. ■

Lemma 1 simply states that at time $k = 0$ we should measure. The next lemma clarifies when we should control.

Lemma 2: Suppose at time $k \in [1,N-1]$ we control, i.e. $a_k = 1$. Then, the cost we obtain by measuring at times $k-1$ and $k+1$ is at least as small as the cost obtained from any other decision policy.

Proof: If $a_k = 1$, the state will evolve according to

$$\begin{aligned} x_{k+1} &= A(x_k - E\{x_k|I_k\}) + w_k \\ x_{k+2} &= A[A(x_k - E\{x_k|I_k\})] + a_{k+1}u_{k+1} + Aw_k + w_{k+1} \end{aligned}$$

Note that u_{k+1} is still a function of I_k, since $I_{k+1} = I_k$ when $a_k = 1$. Now, we can either have $(a_{k+1},u_{k+1}) = (0,0)$, or $a_{k+1} = 1$ and

$$\begin{aligned} u_{k+1} &= -AE\{x_{k+1}|I_{k+1}\} = -AE\{x_{k+1}|I_k\} \\ &= -A^2E\{(x_k - E\{x_k|I_k\})|I_k\} = 0 \end{aligned}$$

Since in both cases we have $u_{k+1} = 0$, similar to Lemma 1, there is no need to control at time $k+1$.

The proof of why we must measure at time $k-1$ follows from a similar argument. If we are given that $a_k = 1$, then we must have had $a_{k-1} = 0$, otherwise we would not have had $a_k = 1$ in the first place by our previous argument. ■

Lemma 2 implies that if $\{a_k\}_{k=0}^{N-1}$ is the optimal scheduling sequence, and if for some $k = k_0$, $a_{k_0} = 1$, then we must have $a_{k_0-1} = a_{k_0+1} = 0$. That is, if we control at time k_0, we must have measured at time $k_0 - 1$, and we should measure at time $k_0 + 1$. A consequence of Lemmas 1 and 2 is that the optimal control u_k must have the form $u_k = -A^2x_{k-1}$ for those $k \in [1,N-1]$ we choose to control. Figure 2 illustrates the possible evolution of the measure (M) or control (C) sequence for a time horizon of length N.

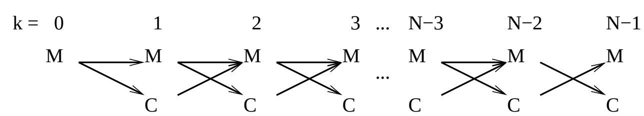

Fig. 2. Possible sample paths of the measure (M) - control (C) sequence.

Equipped with these two Lemmas, we now proceed to show that the best open-loop scheduling sequence is to start with measuring, M, and control afterwards, and repeat the same pattern, i.e. MCMCMC….

By Lemma 1, we know that $a_0 = 0$. Also, we know that if for some k, $u_k = -A^2x_{k-1}$, then $u_{k-1} = u_{k+1} = 0$. Therefore, the state, x_k, will evolve according to

$$x_{k+1} = \begin{cases} Ax_k + w_k & \text{if} \quad a_k = 0 \\ Aw_{k-1} + w_k & \text{if} \quad a_k = 1 \end{cases}$$

and the mean-square value of the state will evolve according to

$$E\{x_{k+1}^2\} = \begin{cases} A^2E\{x_k^2\} + \sigma_w^2 & \text{if} \quad a_k = 0 \\ A^2\sigma_w^2 + \sigma_w^2 & \text{if} \quad a_k = 1 \end{cases} \tag{4}$$

for $k \geq 1$. For $k = 0$, we have $x_1 = Ax_0 + w_0$ and $E\{x_1^2\} = A^2\sigma_{x_0}^2 + \sigma_w^2$,where $\sigma_{x_0}^2$ is the variance of the initial state x_0.

