ABSTRACT

Title of dissertation: PACKET BASED INFERENCE AND CONTROL

Maben Rabi, Doctor of Philosophy, 2006

Dissertation directed by: Professor John S. Baras
Department of Electrical and Computer Engineering

Communication constraints in Networked Control systems are frequently limits on data packet rates. To efficiently use the available packet rate budgets, we have to resort to event-triggered packet transmission. We have to sample signal waveforms and transmit packets not at deterministic times but at random times *adapted* to the signals measured. This thesis poses and solves some basic design problems we face in reaching for the extra efficiency.

We start with an estimation problem involving a single sensor. A sensor makes observations of a diffusion process, the state signal, and has to transmit samples of this process to a supervisor which maintains an estimate of the state. The objective of the sensor is to transmit samples strategically to the supervisor to minimize the distortion of the supervisor's estimate while respecting sampling rate constraints. We solve this problem over both finite and infinite horizons when the state is a scalar linear system. We describe the relative performances of the optimal sampling scheme, the best deterministic scheme and of the suboptimal but simple to implement level-triggered sampling scheme. Apart from the utility of finding the op-

timal sampling strategies and their performances, we also learnt of some interesting properties of the level-triggered sampling scheme.

The control problem is harder to solve for the same setting with a single sensor. In the estimation problem for the linear state signal, the estimation error is also a linear diffusion and is reset to zero at sampling times. In the control problem, there is no equivalent to the error signal. We pay attention to an infinite horizon average cost problem where, the sampling strategy is chosen to be level-triggered. We design piece-wise constant controls by translating the problem to one for discrete-time Markov chain formed by the sampled state. Using results on the average cost control of Markov chains, we are able to derive optimality equations and iteratively compute solutions.

The last chapter tackles a binary sequential hypothesis testing problem with two sensors. The special feature of the problem is the ability of each sensor to hear the transmissions of the other towards the supervisor. Each sensor is afforded on transmission of a sample of its likelihood ratio process. We restrict attention to level-triggered sampling. The results of this chapter remind us not to expect improvements in performance merely beccause of switching to event-triggered sampling. Even though the detection problem is posed over an infinite-time horizon, threshold policies dont measure up.

The chief merits of this thesis are the formulation and solution of some basic problems in multi-agent estimation and control. In the problems we have attacked, we have been able to deal with the differences in information patterns at sensors and supervisors. The main demerits are the ignoring of packet losses and of variable delays in packet transmissions. The situation of packet losses can however be handled at the expense of additional computations. To summarize, this thesis provides valuable generalizations of the works of Åström and Bernhardsson [1] and of Kushner [2] on timing of Control actions and of Sampling observations respectively.

PACKET BASED INFERENCE AND CONTROL

by

MABEN RABI

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2006

Advisory Commmittee:

Professor John S. Baras, Chair/Advisor
Professor Armand M. Makowski
Professor Min Wu
Professor Eyad H. Abed
Professor Eric V. Slud, Dean's representative

ACKNOWLEDGMENTS

I am very glad to thank my advisor Professor John S. Baras for his incredible support, patience and mentoring. He has a habit of rephrasing my thoughts lucidly - pointing out the key issues as well as an ability to infect me with his enthusiasm and energy. His generosity with every resource at his disposal has let me learn from him, explore research topics carefully, grow as a research student, switch fields, attend a lot of conferences and meetings and to pursue with confidence and optimism a career I will be making a living out of. Few other advisors could have enabled me to come this far.

Through his work on the estimation problem, as well as by brain-storming about the other problems, Professors George V. Moustakides revived the dissertation at a time when I was stuck. His input to the collaboration was intense, selfless and uplifting. Professors Armand M. Makowski and Vivek S. Borkar provided a lot of useful questions, pathways and reassurance all of which I am thankful for. Professor Eyad H. Abed has been a constant friend. He also organized a symposium which was very productive for me and provided partial financial support for attending it. I also thank Professors P.S. Krishnaprasad and Pamela Abshire for their support and advice. To Professors Eric V. Slud and Min Wu are thanks for agreeing to sit in my committee.

Althia, Diane and Kim have helped me with real friendliness and smiles on a regular basis while helping me with all kinds of paper-work.

To my brother Reuben are thanks for a lifetime of encouragement and inspiration. To my parents and younger siblings are due my gratitude for being there always with emotional support.

My friends Senni, Pedram, Vahid, Sudhir, Vinay, Xiaobo, Punya, Wei, Matt, George P, Tao, Alvaro, Gelareh, Priya, Huigang and Shahan, Kyriakos, and Maria have been patient listeners, counsellors and teachers. I promise to read their theses in full ! To Alokika, Arash, Arni, Arvind, Ayan, Bata, Behnam, Dinesh, George T, Irena, Jayendu, Karthik, Lux, Mohammed, Munther, Nassir, Naveen, Sreeni, Svetlana, Swami, Thomas, Veera, Velu, and Vishwa go my thanks for refreshing my mind with their good company. At the beginning, there were Navin, Ranjith, Mark, Umesh, Manuela, Dmitris, Ksriti, Shung Loong, Steve, Jason, Kripa and Prahaladh to help me.

I am grateful for financial support of my studies and research by: the Army Research Office under Grant No. DAAD 19021 0319, and under the CIP-URI Grant No. DAAD 190110494, the NASA Cooperative Agreement NCC8235 for the Center for Satellite and Hybrid Communication Networks, the National Science Foundation Grant No. ECS 9727805, and by funds from Boeing, BAE, the Lockheed Martin Chair, the State of Maryland, and the University of Maryland College Park. I have also benefitted from two exciting summer internships at ATT Research Labs (2001) and Army Research Laboratory (2000).

TABLE OF CONTENTS

LIST OF FIGURES

Chapter 1

Estimation and Control with Data-rate Constraints

This work proposes some new ideas for the design of sampling, estimation and control schemes in multi-agent architectures where there is a one-way rate constraint on the flow of information from the sensing agents to a central supervisor. The constraint will essentially be a limit on the rate at which data packets can be transmitted from each sensor.

These situations arise frequently from basic limitations on the information exchange pipelines such as costs on the usage of bandwidth, energy and power. Sometimes, there is a limited ability to process all the information that can be gathered. There could also be a task-induced need to minimize communication in this distributed setting because of reasons like 'keeping the voices low' when using a distributed sensor bed for spying. Sometimes, the chief reason is a 'Decentralized' design philosophy that emphasizes autonomous behaviour at the agent-level so that the overall collaborative effort is less susceptible to problems with information exchange or failure of individual agents. The same information-rate constraints appear when autonomous agents are cast into a team and need to communicate using costly resources in order to achieve the mission assigned. These constraints come-up again, in centralized architectures where the measurements have to be digitized - sampled and quantized, in order to reach the decision-maker. Frequently, the digital communication link over which the measurements travel is a data packet network. The **Networked control problem** is to identify the required data-rates for achieving control tasks and given such data-rates, to prescribe measurement-communication, estimation and control policies for good performance. To solve this, we will actual analyze the reverse of the problem where given prescribed rate-constraints, we seek the best achievable performance and the corresponding estimation and control algorithms.

1.1 A collection of motivating examples.

The design of decision and control algorithms when the information input is rate-constrained requires a simultaneous planning of the real-time information transmitted along with the choice of estimation and control policies. There are many systems-control applications where the resolution of this challenge is essential to the functioning of the system. We list some examples in the following paragraphs.

1.1.1 Controller Area Networks (CANs) in Automobiles:

Today's motor vehicle contains on-board computers that monitor and control various operations in the whole system. A single computer handling the Fuel-injection, cruise-control, and various other monitoring and regulation tasks can process only a limited amount of sensor measurements per unit time. This means that the total rate of all information gathered for processing should correspond to this limit. Frequently, a single Ethernet bus can connect most of the sensors. This makes sense because the computer can listen to only one device at a time and having a single bus eliminates a lot of wiring costs. This bus, the so-called Controller Area Network(CAN) while sufficiently bandwidth-endowed for safety and efficiency, is designed to handle some fraction of its maximum traffic rate. The system design must allocate packet rate quotas (average and peak) to various sensors and actuators in a way that ensures satisfactory performance of the various sensing and control tasks. It must also take care of finer but important details such as assigning priorities for packet traffic from various sources on the CAN and guaranteeing on-demand attention for certain sensors (say, the ones for the Anti-lock Braking system). To solve the rate allocation problem, the designer needs, for every task, a Pareto curve that describes the trade-off between task performance and the available packet rates between the relevant sensors and the computer. Then perhaps, a (hierarchical) multi-objective optimization problem can find the best allocations. For each specific sensing, detection or feedback control problem, we need to find out the best technique of packetizing measurements at the rate available.

1.1.2 The CEO problem:

A CEO who makes various decisions based on the information the employees of her firm gather has to determine when she needs to receive intimation of the various goings-on. For instance, while she may want the newly hired System Administrator to keep her posted on the progress in filtering spam (to prevent him from slacking off), she will not want him to report each day that there has been no virus flood on the network thus far. She has to allocate attention (time) to different charges and make sure that the individual attentions dispensed get the jobs done. Some CEOs might want to find out the limits of the multi-tasking she/he could pull-off. All of them would/should have learnt how many sub-CEOs to hire into what hierarchy and how their individual times are going to be spent. We will not have much to say about this specific problem since our own attention will be devoted to continuous-time problems. Still, the questions and answers we care about for diffusion state processes can, at least in principle, be carried over to Discrete Markov chains. A more academic version of these issues is discussed in [3, 4, 5] and [6].

1.1.3 A splurge of sensors:

It is common and sometimes necessary nowadays to perform some sensing tasks by throwing a large number of cheap sensors at the job. They would all use a common medium to communicate with a computer which acts as the supervisor. Two examples are given below:

MEMS arrays and sensors on a bus: If you are going to keep tabs on something like air quality with a very cheap but not very reliable device, it makes a lot of sense to use a horde of them and make a computer listen to all of them. You may also need to use different devices to look for different things in the environment. If many of these can be cheaply fabricated on a single board (a MEMS device array for example), it is infeasible to wire them individually to the computer (No space for the wires and limited data fan-in for the computer). All of these devices and the computer are going to be shouting on the same data bus and the designer needs to lay out a communication strategy for each of the sensors that gives everyone just enough voice. There is also a desire to keep the data-rate low because energy dissipation and the associated rise of temperature in these micro electronic components are costly or even undesirable. This needs to be done in a way that maximizes (detection) performance. Even if the bus is shared on a TDMA-like fixed scheme guaranteeing periodic access to individual devices, we need to determine the best allocation of sampling rates to individual devices.

Wireless sensor cluster: A cluster of wireless sensors all linked to a common hub in one hop presents the same issues as in the case of the MEMS array. To save energy and power and bandwidth, and perhaps to limit RF activity to cover-up a covert sensing operation, the individual nodes should plan to send as few packets as possible to the supervisor. The situation becomes much more interesting if the sensors are laid out not in a small neighbourhood where everybody is within reach of everyone else, but as a spreading mesh network that is several hops wide. The overall collation of information that helps the supervisor arrive at the best possible decision and control choices is very much like the firm hierarchy problem faced by the CEO. There are many facets of the multi-access wireless communication channels that are not factored into estimation and control design right now. With or without the multi-access channel complications, one would like to know how many sensors can provide a prespecified-level of performance and how to organize the information gathering process to achieve good performance. Some related problems are studied in [7, 8, 9].

1.1.4 Collaborative Sensing and Control

In tracking and control tasks carried out by a team of possibly mobile agents, the issue of how much of the local measurements to relay to others is crucial for team success.

Energy efficient Monitoring: Detection and tracking: Take a sensor minefield replete with acoustic sensors and video/IR cameras. The acoustic sensors are cheap and low-energy devices and hence can be used all the time. But the video/IR cameras have better 'SNR'. To detect and track an intruder without wasting too much energy on the expensive sensors or in communicating local scenarios, the designer needs to solve a joint sensor scheduling, sampling and detection problem.

Pursuit games: A flock of mobile agents trying to zero-in on a possibly mobile target that is not exactly visible (at least until getting up close) need to share their views. Having to do this in a packet efficient manner brings us back to the world of rate-constrained estimation and control.

Each sensor has to transmit to a supervisor, at times it chooses, *data packets* that contain condensed information that will be useful for the supervisor to estimate the state at current and future times. At all times, the supervisor computes an estimate (filter) of the current state given the record of packets (contents of the packets including the sampling times) received thus far from the sensor. The strategy used by the sensor to choose the times at which to sample the observation process is known to the supervisor as well. The real-time estimate by the supervisor could be used to compute a certainty-equivalence continuous feedback control signal that can be relayed to the plant without communication constraints. The only constraint on communication rates in this setup is a limit on the rate of packets sent from the sensor to the supervisor. This limit apart, the data packet link from the sensor to the supervisor is to be considered a lossless, zero-delay packet pipeline.

1.2 Packetization of measurements

The digital representation of observations for communication or other purposes introduces a loss of information that decreases estimation and control performance. In recursive estimation and control problems, the *signal to noise ratio* of the received information as well as the *timeliness* of the information are vital for efficient use. Digitization affects both. The sampling introduces periods of virtual information black-out[1] and a coarse quantization introduces more noise than a fine quantization. The effect of digitization on estimation, detection and control performance has been studied by some researchers basically as the effect of measurements made piece-wise constant with the jumps made at sampling times [10, 11, 12, 13]. In this framework, we can describe what the optimal sampling rate and quantization scheme is, or what the minimal rates should be to guarantee stabilization or boundedness of estimation error measures, especially for linear systems.

The transmission of measurements using a packet-based communication scheme has some features that simplify the analysis. In most scenarios, the packets are of

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

uniform size and even when of variable size, have at least a few bytes of header and trailer files. These segments of the packet carry source and destination node addresses, a time stamp at origin, some error control coding, some higher layer (link and transport layers in the terminology of data networks) data blocks and any other bits/bytes that are essential for the functioning of the packet exchange scheme but which nevertheless constitute what is clearly an overhead. The payload or actual measurement information in the packet should then be at least of the same size as these 'bells and whistles'. It costs only negligibly more in terms of network resources, time, or energy to send a payload of five or ten bytes instead of two bits or one byte when the overhead part of the packet is already 5 bytes. In other words, in essentially all packet-based communication schemes, the right unit of communication cost is the cost of transmitting a single packet whether or not the payload is longer by a few bytes. This means that the samples being packetized can be quantized with very fine detail, say with 4 bytes, a rate at which the quantization noise can be ignored for low dimensional variables. The actual effect of this fine quantization could be investigated perhaps along the lines of [14]. For Markov state processes, dealt with in this paper, this will actually mean that all of these bytes of payload can be used to specify the latest state estimate. An example of an information-constrained problem where this argument fails is the TCP-RED congestion control problem where the state information is carried by a single bit in the whole packet in which the real payload is irrelevant to the congestion state.

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

1.3 Sampling strategy: predetermined or adaptive ?

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

A special situation of such a sampling scheme (or rather a control invocation scheme), called Lebesgue sampling, is studied in [15, 1]. A deterministic problem is treated in [16].

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

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

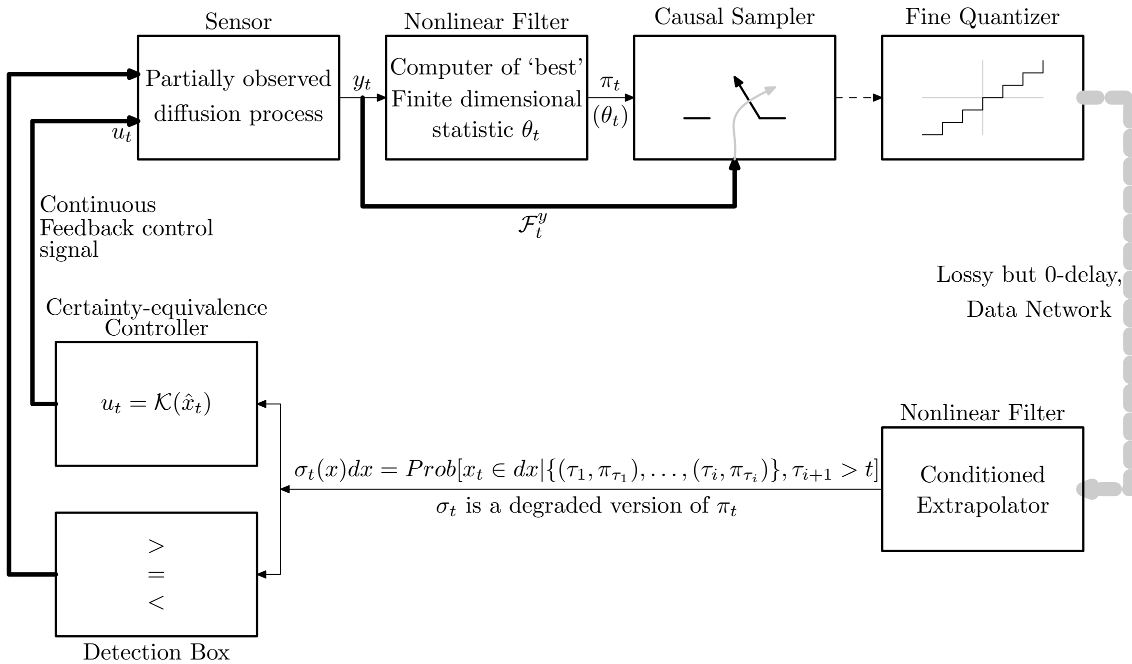

Figure 1.1: Schematic of a general Networked Control and Monitoring system with a single sensor.

1.4 Contributions of this thesis

Control, estimation and detection problems with communication constraints have been usually analysed with a focus on the effects of quantization. In each of the works [10, 11, 17, 18], we have a discrete time linear system (perhaps the result of sampling a continuous time system periodically). The communication constraint is a limit on the number of bits allowed for quantizing sensor transmissions at the discrete time instants. In these works, the task is to keep bounded and perhaps minimize an estimation or control cost in the presence of hard limits on the bit rate.

The work of Åström and Bernhardsson [1] focusses on the sampling problem for a rate-limited impulse control problem. This thesis is an exploration of event-triggered sampling strategies for estimation, control and detection problems.

The study of event-triggered sampling, estimation and control problems involves a choice of sequences of stopping times along with control and estimation waveforms. Results on the joint choice of stopping times and feedback control signals are available in the works [19, 20]. However, we are unable to utilize these rersults because of a key difference between the problems solved in these papers and the ones in this thesis. In this thesis, the problems of joint choice of control and stopping times have a special information pattern. While any stopping (sampling) time has to be adapted to the sensor observations, the feedback control waveform is adapted to the process of samples and sample times of these observations.

Throughout this thesis, packet losses as well as transmission delays will be ignored.

