Appl. Comput. Math. **4** (2005), no.2, pp.

OPTIMAL ESTIMATION WITH SCHEDULED MEASUREMENTS*

ORHAN C. IMER †, TAMER BAŞAR †,§

Abstract. We consider the problem of scheduling one of several sensor measurements in a communication bus for estimation purposes. The sensors prioritize their measurements by assigning priority numbers to each measurement they make, and the bus access is granted using a decentralized arbitration mechanism. In the case of two sensors observing independent discrete, Gaussian, or uniform random variables, we show that the optimum scheduling policy for each sensor is a threshold policy where the thresholds depend on the *a priori* distribution of the sensor measurements.
Key words: Communication networks; sensor networks; optimum scheduling; estimation

Introduction

In this paper, we investigate the trade-off between sending one of several measurements when components of a networked control system are connected over a communication network which does not allow simultaneous transmission of measurement packets. This type of a network constraint arises in controller area networks (CAN), where multiple sensor nodes are connected over a serial bus, and only a single node can access it at any given time [7],[10], [13],[15]. In these types of networks, the question is how to best *schedule* the measurement packets to achieve a desired system performance [8], [14].

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* [6]. Communication sequences quantify the amount of "attention" that the decision maker pays to each component of a control system. In [1] 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 [2]. Earlier work in this field include [3], [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 our setup, we have multiple sensors observing different components of a state vector. All sensors are connected via a serial bus to a central monitoring station. The monitoring station wants to estimate the state of an underlying stochastic process in real time. The sensors communicate their measurements to the monitoring station by accessing the communication bus one at a time. The problem is to determine the scheduling conditions for each sensor so as to minimize the state estimation error at the monitoring station. The sensors make their decisions *decentrally.* The decision as to who will access the communication bus is determined

*This paper is based partially on a plenary talk given by the secoind author at the First International Conference on Control and Optimization with Industrial Applications (COIA-2005), in Baku, Azerbaijan, May 22-25, 2005.

†First author was with Coordinated Science Laboratory, University of Illinois at Urbana-Champaign; he is now with GE Global Research, One Research Circle, Niskayuna, NY 12309. Second author is with Coordi nated Science Laboratory, University of Illinois at Urbana-Champaign, 1308 W. Main Street, Urbana, IL 61801. (imer,tbasar)@control.csl.uiuc.edu .

Research supported by the NSF grant NSF CCR 00-85917
§*Manuscript received October 27, 2005.*

by an arbitration mechanism in which at the beginning of each decision epoch, each sensor communicates a priority number to a central arbiter one bit at a time. The arbitration unit is a binary function that maps the bits it receives from the sensors into 0 or 1. The result of this map is returned to each sensor, and at the end of the arbitration period, the sensor with the highest priority gains access to the bus. Thus, the scheduling problem as to who should access the bus boils down to determining how each sensor should match the set of its potential observations into a priority number.

We solve this problem for the special case when there are exactly two sensors, and each sensor observes a random variable independent from the one observed by the other one. We investigate the three cases when the random variables are discrete, Gaussian, and uniform. Note that, this is a special case of the general problem, where each sensor observes an i.i.d. random sequence which is also independent from the one observed by the other sensor. The case when the sensors observe random processes which are correlated both across time and with each other is an open problem.

The rest of this paper is organized as follows. In Section 1, we formally define the problem for the multi-sensor and two-sensor cases. The solution for the special case of two sensors observing independent random variables is derived in Section 2 for discrete, Gaussian, and uniformly distributed sensor observations. The paper ends with the concluding remarks of Section 3.

1. Problem definition

In this section, we formally define the problem of scheduling measurements of several sensors that are connected via a communication bus to a monitoring station. In a communication bus, the way in which sensors are granted access is specified by the so-called medium access control (MAC) layer of the underlying network protocol. Not all MAC layer protocols allow for a design of implementable scheduling policies for measurement packets. For example, if Ethernet is used for MAC, due to its random access nature [9], [12], it is not possible to design a measurement schedule that can be deterministically adhered to by the network.

There are MAC layer protocols, however, that allow for dynamic arbitration of packets with different priorities. DeviceNET [5] is one such example, where the MAC layer protocol is of type CSMA/BA (carrier sense multiple access with bit-wise arbitration). In CSMA/BA, each packet has an identifier field, which controls the bus arbitration. DeviceNET makes use of wired-OR bus to connect all the nodes. When a sensor node has to send a message, it first calculates the message ID, which may be based on the priority of the message. The ID of each message must be unique. Then, each node writes the ID on the bus, one bit at a time, starting with the most significant bit. After writing each bit, each node waits long enough for signals to propagate along the bus, then it reads the bus. If a node had written a 0 but reads a 1, it means that another node has a message with higher priority. If so, this node drops out of contention. In the end, there is only one winner and it can use the bus.