We can see from (4) that the mean-square cost $E\{x_{k+1}^2\}$ is always smaller when $a_k = 1$, since for any k $E\{x_k^2\} \geq \sigma_w^2$. Therefore, we should pick $a_k = 1$, if we can. Note that, the evolution of the mean-square error (4) relies on the fact that the scheduling policy is in compliance with Lemmas 1 and 2. Therefore, the optimum open-loop scheduling policy is to control as soon as we can, which can only be done after we measure. Since we measure at time $k = 0$, we should control at time $k = 1$, i.e. $a_1 = 1$. Now, at time $k = 2$ we cannot control by Lemma 2, so we measure, i.e. $a_2 = 0$. Since, we can now control at time $k = 3$, we set $a_3 = 1$, and so on. The resulting sequence starts with measure, (M), and alternates between measure, (M), and control, (C), until the end of the decision-horizon. Note that the optimum open-loop scheduling sequence is independent of the time-horizon, N. Thus, for an infinite horizon problem with the average cost criterion

$$\bar{J} = \limsup_{N\to\infty} \frac{1}{N} E\left\{\sum_{k=0}^{N-1} x_k^2\right\} \tag{5}$$

the solution will remain the same. We can actually calculate the optimum average open-loop cost corresponding to this policy as follows. It is clear that with the policy of alternating measure, (M), and control, (C), that starts with measure, (M), at time $k = 0$, the state will evolve according to

$$x_{k+1} = \begin{cases} Aw_{k-1} + w_k & \text{if} \quad k = 1,3,5,\ldots \\ A^2 w_{k-2} + Aw_{k-1} + w_k & \text{if} \quad k = 2,4,6,\ldots \end{cases}$$

and $x_1 = Ax_0 + w_0$. Now the optimum average cost, $\bar{J}^*$, can be calculated by averaging over the sum of mean-square value of the state:

$$\begin{aligned} \bar{J}^* &= \left(\frac{1}{2}(1+A^2) + \frac{1}{2}(1+A^2+A^4)\right)\sigma_w^2 \\ &= \left(1 + A^2 + \frac{1}{2}A^4\right)\sigma_w^2 \end{aligned}$$

In the next section, we will show that the optimum closed-loop decision policy can beat this cost by delaying the decision to measure or control, a_k, until time k.

III. Derivation of the Solution

A. The finite-horizon solution

In order to derive the optimal decision policy, we use a dynamic programming approach. Starting at time $N-1$, we go back in time and find a recursive way to calculate the cost-to-go functions. Now, it can be seen that the optimum cost-to-go from time $k+1$, J_{k+1}, will be of the form:

$$J_{k+1} = \begin{cases} J_{k+1}^{(0)} & \text{if} \quad a_k = 0 \\ J_{k+1}^{(1)} & \text{if} \quad a_k = 1 \end{cases}$$

where $J_{k+1}^{(0)}$, and $J_{k+1}^{(1)}$ are given as

$$\begin{aligned} J_{k+1}^{(0)} &= E\{(Ax_k + w_k)^2 | x_k\} \\ &\quad + \min\left\{E\left\{\Gamma_{k+1}^{(0)}(A(Ax_k + w_k) + w_{k+1}) | x_k\right\}\right. \\ &\qquad \left., E\left\{\Gamma_{k+1}^{(1)}(Aw_k + w_{k+1})\right\}\right\} \\ J_{k+1}^{(1)} &= E\{(Aw_{k-1} + w_k)^2\} \\ &\quad + E\{\Lambda_{k+1}(A(Aw_{k-1} + w_k) + w_{k+1})\} \end{aligned}$$

Here, $\Gamma_{k+1}^{(0)}$, $\Gamma_{k+1}^{(1)}$ and Λ_{k+1} are functions that are initialized as $\Gamma_{N-1}^{(0)}(u) = \Gamma_{N-1}^{(1)}(u) = \Lambda_{N-1}(u) = u^2$. Note that the cost-to-go function, J_{k+1}, depends on our past actions only through a_k. When $a_k = 0$, $J_{k+1}^{(0)}$ contains a term which is the minimum of two numbers. This minimization is between choosing the two alternatives, namely measure, (M), or control, (C), at time $k+1$. More specifically, the first term in the minimum is realized if $a_{k+1} = 0$, (M), and the second term is realized if $a_{k+1} = 1$, (C). This minimum can be found by the comparison:

$$E\{\Gamma_{k+1}^{(0)}(A(Ax_k + w_k) + w_{k+1}) | x_k\} \lesseqgtr E\{\Gamma_{k+1}^{(1)}(Aw_k + w_{k+1})\}$$