In the next chapter, we will address the problem of real-time estimation on a finite time horizon. We will be able to find optimal sampling strategies when the single sensor has perfect observations. We produce comparisons of the performances of key strategies. We should mention that analogous control problems in finite time can be posed and solved. However, their solutions become computationally much more burdensome. Our solution of multiple stopping problems which arise in this chapter will be solved by using standard solution techniques [21, 22] of optimal single stopping problems and a recursive reduction of the solution of multiple stopping problems to a single stopping problem.

In chapter 3, we will deal with a countably repeated sampling problem over the infinite horizon. This will also be for real-time estimation. The literature on average cost optimal multiple stopping [23] provides the tools necessary for determining the optimal sampling policy. The result we have on the optimality of Lebesgue sampling for scalar linear systems seems to be new. The level-triggered sampling whose performance we describe here is the stochastic analogue of level-triggered sampling for (unknown) deterministic bandlimited signals studied by Lazar and Tóth [24].

In chapter 4, we design controls for countably repeated sampling on the infinite horizon. We see how the control problem is inherently more difficult than the estimation problem even when the signal model is linear. Here, the non-triditional information pattern described at the beginning of this section makee the problem different from that descussed in the literature on swtiching control.

In the last chapter 5, we attack a sequential detection problem with two sensors gathering measurements. The problem is one of sampling asynchronously, likelihood ratios once at each sensor with the sample being heard at the other sensor as well. Although, we are unable to prove the overall optimality of the asynchronous threshold-triggered sampling scheme we study, we are able to compare performances of natural candidates for good performance and/or ease of implementation.

Chapter 2

Finite Sampling for Real-Time Estimation

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

$$dx_t = f(x_t)dt + g(x_t)dW_t, \tag{2.1}$$

$$dy_t = h(x_t)dt + dV_t. \tag{2.2}$$

With $x_0 \sim \pi_0(x)dx$, $y_0 = 0$, $x_t \in \mathbb{R}^n$, $y_t \in \mathbb{R}^m$, $W_t \in \mathbb{R}^n$, $V_t \in \mathbb{R}^m$, W and V being standard, independent Wiener processes, with g being positive definite: $g(x)g(x)^T > 0, \ \forall x \in \mathbb{R}^n$, and with f, g, h and π_0 being such that the conditional probability density of x_t given $\{y_s | 0 \leq s \leq t\}$ exists. The sensor has to transmit to a supervisor, at times it chooses in $[0, T]$, *data packets* that contain condensed information that will be useful for the supervisor to estimate the state at current and future times.

2.1 Sampling by a single sensor

We will first describe in general terms, the problem of optimal adaptive sampling that minimizes a filtering distortion. However, we will revert to a specific version of the problem in order to get concrete solutions.

2.1.1 The sampling problem

The state process x_t is a partially observed diffusion process. Any unnormalized version (ρ_t) of the conditional density of the state given the observations so far (π_t), will obey the Duncan-Mortensen-Zakai SPDE:

$$d\rho_t(x) = \mathcal{L}^*\Big(\rho_t(x)\Big) + h(x)\rho_t(x)dy_t \tag{2.3}$$

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

$$\mathcal{L}^*\phi = -\sum_{i=1}^{n} \frac{\partial(f_i\phi)}{\partial x_i} + \frac{1}{2}\sum_{i=1}^{n}\sum_{j=1}^{n} \frac{\partial^2(g_i g_j \phi)}{\partial x_i \partial x_j}. \tag{2.4}$$

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

$$d\theta_t = \Phi(\theta_t)dt + \Psi(\theta_t)dy_t \tag{2.5}$$

$$\pi_t(x) = \kappa(t, x, \theta_t) \tag{2.6}$$

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

The causal sampling problem with a fixed number of samples is to pick an increasing sequence $\mathcal{T}_N(\{y_s|0 \le s \le T\}) : \{y_s|0 \le s \le T\} \to [0,T]^N$ of N stopping times.

$$\mathcal{T}_N(\{y_s|0 \le s \le T\}) = \{\tau_1, \dots \tau_N\}, \tag{2.7}$$

$$0 \le \tau_1 < \tau_2 < \dots < \tau_{N-1} < \tau_N \le T, \tag{2.8}$$

$$\mathbf{1}_{\{\tau_i > t\}} \in \mathcal{F}_t^y \quad \forall\ i \in \{1, 2, \dots, N\}. \tag{2.9}$$

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

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

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

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

$$\sigma_t(x)dx = \mathbb{P}\Big[x_t \in dx \Big| \big\{(\tau_1, \pi_{\tau_1}), \ldots, (\tau_{i(t)}, \pi_{l(t)})\big\}, \tau_{i(t)+1} > t\Big]. \tag{2.10}$$

If $i(t) = N$,

$$\sigma_t(x)dx = \mathbb{P}\Big[x_t \in dx \Big| \big\{(\tau_1, \pi_{\tau_1}), \ldots, (\tau_N, \pi_{\tau_N})\big\}\Big]. \tag{2.11}$$

Right at the sampling instants, the conditional densities at the supervisor are the same as those at the sensor.

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

π_t is a density-valued Markov process. [25]. σ_t is the best extrapolation of π_t available at the supervisor.

$$\mathbb{P}\left[\pi_t \in d\pi | \left\{\pi_s | 0 \le s \le l(t)\right\}, t < \tau_{i(t)+1}\right]$$
$$= \mathbb{P}\left[\pi_t \in d\pi | \pi_{l(t)}, t < \tau_{i(t)+1}\right] \tag{2.12}$$

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

Filtering distortion: Let $\Delta(\cdot,\cdot)$ be a distance operator (positive and semi-definite binary function) in the space of densities. We could also use a pseudo-distance function namely the Kullback Liebler divergence (we will first have to show that σ_t is absolutely conntinuous with respect to π_t), the L_1 distance, and the square of the Euclidean distance between the means of σ_t and π_t. Corresponding to a chosen distance function, we can set-up a filtering distortion measure at the supervisor end:

$$\mathbb{E}\left[\int_0^T \Delta(\sigma_s,\pi_s)ds\right]. \tag{2.13}$$

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

$$\mathcal{T}_N^*(\{y_s|0\le s\le T\}) = \underset{\mathcal{T}_N(\{y_s|0\le s\le T\})}{\arg\min}\ \mathbb{E}\left[\int_0^T \Delta(\sigma_s,\pi_s)ds\right]. \tag{2.14}$$

The discussion above can be summarized as follows: A causal sampling policy is a multiple stopping time policy. Given such a policy $\mathcal{T}_N(\{y_s|0\le s\le T\})$, the optimal filter at the supervisor is derived from it as the conditional density given by (2.10,2.11). We seek the optimal $\mathcal{T}_N^*(\{y_s|0\le s\le T\})$ as the one that minimizes (2.13). It would save a lot of computational effort if for this optimal sampling strategy, the conditional density can be computed in the fashion of eqn. (2.6) or least as a numerical approximation to something like eqn. (2.3).

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

Variable number of samples: We can easily extend the solution of the fixed packet count problem to a slightly better performing *variable packet count* problem. We describe such an extension in section II.C of [26]. This can be done when the sensor has a lot of computing power at its disposal.

A solution to this joint multiple stopping and filtering problem seems difficult because of the complicated relationship between a stopping policy and the corresponding filter at the supervisor. However, the problem formulation itself is a step forward because it can be solved in special cases and because a natural approximation (which is used in [1] for an infinite time interval problem) still outperforms the periodic sampling strategy. It is a proper generalization of the deterministic sampling problem for linear systems studied by Kushner [2].

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

2.1.2 Scheduling a single packet from an ideal sensor - Decoupling the sampling strategy from the filter

We describe here the optimal schedule of a single sample on $[0, T]$ for the special case of a perfectly observed scalar state process with odd drift and either

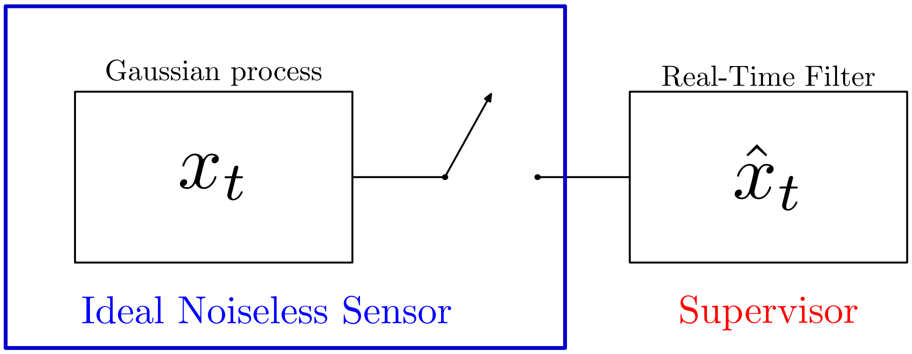

Figure 2.1: The situation of sampling a perfect sensor for real-time estimation based on the sampled stream. The plant here is linear. The discussion in section 2.1.2 covers some nonlinear plants as well.

even or odd diffusion coefficient functions and an even initial probability density function.

$$dx_t = f(x_t)dt + g(x_t)dB_t \tag{2.15}$$

$$dy_t = x_t dt \tag{2.16}$$

With, $x_0 \sim \pi_0(x)dx$, $y_0 = 0$, $x_t \in \mathbb{R}$, f being odd, and g either odd or even.

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

$$\mathcal{T}_1(\{x_s | 0 \le s \le T\}) = \{\tau\}, \tag{2.17}$$

$$0 \le \tau \le T, \tag{2.18}$$

$$\mathbf{1}_{\{\tau > t\}} \in \mathcal{F}_t^x. \tag{2.19}$$

A non-standard optimal stopping problem

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

$$J = \mathbb{E}\left[\int_0^T (\hat{x}_s - x_s)^2 ds\right] \tag{2.20}$$

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

$$\hat{x}_t = \begin{cases} \mathbb{E}[x_t | x_0 = 0, \tau > t] & \text{if} \quad \tau > t \\ \\ \mathbb{E}[x_t | x_\tau] & \text{if} \quad \tau \leq t \end{cases} \tag{2.21}$$

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

$$\mathbb{E}\left[\int_0^\tau (\hat{x}_s - x_s)^2 ds\right] + \mathbb{E}\left[\int_\tau^T (\hat{x}_s - x_s)^2 ds\right]. \tag{2.22}$$

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

$$\frac{dP_t}{dt} = \mathbb{E}\left[2x_t f(x_t) + g^2(x_t) \middle| x_\tau\right] dt \tag{2.23}$$

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

$$\mathbb{E}\left[\int_0^\tau (\hat{x}_s - x_s)^2 ds\right] + \mathbb{E}\left[\mathbb{E}\left[\int_\tau^T (\hat{x}_s - x_s)^2 ds \middle| \tau, x_\tau\right]\right]$$
$$= \mathbb{E}\left[\int_0^\tau (\hat{x}_s - x_s)^2 ds\right] + \mathbb{E}\left[\int_\tau^T C(\tau, s, x_\tau) ds\right]$$

Now, let the cost to go from τ be

$$\int_{\tau}^{T} C(\tau, s, x_\tau) ds = \mathscr{C}(\tau, T, x_\tau). \tag{2.24}$$

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

$$J = \mathbb{E}\left[\int_0^{\tau} (\hat{x}_s - x_s)^2 ds + \mathscr{C}(\tau, T, x_\tau)\right] \tag{2.25}$$

is minimized. For an optimal sampling strategy, if we can somehow know the dependence of $\hat{x}_t$ on t for $t \in [0, \tau)$, we can use the *Snell envelope* (S_t) (see [27] Appendix D) to determine the optimal stopping rule.

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

Then, the smallest time τ^* when the cost of stopping at that time hits the Snell envelope is an optimal stopping time (see [27] Appendix D).

$$\int_0^{\tau^*} (\hat{x}_s - x_s)^2 ds + \mathscr{C}(\tau^*, T, x_{\tau^*}) = S_{\tau^*}. \tag{2.26}$$

Or equivalently,

$$\mathscr{C}(\tau^*, T, x_{\tau^*}) = \underset{\tau \geq \tau^*}{\text{essup}} \quad \mathbb{E}\left[\int_{\tau^*}^{\tau} (\hat{x}_s - x_s)^2 ds + \mathscr{C}(\tau, T, x_\tau) \Big| x_t\right]. \tag{2.27}$$

Since the Snell envelope depends only on the current value of the state and the current time, we get a simple threshold solution for our problem. We can compute the condition to be satisfied by x_t, t for stopping at t by relating this problem to a variational inequality that gives us continuation and stopping regions. In any case, for numerical computation of the solution, we will have to take that route [28]. Now, we will use some properties of the state process that result from our earlier assumptions.

The unobserved x_t is a process with an even density function at all times if the initial density function is even. Basically, the FP operator (2.4) is linear and so, if we split ρ_t into its even and odd parts

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

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

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

Optimization over arbitrary estimate waveforms at the supervisor

The joint optimization problem of filtering and sampling has been cast so far as a non-standard optimal stopping problem with the filter $\hat{x}_t$ being a funtional of the stopping rule being optimized. Now, we will look at this optimization (2.25) as one over different surrogate waveforms ξ_t that the supervisor could use *up to the stopping time*:

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

For example, the supervisor may want to use a piece-wise linear waveform to keep track of x_t until the sampling time but use the least squares estimate $\hat{x}_t$ after the sample has been received.

By arbitrarily using a ξ-waveform instead of $\hat{x}_t$, the supervisor disregards the stopping policy used at the sensor. All the sensor can do now is to tailor its stopping policy $\mathcal{T}_1$ to minimize the aggregate distortion between the state and the supervisor's estimate process:

$$\mathbb{E}\left[\int_0^T \left(\mathbf{1}_{\{\tau>s\}}\xi_s + \mathbf{1}_{\{\tau\le s\}}\mathbb{E}\left[x_t | \tau, x_\tau\right] - x_s\right)^2 ds\right].$$

This cost can be expressed as:

$$J_{\text{TOTAL}}(\xi_t, \mathcal{T}_1) = \mathbb{E}\left[\int_0^\tau (\xi_s - x_s)^2 ds + \mathscr{C}(\tau, T, x_\tau)\right]. \tag{2.28}$$

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

$$\hat{J}_{\text{TOTAL}}(\xi_t) = J_{\text{TOTAL}}\Big(\xi_t, \mathcal{T}_1^*\big(\{\xi\}\big)\Big) = \min_{\mathcal{T}_1} J_{\text{TOTAL}}(\xi_t, \mathcal{T}_1). \tag{2.29}$$

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

$$\hat{J}_{\text{TOTAL}}(\xi_t^*) = \min_{\xi_t} \hat{J}_{\text{TOTAL}}(\xi_t). \tag{2.30}$$

This means that the pair

$$\Big(\xi_t^*, \mathcal{T}_1^*\big(\{\xi^*\}\big)\Big)$$

minimizes (in sequence) the nested optimization problem:

$$\min_{\xi_t}\left(\min_{\mathcal{T}_1}\left\{\mathbb{E}\left[\int_0^\tau (\xi_s - x_s)^2 ds + \mathscr{C}(\tau, T, x_\tau)\right]\right\}\right) \tag{2.31}$$

The optimal waveform ξ_t^* has the property that it is also the least squares estimate (conditional mean of (2.21)) corresponding to the stopping policy $\mathcal{T}_1^*\big(\{\xi^*\}\big)$:

$$\xi_t^* \overset{\text{a.s.}}{=} \mathbb{E}\left[x_t | \tau > t, \mathcal{T}_1^*\big(\{\xi^*\}\big) \rightsquigarrow \tau\right].$$

If not, we could achieve lower cost by retaining the sampling policy $\mathcal{T}_1^*\big(\{\xi^*\}\big)$ and using the conditional mean (of eqn. (2.21)) it generates.

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

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

This is indeed the conditional mean for the corresponding optimal stopping problem because its Snell's envelope S_t depends only on $|x_t|$ and t. In essence, there is no *Timing Channel* between the optimal filter and the optimal stopping policy. We should remember that, although the conditional mean $\hat{x}_t$ under optimal sampling at the supervisor is the same as the mean of the density from the Fokker-Planck equation(FP-ρ), the conditional variance at the supervisor is smaller than that of FP-ρ.

2.1.3 Keeping track of a scalar Ornstein-Uhlenbeck process

For simplicity of exposition, here we consider the signal to be a scalar Ornstein-Uhlenbeck process.

We are interested in keeping track of the state x_t on a prescribed time interval $[0, T]$.

$$dx_t = -ax_t dt + dw_t$$

where w_t is a standard Wiener process and x_0 is a *zero mean* random variable with pdf $f(x_0)$. Positive values of a give rise to a stable process, negative to an unstable and finally $a = 0$ to the Wiener process.

If $\hat{x}_t$ is the estimate, we measure its quality by the following average integral squared error

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

The estimate $\hat{x}_t$ relies on knowledge about x_t acquired during the time interval $[0, T]$. The type of information we are interested in, are *samples* obtained by sampling x_t at k time instances $0 \leq \tau_1 \leq \tau_2 \leq \cdots \leq \tau_k \leq T$. If we use the minimum mean square error estimate given by

$$\hat{x}_t = \begin{cases} 0 & \text{if } \ t \in [0, \tau_1), \\ x_{\tau_n} e^{-a(t-\tau_n)} & \text{if } \ t \in [\tau_n, \tau_{n+1}), \end{cases}$$

the performance measure becomes

$$\mathcal{J}(\tau_1, \ldots, \tau_k) = \mathbb{E}\left[\int_0^{\tau_1} x_t^2 \, dt + \sum_{n=1}^{k-1} \int_{\tau_n}^{\tau_{n+1}} (x_t - \hat{x}_t)^2 \, dt + \int_{\tau_k}^{T} (x_t - \hat{x}_t)^2 \, dt\right]. \quad (2.32)$$

The goal here is to find sampling policies that are optimal in the sense that they solve the following optimization problem:

$$\inf_{\tau_1, \ldots, \tau_k} \mathcal{J}(\tau_1, \ldots, \tau_k).$$

For the remainder of this paper, and in order to clarify the concepts and computations involved, we treat the single sample case. The multiple sample case just described will be treated in sections (2.3,2.4).

