SIAM J. CONTROL OPTIM.
Vol. 0, No. 0, pp. 000–000

ADAPTIVE SAMPLING FOR LINEAR STATE ESTIMATION*

MABEN RABI†, GEORGE V. MOUSTAKIDES‡, AND JOHN S. BARAS§

Abstract. When a sensor has continuous measurements but sends occasional messages over a data network to a supervisor which estimates the state, the available packet rate fixes the achievable quality of state estimation. When such rate limits turn stringent, the sensor's messaging policy should be designed anew. What are good causal messaging policies? What should message packets contain? What is the lowest possible distortion in a causal estimate at the supervisor? Is Delta sampling better than periodic sampling? We answer these questions for a Markov state process under an idealized model of the network and the assumption of perfect state measurements at the sensor. If the state is a scalar, or a vector of low dimension, then we can ignore sample quantization. If in addition we can ignore jitter in the transmission delays over the network, then our search for efficient messaging policies simplifies. First, each message packet should contain the value of the state at that time. Thus a bound on the number of data packets becomes a bound on the number of state samples. Second, the remaining choice in messaging is entirely about the times when samples are taken. For a scalar, linear diffusion process, we study the problem of choosing the causal sampling times that will give the lowest aggregate squared error distortion. We stick to finite horizons and impose a hard upper bound N on the number of allowed samples. We cast the design as a problem of choosing an optimal sequence of stopping times. We reduce this to a nested sequence of problems, each asking for a single optimal stopping time. Under an unproven but natural assumption about the least-square estimate at the supervisor, each of these single stopping problems are of standard form. The optimal stopping times are random times when the estimation error exceeds designed envelopes. For the case where the state is a Brownian motion, we give analytically: the shape of the optimal sampling envelopes, the shape of the envelopes under optimal Delta sampling, and their performances. Surprisingly, we find that Delta sampling performs badly. Hence, when the rate constraint is a hard limit on the number of samples over a finite horizon, we should not use Delta sampling.

Key words. state estimation, optimal sampling, Delta sampling, networked control

AMS subject classifications. AUTHOR MUST PROVIDE

DOI. 10.1137/090757125

1. Introduction.

Networked control systems have some control loops completed over data networks rather than over dedicated analog wires or field buses. In such systems, monitoring and control tasks have to be performed under constraints on the amount of information that can be communicated to the supervisor or control station. These communication constraints limit the rate of packet transmissions from sensor nodes to the supervisor node. Even at these limited rates, the network communications can be less than ideal: the packets can be delayed and sometimes lost. In the

*Received by the editors April 27, 2009; accepted for publication (in revised form) November 28, 2011; published electronically DATE. Some of the work in this paper was reported in the Ph.D. thesis of the first author, *Packet Based Inference and Control*, University of Maryland, 2006, and in *Proceedings of the IEEE Conferences on Decision and Control*, 2004 and 2006. This work was supported by the United States Army Research Office through the Center for Networked Communication and Control Systems at Boston University, grant DAAD 19021 0319, and by CIP-URI grant DAAD 190110494.

http://www.siam.org/journals/sicon/x-x/75712.html

†Department of Signals and Systems, Chalmers University of Technology, 412 96 Gothenburg, Sweden (mabenr@gmail.com). This author received partial financial support from the Swedish Research Council and the European Commission.

‡Department of Electrical and Computer Engineering, University of Patras, 26500 Rio, Greece (moustaki@upatras.gr).

§Institute for Systems Research, Department of Electrical and Computer Engineering, University of Maryland, College Park, MD 20742 (baras@umd.edu).

networked system, all of these communication degradations lower performance, and so these effects must be accounted for during control design. In this paper, we only account for the limit on the packet rates and completely ignore random delays and packet losses.

Sending data packets as per a periodic timetable works well when high data rates are possible. Sending packets aperiodically and at variable times becomes worthwhile only when the packet rate limits get stringent, like in an industrial wireless network. Conceptually, packet rate constraints can be of the following three types: (1) *average rate limit*, a "soft constraint" that calls for an upper limit on the average number of transmissions; (2) *minimum waiting time between transmissions*, under which there is a mandatory minimum waiting time between two successive transmissions from the same node; and (3) *finite transmission budget*, a "hard constraint" that allows only up to a prescribed number of transmissions from the same node over a given time window. In the simplest version of the third type of constraint, we set the constraint's window to be the problem's entire time horizon. In its other variations, we can promote a steadier flow of samples and avoid too many samples being taken in a short time. This we do by cutting the problem's time horizon into many disjoint segments and applying the finite transmission budget constraint on every segment.

Notice that these different types of constraints can be mixed in interesting ways. In this work, we will adopt the simple version of the finite transmission budget, in which the budget window is the same as the problem's time horizon. We study a problem of state estimation which is an important component of distributed control and monitoring systems. Specifically, a scalar linear system is continuously and fully observed at a sensor which generates a limited number of packets. A supervisor receives the causal sequence of packets and, on its basis, maintains a causal estimate. Clearly, the fewer packets allowed, the worse the error in the supervisor's estimate. The design question is, *How should the packets be chosen by the sensor to minimize the estimation distortion?* The answer to this question employs the idea that *packets should be generated only when they contain "sufficiently" new information.* Adaptive sampling schemes, or event-triggered sampling schemes as they are also called, exploit this idea and send samples at times determined by the trajectory of the source signal being sampled. In contrast, deterministic sampling chooses sample times according to an extraneous clock.

But first we will consider possible times when packets should be sent and the allowable payloads they can carry. The times when packets are sent must be causal times which, even if random, are stopping times w.r.t. the sensor's observations process. Likewise, the payloads have to be measurable w.r.t. the filtration generated by the observations process. The above restrictions are merely the demands of causality. When we place some idealized assumptions about the network, a simple and obvious choice of payload emerges.

1.1. Strong Markov property, idealized network, and choice of payload.

For all the problems treated in this paper, we need the two clocks at the sensor and the supervisor to agree and, of course, to report time correctly. We also assume that the state signal x_t is a strong Markov process. This means that for *any stopping time* τ, and any measurable subset A of the range of x, and for any time $t \geq \tau$,

$$\mathbb{P}\left[x_t \in A \,\middle|\, \mathcal{F}^x_\tau\right] = \mathbb{P}\left[x_t \in A \,\middle|\, x_\tau\right].$$

Linear diffusions, of course, have the strong Markov property.

Let the sequence $\{\tau_1, \tau_2, \dots\}$ of positive reals represent the sequence of times when the sensor puts packets on the network. Let the sequence of binary words $\{\pi_1, \pi_2, \dots\}$ denote the corresponding sequence of payloads put out. Let the sequence $\{\sigma_1, \sigma_2, \dots\}$ of nonnegative reals denote the corresponding transmission delays incurred by these packets. We let these delays be random but require that they be independent of the signal process. The packet arrival times at the supervisor, arranged in the order in which they were sent, will be $\{\tau_1 + \sigma_1,\ \tau_2 + \sigma_2,\ \dots\}$. Let the positive integer $l(t)$ denote the number of packets put out by the sensor up to and including the time t. We have

$$l(t) = \sup\{i \,|\tau_i \le t\}.$$

A causal record of the sensor's communication activities is the transmit process defined as the following piecewise constant process:

$$\mathbb{TX}_t = \begin{pmatrix} \tau_{l(t)} \\ \pi_{l(t)} \end{pmatrix}.$$

When a packet arrives, the supervisor can see its time stamp τ_j, its payload π_j, and of course its arrival time $\tau_j + \sigma_j$. We ignore quantization noise in the time stamps, with the result that the supervisor can read both τ_j and $\tau_j + \sigma_j$ with infinite precision. The causal record of what the supervisor receives over the network is described by the random process defined as

$$\mathbb{RX}_t = \sum_j \mathbf{1}_{\left\{\substack{\tau_j + \sigma_j \le t, \text{ and} \\ t < \tau_{j+1} + \sigma_{j+1}}\right\}} \begin{pmatrix} \tau_j \\ \tau_j + \sigma_j \\ \pi_j \end{pmatrix},$$

where we have assumed that no two packets arrive at exactly the same time and that packets are received in exactly the order in which they were sent. If we were to study the general case where packets can arrive out of sequence, then the arguments below will have to be made more delicate, but the conclusion below will still hold.

The supervisor's task is causal estimation. This fact restricts the way in which $\mathbb{RX}_t$ is used by the supervisor. Let the count $r(t)$ denote the number of packets received so far. Then, the data in the hands of the supervisor at time t is the collection

$$r(t), \left\{(\tau_j, \tau_j + \sigma_j, \pi_j)\,|1 \le j \le r(t)\right\}$$

This is to be used to estimate the present and future values of the state.

What should the sensor assign as payloads to maximize information useful for signal extrapolation? Specifically, what should the latest payload $\pi_{r(t)}$ be? If the bit width of payloads is large enough to let us ignore quantization, then the best choice of payload is the sample value at the time of generation, namely, $x_{\tau_{r(t)}}$. Because of the strong Markov property, at times $s \ge t$,

$$\mathbb{P}\Big[x_s \in A \,\Big|\, x_{\tau_{r(t)}}, r(t), \left\{(\tau_j, \tau_j + \sigma_j, \pi_j)\,|1 \le j \le r(t)\right\}\Big] = \mathbb{P}\Big[x_s \in A \,\Big|\, x_{\tau_{r(t)}}\Big],$$

which means that if $\pi_{r(t)}$ carries $x_{\tau_{r(t)}}$ exactly, then the future estimation errors are minimized. Therefore, the ideal choice of payload is the sample value. But what about the practical nonzero quantization noise? Again, the strong Markov property implies that all the bits available should be used to encode the current sample; the encoding scheme depends on the distortion criterion for estimation.

If the packets do not arrive out of turn, the effect of packet delays even when random is not qualitatively different from the ideal case where all transmission delays are zero. Nonzero delays can merely make the estimation performance worse but cannot change the structure of the optimal sampler and estimator. Hence, we will assume all packet transit delays to be zero, and $l(t) = r(t)$ always.

1.2. Ignoring quantization noise in payloads. In most networks [8, 21, 16], the packets are of uniform size and even when of variable size have at least a few header and trailer bytes. 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 and whistles." It costs only negligibly more in terms of network resources, of time, or of energy to send a payload of 5 or 10 bytes instead of 2 bits or 1 byte when the overhead part of the packet is already 5 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. For Markov state processes, this means that all these bytes of payload can be used to specify the latest value of the state. In other words, in essentially all packet-based communication schemes, the right unit of communication cost is the cost of transmitting a single packet. The exact number of bits used to quantize the sample is not important, as long as there are enough to make quantization noise insignificant. There are of course special situations where the quantization rate as well as the sample generation rate matter. An example occurs in the Internet congestion control mechanism called transmission control protocol [15], where a node estimates the congestion state of a link through congestion bits added to regular data packets. In this case, the real payload in packets is irrelevant to the congestion state, and the information on the congestion state is derived from the 1 or 2 bits riding piggyback on the data packets. The developments in this paper do not apply to such problems where quantization is important.

1.3. Infinite Shannon capacity and well-posedness. The continuous time channel from the sensor to the supervisor is idealized and noise-free. Even when a sequence of packets is delivered with delays, the supervisor can recover perfectly the input trajectory $\{\mathbb{TX}\}_0^T$ from the corresponding trajectory of the output $\{\mathbb{RX}\}_0^T$. The supervisor can read each time τ_i and the sample value x_{τ_i} with infinite precision. Since the sensor has an infinite range of choices for each τ_i, the channel has infinite communication capacity in the sense of Shannon.

But this does not render the sampling problem ill-posed. A packet arriving at time τ_i carries the data $l(\tau_i), \tau_i, x_{\tau_i}$. Given (τ_i, x_{τ_i}), the trajectory of x prior to τ_i is of no use for estimating $\{x_s \,|s \geq \tau_i\}$. Therefore, it does not pay to choose τ_i cleverly so as to convey extra news about the past trajectory of x. No such strategy can add to what the supervisor already gets, namely, the pair (τ_i, x_{τ_i}). There is nevertheless scope, and in fact a need for choosing τ_i cleverly so that the supervisor can use the silence *before* τ_i to improve its state estimate *before* τ_i. But for the causal estimation problem the infinite Shannon capacity does not sway the choice of sampling policies.

In summary, our assumptions so far are (1) the state is a strong Markov process, (2) the channel does not delay or lose packets, (3) the time stamps τ_i are available with infinite precision to the supervisor, and (4) the sample value x_{τ_i} is available with infinite precision to the supervisor. Thus we have $\sigma_i = 0 \; \forall \; i$, $\mathbb{RX}_t = \mathbb{TX}_t \; \forall \; t$, and $r(t) = l(t) \; \forall \; t$.

1.4. Relationship to previous works. State estimation problems with communication rate constraints arise in a wide variety of networked monitoring and control setups such as sensor networks, wireless industrial monitoring and control systems, rapid prototyping using a wireless network, and multiagent robotics. A recent overview of research in networked control systems including a variety of specific applications is available from the special issue [3].

