43rd IEEE Conference on Decision and Control
December 14-17, 2004
Atlantis, Paradise Island, Bahamas

Sampling of Diffusion Processes for Real-time Estimation

Maben Rabi and John S. Baras

***Abstract*— This paper addresses the causal sampling of observations of a diffusion process that results in a good quality continuous estimator based upon these samples. The optimal sampling scheme with a fixed number of samples is found by solving an *optimal (multiple) stopping problem*. This is solved explicitly in a special case. A class of threshold approximations is also described.**

I. Estimation and Control with Data-rate Constraints

In this paper, we focus our attention on packetization (sampling) and on estimation based on the generated packets in a special *Networked Control/Estimation System*. We have a sensor that makes continuous observations (y_t) of a diffusion state process (x_t). On $[0,T]$,

$$dx_t = f(x_t)dt + g(x_t)dW_t, \qquad (1)$$
$$dy_t = h(x_t)dt + dV_t. \qquad (2)$$

With $x_0 \sim \pi_0(x)dx$, $y_0 = 0$, $x_t \in \mathbb{R}^n$, $y_t \in \mathbb{R}^m$, $W_t \in \mathbb{R}^n$, $V_t \in \mathbb{R}^m$, W and V being standard Wiener processes, with g being positive definite: $g(x)g(x)^T > 0 \ \forall x \in \mathbb{R}^n$, and with f, g, h and π_0 being such that the conditional probability density of x_t given $\{y_s | 0 \leq s \leq t\}$ exists. The sensor has to transmit to a supervisor, at times it chooses in $[0,T]$, *data packets* that contain condensed information that will be useful for the supervisor to estimate the state at current and future times. At all times in $[0,T]$, the supervisor computes an estimate (filter) of the current state given the record of packets (contents of the packets including the sampling times) received thus far from the sensor. The strategy used by the sensor to choose the times at which to sample the the observation process is known to the supervisor as well. The real-time estimate at the supervisor could be used to compute a certainty-equivalence continuous feedback control signal that can be relayed to the plant without communication constraints. The only constraint on communication rates in this setup is a limit on the rate of packets sent from the sensor to the supervisor. This limit apart, the data packet link from the sensor to the supervisor is to be considered a lossless, zero-delay packet pipeline.

In this paper, we will study the optimal causal packetization of the sensor observation stream for minimizing the distortion between the continuous time filters at the supervisor and at the sensor. This is a problem that can be found in many systems control design charts today. In these situations, frequently, there are basic limitations on the information exchange pipelines such as costs on the usage of bandwidth, energy and power, or a limited ability to process all the information that can be gathered. There could also be a task-induced need to minimize communication in this distributed setting because of reasons like 'keeping the voices low' when using a distributed sensor bed for sensing. Here is a list of practical scenarios where this estimation problem with limited communication appears: a) Allocation of average packet rates to individual nodes on an automobile CAN [5], [6], b) Sampling individual sensors in a MEMS array [1], c) Scheduling packets in a wireless sensor network [2], and d) Multi-agent collaborative tracking and sensing with communication rate constraints [4].

A. Packetization of measurements

The digital representation of observations for communication or other purposes introduces a loss of information that diminishes estimation and control performance. In recursive estimation and control problems, the *signal to noise ratio* of the received information as well as the *timeliness* of the information are vital for efficient use. Digitization affects both. The sampling introduces periods of virtual information black-out[1] and a coarse quantization introduces more noise than a fine quantization. The effect of digitization on estimation, detection and control performance has been studied by some researchers basically as the effect of measurements made piece-wise constant with the jumps made at sampling times [8], [9], [10], [11]. In this framework, we can describe what the optimal sampling rate and quantization scheme is, or what the minimal rates should be to guarantee stabilization or boundedness of estimation error measures, especially for linear systems.

The transmission of measurements using a packet communication scheme has some features that simplify the analysis. In most scenarios, the packets are of uniform size and even when of variable size, have at least a few bytes of header and trailer files. These segments of the packet carry source and destination node addresses, a time stamp at origin, some error control coding, some higher layer (link and transport layers in the terminology of data networks) data blocks and any other bits/bytes that are essential for the functioning of the packet exchange scheme but which nevertheless constitute what is clearly an overhead. The payload or actual measurement information in the packet should then be at least of the same size as these 'bells

This research was supported by the Army Research Office under the ODDR&E MURI 01 Program Grant No. DAAD 19-01-1-0465 to the Center for Networked Communicating Control Systems (through Boston University), and under ARO Grant No. DAAD 190210319.