2.2 The single sample case

Let us limit ourselves to the single sample case where, for simplicity, we drop the subscript from the unique sampling instance τ_1. Material in this section was publishe in [29]. In this special case the performance measure in (2.32) takes the form

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

Now notice that the second term can be written as follows

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

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

$$\begin{aligned}\mathcal{J}(\tau) &= \mathbb{E}\left[\int_0^T x_t^2 \, dt - \int_\tau^T (\hat{x}_t)^2 \, dt\right] \\ &= \frac{e^{-2aT} - 1 + 2aT}{4a^2} + \mathbb{E}\left[x_0^2\frac{1 - e^{-2aT}}{2a} - x_\tau^2\frac{1 - e^{-2a(T-\tau)}}{2a}\right] \\ &= T^2\left\{\frac{e^{-2aT} - 1 + 2aT}{4(aT)^2}\mathbb{E}\left[\frac{x_0^2}{T}\frac{1 - e^{-2aT}}{2(aT)} - \frac{x_\tau^2}{T}\frac{1 - e^{-2(aT)(1-\tau/T)}}{2(aT)}\right]\right\} \\ &= T^2\left\{\frac{e^{-2\bar{a}} - 1 + 2\bar{a}}{4\bar{a}^2} + \mathbb{E}\left[\bar{x}_0^2\frac{1 - e^{-2\bar{a}}}{2\bar{a}} - \bar{x}_{\bar{\tau}}^2\frac{1 - e^{-2\bar{a}(1-\bar{\tau})}}{2\bar{a}}\right]\right\},\end{aligned}$$

where,

$$\bar{t} = \frac{t}{T}; \quad \bar{a} = aT; \quad \bar{x}_{\bar{t}} = \frac{x_{\frac{t}{T}}}{\sqrt{T}}. \tag{2.33}$$

It is interesting to note that

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

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 (2.33). The performance measure we are finally considering is

$$\begin{aligned}\mathcal{J}(\tau) = {} & \frac{e^{-2a} - 1 + 2a}{4a^2} \\ & + \mathbb{E}\left[x_0^2 \frac{1 - e^{-2a}}{2a} - x_\tau^2 \frac{1 - e^{-2a(1-\tau)}}{2a}\right]; \quad \tau \in [0, 1].\end{aligned} \tag{2.34}$$

We will also need the following expression

$$\begin{aligned}\mathcal{J}(\tau, x_0) = {} & \frac{e^{-2a} - 1 + 2a}{4a^2} + x_0^2 \frac{1 - e^{-2a}}{2a} \\ & - \mathbb{E}\left[x_\tau^2 \frac{1 - e^{-2a(1-\tau)}}{2a} \middle| x_0\right]; \quad \tau \in [0, 1].\end{aligned} \tag{2.35}$$

Clearly, $\mathcal{J}(\tau) = \mathbb{E}[\mathcal{J}(\tau, x_0)]$, where the last expectation is with respect to the statistics of the initial condition x_0.

Next we are going to consider three different classes of admissible sampling strategies and we will attempt to find the optimum within each class that minimizes the performance measure in (2.34). The classes we are interested in are: a) deterministic sampling; b) threshold sampling and c) general event-triggered sampling. Our results in this problem have appeared in [29].

2.2.1 Optimum deterministic sampling

Let us first minimize (2.34) over the class of deterministic sampling times i.e. open loop, predetermined sampling times. The performance measure then takes the form

$$\mathcal{J}(\tau) = \frac{e^{-2a} - 1 + 2a}{4a^2} + \sigma^2 \frac{1 - e^{-2a}}{2a} - \frac{1}{4a^2}\left\{1 - (1 - 2a\sigma^2)e^{-2a\tau}\right\}\left\{1 - e^{-2a}e^{2a\tau}\right\}; \quad \tau \in [0, 1] \quad (2.36)$$

where σ^2 denotes the variance of the initial condition. Clearly $\mathcal{J}(\tau)$ is minimized when we maximize the last term in the last expression. It is a simple exercise to verify that the optimum sampling time satisfies

$$\tau_o = \arg\max_{\tau} \left\{1 - (1 - 2a\sigma^2)e^{-2a\tau}\right\}\left\{1 - e^{-2a}e^{2a\tau}\right\} = \begin{cases} \frac{1}{2} + \frac{\log(1-2a\sigma^2)}{4a} & \text{for } \sigma^2 \leq \frac{1-e^{-2a}}{2a}, \\ 0 & \text{otherwise.} \end{cases} \quad (2.37)$$

In other words, if the initial variance is greater than the value $(1 - e^{-2a})/2a$ then it is better to sample at the beginning. The corresponding optimum performance becomes

$$\mathcal{J}(\tau_o) = \frac{e^{-2a} - 1 + 2a}{4a^2} - \frac{1}{4a^2}(e^{-a} - \sqrt{1 - 2a\sigma^2})^2 \mathbf{1}_{\sigma^2 \leq \frac{1-e^{-2a}}{2a}}. \quad (2.38)$$

2.2.2 Optimum threshold sampling

Here we consider a threshold η and we sample the process x_t whenever $|x_t|$ exceeds η for the first time. If we call τ_η the sampling instance

$$\tau_\eta = \inf_{0 \leq t}\{t : |x_t| \geq \eta\}.$$

then it is clear that we can have $\tau_\eta > 1$. We therefore define our sampling time as the minimum of the two, that is, $\tau = \min\{\tau_\eta, 1\}$. Of course sampling at time $\tau = 1$, has absolutely no importance since from (2.34) we can see that such a sampling produces no contribution in the performance measure. Another important detail in threshold sampling is the fact that whenever $|x_0| \geq \eta$ then we must sample at the beginning.

Our goal here is, for a given parameter a and pdf $f(x_0)$ to find the threshold η that will minimize the performance measure $\mathcal{J}(\tau)$. As in the previous case let us analyze $\mathcal{J}(\tau)$. We first need to compute $\mathcal{J}(\tau, x_0)$ for given threshold η. From (2.35) we have

$$\mathcal{J}(\tau, x_0) = \frac{e^{-2a} - 1 + 2a}{4a^2} + \left\{ x_0^2 \frac{1 - e^{-2a}}{2a} - \eta^2 \mathbb{E}\left[\frac{1 - e^{-2a(1-\tau)}}{2a} \Big| x_0 \right] \right\} \mathbf{1}_{|x_0| < \eta}. \quad (2.39)$$

We first note that our expression captures the fact that we sample in the beginning whenever $|x_0| \geq \eta$. Whenever this does not happen, that is, on the event $\{|x_0| < \eta\}$ we apply our threshold sampling. If $|x_t|$ reaches the threshold η before the limit time 1, then we sample and $x_\tau = \pm\eta$, therefore

$$x_\tau^2 \frac{1 - e^{-2a(1-\tau)}}{2a} = \eta^2 \frac{1 - e^{-2a(1-\tau)}}{2a}.$$

If however $|x_t|$ does not reach the threshold before time 1, then we sample at $t = 1$ and we have

$$x_\tau^2 \frac{1 - e^{-2a(1-\tau)}}{2a}\Big|_{\tau=1} = 0 = \eta^2 \frac{1 - e^{-2a(1-\tau)}}{2a}\Big|_{\tau=1},$$

Manipulating the last term in (2.39) we obtain

$$\begin{aligned}\mathcal{J}(\tau, x_0) = {} & \frac{e^{-2a} - 1 + 2a}{4a^2} \\ & - (\eta^2 - x_0^2)\left[\frac{1 - e^{-2a}}{2a}\right]\mathbf{1}_{|x_0|<\eta} \\ & + \eta^2 e^{-2a}\mathbb{E}\left[\int_0^\tau e^{2at}\,dt\Big|x_0\right]\mathbf{1}_{|x_0|<\eta}.\end{aligned}$$

The only term that needs special attention in the previous formula is the last one, for which we must find a computational recipe. Consider a function $U(x, t)$ defined on the orthogonal region $|x| \leq \eta,\ 0 \leq t \leq 1$. We require $U(x, t)$ to satisfy the following PDE and boundary conditions

$$\frac{1}{2}U_{xx} - axU_x + U_t + e^{2at} = 0; \quad U(\pm\eta, t) = U(x, 1) = 0. \tag{2.40}$$

If we apply standard Itô calculus on $U(x_t, t)$ we have

$$\begin{aligned}\mathbb{E}[U(x_\tau, \tau)|x_0] - U(x_0, 0) &= \mathbb{E}\left[\int_0^\tau dU(x_t, t)\Big|x_0\right] \\ &= \mathbb{E}\left[\int_0^\tau \left\{\frac{1}{2}U_{xx} - axU_x + U_t\right\} dt\Big|x_0\right] \\ &= -\mathbb{E}\left[\int_0^\tau e^{2at}dt\Big|x_0\right].\end{aligned}$$

Notice that at the time of sampling, x_τ is either at the boundary $x_\tau = \pm\eta$ in which case $U(x_\tau, \tau) = U(\pm\eta, \tau) = 0$, or we have reached the limit $t = 1$ with $|x_1| < \eta$, thus we sample at $\tau = 1$ which yields $U(x_\tau, \tau) = U(x_1, 1) = 0$. We thus conclude that $\mathbb{E}[\int_0^\tau e^{2at}dt|x_0] = U(x_0, 0)$.

With the help of the function $U(x_0, 0)$ we can write $\mathcal{J}(\tau, x_0)$ as

$$\begin{aligned}\mathcal{J}(\tau, x_0) = {} & \frac{e^{-2a} - 1 + 2a}{4a^2} \\ & - \left\{(\eta^2 - x_0^2)\frac{1 - e^{-2a}}{2a} + \eta^2 e^{-2a}U(x_0, 0)\right\}\mathbf{1}_{|x_0|<\eta}.\end{aligned}$$

Averaging this over x_0 yields the following performance measure

$$\mathcal{J}(\tau) = \frac{e^{-2a} - 1 + 2a}{4a^2} - \frac{1 - e^{-2a}}{2a}\mathbb{E}\left[(\eta^2 - x_0^2)\mathbf{1}_{|x_0|<\eta}\right] - \eta^2 e^{-2a}\mathbb{E}\left[U(x_0, 0)\mathbf{1}_{|x_0|<\eta}\right].$$

To find the optimum threshold and the corresponding optimum performance we need to minimize $\mathcal{J}$ over η. This optimization can be performed numerically as follows: for every η we compute $U(x_0, 0)$ by solving the PDE in (2.40); then we perform the averaging over x_0; we then compute the performance measure for different values of η and select the one that yields the minimum $\mathcal{J}(\tau)$.

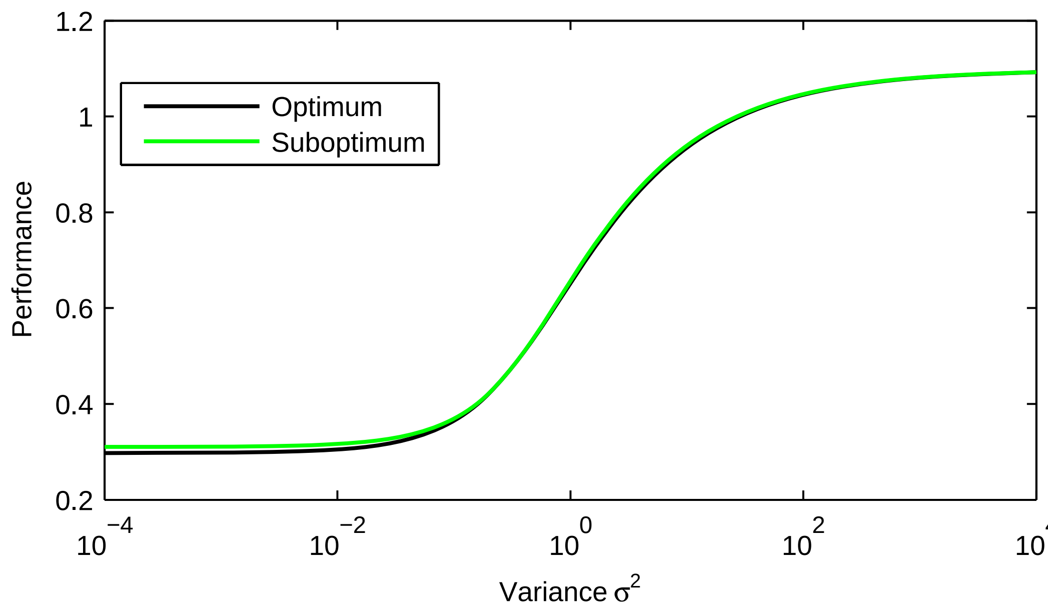

Figure 2.2: Relative performance of optimum (variable) threshold and suboptimum constant threshold sampling scheme, as a function of the variance (σ^2) of the initial condition.

In order to observe certain key properties of the optimum thresholding scheme

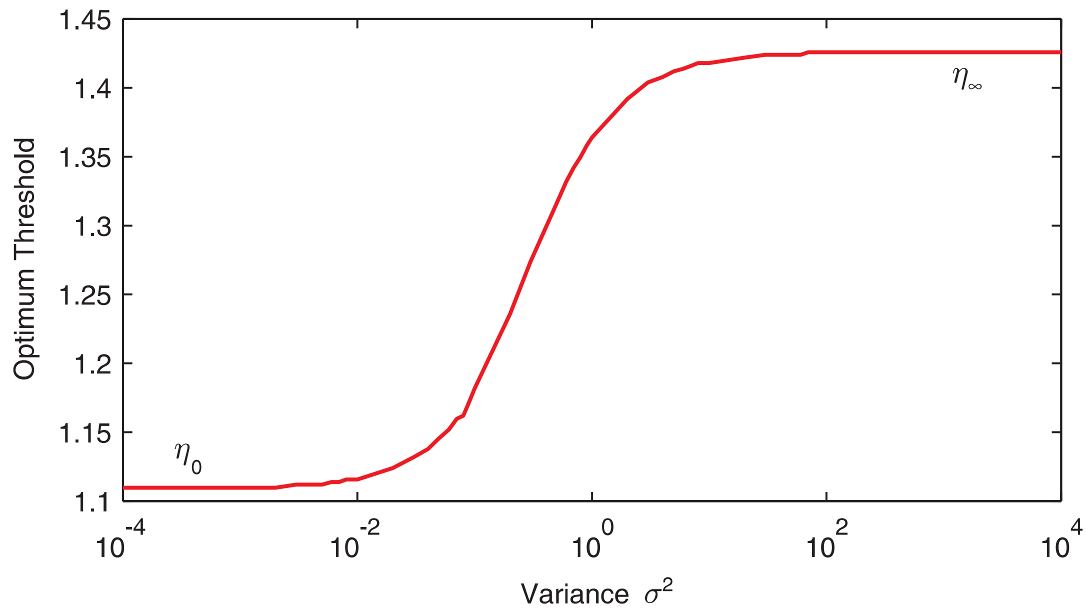

Figure 2.3: Optimum threshold as a function of the initial variance σ^2, with $a = 1$.

let us consider the case $a = 1$ with a zero mean Gaussian initial value x_0 of variance σ^2. Fig. 2.2(a) depicts the optimum performance $\mathcal{J}(\tau)$ as a function of the variance σ^2 and Fig. 2.3(b) the corresponding optimum threshold η. From Fig. 2.3(b) we observe that the optimum threshold is between two limiting values η_0, η_∞. The interesting point is that both these values are *independent* of the actual density function $f(x_0)$, as long as the pdf is from an *unimodal* family of the form: $f(x) = h(x/\sigma)/\sigma,\ \sigma \geq 0$ where, $h(\cdot)$ is a unimodal pdf with unit variance and with both its mean and mode being zero. Indeed for such a pdf, variance tending to 0, means that the density $f(x_0)$ tends to a Dirac delta function at zero. The performance measure in (2.2.2) then takes the simple form

$$\mathcal{J}(\tau) = \frac{e^{-2a} - 1 + 2a}{4a^2} - \eta^2 \left\{ \frac{1 - e^{-2a}}{2a} + e^{-2a} U(0,0) \right\}$$

which, if minimized with respect to η, yields η_0. If now we let the variance $\sigma^2 \to \infty$

then every unimodal function becomes almost flat with value $f(0)$ inside each finite interval $[-\eta, \eta]$. The corresponding performance measure then takes the form

$$\mathcal{J}(\tau) \approx \frac{e^{-2a} - 1 + 2a}{4a^2} - f(0) \int_{-\eta}^{\eta} \frac{1 - e^{-2a}}{2a} (\eta^2 - x_0^2) \, dx_0 - f(0) \int_{-\eta}^{\eta} \eta^2 e^{-2a} U(x_0, 0) \, dx_0.$$

To optimize the previous expression it is sufficient to optimize the last integral, which is independent of the actual pdf $f(x_0)$. This optimization will yield η_∞.

Threshold sampling has another interesting property. If instead of using the optimal threshold η which is a function of the initial pdf and the variance σ^2, we use the *constant* threshold $\eta_o = 0.5(\eta_0 + \eta_\infty)$, then the resulting sampling policy is clearly suboptimal. However as we can see from Fig. 2.2 the performance of the suboptimal scheme is practically indistinguishable from that of the optimal. Having a sampling scheme which is (nearly) optimal for a large variety of pdfs (unimodal functions) and practically any variance value, is definitely a very desirable characteristic. We would like to stress that this property breaks when $f(x_0)$ is not unimodal and also when a takes upon large negative values (i.e. the process is strongly unstable).

2.2.3 Optimal sampling

In this section we are interested in sampling strategies that are optimal in the sense that they minimize the performance measure (2.34) among *all* possible sampling policies (stopping times) τ. Unlike the previous sampling scheme, the optimal sampling rule is completely *independent* of the pdf $f(x_0)$. From (2.34) it is clear that in order to minimize the cost $\mathcal{J}(\tau)$ it is sufficient to perform the following maximization

$$V(\tau) = \sup_{0\le\tau\le 1} \mathbb{E}\left[x_\tau^2 \frac{1-e^{-2a(1-\tau)}}{2a}\right]. \tag{2.41}$$

Using standard optimal stopping theory [21] let us define the optimum cost to go (Snell envelope) as follows

$$V_t(x) = \sup_{t\le\tau\le 1} \mathbb{E}\left[x_\tau^2 \frac{1-e^{-2a(1-\tau)}}{2a}\Big| x_t = x\right]. \tag{2.42}$$

If one has the function $V_t(x)$ then it is straightforward to find the optimal sampling policy. Unfortunately this function is usually very difficult to obtain analytically, we therefore resort to numerical approaches. By discretizing time with step $\delta = 1/N$, we define a sequence of (conditionally with respect to x_0) Gaussian random variables $x_1, \ldots, x_N$, that satisfy the AR(1) model

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

As it is indicated, w_n are i.i.d. Gaussian random variables.

Sampling in discrete time means selecting a sample x_ν from the set of $N+1$ sequentially available random variable $x_0, \ldots, x_N$, with the help of a stopping time $\nu \in \{0, 1, \ldots, N\}$. As in (2.42) we can define the optimum cost to go which can be analyzed as described below. For $n = N, N-1, \ldots, 0$,

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

Equation (2.43) provides a (backward) recurrence relation for the computation of the cost function $V_n(x)$. Notice that for values of x for which the l.h.s. in (2.43) exceeds the r.h.s. we stop and sample, otherwise we continue to the next time instant. We can prove by induction that the optimal policy is a *time-varying threshold* one. Specifically for every time n there exists a threshold λ_n such 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 $V_n(x)$ 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. Fig. 2.4 depicts examples of threshold time evolution for values of the parameter $a = -1, 0, 1$.