Adaptive or event-triggered sampling may also be used to model the functioning of various neural circuits in the nervous systems of animals. After all, the neuron is a threshold-triggered firing device whose operation is closely related to Delta sampling. However, it is not presently clear if the communication rate constraint adopted in this paper occurs in biological neural networks.

Adaptive sampling and adaptive timing of actuation have been used in engineered systems for close to a hundred years. Thermostats use on-off controllers which switch on or off at times when the temperature crosses thresholds (subject to some hysteresis). Delta-Sigma modulation (Delta sampling) is an adaptive sampling strategy used widely in signal processing and communication systems. Nevertheless, theory has not kept up with practice.

Timing of observations via pull sampling and push sampling. The problem of choosing the time instants to sample sensor measurements received early attention in the literature. Kushner [19], in 1964, studied the deterministic, offline choice of measurement times in a discrete-time, finite horizon, linear quadratic Gaussian (LQG) optimal control problem. He showed that the optimal deterministic sampling schedule can be found by solving a nonlinear optimization problem. Skafidas and Nerode [31] allow the online choice of times for sampling sensor measurements, but these times are to be chosen online by the controller rather than by the sensor. Their conclusion is that for linear controlled systems, the optimal choice of measurement times can be made offline. Their offline scheduling problem is the same as Kushner's deterministic one.

A generalization of these problems of deterministic choice of measurement times is the *sensor scheduling problem,* which has been studied for estimation, detection, and control tasks [22, 5, 32]. This problem asks for online schedules for gathering measurements from different available sensors. However, the information pattern for this problem is the same as in the works of Kushner and of Skafidas and Nerode. Under this information pattern, data flow from sensors to their recipients is directed by the recipients. Such sensor sampling is of the "pull" type. An alternative is the "push" type of sampling, where the sensor itself regulates the flow of its data. When only one sensor is available, that sensor has more information than the recipient and hence its decisions on when to communicate its measurements can be better than decisions the supervisor can make. Adaptive sampling is essentially the push kind of sampling.

Lebesgue sampling and its generalizations. The first analytic study of the communication benefits of using event-triggered sampling was presented in the 2002 paper of Åström and Bernhardsson [4]. They treat a minimum variance control problem with the push type of sampling. The control consists of impulses which reset the state to the origin, but there is an upper limit on the average rate at which impulses can be applied. Under such a constraint, the design asks for a schedule of the application times for the impulses. For scalar Gaussian diffusions, they perform explicit calculations to show that the application of impulses triggered by the crossing of fixed, symmetric levels is more efficient than periodic resetting.

This has spurred further research in the area. Our own work [26, 27, 28, 29] generalized the work of Åström and Bernhardsson. Their impulse control problem is equivalent to the problem of sampling for causal estimation. In the setting of discrete time, Imer and Basar [14] study the problem of efficiently using a limited number of discrete-time impulses. For a finite horizon LQG optimal control problem, they use dynamic programming to show that time-varying thresholds are optimal. Henningsson, Johannesson, and Cervin [13] have generalized to delays and transmission constraints imposed by real data networks.

In the setting of discrete time, for infinite horizons, Hajek [11] and Hajek, Mitzel, and Yang [12] have treated essentially the same problem as ours. They were the first to point out that in the sequential decision problem, the two agents have different information patterns. For a general Markov state process, they describe as unknown the jointly optimal choice of sampling policy and estimator. For state processes which are symmetric random walks, they show that the jointly optimal scheme uses adaptive sampling and that the corresponding estimator is the same "centered" estimator one uses for deterministic sampling. We are unable to prove a similar claim about the optimal estimator for our continuous time problem.

The study of optimal adaptive sampling timing leads to *optimal stopping* problems of stochastic control or, equivalently, to impulse control problems. The information pattern of adaptive sampling complicates the picture but methods of solving multiple stopping time problems of standard form which are available in the literature [7] are indeed useful.

The work reported in this paper has been announced previously in [25, 26, 30]. In [27], the single sample case has been dealt with in more detail than here.

1.5. Contributions and outline of the paper. For the finite horizon state estimation problem, we cast the search for efficient sampling rules as sequential optimization problem over a fixed number of causal sampling times. This we do in section 2, where we formulate an optimal multiple stopping problem with the aggregate quadratic distortion over the finite time horizon as its cost function. We restrict the estimate at the supervisor to be that which would be optimal under deterministic sampling. Following Hajek [11] and Hajek, Mitzel, and Yang [12], we conjecture that when the state is a linear diffusion process, this estimate is indeed the least-square estimate corresponding to the optimal sampling strategy.

In section 3, we take the simplified optimal multiple stopping problem and solve it explicitly when the state is the (controlled) Brownian motion process. The optimal sampling policies are first hitting times of time-varying envelopes by the estimation error signal. Our analytical solution shows that for each of the sampling times, the triggering envelopes are symmetric around zero and diminish monotonically in a reverse-parabolic fashion as time nears the end of the horizon. We also describe analytically the performance of the class of modified Delta sampling rules in which the threshold δ varies with the number of remaining samples. We point out a simple and recursive procedure for choosing the most efficient of these Delta sampling policies.

For the Ornstein–Uhlenbeck process, in section 4, we derive dynamic programming equations for the optimal sampling policy. We compute the solution to these equations numerically. We are not able to say whether an explicit analytic solution like for the Brownian motion process is possible. We can say that the optimal sampling times are first hitting times of time-varying envelopes by the estimation error signal. These envelopes are symmetric around zero and diminish monotonically as time nears the end of the horizon. Also derived are the equations governing the performance of modified Delta sampling rules and the most efficient among them is found through a numerical search. Finally, in section 5, we conclude and speculate on extensions to this work for other estimation, control and detection problems.

2. Minimum mean-square-error estimation and optimal sampling. Under a deterministic time-table for the sampling instants, the minimum mean square error (MMSE) reconstruction for linear systems is well known and is straightforward to describe—it is the Kalman filter with intermittent but perfect observations. The error variance of the MMSE estimate obeys the standard Riccati equation. In Delta sampling [23, 9, 10], also called Delta modulation, a new sample is generated when the source signal moves away from the previously generated sample value by a distance δ. By this rule, between successive sample times, the source signal lies within a ball of radius δ centered at the earlier sample. Such news of the state signal during an intersample interval is possible in adaptive sampling but never in deterministic sampling. Because of this, the signal reconstruction under adaptive sampling differs from that under deterministic sampling, and we will see below what the difference is.

We will also set up an optimization problem where we seek an adaptive sampling policy minimizing the distortion of the MMSE estimator subject to a limit on the number of samples. Consider a state process x_t which is a (possibly controlled) scalar linear diffusion. It evolves according to the SDE

$$dx_t = ax_t dt + bdB_t + u_t dt, \; x_0 = x, \tag{2.1}$$

where B_t is a standard Brownian motion process. The control process u_t is right continuous with left limits (RCLL) and of course measurable with respect to the x-process. In fact, the feedback form of u_t is restricted to depend on the sampled information only; we will describe this subsequently. We assume that the drift coefficient a, the noise coefficient $b \neq 0$, and the initial value x are known. Now, we will dwell upon sampling and the estimation process.

The state is sampled at instants $\{\tau_i\}_{i\geq 0}$ which are stopping times w.r.t. the x-process. Recall that the process $\mathbb{RX}_t$ represents the data contained in the packets received at the estimator:

$$\mathbb{RX}_t = \sum_i \mathbf{1}_{\left\{\substack{\tau_i \leq t, \text{ and} \\ t < \tau_{i+1}}\right\}} \begin{pmatrix} \tau_i \\ x_{\tau_i} \end{pmatrix}. \tag{2.2}$$

Notice that the binary process $\mathbf{1}_{\{\tau_i \leq t\}}$ is measurable w.r.t. $\mathcal{F}_t^{\mathbb{RX}}$. The MMSE estimate $\hat{x}_t$ is based on knowledge of the multiple sampling policy and all the information contained in the output of the sensor and so it can be written as

$$\hat{x}_t = \mathbb{E}\left[x_t \middle| \mathbb{RX}_t\right].$$

The control signal u_t is measurable w.r.t. $\mathcal{F}_t^{\mathbb{RX}}$. Typically it is restricted to be of the certainty-equivalence type as depicted in Figure 2.1(a). In that case u_t is, in addition, measurable w.r.t. $\mathcal{F}_t^{\hat{x}}$. The exact form of the feedback control is not important for our work, but it is essential that both the supervisor and the sensor know the feedback control policy (and so can compute the control waveform u_0^t). With this knowledge, the control waveform is a known additive component to the state evolution and hence can be subtracted out. Therefore, there is no loss of generality in considering only the uncontrolled plant.

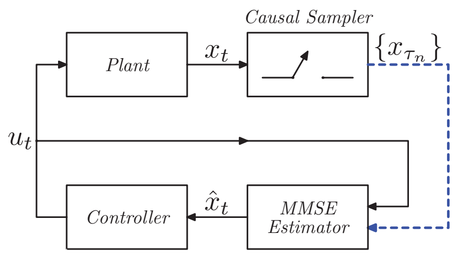

(a) Adaptive sampling for real-time estimation

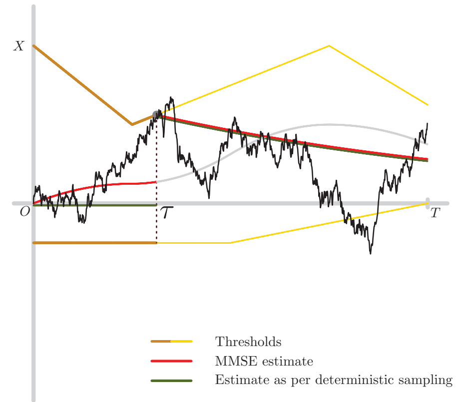

(b) Sampling policy shapes the Estimator

FIG. 2.1. (a) *Setup for the MMSE estimation based on samples arriving at a limited rate.* (b) *Difference between estimators for adaptive sampling and deterministic sampling.*

2.1. MMSE estimation under deterministic sampling. Consider now a deterministic scheme for choosing the sampling times. Let the sequence of nonnegative and increasing sampling times be

$$\mathcal{D} = \{d_0, d_1, \dots\}, \quad d_0 = 0,$$

where the times d_i are all statistically independent of all data about the state received after time zero. They can depend on the initial value of the state x_0.

We will now describe the MMSE estimate and its variance. Consider a time t in the semiopen interval $[d_i, d_{i+1})$. We have

$$\begin{aligned}\hat{x}_t &= \mathbb{E}\left[x_t \middle| \mathbb{R}\mathbb{X}_t\right] \\ &= \mathbb{E}\left[x_t \middle| d_i \le t < d_{i+1}, \left\{\left(d_j, x_{d_j}\right) \middle| 0 \le j \le i\right\}\right] \\ &= \mathbb{E}\left[x_t \middle| d_i \le t < d_{i+1}, d_i, x_{d_i}\right] \\ &= \mathbb{E}\left[x_t \middle| d_i, x_{d_i}\right],\end{aligned}$$

where we have used the Markov property of the state process and the mutual independence, conditioned on x_0, of the state and the sequence $\mathcal{D}$. Furthermore,

$$\begin{aligned}\hat{x}_t &= \mathbb{E}\left[e^{a(t-d_i)} x_{d_i} + \int_{d_i}^{t} e^{a(t-s)} b\, dB_s + \int_{d_i}^{t} e^{a(t-s)} u_s ds \middle| d_i, x_{d_i}\right] \\ &= e^{a(t-d_i)} x_{d_i} + \int_{d_i}^{t} e^{a(t-s)} u_s ds.\end{aligned}$$

Thus, under deterministic sampling, the MMSE estimate obeys a linear ODE with jumps at the sampling times.

$$\frac{d\hat{x}_t}{dt} = a\hat{x}_t + u_t \text{ for } t \notin \mathcal{D} \quad \text{and} \quad \hat{x}_t = x_t \text{ if } t \in \mathcal{D}.$$

The variance $p_t = \mathbb{E}[(x_t - \hat{x}_t)^2]$ is given by the well-known Riccati equation

$$\frac{dp_t}{dt} = 2ap_t + b^2 \text{ for } t \notin \mathcal{D} \quad \text{and} \quad p_t = 0 \text{ if } t \in \mathcal{D}.$$

The above description for the MMSE estimate and its variance is valid even when the sampling times are random provided that these times are independent of the state process except possibly the initial condition. There, too, the evolution equations for the MMSE estimate and its error statistics remain independent of the policy for choosing the sampling times; the solutions to these equations merely get reset with jumps at these random times. On the other hand, adaptive sampling modifies the evolution of the MMSE estimator, as we will see next.