1.1. **Multi-sensor case.**

Let there be 2^N sensors connected via a communication bus of DeviceNET type. Each sensor k makes a measurement x_k, which we model as a random variable with range X, and a finite second moment.[1] The monitoring station wants to estimate $x_k,\ k = 1, \ldots, 2^N$, so as to minimize the average distortion between the random variables x_k, and their estimates $\hat{x}_k$. In case of random variables with continuous density functions, we choose the distortion criterion as the mean-square error:

$$e = E\left\{\sum_{k=1}^{2^N}(x_k - \hat{x}_k)^2\right\}$$

[1]We do not index the random variables x_k with a time index, but the solution derived here for the static one-shot problem can be easily applied to solve the case when each sensor observes an i.i.d. random sequence.

For discrete random variables we choose the probability of error distortion criterion:

$$e = 1 - E\left\{ I_{\hat{x}_k = x_k, k=1,\ldots,2^N} \right\}$$

where I_S denotes the indicator function of the set S.

We assume that the random variables $\{x_k\}$ are independent. Each sensor, after observing the realized value of the random variable x_k, comes up with an N-bit *priority number* that is indicative of the importance of its current observation.

We also assume that 2^M measurements can be scheduled to be transmitted to the monitoring station. To avoid the triviality, we let $M < N$.

The arbitration between the sensors is carried out sequentially, where each sensor starts by writing the most significant bit of its priority number onto the bus. We assume that the bus acts as an arbiter by mapping the 2^N sensor bits written on it to either 0 or 1. Note that this is a static mapping that is determined *a priori*. The bus may, for example, take the exclusive-or (EX-OR) of the 2^N bits written on it. The result of the mapping is returned to each sensor, and then the whole process is repeated. In the end, each sensor will have written and read N bits onto and from the bus. The map of each sensor's priority number as to what it has read uniquely determines who will access the bus first. This procedure is then repeated 2^M times to determine all sensors that will gain access to the bus. In each repetition of the procedure, the sensor that gained access to the bus in the previous arbitration period is excluded from competing again. Also, at every other repetition of the process, the total number of bits written on the bus by each sensor is decreased by 1.

Thus, given this arbitration procedure, the problem becomes one of mapping each sensor's range space X onto a binary partition so as to minimize the average distortion between the sensor measurements and their estimates. Note that for each given partition, and arbitration mapping, the best estimate of the monitoring station is given by the maximum *a posteriori* (MAP) estimate or the conditional expectation of the random variables depending on the distortion criterion.

In the next section, we consider the special case of the general problem formulated above, in which $N = 1$, and $M = 0$. So, there are two sensors, and one of them will gain access to the bus. The priority number of each sensor is just 1-bit, and the arbitration phase has only one iteration.

1.2. **Two-sensor case.** Let x_1 and x_2 be two independent random variables that are observed by sensors 1 and 2, respectively; also see Figure 1.

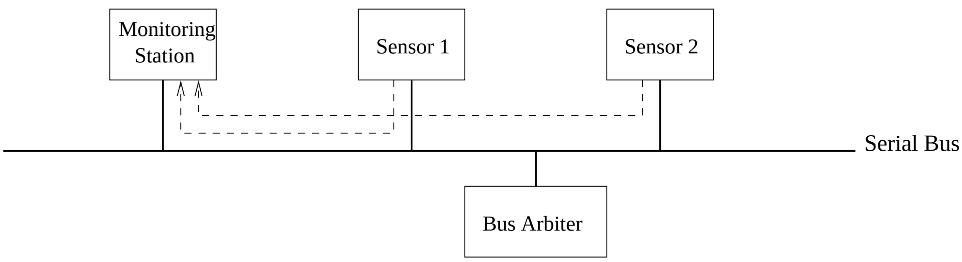

FIGURE 1. Illustration of a simple sensor bus configuration.

In the arbitration phase, each sensor comes up with a priority number, 0 or 1, depending on its current observation. Suppose the bus arbitration is carried out by an EX-OR function. Then, for each possible priority bit, 0 or 1, each sensor may potentially read back a 0 or 1. See Table 1 for a listing of all possible combinations for sensors 1 and 2 priority numbers p_1 and p_2, and what the bus will return in each case to both sensors. In Table 1, p_b denotes the bus output and is determined by taking the EX-OR of the priority bits, p_1 and p_2.