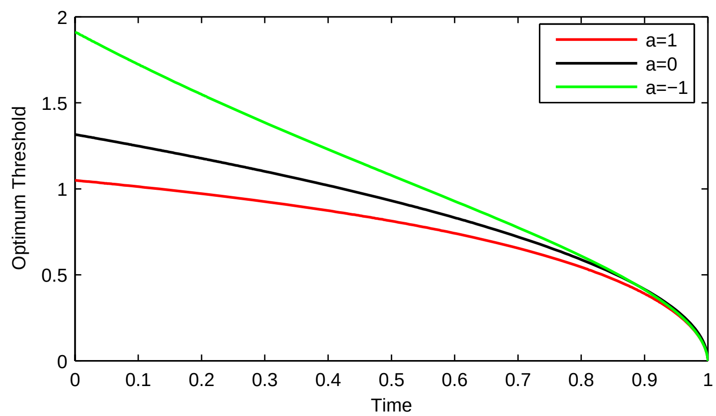

Figure 2.4: Time evolution of the optimum threshold λ_t for parameter values $a = 1, 0, -1$.

Using $V_n(x)$ the final optimum cost can be computed from (2.34) as

$$\mathcal{J}(\tau) = \frac{e^{-2a} - 1 + 2a}{4a^2} - \mathbb{E}\left[V_0(x_0) - x_0^2 \frac{1 - e^{-2a}}{2a}\right].$$

Since from the recursion we know that $V_0(x_0) = x_0^2(1 - e^{-2a})/2a$ for $|x_0| \geq \lambda_0$, we conclude that we can also write

$$\mathcal{J}(\tau) = \frac{e^{-2a} - 1 + 2a}{4a^2} - \mathbb{E}\left[\left\{V_0(x_0) - x_0^2 \frac{1 - e^{-2a}}{2a}\right\} \mathbf{1}_{|x_0| \le \lambda_0}\right]. \tag{2.43}$$

The Wiener case

Let us now focus on the case $a = 0$ which gives rise to a Wiener process. We consider this special case because it is possible to obtain an analytic solution for the optimization problem. For $a = 0$ the optimization in (2.41) takes the form

$$V(\tau) = \sup_{0 \le \tau \le 1} \mathbb{E}\left[x_\tau^2 (1 - \tau)\right].$$

Consider the following function of t and x

$$\mathcal{V}_t(x) = A\left\{\frac{1}{2}(1-t)^2 + x^2(1-t) + \frac{x^4}{6}\right\} \tag{2.44}$$

where $A = \sqrt{3}/(1+\sqrt{3})$. Using standard Itô calculus, if x_t is a standard Wiener process, we can show that

$$\mathbb{E}[\mathcal{V}_\tau(x_\tau)|x_0] - \mathcal{V}_0(x_0) = \mathbb{E}\left[\int_0^\tau d\mathcal{V}_t(x_t)|x_0\right] = 0 \tag{2.45}$$

for any stopping time τ. Notice now that

$$\mathcal{V}_t(x) - x^2(1-t) = A\left(\frac{x^2}{\sqrt{6}} - \frac{1-t}{\sqrt{2}}\right)^2 \ge 0. \tag{2.46}$$

Combining (2.45) and (2.46) we conclude that for any stopping time τ

$$\mathcal{V}_0(x_0) = \mathbb{E}[\mathcal{V}_\tau(x_\tau)|x_0] \ge \mathbb{E}[x_\tau^2(1-\tau)|x_0].$$

This relation suggests that the performance of any stopping time τ is upper bounded by $\mathcal{V}_0(x_0)$. Consequently if we can find a stopping time with performance equal to this value then it will be optimal. In fact such a stopping time exists. From the previous relation the last inequality becomes an equality if at the time of sampling τ we have $\mathcal{V}_\tau(x_\tau) = x_\tau^2(1-\tau)$. From (2.46) we conclude that this can happen iff $|x_\tau|$ is such that the rhs in (2.46) is exactly 0 which happens if $x_\tau^2/\sqrt{6} = (1-\tau)/\sqrt{2}$. This suggests that the optimal threshold for the Wiener process is the following function of time

$$\lambda_t = \sqrt[4]{3}\sqrt{1-t}.$$

The optimum performance measure, from (2.43) and letting $a \to 0$, becomes

$$\mathcal{J}(\tau) = \frac{1}{2} - \mathbb{E}\left[\left\{\mathcal{V}_0(x_0) - x_0^2\right\} \mathbf{1}_{|x_0| \le \lambda_0}\right],$$

where $\mathcal{V}_t(x)$ is defined in (2.44).

2.2.4 Comparisons

We have seen that the best sampling strategy is an *event-triggered* one. Below, we will see graphically that a simpler event-triggered strategy based on a constant threshold, provides almost as good performance compared to the time-triggered one, thus providing more ammunition to the ideas of [1] Let us now compare the performance of the three sampling schemes (deterministic, constant thresholding and optimal) for values of the parameter $a = 10, 1, 0, -1$. Regarding threshold sampling we apply the suboptimal version, which uses a constant threshold. For the pdf of the initial value x_0 we assume zero mean Gaussian with variance σ^2 ranging from 10^{-4} to 10^4. We depict the relative performances of the three schemes with the graphs being normalized so that maximum is 1. In (a),(b) where a is positive (stable

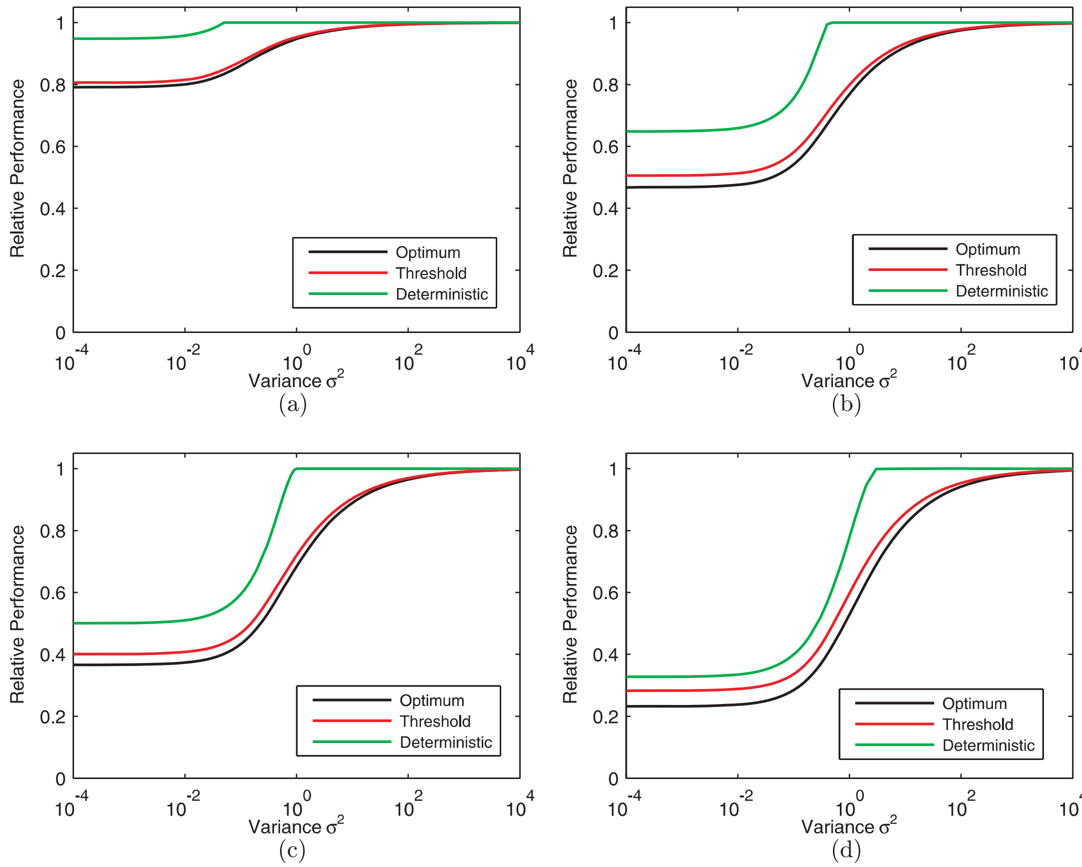

Figure 2.5: Relative performance of Optimum, Threshold and Deterministic samplers as a function of initial variance σ^2 and parameter values (a) $a = 10$, (b) $a = 1$, (c) $a = 0$ and (d) $a = -1$.

process) the performance of the threshold policy is very close to the optimal and the gain, compared to deterministic sampling, is more important. When however we go to values of a that give rise to unstable processes, threshold sampling starts diverging from the optimal, as in (c) and (d) and, although not shown here, when a is less than -5 (strongly unstable process) deterministic sampling can even perform better than threshold sampling.

In the rest of this chapter, we will address multiple sampling policies on a finite interval. Henceforth, we will assume that the initial value of the state is known exactly to the supervisor.

2.3 Multiple samples for a Wiener process

We will now characterize the performance of the three sampling strategies for the Wiener process when the allowed number of samples is more than one. The material in the and the next section has been published in [30].

2.3.1 Deterministic sampling

We will show through induction that uniform sampling on the interval $[0, T]$ is the optimal deterministic choice of N samples

$$\{\tau_1, \tau_2, \ldots, \tau_N \mid 0 \leq \tau_i \leq T, \quad \tau_i \geq \tau_{i-1} \text{ for } i = 1, 2, \ldots, N\}$$

given that the initial value of the signal is zero.

When the number of samples permitted is N, the distortion takes the form:

$$J_{[0,T]}\left(\{\tau_1, \tau_2, \ldots \tau_N\}\right) = \int_0^{\tau_1} \mathbb{E}\left(x_s - \hat{x}_s\right)^2 ds \\ + \int_{\tau_1}^{\tau_2} \mathbb{E}\left(x_s - \hat{x}_s\right)^2 ds + \ldots + \int_{\tau_N}^{T} \mathbb{E}\left(x_s - \hat{x}_s\right)^2 ds.$$

For the induction step, we assume that the optimal choice of $N-1$ deterministic samples over $[T_1, T_2]$ is the uniform one:

$$\{\tau_1, \tau_2, \ldots \tau_{N-1}\} = \left\{T_1 + i\frac{T_2 - T_1}{N} \middle| i = 1, 2, \ldots, N-1\right\},$$

and then the corresponding minimum distortion becomes:

$$N\frac{\left(T_2 - T_1\right)^2}{2N^2} = \frac{\left(T_2 - T_1\right)^2}{2N}.$$

Then, the minimum distortion over the set of N sampling times is:

$$
\begin{aligned}
&\min_{\{\tau_1,\tau_2,\ldots\tau_N\}} J_{[0,T]}\left(\{\tau_1,\tau_2,\ldots\tau_N\}\right) \\
&= \min_{\tau_1}\left\{\int_0^{\tau_1}(x_s-\hat{x}_s)^2 ds + \min_{\{\tau_2,\tau_2,\ldots\tau_N\}} J_{[\tau_1,T]}\left(\{\tau_1,\tau_2,\ldots\tau_N\}\right)\right\}, \\
&= \min_{\tau_1}\left\{\frac{{\tau_1}^2}{2}+\frac{(T-\tau_1)^2}{2N}\right\}, \\
&= \min_{\tau_1}\left\{\frac{N\tau_1^2+\tau_1^2-2\tau_2 T+T^2}{2N}\right\}, \\
&= \min_{\tau_1}\left\{\frac{(N+1)(\tau_1-T/(N+1))^2+T^2-T^2/(N+1)}{2N}\right\}, \\
&= \frac{T^2}{2(N+1)},
\end{aligned}
$$

the minimum being achieved for $\tau_1 = T/(N+1)$. This proves the assertion about the optimality of uniform sampling.

2.3.2 Level triggered sampling

Here, the sampling times are defined through: For $i = 1, 2, \ldots, N$

$$
\begin{aligned}
\tau_0 &= 0, \\
\eta_i &\geq 0, \\
\tau_{i,\eta_i} &= \inf\left\{t : t \geq \tau_{i-1}, \left|x_t - x_{\tau_{i-1}}\right| \geq \eta_i\right\}, \\
\tau_i &= \min\left\{\tau_{i,\eta_i}, T\right\}.
\end{aligned}
$$

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_η be the level-crossing time of the last set of equations. Then, given a positive real number α, the following minimal cost

$$\min_{\eta\geq 0} \mathcal{J}(\eta) = \min_{\eta\geq 0} \mathbb{E}\left[\int_0^{\tau_\eta\wedge T} x_s^2 ds + \alpha\big[(T-\tau_\eta)^+\big]^2\right] \tag{2.47}$$

turns out to be:

$$\beta\big[(T-\tau_\eta)^+\big]^2,$$

where $\beta > 0$ depends only on α. We will now prove this useful fact.

Notice that:

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

and that,

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

Thus, the cost (2.47) becomes:

$$\mathcal{J}(\eta) = \frac{T^2}{2} - \eta^2\mathbb{E}\big[(T-\tau_\eta)^+\big] - \left(\frac{1}{2}-\alpha\right)\mathbb{E}\left[\big[(T-\tau_\eta)^+\big]^2\right].$$

The above expression is convenient because, we can rewrite this in terms of T and λ alone. We have:

$$\mathcal{J}(\eta) = \frac{T^2}{2}\left[\varphi(\lambda) + \left(\frac{1}{2}-\alpha\right)\psi(\lambda)\right],$$

where, φ, ψ are defined through:

$$\begin{aligned}
\varphi(\lambda) &= 1 - \frac{\pi}{\lambda^2}\sum_{k=0}^{\infty}(-1)^k \frac{e^{-(2k+1)^2\lambda} - 1 + (2k+1)^2\lambda}{(2k+1)^2} \\
\psi(\lambda) &= \frac{16}{\pi\lambda^2}\sum_{k=0}^{\infty}(-1)^k \frac{e^{-(2k+1)^2\lambda} - 1 + (2k+1)^2\lambda - (1/2)(2k+1)^4 x^2}{(2k+1)^5}.
\end{aligned}$$

Then we have the optimal cost (2.47) as:

$$\min_{\eta \geq 0} \mathcal{J}(\eta) = \frac{T^2}{2} \inf_{\lambda} \left\{ \varphi(\lambda) + (0.5 - \alpha)\psi(\lambda) \right\}.$$

The minimal distortion for the level-triggered scheme with a single sample allowed is:

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

Based on the above discussion, we can define c_k recursively as follows: For $k \geq 2$,

$$\begin{aligned}
c_k &= \inf_{\lambda} \left\{ \varphi(\lambda) + (0.5 - c_{k-1})\psi(\lambda) \right\}, \\
\lambda_k^* &= \arg\inf_{\lambda} \left\{ \varphi(\lambda) + (0.5 - c_{k-1})\psi(\lambda) \right\}, \\
\rho_k &= \frac{\pi}{2\sqrt{2\lambda_k^*}}.
\end{aligned}$$

The optimal set of thresholds are given by:

$$\eta_k^* = \rho_{N-k+1} \sqrt{T - \tau_{k-1}}.$$

2.3.3 Optimal multiple sampling

Exactly like in the discussion of the previous section on multiple level-triggered sampling, we will obtain a parametric expression for the minimal expected distortion given at most k samples. Analogous to equation (2.47), consider the stopping cost:

$$\mathcal{J}(\tau) = \mathbb{E}\left[\int_0^{\tau \wedge T} x_s^2 ds + \frac{\alpha}{2} \left[(T - \tau)^+ \right]^2 \right] \tag{2.48}$$

where $\alpha \geq 0$ is a given constant. We can rewrite this as:

$$\frac{1}{2} \left\{ T^2 - \mathbb{E}\left[2x_{\tau \wedge T}^2 (T - \tau)^+ + (1 - \alpha) \left[(T - \tau)^+ \right]^2 \right] \right\}.$$

Note that there is no change in optimality by permitting τ to take values bigger than T. In fact the optimal τ even with this relaxation will a.s. be less than T. Like in the single sample case, let us pay attention to the part of the above expression which depends on τ and define the following optimal stopping problem:

$$\min_{\tau} \mathbb{E}\left[2x_\tau^2(T-\tau)+(1-\alpha)(T-\tau)^2\right].$$

Consider the candidate *maximum expected reward function*:

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

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

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

Then the optimal stopping time is given by:

$$\begin{aligned}\tau^* &= \inf_t \left\{t : g(x_t,t) \le 2x_t^2(T-t)+(1-\alpha)(T-t)^2\right\},\\ &= \inf_t \left\{t : x_t^2 \ge \sqrt{\frac{3(A-1+\alpha)}{A}}(T-t)\right\},\end{aligned}$$

and the corresponding optimal distortion $\mathcal{J}$ becomes

$$\mathcal{J}=(1-A)\frac{T^2}{2}.$$

Now, we obtain the explicit stopping rules and the corresponding minimal distortions for different values of the sample budget N by defining recursively κ_N, γ_N:

$$\begin{aligned}\kappa_N &= 1-\frac{(5+\kappa_{N-1})-\sqrt{(5+\kappa_{N-1})^2-24}}{4},\\ \gamma_N &= \sqrt{\frac{3(\kappa_{N-1}-\kappa_N)}{1-\kappa_N}}.\end{aligned}$$

The $(k+1)^{\text{th}}$ sampling time is chosen as:

$$\tau_{k+1} = \inf_{t \geq \tau_k} \left\{ t : (x_t - x_{\tau_k})^2 \geq \gamma_{N-k+1} T - t \right\}.$$

2.3.4 Comparisons

We now list a numerical comparison of the aggregate filtering distortions incurred by the three sampling strategies on the same time interval $[0, T]$. We obtained the distortions for all sampling strategies as product of $T^2/2$ and a positive coefficient. The numbers listed in the table are the values of these coefficients.

N	1	2	3	4
Deterministic	0.5	0.333	0.25	0.2
Level-triggered	0.3953	0.3471	0.3219	0.3078
Optimal	0.3660	0.2059	0.1388	0.1032

2.4 Sampling an Ornstein-Uhlenbeck process N-times

Now we turn to the case when the signal is an Ornstein-Uhlenbeck process:

$$dx_t = ax_t dt + dW_t, \quad t \in [0, T], \tag{2.49}$$

with $x_0 = 0$ and W_t being a standard Brownian motion.Again, the sampling times 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

$$\hat{x}_t = \begin{cases} 0 & \text{if } 0 \leq t < \tau_1, \\ x_{\tau_i} e^{a(t-\tau_i)} & \text{if } \tau_i \leq t < \tau_{i+1} \leq \tau_N, \\ x_{\tau_N} e^{a(t-\tau_N)} & \text{if } \tau_N \leq t \leq T. \end{cases} \tag{2.50}$$

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

$$\begin{aligned} J(S) &= \mathbb{E}\left[\int_0^T (x_s - \hat{x}_s)^2 ds\right] \\ &= \mathbb{E}\left[\int_0^{\tau_1} (x_s - \hat{x}_s)^2 ds + \sum_{i=2}^{N} \int_{\tau_{i-1}}^{\tau_i} (x_s - \hat{x}_s)^2 ds \right. \\ &\quad \left. + \int_{\tau_N}^{T} (x_s - \hat{x}_s)^2 ds\right]. \end{aligned}$$

2.4.1 Optimal deterministic sampling

We will show through induction that uniform sampling on the interval $[0, T]$ is the optimal deterministic choice of N samples

$$\{\tau_1, \tau_2, \ldots, \tau_N \mid 0 \leq \tau_i \leq T, \quad \tau_i \geq \tau_{i-1} \text{ for } i = 1, 2, \ldots, N\}$$

given that the initial value of the signal is zero.

