Optimal estimation with limited measurements

Orhan C. Imer and Tamer Başar*

Coordinated Science Laboratory,
University of Illinois at Urbana Champaign,
Urbana, IL 61801, USA
E-mail: orhan.imer@gmail.com
E-mail: tbasar@control.csl.uiuc.edu
*Corresponding author

Abstract: This paper introduces a sequential estimation problem with two decision makers, or agents, who work as members of a team. One of the agents sits at an observation post, and makes sequential observations about the state of an underlying stochastic process for a fixed period of time. The observer agent upon observing the process makes a decision as to whether or not to disclose some information about the process to the other agent who acts as an estimator. The estimator agent sequentially estimates the state of the process. The agents have the common objective of minimising a performance criterion with the constraint that the observer agent may only act a limited number of times.

Keywords: networked control systems; optimal estimation; limited information; wireless sensor networks.

Reference to this paper should be made as follows: Imer, O.C. and Başar, T. (2010) 'Optimal estimation with limited measurements', *Int. J. Systems, Control and Communications*, Vol. 2, Nos. 1/2/3, pp.5–29.

Biographical notes: Orhan C. Imer received his BS Degree in Electrical Engineering from the Middle East Technical University, Ankara, in 1997; and his MS and PhD Degrees in Electrical and Computer Engineering from the University of Illinois at Urbana-Champaign in 2000 and 2005. 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. He is a member of the IEEE, and is the recipient of the C.R. Allen International Student Award at the Electrical and Computer Engineering Department of University of Illinois.

Tamer Başar is with the University of Illinois at Urbana-Champaign, where he joined in 1981, and is currently Swanlund Endowed Chair and Center for Advanced Study Professor of Electrical and Computer Engineering. He received BSEE from Robert College, Istanbul, and MS, MPhil, and PhD from Yale University. He has published extensively in the areas of systems, control, communications, and dynamic games, and has current research interests in networks, security, pricing, sensing, and control. He is a member of the National Academy of Engineering, a fellow of IEEE, a fellow of IFAC, and the incoming President of the American Automatic Control Council.

1 Introduction

Recursive estimation of a linear stochastic process with full and partial state information has been extensively studied in the literature (Anderson and Moore, 1979). In this paper, we introduce the problem of recursive estimation with *limited* information. More specifically, we consider estimating a one-dimensional discrete-time stochastic process over a decision horizon of length N using only $M \leq N$ measurements. Both the measurement and the estimation of the process are carried out sequentially by two different decision makers, or agents, called the *observer* and the *estimator*,[1] respectively. Over the decision horizon of length N, the observer agent has exactly M opportunites to disclose some information about the process to the estimator. These information disclosures, or *transmissions*, are assumed to be error and noise free, and the problem is to jointly determine the best observation and estimation policies that minimise the average estimation error between the process and its estimate.

Estimation problems of this nature arise in real-time monitoring and control applications over wireless sensor networks (Imer, 2005; Imer and Başar, 2005, 2006a, 2006b, 2007; US Department of Energy, 2004). Due to the power-limited nature of the wireless sensors, in most sensor net applications the wireless devices can only make a *limited* number of transmissions (US Department of Energy, 2004). The limited battery power of the wireless device can be modelled by imposing a hard constraint on either the number of available transmissions it can make or the number of cycles it can stay awake (Imer, 2005; Imer and Başar, 2005, 2006a, 2006b, 2007). This hard constraint can be thought of as a measurement budget, and the problem is to determine the best way to spend this budget by scheduling the measurements over a decision horizon.

Earlier work on this general topic includes (Athans, 1972; Mehra, 1976), where the focus was 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. Another paper (Bansal and Başar, 1989) introduced and solved a joint control and sensor design problem with soft constraints on both sensing and control. In a more recent paper (Rabi et al., 2006), the authors study sampling policies for estimation on a finite-time horizon for continuous-time scalar systems. Also relevant are (Cogill et al., 2007; Xu and Hespanha, 2004; Xu and Hespanha, 2005), where communication costs are introducted as a *soft-constraint* into the linear sequential estimation problem.

The rest of this paper is organised as follows. In Section 2, we formally define the problem, and breifly discuss some of the potential applications. Section 3 discusses estimating a sequence of independent identically distributed (i.i.d.) random variables under a given bound on the number of measurements. In Section 4, the results of Section 3 are extended to the Gauss-Markov case. We present some illustrative examples in Section 5, and the paper ends with the concluding remarks of Section 6.

2 Problem statement

2.1 Problem definition

The problem of optimal estimation with limited measurements can be treated in the more general framework of a communication system with limited channel uses.

For this purpose, consider the generic communication system whose block diagram is given in Figure 1 (Cover and Thomas, 1991). The source outputs some data b_k for $0 \leq k \leq N-1$, that needs to be communicated to the user over a channel. The data b_k are generated according to some *a priori* known stochastic process, $\{b_k\}$, which may be i.i.d., or correlated as in a Markov process. An encoder (or an observer) and a decoder (or an estimator) is placed after the source output and the channel output, respectively, to communicate the data to the user efficiently. In the most general case, the encoder/observer may have access to a noise-corrupted version of the source output:

$$z_k = b_k + v_k, \quad 0 \leq k \leq N-1$$

where $\{v_k\}$ is an independent[2] noise process.

Figure 1 Communication with limited channel use

The main constraint is that the encoder/observer can access the channel only a *limited*, $M < N$, number of times. The goal is to design an observer-estimator pair,[3] $(\mathcal{O}, \mathcal{E})$, that will 'causally' (or sequentially) observe/encode the data measurements, z_k, and estimate/decode the channel output, y_k, so as to minimise the *average distortion* or *error* between the observed data, b_k, and estimated data, $\hat{b}_k$.

The channel is assumed to be memoryless, and is completely characterised by the conditional probability distribution $P_c(y \mid x)$ on $y \in \mathcal{Y}$ for each $x \in \mathcal{X}$, where $\mathcal{X}$ and $\mathcal{Y}$ are the set of allowable channel inputs, and the set of possible channel outputs, respectively. In this paper, we assume that the channel is noiseless; hence it is characterised by the probability distribution $P_c(y \mid x) = \delta(y - x)$.

The average distortion, $D_{(M,N)}$, depends on the distortion measure and may be picked differently depending on the underlying application. Some examples are the average mean-square error

$$D_{(M,N)} = E\left\{\frac{1}{N}\sum_{k=0}^{N-1}(b_k - \hat{b}_k)^2\right\} \tag{1}$$

or the Hamming (probability of error) distortion measure

$$D_{(M,N)} = E\left\{\frac{1}{N}\sum_{k=0}^{N-1}\mathcal{I}_{b_k \neq \hat{b}_k}\right\} \tag{2}$$

where $\mathcal{I}_S$ denotes the indicator function of the set S.

From a communication-theoretic standpoint, with the channel, source, and the distortion measure defined, we can formally state our main problem: Given a source and a memoryless channel, for a given decision-horizon N, and number of channel uses M, what is the minimum attainable value of the average distortion $D_{(M,N)}$? This minimisation is carried out over the choice of possible encoder-decoder (observer-estimator) pairs which are *causal*.

In what follows, we present a solution to this problem when the source process is i.i.d. with a continuous or discrete probability density function, and the encoder/observer has access to the noiseless or a noisy version of the source output. We also present the solution to the case when the source process is Gauss-Markov.

2.2 Example applications

Recent advances in wireless technology and standards, such as ZigBee and IEEE 802.15.4, have made wireless sensing solutions feasible for industrial applications (US Department of Energy, 2004; Gutierrez et al., 2003; Stralen et al., 2006; Azimi et al., 2006). Most of these applications use battery-powered integrated wireless sensing/communication devices, also called *motes*, for data logging and monitoring purposes (Gutierrez et al., 2003). Often times, data collected from sensors is relayed back to a central processing unit where it is analysed for trends. In most monitoring applications, the data is collected on a near real-time basis. An important issue that needs to be addressed in designing these systems is the power-limited nature of the wireless devices.

In most industrial applications a battery lifespan in the order of several years is required for feasible commercial operation (US Department of Energy, 2004). This requirement imposes severe restrictions on the duration of time the wireless device can be on/awake and the number of transmissions it can make. This is because the Radio Frequency (RF) communication consumes a significant portion of the battery power when the wireless unit is awake. Therefore, life of the wireless device can be lengthened by optimising the duty cycle (or reporting frequency) of the unit as well as by transmitting data only when it is necessary.

The desired length of time the wireless device will be in operation can be related to the decision horizon N in some appropriate time unit, and the size of the battery installed in the sensor can be related to the possible number of transmissions or channel uses M (see Figure 2).

Figure 2 Optimal transmission scheduling with limited channel access

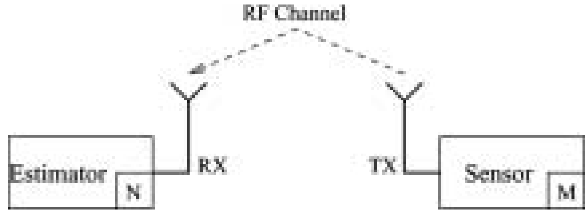

Hence, given an underlying performance criterion $D_{(M,N)}$, the problem is to design the best transmission schedule and estimation policies for the wireless device and the remote monitoring station, respectively.