2.2. The MMSE estimate under adaptive sampling. Between sample times, an estimate of the state is an estimate up to a stopping time, and this is the crucial difference from deterministic sampling. Denote this estimate by $\tilde{x}_t$. At time t within the sampling interval $[\tau_i, \tau_{i+1})$, the MMSE estimate is given by

$$\begin{aligned}
\tilde{x}_t &= \mathbb{E}\left[x_t \middle| \mathbb{R}\mathbb{X}_t\right] \\
&= \mathbb{E}\left[x_t \middle| \tau_i \le t < \tau_{i+1}, \left\{\left(\tau_j, x_{\tau_j}\right) \middle| 0 \le j \le i\right\}\right] \\
&= \mathbb{E}\left[x_t \middle| \tau_i \le t < \tau_{i+1},\ \tau_i,\ x_{\tau_i}\right] \quad \text{(strong Markov property)} \\
&= x_{\tau_i} + \mathbb{E}\left[x_t - x_{\tau_i} \middle| t - \tau_i < \tau_{i+1} - \tau_i,\ \tau_i,\ x_{\tau_i}\right].
\end{aligned}$$

Similarly, its variance p_t can be written as

$$p_t = \mathbb{E}\left[(x_t - \hat{x}_t)^2 \middle| \tau_i \le t < \tau_{i+1},\ \tau_i,\ x_{\tau_i}\right].$$

Between samples, the MMSE estimate is an estimate up to a stopping time because the difference of two stopping times is also a stopping time. In general, it is different from the MMSE estimate under deterministic sampling (see Appendix A). This simply means that in addition to the information contained in previous sample times and samples, there are extra partial observations about the state. This information is the fact that the next stopping time τ_{i+1} has not arrived. Thus, in adaptive schemes, the evolution of the MMSE estimator is dependent on the sampling policy. This opens the possibility of a *timing channel* [2] for the MMSE estimator.

Figure 2.1(b) describes a particular (suboptimal) scheme for picking a single sample. There are two time-varying thresholds for the state signal, an upper one and a lower one. The initial state is zero and within the two thresholds. The earliest time within $[0, T]$, when the state exits the zone between the thresholds, is the sample time. The evolution of MMSE estimator is dictated by the shape of the thresholds, thus utilizing information available via the timing channel.

2.3. An optimal stopping problem. We formalize a problem of sampling for optimal estimation over a finite horizon. We seek to minimize the distortion between the state and the estimate $\hat{x}$. We conjecture that under optimal sampling,

$$\tilde{x}_t = \hat{x}_t \quad \text{almost surely.} \tag{2.3}$$

If increments of the state process are not required to have symmetric PDFs, clearly the conjecture is false (see Appendix A).

On the interval $[0, T]$, for the state process x_t obeying (2.1), with the initial condition x_0, we seek an increasing and causal sequence of at most N sampling times $\{\tau_1, \ldots, \tau_N\}$ to minimize the aggregate squared error distortion

$$J(T, N) = \mathbb{E}\left[\int_0^T (x_s - \hat{x}_s)^2 ds\right]. \tag{2.4}$$

Notice that the distortion measure does not depend on the initial value of the state because it operates only on the error signal $(x_t - \hat{x}_t)$, which is zero at time zero no matter what x_0 is. Notice also that the communication constraint is captured by an upper limit on the number of samples. In this formulation, we do not get any reward for using fewer samples than the budgeted limit.

The optimal sampling times can be chosen one at a time using a nested sequence of solutions to optimal single stopping time problems. This is because for a sampling time τ_{i+1} which succeeds the time τ_i, using knowledge of how to choose the sequence $\{\tau_{i+1}, \ldots, \tau_N\}$ optimally, we can obtain an optimal choice for τ_i by solving over $[0, T]$ the optimal single stopping time problem

$$\inf_{\tau \geq 0} \mathbb{E}\left[\int_0^{\tau_i} (x_s - \hat{x}_s)^2 ds + J^*(T - \tau_i, N - i)\right],$$

where $J^*(T - \tau_i, N - i)$ is the minimum distortion obtained by choosing $N - i$ sample times $\{\tau_{i+1}, \ldots, \tau_N\}$ over the interval $[\tau_i, T]$. The best choice for the terminal sampling time τ_N is based on solving a single stopping problem. Hence we can inductively find the best policies for all earlier sampling times. Without loss of generality, we can examine the optimal choice of the first sampling time τ_1 and drop the subscript 1 in the rest of this section.

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

$$F(T, 1) = \mathbb{E}\left[\int_0^{\tau} (x_s - \hat{x}_s)^2 ds + J^*(T - \tau, N - 1)\right],$$

where

$$J^* = \operatorname*{ess\,inf}_{\{\tau_2, \ldots, \tau_N\}} \mathbb{E}\left[J(T - \tau, N - 1)\right].$$

This is a stopping problem in standard form, and to solve it we can use the so-called Snell envelope (see [18, Appendix D] and [24]):

$$\begin{aligned} S_t &= \operatorname*{ess\,inf}_{\tau \geq t} \quad \mathbb{E}\left[\int_0^{\tau} (x_s - \hat{x}_s)^2 ds + J^*(T - \tau, N - 1) \Big| \mathcal{F}_t^x\right], \\ &= \int_0^t (x_s - \hat{x}_s)^2 ds + \operatorname*{ess\,inf}_{\tau \geq t} \quad \mathbb{E}\left[\int_t^{\tau} (x_s - \hat{x}_s)^2 ds + J^*(T - \tau, N - 1) \Big| x_t\right]. \end{aligned}$$

Then, the earliest time when the cost of stopping does not exceed the Snell envelope is an optimal stopping time. Thus we get a simple threshold solution for our problem.

2.4. Extensions to nonlinear and partially observed systems. When the plant is nonlinear, the MMSE estimate under deterministic sampling is the mean of the Fokker–Planck equation and is given by

$$\xi_t = \mathbb{E}\left[x_t \Big| \tau_{latest}, x_{\tau_{latest}}\right], \quad \text{where}, \tau_{latest} = \sup\{d_i \leq t\}.$$

Under adaptive sampling, this may not be the optimal choice of estimate. To obtain a tractable optimization problem we can restrict the kind of estimator waveforms allowed at the supervisor. Using the Fokker–Planck mean above leads to a tractable stopping problem, as does use of the following special zero-order hold waveform:

$$\xi_t = x_{\tau_{latest}}.$$

However, even a slightly more general piecewise constant estimate

$$\xi_t = \mu\left(x_{\tau_{latest}}, \tau_{latest}\right)$$

leads to a stopping problem of nonstandard form because τ and μ have to be chosen in concert.

When the plant sensor has noisy observations, or in the vector case, noisy partial observations, the sampling problem remains unsolved. The important question now is, *What signal at the sensor should be sampled? Should the raw sensor measurements be sampled and transmitted, or is it profitable to process them first?* We propose a solution with a separation into local filtering and sampling. Accordingly, the sensor should compute a continuous filter for the state. The sufficient statistics for this filter should take the role of the state variable. This means that the sensor should transmit current samples of the sufficient statistics, at sampling times that are stopping times w.r.t. the sufficient statistics process.

In the case of a scalar linear system with observations corrupted by white noise, the local Kalman filter at the sensor $\hat{x}_t^{sensor}$ plays the role of the state signal. The Kalman filter obeys a linear evolution equation and so the optimal sampling policies presented in this paper should be valid. In the rest of the paper, we will investigate and solve the sampling problem, first for the Brownian motion process and then for the Ornstein–Uhlenbeck process.

3. Sampling Brownian motion. The sampling problem for Brownian motion with a control term added to the drift is no different from the problem without it. This is because the control process $\{u_t\}_{t\geq 0}$ is measurable w.r.t. $\mathcal{F}_t^{\mathbb{R}\mathbb{X}}$, whether it is a deterministic feed-forward term or a feedback based on the sampled information. Thus for the estimation problem, we can safely set the control term to be zero to get

$$dx_t = bdB_t,\ x_0 = x.$$

The diffusion coefficient b can be assumed to be unity. If it is not, we can simply scale time, and in the $\frac{t}{b^2}$-time, the process obeys an SDE driven by a Brownian motion with a unit diffusion coefficient. We study the sampling problem under the assumption that the initial state is known to the MMSE estimator. Under deterministic sampling, the MMSE estimate for this process is a zero-order hold extrapolation of received samples.

We study three important classes of sampling. The optimal deterministic one is traditionally used, and it provides an upper bound on the minimum distortion possible. The first adaptive scheme we study is Delta sampling, which is based on first hitting times of symmetric levels by the error process. Finally, we completely characterize the optimal sampling scheme by recursively solving an optimal multiple stopping problem.

3.1. Optimal deterministic sampling. Given that the initial value of the error signal is zero, we will show through induction that uniform sampling on the interval

$[0,T]$ is the optimal choice of N deterministic sample times. Call the deterministic set of sample times

$$\mathcal{D} = \{d_1, d_2, \ldots, d_N \mid 0 \leq d_i \leq T, \quad d_{i-1} \leq d_i \text{ for } i = 2, \ldots, N\}.$$

Then, the distortion takes the form

$$J_{Deter}(T,N) = \int_0^{d_1} \mathbb{E}(x_s - \hat{x}_s)^2 ds + \int_{d_1}^{d_2} \mathbb{E}(x_s - \hat{x}_s)^2 ds + \ldots + \int_{d_N}^{T} \mathbb{E}(x_s - \hat{x}_s)^2 ds.$$

Consider the situation of having to choose exactly one sample over the interval $[T_1, T_2]$ with the supervisor knowing the state at time T_1. The best choice of the sample time which minimizes the cost $J_{Deter}(T_2 - T_1,\ 1)$ is the midpoint $\frac{1}{2}(T_2 + T_1)$ of the given interval. On this basis, we propose for $N > 2$ that the optimal choice of $N-1$ deterministic times over $[T_1, T_2]$ is the uniform one:

$$\{d_1, d_2, \ldots d_{N-1}\} = \left\{ T_1 + i\frac{T_2 - T_1}{N} \middle| i = 1, 2, \ldots, N-1 \right\}.$$

This gives a distortion equaling $\frac{1}{2N}(T_2 - T_1)^2$. Let $J^*_{Deter}(T_2 - T_1,\ N)$ be the minimum distortion over $[0, T_2 - T_1]$ using N samples generated at deterministic times. Now, we carry out the induction step and obtain the minimum distortion over the set of N sampling times over $[T_1, T_2]$ to be

$$\begin{aligned}
&\min_{d_1} \left\{ \int_0^{d_1} (x_s - \hat{x}_s)^2 ds + \min_{\{d_2, d_2, \ldots d_N\}} J_{Deter}(T_2 - T_1 - d_1,\ N-1) \right\} \\
&= \min_{d_1} \left\{ \frac{{d_1}^2}{2} + \frac{(T_2 - T_1 - d_1)^2}{2N} \right\} \\
&= \min_{d_1} \left\{ \frac{N d_1^2 + d_1^2 - 2 d_2 (T_2 - T_1) + (T_2 - T_1)^2}{2N} \right\} \\
&= \min_{d_1} \left\{ \frac{(N+1)\Big(d_1 - \frac{1}{(N+1)}(T_2 - T_1)\Big)^2 + \Big(1 - \frac{1}{(N+1)}\Big)(T_2 - T_1)^2}{2N} \right\} \\
&= \frac{1}{2(N+1)}(T_2 - T_1)^2,
\end{aligned}$$

the minimum being achieved for $d_1 = \frac{1}{N+1}(T_2 - T_1)$. This proves the assertion about the optimality of uniform sampling among all deterministic schemes provided that the supervisor knows the value of the state at the start time.

3.2. Optimal Delta sampling. As described before, Delta sampling is a simple event-triggered sampling scheme which generates a new sample whenever the input signal differs from the last sample by a prespecified threshold. Delta sampling is really meant for infinite horizon problems as it produces intersample intervals that are unbounded. Since we have on our hands a finite horizon problem, we will use a time-out at the end time of the problem's horizon. To make the most of this class of rules, we allow the thresholds to vary with the past history of sample times. Thus the supervisor can compute the sequence of thresholds from the record of samples received previously. Only the sensor can find the actual sample time since it also has full access to the state and error signals.

More precisely, at any sampling time as well as at the start of the horizon, the threshold for the next sampling time is chosen. This choice is allowed to depend on the number of samples remaining as well as the amount of time left till the end of the horizon. We set $\tau_0 = 0$ and define thresholds and sampling times times recursively. The threshold for the ith sampling time is allowed to depend on the values of the previous sampling times, and so it is measurable w.r.t. $\mathcal{F}_t^{\mathbb{RX}}$. Assume that we are given the policy for choosing causally a sequence of nonnegative thresholds $\{\delta_1, \delta_2, \ldots, \delta_N\}$. Then for $i = 1, 2, \ldots, N$, we can characterize the sampling times $\{\zeta_1, \zeta_2, \ldots, \zeta_N\}$ as follows:

$$\begin{aligned} \mathcal{F}^{\delta_i} &\subset \mathcal{F}^{(\tau_1,\ldots,\tau_{i-1})} \text{ if } i \ > \ 1, \\ \tau_{i,\delta_i} &= \inf\left\{t \middle| t \geq \tau_{i-1,\delta_{i-1}}, \left|x_t - x_{\tau_{i-1}}\right| \geq \delta_i\right\}, \\ \zeta_i &= \min\left\{\tau_{i,\delta_i}, T\right\}. \end{aligned}$$

The first threshold δ_1 depends only on the length of the horizon, namely, T.