When the number of samples permitted is N, the distortion takes the form:

$$\begin{aligned} J_{[0,T]}\left(\{\tau_1, \tau_2, \ldots \tau_N\}\right) = &\int_0^{\tau_1} \mathbb{E}\left(x_s - \hat{x}_s\right)^2 ds \\ &+ \int_{\tau_1}^{\tau_2} \mathbb{E}\left(x_s - \hat{x}_s\right)^2 ds + \ldots + \int_{\tau_N}^{T} \mathbb{E}\left(x_s - \hat{x}_s\right)^2 ds. \end{aligned}$$

For the induction step, we assume that the optimal choice of $N-1$ deterministic samples over $[T_1, T_2]$ is the uniform one:

$$\{\tau_1, \tau_2, \ldots \tau_{N-1}\} = \left\{T_1 + i\frac{T_2 - T_1}{N} \middle| i = 1, 2, \ldots, N-1\right\}.$$

The corresponding minimum distortion becomes:

$$\frac{N}{4a^2}\left(e^{2a\frac{T_2-T_1}{N}}-1\right)-\frac{1}{2a}\left(T_2-T_1\right)$$

The minimum distortion over the set of N sampling times is:

$$\begin{aligned}
&\min_{\{\tau_1,\tau_2,\ldots\tau_N\}} J_{[0,T]}\left(\{\tau_1,\tau_2,\ldots\tau_N\}\right)\\
&=\min_{\tau_1}\left\{\int_0^{\tau_1}\left(x_s-\hat{x}_s\right)^2 ds+\min_{\{\tau_2,\tau_2,\ldots\tau_N\}} J_{[\tau_1,T]}\left(\{\tau_1,\tau_2,\ldots\tau_N\}\right)\right\}\\
&=\min_{\tau_1}\left\{\frac{1}{4a^2}\left(e^{2a\tau_1}-1\right)+\frac{N}{4a^2}\left(e^{2a\frac{T-\tau_1}{N}}-1\right)-\frac{1}{2a}\left(T\right)\right\}\\
&=\frac{N+1}{4a^2}\left(e^{2a\frac{T}{N+1}}-1\right)-\frac{1}{2a}\left(T\right),
\end{aligned}$$

the minimum being achieved for $\tau_1 = T/(N+1)$. Thus, we have the uniform sampling scheme being the optimal one here too.

2.4.2 Optimal Level-triggered sampling

Let us first address the single sample case. The performance measure then takes the form

$$\begin{aligned}
\mathcal{J}(\tau_1) &= \mathbb{E}\left[\int_0^{\tau_1} x_t^2+\int_{\tau_1}^{T}\left(x_t-\hat{x}_t\right)^2 dt\right]\\
&= \mathbb{E}\left[\int_0^{T} x_t^2-2\int_{\tau_1}^{T} x_t\hat{x}_t\, dt+\int_{\tau_1}^{T}(\hat{x}_t)^2\, dt\right].
\end{aligned}$$

Now notice that the second term can be written as follows

$$\mathbb{E}\left[\int_{\tau_1}^{T} x_t\hat{x}_t\, dt\right]=\mathbb{E}\left[\int_{\tau_1}^{T}\mathbb{E}[x_t|\mathcal{F}_{\tau_1}]\hat{x}_t\, dt\right]=\mathbb{E}\left[\int_{\tau_1}^{T}(\hat{x}_t)^2\, dt\right],$$

where we have used the strong Markov property of x_t, and that for $t > \tau_1$ we have $\mathbb{E}[x_t|\mathcal{F}_{\tau_1}] = x_\tau e^{-a(t-\tau_1)} = \hat{x}_t$. Because of this observation the performance measure

$\mathcal{J}(\tau_1)$ takes the form

$$
\begin{aligned}
\mathcal{J}(\tau_1) &= \mathbb{E}\left[\int_0^T x_t^2\, dt - \int_{\tau_1}^T (\hat{x}_t)^2\, dt\right] \\
&= \frac{e^{2aT} - 1 - 2aT}{4a^2} - \mathbb{E}\left[x_{\tau_1}^2 \frac{e^{2a(T-\tau_1)} - 1}{2a}\right] \\
&= T^2 \left\{ \frac{e^{2aT} - 1 - 2aT}{4(aT)^2} - \mathbb{E}\left[\frac{x_{\tau_1}^2}{T} \frac{e^{2(aT)(1-\tau_1/T)} - 1}{2(aT)}\right]\right\} \\
&= T^2 \left\{ \frac{e^{-2\bar{a}} - 1 + 2\bar{a}}{4\bar{a}^2} - \mathbb{E}\left[-\bar{x}_{\bar{\tau}_1}^2 \frac{e^{2\bar{a}(1-\bar{\tau}_1)} - 1}{2\bar{a}}\right]\right\}
\end{aligned}
$$

where

$$
\bar{t} = \frac{t}{T}; \quad \bar{a} = aT; \quad \bar{x}_{\bar{t}} = \frac{x_{\frac{t}{T}}}{\sqrt{T}}. \tag{2.51}
$$

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 (2.51). In fact, we can solve the multiple sampling problem on $[0, T]$ without loss of generality.

We carry over the definitions for threshold sampling times from section 2.3.2. We do not have series expansions like for the case of the Wiener process. Instead we have a computational procedure that involves solving a PDE initial and boundary value problem. We have a nested sequence of optimization problems. The choice at each stage being the non-zero level η_i. For $N = 1$, the distortion corresponding to a chosen η_1 is given by:

$$\frac{1}{4a^2}\left(e^{2aT}-1\right)-\frac{1}{2a}T-\frac{\eta_1^2}{2a}\mathbb{E}\left[e^{2a(T-\tau_1)}-1\right]$$
$$=\frac{1}{4a^2}\left(e^{2aT}-1\right)-\frac{1}{2a}T-\frac{\eta_1^2}{2a}\left(e^{2aT}\left(1+2aU^1(0,0)\right)-1\right),$$

where the function $U^1(x,t)$ satisfies the PDE:

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

along with the boundary and initial conditions:

$$\begin{cases} U^1(-\eta_1,t)=U^1(\eta_1,t)=0 & \text{for } \ t\in[0,T], \\ U^1(x,T)=0 & \text{for } \ x\in[-\eta_1,\eta_1]. \end{cases}$$

We choose the optimal η_1 by computing the resultant distortion for increasing values of η_1 and stopping when the cost stops decreasing and starts increasing. Note that the solution $U(0,t)$ to the PDE furnishes us with the performance of the η_1-triggered sampling over $[t,T]$. We will use this to solve the multiple sampling problem.

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

$$\frac{1}{2}U_{xx}+axU_x+U_t=0,$$

along with the boundary and initial conditions:

$$\begin{cases} U^{N+1}(-\eta_1,t)=U^{N+1}(\eta_1,t)=G_N(T-t) & \text{for } \ t\in[0,T], \\ U^{N+1}(x,T)=0 & \text{for } \ x\in[-\eta_1,\eta_1]. \end{cases}$$

Then the distortion we are seeking to minimize over η_1 is given by:

$$\begin{aligned}
&\frac{1}{4a^2}\left(e^{2aT}-1\right)-\frac{1}{2a}T \\
&\qquad -\frac{\eta_1^2}{2a}\mathbb{E}\left[e^{2a(T-\tau_1)}-1+\frac{1}{4a^2}\left(e^{2a(T-\tau_1)}-1\right)-\frac{1}{2a}\left(T-\tau_1\right)\right] \\
&\qquad\qquad +\mathbb{E}\left[G_N(T-\tau_1)\right] \\
&\qquad =\frac{1}{4a^2}\left(e^{2aT}-1\right)-\frac{1}{2a}T-\frac{\eta_1^2}{2a}\left(e^{2aT}\left(1+2aU^1(0,0)\right)-1\right)-U^{N+1}.
\end{aligned}$$

We choose the optimal η_1 by computing the resultant distortion for increasing values of η_1 and stopping when the distortion stops decreasing.

2.4.3 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 $1 \leq n \leq M$

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

$\{w_n\}$ is an i.i.d. Gaussian sequence.

Sampling 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 $n = M, M-1, \ldots, 0$,

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

The above equation provides a (backward) recurrence relation for the computation of the single sampling cost function $V_n^1(x)$. Notice that for values of x for which the l.h.s. exceeds the r.h.s. 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 that if $|x_n| \ge \lambda_n$ we sample, otherwise we go to the next time instant. The numerical solution of the recursion presents no special difficulty if $a \le 1$. 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 $n = M, M-1, \ldots 0$,

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

We provide graphs describing the relative performances of the three sampling schemes in Figures 2.6.

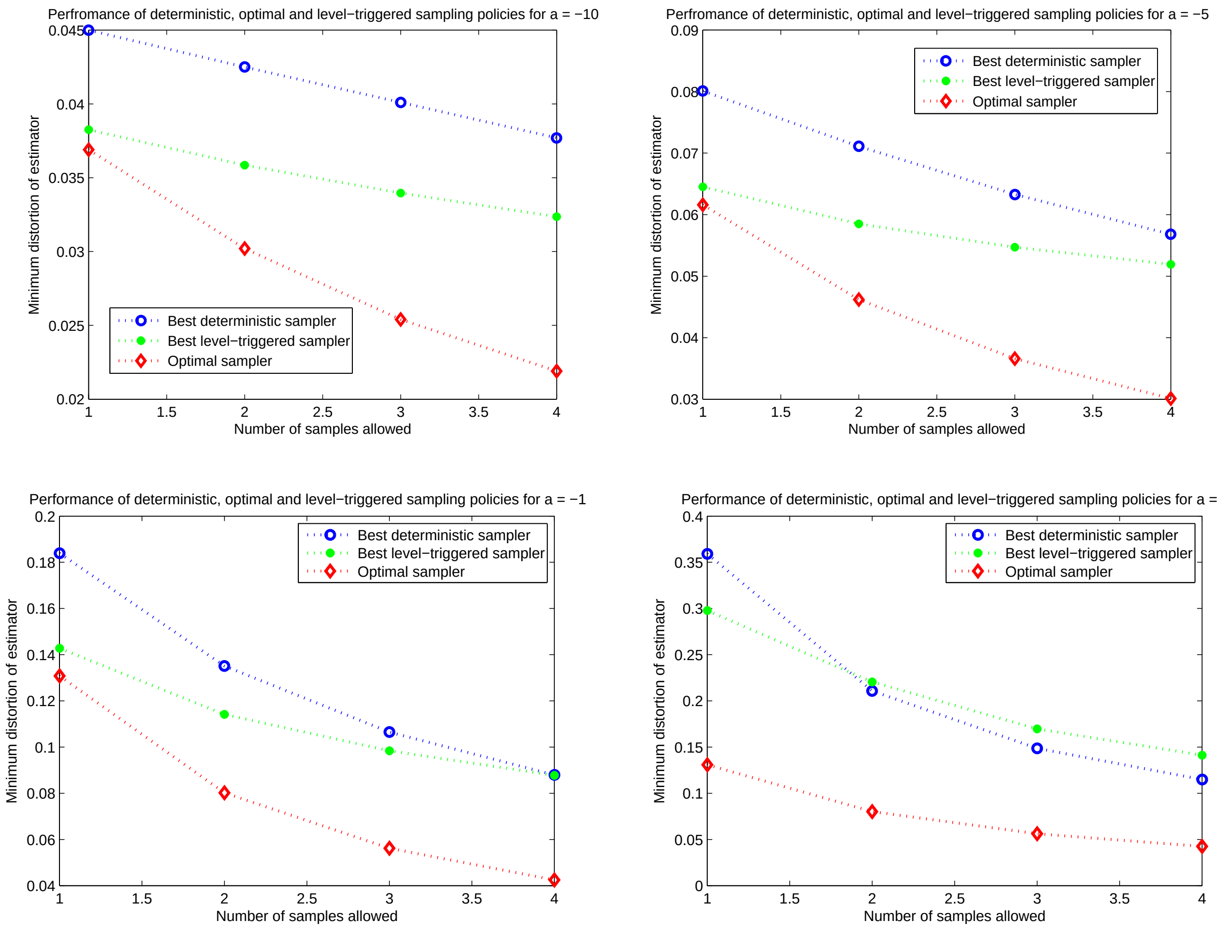

Figure 2.6: Relative performance of Optimal, Threshold and Deterministic sampling schemes as a function of initial variance σ^2 and parameter values (a) $a = 10$, (b) $a = 1$, (c) $a = 0$ and (d) $a = -1$.

Chapter 3

Average Cost Repeated Sampling for Filtering

3.1 Introduction

A sensor makes continuous observations of a Gaussian signal process. It transmits at times it chooses, samples of its observations to a Supervisor which uses this stream of samples to maintain a filtered (real-time) estimate of the signal. We study the tracking performance of an efficient sampling scheme which is event-triggered. Such problems arise in sensor networks because of the limited capacity of a remote sensor node to communicate to the supervisor. For simplicity of exposition, we take the signal to be the scalar Ornstein-Uhlenbeck process. We will see that the constant threshold sampling strategy will be optimal. In the presence of packet losses with the loss events independent of the signal, we conjecture that the constant threshold policy is still optimal.

The work of Sinopoli et.al. [31] discusses the filtering performance of a periodic sampling policy when the samples could be lost according to an IID Bernoulli sequence which is also independent of the signal process.

By establishing the optimality of threshold sampling schemes for the infinite horizon filtering problem, we will also prove the optimality of Lebesgue sampling for an the infinite horizon impulse control problem of Åström and Bernhardsson [15, 1].

We study the repeated sampling problem when the signal process we are interested in is perfectly observed by the sensor. When the signal process is only partially observed, the least-squares filtered estimate or a risk-sensitive filtered estimate can take the role of the fully observed signal.

3.2 Real-time Estimation

Our signal process $x(\cdot)$ is defined on $[0, \infty)$ and is governed by

$$dx(t) = ax(t)dt + bdW(t), \quad x(0) = x_0, \tag{3.1}$$

where $W(t)$ is a standard Brownian motion process. A sensor observes this process perfectly. It has to pick an increasing sequence of *sampling times*

$$\Theta(x_0) = \{\tau_i | i = 0, 1, 2, 3, \ldots\},$$

such that

$$\begin{aligned} &\tau_0 = 0, \\ &\tau_i < \tau_{i+1} \ \text{ a.s. for } \ i \geq 0, \ \text{and} \\ &\tau_i \ \text{ is measurable w.r.t. } \ \mathcal{F}_t^x. \end{aligned}$$

At these times, the sensor transmits the value of the x-process to the supervisor which receives these samples reliably and with negligible delay. The supervisor maintains the estimate waveform $\hat{x}(t)$ based upon the point process of received samples.

The performance of the sampling scheme will be measured by its communication cost - the *average sampling rate* $R_{av}(\Theta)$ and the resultant *average distortion*

$D_{av}(\Theta)$ of the supervisor's estimate. We seek a sampling policy that minimizes the average distortion while not exceeding a prescribed average sample rate. The average sampling rate is computed as

$$R_{av}(\Theta) = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\Big[\#\text{ofSamplesIn}(0,T]\Big] = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{i=1}^{\infty} \mathbf{1}_{\{\tau_i \le T\}}\right].$$

We will exclusively deal with the squared error distortion and so

$$D_{av}(\Theta) = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T (x(s) - \hat{x}(s))^2 ds\right].$$

Define the *Last Sample Count* $l(t)$, which is $\mathcal{F}_t^x$-adapted, through

$$l(t) = \max\{i \ge 0 | \tau_i \le t\} \quad \text{for} \quad t \in [0, \infty).$$

Then, the least-squares estimate $\hat{x}(t)$ is computed by the formulae:

$$\hat{x}(\tau_{l(t)}) = x(\tau_{l(t)}),$$

$$\hat{x}(t) = e^{s\left(t-\tau_{l(t)}\right)} x(\tau_{l(t)}).$$

The error process ϵ_t which is defined as $x_t - \hat{x}_t$ is a jump diffusion process governed by the equations:

$$d\epsilon_t = a\epsilon_t dt + bdW_t, \quad \text{on } [\tau_i, \tau_{i+1}), \ \forall\, i \ge 0,$$

$$\epsilon_{\tau_i} = 0.$$

Notice that the x-process can be reconstructed causally from the trajectory of the error process using the exact values of the jumps in $\epsilon_.$ as in:

$$x(t) = \epsilon_t + \sum_{i=1}^{\infty} \mathbf{1}_{\{\tau_i \le t\}} \left(\lim_{\delta \downarrow 0} \epsilon_{\tau_i - \delta}\right). \tag{3.2}$$

Hence,

$$\mathcal{F}_t^x = \mathcal{F}_t^\epsilon.$$

3.3 Optimal repeated sampling

Stated in terms of the error process, the *optimal sampling problem* is to pick a sequence of stopping times Θ, adapted to the ϵ-process such that the Lagrangian

$$\begin{aligned} J(\Theta) &= \lambda R_{av}(\Theta) + D_{av}(\Theta) \\ &= \lambda \times \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{i=1}^{\infty} \mathbf{1}_{\{\tau_i \le T\}}\right] + \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T \epsilon^2(s)ds\right], \end{aligned}$$

is minimized. Here $\lambda > 0$ is a Lagrange multiplier. The sampling problem can be viewed as one of resetting the new state process $\epsilon(\cdot)$ at stopping times to minimize an average cost. In what follows, we will restrict our attention to Markov times that are time-homogeneous:

$$\begin{aligned} \tau_{i+1} &= \tau_i + \delta_i\left(\epsilon_{\tau_i}\right), \\ &= \tau_i + \delta_i(0) \text{ for } \forall i \ge 1, \\ &= \tau_i + \delta_i, \end{aligned}$$

where, $\{\delta_i\}_1^\infty$ is an IID sequence. The condition that the average sample rate be finite implies

$$\mathbb{E}[\delta_i] < \infty.$$

This makes the error process ergodic. However, we don't need to use (and hence prove) this ergodicity explicitly. Notice that

$$\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{i=1}^{\infty} (\tau_i - \tau_{i-1})\,\mathbf{1}_{\{\tau_i \le T\}}\right] \le 1 \le \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\delta + \sum_{i=1}^{\infty} (\tau_i - \tau_{i-1})\,\mathbf{1}_{\{\tau_i \le T\}}\right],$$

where δ has the same law as any individual random variable in the IID sequence $\{\delta_i\}$. This leads us to:

$$\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{i=1}^{\infty} \delta_{i-1}\mathbf{1}_{\{\tau_i \le T\}}\right] \le 1 \le \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\delta + \sum_{i=1}^{\infty} \delta_{i-1}\mathbf{1}_{\{\tau_i \le T\}}\right],$$

giving us

$$\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{i=1}^{\infty} \delta_{i-1}\mathbf{1}_{\{\tau_i \le T\}}\right] = 1,$$

$$\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\sum_{i=1}^{\#\text{ofSamplesIn}(0,T]} \delta_{i-1}\right] = 1,$$

$$\limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\delta\right] \times \mathbb{E}\Big[\#\text{ofSamplesIn}(0,T]\Big] = 1,$$

$$R_{av}(\Theta) = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\Big[\#\text{ofSamplesIn}(0,T]\Big] = \frac{1}{\mathbb{E}\left[\delta\right]}. \tag{3.3}$$

Similarly,

$$D_{av}(\Theta) = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T \epsilon^2(s)ds\right] = \frac{\mathbb{E}\left[\int_0^\delta \epsilon_s^2 ds\right]}{\mathbb{E}\left[\delta\right]}. \tag{3.4}$$

Now the solution to the stopping problem is obtained by picking a stopping time (adapted to the ϵ-process) δ^* which minimizes the cost

$$\tilde{J}(\delta) = \frac{\mathbb{E}\left[\int_0^\delta \epsilon_s^2 ds\right]}{\mathbb{E}\left[\delta\right]} + \frac{\lambda}{\mathbb{E}\left[\delta\right]}.$$

This cost has the interpretation that by picking a positive value for λ, we hold the average packet rate at a corresponding fixed value and minimize the average distortion. If δ^* is such that

$$\tilde{J}(\delta) \ge \tilde{J}(\delta^*) = \gamma(\lambda) > 0, \tag{3.5}$$

then, it follows that,

$$\mathbb{E}\left[\int_0^\delta \epsilon_s^2 ds\right] + \lambda - \gamma(\lambda)\mathbb{E}\left[\delta\right] \;\ge\; \mathbb{E}\left[\int_0^{\delta^*} \epsilon_s^2 ds\right] + \lambda - \gamma(\lambda)\mathbb{E}\left[\delta^*\right] \;=\; 0. \tag{3.6}$$

In fact, the existence of a δ^* satisfying

$$\mathbb{E}\left[\delta^*\right] > 0,$$

and (3.6) supplies us with one satisfying (3.5).