M. Rabi and J. S. Baras are with the Institute for Systems Research and the Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742.

[1]In packetizing schemes that are adapted to the measurements, there could actually be some useful information even in the non-arrival of packets (resulting in a *Timing channel*).

0-7803-8682-5/04/$20.00 ©2004 IEEE

and whistles'. It costs only negligibly more in terms of network resources, time, or energy to send a payload of five or ten bytes instead of two bits or one byte when the overhead part of the packet is already 5 bytes. In other words, in essentially all packet communication schemes, the right unit of communication cost is the cost of transmitting a single packet whether or not the payload is longer by a few bytes. This means that the samples being packetized can be quantized with very fine detail, say with 4 bytes, a rate at which the quantization noise can be ignored for low dimensional variables. The actual effect of this fine quantization could be investigated perhaps along the lines of [12]. For Markov state processes, dealt with in this paper, this will actually mean that all of these bytes of payload can be used to specify the latest state estimate. An example of an information-constrained problem where this argument fails is the TCP-RED congestion control problem where the state information is carried by a single bit in the whole packet in which the real payload is irrelevant to the congestion state.

In these packetized schemes, the other design variable left then is the sampling scheme.

B. Sampling strategy: predetermined or adapted ?

The question of when to sample and packetize is quite important for the resulting performance. Periodic sampling, or, more generally sampling at times determined independently of the actual observation process brings an element of simplicity to the sampling scheme. But an *adapted* scheme, that chooses the sampling instants causally based on past measurements at the sensor (and any other information granted to it by the supervisor), is better. The adapted schemes include the predetermined ones trivially.

A special situation of such a sampling scheme (or rather a control invocation scheme), called Lebesgue sampling, is studied in [13]. A deterministic problem is treated in [14]. The goal of this report is to study the use of adapted sampling schemes in a basic estimation problem. Consider a particular adapted scheme: sampling at some hitting times of the measurement process. There is information transfer through the packets as well as an additional information transmitted when there is no packet transmitted: the fact that the hitting time hasn't arrived yet. For a practical set-up to take advantage of this, packets should be transmitted reliably and with negligible delay (transmission delays are fine ! our packets travel at the speed of light !) and the clocks at the various nodes should be reasonably synchronized. The synchronization condition is required for all sampling schemes to work well in real-time applications. We also require that all nodes work reliably and that they do not die out during operation. This condition can be relaxed if we are presented with a probabilistic model for node failure. Note that in an non-adpated sampling scheme, at least when the sampling instants are deterministic, non-arrival of a packet at a designated time would automatically signal failure. So, an efficient and robust scheme, especially for networks made up of a horde of cheap sensors, is a combination of open-loop and closed-loop policies: that of predetermined and adapted policies. For example, some engineers are introducing TDMA-style packetizing in CANs (goes by the name of *Time-triggered CAN*) to guarantee access to some sensors in the midst of packet-collisions/reliability-issues etc.

An altogether separate aspect of the multi-sensor case arises when the sensor network has a 'star' topology. Here, all nodes are able to listen to the packets their peers send to the decision-maker because they use a common medium. They can coordinate their message transmissions. The trick then is to come up with decentralized schemes which provide each node with a packetizing policy which takes into account the information fed to the decision maker by peers. This model of information exchange with full listening applies to wireless or Ethernet networks operating under something like CSMA-CD. We should add that, at this stage, we will disregard collisions or multiuser detection possibilities and other issues to do with multi access communication. The same model works when the nodes are communicating with the decision-maker in a way inaudible to their peers but have access to a continuous broadcast of estimates and viewpoints of the resource-rich decision maker. The controlled version of this model is the multi-agent co-ordinated control problem with the information-rate constraints built into it. For problems where the control enters the dynamics in an affine way, the optimal controller, we hope, will turn out to be the certainty-equivalence controller [15], [16].