The last column in Table 1 shows who will access the bus for transmission. Access is granted by a fixed mapping, $a(p_1, p_2)$, that both sensors agree upon *a priori*. The arbitration map, $a(p_1, p_2)$,

TABLE 1. List of possible priority numbers of sensors 1 and 2, and the corresponding bus output.

$p_1 p_b$	$p_2 p_b$	access
0 0	0 0	0, x_1
0 1	1 1	1, x_2
1 1	0 1	1, x_2
1 0	1 0	0, x_1

simply maps each combination of priority bits p_1 and p_2 into 0 or 1, allowing respectively sensor 1 or 2 access the bus. Note that the exclusive-or map of priority bits, which is returned to each sensor by the bus, allows both sensors to unambiguously determine who will access the bus. In the above example, sensor 1 knows it can access the channel when it writes a 0 and reads a 0, or when it writes a 1 and reads a 0. Furthermore, sensor 1 knows that sensor 2 must have written a 0 when it writes and reads a 0, and sensor 2 must have written a 1 when it writes a 1 and reads a 0. Similarly, when sensor 2 gains access to the channel, it can obtain a 1-bit information about sensor 1.[2]

Therefore, the problem of scheduling one of the two random variables can be thought of as sensor k dividing the set of potential values the random variable x_k can take into two sets. When the random variable x_k becomes available, sensor k should send $p_k = 0$ or 1 depending on what set the realized value of x_k falls into.

This partitioning by both random variables into their individual sets will in turn partition the x_1-x_2 plane into four possible types of sets corresponding to the combinations 00, 01, 10, and 11. In each of these sets, the sensor that should gain access to the bus is determined by the arbitration map, $a(p_1, p_2)$. When the monitoring station receives the transmitted signal, it will not only determine the value of the random variable observed by the sensor that gained access to the bus, but also potentially a 1-bit information about the other sensor's measurement. This information will be generated by the monitoring station knowing the value of the arbitration function $a(p_1, p_2)$, as well as one of the priority numbers p_1 or p_2. If the arbitration function is such that for a fixed value of p_1 (or p_2), $a(p_1, p_2)$ is not the same for different values of p_2 (or p_1), then knowing p_1 (or p_2) and $a(p_1, p_2)$ will uniquely determine p_2 (or p_1). Basically, we want $a(p_1, p_2)$ to be invertible in the sense that knowing $a(p_1, p_2)$ and one of its arguments should uniquely determine the other argument.

The search for an optimal partitioning of sample spaces of random variables is in general a difficult problem when the random variables have continuous density functions. Thus, in the case when the sensor observations x_k are Gaussian or uniformly distributed, we simplify the search for an optimal solution by restricting ourselves to a class of partitions where each sensor can divide its range space into at most three simple intervals. This can be accomplished by sensor k picking two thresholds

$$-\infty \leq \tau_1^{(k)} \leq \tau_2^{(k)} \leq +\infty, \quad k = 1, 2$$

The choice of these thresholds by each sensor will in turn divide the x_1-x_2 plane into at most nine boxes. In each box, either sensor 1 or sensor 2 will gain access to the bus, which is determined by our choice of the arbitration function. In Sections 2.2 and 2.3, we give the solution of this optimization problem in two special cases, which we found by exhaustive search over all possible arbitration maps, and a parameterized optimization over the thresholds, $\tau_1^{(k)} \leq \tau_2^{(k)}, k = 1, 2$, for each fixed choice of an arbitration function.

When the random variables are discrete, we can solve the problem of optimally partitioning each sensor's set of outcomes. The solution to this case is given in Section 2.1.

[2]As we will see next, this is not always the case, and is determined by the *invertibility* of the arbitration function.