3.3.1 Solving the single stopping problem

Here, we carry-out the calculations needed to find a candidate δ^* satisfying equation (3.6). For positive γ, let

$$H(\delta, \gamma) = \lambda + \text{ess inf}_{\delta}\, \mathbb{E}\left[\int_0^{\delta} \left(\epsilon_s^2 - \gamma\right) ds\right].$$

To solve this optimal stopping problem, we introduce its *Snell envelope.* We seek a twice differentiable function $g(x) : \mathbb{R} \to \mathbb{R}$, which satisfies the following PDE:

$$\frac{b^2}{2} g\prime\prime(x) + axg\prime(x) + x^2 - \gamma = 0, \tag{3.7}$$

Then, by the results of A we have:

$$\begin{aligned}
\mathbb{E}\left[g\left(\epsilon_\tau\right) - g\left(\epsilon_0\right)\right] &= \mathbb{E}\left[\int_0^\tau dg(\epsilon_s)\right] \\
&= \mathbb{E}\left[\int_0^\tau \left\{\frac{b^2}{2} g\prime\prime\left(\epsilon_s\right) + a\epsilon_s g\prime\left(\epsilon_s\right)\right\} ds\right] \\
&= \mathbb{E}\left[\int_0^\tau -\left\{\epsilon_s^2 - \gamma\right\} ds\right].
\end{aligned}$$

Notice that the general solution to (3.7) has two parameters. If we can select the parameters on which g depends so that

$$g(\cdot) \leq 0; \quad g(\pm\eta) = 0 \tag{3.8}$$

for some η (dependent on λ), then for any stopping time τ we can conclude that for any stopping time τ,

$$\mathbb{E}\left[\int_0^{\tau}\left\{\epsilon_s^2-\gamma\right\}ds\right] \geq \mathbb{E}\left[g(e_\tau)+\int_0^{\tau}\left\{\epsilon_s^2-\gamma\right\}ds\right] = g(0). \tag{3.9}$$

Hence, if we can find a solution $g(\cdot)$ to (3.7) that satisfies (3.8), the above relation suggests that the performance of any stopping time will be lower bounded by $g(0)$. We have equality iff the first inequality in (3.9) becomes an equality. This happens when the stopping time τ stops only when ϵ_t hits either of the two values $\pm\eta$. Denote this (optimal) stopping time as τ_0. To complete the proof, we need to the existence of a nonpositive function $g(\cdot)$ which is a solution to (3.7) and the existence of a constant η that satisfies (3.8).

The general solution to (3.7)cannot be found in closed form. Because of the symmetry of the process ϵ_t, we must look for an even symmetric function $g(\cdot)$. We can then verify that

$$g(x) = \frac{b^2}{2a}\left\{2\left(1+2a\frac{\lambda}{b^2}\right)\left(\int_{\frac{\eta}{b}}^{\frac{x}{b}} e^{-az^2}\int_0^z e^{aw^2}dwdz\right) - \left(\frac{x^2}{b^2}\right) + A\right\},$$

is the general, even symmetric solution to (3.7). In order to satisfy (3.8), we have two parameters namely, A and η, and two equations

$$g(\eta) = g'(\eta) = 0.$$

The second equation is needed because $g(\cdot)$ must "touch" the value 0 at $e = \eta$ (otherwise, due to continuity, $g(\cdot)$ will assume positive values).

Consider the first equation $g'(\eta) = 0$. Due to symmetry, we limit ourselves to

$\eta \geq 0$. We then obtain

$$0 = -ag'(\eta) = \frac{\eta}{b} - \left(1 + \frac{2a\lambda}{b^2}\right) e^{-a\eta^2} \int_0^{\frac{\eta}{b}} e^{aw^2} dw \tag{3.10}$$

which is the equation that defines η as a function of λ. It is easy to show that when $0 \leq 1 + \frac{2a\lambda}{b^2} < 1$ and $g'(\cdot)$ is convex and the equation $g'(\eta) = 0$ has a unique positive solution $\eta > 0$. Note that 0 is also a solution to this equation but not a positive one. Using this uniqueness of η and the convexity of $g'(\cdot)$, we can verify that for $x \geq 0$, the derivative $g'(x)$ has the same sign as $\eta - x$. this suggests that $g(e)$ has a maximum at $x = \eta$ (and a local minimum at $x = 0$) meaning that $g(x) \leq g(\eta)$.

We now produce the desired function $g(\cdot)$ which satisfies the conditions of (3.7) and permits the existence of the stopping time τ^* with bounded expectation: where $\eta \geq 0$ is determined by γ.

It is easy to see that $g(\eta) = 0$ and $g(x) > 0$ when $x \neq \eta$. Hence, the stopping time:

$$\begin{aligned} \tau^* &= \inf \left\{ t \middle| g(x(t)) = 0 \right\} \\ &= \inf \left\{ t \middle| |x(t)| = \eta \right\} \end{aligned} \tag{3.11}$$

is optimal.

3.3.2 Performance gains

Without loss of generality, we can assume that $b = 1$. The average distortion (3.4) and sampling rate (3.3) incurred by using threshold η can be found by solving ODEs [32]:

$$R_{av}(\eta) = \frac{1}{\mathbb{E}[\tau^*]} = \frac{1}{h(0)},$$

where $h(\cdot)$ satisfies (see chapter 9, page 175 of [32]):

$$ax\frac{dh}{dx} + \frac{1}{2}\frac{d^2h}{dx^2} = -1,$$

with the boundary conditions:

$$h(\eta) = 0, h(-\eta) = 0.$$

Similarly,

$$D_{av}(\eta) = R_{av}\mathbb{E}\left[\int_0^{\tau^*} \epsilon_s^2 ds\right] = R_{av}f(0),$$

where $f(\cdot)$ satisfies [32]:

$$ax\frac{df}{dx} + \frac{1}{2}\frac{d^2f}{dx^2} = -x^2,$$

with the boundary conditions:

$$f(\eta) = 0, f(-\eta) = 0.$$

Solving for the functions $h(\cdot), f(\cdot)$ gives us the performance parameters for optimal sampling to be:

$$\begin{aligned} R_{av}(\eta) &= \frac{1}{\int_0^\eta e^{-az^2}\int_0^z e^{ay^2}dydz}, \\ D_{av}(\eta) &= \frac{\int_0^\eta e^{-az^2}\int_0^z y^2e^{ay^2}dydz}{\int_0^\eta e^{-az^2}\int_0^z e^{ay^2}dydz}. \end{aligned}$$

For periodic sampling with a sampling period of Δ, we get:

$$\begin{aligned} R_{av}^{per}(\Delta) &= \frac{1}{\Delta}, \\ D_{av}^{per}(\Delta) &= \frac{1}{\Delta}\left(\frac{e^{2a\Delta}-1}{4a^2} - \frac{\Delta}{2a}\right). \end{aligned}$$

In Figure 3.1, we have a graphical comparison of the performance gains of using the optimal sampling scheme:

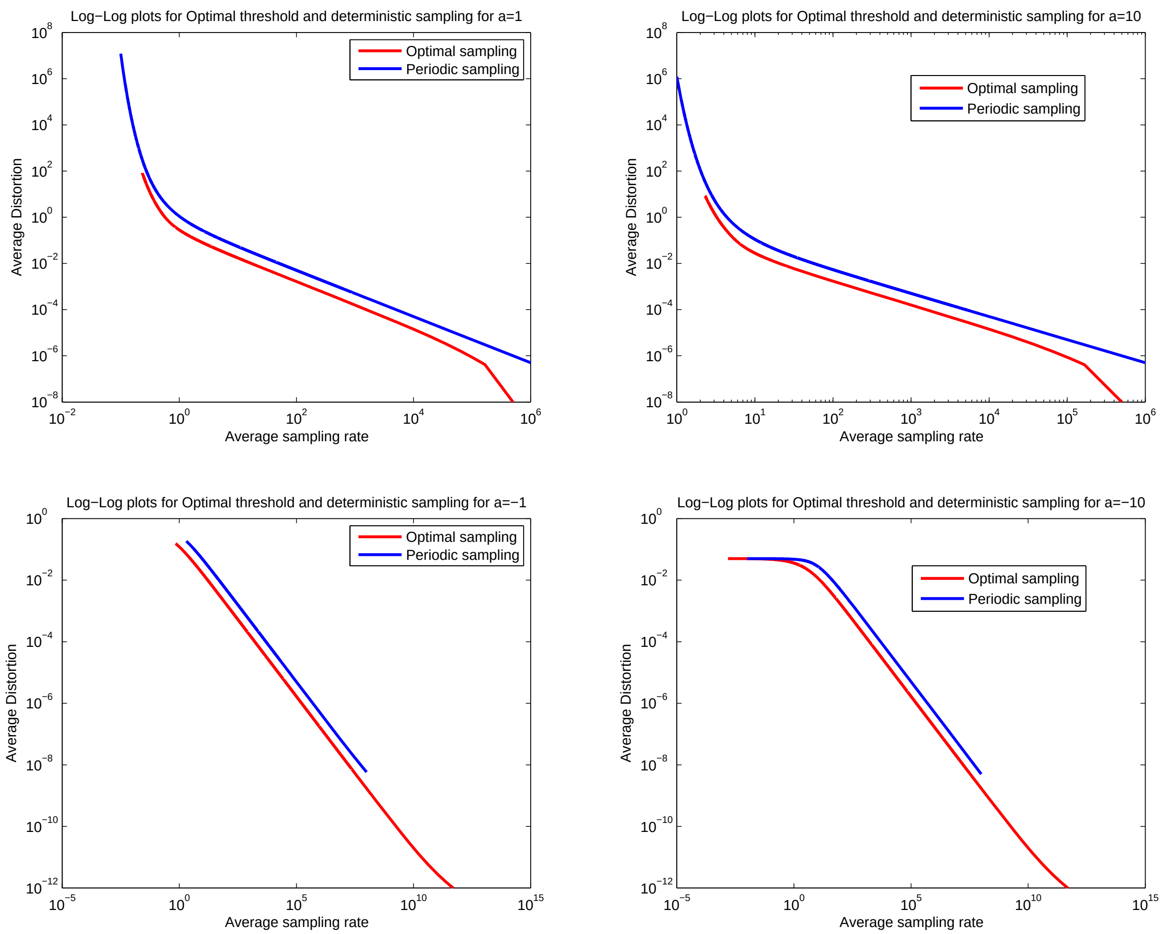

Figure 3.1: Relative performance of Optimal (Threshold) and Periodic samplers as a function of average sampling rate R_{av} and parameter values (a) $a = 1$, (b) $a = 10$, (c) $a = -1$ and (d) $a = -10$.

Chapter 4

Average Cost Control with Level-triggered Sampling

4.1 Introduction

For the single sensor configuration used in the previous chapters, we focus on Average cost optimal control in this chapter. The methods of this chapter apply only

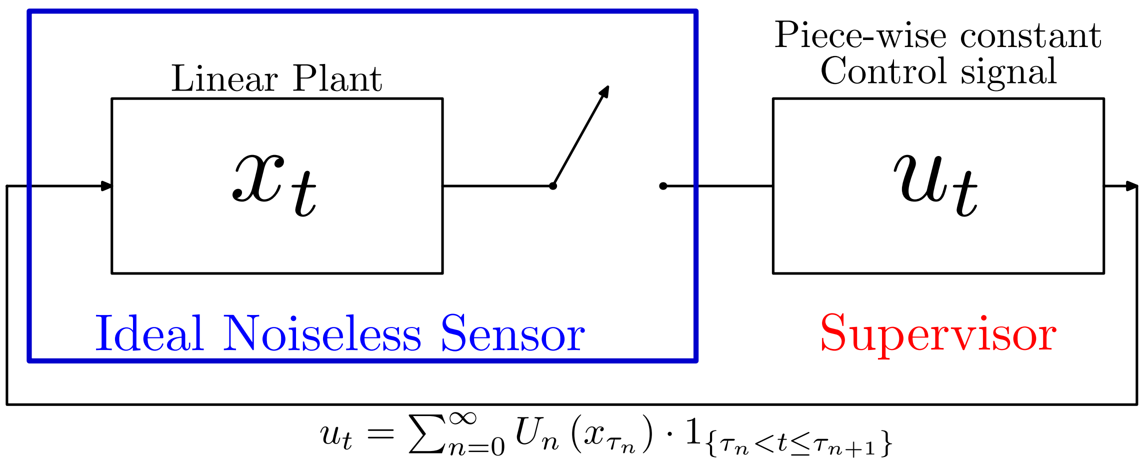

Figure 4.1: Sample and hold control

to scalar systems. Although the plant we stabilize will be linear, extension to the nonlinear case is straightforward albeit computationally much more burdensome.

On an infinite horizon, the sensor sends samples to the supervisor which issues a piece-wise constant control signal. The values of the control signal are allowed to change only at times when the supervisor receives samples. The control objective is to minimize the average power of the state signal. The sensor's objective is to aid in the control task and send samples as often as it can while also respecting a bound on the average sampling rate.

The problem of jointly optimal sampling and control for the Linear System is much more difficult than the jointly optimal sampling and filtering problem tackled in the preceding chapter. In the estimation problem, the error signal was reset to zero at sampling times and so, the repeated sampling problem was reduced to the problem of choosing a single sampling policy to be repeated. In the control problem however, the state signal does not get reset to zero like the error signal does for the estimation problem. Thus, no reduction to repeating the same sampling policy is possible. In practical terms, this means that the feedback control signal as well as the sampling policy should be 'aggressive' when the state wanders away from the origin.

This problem differs in its information pattern from similar ones addressed in the Stochastic Control literature. The works [20, 37, 19] seek combined control and stopping policies with both of these adapted to the same signal process. In our problem on the other hand, the stopping policy is allowed to depend on the x-signal process while the control signal is adapted to the sampled sequence. The work of [38] discusses the LQG control performance under Poisson sampling. A deterministic version of control for event-triggered sampling is presented in [16].

In this chapter, we will seek optimal control policies corresponding to a chosen sampling strategy. We will study the performances of the optimal controls corresponding to two types of level-triggered sampling strategies.

4.2 Average cost control problem

The state signal obeys:

$$dx_t = ax_t dt + dW_t + u_t dt, \quad x(0) = x_0, \tag{4.1}$$

where W_t is a standard Brownian motion process and the control signal is piecewise constant and adapted to the sampled stream. Let $\mathcal{T}$ be the sequence of sampling times:

$$\mathcal{T} = \{\tau_0, \tau_1, \tau_2, \ldots\},$$

with

$$\begin{aligned} &\tau_0 = 0, \\ &\tau_i < \tau_{i+1} \text{ a.s. for } i \geq 0, \text{ and} \\ &\tau_i \text{ is measurable w.r.t. } \mathcal{F}_t^x. \end{aligned}$$

Then the (stationary feedback) control policy U should be adapted to the sample sequence. Define the process $\{\Psi_t\}$ describing information available at the controller from the last received data packet:

$$\Psi_t = \begin{pmatrix} \chi_t \\ \nu_t \end{pmatrix} = \begin{pmatrix} \sum_{n=0}^{\infty} 1_{\{\tau_n < t \leq \tau_{n+1}\}} \cdot x_{\tau_n} \\ \sum_{n=0}^{\infty} 1_{\{\tau_n < t \leq \tau_{n+1}\}} \cdot \tau_n \end{pmatrix}.$$

Let $\mathscr{U}$ stand for the set of control policies U that are adapted to the Ψ-process. The actual control signal generated by U is given by:

$$u_t = \sum_{n=0}^{\infty} 1_{\{\tau_n < t \leq \tau_{n+1}\}} \cdot U_n(x_{\tau_n}),$$

where $U_n(x_{\tau_n})$ is the value of the control signal after the sample at time τ_n has been received.

Stabilization performance is measured through the average power of the state signal:

$$J_u = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T x_s^2 ds\right],$$

while the average sampling rate:

$$R = \limsup_{T\to\infty} \frac{1}{T}\mathbb{E}\left[\int_0^T \sum_{n=0}^{\infty} \mathbf{1}_{\{\tau_n\le T\}}\delta\left(s-\tau_n\right) ds\right]$$

is kept less than or equal to a desired bound. Here, $\delta(\cdot)$ is the Dirac-delta function.

Since we use stationary feedback controls, the sampled stream forms a controlled Markov chain in discrete-time. We will translate the continuous-time optimal control problem into one in discrete time and seek solutions. In the next section, we do this for the case of periodic sampling.