The optimal thresholds can be chosen one at a time using solutions to a nested sequence of optimization problems each with a single threshold as its decision variable. This is because, knowing how to choose the sequence $\{\zeta_{i+1}, \ldots, \zeta_N\}$ optimally, we can obtain an optimal choice for ζ_i by solving over the optimization problem:

$$\inf_{\delta_i \geq 0} \mathbb{E}\left[\int_0^{\zeta_i} \left(x_s - \hat{x}_s\right)^2 ds + J_\delta^*\left(T - \zeta_i,\ N - i\right)\right],$$

where the cost function $J_\delta^*(T - \zeta_i,\ N - i)$ is the minimum aggregated distortion over $[T - \zeta_i, T]$ achievable using at most $N - i$ samples generated using thresholds for the magnitude of the error signal. Hence, if we know how to generate the last sample efficiently, we can inductively figure out rules governing the best thresholds for earlier sampling times.

3.2.1. Optimal level for a single sample. These computations are carried out in Appendices B and C. In particular, (C.1) gives the expression

$$J_\delta\left(T,\ 1\right)(\lambda) = \frac{T^2}{2}\left\{1 + \frac{\pi^4}{32\lambda^2} - \frac{\pi^2}{4\lambda} - \frac{\pi}{\lambda^2}\sum_{k\geq 0}(-1)^k \frac{e^{-(2k+1)^2\lambda}}{(2k+1)^3}\right\},$$

where $\lambda = \frac{T\pi^2}{8\delta^2}$. Parametrizing in terms of λ reveals some structural information about the solution. First, note that the length of the time horizon does not directly affect the optimum choice of λ. The function $J_\delta\left(T,\ 1\right)$ has a shape that does not depend on T. It is merely scaled by the factor $\frac{T^2}{2}$. The behavior of the distortion as λ is varied can be seen in Figure 3.1(b). The minimum distortion incurred turns out to be

$$c_1 \frac{T^2}{2} = 0.3952 \frac{T^2}{2},$$

this being achieved by the choice $\delta^* = 0.9391\sqrt{T}$. As compared to deterministic sampling, whose optimum performance is $0.5\frac{T^2}{2}$, we have slightly more than 20% improvement by using the optimum thresholding scheme.

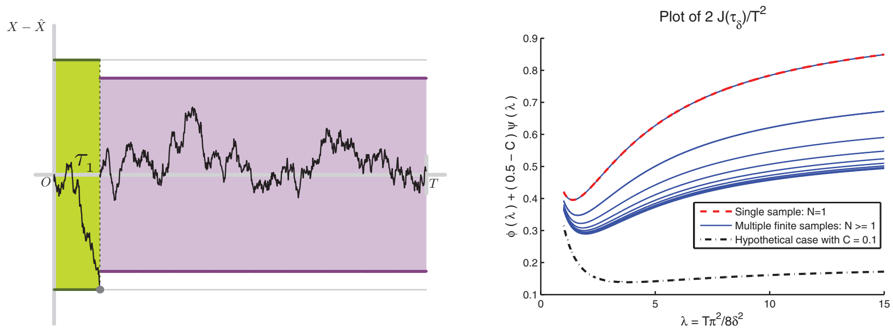

(a) Only one sample used

(b) Distortion as a function of threshold.

FIG. 3.1. (a) *Delta sampling.* (b) *Estimation distortion due to Delta sampling as a function of the threshold used. Notice that for a fixed δ, the distortion decreases steadily as the number of samples remaining (N) grows. The distortion, however, never reaches zero. The minimum distortion reaches its lower limit of* $0.287\frac{T^2}{2}$.

How often does the Delta sampler actually generate a sample? To determine that, we need to compute the probability that the estimation error signal reaches the threshold before the end time T. Equation (C.3) provides the answer: 98%. Note that this average sampling rate of the optimal Delta sampler is independent of the length of the time horizon.

We have the performance of the Delta sampler when the sample budget is one. Now we will compute the performance for larger budgets, and we will find that for budgets larger than one, it is actually more efficient to sample at deterministic times.

3.2.2. Multiple Delta sampling. Like in the single sample case, we will show that the expected distortion over $[0, T]$ given at most N samples is of the form

$$c_N \frac{T^2}{2}.$$

Let τ_δ be the level-crossing time as before. Then, given a positive real number α, consider the following cost:

$$\Upsilon\left(T, \alpha, \delta\right) \triangleq \mathbb{E}\left[\int_0^{\tau_\delta \wedge T} x_s^2 ds + \alpha\Big[\left(T-\tau_\delta\right)^+\Big]^2\right].$$

Using the same technique as in the single sample case (precisely, the calculations between and including (B.1), (B.3)), we get

$$\Upsilon\left(T, \alpha, \delta\right) = \frac{T^2}{2} - \delta^2 \mathbb{E}\Big[\left(T-\tau_\delta\right)^+\Big] - \left(\frac{1}{2} - \alpha\right) \mathbb{E}\left[\Big[\left(T-\tau_\delta\right)^+\Big]^2\right].$$

Using calculations presented in Appendices B and C we can write (C.2)

$$\Upsilon\left(T, \alpha, \delta\right) = \frac{T^2}{2}\left\{\phi(\lambda) + \left[\frac{1}{2} - \alpha\right]\psi(\lambda)\right\},$$

where, $\lambda = \frac{T\pi^2}{8\delta^2}$, and we define the functions ϕ, ψ as follows:

$$\phi(\lambda) \triangleq 1 + \frac{\pi^4}{32\lambda^2} - \frac{\pi^2}{4\lambda} - \frac{\pi}{\lambda^2}\sum_{k\geq 0}\frac{(-1)^k e^{-(2k+1)^2\lambda}}{(2k+1)^3}$$

and

$$\psi(\lambda) \triangleq -\frac{5\pi^4}{96\lambda^2} - \frac{\pi^2}{2\lambda} - 2 + \frac{16}{\pi\lambda^2}\sum_{k\geq 0}\frac{(-1)^k e^{-(2k+1)^2\lambda}}{(2k+1)^5}.$$

The choice of λ that minimizes the cost Υ can be determined by performing a grid search for the minimum of the scalar function $\phi(\lambda) + [\frac{1}{2} - \alpha]\psi(\lambda)$. Since this sum is a fixed function, we conclude that the minimum cost is a *fixed* percentage of $\frac{T^2}{2}$ exactly as in the case of the single sample. This property of this optimization problem is what lets us compute the optimal sequence of thresholds by induction.

Consider the distortion when N samples are generated using a Delta sampler, with N being at least 2. If we have the optimal Delta samplers for utilizing a budget of $N-1$ or less, then the minimum distortion with a budget of N takes the form

$$\begin{aligned} J^*_\delta(T,\ N) &= \inf_{\delta_N \geq 0} \mathbb{E}\left[\int_0^{\zeta_N} (x_s - \hat{x}_s)^2 ds + J^*_{delta}(T - \zeta_N,\ N-1)\right] \\ &= \inf_{\delta_N \geq 0} \mathbb{E}\left[\int_0^{\tau_{\delta_N}\wedge T} (x_s - \hat{x}_s)^2 ds + J^*_\delta\Big(\big(T - \tau_{\delta_N}\big)^+,\ N-1\Big)\right]. \end{aligned}$$

When the budget is zero, the distortion at the supervisor is $\frac{T^2}{2}$. When the budget is one, the minimum distortion is a fixed fraction of $\frac{T^2}{2}$, namely, $c_1\frac{T^2}{2}$. Similarly, by mathematical induction, we find the minimum distortions under higher budgets to be smaller fractions of $\frac{T^2}{2}$. Let the positive coefficient c_k stand for the hypothetical fraction whose product with $\frac{T^2}{2}$ is the minimum distortion $J^*_\delta(T,\ k)$. Continuing the previous set of equations, we get

$$\begin{aligned} J^*_\delta(T,\ N) &= \inf_{\delta_N \geq 0} \mathbb{E}\left[\int_0^{\tau_{\delta_N}\wedge T} (x_s - \hat{x}_s)^2 ds + c_{N-1}\left[\Big(T - \tau_{\delta_N}\Big)^+\right]^2\right] \\ &= \inf_{\delta_N \geq 0} \Upsilon\Big(T, c_{N-1}, \delta_N\Big) \\ &= \frac{T^2}{2} \inf_{\lambda_N = \frac{T\pi^2}{8\delta_N^2}} \left\{\phi(\lambda_N) + \left[\frac{1}{2} - c_{N-1}\right]\psi(\lambda_N)\right\}. \end{aligned}$$

Because of the scale-free nature of the functions ϕ, ψ, we have proved that the minimum distortion is indeed a fixed fraction of $\frac{T^2}{2}$. Figure 3.1(b) shows for different values of N the graph of $J_\delta(T,\ N)$ as a function of λ. The last equation gives us the following recursion for $k = 1, 2, \ldots, N$:

$$\boxed{\begin{aligned} c_k &= \inf_\lambda \{\phi(\lambda) + (0.5 - c_{k-1})\psi(\lambda)\}, & \rho_k &= \frac{\pi}{2\sqrt{2\lambda^*_k}} \text{ and} \\ \lambda^*_k &= \arg\inf_\lambda \{\phi(\lambda) + (0.5 - c_{k-1})\psi(\lambda)\}, & \delta^*_k &= \rho_{N-k+1}\sqrt{T - \zeta_{k-1}}. \end{aligned}} \quad (3.1)$$

TABLE 3.1
Characteristics of optimal multiple Delta sampling for small values of the sample budget.

N	1	2	3	4	5
c_N	0.3953	0.3471	0.3219	0.3078	0.2995
ρ_N	0.9391	0.8743	0.8401	0.8208	0.8094
$\mathbb{E}[\Xi_N]$	0.9767	1.9306	2.8622	3.7541	4.4803

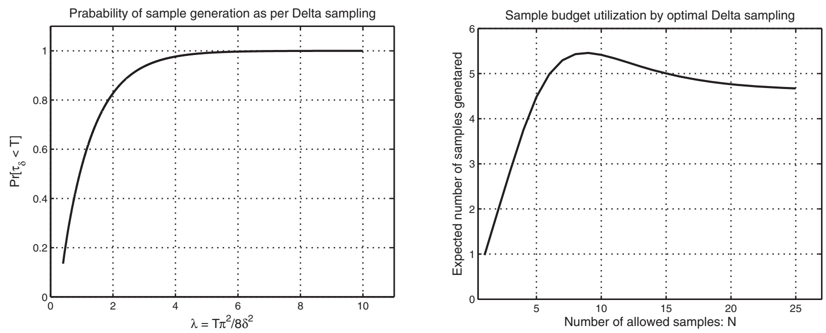

(a) Firing rate of Delta sampling (b) Average rate of sample generation

FIG. 3.2. (a) *Probability that a sample is generated (Ξ) as a function of the parameter λ, which is inversely related to the square of the threshold δ.* (b) *Delta sampling is shown to be ill-suited for repeated sampling over finite horizons. The average sample usage of optimal delta sampling does not rise monotonically with the budget and is actually counterintuitive. In fact, for any finite budget, the average sample usage is less than six.*

Now we determine the expected sampling rate of the optimal Delta sampler. Let Ξ_k be the random number of samples generated before T by the Delta sampler with a budget of k samples. Then almost surely, Ξ_k equals the number of threshold crossings generated by this sampler. Clearly, we have the bounds $0 \le \Xi_k \le k$. Also, under optimal Delta sampling, the statistics of the sampling rate do not depend on the length of the time interval T as long as the latter is nonzero. This gives us the recursion

$$\mathbb{E}[\Xi_k] = 0 \cdot \mathbb{P}\left[\tau_{\delta_1^*} \ge T\right] + (1 + \mathbb{E}[\Xi_{k-1}]) \cdot \mathbb{P}\left[\tau_{\delta_1^*} < T\right], \tag{3.2}$$

where δ_1^* is the optimal threshold for the first sample when the budget is k. The performance of optimal Delta sampling for small values of k are given in Table 3.1. To understand the behavior of optimal Delta sampling when the sample budget is larger than five, look at Figures 3.2(b) and 3.4. The minimum distortion decreases with increasing sample budgets but it does not decay to zero. It stagnates at approximately $0.3\frac{T^2}{2}$ no matter how large a budget is provided. The expected number of samples does not monotonically rise with the budget. It settles at a value close to 4.5. Clearly, Delta sampling is far from optimal over finite horizons. In fact, if the sample budget is at least two, even deterministic sampling performs better.