2. Solution of the two-sensor case

2.1. **Discrete random variables case.** Let x_1 and x_2 be independent discrete random variables with probability distributions $p_i^{(1)} = P[x_1 = m_i^{(1)}]$ and $p_i^{(2)} = P[x_2 = m_i^{(2)}]$, respectively. Without any loss of generality, we can assume that the probabilities $p_i^{(k)}, k = 1, 2$ are ordered in such a way that

$$p_1^{(k)} \geq p_2^{(k)} \geq \cdots$$

Suppose the goal is to minimize the probability of error distortion criterion:

$$e = P[\hat{x}_1 \neq x_1, \hat{x}_2 \neq x_2]$$

Given a particular partitioning of the range spaces of the random variables x_k, and an arbitration function $a(p_1, p_2)$, the optimal estimate of the random variable that is not transmitted to the monitoring station is given by its MAP estimate. Thus, to minimize the probability of error each sensor should distinguish its most likely outcome from all the other outcomes. In fact, it can be shown that the optimal scheduling policy of sensor k is to partition the set of outcomes, $m_i^{(k)}$, into two sets: a singleton with the most likely outcome, i.e. $\{m_1^{(k)}\}$, and a set with all the other outcomes, i.e. $\{m_i^{(k)}, i \geq 2\}$. With this partitioning the transmission policy of each sensor is as follows:

- Sensor 1 should transmit when $x_2 = m_1^{(2)}$ and $x_1 = m_i^{(1)}, i \geq 2$.
- Sensor 2 should transmit when $x_1 = m_1^{(1)}$ and $x_2 = m_i^{(2)}, i \geq 2$.
- Either Sensor 1 or Sensor 2 should transmit when $x_1 = m_1^{(1)}$ and $x_2 = m_1^{(2)}$.
- Sensor 1 should transmit when $x_1 = m_i^{(1)}, i \geq 2$ and $x_2 = m_i^{(2)}, i \geq 2$ if $p_2^{(1)}(1 - p_1^{(2)}) \leq p_2^{(2)}(1 - p_1^{(1)})$.
- Sensor 2 should transmit when $x_1 = m_i^{(1)}, i \geq 2$ and $x_2 = m_i^{(2)}, i \geq 2$ if $p_2^{(1)}(1 - p_1^{(2)}) > p_2^{(2)}(1 - p_1^{(1)})$.

Note that, with this transmission and arbitration policy the monitoring station can determine both sensor measurements if

- $x_1 = m_1^{(1)},\ x_2 = m_i^{(2)}, i \geq 1$.
- $x_2 = m_1^{(2)},\ x_1 = m_i^{(1)}, i \geq 1$.
- either $x_1 = m_2^{(1)},\ x_2 = m_i^{(2)}, i \geq 2$ or $x_2 = m_2^{(2)},\ x_1 = m_i^{(1)}, i \geq 1$.

2.2. **Gaussian random variables case.** Let x_1 and x_2 be independent Gaussian with zero-mean and variances $\sigma_{x_1}^2$, $\sigma_{x_2}^2$, respectively. In the class of two-threshold partitions, the solution of the scheduling optimization problem is determined by the ratio of the standard deviations

$$\rho := \frac{\sigma_{x_1}}{\sigma_{x_2}}$$

We have three cases depending on what the value of ρ is. If

$$l_g \leq \rho \leq \frac{1}{l_g}, \tag{1}$$

then the optimal solution for each sensor is to divide the real line into two sets around zero. In other words, sensor k should set $p_k = 0$ or 1 depending on the sign of its observation. This essentially corresponds to setting

$$\tau_1^{(k)} = \tau_2^{(k)} = 0, \quad k = 1, 2$$

Then, if we use any invertible arbitration function to decide on who should access the bus, the monitoring station, upon receiving one of the random variables, will be able to uniquely determine the sign of the other random variable that has not been transmitted. Hence, we can basically have either one of the following two policies:

$$P_1 = \begin{cases} \{x_1 > 0, x_2 > 0\} \text{ OR } \{x_1 < 0, x_2 < 0\} \Rightarrow \text{TX } x_1 \\ \{x_1 < 0, x_2 > 0\} \text{ OR } \{x_1 > 0, x_2 < 0\} \Rightarrow \text{TX } x_2 \end{cases}$$

$$P_2 = \begin{cases} \{x_1 > 0, x_2 > 0\} \text{ OR } \{x_1 < 0, x_2 < 0\} \Rightarrow \text{TX } x_2 \\ \{x_1 < 0, x_2 > 0\} \text{ OR } \{x_1 > 0, x_2 < 0\} \Rightarrow \text{TX } x_1 \end{cases}$$

where TX stands for transmit. The decision regions describing the policy P_1 is depicted in Figure 2. A similar figure can be drawn for the policy P_2.