Later, we will show that this comparison gives us a threshold on the previous state, x_k, which has been measured since $a_k = 0$. The threshold, τ_k, is such that if $|x_k| \leq \tau_k$ the minimum will be given by the first term, i.e. we should measure (M) at time $k+1$, otherwise the minimum will be given by the second term, i.e. we should control (C) at time $k+1$. The proof of this fact relies on the fact that as a function of x_k, the conditional expectation $E\{\Gamma_{k+1}^{(0)}(A(Ax_k + w_k) + w_{k+1}) | x_k\}$ achieves its global minimum at $x_k = 0$. Also note that if $a_k = 1$, the previous state, x_k, must have been of the form $x_k = Aw_k + w_{k-1}$, which is independent of the past states, x_0^{k-1}. Therefore, when $a_k = 1$, the expectations do not involve any conditioning on the past. We proceed by finding the functional recursions $\Gamma_{k+1}^{(0)}$, $\Gamma_{k+1}^{(1)}$, and Λ_{k+1} must satisfy. First, we write J_k as

$$J_k = \begin{cases} J_k^{(0)} & \text{if} \quad a_{k-1} = 0 \\ J_k^{(1)} & \text{if} \quad a_{k-1} = 1 \end{cases}$$

After some algebra, we can obtain $J_k^{(0)}$ and $J_k^{(1)}$ in terms of $\Gamma_{k+1}^{(0)}$, $\Gamma_{k+1}^{(1)}$, and Λ_{k+1} as follows:

$$\begin{aligned} J_k^{(0)} &= E\{(Ax_{k-1} + w_{k-1})^2 | x_{k-1}\} \\ &\quad + \min\{E\{(A(Ax_{k-1} + w_{k-1}) + w_k)^2 | x_{k-1}\} \\ &\quad + \int_{|x_k| \leq \tau_k} E\{\Gamma_{k+1}^{(0)}(A(Ax_k + w_k) + w_{k+1}) | x_k\} \\ &\qquad \times \frac{e^{-\frac{(x_k - Ax_{k-1})^2}{2\sigma_w^2}}}{\sqrt{2\pi\sigma_w^2}} dx_k \\ &\quad + \int_{|x_k| > \tau_k} E\{\Gamma_{k+1}^{(1)}(Aw_k + w_{k+1})\} \frac{e^{-\frac{(x_k - Ax_{k-1})^2}{2\sigma_w^2}}}{\sqrt{2\pi\sigma_w^2}} dx_k \\ &\quad , E\{(Aw_{k-1} + w_k)^2\} \\ &\quad + E\{\Lambda_{k+1}(A(Aw_{k-1} + w_k) + w_{k+1})\}\} \end{aligned}$$

$$
\begin{aligned}
J_k^{(1)} &= E\{(Aw_{k-2}+w_{k-1})^2\} \\
&+\int_{-\infty}^{+\infty} E\{(Ax_k+w_k)^2|x_k\}\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k \\
&+E\{\Gamma_{k+1}^{(0)}(A(Ax_k+w_k)+w_{k+1})|x_k\} \\
&\qquad\times\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k \\
&+E\{\Gamma_{k+1}^{(1)}(Aw_k+w_{k+1})\}\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k
\end{aligned}
$$