In optimal Delta sampling, the sensor chooses a sequence of thresholds to be applied on the estimation error signal. The choice of a particular threshold is made at the time of the previous sample and is allowed to depend on the past history of sample times. *Suppose now that the sensor is allowed to modify this choice causally and continuously at all time instants.* Then we get a more general class of sampling policies with a family of continuously varying envelopes for the estimation error signal. This class of policies happens to contain the optimal sampling policy which achieves the minimum possible distortion. Next, we will obtain the optimal family of envelopes by studying the problem of minimum distortion as an optimal multiple stopping problem.

3.3. Optimal sampling. Consider the nondecreasing sequence $\{\tau_1, \tau_2, \ldots, \tau_N\}$ with each element lying within $[0, T]$. For this to be a valid sequence of sampling times, its elements have to be stopping times w.r.t. the x-process. We will look for the best choice of these times through the optimization

$$J^*(T,\ N) = \inf_{\{\tau_1,\ldots,\tau_N\}} \mathbb{E}\left[\int_0^{\tau_1} {x_s}^2 ds + \int_{\tau_1}^{\tau_2} (x_s - \hat{x}_{\tau_1})^2 ds + \cdots + \int_{\tau_N}^{T} (x_s - \hat{x}_{\tau_N})^2 ds\right].$$

The solution to this optimization parallels the developments for Delta sampling. In particular, the minimum distortion obtained by optimal sampling will turn out to be a fraction of $\frac{T^2}{2}$. We will recursively obtain optimal sampling policies by utilizing the solution to the following optimal (single) stopping problem concerning the objective function χ:

$$\inf_\tau \chi(T, \beta, \tau) = \inf_\tau \mathbb{E}\left[\int_0^\tau {x_s}^2 ds + \frac{\beta}{2}(T-\tau)^2\right],$$

where τ is a stopping time w.r.t. the x-process that lies in the interval $[0, T]$ and β is a positive real number. We reduce this stopping problem into one having just a terminal cost using the calculations between and including (B.1), (B.3),

$$\chi(T, \beta, \tau) = \frac{T^2}{2} - \mathbb{E}\left[2x_\tau^2 (T-\tau) + (1-\beta)\left[(T-\tau)\right]^2\right],$$

which can be minimized by solving the following optimal stopping problem:

$$\operatorname*{ess\,inf}_\tau \mathbb{E}\left[2x_\tau^2 (T-\tau) + (1-\beta)\left[(T-\tau)\right]^2\right].$$

This stopping problem can be solved explicitly by determining its Snell envelope process. We look for a C^2 function $g(x, t)$ which satisfies the free boundary PDE system:

$$\frac{1}{2} g_{xx} + g_t = 0 \text{ and } g(x, t) \geq 2x^2 (T-t) + (1-\beta)(T-t)^2. \tag{3.3}$$

Given a solution g, consider the process

$$S_t \triangleq g(x_t, t).$$

This is in fact the Snell envelope. To see that, fix a deterministic time t within $[0, T]$ and verify using Itô's formula that

$$\mathbb{E}\left[S_\tau(x_\tau)|x_t\right] - S_t = \mathbb{E}\left[\int_t^\tau dS_t | x_t\right] = 0$$

for any stopping time $\tau \in [t, T]$, and hence,

$$S_t = \mathbb{E}\left[S_\tau | x_t\right] \ \geq \ \mathbb{E}\left[x_\tau^2 (1-\tau) | \tau \geq t, x_t\right].$$

The last equation confirms that S_t is indeed the Snell envelope. Consider the following solution to the free-boundary PDE system:

$$g(x, t) = A\left\{(T-t)^2 + 2x^2 (T-t) + \frac{x^4}{3}\right\}.$$

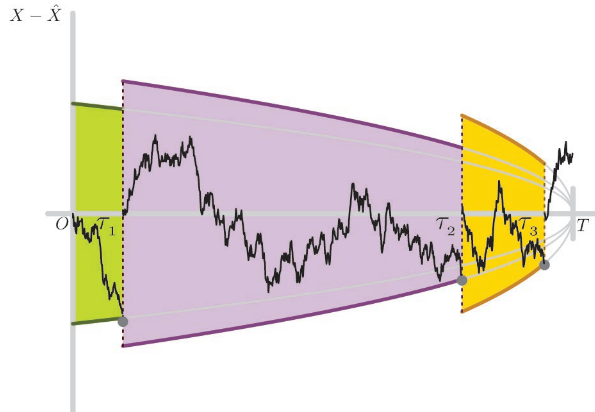

FIG. 3.3. *Optimal envelopes for the estimation error when the signal is a Brownian motion.*

where A is a constant chosen such that $g(x,t) - 2x^2(T-t) - (1-\beta)(T-t)^2$ becomes a perfect square. The only possible value for A then is

$$\frac{(5+\beta) - \sqrt{(5+\beta)^2 - 24}}{4}.$$

Then the first time when the reward equals or exceeds the Snell envelope is optimal

$$\begin{aligned}\tau^* &= \inf_t \left\{ t \,\middle|\, S_t \leq 2x_t^2(T-t) + (1-\beta)(T-t)^2 \right\}, \\ &= \inf_t \left\{ t \,\middle|\, x_t^2 \geq \sqrt{\frac{3(A-1+\beta)}{A}}(T-t) \right\}\end{aligned}$$

and the corresponding minimum distortion becomes

$$\chi^* = (1-A)\frac{T^2}{2}.$$

We now examine the problem of choosing optimally a single sample.

3.3.1. Optimal choice of a single sample. The minimum distortion due to using exactly one sample is

$$\begin{aligned}J^*(T,\ 1) &= \inf_{\tau_1} \mathbb{E}\left[\int_0^{\tau_1} {x_s}^2 ds + \int_{\tau_1}^{T} \left(x_s - \hat{x}_{\tau_1}\right)^2 ds\right] \\ &= \inf_{\tau_1} \mathbb{E}\left[\int_0^{\tau_1} {x_s}^2 ds + \frac{1}{2}(T-\tau_1)^2 ds\right] \\ &= \inf_{\tau_1} \chi\left(T, 1, \tau_1\right).\end{aligned}$$

We have thus reduced the optimization problem to one whose solution we already know. Hence, we have

$$\tau_1^* = \inf_{t\geq 0}\left\{ t \,\middle|\, {x_t}^2 \geq \sqrt{3}\,(T-t) \right\} \text{ and } J^*(T,\ 1) = \left(\sqrt{3}-1\right)\frac{T^2}{2}.$$

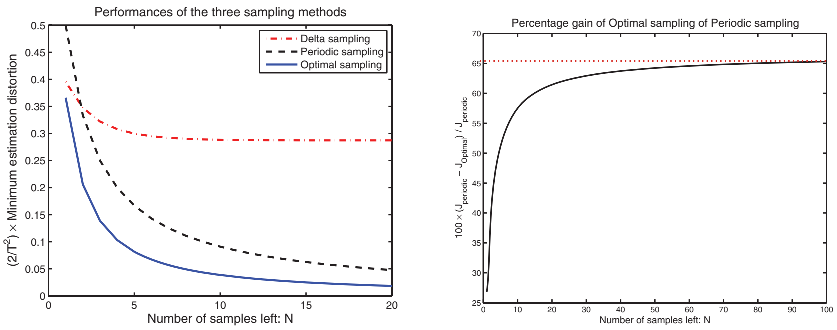

(a) Comparisons

(b) Optimal vs. periodic

FIG. 3.4. *The minimum distortions of the three sampling methods for the Brownian motion process. As the budget grows, so does the relative efficiency of optimal sampling over periodic sampling, and this efficiency aymptotically reaches a limit of* 67%.

3.3.2. Optimal multiple sampling. We obtain the family of policies for optimal multiple sampling by mathematical induction. Suppose that the minimum distortions due to using no more than $k-1$ samples over $[0,T]$ is given by the sequence of values $\{\theta_1 \frac{T^2}{2}, \ldots, \theta_{k-1}\frac{T^2}{2}\}$. Then consider the minimal distortion due to using up to k samples:

$$\begin{aligned} J^*(T,\ k) &= \inf_{\tau_1} \mathbb{E}\left[\int_0^{\tau_1} {x_s}^2 ds + J^*(T-\tau_1,\ k-1)\right] \\ &= \inf_{\tau_1} \mathbb{E}\left[\int_0^{\tau_1} {x_s}^2 ds + \frac{\theta_{k-1}}{2}(T-\tau_1)^2 ds\right] \\ &= \inf_{\tau_1} \chi(T, \theta_{k-1}, \tau_1). \end{aligned}$$

This proves the hypothesis that the minimum distortions for increasing values of the sample budget form a sequence with the form $\{\theta_k \frac{T^2}{2}\}_{k\ge 1}$. The last equation also provides us with the recursion which is started with $\theta_0 = 1$:

$$\boxed{\begin{aligned} \theta_k &= 1 - \frac{(5+\theta_{k-1}) - \sqrt{(5+\theta_{k-1})^2 - 24}}{4}, \\ \gamma_k &= \sqrt{\frac{3(\theta_{k-1}-\theta_k)}{1-\theta_k}}, \\ \tau_k^* &= \inf_{t\ge\tau_k} \left\{ t : (x_t - x_{\tau_k})^2 \ge \gamma_{N-k+1} T - t \right\}. \end{aligned}} \tag{3.4}$$

3.4. Comparisons. In Figure 3.4 we have a comparison of the estimation distortions incurred by the three sampling strategies. The remarkable news is that Delta sampling which is optimal for the infinite horizon version of the estimation problem is easily beaten by the best deterministic sampling policy. There is something intrinsic to Delta sampling which makes it ill-suited for finite horizon problems with hard budget limits. This also means that it is not safe to settle for "natural" event-triggered sampling policies such as Delta sampling. Also, notice that the relative gain of optimal sampling over periodic sampling consistently grows to about 67%.

4. Sampling the Ornstein–Uhlenbeck process. Now we turn to the case when the signal is an Ornstein–Uhlenbeck process,

$$dx_t = ax_t dt + dW_t, \quad t \in [0,T], \tag{4.1}$$

with $x_0 = 0$ and W_t being a standard Brownian motion. Again, the sampling times $\mathrm{S} = \{\tau_1, \ldots, \tau_N\}$ have to be an increasing sequence of stopping times with respect to the x-process. They also have to lie within the interval $[0,T]$. Based on the samples and the sample times, the supervisor maintains an estimate waveform $\hat{x}_t$ given by

$$\hat{x}_t = \begin{cases} 0 & \text{if } 0 \le t < \tau_1, \\ x_{\tau_i} e^{a(t-\tau_i)} & \text{if } \tau_i \le t < \tau_{i+1} \le \tau_N, \\ x_{\tau_N} e^{a(t-\tau_N)} & \text{if } \tau_N \le t \le T. \end{cases} \tag{4.2}$$

The quality of this estimate is measured by the aggregate squared error distortion:

$$J^*(T,\ N) = \mathbb{E}\left[\int_0^T (x_s - \hat{x}_s)^2 ds\right].$$

4.1. Optimal deterministic sampling. Just like in the case of Brownian motion, we can show through mathematical induction that uniform sampling on the interval $[0,T]$ is the optimal deterministic choice of N samples. For the induction step, we assume that the optimal choice of $N-1$ deterministic samples over $[T_1, T_2]$ is the uniform one:

$$\{d_1, d_2, \ldots d_N\} = \left\{T_1 + i\frac{T_2 - T_1}{N+1}\middle| i = 1, 2, \ldots, N\right\}.$$

Then, the corresponding minimum distortion becomes

$$(N+1)\,\frac{e^{2a\frac{T_2-T_1}{N+1}} - 1}{4a^2} - \frac{1}{2a}\,(T_2 - T_1)\,.$$

4.2. Optimal Delta sampling. We do not have an analytical characterization of the performance of Delta sampling. Let us first address the single sample case. The performance measure then takes the form

$$\begin{aligned} J_\delta(T,\ 1) &= \mathbb{E}\left[\int_0^{\zeta_1} x_t^2 + \int_{\zeta_1}^T (x_t - \hat{x}_t)^2\ dt\right] \\ &= \mathbb{E}\left[\int_0^T x_t^2 - 2\int_{\zeta_1}^T x_t \hat{x}_t\, dt + \int_{\zeta_1}^T (\hat{x}_t)^2\, dt\right]. \end{aligned}$$