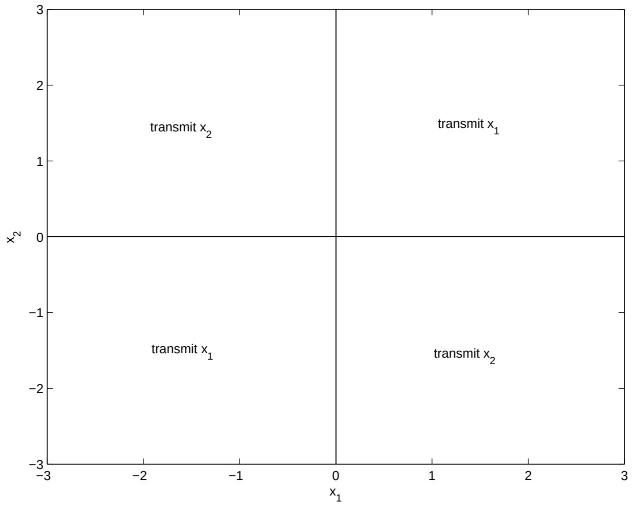

FIGURE 2. Decision regions for $l_g \leq \rho \leq \frac{1}{l_g}$.

In (1) above, l_g is the solution of the nonlinear equation

$$(l_g^2 + 1)\left(\frac{1}{2} - \frac{1}{\pi}\right) = 2l_g^2(1 - \Phi(l_g)) + (2\Phi(l_g) - 1) - \frac{2l_g^2}{\sqrt{2\pi}} e^{-\frac{1}{2}l_g^2}$$

and is approximately equal to $l_g \approx 0.5955$.

We should expect this symmetric solution for the case when both random variables have equal variance, but it is interesting to note that the solution remains the same even when $\rho \neq 1$, i.e. $\sigma_{x_1} \neq \sigma_{x_2}$, up until a certain threshold.

When the ratio ρ exceeds $\frac{1}{l_g}$, the optimal solution changes its symmetric structure. Now, the sensor that is observing the random variable with the larger variance, x_1, dominates most of the observation space. Sensor 1 partitions the x_1 axis into three intervals, one around zero from $-\sigma_{x_2}$ to $+\sigma_{x_2}$, and the other two, $(-\infty, -\sigma_{x_2})$ and $(+\sigma_{x_2}, \infty)$. Sensor 2, on the other hand, does not do any partitioning of its own observation space. This partitioning of the x_1-x_2 plane corresponds to the following selection of the thresholds:

$$-\tau_1^{(1)} = \tau_2^{(1)} = \sigma_{x_2}, \quad \tau_1^{(2)} = -\infty, \tau_2^{(2)} = +\infty$$

With this partitioning of the observation space, the optimal transmission policy is of the form

$$P = \begin{cases} \{x_1 > +\sigma_{x_2}, -\infty < x_2 < +\infty\} \Rightarrow \text{TX } x_1 \\ \{x_1 < -\sigma_{x_2}, -\infty < x_2 < +\infty\} \Rightarrow \text{TX } x_1 \\ \{-\sigma_{x_2} < x_1 < +\sigma_{x_2}, -\infty < x_2 < +\infty\} \Rightarrow \text{TX } x_2 \end{cases}$$

which can be realized with a number of arbitration functions. In Figure 3, we show the decision regions in the x_1-x_2 plane described by the policy P when $\sigma_{x_2} = 1$.

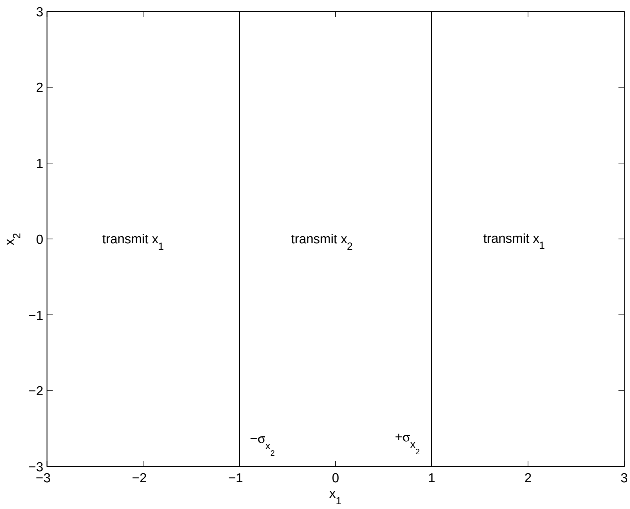

FIGURE 3. Decision regions for $\rho \geq \frac{1}{l_g}$, $\sigma_{x_2} = 1$.

Note that, when the monitoring station receives a transmission from x_1, this will not reveal anything about x_2, but when it receives a transmission from x_2, this will tell him that $|x_1| \leq \sigma_{x_2}$.