Let

$$
\begin{aligned}
\overline{\Gamma}_{k+1}^{(0)}(x_k) &:= E_{w_{k+1},x_{k+1}}\{\Gamma_{k+1}^{(0)}(Ax_{k+1}+w_{k+1})|x_k\} \\
\overline{\Gamma}_{k+1}^{(1)} &:= E_{w_k,w_{k+1}}\{\Gamma_{k+1}^{(1)}(Aw_k+w_{k+1})\} \\
\overline{\Lambda}_{k+1} &:= E_{w_{k-1},w_k,w_{k+1}}\{\Lambda_{k+1}(A(Aw_{k-1}+w_k) \\
&\qquad +w_{k+1})\}
\end{aligned}
$$

where $x_{k+1}\sim N(Ax_k,\sigma_w^2)$ in the definition of $\overline{\Gamma}_{k+1}^{(0)}$. Then, $\overline{\Gamma}_k^{(0)}(x_{k-1})$, $\overline{\Gamma}_k^{(1)}$, and $\overline{\Lambda}_k$ can be written as

$$
\begin{aligned}
\overline{\Gamma}_k^{(0)}(x_{k-1}) &= E_{x_k,w_k}\{(Ax_k+w_k)^2|x_{k-1}\} \\
&+\int_{|x_k|\le\tau_k}\overline{\Gamma}_{k+1}^{(0)}(x_k)\frac{e^{-\frac{(x_k-Ax_{k-1})^2}{2\sigma_w^2}}}{\sqrt{2\pi\sigma_w^2}}dx_k \\
&+\int_{|x_k|>\tau_k}\overline{\Gamma}_{k+1}^{(1)}\frac{e^{-\frac{(x_k-Ax_{k-1})^2}{2\sigma_w^2}}}{\sqrt{2\pi\sigma_w^2}}dx_k
\end{aligned}
$$

where x_k in the first term is $x_k\sim N(Ax_{k-1},\sigma_w^2)$.

$$
\begin{aligned}
\overline{\Gamma}_k^{(1)} &= E\{(Aw_{k-1}+w_k)^2\}+\overline{\Lambda}_{k+1} \\
\overline{\Lambda}_k &= E\{(Ax_k+w_k)^2\} \\
&+\int_{|x_k|\le\tau_k}\overline{\Gamma}_{k+1}^{(0)}(x_k)\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k \\
&+\int_{|x_k|>\tau_k}\overline{\Gamma}_{k+1}^{(1)}\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k
\end{aligned}
$$

where x_k in the first term is $x_k\sim N(0,(1+A^2)\sigma_w^2)$.

Simplifying, we obtain

$$
\begin{aligned}
\overline{\Gamma}_k^{(0)}(x_{k-1}) &= A^4x_{k-1}^2+(1+A^2)\sigma_w^2 \\
&+\int_{|x_k|\le\tau_k}\overline{\Gamma}_{k+1}^{(0)}(x_k)\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k \\
&+\int_{|x_k|>\tau_k}\overline{\Gamma}_{k+1}^{(1)}\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k \\
\overline{\Gamma}_k^{(1)} &= (1+A^2)\sigma_w^2+\overline{\Lambda}_{k+1}
\end{aligned}
$$

$$
\begin{aligned}
\overline{\Lambda}_k &= A^4\sigma_w^2+(1+A^2)\sigma_w^2 \\
&+\int_{|x_k|\le\tau_k}\overline{\Gamma}_{k+1}^{(0)}(x_k)\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k \\
&+\int_{|x_k|>\tau_k}\overline{\Gamma}_{k+1}^{(1)}\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2\pi(1+A^2)\sigma_w^2}}dx_k
\end{aligned}
$$

Note that the cost-to-go from time k, J_k, can be written as

$$
J_k=\begin{cases} A^2x_{k-1}^2+\sigma_w^2+\min\left\{\overline{\Gamma}_k^{(0)}(x_{k-1}),\overline{\Gamma}_k^{(1)}\right\} & ,a_{k-1}=0 \\ A^2\sigma_w^2+\sigma_w^2+\overline{\Lambda}_k & ,a_{k-1}=1\end{cases}
$$