3 Estimating an i.i.d. random sequence with limited measurements

3.1 Problem definition

Consider the special case of the general problem defined in Section 2, where the source outputs a zero-mean[4] i.i.d. random sequence $b_k,\ 0 \le k \le N-1$. Let $\mathcal{B}$ denote the range of the random variable b_k. We assume that b_ks have a finite second moment, $\sigma_b^2 < \infty$, but their probability distribution remains unspecified for now. At time k, the encoder/observer makes a sequential measurement of b_k, and determines whether to access the channel for transmission, which it can only do a limited, $M \le N$, number of times. The channel is noiseless and thus has a capacity to transmit the source output error-free when it is used to transmit. Note that, even when it decides not to use the channel for transmission, the observer/encoder may still convey a 1-bit information to the estimator/decoder. In view of this, the channel input x_k belongs to the set $\mathcal{X} := \mathcal{B} \cup \{\text{NT}\}$, where NT stands for 'no transmission'.

More precisley, we let s_k denote the number of channel uses (or transmissions) left at time k. Now if $s_k \ge 1$, we have $y_k = x_k$ for $x_k \in \mathcal{B} \cup \{\text{NT}\}$. If $s_k = 0$, on the other hand, the channel is useless, since we have exhausted the allocated number of channel uses. Note that, when the channel is noiseless, both the encoder and the decoder can keep track of s_k by initialising $s_0 = M$ and decrementing it by 1 every time a transmission decision is taken.

We want to design an estimator/decoder

$$\hat{b}_k = \hat{\mu}_k(I_k^d) \quad \text{for } 0 \le k \le N-1$$

based on the available information I_k^d at time k. Clearly, the information available to the estimator is controlled by the observer. The average distortion between the observed and estimated processes can be taken to be the average mean square error as given by equation (1), or the probability of error distortion measure which is given by equation (2).

The information I_k^d available to the estimator at time k is a result of an outcome of decisions taken by the observer up until time k. Let the observer's decision at time k be

$$x_k = \mu_k(I_k^e)$$

where I_k^e is the information available to the observer at time k. Assuming perfect recall, we have

$$I_0^e = \{(s_0, t_0); b_0\}$$
$$I_k^e = \{(s_k, t_k); b_0^k; x_0^{k-1}\}, \quad 1 \le k \le N-1$$

where t_k denotes the number of time, or decision slots left at time k. We have

$$t_{k+1} = t_k - 1, \quad 0 \le k \le N-2$$

with $t_0 = N$.

The range of $\mu_k(\cdot)$ is the space $\mathcal{X} = \mathcal{B} \cup \{\text{NT}\}$. Let σ_k denote the decision whether the observer has decided to transmit or not. Assume $s_k \geq 1$, and let $\sigma_k = 1$ if a transmission takes place; i.e., $x_k \in \mathcal{B}$, and $\sigma_k = 0$ if no transmission takes place. We have

$$s_{k+1} = s_k - \sigma_k, \quad 0 \leq k \leq N-2$$

with $s_0 = M$.

The observer's decision at time k is a function of its k past measurements, and $k-1$ past decisions, i.e.,

$$\mu_k(I_k^e) : \mathcal{B}^k \times \mathcal{X}^{k-1} \to \mathcal{X}, \quad 0 \leq k \leq N-1.$$

Now, the information I_k^d available to the estimator at time k can be written as

$$I_k^d = \{(s_k, t_k); y_0^k\}, \quad 0 \leq k \leq N-1.$$

By definition, the channel output y_k satisfies $y_k = x_k$ if $s_k \geq 1$, and $y_k \in \emptyset$ (i.e., no information) if $s_k = 0$.

Consider the class of observer-estimator (encoder-decoder) policies consisting of a sequence of functions

$$\Pi = \{\mu_0, \hat{\mu}_0, \ldots, \mu_{N-1}, \hat{\mu}_{N-1}\}$$

where each function μ_k maps I_k^e into $\mathcal{X}$, and $\hat{\mu}_k$ maps I_k^d into $\mathcal{B}$,[5] with the additional restriction that μ_k can map to $\mathcal{B}$ at most M times. Such policies are called *admissible*.

We want to find an admissible policy $\pi^* \in \Pi$ that minimises the average N-stage distortion, or estimation error:

$$e^{\pi}_{(M,N)} = E\left\{ \sum_{k=0}^{N-1} (b_k - \hat{\mu}_k(I_k^d))^2 \right\} \tag{3}$$

or for source processes, b_k, with discrete probability densities:

$$e^{\pi}_{(M,N)} = E\left\{ \sum_{k=0}^{N-1} \mathcal{I}_{b_k \neq \hat{\mu}_k(I_k^d)} \right\}. \tag{4}$$

That is

$$e^{*}_{(M,N)} = \min_{\pi \in \Pi} e^{\pi}_{(M,N)}. \tag{5}$$

Note that, we omitted the factor of $\frac{1}{N}$ from the average error expressions for convenience.

If $M \geq N$, this problem has the trivial solution where the observer writes the source output b_k directly into the channel at each time k (i.e., $\mu_k^*(b_k) = b_k$), and since the channel is noiseless, the estimator can use an identity mapping (i.e., $\hat{\mu}_k^*(I_k^d) = b_k$), resulting in zero distortion. Therefore, we only consider the case when $M < N$.

Before closing our account on this section, we would like to note the nonclassical nature of the information in this problem. Clearly, the observer's action affects the information available to the estimator, and there is no way in which the estimator can infer the information available to the observer. Also note the order of actions between the decision makers in the problem. At time k, first the random variable b_k becomes available, then the observer acts by transmitting some data or not, and finally, the estimator acts by estimating the state with $\hat{\mu}_k$, the cost is incurred, and we move to the next time $k+1$.

3.2 Structure of the solution

We first consider the problem of finding the optimal estimator $\hat{\mu}_k^*$ at time k. Note that the estimator $\hat{\mu}_k$ appears only in a single term in the error expressions (3) and (4). Thus, for the mean-square error criterion, the optimal estimator is simply the solution of the quadratic minimisation problem

$$\min_{\hat{\mu}_k(I_k^d)} E\{(b_k - \hat{\mu}_k(I_k^d))^2 \mid I_k^d\} \tag{6}$$

which is given by the conditional expectation of b_k given the available information at time k:

$$\hat{\mu}_k^*(I_k^d) = E\{b_k \mid I_k^d\} = E\{b_k \mid (s_k, t_k); y_0^k\}. \tag{7}$$

Similarly, for the probability of error distortion criterion, the optimal estimator is the solution of the minimisation problem

$$\min_{\hat{\mu}_k(I_k^d)} E\{\mathcal{I}_{b_k \neq \hat{\mu}_k(I_k^d)} | I_k^d\}.$$

If at time k the channel can still be used ($s_k \geq 1$), the solution to this problem is given by the maximum *a posteriori* probability (MAP) estimate of the random variable b_k given the available information at time k:

$$\hat{\mu}_k^*(I_k^d) = \arg \max_{m_i \in \mathcal{B}_k(I_k^d)} \delta(y_k - i) p_i = \arg \max_{m_i \in \mathcal{B}_k((s_k, t_k); y_0^k)} p_i \tag{8}$$

where $\mathcal{B}_k(I_k^d) \subset \mathcal{B}$ is some subset of the range of the random variable b_k, which we assume is countable. Let $\{m_i\}$ denote the finite set of values the random variable b_k takes. Then, p_i's make up the probability mass function of the random variable b_k, i.e., $p_i = P[b_k = m_i]$.

Note that, for the probability of error distortion criterion, if the channel is useless at time k (i.e., $s_k = 0$), the best estimate of b_k is simply given by

$$\hat{\mu}_k^*(I_k^d) = \arg \max_{m_i \in \mathcal{B}} p_i \tag{9}$$

since the past channel outputs, y_0^{k-1}, are independent of b_k.

Similarly, for the mean-square error criterion, the channel output y_k has no information on b_k if $s_k = 0$. Thus, in this case, the conditional expectation in equation (7) is given by

$$\hat{\mu}_k^*(I_k^d) = E\{b_k \mid (0, t_k); y_0^{k-1}, y_k\} = E\{b_k\} = 0 \tag{10}$$

since again the past channel outputs, y_0^{k-1}, are generated by the σ-algebra of random variables b_0^{k-1}, and hence are independent from b_k.

If $s_k \geq 1$, the channel output $y_k = x_k$, but since $y_0^{k-1} = x_0^{k-1}$ is the outcome of a Borel-measurable function defined on the σ-algebra generated by b_0^{k-1}, the conditional expectation in equation (7) is equivalent to

$$\hat{\mu}_k^*(I_k^d) = E\{b_k \mid (s_k, t_k); x_k\}. \tag{11}$$

By a similar argument, we can write equation (8) as

$$\hat{\mu}_k^*(I_k^d) = \arg \max_{m_i \in \mathcal{B}_k((s_k, t_k); x_k)} p_i. \tag{12}$$

Now, substituting the optimal estimators (11)–(12) back into the estimation error expressions (3)–(4) yields

$$e^{\pi}_{(M,N)} = E\left\{ \sum_{k=0}^{N-1} (b_k - E\{b_k \mid (s_k, t_k); x_k\})^2 \right\} \tag{13}$$

and

$$e^{\pi}_{(M,N)} = E\left\{ \sum_{k=0}^{N-1} \mathcal{I}_{b_k \neq \arg\max_{m_i \in \mathcal{B}_k((s_k, t_k); x_k)} p_i} \right\} \tag{14}$$

which we seek to minimise over the observer/encoder policies $\mu_k(I_k^e), 0 \leq k \leq N-1$. Since $x_k = \mu_k(I_k^e)$, we see that the choice of an observer policy affects the cost only through the information made available to the estimator.