4.3 Optimal control under periodic sampling

Under periodic sampling, the sample times are given by

$$\tau_n = n\Delta \quad \text{for} \ \ n \ge 0.$$

The sampled state takes the form of a discrete time linear system:

$$X_{n+1} = e^{aT}X_n + \int_{n\Delta}^{(n+1)\Delta} e^{a((n+1)\Delta-s)}dW_s + \int_{n\Delta}^{(n+1)\Delta} e^{a((n+1)\Delta-s)}u_s ds \tag{4.2}$$

$$= e^{aT}X_n + \sqrt{\frac{\mu^2-1}{2a}}V_n + \frac{\mu-1}{a}U_n, \tag{4.3}$$

where $\{V_n\}$ is an IID sequence of standard Gaussian random variables and $\mu = e^{a\Delta}$.

It is easy to find feedback control policies $\{U_n\}$ that stabilize (in the mean square sense) the sampled linear system. For example, linear feedback controls of the form

$$U_n = kX_n, \quad n \geq 0,$$

with $-1 < k + \frac{\mu^2-1}{2a} < 1$ will stabilize $\{X_n\}$. It can also be seen that stability of the sampled state sequence implies stability of the original continuous time system. We shall restrict our attention to mean-square stabilizing control policies that also make the controlled process (4.3) ergodic with a p.d.f such that the fourth moment at steady state is finite. We will need this restriction to translate the continuous time optimization problem into an equivalent one in discrete time. The class of linear feedback policies which are described above are included in our restricted policy space. Let ρ_U be the steady state p.d.f. corresponding to the control policy $\{U_n\}$. When the sampled state sequence is ergodic, so is the actual continuous time state waveform.

4.3.1 Equivalent discrete time ergodic control problem

The expected integral cost of (4.2) is the sum of expected integrals over the inter-sample intervals. We want to be able to write the expected integral costs during such intervals as functions of the state of the chain at the beginning (or end) of the interval instead of being functions of the chain states at both end-points.

Because of the assumed ergodicity, we can replace the average cost (4.2) with the long run average cost (ergodic cost). Then, along the lines of lemma 3.4 of [39], we have:

$$
\begin{aligned}
J_u &\overset{\text{a.s.}}{=} \limsup_{T\to\infty} \frac{1}{T}\int_0^T x_s^2 ds \\
&= \limsup_{N\to\infty} \frac{1}{N\Delta}\sum_{n=0}^{N-1}\int_{\tau_n}^{\tau_{n+1}} x_s^2 ds \\
&= \limsup_{N\to\infty} \frac{1}{N\Delta}\sum_{n=0}^{N-1}\left\{\int_{\tau_n}^{\tau_{n+1}} x_s^2 ds - \mathbb{E}\left[\int_{\tau_n}^{\tau_{n+1}} x_s^2 ds\middle| X_n\right]\right\} \\
&\quad + \limsup_{N\to\infty} \frac{1}{N\Delta}\sum_{n=0}^{N-1}\mathbb{E}\left[\int_{\tau_n}^{\tau_{n+1}} x_s^2 ds\middle| X_n\right].
\end{aligned}
$$

The first part of the last expression is zero according to the Martingale stability theorem (page 105 of [40]). We are able to use this theorem because of the finiteness of the fourth moment of the state signal.

Let $\delta(\cdot)$ denote the Dirac-delta function. We have:

$$
\begin{aligned}
J_u &\overset{\text{a.s.}}{=} \limsup_{N\to\infty} \frac{1}{N\Delta}\sum_{n=0}^{N-1}\mathbb{E}\left[\int_{\tau_n}^{\tau_{n+1}} x_s^2 ds\middle| X_n\right] \\
&= \limsup_{N\to\infty} \frac{1}{N\Delta}\sum_{n=0}^{N-1}\int_{\mathbb{R}}\mathbb{E}\left[\int_{\tau_n}^{\tau_{n+1}} x_s^2 ds\right]\delta(x - X_n)dx \\
&= \frac{1}{\Delta}\int_{\mathbb{R}}\mathbb{E}\left[\int_0^{\Delta} x_s^2 ds\middle| X_0 = x\right]\rho_{U}(x)dx \\
&= \frac{1}{\Delta}\int_{\mathbb{R}}\left\{AX_n^2 + 2BX_nU_n + CU_n^2\right\}\rho_{U}(x)dx,
\end{aligned}
$$

where,

$$
\begin{aligned}
A &= \frac{\mu^2 - 1}{2a}, \\
B &= \frac{1}{a}\left\{\frac{\mu^2 - 1}{2a} - \frac{\mu - 1}{a}\right\}, \\
C &= \frac{1}{a^2}\left\{\frac{\mu^2 - 1}{2a} - 2\frac{\mu - 1}{a} + T\right\}.
\end{aligned}
$$

In fact, we can write:

$$J_u \stackrel{\text{a.s.}}{=} \limsup_{N\to\infty} \frac{1}{N\Delta} \sum_{n=0}^{N-1} AX_n^2 + 2BX_nU_n + CU_n^2$$
$$\triangleq G_U.$$

The solution to this fully observed, average cost ($\equiv$ ergodic cost) control problem is well known [41]. The optimal controls are linear feedback controls. To find the optimal feedback gain, we use the Average Cost optimality equation:

$$\alpha^* + h(x) = \inf_{u\in\mathbb{R}} \left\{ Ax^2 + 2Bux + Cu^2 + \mathbb{E}\left[h\left(\mu x + \frac{\mu-1}{a}u + \sqrt{\frac{\mu^2-1}{2a}}W_0 \right) \right] \right\}.$$

Here, α^* is the minimum cost and a quadratic choice for $h(\cdot)$ verifies the optimality of linear feedback. The optimal control is given by:

$$U_n^* = -\frac{a(\mu^2-1)^2 + 2a^2\theta\mu(\mu-1)}{(\mu-2)^2 - 1 + 2a\Delta + 2a\theta(\mu-1)^2} X_n,$$

where,

$$\theta = \frac{D}{2} + \sqrt{\frac{D}{2}^2 + \frac{D}{2a}},$$
$$D = \frac{\mu+1}{\mu-1}\left(\frac{\mu^2-1}{2a} - 2\frac{\mu-1}{a} + \Delta \right) - \frac{(\mu-1)^2}{2a}.$$

The minimum average cost is:

$$\alpha^* = \frac{\theta}{\Delta}\frac{(\mu^2-1)^2}{2a}.$$

The sampling rate is of course equal to $\frac{1}{\Delta}$.

4.4 Level-triggered sampling

Let $\mathcal{L}$ be a given infinite set of levels:

$$\mathcal{L} = \{\ldots, l_{-2}, l_{-1}, l_0, l_1, l_2, \ldots\},$$

$$
\text{with,} \begin{cases} l_i \in \mathbb{R} & \forall\ i \in \mathbb{Z}, \\ l_i < l_{i+1} & \forall\ i \in \mathbb{Z}, \\ l_0 = 0. \end{cases}
$$

If we want a set of levels all non-zero, we just remove l_0 from the set $\mathcal{L}$. We need an infinite set in order to be able to stabilize unstable plants

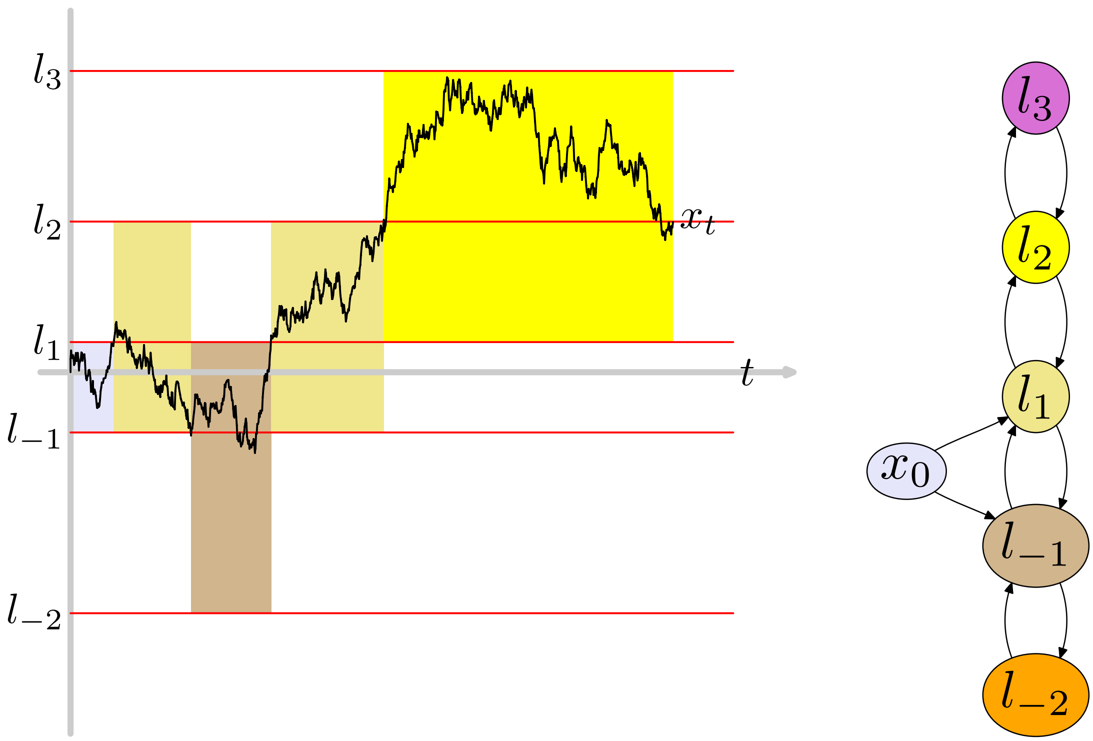

Figure 4.2: Level-triggered sampling and the associated Markov chain. All levels are non-zero. The initial state does not belong to the set of levels $\mathcal{L}$. This gives rise to the only transient state 'x_0'.

The sampling times triggered by $\mathcal{L}$ are defined through fresh crossings of levels:

$$
\tau = 0,
$$

$$
\tau = \inf \left\{ \tau \middle| \tau > \tau_i, x_\tau \in \mathcal{L}, x_\tau \notin x_{\tau_i} \right\}.
$$

We have to use fresh crossings instead of arbitrary crossings to keep the sampling rate finite. The expected inter-sample times depend on the state at the beginning of the interval as well as the control policy. We shall assume that the levels in $\mathcal{L}$ as well as the control policy are such that the expected inter-sample times are finite and bounded away from zero. When the plant is unstable, this means that the levels in $\mathcal{L}$ go up to ∞ and $-\infty$.

4.4.1 Equivalent Discrete-time Markov chain

As with periodic sampling, the sequence

$$\left\{X_n \middle| n = 0, 1, 2, \ldots; X_n = x_{\tau_n}\right\}$$

forms a discrete-time controlled Markov chain. Here, it takes values in the finite set $\mathcal{L}$. As before, we will assume that the discrete-time control sequence $\{U_n\}$ is such that the resultant Markov chain is ergodic and also stable in the following sense:

$$\mathbb{E}\left[h^4\left(X\right)\right] < \infty,$$

where, $h : \mathcal{L} \to \mathbb{R}$ is defined by

$$h(l) = |l|.$$

Like in section 4.3.1, we can express the average quantities for the continuous time problem in terms of the ones for a related discrete-time controlled Markov chain $\{x_n\}$.

The average sampling rate is given by [39]:

$$R_{\mathcal{L},U} = \frac{1}{\sum_{l_i \in \mathcal{L}} \pi_U(l_i) \zeta\left(l_i, U(l_i)\right)},$$

where, $\zeta : \mathcal{L} \times \mathbb{R} \to \mathbb{R}^+$ is the expected first exit time defined by:

$$\zeta\left(l_i, U(l_i)\right) = \mathbb{E}\left[\tau_1 \left| \tau_0 = 0, x_0 = l_i \in \mathcal{L}, U_0 = U(l_i)\right.\right],$$

and $\{\pi_U(l_i)\}_{i=-\infty}^{\infty}$ is the steady state distribution for $\{x_n\}$.

The average stabilization error is given by [39]:

$$J_{\mathcal{L},U} = \frac{\sum_{l_i \in \mathcal{L}} \pi_U(l_i) g\left(l_i, U(l_i)\right)}{\sum_{l_i \in \mathcal{L}} \pi_U(l_i) \zeta\left(l_i, U(l_i)\right)},$$

where, $g : \mathcal{L} \times \mathbb{R} \to \mathbb{R}^+$ is defined by:

$$g(l) = \mathbb{E}\left[\int_0^{\tau_1} x_s^2 ds \middle| \tau_0 = 0, x_0 = l_i \in \mathcal{L}, U_0 = U(l_i)\right].$$

Both the numerator and the denominator in the last expression for the average stabilization error are dependent on U. But the denominator is just the reciprocal of the average sampling rate which is constrained. Define the cost $\tilde{J}_{\mathcal{L},U}$ by:

$$\tilde{J}_{\mathcal{L},U} = \sum_{l_i \in \mathcal{L}} \pi_U(l_i) g\left(l_i, U(l_i)\right).$$

Then, minimizing $J_{\mathcal{L},U}$ while respecting the constraint on the average sampling rate is the same as minimizing

$$\frac{\tilde{J}_{\mathcal{L},U}}{\sum_{l_i \in \mathcal{L}} \pi_U(l_i) \zeta\left(l_i, U(l_i)\right)}$$

while ensuring that

$$R_{\mathcal{L},U} = \frac{1}{\sum_{l_i \in \mathcal{L}} \pi_U(l_i) \zeta\left(l_i, U(l_i)\right)} \leq R_{\text{desired}}.$$

It is clear from the problem setup that under optimal U, both $\tilde{J}_{\mathcal{L},U}$ and $R_{\mathcal{L},U}$ are finite, positive and non-zero. If $\Gamma^* > 0$ is the minimum value of $J_{\mathcal{L},U}$, then, while respecting the constraint on average sampling rate (average inter-sample interval),

$$\begin{cases} \frac{\tilde{J}_{\mathcal{L},U}}{\sum_{l_i\in\mathcal{L}} \pi_U(l_i)\zeta(l_i,U(l_i))} \geq \Gamma^* > 0, \quad \forall\ U \in \mathscr{U} \\ \frac{\tilde{J}_{\mathcal{L},U^*}}{\sum_{l_i\in\mathcal{L}} \pi_{U^*}(l_i)\zeta(l_i,U^*(l_i))} = \Gamma^*. \end{cases}$$

$$\Leftrightarrow \begin{cases} \tilde{J}_{\mathcal{L},U} - \Gamma^* \sum_{l_i\in\mathcal{L}} \pi_U(l_i)\zeta\left(l_i, U(l_i)\right) \geq 0, \quad \forall\ U \in \mathscr{U} \\ \tilde{J}_{\mathcal{L},U^*} - \Gamma^* \sum_{l_i\in\mathcal{L}} \pi_{U^*}(l_i)\zeta\left(l_i, U^*(l_i)\right) = 0. \end{cases}$$

This means we only have to worry about minimizing $\tilde{J}_{\mathcal{L},U}$ subject to the sampling rate constraint. This is the same as minimizing the Lagrangian:

$$\tilde{J}_{\mathcal{L},U} - \Gamma \sum_{l_i\in\mathcal{L}} \pi_U(l_i)\zeta\left(l_i, U(l_i)\right).$$

Denote the second sum in the above Lagrangian, the average inter-sample time, by $S_{\mathcal{L},U}$.

We will now turn to the calculation of the transition probability kernel of $\{x_n\}$, and the average quantities $\tilde{J}_{\mathcal{L},U}$, $S_{\mathcal{L},U}$. To do so, we will appeal to the results of A. Because the state signal is scalar, there are only two possible transitions from any state in $\mathcal{L}$. The transition probabilities

$$p\left(l', l, U\right) = \mathbb{P}\left[X_{n+1} = l' \middle| X_n = l, U_n = U(l)\right], \quad \forall\ \left(l', l, U\right) \in \mathcal{L} \times \mathcal{L} \times \mathscr{U},$$

are found by solving an ODE [32]:

$$p\left(l_{i+1}, l_i, U\right) = \eta(l_i),$$

where $\eta(\cdot)$ satisfies:

$$(u + ax)\,\frac{d\eta}{dx} + \frac{1}{2}\frac{d^2\eta}{dx^2} = 0,$$

with the boundary conditions:

$$\eta\left(l_{i+1}\right)=1, \eta\left(l_{i-1}\right)=0.$$

Then we have $\forall\ l_i \in \mathcal{L}$:

$$p\left(l, l_i, U\right)=\begin{cases}\frac{\int_{l_{i-1}}^{l_i} e^{-2us-as^2} ds}{\int_{l_{i-1}}^{i+1} e^{-2ur-ar^2} dr} & \text{if } l=l_{i+1}, \\ \frac{\int_{l_i}^{i+1} e^{-2us-as^2} ds}{\int_{l_{i-1}}^{i+1} e^{-2ur-ar^2} dr} & \text{if } l=l_{i-1}, \\ 0 & \text{otherwise.}\end{cases}$$

The steady-state occupation measure π_U can be calculated using the equations:

$$\pi_U\left(l_i\right)=\pi_U\left(l_{i-1}\right) p\left(l_{i-1}, l_i, U\right)+\pi_U\left(l_{i+1}\right) p\left(l_{i+1}, l_i, U\right) \quad \forall\ l_i \in \mathcal{L}$$

The expected stabilization error starting at level l_i up to the first exit time out of $\left(l_{i-1}, l_{i+1}\right)$ is given by:

$$g\left(l_i, U\right)=\eta(l_i),$$

where $\eta(\cdot)$ satisfies the ODE:

$$(u+ax) \frac{d\eta}{dx}+\frac{1}{2} \frac{d^2\eta}{dx^2}=-x^2,$$

with the boundary conditions:

$$\eta\left(l_{i+1}\right)=0, \eta\left(l_{i-1}\right)=0.$$

Let

$$q_1(x)=\int_0^x e^{-az^2-2uz} dz, \quad \text{and,}$$

$$q_1(x)=\int_0^x e^{-az^2-2uz} \int_0^z y^2 e^{ay^2+2uy} dy dz.$$

Then we have $\forall\ l_i \in \mathcal{L}$:

$$g\left(l_i, U\right) = \frac{q_1\left(l_i\right) q_2\left(l_{i+1}\right) - q_1\left(l_i\right) q_2\left(l_{i-1}\right) - q_1\left(l_{i+1}\right) q_2\left(l_i\right)}{q_1\left(l_{i+1}\right) - q_1\left(l_{i-1}\right)} + \frac{q_1\left(l_{i-1}\right) q_2\left(l_i\right) - q_1\left(l_{i+1}\right) q_2\left(l_{i+1}\right) + q_1\left(l_{i+1}\right) q_2\left(l_{i-1}\right)}{q_1\left(l_{i+1}\right) - q_1\left(l_{i-1}\right)}.$$

The expected first exit time $\mathcal{E}\left(l_i, U\right)$ is given by :

$$\mathcal{E}\left(l_i, U\right) = \eta(l_i),$$

where $\eta(\cdot)$ satisfies the ODE:

$$\left(u + ax\right) \frac{d\eta}{dx} + \frac{1}{2} \frac{d^2\eta}{dx^2} = -1,$$

with the boundary conditions:

$$\eta\left(l_{i+1}\right) = 0, \eta\left(l_{i-1}\right) = 0.$$

Let

$$q_3(x) = \int_0^x e^{-az^2 - 2uz} \int_0^z e^{ay^2 + 2uy} dydz.$$

Then we have $\forall\ l_i \in \mathcal{L}$:

$$\mathcal{E}\left(l_i, U\right) = \frac{q_1\left(l_i\right) q_3\left(l_{i+1}\right) - q_1\left(l_i\right) q_3\left(l_{i-1}\right) - q_1\left(l_{i+1}\right) q_3\left(l_i\right)}{q_1\left(l_{i+1}\right) - q_1\left(l_{i-1}\right)} + \frac{q_1\left(l_{i-1}\right) q_3\left(l_i\right) - q_1\left(l_{i+1}\right) q_3\left(l_{i+1}\right) + q_1\left(l_{i+1}\right) q_3\left(l_{i-1}\right)}{q_1\left(l_{i+1}\right) - q_1\left(l_{i-1}\right)}.$$

4.4.2 Existence of Optimal Controls and their Computation

The Markov chain $\{X_n\}$ has the property that, independent of U, only a finite number of elements of $\mathcal{L}$ can be reached from any member of $\mathcal{L}$ in one step. The per stage cost in the average cost formulation is an unbounded function of the state. For such situations, Borkar [42] shows the existence of optimal (non-randomized) stationary policies and proves the validity of the Average Cost optimality equations:

$$\alpha^* = \inf_{u_i \in \mathbb{R}} \left\{ g\left(l_i, u_i\right) - \Gamma \mathcal{E}\left(l_i, u_i\right) - v_i + p\left(l_{i+1}, l_i, u_i\right) v_{i+1} + p\left(l_{i-1}, l_i, u_i\right) v_{i-1} \right\}$$

$$\forall\ l_i \in \mathcal{L}$$

We will use value iteration based on the above equations to determine the optimal controls and their performance for fixed $\mathcal{L}$. We will next consider some natural classes of level-triggered sampling schemes.