Thus, the decision to choose $a_k=0$ or $a_k=1$, can be determined by checking $\overline{\Gamma}_k^{(0)}(x_{k-1})\lesseqgtr\overline{\Gamma}_k^{(1)}$, which gives a threshold, τ_{k-1}, on the value of the state, x_{k-1}, which we have access to since $a_{k-1}=0$. The threshold, τ_{k-1}, determines whether we should measure or control at time k depending on the value of the measurement, x_{k-1}. More specifically,

$$
a_k=\begin{cases}0 & \text{if } |x_{k-1}|\le\tau_{k-1} \\ 1 & \text{if } |x_{k-1}|>\tau_{k-1}\end{cases}
$$

The existence of such a threshold is the consequence of the fact that the function $\Delta_k^{(0)}(x_{k-1}):=\overline{\Gamma}_k^{(0)}(x_{k-1})-\overline{\Gamma}_k^{(1)}$ is even, and it achieves its unique global minimum at $x_{k-1}=0$. Later, we will show inductively that $\Delta_k^{(0)}(x_{k-1})$'s have these properties for all $k\le N-1$. Furthermore, we will show that the minimum value of $\Delta_k^{(0)}(x_{k-1})$, which is achieved at $x_{k-1}=0$, has the property that $\Delta_k^{(0)}(0)\le 0$. Therefore, there exists a threshold, $\tau_k>0$, which is characterized by the solution of the nonlinear equation $\Delta_k^{(0)}(\pm\tau_k)=0$. For ease of notation, let us also introduce a new sequence $\Delta_k^{(1)}:=\overline{\Lambda}_k-\overline{\Gamma}_k^{(1)}$. Using the recursions for $\overline{\Gamma}_k^{(0)}(x_{k-1})$, $\overline{\Gamma}_k^{(1)}$, and $\overline{\Lambda}_k$, we can derive the recursive relations $\Delta_k^{(0)}(x_{k-1})$ and $\Delta_k^{(1)}$ must satisfy as follows:

$$
\begin{aligned}
\Delta_k^{(0)}(x_{k-1}) &= A^4x_{k-1}^2-\Delta_{k+1}^{(1)} \\
&+\int_{|x_k|\le\tau_k}\Delta_{k+1}^{(0)}(x_k)\frac{e^{-\frac{(x_k-Ax_{k-1})^2}{2\sigma_w^2}}}{\sqrt{2\pi\sigma_w^2}}dx_k \\
\Delta_k^{(1)} &= A^4\sigma_w^2-\Delta_{k+1}^{(1)} \\
&+\int_{|x_k|\le\tau_k}\Delta_{k+1}^{(0)}(x_k)\frac{e^{-\frac{x_k^2}{2(1+A^2)\sigma_w^2}}}{\sqrt{2(1+A^2)\pi\sigma_w^2}}dx_k
\end{aligned}
$$

along with the boundary conditions

$$
\begin{aligned}
\Delta_{N-1}^{(0)}(x_{N-2}) &= A^4x_{N-2}^2 \\
\Delta_{N-1}^{(1)} &= A^4\sigma_w^2 \\
\tau_{N-2} &= 0
\end{aligned}
$$

This leads to the following theorem whose proof can be found in [2].

Theorem 1: Let $N\ge 2$ be given. For $0\le k\le N-1$, the sequence of functions $\Delta_k^{(0)}(u)$ are even, differentiable with a unique critical point at $u = 0$, i.e.

$$\left.\frac{\partial \Delta_k^{(0)}(u)}{\partial u}\right|_{u=0} = 0$$

Furthermore, we have

$$\frac{\partial \Delta_k^{(0)}(u)}{\partial u} > 0, \quad \text{if} \quad u > 0$$
$$\frac{\partial \Delta_k^{(0)}(u)}{\partial u} < 0, \quad \text{if} \quad u < 0$$

Thus, $\Delta_k^{(0)}(u)$ achieves its global minimum at the critical point $u = 0$. Also, the minimum value of $\Delta_k^{(0)}(u)$ achieved at $u = 0$ is non-positive, i.e. $\Delta_k^{(0)}(0) \leq 0$.