The solution when the ratio ρ becomes smaller than l_g is similar, with the roles of x_1 and x_2 interchanged. This case is illustrated in Figure 4 when $\sigma_{x_1} = 1$.

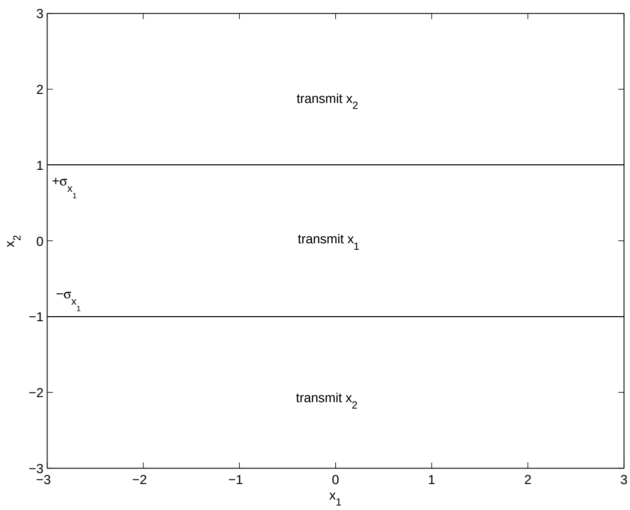

FIGURE 4. Decision regions for $\rho \leq l_g$, $\sigma_{x_1} = 1$.

Note that the solution is completely symmetric in x_1 and x_2. Also, as $\sigma_{x_2} \to 0$, the set over which we transmit x_2 (i.e., $\{|x_1| \leq \sigma_{x_2}, -\infty < x_2 < +\infty\}$) becomes arbitrarily small. However, it is interesting to note that for a fixed value of σ_{x_2}, no matter how large σ_{x_1} is, we still transmit x_2 in a set with nonzero measure.

2.3. **Uniform random variables case.** We next present the solution to the uniform random variable case. Let x_1 and x_2 be zero-mean, independent, and uniformly distributed on the intervals $[-\sqrt{3}\sigma_{x_1}, +\sqrt{3}\sigma_{x_1}]$, and $[-\sqrt{3}\sigma_{x_2}, +\sqrt{3}\sigma_{x_2}]$, respectively.

The solution here turns out to be very similar to the Gaussian case. Let again

$$\rho := \frac{\sigma_{x1}}{\sigma_{x2}}$$

We have three cases depending on what the value of ρ is. If

$$l_u \leq \rho \leq \frac{1}{l_u}$$

where $l_u \approx 0.4184$, the optimal solution for sensor k is to divide the interval $[-\sqrt{3}\sigma_{x_k}, +\sqrt{3}\sigma_{x_k}]$ into two symmetric sets around zero, and the exact same analysis as in the Gaussian case follows. Figure 5 shows the decision regions in this case when $\sigma_{x_1} = \sigma_{x_2} = 1$.

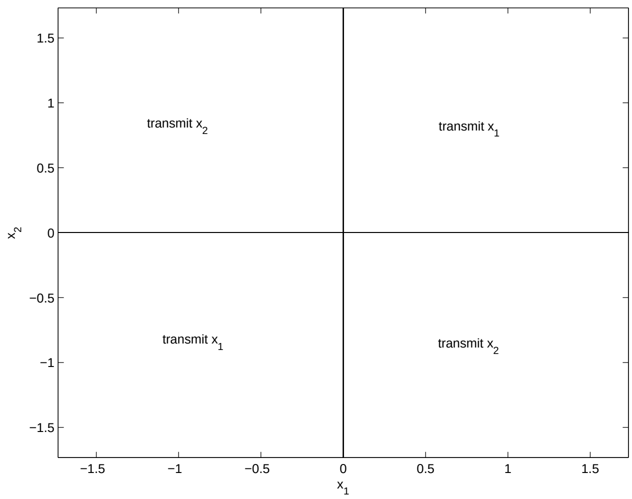

FIGURE 5. Decision regions for $l_u \leq \rho \leq \frac{1}{l_u}$, $\sigma_{x_1} = \sigma_{x_2} = 1$.