In general, the observer's decision μ_k at time k depends on (s_k, t_k), all past measurements b_0^{k-1}, the present measurement b_k, and its past actions x_0^{k-1}. However, as we show next, there is nothing the observer can gain by having access to its past measurements, b_0^{k-1}, and its past actions, x_0^{k-1}, as far as the optimisation of the criteria (13)–(14) are concerned. Thus, *sufficient statistics* for the observer are the current measurement b_k and the remaining number of channel uses (transmission opportunities) and decision instances, i.e. (s_k, t_k).

Proposition 1: *The set* $S_k^e = \{(s_k, t_k); b_k\}$ *constitutes sufficient statistics* $S_k^e(I_k^e)$ *for the optimal policy* μ_k^* *of the observer. In other words,*

$$\mu_k^*(I_k^e) = \bar{\mu}(S_k^e(I_k^e))$$

for some function $\bar{\mu}$.

Proof: Suppose we would like to determine the optimal observer policy $\mu_k^*(I_k^e)$ at time k, where $0 \le k \le N-1$ is arbitrary. Due to the sequential nature of the decision problem, any observer policy we decide on at time k will only affect the error e_k incurred after time k, i.e.,[6]

$$e_k = E\left\{ \sum_{n=k}^{N-1} (b_n - E\{b_n \mid (s_n, t_n); x_n\})^2 \right\}.$$

Taking the conditional expectation given the available information I_k^e, under any observer policy $\mu_k(I_k^e)$ we have

$$E\{e_k \mid (s_k, t_k); b_0^k; x_0^{k-1}\} = E\{e_k \mid (s_k, t_k); b_k\} \tag{15}$$

because b_k^{N-1} is independent of b_0^{k-1}, and x_0^{k-1} is the outcome of a Borel-measurable function defined on the σ-algebra generated by b_0^{k-1}. Hence, at time k, the knowledge of b_0^{k-1} and x_0^{k-1} is redundant. □

A consequence of Proposition 1 is that the observer's decision whether to use the channel to transmit a source measurement or not is based purely on the current observation b_k, and its past actions only through (s_k, t_k).

Since μ_k depends explicitly only on the current source output b_k, the search for an optimal observer policy can be narrowed down to the class of policies of the form[7]

$$\mu_k(I_k^e) = \bar{\mu}((s_k, t_k); b_k) = \begin{cases} b_k & \text{if } b_k \in \mathcal{T}_{(s_k,t_k)} \\ \text{NT} & \text{if } b_k \in \mathcal{T}^c_{(s_k,t_k)} \end{cases} \tag{16}$$

where $\mathcal{T}_{(s_k,t_k)}$ is a measurable set on $\mathcal{B}$ and is a function of (s_k, t_k). The complement of the set $\mathcal{T}_{(s_k,t_k)}$ is taken with respect to $\mathcal{B}$, i.e., $\mathcal{T}^c_{(s_k,t_k)} = \mathcal{B} \backslash \mathcal{T}_{(s_k,t_k)}$. When probability of error distortion criterion is used, Proposition 1 implies that $\mathcal{B}_k((s_k, t_k); \text{NT}) = \mathcal{T}^c_{(s_k,t_k)}$, and $\mathcal{B}_k((s_k, t_k); m_i) = m_i$.

Note that the optimal estimators (11) and (12) have access to (s_k, t_k) as well. Thus, even when the observer chooses not to transmit b_k, it can still pass a 1-bit flag about b_k to the estimator provided that $s_k \ge 1$. If k is such that all M transmissions are concluded prior to time k (i.e., $s_k = 0$), the estimators are given by equations (9) and (10), irrespective of b_k.

Now, observe that the optimisation over the observer policies is equivalent to optimisation over the sets $\mathcal{T}_{(s_k,t_k)}$ for all k such that

$$\max\{0, M-k\} \le s_k \le \min\{t_k, M\}$$

and $t_k = N - k$. The nonnegativity of s_k is a result of the limited channel use constraint. Note that if $s_{k_0} = 0$ for some k_0, then $s_k = 0$ for all k such that $k_0 \le k \le N-1$. At the other extreme, we must have $s_k \le N - k$, since if $s_k = N - k$, this means that there are as many channel uses left as there are decision instances, and the optimal observer and estimator policies in this case are obvious.

3.3 The solution with the mean-square error criterion

Let $(s_k, t_k) = (s, t)$, and $e^*_{(s,t)}$ denote the optimal value of the estimation error (or distortion) (13) when the decision horizon is of length t, and the observer is limited to s channel uses, where $s \le t$. We know that at time k, the optimal observation policy will be of the form (16).

Now, at time $k+1$, depending on the realisation of the random variable b_k, the remaining $(t-1)$-stage estimation error is either $e^*_{(s-1,t-1)}$, or $e^*_{(s,t-1)}$. Thus, inductively by the DP equation (Bertsekas, 1995), we can write[8]

$$e^*_{(s,t)} = \min_{\mathcal{T}_{(s,t)}} \left\{ e^*_{(s-1,t-1)} \int_{b \in \mathcal{T}_{(s,t)}} f(b)\,\mathrm{d}b + e^*_{(s,t-1)} \int_{b \in \mathcal{T}^c_{(s,t)}} f(b)\,\mathrm{d}b \right.$$
$$\left. + \int_{b \in \mathcal{T}^c_{(s,t)}} \left[b - E\left\{ b \,|\, b \in \mathcal{T}^c_{(s,t)} \right\} \right]^2 f(b)\,\mathrm{d}b \right\}$$

where $f(b)$ is the probability density function (pdf) of the random variable b_k. If b_k's are discrete random variables with a probability mass function (pmf), one has to replace the integrals in the above expression with sums. Expanding out the expectation yields

$$e^*_{(s,t)} = \min_{\mathcal{T}_{(s,t)}} \left\{ e^*_{(s-1,t-1)} \int_{b \in \mathcal{T}_{(s,t)}} f(b)\,\mathrm{d}b + e^*_{(s,t-1)} \int_{b \in \mathcal{T}^c_{(s,t)}} f(b)\,\mathrm{d}b \right.$$
$$\left. + \int_{b \in \mathcal{T}^c_{(s,t)}} \left[b - \frac{\int_{b \in \mathcal{T}^c_{(s,t)}} b f(b)\,\mathrm{d}b}{\int_{b \in \mathcal{T}^c_{(s,t)}} f(b)\,\mathrm{d}b} \right]^2 f(b)\,\mathrm{d}b \right\}. \tag{17}$$

To solve for $e^*_{(s,t)}$, we first note the boundary conditions $e^*_{(t,t)} = 0$, and $e^*_{(0,t)} = t\sigma_b^2$, $\forall t \ge 0$, where σ_b^2 is the variance of b_k. The term $e^*_{(s,t)}$ remains undefined for $s > t$. The optimal sets satisfy the boundary conditions $\mathcal{T}^*_{(t,t)} = \mathcal{B}$, and $\mathcal{T}^*_{(0,t)} = \emptyset$, $\forall t \ge 0$. The recursion of equation (17) needs to be solved offline and the optimal sets $\mathcal{T}^*_{(s,t)}$ must be tabulated starting with smaller values of (s,t).[9] The solution to the original problem can then be determined as follows:

Initialise $s_0 = M$, $t_0 = N$. For each k in $0 \le k \le N-1$, do the following:

1 Look up the optimal set $\mathcal{T}^*_{(s_k,t_k)}$ from the table that was determined offline.

2 Observe b_k, and apply the observation policy

$$\bar{\mu}^*((s_k, t_k); b_k) = \begin{cases} b_k & \text{if } b_k \in \mathcal{T}^*_{(s_k,t_k)} \\ \text{NT} & \text{if } b_k \in \mathcal{T}^{*c}_{(s_k,t_k)} \end{cases}$$

3 Apply the estimation policy

$$\hat{\mu}^*_k(\mathcal{T}^*_{(s_k,t_k)}) = E\{b_k \,|\, b_k \in \mathcal{T}^{*c}_{(s_k,t_k)}\} = \frac{\int_{b \in \mathcal{T}^{*c}_{(s_k,t_k)}} b f(b)\,\mathrm{d}b}{\int_{b \in \mathcal{T}^{*c}_{(s_k,t_k)}} f(b)\,\mathrm{d}b}$$

4 Update

$$s_{k+1} = s_k - \sigma_k, \quad t_{k+1} = t_k - 1.$$

In tabulating $\mathcal{T}^*_{(s,t)}$ one should start with solving for $\mathcal{T}^*_{(1,2)}$, and the corresponding estimation error $e^*_{(1,2)}$. To determine the optimal set at (s,t), we need to know the optimal costs at $(s, t-1)$, and $(s-1, t-1)$. Hence, we can propagate our calculations as shown in Figure 3 starting with $(s,t) = (1,2)$.

Figure 3 Recursive calculation of $e^*_{(s,t)}$

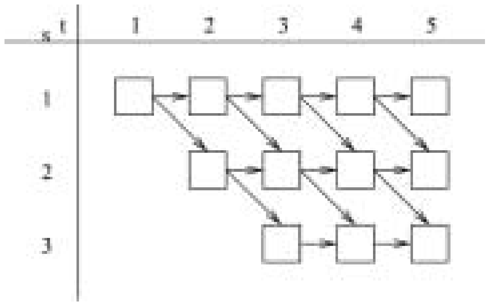

Now, we come back to the problem of minimising equation (17) over $\mathcal{T}_{(s,t)}$. Expanding out the expression inside the minimisation we get

$$e^*_{(s,t)} = e^*_{(s-1,t-1)} + \min_{\mathcal{T}^c_{(s,t)}} \left\{ -\left(e^*_{(s-1,t-1)} - e^*_{(s,t-1)}\right) \int_{b \in \mathcal{T}^c_{(s,t)}} f(b)\,\mathrm{d}b \right.$$
$$\left. + \int_{b \in \mathcal{T}^c_{(s,t)}} b^2 f(b)\,\mathrm{d}b - \frac{\left[\int_{b \in \mathcal{T}^c_{(s,t)}} b f(b)\,\mathrm{d}b\right]^2}{\int_{b \in \mathcal{T}^c_{(s,t)}} f(b)\,\mathrm{d}b} \right\} \tag{18}$$

where we used the fact that $\int_{b \in \mathcal{T}_{(s,t)}} f(b)\,\mathrm{d}b = 1 - \int_{b \in \mathcal{T}^c_{(s,t)}} f(b)\,\mathrm{d}b$.