In order to illustrate the way the recursions $\Delta_k^{(0)}(x_{k-1})$, $\Delta_k^{(1)}$ can be used to determine the thresholds, τ_k, we find τ_{N-3}. Since $\tau_{N-2} = 0$, we do not need to carry out any integration. We simply have

$$\begin{aligned} \Delta_{N-2}^{(0)}(x_{N-3}) &= A^4 x_{N-3}^2 - A^4\sigma_w^2 = A^4(x_{N-3}^2 - \sigma_w^2) \\ \Delta_{N-2}^{(1)} &= A^4\sigma_w^2 - A^4\sigma_w^2 = 0 \end{aligned}$$

The threshold, τ_{N-3}, can be determined from the positive solution of $\Delta_{N-2}^{(0)}(\tau_{N-3}) = 0$, which is given by $\tau_{N-3} = \sigma_w$. The remaining thresholds can be calculated recursively by using numerical integration. In Section IV, we present some simulation results where we calculate the thresholds as we vary various parameters.

To summarize, the optimal closed-loop decision policy is a threshold policy in the sense that at time k, we decide to measure or control based on our measurement at the previous time instance, $k-1$. The choice at time k between measurement and control can only be made when at time $k-1$ we have measured, since otherwise we must measure at time k. Furthermore, at time $k = 0$, we must start with measuring by Lemma 1. If we need to choose between measurement and control, we compare the measurement x_{k-1} against a threshold value, τ_{k-1}, which can be determined offline by the recursive methodology outlined in this section. Algorithmically, we can describe the optimal decision policy as follows.

At time $k = 0$, start by measuring (M), i.e. $a_0 = 0$. Starting with $k = 1$, do the following:

1) Lookup the threshold τ_{k-1} corresponding to the current stage from the table.
2) Apply the measurement/control policy

$$(a_k, u_k) = \begin{cases} (0,0) & , a_{k-1} = 0, \ |x_{k-1}| \leq \tau_{k-1} \\ (1, -A^2 x_{k-1}) & , a_{k-1} = 0, \ |x_{k-1}| > \tau_{k-1} \\ (0,0) & , a_{k-1} = 1 \end{cases}$$

3) Update $k = k+1$, and go to the first step unless $k = N$.

B. The limiting solution

In this section, we investigate the limiting behavior of the thresholds, τ_k, as the number of stages in the problem, N, becomes arbitrarily large. Note that, in this case the cost function (2) must be replaced with the average cost function (5) to avoid an infinite cost as $N \to \infty$.

Finding an analytical expression for the limit of the sequence $\{\tau_k\}$ is difficult. As we show in the next section, the convergence of τ_k to its limit is not monotonic, which makes the analytical calculation of the limit as a function of A, and σ_w^2 very difficult. In the next section, we will calculate this limit for some select values of A and σ_w^2, and discuss how it changes as these parameters are varied.

The purpose of this section is to show sort of a negative result in the sense that the limit of the sequence $\{\tau_k\}$, if it exists, cannot be zero. This is important, because otherwise, for an infinite-horizon problem with the average cost criterion, the optimal closed-loop scheduling sequence in the limit would have coincided with the optimal open-loop scheduling sequence, where essentially $\tau_k = 0$, for all k. We have the following proposition whose proof can be found in [2].

Proposition 1: Suppose the limit

$$\tau_\infty = \lim_{k\to\infty} \tau_k$$

exists. Then, $\tau_\infty \neq 0$.

IV. Numerical Solutions

In this section, we use numerical integration and Matlab to calculate the recursions of $\Delta_k^{(0)}(x_{k-1})$ and $\Delta_k^{(1)}$, with the aim of determining the sequence of thresholds $\{\tau_k\}$. We investigate how the threshold sequence behaves as the plant parameter, A, and the noise variance, σ_w^2, are varied. As we mentioned before, the threshold sequence converges to some value τ_∞ regardless of how large or small A and σ_w^2 are. The plant parameter, A, affects the rate of convergence while the noise variance σ_w^2 affects the limiting value of the thresholds. As A gets larger the convergence becomes slower, and the convergence is faster as $A \to 0$. For a fixed A, the limiting value of the threshold, τ_∞, scales linearly with the standard deviation of the noise process, σ_w.