Now notice that the second term can be written as

$$\mathbb{E}\left[\int_{\zeta_1}^T x_t \hat{x}_t\, dt\right] = \mathbb{E}\left[\int_{\zeta_1}^T \mathbb{E}[x_t|\mathcal{F}_{\zeta_1}]\hat{x}_t\, dt\right] = \mathbb{E}\left[\int_{\zeta_1}^T (\hat{x}_t)^2\, dt\right],$$

where we have used the strong Markov property of x_t, and that for $t > \zeta_1$ we have $\mathbb{E}[x_t|\mathcal{F}_{\zeta_1}] = x_\zeta e^{-a(t-\zeta_1)} = \hat{x}_t$. Because of this observation the performance measure takes the form

$$\begin{aligned} J_\delta(T,\ 1) &= \mathbb{E}\left[\int_0^T x_t^2\,dt - \int_{\zeta_1}^T (\hat{x}_t)^2\,dt\right] \\ &= \frac{e^{2aT}-1-2aT}{4a^2} - \mathbb{E}\left[x_{\zeta_1}^2 \frac{e^{2a(T-\zeta_1)}-1}{2a}\right] \\ &= T^2\left\{\frac{e^{2aT}-1-2aT}{4(aT)^2} - \mathbb{E}\left[\frac{x_{\zeta_1}^2}{T}\frac{e^{2(aT)(1-\zeta_1/T)}-1}{2(aT)}\right]\right\} \\ &= T^2\left\{\frac{e^{-2\bar{a}}-1+2\bar{a}}{4\bar{a}^2} - \mathbb{E}\left[-\bar{x}_{\bar{\zeta}_1}^2\frac{e^{2\bar{a}(1-\bar{\zeta}_1)}-1}{2\bar{a}}\right]\right\}, \end{aligned}$$

where

$$\bar{t} = \frac{t}{T};\quad \bar{a} = aT;\quad \bar{x}_{\bar{t}} = \frac{x_{\frac{t}{T}}}{\sqrt{T}}. \tag{4.3}$$

We have $\bar{x}$ satisfying the following SDE:

$$d\bar{x}_{\bar{t}} = -\bar{a}\bar{x}_{\bar{t}}d\bar{t} + dw_{\bar{t}}.$$

This suggests that without loss of generality, we can limit ourselves to the normalized case $T = 1$ since the case $T \neq 1$ can be reduced to the normalized one by using the transformations in (4.3). In fact, we can solve the single sampling problem on $[0, 1]$ to minimize

$$J_\delta(1,\ 1) = \left\{\frac{e^{-2a}-1+2a}{4a^2} - \mathbb{E}\left[-x_{\zeta_1}^2\frac{e^{2a(1-\zeta_1)}-1}{2a}\right]\right\}. \tag{4.4}$$

We carry over the definitions for threshold sampling times from section 3.2. We do not have series expansions like for the case of the Brownian motion process. Instead we have a computational procedure that involves solving a PDE initial and boundary value problem [20]. We have a nested sequence of optimization problems, the choice at each stage being the nonzero level δ_i. For $N = 1$, the distortion corresponding to a chosen δ_1 is given by

$$\frac{e^{2a}-1-2a}{4a^2} - \frac{\delta_1^2}{2a}\mathbb{E}\left[e^{2a(1-\zeta_1)}-1\right] = \frac{e^{2a}-1-2a}{4a^2} - \frac{\delta_1^2}{2a}\left\{e^{2a}\left(1+2aU^1(0,0)\right)-1\right\},$$

where the function $U^1(x,t)$ defined on $[-\delta_1, \delta_1]\times[0,1]$ satisfies the PDE

$$\frac{1}{2}U_{xx}^1 + axU_x^1 + U_t^1 + e^{-2at} = 0,$$

along with the boundary and initial conditions

$$\begin{cases} U^1(-\delta_1,t) = U^1(\delta_1,t) = 0 & \text{for}\ \ t\in[0,1], \\ U^1(x,1) = 0 & \text{for}\ \ x\in[-\delta_1,\delta_1]. \end{cases}$$

We choose the optimal δ_1 by computing the resultant distortion for increasing values of δ_1 and stopping when the cost stops decreasing and starts increasing. Note that the solution $U(0,t)$ to the PDE also furnishes us with the performance of the δ_1-triggered sampling over $[t, 1]$. We will use this to solve the multiple sampling problem.

Let the optimal policy of choosing N levels for sampling over $[T_1, 1]$ be given where $0 \leq T_1 \leq 1$. Let the resulting distortion also be known as a function of T_1. Let this known distortion over $[T_1, 1]$ given N level-triggered samples be denoted $J^*_\delta(1 - T_1,\ N)$. Then, the $N+1$ sampling problem can be solved as follows. Let $U^{N+1}(x, t)$ satisfy the PDE

$$\frac{1}{2}U^{N+1}_{xx} + axU^{N+1}_x + U^{N+1}_t = 0,$$

along with the boundary and initial conditions

$$\begin{cases} U^{N+1}(-\delta_1, t) = U^{N+1}(\delta_1, t) = J^*_\delta(1 - t,\ N) & \text{for } \ t \in [0, 1], \\ U^{N+1}(x, 1) = 0 & \text{for } \ x \in [-\delta_1, \delta_1]. \end{cases}$$

Then the distortion is given by

$$\begin{aligned} &\frac{e^{2a} - 1 - 2a}{4a^2} - \frac{\delta_1^2}{2a}\mathbb{E}\left[e^{2a(1-\zeta_1)} - 1 + \frac{e^{2a(1-\zeta_1)} - 1}{4a^2} - \frac{1 - \zeta_1}{2a}\right] + \mathbb{E}\left[J^*_\delta(1 - \zeta_1,\ N)\right] \\ &\qquad = \frac{e^{2a} - 1}{4a^2} - \frac{1}{2a} - \frac{\delta_1^2}{2a}\left\{e^{2a}\left(1 + 2aU^1(0, 0)\right) - 1\right\} + U^{N+1}(0, 0). \end{aligned}$$

We choose the optimal δ_1 by computing the resultant distortion for increasing values of δ_1 and stopping when the distortion stops decreasing.

4.3. Optimal sampling. We do not have analytic expressions for the minimum distortion like in the Brownian motion case. We have a numerical computation of the minimum distortion by finely discretizing time and solving the discrete-time optimal stopping problems.

By discretizing time, we get random variables $x_1, \ldots, x_M$ that satisfy the AR(1) model below. For $1 \leq n \leq M$ with $h = T/(M + 1)$,

$$x_n = e^{ah}x_{n-1} + w_n, \quad w_n \sim \mathcal{N}\left(0, \frac{e^{2ah} - 1}{2a}\right); \quad 1 \leq n \leq M.$$

The noise sequence $\{w_n\}$ is independently and identically distributed (i.i.d.) and Gaussian.

Sampling exactly once in discrete time means selecting a sample x_ν from the set of $M + 1$ sequentially available random variables $x_0, \ldots, x_M$ with the help of a stopping time $\nu \in \{0, 1, \ldots, M\}$. We can define the optimum cost to go which can be analyzed as follows. For $n = M, M - 1, \ldots, 0$, using (4.4),

$$\begin{aligned} V^1_n(x) &= \sup_{n \leq \nu \leq M} \mathbb{E}\left[x_\nu^2 \frac{e^{2ah(M-\nu)} - 1}{2a} \Big| x_n = x\right] \\ &= \max\left\{x^2 \frac{e^{2ah(M-n)} - 1}{2a},\ \mathbb{E}[V^1_{n+1}(x_{n+1})|x_n = x]\right\}. \end{aligned}$$

The above equation provides a (backward) recurrence relation for the computation of the single sampling value function $V^1_n(x)$. Notice that for values of x for which the left-hand side exceeds the right-hand side we stop and sample; otherwise we continue to the next time instant. We can prove by induction that the optimum policy is a *time-varying threshold* one. Specifically, for every time n there exists a threshold λ_n such

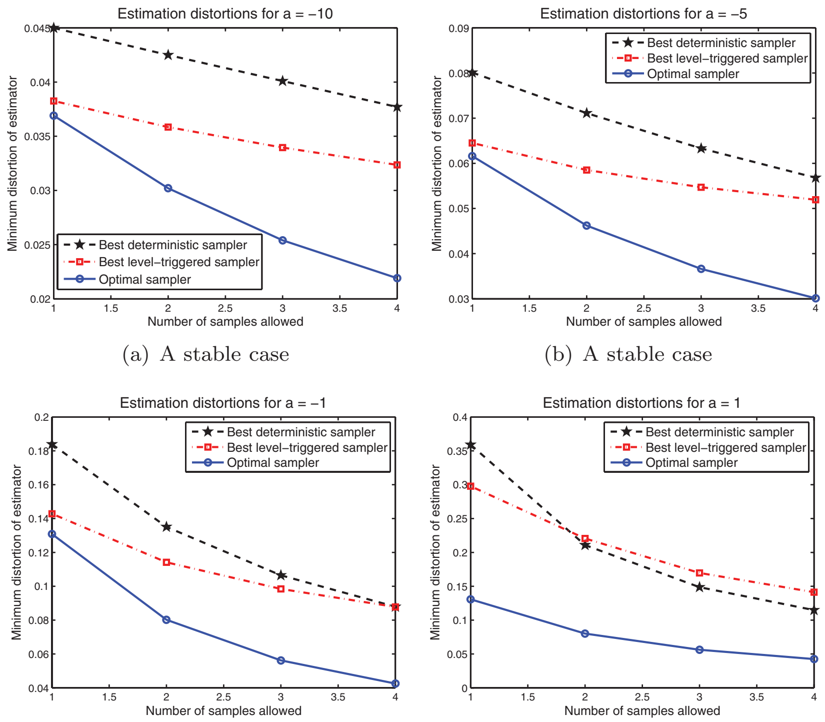

(a) A stable case (b) A stable case

(c) A stable case (d) An unstable case

FIG. 4.1. *The minimum distortions of the three sampling methods for the Ornstein–Uhlenbeck process. In the stable regime, for small budgets, Delta sampling is more efficient than deterministic sampling. In the unstable regime, deterministic sampling always beats Delta sampling.*

that if $|x_n| \geq \lambda_n$ we sample; otherwise we go to the next time instant. The numerical solution of the recursion presents no special difficulty if $a \leq 1$. For $a > 1$, we need to use careful numerical integration schemes in order to minimize the computational errors [20]. If $V_n^1(x)$ is sampled in x, then this function is represented as a vector. In the same way we can see that the conditional expectation is reduced to a simple matrix-vector product. Using this idea we can compute numerically the evolution of the threshold λ_t with time. The minimum expected distortion for this single sampling problem is

$$\frac{e^{2aT} - 1 - 2aT}{4a^2} - V_0^1(0).$$

For obtaining the solution to the $N+1$-sampling problem, we use the solution to the N-sampling problem. For $n = M, M-1, \ldots 0$,

$$\begin{aligned} V_n^{N+1}(x) &= \sup_{n \leq \nu \leq M} \mathbb{E}\left[V_\nu^N(0) + x_\nu^2 \frac{e^{2ah(M-\nu)} - 1}{2a} \middle| x_n = x \right] \\ &= \max\left\{ V_n^N(0) + x^2 \frac{e^{2ah(M-n)} - 1}{2a}, V_{n+1}^N(0) + \mathbb{E}\left[V_{n+1}^1(x_{n+1}) | x_n = x \right] \right\}. \end{aligned}$$

4.4. Comparisons. Figure 4.1 shows the result of the numerical computations for a few stable plants and a single unstable plant. Again, Delta sampling is not competitive. But in the stable cases, it provides a distortion lower than periodic sampling, when the sample budget is small.

5. Summary and extensions. We have set up the problem of efficient sampling as an optimal sequential sampling problem. We conjecture that the estimator under optimal sampling is the simple least-squares one under deterministic sampling. By fixing the estimate to be the same as the MMSE estimate under deterministic sampling, we reduce the optimization into a tractable stopping time problem. Our conjecture, of course, needs to be proved or disproved.

We have furnished methods to obtain good sampling policies for the finite horizon state estimation problem. When the signal is a Brownian motion, we have analytic solutions. When the signal is an Ornstein–Uhlenbeck process, we have provided computational recipes to determine the best sampling policies and their performances. In both cases, Delta sampling performs poorly with its distortion staying boundedly away from zero even as the sample budget increases to infinity. This means that the designer cannot just settle for “natural” event-triggered schemes without further investigation. In particular, a scheme optimal in the infinite horizon may perform badly on finite horizons.

The approach adopted in this paper leads us to also consider some sampling and filtering problems with multiple sensors. These can possibly be solved in the same way as the single sensor problem. The case where the samples are not reliably transmitted but can be lost in transmission is computationally more involved. There, the relative performances of the three sampling strategies is unknown. However, in principle, the best policies and their performances can be computed using nested optimization routines like we have used in this paper.

Another set of unanswered questions involves the performance of these sampling policies when the actual objective is not filtering but control or signal detection based on the samples. It will be very useful to know the extent to which the overall performance suffers when we minimize filtering error rather than the true cost. The communication constraint we treated in this paper was a hard limit on the number of allowed samples. Instead, we could use a soft constraint: a limit on the expected number of samples. We could also study the effect of mandatory minimum intervals between successive sampling times. Extension to nonlinear systems is needed, as are extensions to the case of partial observations at the sensor. One could follow the attack line sketched at the end of section 2.