II. Results using Adapted sampling

A. Sampling a single sensor

We will describe in general terms, the problem of optimal adapted sampling that minimizes a filtering distortion. The state process x_t is a controlled partially observed diffusion process. Any unnormalized version (ρ_t) of the conditional density of the state given the observations so far (π_t), will obey the Duncan-Mortensen-Zakai SPDE:

$$d\rho_t(x) = \mathscr{L}^*\left(\rho_t(x)\right) + h(x)\rho_t(x)dy_t \tag{3}$$

where $\mathscr{L}^*$ is the Fokker-Planck (FP) operator given by:

$$\mathscr{L}^*\phi = -\sum_{i=1}^{n}\frac{\partial(f_i\phi)}{\partial x_i} - \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n}\frac{\partial^2(g_ig_j\phi)}{\partial x_i\partial x_j}. \tag{4}$$

for $\phi \in C^2(\mathbb{R}^n)$. We will assume that a finite dimensional sufficient statistic (θ_t) exists for π_t so that

$$\begin{aligned} d\theta_t &= \Phi(\theta_t)dt + \Psi(\theta_t)dy_t & (5)\\ \pi_t(x) &= \kappa(t,x,\theta_t) & (6)\end{aligned}$$

with $\theta_t \in \mathbb{R}^k$. We will further assume that the sensor is able to compute with high accuracy, a numerical approximation of θ_t resulting in a high accuracy computation of π_t.

The causal sampling problem with a fixed number of samples is to pick an increasing sequence $\mathscr{T}_N(\{y_s|0 \leq s \leq$

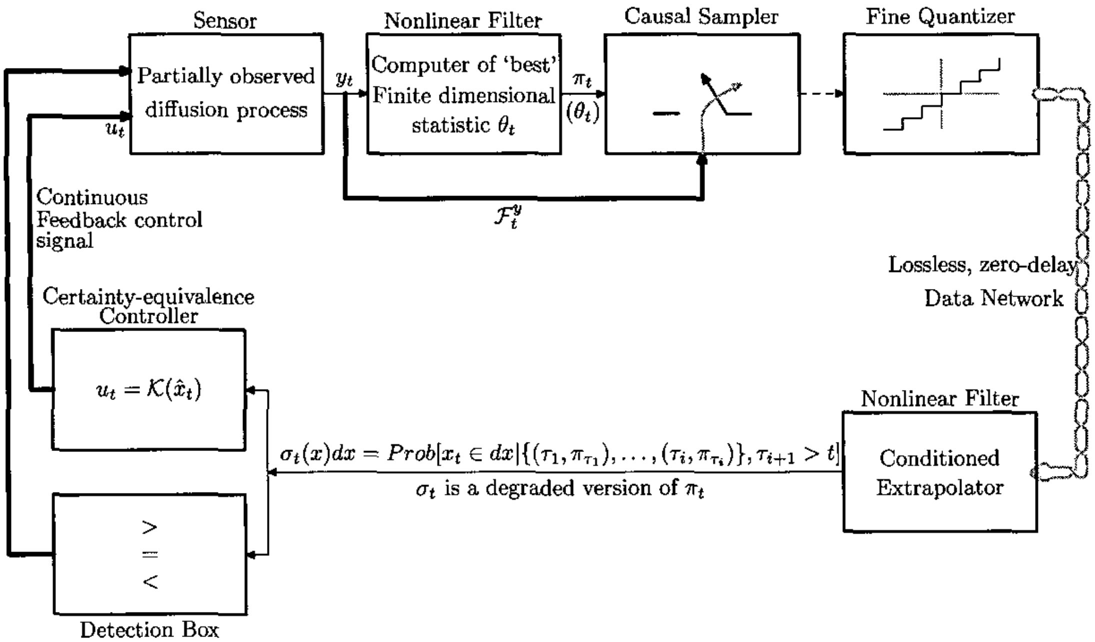

Fig. 1. Schematic of a general Networked control and monitoring system with a single sensor.

$T\}) : \{y_s|0 \leq s \leq T\} \to [0,T]^N$ of N stopping times.

$$\mathscr{T}_N(\{y_s|0 \leq s \leq T\}) = \{\tau_1, \ldots \tau_N\}, \quad (7)$$

$$0 \leq \tau_1 < \tau_2 < \ldots < \tau_{N-1} < \tau_N \leq T, \quad (8)$$

$$\mathbf{1}_{\{\tau_i > t\}} \in \mathscr{F}_t^y \quad \forall\ i \in \{1, 2, \ldots, N\}. \quad (9)$$

At these stopping times, the supervisor receives instantaneously, the current values of θ_t (equivalently, the current values of π_t). Notice that we have basically neglected the noise introduced through quantization of θ_t and through the actual numerical computation of $\{\theta_t\}$ itself.