In Figure 3 we demonstrate how the sequence of thresholds converges for a fixed value of A, as σ_w is varied. The limiting values in Figure 3 are multiples of $\tau_\infty(A=1) \approx 0.7767$. That is

$$\begin{aligned} \sigma_w = 1 \quad &\Rightarrow \quad \tau_\infty = \tau_\infty(A=1) \\ \sigma_w = 2 \quad &\Rightarrow \quad \tau_\infty = 2 \times \tau_\infty(A=1) \\ \sigma_w = 3 \quad &\Rightarrow \quad \tau_\infty = 3 \times \tau_\infty(A=1) \end{aligned}$$

This pattern is repeated for other, but fixed values of A. Note that, in the first iteration $\tau_1 = \sigma_w$, as predicted by our example calculation in Section III.

Also note that the convergence is not monotonic, which turns out to be the case regardless of what A or σ_w^2 is. We next fix the noise variance $\sigma_w^2 = 1$, and vary the plant parameter A. Figure 4 illustrates how the convergence rate of τ_k is affected by the variation of the plant constant A. In Figure 4, A is varied from 0.5 to 2. Note the decrease in convergence rate as A is increased. Also, observe that the limiting value of the threshold, τ_∞, is insensitive to the plant parameter A.

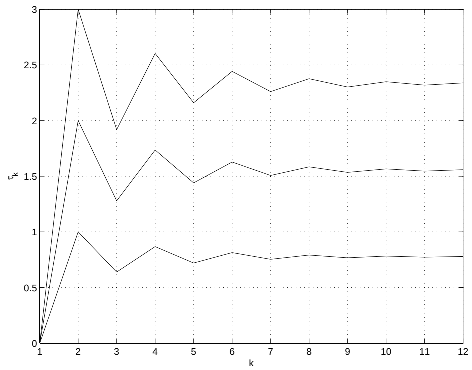

Fig. 3. Convergence of thresholds, τ_k, for $A = 1$, and $\sigma_w = 1, 2, 3$.

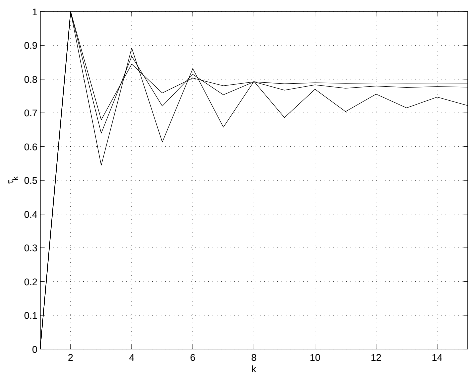

Fig. 4. Convergence of thresholds, τ_k, for $\sigma_w^2 = 1$, and $A = 0.5, 1, 2$.

We also investigate the extreme cases, as $A \to 0$, and $A \to \infty$. Increasing A indefinitely makes the convergence rate infinitely small, but τ_k still converges. Making A smaller, on the other hand, makes τ_k converge faster. For $\sigma_w^2 = 1$, in the limit as $A \to 0$, τ_∞ goes to approximately 0.7894. Therefore, if $\sigma_w^2 = 1$, for the entire range of plant parameters, A, the limiting value of threshold, τ_∞, lies in the interval $[0, 1]$.

Having derived the optimal values of the thresholds for various plants, we next turn our attention to comparing the optimal average cost that can be achieved with a closed-loop control with the optimal open-loop cost. Recall from Section II-B that the optimal open-loop scheduling policy is to start with measuring, and alternate between measure (M) and control (C) indefinitely. As derived in Section II-B, this policy results in an average cost:

$$\bar{J}_{OL}^* = \left(1 + A^2 + \frac{1}{2}A^4\right)\sigma_w^2 \qquad (6)$$

Now, if we implement the optimal closed-loop controller, we expect the average cost to be smaller, but exactly how much smaller? Can we in effect quantify the gain achieved by optimal closed-loop scheduling of measurement and control packets? For this purpose, we basically run several instances of the plant process with optimal open-loop and closed-loop controls, and calculate the sample-path average of the performance criterion. We set $A = 1$ and $\sigma_w^2 = 1$, and average the sample path costs over 100 instances of the plant process. The results are tabulated in Table I, where $\bar{J}_{CL}^*$ is the long-run average optimal closed-loop cost, $\bar{J}_{OL(e)}^*$ denotes the long-run average open-loop cost, $\bar{J}_{OL(t)}^*$ is the theoretical value of the open-loop cost obtained from (6), and % column indicates the percentage improvement in the long-run average cost of the closed-loop control over the open-loop one. As can bee seen from Table I, the percentage decrease in the average cost in going from the open-loop scheduling to closed-loop scheduling is in the range $0.5 - 7\%$. This may not seem like a huge improvement; however, note that the amount of online computation we need to do to achieve this gain is minimal. For large N, since the threshold sequences, $\{\tau_k\}$, converge, all we need to do is compare the last measured value of the state against a constant threshold, τ_∞.

TABLE I

COMPARISON OF AVERAGE SAMPLE-PATH COSTS.

A	σ_w	$\bar{J}_{CL}^*$	$\bar{J}_{OL(e)}^*$	$\bar{J}_{OL(t)}^*$	%
0.5	1	1.2753	1.2827	1.2813	0.5790
1	2	9.6073	9.9822	10.000	3.7556
1	1	2.4040	2.5015	2.5000	3.8946
2	1	12.0999	13.0001	13.0000	6.9246
4	1	137.3523	144.8553	145	5.1797

V. CONCLUSIONS

In this paper, we introduced the optimal control problem with scheduled measurements and controls, where the controller has to choose between controlling the state now and bringing it closer to the desired set point, and making a measurement to improve its estimate of the next state. We showed that the solution of this problem is a linear threshold policy, where the controller decides to measure or control by comparing its last measurement of the plant state against a pre-determined threshold.

REFERENCES

[1] D. P. Bertsekas, *Dynamic Programming and Optimal Control*. Belmont, MA: Athena Scientific, 1995.

[2] O. C. Imer, "Optimal estimation and control under communication network constraints," *Ph.D. Dissertation*, University of Illinois, 2005.

[3] F.-L. Lian, J. R. Moyne, and D. M. Tilbury, "Performance evaluation of control networks: Ethernet, ControlNet, and DeviceNet," *IEEE Control Systems Magazine*, vol. 21, pp. 66–83, February 2001.

[4] K. M. Zuberi and K. G. Shin, "Design and implementation of efficient message scheduling for controller area network," *IEEE Transactions on Computers*, vol. 49, no. 2, pp. 182–188, February 2000.

[5] G. C. Walsh and H. Ye, "Scheduling of networked control systems," *IEEE Control Systems Magazine*, vol. 21, pp. 57–65, February 2001.

[6] D. Hristu and P. R. Kumar, "Interrupt-based feedback control over a shared communication medium," in *Proc. IEEE Conference on Decision and Control (CDC)*, December 2002, pp. 1787–1792.

[7] R. Brockett "Stabilization of motor networks,' in *Proc. IEEE Conference on Decision and Control (CDC)*, December 1995, pp. 1484–1488.

[8] D. Hristu "Stabilization of LTI systems with communication constraints," in *Proc. ACC*, June 2000, pp. 2342–2346.

[9] H. Rehbinder and M. Sanfridson, "Scheduling a limited communication channel for optimal control," in *Proc. IEEE Conference on Decision and Control (CDC)*, December 2000, pp. 1011–1016.

[10] M. Athans, "On the determination of optimal costly measurement strategies for linear stochastic systems," *Automatica*, vol. 8, pp. 397–412, 1972.

[11] R. K. Mehra "Optimization of measurement schedules and sensor designs for linear dynamic systems," *IEEE Transactions on Automatic Control*, vol. 21, no. 1, pp. 55–64, February 1976.