When the ratio ρ exceeds $\frac{1}{l_u}$, the optimal solution changes its symmetric structure just like in the Gaussian case. However, unlike the Gaussian case, now the sensor that is observing the random variable with the larger variance x_1 can divide the interval $[-\sqrt{3}\sigma_{x_1}, +\sqrt{3}\sigma_{x_1}]$ into three intervals in infinitely many ways as long as the length of the interval where x_1 will not transmit is fixed at $2\sigma_{x_2}$. Sensor 2, on the other hand, does not do any partitioning of its own observation space. This partitioning of the x_1-x_2 plane corresponds to the following selection of the thresholds:

$$|\tau_2^{(1)} - \tau_1^{(1)}| = 2\sigma_{x_2}, \quad \tau_1^{(2)} = -\sqrt{3}\sigma_{x_2}, \tau_2^{(2)} = +\sqrt{3}\sigma_{x_2}$$

This partitioning is illustrated in Figure 6 when $\sigma_{x_1} = 3$, $\sigma_{x_2} = 1$.

Finally, like in the Gaussian case, the solution when the ratio ρ becomes smaller than l_u is similar, with the roles of x_1 and x_2 interchanged. Figure 7 illustrates the partitioning of the x_1-x_2 plane in this case.

3. CONCLUSIONS

In this paper, we introduced the problem of decentralized scheduling of random variables with an arbitration mechanism for estimation purposes. We note that the problem resembles the multiple access channel problem of network information theory [4], but with additional constraints due to the arbitration mechanism. When there are only two random observations that needs to be scheduled, we showed that a simple solution can be obtained for the discrete random variables case. When the random variables have continuous density functions a threshold

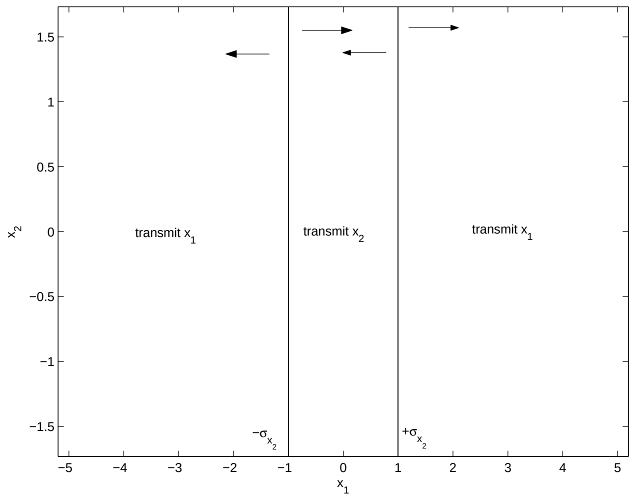

FIGURE 6. Decision regions for $\rho \geq \frac{1}{l_u}$, $\sigma_{x_1} = 3$, $\sigma_{x_2} = 1$.

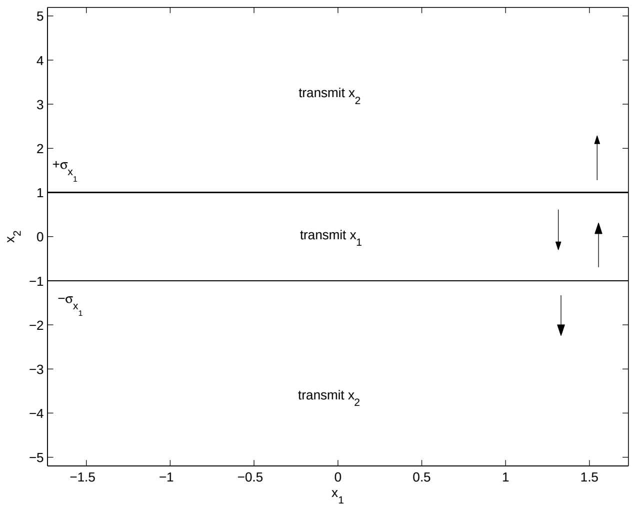

FIGURE 7. Decision regions for $\rho \leq l_u$, $\sigma_{x_1} = 1$, $\sigma_{x_2} = 3$.

solution can be obtained by restricting the search space, and solving a parametric optimization problem.

We note that the results in this paper complement those in [13] where optimal control of a discrete time linear system with scheduled measurements and controls is discussed. In [13] the controller is connected with a sensor and an actuator over a serial bus, and it has to tradeoff between receiving a measurement signal from the sensor and sending a control signal to the actuator. This leads to a scheduling between measurement and control signals, as opposed to a scheduling of measurement signals as in this paper.

References

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

[2] Bertsekas, D. and Gallegar, R., *Data Networks*, 2nd ed. Englewood Cliffs, NJ: Prentice-Hall Inc., 1992.

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

[4] Cover, T. M. and Thomas, J.A., *Elements of Information Theory.* New York, NY: John Wiley & Sons, Inc., 1991.