This is an optimisation problem over measurable sets $\mathcal{T}^c_{(s,t)}$ on the real line, and since these sets are not countable, there is no known method for carrying out this minimisation in a systematic manner. One reasonable assumption is to restrict the search to sets that are in the form of simple symmetric intervals, i.e., $\mathcal{T}^c_{(s,t)} = [-\beta_{(s,t)}, \beta_{(s,t)}]$, where $0 \le \beta_{(s,t)} \le \infty$.

Now, because of symmetry, the last term on the right-hand side of equation (18) disappears from the minimisation. Differentiating the remaining terms inside the curly brackets, we obtain the first-order necessary condition:

$$-\left(e^*_{(s-1,t-1)} - e^*_{(s,t-1)}\right) f(\beta_{(s,t)}) + \beta^2_{(s,t)} f(\beta_{(s,t)}) = 0. \tag{19}$$

From which the critical point $\beta^*_{(s,t)}$ can be determined as[10]

$$\beta^*_{(s,t)} = \sqrt{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}}. \tag{20}$$

Note that, we always have $e^*_{(s,t-1)} \leq e^*_{(s-1,t-1)}$, since for the same decision horizon, $t-1$, the minimum average distortion achieved by s channel uses, is always less than that achieved by $s-1$ channel uses. So, $\beta^*_{(s,t)}$ always exists.

From the first-order condition, we observe that the objective function is strictly decreasing on the interval $[0, \beta^*_{(s,t)})$, and it is strictly increasing on the interval $(\beta^*_{(s,t)}, \infty)$. Thus, $\beta^*_{(s,t)}$ must be a strict global minimiser. Thus, in the class of symmetric intervals, the best set $\mathcal{T}^c_{(s,t)}$ is given by the interval

$$\mathcal{T}^{*c}_{(s,t)} = \Big[-\sqrt{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}}, \sqrt{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}} \Big]. \tag{21}$$

3.4 The solution with the probability of error criterion

As in Section 3.3, let $(s_k, t_k) = (s, t)$, and let $e^*_{(s,t)}$ denote the optimal value of the estimation error (or distortion) (14) when the decision horizon is of length t, and the observer is limited to s channel uses, where $s \leq t$. We know that at time k, the optimal observation (transmission) policy will be of the form (16).

Now, at time $k+1$, depending on the realisation of the random variable b_k, the remaining $(t-1)$-stage estimation error is either $e^*_{(s-1,t-1)}$ or $e^*_{(s,t-1)}$. Thus, assuming that $s \geq 1$, inductively by the DP equation, we can write

$$e^*_{(s,t)} = \min_{\mathcal{T}_{(s,t)}} \Big\{ P\big[b_k \in \mathcal{T}_{(s,t)}\big] e^*_{(s-1,t-1)} + P\big[b_k \in \mathcal{T}^c_{(s,t)}\big] e^*_{(s,t-1)} + P\big[b_k \in \mathcal{T}^c_{(s,t)}\big] - \max_{m_j \in \mathcal{T}^c_{(s,t)}} p_j \Big\}.$$

or equivalently

$$e^*_{(s,t)} = \min_{\mathcal{T}^c_{(s,t)}} \Big\{ \big(1 - P\big[b_k \in \mathcal{T}^c_{(s,t)}\big]\big) e^*_{(s-1,t-1)} + P\big[b_k \in \mathcal{T}^c_{(s,t)}\big] e^*_{(s,t-1)} + P\big[b_k \in \mathcal{T}^c_{(s,t)}\big] - \max_{m_j \in \mathcal{T}^c_{(s,t)}} p_j \Big\}.$$

Plugging in $P[b_k \in \mathcal{T}^c_{(s,t)}] = \sum_{m_i \in \mathcal{T}^c_{(s,t)}} p_i$, and rearranging the terms, we obtain the following error recursion:

$$e^*_{(s,t)} = e^*_{(s-1,t-1)} + \min_{\mathcal{T}^c_{(s,t)}} \Big\{ -\big(e^*_{(s-1,t-1)} - e^*_{(s,t-1)}\big) \sum_{m_i \in \mathcal{T}^c_{(s,t)}} p_i + \sum_{m_i \in \mathcal{T}^c_{(s,t)}} p_i - \max_{m_j \in \mathcal{T}^c_{(s,t)}} p_j \Big\}. \tag{22}$$

We next show that the error difference, $e^*_{(s-1,t-1)} - e^*_{(s,t-1)}$, can be bounded from below and above.

Proposition 2: *Suppose $1 \leq s \leq t$. Then, the error difference $e^*_{(s-1,t-1)} - e^*_{(s,t-1)}$ satisfies:*

$$0 \leq e^*_{(s-1,t-1)} - e^*_{(s,t-1)} \leq 1.$$

Proof: The lower bound can be established by observing that for the same decision horizon, $t-1$, the minimum average distortion achieved by s channel uses, is always at least as small as the one that can be achieved by $s-1$ channel uses. For the upper bound, one needs to observe that the maximum stage-wise estimation error is bounded by 1. $\square$

Using Proposition 2, we will next show that the optimum choice for the sets $\mathcal{T}^c_{(s,t)}$ is the singleton $\mathcal{T}^{c*}_{(s,t)} = \{m_{i^*}\}$, where $i^* = \arg\max_{m_i \in \mathcal{B}} p_i$. In other words, the optimal solution is not to transmit the *most likely* outcome, and transmit all the other outcomes of the source process b_k. Moreover, this policy is independent of the number of decision instances left, t_k, and the number of transmission opportunities left, s_k, provided that $s_k \geq 1$. Recall that, the optimum estimator is the MAP estimator and is given by equation (12).

In order to show that this is indeed the optimal observer (or transmission) policy, we first set the cardinality of the set $\mathcal{T}_{(s,t)}$ to $|\mathcal{T}^c_{(s,t)}| = 0$, and determine that the expression inside the curly brackets in equation (22) is just 0.

We next set $|\mathcal{T}^c_{(s,t)}| = 1$, and note that the minimisation of the function inside the curly brackets in equation (22) is equivalent to the following minimisation:

$$\min_i \left\{ (1-p_i)e^*_{(s-1,t-1)} + p_i e^*_{(s,t-1)} + p_i - p_i \right\}.$$

Here i is such that $m_i \in \mathcal{B}$. Cancelling p_i's and rearranging, we obtain an equivalent minimisation problem:

$$\min_i -p_i(e^*_{(s-1,t-1)} - e^*_{(s,t-1)}).$$

By Proposition 2, the error difference, $e^*_{(s-1,t-1)} - e^*_{(s,t-1)}$, is nonnegative; thus, the minimum is achieved by picking i as i^* given by

$$m_{i^*} = \arg\max_{m_i \in \mathcal{B}} p_i.$$

This choice yields a minimum value of $-(e^*_{(s-1,t-1)} - e^*_{(s,t-1)})p_{i^*}$. Note that this value is at least as good as the value we obtained when we set the cardinality of the set $|\mathcal{T}^c_{(s,t)}| = 0$. Thus, we never pick $\mathcal{T}^c_{(s,t)}$ such that it has zero cardinality.

Finally, we let $|\mathcal{T}^c_{(s,t)}| \geq 2$, and let $p_{\max}$ denote the particular element of $\mathcal{T}^c_{(s,t)}$ with maximal probability. That is,

$$p_{\max} = \max_{m_j \in \mathcal{T}^c_{(s,t)}} p_j.$$

Since the number of elements of $\mathcal{T}^c_{(s,t)}$ is at least 2, the minimisation problem inside the curly brackets in equation (22) can be written as

$$\min_{\mathcal{T}^c_{(s,t)}} \left\{ -(e^*_{(s-1,t-1)} - e^*_{(s,t-1)})p_{\max} + (1 - (e^*_{(s-1,t-1)} - e^*_{(s,t-1)})) \sum_{m_i \in \mathcal{T}^c_{(s,t)} \setminus m_{j^*}} p_i \right\}$$

where

$$m_{j^*} = \arg\max_{m_i \in \mathcal{T}^c_{(s,t)}} p_i.$$

Now, by Proposition 2, the term multiplying the sum $\sum_{m_i \in \mathcal{T}^c_{(s,t)} \setminus m_{j^*}} p_i$ is always nonnegative; hence, we can conclude that the above minimum is bounded from below by

$$-(e^*_{(s-1,t-1)} - e^*_{(s,t-1)})p_{\max}$$

for any choice of the set $\mathcal{T}^c_{(s,t)}$ with cardinality $|\mathcal{T}^c_{(s,t)}| \geq 2$. However the expression $-(e^*_{(s-1,t-1)} - e^*_{(s,t-1)})p_{\max}$ satisfies

$$-(e^*_{(s-1,t-1)} - e^*_{(s,t-1)})p_{\max} \geq -(e^*_{(s-1,t-1)} - e^*_{(s,t-1)})p_{i^*}$$

since $p_{i^*} \geq p_{\max}$. Therefore, the minimum value of the function inside the curly brackets in equation (22) is achieved when $\mathcal{T}^c_{(s,t)} = \{m_{i^*}\}$, as claimed.