At all times, the supervisor computes the conditional density of the state (σ_t) given the packet record thus far

$$\{(\tau_1, \pi_{\tau_1}), \ldots, (\tau_{i(t)}, \pi_{l(t)})\},$$

where, $l(t)$ is the last sampling time and $i(t)$ the corresponding packet index. Note that σ_t could be discontinuous at sampling times. If $i(t) < N$,

$$\sigma_t(x)dx = Prob\Big[x_t \in dx \Big| \{(\tau_1, \pi_{\tau_1}), \ldots, (\tau_{i(t)}, \pi_{l(t)})\}, \tau_{i(t)+1} > t\Big]. \quad (10)$$

If $i(t) = N$,

$$\sigma_t(x)dx = Prob\Big[x_t \in dx \Big| \{(\tau_1, \pi_{\tau_1}), \ldots, (\tau_N, \pi_{\tau_N})\}\Big]. \quad (11)$$

Right at the sampling instants, the conditional densities at the supervisor is the same as that at the sensor.

$$\sigma_{l(t)} = \pi_{l(t)}.$$

π_t is a density-valued Markov process. σ_t is the best extrapolation of π_t available at the supervisor.

$$Prob[\pi_t \in d\pi|\{\pi_s|0 \leq s \leq l(t)\}, t < \tau_{i(t)+1}] = Prob[\pi_t \in d\pi|\pi_{l(t)}, t < \tau_{i(t)+1}] \quad (12)$$

This justifies our decision to packetize π_t (or actually, its finite dimensional statistic θ_t). When there is no known finite dimensional sufficient statistic for π_t, it is not clear whether it is optimal to packetize y_t or $E[x_t|\mathscr{F}_t^y]$ or some finite dimensional approximation of a sufficient statistic.

Filtering distortion: Let $\Delta(\cdot,\cdot)$ be a distance operator (positive and semi-definite binary function) in the space of densities. Examples of a suitable distance function would be the Kullback Liebler divergence (we will first have to show that $\sigma_t \ll \pi_t$), the L_1 distance, and the square of the Euclidean distance between the means of σ_t and π_t. Corresponding to a chosen distance function, we can set-up a filtering distortion measure at the supervisor end:

$$E\left[\int_0^T \Delta(\sigma_s, \pi_s)ds\right]. \quad (13)$$

The communication cost is the total number of packets sent: N. The *Optimal sampling problem for Filtering with a fixed sample count* is to choose a sequence $\mathscr{T}_N^*(\{y_s|0 \leq s \leq T\})$ of stopping times that minimizes the filtering distortion and to provide a recipe for computing σ_t.

$$\mathscr{T}_N^*(\{y_s|0 \leq s \leq T\}) = \underset{\mathscr{T}_N(\{y_s|0 \leq s \leq T\})}{\arg\min}\ E\left[\int_0^T \Delta(\sigma_s, \pi_s)ds\right]. \quad (14)$$

The discussion above can be summarized as follows: A causal sampling policy is a multiple stopping policy. Given such a policy $\mathcal{T}_N(\{y_s|0 \le s \le T\})$, the optimal filter at the supervisor is derived from it as the conditional density given by (10,11). We seek the optimal $\mathcal{T}_N^*(\{y_s|0 \le s \le T\})$ as one that minimizes (13). It would save a lot of computational effort if for this optimal sampling strategy, the conditional density can be computed in the fashion of (6) or least as a numerical approximation to something like (3).

This optimization problem can be also posed as a joint optimization over causal sampling policies and causal estimators. We will have occasion to do that in the special case of sampling an ideal sensor once.

Variable number of samples: We can easily extend the solution of the fixed packet count problem to a slightly better performing *Variable packet count* problem. This can be done when the sensor has a lot of computing power at its disposal.

A solution to this joint multiple stopping and filtering problem seems difficult because of the complicated relationship between a stopping policy and the corresponding filter at the supervisor. However, the problem formulation itself is a step forward because it can be solved in special cases and because a natural approximation (which is used in [13] for an infinite time interval problem) still outperforms the periodic sampling strategy. It is a proper generalization of the sampling problem for linear systems studied by Kushner [17].

In what follows, we will (almost) solve this problem for a very special case in which there is a decoupling between the optimal stopping policy and the matching least squares estimate at the supervisor.

B. Scheduling a single packet from an ideal sensor

We describe here the optimal schedule of a single sample on $[0,T]$ for the special case of a perfectly observed scalar state process with odd drift and either even or odd diffusion coefficient functions and an even initial probability density function.

$$dx_t = f(x_t)dt + g(x_t)dB_t \tag{15}$$

$$dy_t = x_t dt \tag{16}$$

With, $x_0 = 0dx$, $y_0 = 0$, $x_t \in \mathbb{R}$, f being odd, and g either odd or even.