Appendix A. Optimal sampling with nonstandard MMSE estimates. The conjecture expressed by (2.3) conforms to our intuition about scalar Gaussian diffusions. Here, we give an example of a well-behaved and widely used stochastic process for which the optimal sampling policy leads to an MMSE estimate which is different from that under deterministic sampling. Its increments do not have symmetric PDFs. For convenience, we consider an infinite horizon repeated sampling problem where the communication constraint is a limit on the average sampling rate.

Choose the state process to be the Poisson counter N_t, a continuous time Markov chain. This is a nondecreasing process which starts at zero and takes integer values. Its sample paths are piecewise constant and RCLL. The sequence of times between successive jumps is i.i.d., with an exponential distribution of parameter λ.

Under any deterministic sampling rule, the MMSE estimate is piecewise linear with slope λ and has the form

$$\hat{N}_t = N_{d_{latest}} + \lambda\left(t - d_{latest}\right), \tag{A.1}$$

where d_{latest} is the latest sampling instant as of time t. The optimal sampling policy leads to an MMSE estimate which is different.

Stipulate that the constraint on the average sampling rate is greater than or equal to λ the parameter of the Poisson process. Consider the following sampling policy whose MMSE estimate is of the zero-order hold type:

$$\hat{N}_t = N_{\tau_{latest}}, \tag{A.2}$$

$$\begin{cases} \tau_0 = 0, \\ \tau_{i+1} = \inf\left\{t \,|t > \tau_i, N_t > N_{\tau_i}\right\} \; \forall i \;\geq\; 0, \\ \tau_{latest} = \max\left\{\tau_i \,|\tau_i \leq t\right\}. \end{cases} \tag{A.3}$$

This sampling rule with its MMSE estimate $\hat{N}_t$ leads to an error signal which is identically zero. We also have that

$$\mathbb{E}\left[\tau_{i+1} - \tau_i\right] = \frac{1}{\lambda} \;\forall\; i \;\geq\; 0,$$

and so the communication constraint is met. On the other hand, the conventional MMSE estimate (A.1) would result in a nonzero average squared error distortion.

Suppose now that the distortion criterion is not the average value of the squared error but of the lexicographic distance:

$$D\left(N_t, \hat{N}_t\right) = \begin{cases} 0 & \text{if } N_t = \hat{N}_t, \\ 1 & \text{otherwise.} \end{cases}$$

Then, under deterministic sampling, the maximum likelihood estimate

$$\bar{N}_t = N_{d_{latest}} + \left\lfloor \lambda\left(t - d_{latest}\right)\right\rfloor$$

minimizes the average lexicographic distortion which will be nonzero. However, the adaptive policy (A.3) provides zero error reconstruction if the allowed average sampling rate is at least λ.

Appendix B. Threshold sampling once. We drop the subscript N for the terminal sample time

$$\tau_\delta = \inf_t \left\{t : |x_t - \hat{x}_t| = \delta\right\}$$

and its corresponding threshold δ. Here, δ is a threshold independent of the data acquired after time 0. Our goal is to compute the estimation distortion for any non-negative choice of the threshold and then select the one that minimizes the distortion:

$$J_\delta\left(T,\; 1\right)(\delta) = \mathbb{E}\left[\int_0^{\tau_\delta \wedge T} x_s^2 ds + \int_{\tau_\delta \wedge T}^{T} \left(x_s - x_{\tau_\delta \wedge T}\right)^2 ds\right].$$

By using iterated expectations on the second term, we get

$$\begin{aligned}
J_\delta(T,\ 1)(\delta) &= \mathbb{E}\left[\int_0^{\tau_\delta\wedge T} x_s^2 ds + \mathbb{E}\left[\int_{\tau_\delta\wedge T}^T (x_s - x_{\tau_\delta\wedge T})^2 ds \middle| \tau_\delta\wedge T, x_{\tau_\delta\wedge T}\right]\right] \\
&= \mathbb{E}\left[\int_0^{\tau_\delta\wedge T} x_s^2 ds + \int_{\tau_\delta\wedge T}^T \mathbb{E}\left[(x_s - x_{\tau_\delta\wedge T})^2 \middle| \tau_\delta\wedge T, x_{\tau_\delta\wedge T}\right] ds\right] \\
&= \mathbb{E}\left[\int_0^{\tau_\delta\wedge T} x_s^2 ds + \int_{\tau_\delta\wedge T}^T (s - \tau_\delta\wedge T)\, ds\right] \\
&= \mathbb{E}\left[\int_0^{\tau_\delta\wedge T} x_s^2 ds + \frac{1}{2}\left[(T-\tau_\delta)^+\right]^2\right]. \qquad \text{(B.1)}
\end{aligned}$$

We have thus reduced the distortion measure to a standard form with a running cost and a terminal cost. We will now take some further steps and reduce one with a terminal part alone. Notice that

$$d\left[(T-t)\,x_t^2\right] = -x_t^2 dt + 2\,(T-t)\,x_t dx_t + (T-t)\,dt,$$

which leads to the following representation for the running cost term:

$$\begin{aligned}
\mathbb{E}\left[\int_0^{\tau_\delta\wedge T} x_s^2 ds\right] &= \mathbb{E}\left[(T-\tau_\delta\wedge T)\,x_{\tau_\delta\wedge T}^2 + \frac{T^2}{2} - \frac{1}{2}(T-\tau_\delta\wedge T)^2\right] \\
&= \frac{T^2}{2} - \mathbb{E}\left[x_{\tau_\delta\wedge T}^2 (T-\tau_\delta)^+ + \frac{1}{2}\left[(T-\tau_\delta)^+\right]^2\right] \qquad \text{(B.2)} \\
&= \frac{T^2}{2} - \mathbb{E}\left[\delta^2 (T-\tau_\delta)^+ + \frac{1}{2}\left[(T-\tau_\delta)^+\right]^2\right].
\end{aligned}$$

Note that (B.2) is valid even if we replace τ_δ with a random time that is a stopping time w.r.t. the x-process. Thus, the cost (B.1) becomes

$$J_\delta(T,\ 1)(\delta) = \frac{T^2}{2} - \delta^2 \mathbb{E}\left[(T-\tau_\delta)^+\right]. \qquad \text{(B.3)}$$

If we can describe the dependence of the expected residual time $\mathbb{E}\left[(T-\tau_\delta)^+\right]$ on the threshold δ, then we can parametrize the cost purely in terms of δ. Had we known the PDF of τ_δ the computation of the expectation of the difference $(T-\tau_\delta)^+$ would have been easy. Unfortunately the PDF of the hitting time τ_δ does not have a closed-form solution. There exists a series representation [17, p. 99] which is

$$f_{\tau_\delta}(t) = \delta\sqrt{\frac{2}{\pi t^3}} \sum_{k=-\infty}^{\infty} (4k+1) e^{-\frac{(4k+1)^2\delta^2}{2t}}.$$

This series is not integrable and so it cannot meet our needs. Instead we compute the moment generating function of $(T-\tau_\delta)^+$ and thereby compute the expected distortion.

Appendix C. Statistics of an exit time curtailed by a time-out T. We start by deriving the moment generating function of the first hitting time τ_δ:

$$\tau_\delta = \inf_t \left\{ t \mid x_0 - \hat{x}_0 = w_0,\ \ |x_t - \hat{x}_t| = \delta \right\}.$$

LEMMA C.1. *If τ_δ is the first hitting time of $|x_t - \hat{x}_t|$ at the threshold δ, then*

$$\mathbb{E}[e^{-s\tau_\delta}] = \frac{\cosh(w_0\sqrt{2s})}{\cosh(\delta\sqrt{2s})} = F_\tau(s).$$

Proof. Consider the C^2 function $h(w,t) = e^{-st}[1 - \cosh(\sqrt{2s}w)/\cosh(\sqrt{2s}\delta)]$ and apply the Itô calculus on $h(w_t, t)$. We can then conclude that

$$\mathbb{E}\left[h\left(w_{\tau_\delta}, \tau_\delta\right)\right] - h(w_0, 0) = \mathbb{E}\left[\int_0^{\tau_\delta} [h_t(w_t,t) + 0.5h_{ww}(w_t,t)]\,dt\right] = \mathbb{E}[e^{-s\tau_\delta}] - 1,$$

from which we immediately obtain the desired relation because of the boundary condition: $h(w_{\tau_\delta}, \tau_\delta) = 0$. □

Lemma C.1 suggests that the PDF of the random variable τ_δ can be computed as $f_\tau(t) = \mathcal{L}^{-1}(F_\tau(s))$, that is, the inverse Laplace transform of $F_\tau(s)$. Invoking the initial condition $w_0 = 0$, we can then write

$$\begin{aligned}\mathbb{E}[(T-\tau_\delta)^+] &= \int_0^T (T-t) f_\tau(t)dt = \int_0^T (T-t)\left[\frac{1}{2\pi j}\oint F_\tau(s)e^{st}\,ds\right]dt \\ &= \frac{1}{2\pi j}\oint F_\tau(s)\left[\int_0^T (T-t)e^{st}\,dt\right]ds \\ &= \frac{1}{2\pi j}\oint \frac{e^{sT} - 1 - sT}{s^2\cosh(\delta\sqrt{2s})}\,ds,\end{aligned}$$

this contour integral being evaluated along a path that encloses the whole left half of the complex plane.

In order to compute this line integral over the complex plane, we need to find the poles of the integrand and then apply the residue theorem. Notice first that $s = 0$ is not a pole since the numerator has a double zero at zero. The only poles come from the zeros of the function $\cosh(\delta\sqrt{2s})$. Since $\cosh(x) = \cos(jx)$ we conclude that the zeros of $\cosh(\delta\sqrt{2s})$ which are also the poles of the integrand are

$$s_k = -(2k+1)^2\frac{\pi^2}{8\delta^2},\ k = 0, 1, 2, \ldots,$$

and they all belong to the negative half plane. This of course implies that they all contribute to the integral. We can now apply the residue theorem to conclude that

$$\mathbb{E}[(T-\tau)^+] = \frac{1}{2\pi j}\oint \frac{e^{sT} - 1 - sT}{s^2\cosh(\delta\sqrt{2s})}\,ds = \sum_{k\geq 0}\frac{e^{s_kT} - 1 - s_kT}{s_k^2}\lim_{s\to s_k}\frac{s - s_k}{\cosh(\delta\sqrt{2s})}.$$

In order to find the last limit we can assume that $s = s_k(1+\epsilon)$ and let $\epsilon \to 0$. Then we can show that

$$\lim_{s\to s_k}\frac{s - s_k}{\cosh(\delta\sqrt{2s})} = (-1)^{(k+1)}\frac{4s_k}{(2k+1)\pi}.$$

Using this expression, the performance measure of the stopping time τ_δ takes the following form:

$$\begin{aligned} J_\delta(T,\ 1) &= \frac{T^2}{2} - \delta^2 \mathbb{E}\Big[(T-\tau_\delta)^+\Big] \\ &= \frac{T^2}{2}\left\{1 - \frac{8\delta^2}{\pi T}\sum_{k\ge 0}(-1)^{(k+1)}\frac{1}{2k+1}\frac{e^{s_k T}-1-s_k T}{s_k T}\right\} \\ &= \frac{T^2}{2}\phi(\lambda), \end{aligned}$$

where with the change of variables $\lambda = \frac{T\pi^2}{8\delta^2}$, we have

$$\begin{aligned} \phi(\lambda) &\triangleq 1 - \frac{\pi}{\lambda^2}\sum_{k\ge 0}(-1)^k \frac{e^{-(2k+1)^2\lambda} - 1 + (2k+1)^2\lambda}{(2k+1)^3}, \\ &= 1 - \frac{\pi}{\lambda^2}\sum_{k\ge 0}\frac{(-1)^k e^{-(2k+1)^2\lambda}}{(2k+1)^3} + \frac{\pi}{\lambda^2}\sum_{k\ge 0}\frac{(-1)^k}{(2k+1)^3} - \frac{\pi}{\lambda}\sum_{k\ge 0}\frac{(-1)^k}{2k+1}. \end{aligned}$$

The final two series in the last equation can be summed explicitly. To do so, we adopt a summation technique described in the book of Aigner and Ziegler [1]. Consider

$$\int_0^1 \frac{dx}{1+x^2} \ = \ \int_0^1 \left(\sum_{k\ge 0}(-1)^k x^{2k}\right) dx \ = \ \sum_{k\ge 0}(-1)^k \int_0^1 x^{2k} dx \ = \ \sum_{k\ge 0}\frac{(-1)^k}{2k+1}.$$

By an easy evaluation of the definite integral we started with, we get a sum of $\frac{\pi}{4}$ for the series $\sum_{k\ge 0}\frac{(-1)^k}{2k+1}$; this result is useful because the series converges slowly. Proceeding along similar lines [6] and working with the multiple integral

$$\int_A \cdots \int \frac{dx_1 \cdots dx_n}{1 + (x_1 x_2 \cdots x_n)^2}$$

over the unit hypercube $A = [0,1]^n$ in $\mathbb{R}^n$, we get an explicit expression for the sum $\sum_{k\ge 0}\frac{(-1)^k}{(2k+1)^n}$ whenever n is an odd number. In particular,

$$\sum_{k\ge 0}\frac{(-1)^k}{(2k+1)^3} = \frac{\pi^3}{32}; \qquad\qquad \sum_{k\ge 0}\frac{(-1)^k}{(2k+1)^5} = \frac{5\pi^5}{1536}.$$

This reduces the distortion to

(C.1)
$$\boxed{J_\delta(T,\ 1) = \frac{T^2}{2}\phi(\lambda) = \frac{T^2}{2}\left\{1 + \frac{\pi^4}{32\lambda^2} - \frac{\pi^2}{4\lambda} - \frac{\pi}{\lambda^2}\sum_{k\ge 0}\frac{(-1)^k e^{-(2k+1)^2\lambda}}{(2k+1)^3}\right\}},$$

where $\lambda = \frac{T\pi^2}{8\delta^2}$.