In summary, when the distortion criterion is the probabilty of error, at time k, the optimal observer first observes the source output b_k. Then, it checks to see if $s_k \geq 1$; if so, it transmits b_k unless $b_k = m_{i^*}$, i.e., the most likely outcome. The estimator (or decoder), on the other hand, employs the MAP estimation rule given the output of the channel.

3.5 The Gaussian case

We now come back to the mean-square error distortion criterion, and study the case when the source is Gaussian in more detail. Suppose b_k's are zero-mean, i.i.d. Gaussian with the pdf

$$f(b) = \frac{1}{\sqrt{2\pi\sigma_b^2}} e^{-\frac{b^2}{2\sigma_b^2}}. \tag{23}$$

Let $\Phi(\cdot)$ denote the cumulative density function (CDF) of the standard Gaussian random variable with zero mean and unit variance. We list some properties of the Gaussian pdf here for completeness, as the derivations that follow use these expressions extensively.

- $$\int_a^b f(x)\,\mathrm{d}x = \Phi\left(\frac{b}{\sigma_x}\right) - \Phi\left(\frac{a}{\sigma_x}\right)$$
- $$\int_a^b x f(x)\,\mathrm{d}x = \frac{\sigma_x^2}{\sqrt{2\pi\sigma_x^2}}\left(e^{-\frac{a^2}{2\sigma_x^2}} - e^{-\frac{b^2}{2\sigma_x^2}}\right)$$
- $$\int_a^b x^2 f(x)\,\mathrm{d}x = \sigma_x^2\left[\Phi\left(\frac{b}{\sigma_x}\right) - \Phi\left(\frac{a}{\sigma_x}\right)\right] + \frac{\sigma_x^2}{\sqrt{2\pi\sigma_x^2}}\left(ae^{-\frac{a^2}{2\sigma_x^2}} - be^{-\frac{b^2}{2\sigma_x^2}}\right).$$

In the Gaussian case, we can generalise our search for an optimum in equation (18) to more general intervals of the form $\mathcal{T}^c_{(s,t)} = [\alpha_{(s,t)}, \beta_{(s,t)}]$, where $-\infty \leq \alpha_{(s,t)} \leq \beta_{(s,t)} \ +\infty$.

Figure 4 shows the plot of the objective function on the right-hand side of equation (18) for the case when $\mathcal{T}_{(s,t)} = [a,b]$, $\sigma_b^2 = 1$, $e^*_{(s-1,t-1)} = 3$, and $e^*_{(s,t-1)} = 1$. Note that the minimum occurs at $b^* = -a^* = \sqrt{3-1} = \sqrt{2} = 1.4142$. Thus, even though we did not restrict ourselves to symmetric intervals, the solution is still a symmetric interval around zero. To show that this is indeed the case in general, one needs to differentiate the objective function inside the curly brackets in equation (18) with respect to both $\alpha_{(s,t)}$ and $\beta_{(s,t)}$, and show that the minimum occurs at $\beta^*_{(s,t)} = -\alpha^*_{(s,t)}$ when $f(b)$ is the Gaussian pdf (Imer, 2005).

Figure 4 Plot of the objective function in the Gaussian case with $\mathcal{T}_{(s,t)} = [a, b]$ when $\sigma_b^2 = 1$, $e^*_{(s-1,t-1)} = 3$, and $e^*_{(s,t-1)} = 1$ (see online version for colours)

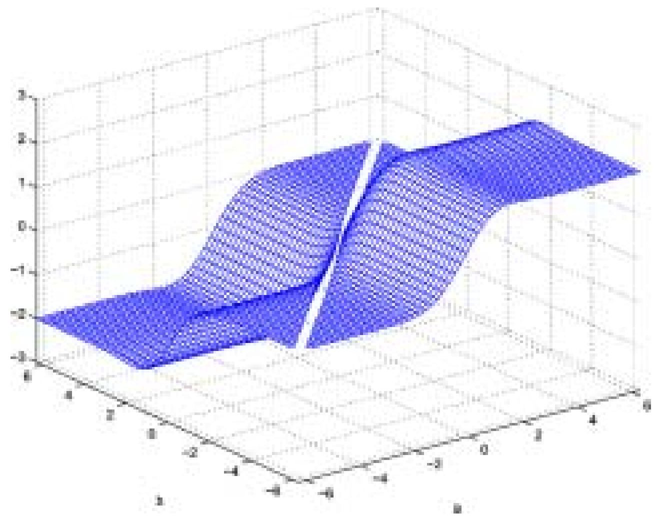

To evaluate the optimum estimation error $e^*_{(s,t)}$ in terms of $e^*_{(s-1,t-1)}$ and $e^*_{(s,t-1)}$, we substitute the optimum interval solution (21) into the right-hand side of equation (18), and use the standard properties of the Gaussian density that we listed above to obtain

$$\begin{aligned} e^*_{(s,t)} = e^*_{(s-1,t-1)} &- \left[e^*_{(s-1,t-1)} - e^*_{(s,t-1)} - \sigma_b^2\right] \\ &\times \left[2\Phi\left(\sqrt{\frac{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}}{\sigma_b^2}}\right) - 1\right] \\ &- \frac{2\sigma_b^2}{\sqrt{2\pi\sigma_b^2}} \sqrt{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}}\, e^{-\frac{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}}{2\sigma_b^2}}. \end{aligned} \tag{24}$$

We can normalise the optimal estimation error by letting

$$\epsilon_{(s,t)} = \frac{e^*_{(s,t)}}{\sigma_b^2}. \tag{25}$$

and rewrite the recursion (24) in a simpler form:

$$\begin{aligned} \epsilon_{(s,t)} = \epsilon_{(s-1,t-1)} &- \left[\epsilon_{(s-1,t-1)} - \epsilon_{(s,t-1)} - 1\right]\left[2\Phi\left(\sqrt{\epsilon_{(s-1,t-1)} - \epsilon_{(s,t-1)}}\right) - 1\right] \\ &- \frac{2}{\sqrt{2\pi}} \sqrt{\epsilon_{(s-1,t-1)} - \epsilon_{(s,t-1)}}\, e^{-\frac{\epsilon_{(s-1,t-1)} - \epsilon_{(s,t-1)}}{2}} \end{aligned} \tag{26}$$

with the initial conditions

$$\epsilon(t,t) = 0, \quad \epsilon(0,t) = t, \quad \forall t \geq 0$$

and $\epsilon(s,t)$ is undefined for $s > t$.

Hence, we can provide a solution to the problem of optimal sequential estimation of an i.i.d. Gaussian process of finite length over a noiseless channel that can only be used a limited number of times. First, a table has to be formed by an offline numerical computation of the recursion equation (26). Then, the table can be scaled, if needed, to the actual variance of the process through equation (25). Next, the transmission intervals for the observer are determined via equation (20), and tabulated for all feasible pairs (s,t). For the online computation, as illustrated in Section 3.3, the observer has to keep two states, (s_k, t_k). Each time unit k, after observing the realisation of the random variable b_k, the observer compares the realised value of the random variable to the optimum decision interval corresponding to the current state (s_k, t_k), and makes a transmission decision. The estimator, on the other hand, has access to the same tabulated values of the transmission intervals, $\mathcal{T}^*_{(s,t)}$, and it keeps track of the states (s_k, t_k) in the same way the observer does. Upon receiving the transmitted data, y_k, from the channel, the estimator simply applies the estimation policy given in Section 3.3.

3.6 Gaussian case with noisy measurements

Let the source process $\{b_k\}$ be i.i.d. Gaussian. If the observer has access to a noisy version of the source output, i.e.,

$$z_k = b_k + v_k$$

where $\{v_k\}$ is zero-mean, i.i.d. Gaussian[11] with variance σ_v^2, the optimisation problem with the mean-square distortion measure can be solved using a similar approach. In this case, the observer's decision as to whether to use the channel to transmit or not depends on the available data z_k. In the derivation of the optimal observer-estimator pair, most of the analysis of Section 3.5 carries over.

In order to see that the structure of the solution is preserved, first observe that when $s_k = 0$, $\hat{\mu}_k^* = 0$, and for $s_k \geq 1$, the optimal estimator has the form:

$$\hat{\mu}_k^*(I_k^d) = E\{b_k \mid (s_k, t_k); x_k\}.$$

Substituting this into the error expression, and following along the lines of Proposition 1, one can see that the optimal observer policy has the form

$$\mu_k(I_k^e) = \bar{\mu}((s_k, t_k); z_k).$$

In other words $\{(s_k, t_k); z_k\}$ is a sufficient statistics for the optimal policy $\mu_k^*(I_k^e)$.