The sampling problem is to choose a single $\mathcal{F}_t^x$-stopping time τ on $[0,T]$.

$$\mathcal{T}_1(\{x_s|0 \le s \le T\}) = \{\tau\}, \tag{17}$$

$$0 \le \tau \le T, \tag{18}$$

$$\mathbf{1}_{\{\tau > t\}} \in \mathcal{F}_t^x. \tag{19}$$

Since x_t is fully observed at the sensor, the relevant distortion at the supervisor is now:

$$J = E\left[\int_0^T (\hat{x}_s - x_s)^2 ds\right] \tag{20}$$

where $\hat{x}_t$ is the conditional mean of the state computed by the supervisor based on the initial density, the knowledge of the sampling strategy and either the received single sample or the fact that the sampling has not happened yet.

$$\hat{x}_t = \begin{cases} E[x_t|x_0 = 0, \tau > t] & \text{if} \quad \tau > t \\ E[x_t|x_\tau] & \text{if} \quad \tau \le t \end{cases} \tag{21}$$

On $[0,\tau)$, $\hat{x}_t$ is determined entirely by t. On $[\tau,T]$, $\hat{x}_t$ is determined by the sample received : x_τ. The filtering distortion splits into two parts.

$$E\left[\int_0^\tau (\hat{x}_s - x_s)^2 ds\right] + E\left[\int_\tau^T (\hat{x}_s - x_s)^2 ds\right]. \tag{22}$$

The second part is entirely determined by x_τ and $T - \tau$. On $[\tau,T]$, the variance $E[(\hat{x}_t - x_t)^2] = P_t$ obeys the ODE:

$$\frac{dP_t}{dt} = E\left[2x_t f(x_t) + g^2(x_t)|x_\tau\right] dt \tag{23}$$

with zero as the initial condition : $P_\tau = 0$. Let $C(\tau,t,x_\tau)$ be the solution to this ODE on $[\tau,T]$. Then, the supervisor's distortion becomes:

$$\begin{aligned} &E\left[\int_0^\tau (\hat{x}_s - x_s)^2 ds\right] + E\left[E\left[\int_\tau^T (\hat{x}_s - x_s)^2 ds \Big| \tau, x_\tau\right]\right] \\ &\quad = E\left[\int_0^\tau (\hat{x}_s - x_s)^2 ds\right] + E\left[\int_\tau^T C(\tau,s,x_\tau) ds\right] \end{aligned}$$

Now, let the cost to go from τ be

$$\int_\tau^T C(\tau,s,x_\tau) ds = \mathscr{C}(\tau,T,x_\tau). \tag{24}$$

Then, the overall optimization problem is to choose a stopping policy $\mathcal{T}_1(\{x_s|0 \le s \le T\})$ such that the cost

$$J = E\left[\int_0^\tau (\hat{x}_s - x_s)^2 ds + \mathscr{C}(\tau,T,x_\tau)\right] \tag{25}$$

is minimized. For an optimal sampling strategy, if we can somehow know the dependence of $\hat{x}_t$ on t for $t \in [0,\tau)$, we can use the *Snell envelope* (S_t) (see [19] Appendix D) to determine the optimal stopping rule.

$$\begin{aligned} S_t &= \operatorname*{essup}_{\tau \ge t} \; E\left[\int_0^\tau (\hat{x}_s - x_s)^2 ds + \mathscr{C}(\tau,T,x_\tau) \Big| \mathcal{F}_t^x\right], \\ &= \int_0^t (\hat{x}_s - x_s)^2 ds \\ &\quad + \operatorname*{essup}_{\tau \ge t} \; E\left[\int_t^\tau (\hat{x}_s - x_s)^2 ds + \mathscr{C}(\tau,T,x_\tau) \Big| x_t\right]. \end{aligned}$$

Then, the smallest time τ^* when the cost of stopping at that time hits the Snell envelope is an optimal stopping time (see [19] Appendix D).

$$\int_0^{\tau^*} (\hat{x}_s - x_s)^2 ds + \mathscr{C}(\tau^*,T,x_{\tau^*}) = S_{\tau^*}. \tag{26}$$

Or equivalently,

$$\mathscr{C}(\tau^*,T,x_{\tau^*}) = \operatorname*{essup}_{\tau \ge \tau^*} \; E\left[\int_{\tau^*}^\tau (\hat{x}_s - x_s)^2 ds + \mathscr{C}(\tau,T,x_\tau) \Big| x_t\right]. \tag{27}$$

Since the Snell envelope depends only on the current value of the state and the current time, we get a simple threshold solution for our problem. We can compute the condition to be satisfied by x_t, t for stopping at t by relating this problem to a variational inequality that gives us continuation and stopping regions. In any case, for numerical computation of the solution, we will have to take that route [20]. Now, we will use some properties of the state process that result from our earlier assumptions.