The estimation distortion due to using $N+1$ samples when N is nonnegative is given through the recursion

$$J_\delta(T,\ N+1) = \frac{T^2}{2} - \delta^2\mathbb{E}\Big[(T-\tau_{\delta_1})^+\Big] - \left(\frac{1}{2} - J_\delta(T,\ N)\right)\mathbb{E}\left[\Big[(T-\tau_{\delta_1})^+\Big]^2\right].$$

To use this we need to know the statistics of the first sample time τ_{δ_1} and how this time determines the average distortion incurred by the remaining samples over the remainder of the horizon. Regardless of actual budget, J_δ takes the generic form

$$\Upsilon(T,\alpha,\delta) \triangleq \frac{T^2}{2} - \delta^2\mathbb{E}\Big[(T-\tau_\delta)^+\Big] - \left(\frac{1}{2}-\alpha\right)\mathbb{E}\left[\Big[(T-\tau_\delta)^+\Big]^2\right],$$

where α is positive but no greater than 0.5. This requires an evaluation of the second moment: $\mathbb{E}[[(T-\tau_\delta)^+]^2]$. We can calcuate it like we did the first moment:

$$\mathbb{E}\left[\Big[(T-\tau)^+\Big]^2\right] = \frac{1}{\pi j}\oint \frac{e^{sT}-1-sT-\frac{1}{2}s^2T^2}{s^3\cosh\left(\delta\sqrt{2s}\right)}\,ds.$$

This gives the expression for the cost $\Upsilon(T,\alpha,\delta)$

$$\Upsilon(T,\alpha,\delta) = \frac{T^2}{2}\left\{\phi(\lambda) + \left[\frac{1}{2}-\alpha\right]\psi(\lambda)\right\},$$

where $\lambda = \frac{T\pi^2}{8\delta^2}$, and we define functions ϕ,ψ with ϕ being the same as it was earlier in appendix:

$$\begin{aligned}\phi(\lambda) &\triangleq 1 - \frac{\pi}{\lambda^2}\sum_{k\geq 0}(-1)^k\frac{e^{-(2k+1)^2\lambda}-1+(2k+1)^2\lambda}{(2k+1)^3}\\ &= 1 - \frac{\pi}{\lambda^2}\sum_{k\geq 0}\frac{(-1)^k e^{-(2k+1)^2\lambda}}{(2k+1)^3} + \frac{\pi}{\lambda^2}\sum_{k\geq 0}\frac{(-1)^k}{(2k+1)^3} - \frac{\pi}{\lambda}\sum_{k\geq 0}\frac{(-1)^k}{2k+1}\end{aligned}$$

and

$$\begin{aligned}\psi(\lambda) &\triangleq \frac{16}{\pi\lambda^2}\sum_{k\geq 0}(-1)^k\frac{e^{-(2k+1)^2\lambda}-1+(2k+1)^2\lambda-0.5(2k+1)^4\lambda^2}{(2k+1)^5}\\ &= \frac{16}{\pi\lambda^2}\sum_{k\geq 0}\frac{(-1)^k\left(e^{-(2k+1)^2\lambda}-1\right)}{(2k+1)^5} + \frac{16}{\pi\lambda}\sum_{k\geq 0}\frac{(-1)^k}{(2k+1)^3} - \frac{8}{\pi}\sum_{k\geq 0}\frac{(-1)^k}{2k+1}.\end{aligned}$$

After replacing the summable series with their sums, the distortion due to multiple samples based on thresholds reduces to the boxed expression below. With $\lambda = \frac{T\pi^2}{8\delta^2}$,

(C.2)

$$\boxed{\begin{aligned}J_\delta(T,\ N+1) = \frac{T^2}{2}\Bigg\{1 + \frac{\pi^4}{32\lambda^2} - \frac{\pi^2}{4\lambda} - \frac{\pi}{\lambda^2}\sum_{k\geq 0}\frac{(-1)^k e^{-(2k+1)^2\lambda}}{(2k+1)^3}\\ + \left(0.5 - J_\delta(T,\ N)\right)\Bigg[\frac{-5\pi^4}{96\lambda^2} - \frac{\pi^2}{2\lambda} - 2\\ + \frac{16}{\pi\lambda^2}\sum_{k\geq 0}\frac{(-1)^k e^{-(2k+1)^2\lambda}}{(2k+1)^5}\Bigg]\Bigg\}.\end{aligned}}$$

To characterize the statistics of sample budget utilization by multiple Delta sampling, we need to find the probabilities of threshold crossings before the time-out T. Given a budget of N samples, let Ξ_N be the random number of samples generated under any timescale-free multiple Delta sampling scheme. Then we have

$$\begin{aligned}\mathbb{E}[\Xi_N] &= 0 \cdot \mathbb{P}\big[\tau_{\delta_1} \geq T\big] + (1 + \mathbb{E}[\Xi_{N-1}]) \cdot \mathbb{P}\big[\tau_{\delta_1} < T\big], \\ &= (1 + \mathbb{E}[\Xi_{N-1}]) \cdot \mathbb{P}\big[\tau_{\delta_1} < T\big],\end{aligned}$$

where δ_1 is the threshold for the first sample when the budget is N. As before, we use the moment generating function of the hitting time to obtain

$$\mathbb{E}\left[\Xi_1\right] = \mathbb{E}\left[\mathbf{1}_{\{\tau_{\delta_1} > T\}}\right] = \frac{1}{\pi j} \oint \frac{e^{sT} - 1}{s \cdot \cosh\left(\delta\sqrt{2s}\right)} \, ds.$$

With the notation $\lambda = \frac{T\pi^2}{8\delta_1^2}$, and evaluating this complex line integral as in previous cases, we obtain

$$\boxed{\mathbb{E}\left[\Xi_1\right] = 1 - \frac{4}{\pi} \sum_{k \geq 0} (-1)^k \frac{e^{-(2k+1)^2 \lambda}}{2k+1}.} \tag{C.3}$$

Acknowledgments. We thank the referees for helping us weed out errors large and small.

REFERENCES

[1] M. Aigner and G. M. Ziegler, *Proofs from The Book*, 2nd ed., Springer-Verlag, Berlin, 2001.
[2] V. Anantharam and S. Verdu, *Bits through queues*, IEEE Trans. Inform. Theory, 42 (1996), pp. 4–16.
[3] P. Antsaklis and J. Baillieul, *Special issue on technology of networked control systems*, Proc. IEEE, 95 (2007).
[4] K. J. Åström and B. Bernhardsson, *Comparison of Riemann and Lebesgue sampling for first order stochastic systems*, in Proceedings of the 41st IEEE Conference on Decision and Control, Las Vegas NV, 2002, pp. 2011–2016.
[5] J. S. Baras and A. Bensoussan, *Optimal sensor scheduling in nonlinear filtering of diffusion processes*, SIAM J. Control Optim., 27 (1989), pp. 786–813.
[6] F. Beukers, J. A. C. Kolk, and E. Calabi, *Sums of generalized harmonic series and volumes*, Nieuw Arch. Wiskd., 11 (1993), pp. 217–224.
[7] R. Carmona and S. Dayanik, *Optimal multiple stopping of linear diffusions*, Math. Oper. Res., 33 (2008), pp. 446–460.
[8] *WirelessHART Protocol Specification*, Technical report, HART Communication Foundation, Austin, TX, April 2008.
[9] G. Gabor and Z. Győrfi, *A theory for the practice of waveform coding*, in Recursive Source Coding, Springer-Verlag, New York, 1986.
[10] R. M. Gray, *Source Coding Theory*, Kluwer Academic, Norwell, MA, 1990.
[11] B. Hajek, *Jointly optimal paging and registration for a symmetric random walk*, in Proceedings of the 2002 IEEE Information Theory Workshop, 2002, pp. 20–23.
[12] B. Hajek, K. Mitzel, and S. Yang, *Paging and registration in cellular networks: Jointly optimal policies and an iterative algorithm*, IEEE Trans. Inform. Theory, 54 (2008), pp. 608–622.
[13] T. Henningsson, E. Johannesson, and A. Cervin, *Sporadic event-based control of first-order linear stochastic systems*, Automatica, 44 (2008), pp. 2890–2895.
[14] O. C. Imer and T. Basar, *Optimal control with limited controls*, in Proceedings of the 2006 American Control Conference, 2006, pp. 298–303.
[15] V. Jacobson, *Congestion avoidance and control*, SIGCOMM Computer Comm. Rev., 25 (1995), pp. 157–187.

[16] K. H. Johansson, M. Törngren, and L. Nielsen, *Vehicle applications of controller area network*, in Handbook of Networked and Embedded Control Systems, W. S. Levine and D. Hristu-Varsakelis, eds., Birkhäuser, Basel, 2005.

[17] I. Karatzas and S. E. Shreve, *Brownian Motion and Stochastic Calculus*, 2nd ed., Grad. Texts in Math. 113, Springer-Verlag, New York, 1991.

[18] I. Karatzas and S. E. Shreve, *Methods of Mathematical Finance*, Appl. Math. (N.Y.), 39, Springer-Verlag, New York, 1998.

[19] H. J. Kushner, *On the optimum timing of observations for linear control systems with unknown initial state*, IEEE Trans. Automat. Control, 9 (1964), pp. 144–150.

[20] H. J. Kushner and P. Dupuis, *Numerical Methods for Stochastic Control Problems in Continuous Time*, 2nd ed., Springer-Verlag, New York, 2001.

[21] F.-L. Lian, J. R. Moyne, and D. M. Tilbury, *Network protocols for networked control systems*, in Handbook of Networked and Embedded Control Systems, D. Hristu-Varsakelis and W. S. Levine, eds., Birkhäuser, Basel, 2005, pp. 651–675.

[22] L. Meier III, J. Peschon, and R. M. Dressler, *Optimal control of measurement subsystems*, IEEE Trans. Automat. Control, 12 (1967), pp. 528–536.

[23] S. R. Norsworthy, R. Schreier, and G. C. Temes, eds., *Delta-Sigma Data Converters: Theory, Design, and Simulation*, Wiley-IEEE Press, New York, 1996.

[24] G. Peskir and A. Shiryaev, *Optimal stopping and free-boundary problems*, Lectures Math. ETH Zürich, Birkhäuser, Basel, 2006.

[25] M. Rabi, *Packet Based Inference and Control*, Ph.D. thesis, University of Maryland, College Park, MD, 2006.

[26] M. Rabi and J. S. Baras, *Sampling of diffusion processes for real-time estimation*, in Proceedings of the 43rd IEEE Conference on Decision and Control, Paradise Island, Bahamas, 2004, pp. 4163–4168.

[27] M. Rabi, J. S. Baras, and G. V. Moustakides, *Efficient sampling for keeping track of a gaussian process*, in Proceedings of the 14th Mediterranean Conferences on Control and Automation, 2006.

[28] M. Rabi and K. H. Johansson, *Optimal stopping for updating controls*, in Proccedings of the 2nd International Workshop on Sequential Methods, UTT Troyes, France, 2009.

[29] M. Rabi, K. H. Johansson, and M. Johansson, *Optimal stopping for event-triggered sensing and actuation*, in Proceedings of the 47th IEEE Conference on Decision and Control, Cancun, Mexico, 2008.

[30] M. Rabi, G. V. Moustakides, and J. S. Baras, *Multiple sampling for estimation on a finite horizon*, in Proceedings of the 45th IEEE Conference on Decision and Control, San Diego, CA, 2006, pp. 1351–1357.

[31] E. Skafidas and A. Nerode, *Optimal measurement scheduling in linear quadratic Gaussian control problems*, in Proceedings of the 1998 IEEE International Conference on Control Applications, vol. 2, 1998, pp. 1225–1229.

[32] W. Wu and A. Arapostathis, *Optimal sensor querying: General Markovian and lqg models with controlled observations*, IEEE Trans. Automat. Control, 53 (2008), pp. 1392–1405.