Since μ_k depends explicitly only on the current measurement z_k, for $s_k \geq 1$, the search for an optimal encoder policy can be narrowed down to the class of policies of the form

$$\mu_k(I_k^e) = \bar{\mu}((s_k, t_k); z_k) = \begin{cases} z_k & \text{if } z_k \in \mathcal{T}_{(s_k,t_k)} \\ \text{NT} & \text{if } z_k \in \mathcal{T}^c_{(s_k,t_k)} \end{cases}. \tag{27}$$

where $\mathcal{T}_{(s_k,t_k)}$ is a measurable set on $\mathcal{B}$, and is a function of (s_k, t_k). Since $x_k = \mu_k(I_k^e)$, for $s_k \geq 1$, we can write the optimal estimator as

$$\hat{\mu}_k((s_k,t_k);x_k) = \begin{cases} \frac{\sigma_b^2}{\sigma_b^2+\sigma_v^2} z_k & \text{if } z_k \in \mathcal{T}_{(s_k,t_k)} \\ E\{b_k \mid z_k \in \mathcal{T}^c_{(s_k,t_k)}\} & \text{if } z_k \in \mathcal{T}^c_{(s_k,t_k)} \end{cases}. \tag{28}$$

We proceed as in Section 3.3, and write the dynamic programming recursion governing the evolution of the optimal estimation error as follows:

$$\begin{aligned} e^*_{(s,t)} = \min_{\mathcal{T}^c_{(s,t)}} \Big\{ & e^*_{(s-1,t-1)} P[z \in \mathcal{T}_{(s,t)}] + \sigma_b^2 \\ & + \left(\frac{\sigma_b^2}{\sigma_b^2+\sigma_v^2}\right)^2 \int_{z\in\mathcal{T}_{(s,t)}} z^2 f_Z(z)\,\mathrm{d}z + e^*_{(s,t-1)} P[z \in \mathcal{T}^c_{(s,t)}] \\ & - 2\frac{\sigma_b^2}{\sigma_b^2+\sigma_v^2} \int_{z\in\mathcal{T}_{(s,t)}} zE[b \mid z] f_Z(z)\,\mathrm{d}z \Big\} \end{aligned}$$

where $f_Z(z) \sim N(0, \sigma_b^2 + \sigma_v^2)$, and $f_{B\,|\,Z}(b\,|\,z) \sim N\big(\frac{\sigma_b^2}{\sigma_b^2+\sigma_v^2} z, \frac{\sigma_b^2\sigma_v^2}{\sigma_b^2+\sigma_v^2}\big)$. The recursion can be simplified as

$$\begin{aligned} e^*_{(s,t)} = \min_{\mathcal{T}^c_{(s,t)}} \Big\{ & e^*_{(s-1,t-1)} + \sigma_b^2 - (e^*_{(s-1,t-1)} - e^*_{(s,t-1)}) \\ & \int_{z\in\mathcal{T}^c_{(s,t)}} f_Z(z)\,\mathrm{d}z - \left(\frac{\sigma_b^2}{\sigma_b^2+\sigma_v^2}\right)^2 (\sigma_b^2+\sigma_v^2) \\ & + \left(\frac{\sigma_b^2}{\sigma_b^2+\sigma_v^2}\right)^2 \int_{z\in\mathcal{T}^c_{(s,t)}} z^2 f_Z(z)\,\mathrm{d}z \Big\}. \end{aligned}$$

Following along the lines of Section 3.5, we restrict our search for an optimum set to simple intervals, i.e., $\mathcal{T}^c_{(s,t)} = [\alpha_{(s,t)}, \beta_{(s,t)}]$. The same analysis gives the optimum choice for $\beta_{(s,t)}$

$$\beta^*_{(s,t)} = \frac{\sigma_b^2+\sigma_v^2}{\sigma_b^2} \sqrt{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}}$$

and $\alpha^*_{(s,t)} = -\beta^*_{(s,t)}$. Substituting these values into the error recursion, we obtain the two-dimensional recursion for the estimation error:

$$\begin{aligned} e^*_{(s,t)} = e^*_{(s-1,t-1)} & - \left[e^*_{(s-1,t-1)} - e_{(s,t-1)} - \frac{(\sigma_b^2)^2}{\sigma_b^2+\sigma_v^2} \right] \\ & \times \left[2\Phi\left(\frac{\sqrt{\sigma_b^2+\sigma_v^2}\sqrt{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}}}{\sigma_b^2} \right) - 1 \right] + \sigma_b^2 - \frac{(\sigma_b^2)^2}{\sigma_b^2+\sigma_v^2} \\ & - \frac{2\sigma_b^2}{\sqrt{2\pi(\sigma_b^2+\sigma_v^2)}} \sqrt{e^*_{(s-1,t-1)} - e^*_{(s,t-1)}} \; e^{-\frac{\left(\frac{\sigma_b^2+\sigma_v^2}{\sigma_b^2}\right)^2 (e^*_{(s-1,t-1)} - e^*_{(s,t-1)})}{2\left(\sigma_b^2+\sigma_v^2\right)}}. \end{aligned}$$

Note that, for $\sigma_v^2 = 0$ this recursion simplifies to equation (24), which is the recursion for the perfect state measurements.

We can normalise the optimal estimation error by letting

$$\epsilon_{(s,t)} = \frac{\sigma_b^2 + \sigma_v^2}{(\sigma_b^2)^2} e^*_{(s,t)} \tag{29}$$

and rewrite the above recursion in a simpler form:

$$\begin{aligned}\epsilon_{(s,t)} = \epsilon_{(s-1,t-1)} - \left[\epsilon_{(s-1,t-1)} - \epsilon_{(s,t-1)} - 1\right]\left[2\Phi\left(\sqrt{\epsilon_{(s-1,t-1)} - \epsilon_{(s,t-1)}}\right) - 1\right] \\ + \frac{\sigma_v^2}{\sigma_b^2} - \frac{2}{\sqrt{2\pi}}\sqrt{\epsilon_{(s-1,t-1)} - \epsilon_{(s,t-1)}} \quad e^{-\frac{\epsilon_{(s-1,t-1)} - \epsilon_{(s,t-1)}}{2}}\end{aligned} \tag{30}$$

with the initial conditions

$$\epsilon(t,t) = \frac{\sigma_v^2}{\sigma_b^2} t, \quad \epsilon(0,t) = \left(1 + \frac{\sigma_v^2}{\sigma_b^2}\right) t, \quad \forall t \geq 0$$

and $\epsilon(s,t)$ is undefined for $s > t$.

We note that the recursion (30) reduces to the recursion (26), as the noise variance $\sigma_v^2 \to 0$.

4 Estimating a Gauss-Markov process with limited measurements

In this section, we discuss the case when the source process is Markov

$$b_{k+1} = Ab_k + w_k$$

driven by an i.i.d. Gaussian process $\{w_k\}$ with zero-mean. The solution to this case is similar to the Gaussian i.i.d. case when the observer has access to the source output b_k without noise. The only difference is that, now the observer-estimator pair has to keep track of three variables (r_k, s_k, t_k), where r_k keeps track of the number of time units passed since the last use of the channel for transmission. A similar DP recursion, now in three dimensions, can be obtained.

Let r denote the number of time units passed since the last transmission of a source output. Reasoning as in Section 3.4, we can deduce that for $s \geq 1$, the optimal estimator has the form

$$\hat{\mu}((r,s,t); b_{N-t}) = \begin{cases} b_{N-t} & b_{N-t} \in \mathcal{T}_{(r,s,t)} \\ E\{b_{N-t} \mid b_{N-t} \in \mathcal{T}^c_{(r,s,t)}\} & b_{N-t} \in \mathcal{T}^c_{(r,s,t)} \end{cases}. \tag{31}$$

With the estimator structure in place, the error recursion can be derived following along the lines of previous sections:

$$\begin{aligned} e^*_{(r,s,t)} = \min_{\mathcal{T}_{(r,s,t)}} \Big\{ & e^*_{(1,s-1,t-1)} P[b_{N-t} \in \mathcal{T}_{(r,s,t)}] + e^*_{(r+1,s,t-1)} P[b_{N-t} \in \mathcal{T}^C_{(r,s,t)}] \\ & + \int_{b_{N-t} \in \mathcal{T}^c_{(r,s,t)}} \left[b_{N-t} - A^r b_{N-t-r}\right]^2 f_{b_{N-t}}(b_{N-t}) \, \mathrm{d}b_{N-t} \Big\} \end{aligned}$$

where $b_{N-t} \sim N(A^r b_{N-t-r}, \left(\sum_{k=1}^{r} A^{2(k-1)}\right)\sigma_b^2)$.

Now if we let $\mathcal{T}^c_{(r,s,t)} = [\alpha_{(r,s,t)}\ \beta_{(r,s,t)}]$, the optimal choices for the parameters $\alpha_{(r,s,t)}$ and $\beta_{(r,s,t)}$ are

$$\alpha^*_{(r,s,t)} = A^r b_{N-t-r} + \sqrt{e^*_{(1,s-1,t-1)} - e^*_{(r+1,s,t-1)}}$$
$$\beta^*_{(r,s,t)} = A^r b_{N-t-r} - \sqrt{e^*_{(1,s-1,t-1)} - e^*_{(r+1,s,t-1)}}.$$

Substituting these choices back into the error recursion and simplifying yields

$$\begin{aligned} e^*_{(r,s,t)} &= e^*(1,s-1,t-1) - \left[e^*_{(1,s-1,t-1)} - e^*_{(r+1,s,t-1)} - \left(\sum_{k=1}^{r} A^{2(k-1)}\right)\sigma_b^2\right] \\ &\quad \times \left[2\Phi\left(\sqrt{\frac{e^*_{(1,s-1,t-1)} - e^*_{(r+1,s,t-1)}}{\sum_{k=1}^{r} A^{2(k-1)}\sigma_b^2}}\right) - 1\right] - \frac{2\sqrt{\sum_{k=1}^{r} A^{2(k-1)}\sigma_b^2}}{2\pi} \\ &\quad \times \sqrt{e^*_{(1,s-1,t-1)} - e^*_{(r+1,s,t-1)}}\ \ e^{-\frac{e^*_{(1,s-1,t-1)} - e^*_{(r+1,s,t-1)}}{2\sum_{k=1}^{r} A^{2(k-1)}\sigma_b^2}} \end{aligned} \tag{32}$$

where we have made use of the fact that for a Gaussian random variable x with mean m and variance k^2, and for $a \geq 0$, we have the expression

$$\int_{m-\sqrt{a}}^{m+\sqrt{a}} (x-m)^2 f_x(x)\,\mathrm{d}x = -\frac{2\sqrt{a}k}{\sqrt{2\pi}} e^{-\frac{a}{2k^2}} + k^2\left(2\Phi\left(\frac{\sqrt{a}}{k}\right) - 1\right).$$