The unobserved x_t is a process with an even density function at all times if the initial density function is even. Basically, the FP operator (4) is linear and so, if we split ρ_t into its even and odd parts

$$\rho_t^+(x) = \frac{\rho_t(x)+\rho_t(-x)}{2},$$
$$\rho_t^-(x) = \frac{\rho_t(x)-\rho_t(-x)}{2},$$

the separate parts obey the FP equation which, with our assumptions, preserves their even and odd properties respectively. Since the initial density function is even, ρ_t is even at all times. $\mathscr{C}(\tau,x,T)$ is also an even function of x.

The joint optimization problem of filtering and sampling has been cast so far as a stopping problem with the optimal filter $(\hat{x}_t)$ determined by the stopping rule. Now, we will look at this optimization (25) as one over different estimators ξ_t

$$\xi_t : [0,T] \to \mathbb{R}$$

for each of which, a stopping rule is devised to minimize the filtering distortion for that estimator. Remember that the estimate before the stopping time is based entirely on π_0, t. There is no ambiguity about the estimator after the packet has arrived. The cost of (25) can be re-interpreted as:

$$J_{\text{TOTAL}}(\xi_t, \mathscr{T}_1) = E\left[\int_0^\tau (\xi_s - x_s)^2 + \mathscr{C}(\tau,T,x_\tau)\right], \quad (28)$$

Given an estimator ξ_t, let $\mathscr{T}_1^*(\{\xi_t\})$ be an optimal stopping rule that minimizes J_{TOTAL} i.e.

$$\hat{J}_{\text{TOTAL}}(\xi_t) = J_{\text{TOTAL}}\left(\xi_t, \mathscr{T}_1^*(\{\xi_t\})\right) = \min_{\mathscr{T}_1} J_{\text{TOTAL}}(\xi_t, \mathscr{T}_1). \quad (29)$$

Let ξ_t^* be an estimator that minimizes $\hat{J}_{\text{TOTAL}}$, i.e.

$$\hat{J}_{\text{TOTAL}}(\xi_t^*) = \min_{\xi_t} \hat{J}_{\text{TOTAL}}(\xi_t). \quad (30)$$

this means that the pair

$$\left(\xi_t^*, \mathscr{T}_1^*(\{\xi_t\})\right)$$

minimizes (in sequence) the nested optimization problem:

$$\min_{\xi_t}\left(\min_{\mathscr{T}_1}\left\{E\left[\int_0^\tau (\xi_s - x_s)^2 ds + \mathscr{C}(\tau,T,x_\tau)\right]\right\}\right) \quad (31)$$

It turns out that combining the estimator $-\xi_t^*$ with the best stopping rule for ξ_t^* does not increase the cost ! This is because the process $-x_t$ has the same statistics as x_t. If $\hat{J}_{\text{TOTAL}}$ has a unique minimizer, then, $\xi_t^* = -\xi_t^*$ a.s. This means that

$$\xi_t^* \equiv 0.$$

This is indeed the conditional mean for the corresponding optimal stopping problem because its Snell's envelope S_t depends only on $|x_t|$ and t. In essence, there is no *Timing channel* between the optimal filter and the optimal stopping policy.

C. A sub-optimal strategy better than any predetermined

We will now describe a rather complicated sampling method adapted to the fully observed state x_t for the general multiple sampling problem. For simplicity, we will assume x_t to be scalar. The resulting sampling policy has some thresholds that are parameters which can be optimized to perform better than the best predetermined sampling strategy. Given any predetermined strategy, we combine it with a distortion threshold strategy that results in a hybrid, variable packet count sampler that can be optimized for performance in several ways. The hybrid sampler waits for predetermined sampling times unless the supervisor's estimate corresponding to the predetermined sampler deviates from the actual state process by more than a threshold. Whatever triggers the sampling, the same wait and watch game is carried on again for the next sample and so on. Note that the supervisor can actually compute a better estimate based on the fact that, until the sample arrives, the estimate for predetermined sampling alone has not deviated too much away from the true state. We are unable to 'close the loop' and use a single estimate corresponding to the hybrid stopping policy. But, we expect to better the predetermined one because, when all the deviation thresholds are very high, the hybrid is simply the same as the predetermined policy. We will fix these ideas in what follows.