[5] *DeviceNet Specifications*, 2.0 ed. Open DeviceNet Vendors Association, Boca Raton, Florida, 1997.
[6] Hristu, D., Stabilization of LTI systems with communication constraints, in *Proc. American Control Conference (ACC)*, June 2000, pp. 2342–2346.
[7] Hristu, D. and Kumar, P.R., Interrupt-based feedback control over a shared communication medium, in *Proc. IEEE Conference on Decision and Control (CDC)*, December 2002, pp. 1787–1792.
[8] Imer, O.C., Optimal estimation and control under communication network constraints,*Ph.D. Thesis*, University of Illinois, 2005.
[9] Imer, O.C. and Başar, T., To Measure or to control: Optimal control with scheduled measurements and controls, submitted to *ACC 2006.*
[10] Lian, F.-L., Moyne, J.R. and Tilbury, D.M., Performance evaluation of control networks: Ethernet, ControlNet, and DeviceNet, *IEEE Control Systems Magazine*, vol. 21, pp. 66–83, February 2001.
[11] Mehra, R.K., 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.
[12] Rehbinder, H. and Sanfridson, M., Scheduling a limited communication channel for optimal control, in *Proc. IEEE Conference on Decision and Control (CDC)*, December 2000, pp. 1011–1016.
[13] Tanenbaum, A. S., *Computer Networks*, 3rd ed. Upper Saddle River, NJ: Prentice-Hall Inc., 1996.
[14] Walsh, G. C. and Ye, H., Scheduling of networked control systems, *IEEE Control Systems Magazine*, vol. 21, pp. 57–65, February 2001.
[15] Zuberi, K.M. and Shin, K. G., Design and implementation of efficient message scheduling for controller area network, *IEEE Transactions on Computers*, vol. 49, no. 2, pp. 182–188, February 2000.

Tamer Başar - is with the University of Illinois at Urbana-Champaign (UIUC), where he holds the positions of the Fredric G. and Elizabeth H. Nearing Endowed Professor of Electrical and Computer Engineering, Center for Advanced Study Professor, and Research Professor at the Coordinated Science Laboratory. He received the B.S.E.E. degree from Robert College, Istanbul, in 1969, and the M.S., M.Phil, and Ph.D. degrees from Yale University during the period 1970-1972.

He joined UIUC in 1981 after holding positions at Harvard University and Marmara Research Institute (Turkey). He has published extensively in systems, control, communications, and dynamic games, and has current research interests in modeling and control of communication networks; control over heterogeneous networks; resource allocation, management and pricing in networks; mobile computing; security issues in computer networks; and risk-sensitive estimation and control. He is the Editor-in-Chief of *Automatica*, Editor of the Birkhäuser Series on *Systems & Control*, Managing Editor of the *Annals of the International Society of Dynamic Games* (ISDG), and member of editorial and advisory boards of several international journals in control, wireless networks, and applied mathematics. He has received several awards and recognitions over the years, among which are the Medal of Science of Turkey (1993); Distinguished Member Award (1993), Axelby Outstanding Paper Award (1995) and Bode Lecture Prize (2004) of the IEEE Control Systems Society (CSS); Millennium Medal of IEEE (2000); Tau Beta Pi Drucker Eminent Faculty Award of UIUC (2004); and the Outstanding Service Award (2005) and the Giorgio Quazza Medal (2005) of the International Federation of Automatic Control (IFAC). He is a member of the National Academy of Engineering, a member of the European Academy of Sciences, a Fellow of IEEE, a Fellow of IFAC, a past president of CSS, and the founding president of ISDG.

Orhan Cagri Imer was born in Berlin, Germany. He received his BS degree in Electrical Engineering in 1997 from the Middle East Technical University in Ankara; and his MS and PhD degrees in Electrical and Computer Engineering in 2000 and 2005 from the University of Illinois at Urbana-Champaign. He was a research assistant in the Decision and Control group at the Coordinated Science Laboratory of University of Illinois from 1997 to 2005. During the same period, he was also a teaching assistant and an instructor in the Electrical and Computer Engineering department at the University of Illinois at Urbana-Champaign. In the summer of 2001, he had experience as an intern at the Tellabs Research Center in Mishawaka, Indiana. Currently, he is a research engineer at the GE Global Research Center of the General Electric Corporation in Niskayuna, New York. Dr. Imer is a member of the IEEE, and is the recipient of the C. R. Allen Outstanding International Student award at the Electrical and Computer Engineering department of University of Illinois. His current research interests are in communication and control theories with applications in networking.