The recursion (32) is defined for $r \geq 1$, and $0 \leq s \leq t$ with the boundary conditions given by

$$e^*_{(r,t,t)} = 0, \quad e^*_{(r,0,t)} = \left(\sum_{l=r}^{r+t-1}\sum_{k=1}^{l} A^{2(k-1)}\right)\sigma_b^2. \tag{33}$$

As in Section 3.5, one can define the normalised estimation error by

$$\epsilon_{(r,s,t)} = \frac{1}{\left(\sum_{k=1}^{r} A^{2(k-1)}\right)\sigma_b^2} e^*_{(r,s,t)} \tag{34}$$

and simplify the recursion (32) further as follows:

$$\begin{aligned} \epsilon_{(r,s,t)} &= \epsilon_{(1,s-1,t-1)} - \left[\epsilon_{(1,s-1,t-1)} - \epsilon_{(r+1,s,t-1)} - 1\right] \\ &\quad \times \left[2\Phi\left(\sqrt{\epsilon_{(1,s-1,t-1)} - \epsilon_{(r+1,s,t-1)}}\right) - 1\right] \\ &\quad - \frac{2}{\sqrt{2\pi}}\sqrt{\epsilon_{(1,s-1,t-1)} - \epsilon_{(r+1,s,t-1)}}\, e^{-\frac{\epsilon_{(1,s-1,t-1)} - \epsilon_{(r+1,s,t-1)}}{2}}. \end{aligned}$$

Note that when $r = 0$, this is the exact same recursion as in the case of estimating an i.i.d. Gaussian process with no measurement noise. The only difference between this case and the i.i.d. case is in scaling back into the original estimation error via equation (34). However, unlike the i.i.d. case, this recursion must be solved offline for all feasible (r, s, t) triplets, and a three-dimensional table has to be formed.

5 Illustrative examples

5.1 Example 1

As an example for the case when the source is binary, i.e., $b_k \in \{0, 1\}$, consider the problem of sequentially estimating a Bernoulli process of length N with M opportunites to transmit over a noiseless binary channel. This problem is a special case of the general problem we solved in Section 3.4. The probability distribution of the source is given, and say, without loss of any generality, that 1 is a more likely outcome than 0. In this case, the best observation policy is to start at time $k = 0$ not transmit the likely outcome 1, and to use the channel to transmit only the unlikely outcome 0. And the best estimation scheme is to employ the MAP estimator which estimates NT as 1, and 0 as 0, as long as $s_k \geq 1$. If $s_k = 0$, on the other hand, then the best estimator should estimate 1 regardless of the channel output.

5.2 Example 2

The second example is just solving the problem of Section 3.5 for $(s, t) = (1, 2)$. So, the observer can use the channel for tranmission only once, at time $k = 0$ or 1, and the observer and the estimator are jointly trying to minimise the average distortion (or estimation error):

$$e = E\{(b_0 - \hat{b}_0)^2 + (b_1 - \hat{b}_1)^2\}$$

where b_0, b_1 are i.i.d. Gaussian with zero mean, and variance σ_b^2. If we arbitrarily choose to transmit the first source output, or the second one, the estimation error would be

$$e^*_{\text{no-observer}} = \sigma_b^2$$

which is the best error that can be achieved without a decision maker that observes the source output. Now, suppose the observer is aware of the fact that the estimator knows the *a priori* distribution of b_0. So, it makes sense for the observer not to transmit the realised value of b_0 if this value happens to be close to the *a priori* estimate of it, which in this case is the mean value of b_0, i.e., zero.

Motivated by this intuition, the observer decides to adopt a policy in which it will not use the channel to transmit b_0 if it lies in an interval $[\alpha, \beta]$ around zero. Note that the decision for the second stage would already have been made once α and β are determined, because, if $b_0 \in [\alpha, \beta]$, then the observer cannot use the channel to transmit at time 1, and if $b_0 \notin [\alpha, \beta]$, there is no reason why it should not transmit at time 1.

Now, the optimisation problem faced by the observer is to choose α and β such that the following error is minimised:

$$e_{(\alpha,\beta)} = \int_{\alpha}^{\beta} (b - E\{b \,|\, b \in [\alpha, \beta]\})^2 f(b)\,\mathrm{d}b + \sigma_b^2 P\{b_0 \notin [\alpha, \beta]\}$$

where $f(b)$ is the Gaussian density given by equation (23). The solution can easily be obtained by checking the first- and second-order optimality conditions, and is given by

$$(\alpha^*, \beta^*) = (-\sigma_b, \sigma_b).$$

Thus, the observer should not use the channel to transmit the source output b_0 if it falls within one standard deviation of its mean. For these values of α and β, the optimal value of the estimation error can be calculated as

$$e_{(\alpha^*,\beta^*)} = \sigma_b^2\left[1 - \sqrt{\frac{2}{\pi e}}\right]. \tag{35}$$

Comparing this error to the no-observer policy, $e^*_{\text{no-observer}} = \sigma_b^2$, we see that there is an approximately $\sqrt{\frac{2}{\pi e}} \approx 48\%$ improvement in the estimation error.

5.3 Example 3

The third and final example we will discuss involves the following design problem. We are given a time horizon of a fixed length N, say 100 For this $N = 100$ time units, we would like to sequentially estimate the state of a zero-mean, i.i.d. Gaussian process with unit variance. We have a design criterion which says that the *aggregate* estimation error should not exceed 20. The solution to this problem without an observer agent is to reveal 80 arbitrary observations to the estimator and achieve an aggregate estimation error of 20. Suppose, now we use the optimal observer-estimator pair. In Figure 5, we plot the optimal value of the 100-stage estimation error for different values of M.

Figure 5 Optimal 100-stage estimation error vs. the number of allowed channel uses

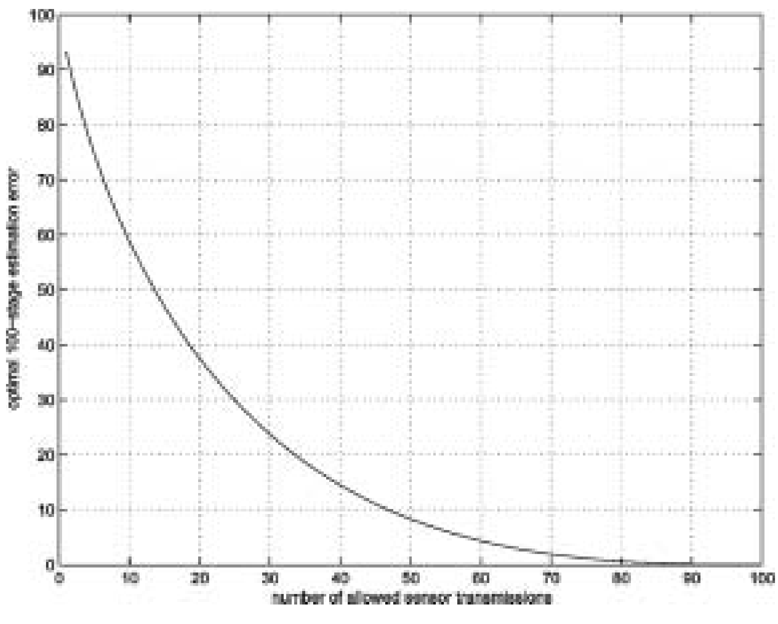

It is striking that a cumulative estimation error of 20 can be achieved with only 34 transmissions. This is approximately a $\frac{80-34}{80} \times 100 \approx 58\%$ improvement over the no-observer policy.

In order to verify our design, we simulate the optimal observer and estimator policies in Matlab. Figures 6 and 7 show typical sample paths of the optimal number of channel uses left for a decision horizon of length $N = 100$, and a limited, $M = 34$, number of channel uses. The sample paths depend on the realisation of the random sequence $\{b_k\}_0^{N-1}$.

Figure 6 A typical sample path of the number of channel uses left under the optimal observer-estimator policies $(N, M) = (100, 34)$

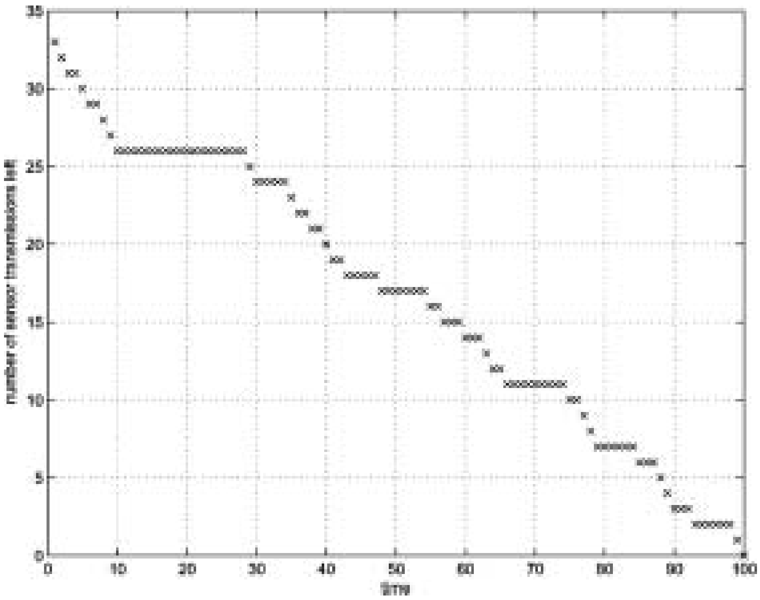

Figure 7 Another typical sample path of the number of channel uses left under the optimal observer-estimator policies $(N, M) = (100, 34)$