Let PD-S be a good predetermined sampler, which, given $[T_1,T_2]$ and the initial state x_{T_1}, prescribes the number of sampling points $N(T_1,T_2,x_{T_1})$ and the actual sampling points

$$PD(T_1,T_2,x_{T_1}) = \left(pd_1(T_1,T_2,x_{T_1}),\ldots,pd_{N(T_1,T_2,x_{T_1})}(T_1,T_2,x_{T_1})\right) \quad (32)$$

that are either deterministic or if random, independent of the state process x_t. A non-imaginative PD-S could just prescribe

$$\left(T_1+\delta,\ T_1+2\delta,\ldots,T_1+\left\lfloor\frac{T_2-T_1}{\delta}\right\rfloor\delta.\right)$$

Let Th-S be the distortion threshold sampler dependent on PD-S. At any time T_1, given a schedule $PD(T_1,T_2,x_{T_1})$ from PD-S, the threshold sampler prescribes a threshold $\eta(T_1,T_2,x_{T_1})$ (the dependence on $PD(T_1,T_2,x_{T_1})$ is suppressed) for the distortion of a simple estimate $\tilde{x}_t$ that can be computed at the supervisor from the true state x_t:

$$J_{FP}(T_1,t) = \int_{T_1}^t (\tilde{x}_s - x_s)^2 ds$$

where, $\tilde{x}_t$ is the estimate on $[T_1,t]$ that would be optimal at the supervisor end until the first packet is received under the purely predetermined sampling strategy PD-S. It is the mean of the density evolving under the FP operator (4).

$$\tilde{x}_t = \begin{cases} E\left[x_t|x_{T_1}\right] & \text{if } T_1 \text{ is a previous sampling time,} \\ E\left[x_t|\pi_0\right] & \text{if } T_1 = 0. \end{cases} \quad (33)$$

What Th-S does is to offer an adapted modification of PD-S. It recommends that the first sample after T_1 be made at the earliest τ_η for which

$$J_{FP}(T_1,\tau_\eta) = \int_{T_1}^{\tau_\eta} (\tilde{x}_s - x_s)^2 ds = \eta(T_1,T_2,x_{T_1}).$$

Th-S could for simplicity, repeat the same η for all initial conditions and time-interval. Or it could do a little computation and produce an η that is inversely dependent on $T_2 - T_1$ and on the largest Lyapunov exponents of the variance ODE (23) and directly proportional to $N(T_1,T_2,x_{T_1})$.

At time zero or at any sampling time, we seek $PD(\tau_i,T_2,x_{\tau_i})$ and $\eta(\tau_i,T_2,x_{\tau_i})$ from the two samplers. The next sampling instant τ_{i+1} is the random time which is the earlier of the two:

$$\tau_{i+1} = \tau_\eta \bigwedge pd_1(\tau_i,T_2,x_{\tau_i}).$$

If Th-S always sets $\eta = \infty$, it will be agreeing with PD-S on when to sample and $PD(T_1,T_2,x_{T_1})$ is the sampling strategy implemented over $[T_1,T_2]$.

Having resolved this scheduling policy that performs (by a judicious choice of the ηs) as good as and possibly better than any predesigned scheduler[2], we will describe a least-squares estimate $\hat{x}_t$ that is better than $\tilde{x}_t$, in fact the best one at the supervisor.

Given the hybrid sampling strategy, the least-squares estimate at the supervisor at any time t between two samples ($\tau_i \le t < \tau_{i+1}$) is:

$$\begin{aligned} \hat{x}_t &= E\left[x_t \middle| x_{\tau_i}, pd_1(\tau_i,T,x_{T_1}) > t, \tau_\eta > t\right] \\ &= E\left[x_t \middle| x_{\tau_i}, \tau_\eta > t\right] \\ &= E\left[x_t \middle| x_{\tau_i}, \int_{T_1}^{\tau_\eta} (\tilde{x}_s - x_s)^2 ds < \eta(\tau_i,T,x_{T_1})\right] \end{aligned}$$

To compute this, look at the 2-dimensional diffusion process:

$$\begin{pmatrix} x_t \\ z_t \end{pmatrix} = \begin{pmatrix} x_t \\ \int_{\tau_i}^t (\tilde{x}_s - x_s)^2 ds \end{pmatrix}.$$

This obeys

$$\begin{pmatrix} dx_t \\ dz_t \end{pmatrix} = \begin{pmatrix} f(x_t) \\ 2(\tilde{x}_t - x_t)\{E[f(x_t)|x_{\tau_i}] - f(x_t)\} + 1 \end{pmatrix} dt + \begin{pmatrix} g(x_t) \\ -2(\tilde{x}_t - x_t) g(x_t) \end{pmatrix} dW_t. \quad (34)$$

Then $\hat{x}_t$ is the first component of the n-long subvector of the conditional mean of this vector process given that it started from the point $(x_{\tau_i},0)$ at time τ_i and is until now, within the open set

$$A = \left\{(x,z) \middle| (x,z) \in \mathbb{R}^2, -\eta(\tau_i,T,x_{T_1}) < z < \eta(\tau_i,T,x_{T_1})\right\}.$$

This can be computed using the Feynman-Kac formula [18].

It seems very much possible that in the partially observed case, with a finite dimensional sufficient statistic θ_t for the statistics of x_t given $\mathscr{F}_t^y$, θ_t which fully determines the information state plays the role of x_t for a fully observed sensor.

[2]We do need a good family of predesigned strategies that work on $(\tau, T]$ to be fed to Th-S at sampling times.

References

[1] A. Mason, N. Yazdi, A. V. Chavan, K. Najafi, and K. D. Wise: A generic multielement microsystem for portable wireless applications, *Proceedings of the IEEE* - Special issue on MEMS , v: 86, no: 3, (August 1998).

[2] A. Hać: *Wireless Sensor Network Designs*, John Wiley (2003).

[3] J. Liu, N. Chu, J. Liu, J. Reich, J. Zhao: State-Centric Programming for Sensor-Actuator Network Systems, *IEEE Pervasive Computing magazine*, vol:2, no:4 pp:50-62 (2003).

[4] R. Holman, J. Stanley, and T. Özkan-Haller: Applying Video Sensor Networks to Nearshore Environment Monitoring, *IEEE Pervasive Computing magazine*, vol:2, no:4 pp:14-21 (2003).

[5] W. F. Powers and P. S. Nicastri: Automotive vehicle control challenges in the 21st century, *Control engineering practice* v:8 pp:605-618 (2000).

[6] Bosch company web-site on Literature of Controller area networks, http://www.can.bosch.com/content/Literature.html .

[7] G. Leen and D. Heffernan: Expanding Automotive Electronic Systems. *IEEE Computer magazine* v:35 n:1 pp:88-93 (2002).

[8] W. S. Wong and R. W. Brockett: Systems with finite communication bandwidth constraints I. State estimation problems, *IEEE Transactions on Automatic Control* v:42 n:9 pp:1294-1299 (1997).

[9] W. S. Wong and R. W. Brockett: Systems with finite communication bandwidth constraints II. Stabilization with limited information feedback, *IEEE Transactions on Automatic Control* v:44 n:5 pp:1049-1053 (1999).

[10] R. W. Brockett and D. Liberzon, Quantized feedback stabilization of linear systems, *IEEE Transactions on Automatic Control* v:45 n:7 pp:1279-1289 (2000).

[11] V. S. Borkar, S. K. Mitter, and S. Tatikonda: Optimal sequential vector quantization of Markov sources, *SIAM Journal on Control and Optimization* v:40 n:1 (2001).

[12] H. V. Poor: Fine quantization in signal detection and estimation, *IEEE Transactions on Information Theory* v:34 no:5 part1 pp:960-972 (1988).

[13] K. J. Åström and B. M. Bernhardsson: Comparison of Riemann and Lebesgue sampling for first order stochastic systems, *Proceedings of the IEEE Conf. on Decision and Control at Las Vegas, NV* (December 2002).

[14] D. Hristu-Varsekalis and P. R. Kumar, Interrupt-based feedback control over a shared communication medium, IEEE Conference on Decision and Control (2002).

[15] H. S. Witsenhausen: Separation of estimation and control for discrete time systems, *Proceedings of the IEEE* v:59 n:11 pp:1557-1566 (1971).

[16] J-L. Menaldi: The separation principle for impulse control problems, *Proceedings of the American Mathematical Society*, v:82, no:3, pp:439-445 (1981).

[17] H. J. Kushner: On the optimum timing of observations for linear control systems with unknown initial state, *IEEE Transactions on Automatic Control* v:9 pp:144-150 (1964).

[18] V. Krishnamurthy and F. LeGland, State and parameter estimation from boundary-crossings, *Proceedings of the IEEE Conference on Decision and Control at Tampa, FL* (December 1998).

[19] I. Karatzas and S. E. Shreve: *Methods of Mathematical Finance*, Springer-Verlag (1998).

[20] R. Glowinski: *Numerical methods for nonlinear variational problems*, Springer Verlag (1984).