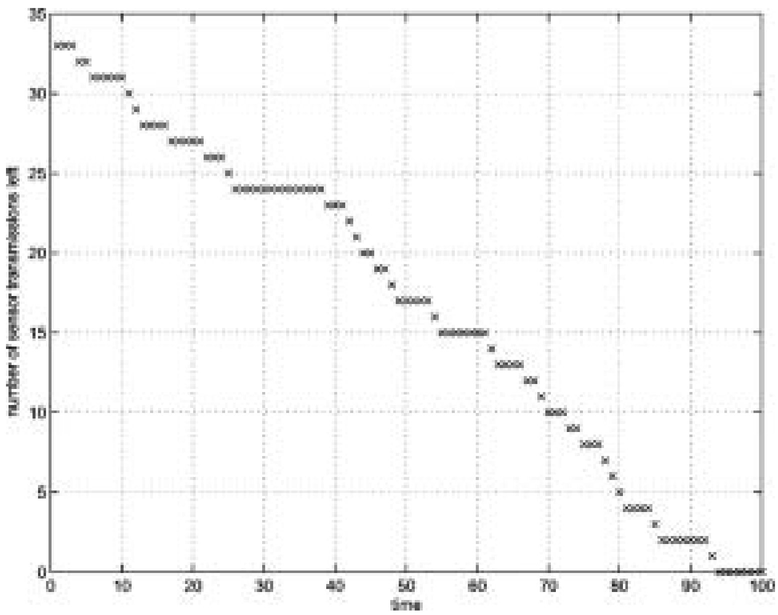

Also, to verify that $M = 34$ channel uses suffice to guarantee a cumulative average estimation error of 20, we simulate 500 instances of the estimation problem with $(N, M) = (100, 34)$. In Figure 8, we plot the realised value of the cumulative estimation error at each run for the optimal observer-estimator case and the no-observer case. Note that the sample path average of the cumulative estimation error (shown as the dotted line) is just below 20, in agreement with the theoretical findings.

Figure 8 Comparison of the sample path cumulative estimation errors between the cases when no observer is present, and when the optimal observer-estimator is used $(N, M) = (100, 34)$

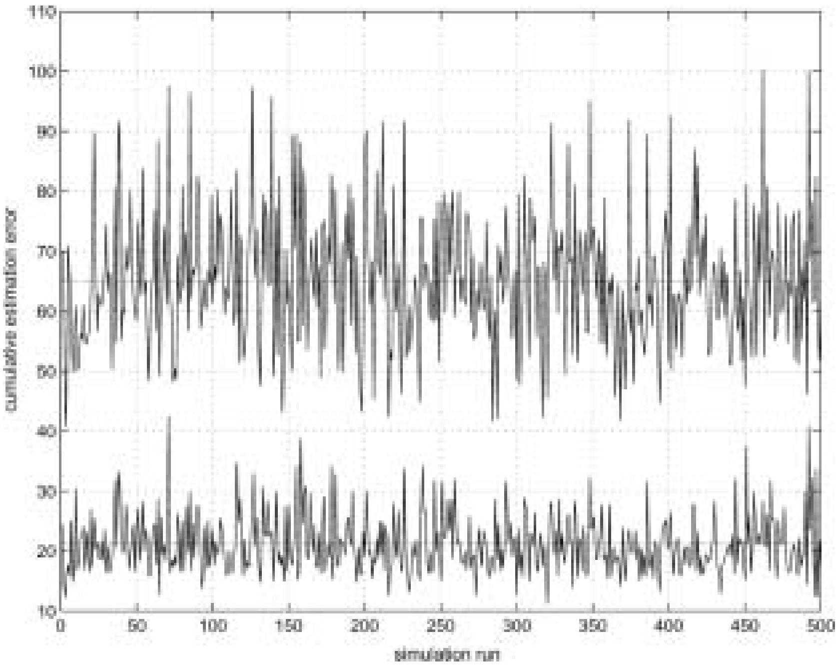

6 Conclusions

In this paper, we introduced a new hard-constrained sequential estimation problem with potential applications in wireless sensing and monitoring systems. The set-up is applicable to other usage limited decision-making scenarios as well, such as a remote control problem where the control signal is sent to the remote plant over a wireless channels which has to be used sparingly (Imer and Başar, 2006a), or when sensing and control are conducted over the same channel and hence one is faced with a tradeoff between the two (Imer and Başar, 2006b). Extensions to other types of scenarios are also possible, such as when the processes are correlated across time in a usage-limited estimation context or when multiple controllers have to share a common communication link in communicating with the plant in a remote control context, which are topics of ongoing research.

Acknowledgements

Research that led to this paper was supported by the NSF grant NSF CCR 00-85917.

References

Anderson, B.D. and Moore, J.B. (1979) *Optimal Filtering*, Prentice-Hall Inc., Englewood Cliffs, NJ.

Athans, M. (1972) 'On the determination of optimal costly measurement strategies for linear stochastic systems', *Automatica*, Vol. 8, pp.397–412.

Azimi, B., Sexton, D., Liu, P. and Mahony, M. (2006) 'Interference effect on IEEE 802.15.4 performance', *Proceedings of the 3rd International Conference on Networked Sensing Systems (INSS'06)*, Chicago, IL, Available online at www.cs.berkeley.edu/~culler/AIIT/papers/ardio/Azimi-Sadjadi%20Interference.pdf

Bansal, R. and Başar, T. (1989) 'Simultaneous design of measurement and control strategies in stochastic systems with feedback', *Automatica*, Vol. 45, No. 5, pp.679–694.

Bertsekas, D.P. (1995) *Dynamic Programming and Optimal Control*, Athena Scientific, Belmont, MA.

Cogill, R., Lall, S. and Hespanha, J. (2007) 'A constant factor approximation algorithm for optimal estimation subject to communication costs', *American Control Conference (ACC)*, New York, NY, pp.305–311.

Cover, T.M. and Thomas, J.A. (1991) *Elements of Information Theory*, John Wiley and Sons, Inc., New York, NY.

Gutierrez, J.A., Callaway, E.H. and Barrett, R.L. (2003) *Low-Rate Wireless Personal Area Networks: Enabling Wireless Sensors with IEEE 802.15.4*, IEEE Press, New York, NY.

Imer, O.C. (2005) *Optimal Estimation and Control under Communication Network Constraints*, PhD Dissertation, University of Illinois, Urbana, IL.

Imer, O.C. and Başar, T. (2005) 'Optimal estimation with limited measurements', *44th IEEE Conference on Decision and Control (CDC) and European Control Conference (ECC)*, Seville, Spain, pp.1029–1034.

Imer, O.C. and Başar, T. (2006a) 'Optimal control with limited controls', *American Control Conference (ACC)*, Minneapolis, MN, pp.298–303.

Imer, O.C. and Başar, T. (2006b) 'To measure or to control: optimal control with scheduled measurements and controls', *American Control Conference (ACC)*, Minneapolis, MN, pp.1003–1008.

Imer, O.C. and Başar, T. (2007) 'Wireless sensing with power constraints', *Advances in Control Theory and Applications*, Lecture Notes in Control and Information Sciences, Vol. 353, Springer-Verlag, pp.129–160.

Mehra, R.K. (1976) 'Optimization of measurement schedules and sensor designs for linear dynamic systems', *IEEE Transactions on Automatic Control*, Vol. 21, No. 1, pp.55–64.

Rabi, M., Baras, J.S. and Moustakides, G. (2006) 'Multiple sampling for estimation on a finite horizon', *45th IEEE Conference on Decision and Control (CDC)*, San Diego, CA, pp.1351–1357.

Stralen, N.V., Imer, O.C., Mitchell, R., Evans, S. and Iyer, S. (2006) 'A multiband random access messaging protocol', *Military Communications Conference (MILCOM)*, Washington DC, pp.1–6.

US Department of Energy (DOE) (2004) 'Industrial wireless technology for the 21st century', based on views of the industrial wireless community, in collaboration with U.S. Department of Energy (DOE).

Xu, Y. and Hespanha, J. (2004) 'Communication logics for networked control systems', *American Control Conference (ACC)*, Boston, MA, pp.572–577.

Xu, Y. and Hespanha, J. (2005) 'Estimation under uncontrolled and controlled communications in networked control systems', *44th IEEE Conference on Decision and Control (CDC)*, Seville, Spain, pp.842–847.

Notes

[1] As we show next, in a communication-theoretic setting we may call them an *encoder* and a *decoder*, respectively.

[2] Independent across time and from the source output process $\{b_k\}$.

[3] Or depending on the application, an encoder-decoder pair $(\mathcal{E}, \mathcal{D})$.

[4] This is not restrictive, as the known mean can be subtracted out by the estimator.

[5] Note that we do not distinguish between the source and user sets.

[6] Here, we provide a proof for the error criterion (13). An identical proof can be constructed for the probability of error distortion criterion (14).

[7] As long as k is such that all M measurements are not exhausted, i.e., $s_k \geq 1$.

[8] Assuming that the random variables $\{b_k\}$ are continuous with a well-defined probability density function (pdf), $f(b)$.

[9] Note that $(1, 2)$ is the smallest possible nontrivial value.

[10] The other critical point, namely $\beta_{(s,t)} = +\infty$, yields a larger cost.

[11] We also assume that the processes $\{b_k\}$ and $\{v_k\}$ are independent.