4.5 Comparisons

We will consider two level-triggered sampling schemes. One will be the Lattice-triggered sampling scheme. Let

$$\mathsf{Latt}_0 = \left\{\ldots, -2\kappa, -\kappa, 0, \kappa, 2\kappa, \ldots\right\}.$$

Choosing $\mathcal{L}$ to be be Latt_0 gives a set of equi-spaced levels. Choosing $\mathcal{L}$ to be

$$\mathsf{Latt}_1 = \left\{\ldots, -2\kappa, -\kappa, \kappa, 2\kappa, \ldots\right\},$$

which does not have zero as a level leads to a variant of the equi-spaced set.

On the other hand, choosing $\mathcal{L}$ to be

$$\mathsf{Log}_0 = \{\ldots, -\log(1+2\kappa), -\log(1+\kappa), 0, \log(1+\kappa), \log(1+2\kappa), \ldots\}$$

gives us a logarithmic set of levels and choosing $\mathcal{L}$ to be

$$\mathsf{Log}_1 = \{\ldots, -\log(1+2\kappa), -\log(1+\kappa), \log(1+\kappa), \log(1+2\kappa), \ldots\}$$

gives us a variant.

Below(fig.4.3), we have sketched the performances of level-triggered schemes with these levels as well as the periodic sampling scheme:

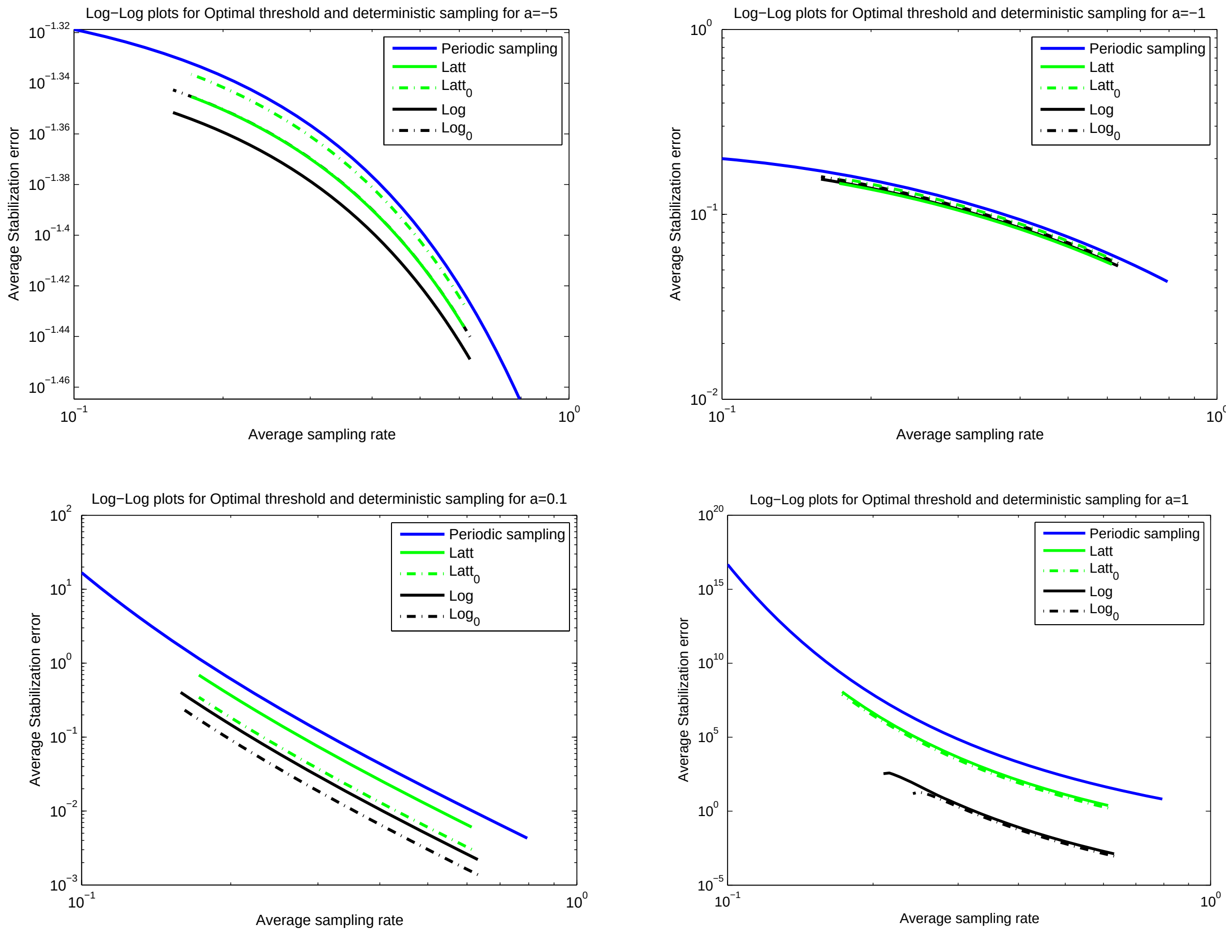

Figure 4.3: Relative performance of Threshold and Periodic sampling as a function of initial variance σ^2 and parameter values (a) $a = 1$, (b) $a = 10$, (c) $a = -1$ and (d) $a = -10$.

Remark: In the optimal control problems dealt so far, we could relax the restriction to controls that render the system ergodic with a finite steady state fourth moment. But doing so forces us to modify the state space. The new Markov chain to be worked with has the state:

$$Z_n = \begin{pmatrix} X_n \\ \\ Y_n \end{pmatrix} = \begin{pmatrix} X_n \\ \\ X_{n-1} \end{pmatrix}, \quad \forall\, n \geq 1, \quad S_0 = \begin{pmatrix} X_0 \\ \\ x \end{pmatrix},$$

where, x is any valid element of the chain's state space that can have X_0 as a successor state. Now the expected integral running cost over the inter-sample interval $[\tau_{n-1}, \tau_n)$ is purely a function of Z_n. However, the computation of the parameters of the Markov chain and the solution of the average cost control problem are more involved.

Chapter 5

Sampling in Teams for Sequential Hypothesis Testing

5.1 Event-triggered sampling in a team of sensors

We have a team of sensors deployed to gather measurements for a Sequential Binary Hypothesis Testing problem in continuous time. For ease of exposition, we will restrict the size of the team to be two. There is also a supervisor to whom the two agents communicate over a common medium. Under the null hypothesis, each sensor observes independent Wiener processes. Under the other hypothesis, each sensor observes the Wiener process with a constant drift. The objective of the supervisor is to pick a hypothesis in reasonable time. To capture this objective, we will use a Bayesian performance measure consisting of a sum of the probability of error in deciding for a hypothesis and the expected time taken to make a decision.

The crux of the problem is a stringent constraint on communications from the sensors to the supervisor. Each sensor is permitted to transmit exactly one data packet to the supervisor. Because the packet link is a medium shared between the two sensors, a packet transmitted by either sensor will be heard by the other. This link has the idealized property that it delivers data packets reliably and with negligible delay. The packets are also of sufficient bit-width so that quantization noise of the low dimensional variables represented in them can be ignored. We will discuss subsequently what kind of variables will need to be packetized. From that

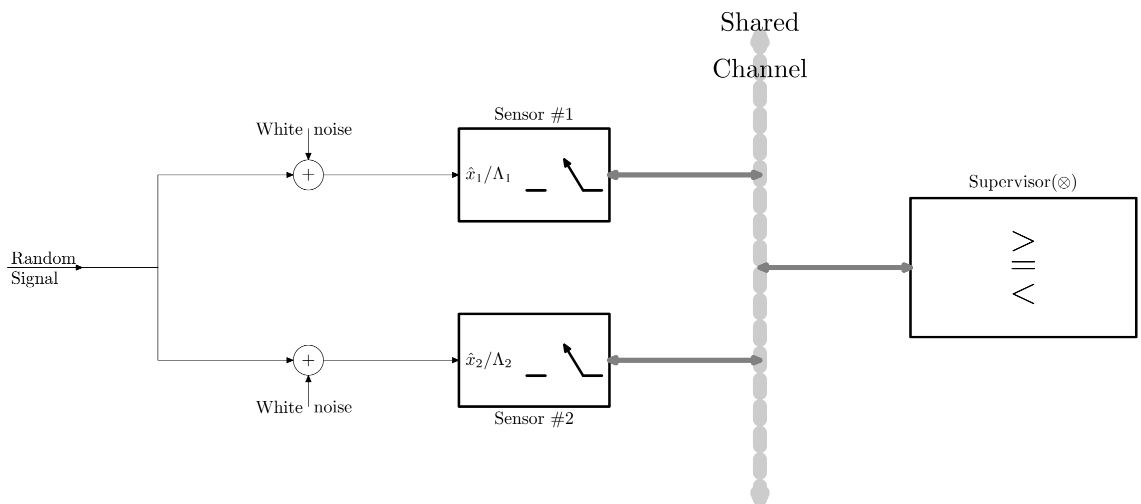

Figure 5.1: The two-sensor sequential detection set-up.

discussion, it will be clear the we need to packetize only likelihoods or likelihood ratios. In the remainder of this section, we will assume that the payloads for data packets will be likelihood ratios.

The information available to the supervisor to make its decision vas well as to choose the decision time consists of the two received data packets with their time-stamps. The ability of each sensor to 'listen' to the packet transmitted by the other provides scope for cooperation in transmission of the packets.

The best time for the supervisor to make a decision is immediately after the second packet has arrived. If any earlier time were to be optimal, we could find a new sampling policy that ensures that the second packet arrives before that time and perform no worse. Also, the sensor that transmits last has the information privy to the supervisor and more. So, a decision made by the supervisor at the time of reception of the second packet can be no better than one reached by the sensor that transmits last. However, we will show that the best decision that can be made at the latter sensor is also based on the likelihood ratio, which is what is available to the supervisor as well.

Our contribution through this work is to describe some natural event-triggered sampling strategies and to graph their relative performances. There are four event-triggered sampling strategies and one simple deterministic sampling strategy in this study.

5.1.1 Related Works

Multi-agent decentralized sequential detection with information constraints is usually cast as a single-shot hypothesis testing problem with constraints on the amount of information shipped from each sensor to the supervisor (fusion-center) [43, 44]. In such problems times of transmission of information from sensors is not subject to design. Instead, it is prescribed to be a periodic sequence. The time to be chosen is a decision time at the supervisor, For the sensors, the emphasis is on how to best satisfy the quantization constraint at each transmission time.

The papers [45, 46] treat a two agent sequential detection problem close in spirit to ours. A key difference in our problem setup, one that is appropriate for the Networked Decision problems arising in sensor networks, is the listening among sensors and cooperative communication [7]. Another difference is in the way the decision performance is measured. In our work, the expectation of the true delay is what is penalized as opposed to a sum of the expectations of the two transmit times.

The use of a common communication medium is typical in sensor networks and in Networked Control/Monitoring systems. Some motivation for this two agent problem comes from a video surveillance using a couple of cameras or acoustic monitoring vehicle sounds using a couple of microphone nodes. The signal model we use is perhaps more appropriate for the microphone outputs. On the other hand, we can use telegraph signal-like models for the output for a video processing and object recognition system. The qualitative results we obtain for our simple signal model would assist us in the case of Markov chain models such as the random telegraph signal.

5.2 The Optimal Sampling Problem

On the probability space $(\Omega, \mathcal{F}, \mathbb{P})$ there are defined a $\{0,1\}$-valued random variable H and two independent standard Wiener processes V_t and W_t. The hypothesis H is described statistically:

$$H = \begin{cases} 1 & \text{with probability} \quad \pi \\ 0 & \text{with probability} \quad 1-\pi \end{cases}$$

The two sensors, labelled as sensor #1 and sensor #2, have observation processes $\{x_t\}$ and $\{y_t\}$ respectively. Under the null hypothesis ($H = 0$), these observations are just the Wiener processes with some scaling.

$$dx_t = \sigma_1 dV_t, \qquad x_0 = 0, \quad \text{and,} \tag{5.1}$$

$$dy_t = \sigma_2 dW_t, \qquad y_0 = 0. \tag{5.2}$$

Under the other hypothesis, the sensor observations obey the equations for Brownian motion with unit drift:

$$dx_t = dt + \sigma_1 dV_t, \qquad x_0 = 0, \quad \text{and}, \tag{5.3}$$

$$dy_t = dt + \sigma_2 dW_t, \qquad y_0 = 0. \tag{5.4}$$

Let τ_1, τ_2 be the packet transmit times of the sensors # 1, 2 respectively. Let F be the earlier of the two times and S the latter.

$$F = \min\{\tau_1, \tau_2\}, \tag{5.5}$$

$$S = \max\{\tau_1, \tau_2\}. \tag{5.6}$$

Let e, l denote the indices of the sensors that transmits earlier and later respectively. We will discuss ties later on. In fact, ties for choice of e do no present any conceptual issue at all.

$$e, l \in \{1, 2\}.$$

5.2.1 The Likelihood ratio processes

Let $\{\lambda_{i,t}\}$ denote the local likelihood ratio process computed at sensor # i.

$$\lambda_{i,t} = \frac{d\mathbb{P}_{H=1}}{d\mathbb{P}_{H=0}} \left(\mathcal{F}_t^{\text{Information at sensor}\#i}\right),$$

where $\mathbb{P}_{H=1}$ is the conditional probability measure given that hypothesis 1 is true. Because of the possible information sharing, the likelihood ratio process at one sensor could be altered after the other has transmitted its packet. At sensor # 1, its likelihood ratio process $\{\lambda_{1,t}\}$ obeys over the intervals $[0, F)$ and (F, S) the SDE[21]:

$$d\lambda_{1,t} = \frac{1}{\sigma_1^2}\lambda_{1,t} dx_t. \tag{5.7}$$

At $t = 0$, we have the initial condition: $\lambda_{1,0} = 1$. At time F^+ we have the modification:

$$\lambda_{1,F^+} = \lambda_{1,F^-} \times \left(\mathbf{1}_{\{e=1\}} + \lambda_{2,F^-} \cdot \mathbf{1}_{\{e=2\}}\right). \tag{5.8}$$

This gives the expressions:

$$\lambda_{1,t} = \begin{cases} \exp\left\{\frac{1}{\sigma_1^2}\left(x_t - \frac{t}{2}\right)\right\} & \text{for } \ t \in [0, F), \\ \exp\left\{\frac{1}{\sigma_1^2}\left(x_t - \frac{t}{2}\right)\right\} & \text{for } \ t \in [F, S) \ \text{ if } \ e = 1, \\ \exp\left\{\frac{1}{\sigma_1^2}\left(x_t - \frac{t}{2}\right) + \frac{1}{\sigma_2^2}\left(y_F - \frac{F}{2}\right)\right\} & \text{for } \ t \in [F, S) \ \text{ if } \ e = 2. \end{cases}$$

We have a similar set of expressions for $\lambda_{1,t}$.

Likelihood ratio used by Supervisor

The likelihood ratio(λ_{Sup}) used by the supervisor is just that of the sensor that transmits last, at it transit time.

$$\begin{aligned} \lambda_{Sup} &= \lambda_{e,S} \\ &= \begin{cases} \exp\left\{\frac{1}{\sigma_1^2}\left(x_F - \frac{F}{2}\right)\right\} \times \exp\left\{\frac{1}{\sigma_2^2}\left(y_S - \frac{S}{2}\right)\right\} & \text{if } \ e = 1, \\ \exp\left\{\frac{1}{\sigma_1^2}\left(x_S - \frac{S}{2}\right)\right\} \times \exp\left\{\frac{1}{\sigma_2^2}\left(y_F - \frac{F}{2}\right)\right\} & \text{if } \ e = 2. \end{cases} \end{aligned} \tag{5.9}$$

Suppose for a moment that there was no sharing of samples amongst the two sensors. Then the likelihood ratio used by the supervisor will still be given by the product expression of eqn. (5.9). The sharing of samples does not change the form of the sufficient statistic for the supervisor's decision. However, it provides more information at the disposal of the sensor that transmits last. That sensor can use the extra information to choose its transmit time so as to improve the efficiency of the team.

5.2.2 Sampling strategies allowed

The signal model and the prior probabilities are known to the two sensors as well as to the supervisor. This means that the values of the parameters π, σ_1, σ_2 are available at all three agents. In addition, each sensor and the supervisor is also aware of the policy employed by the other sensor to choose the packet transmit time.

The measured information available at a sensor consists of the sensed observations and any data heard over the common communication medium. So, a sensor's decision at a time instant to transmit or to wait is made using different information depending on whether or not the other sensor has already transmitted. If the other sensor has not transmitted yet, the decision is based on the local observations and the fact that the other sensor has not transmitted yet. If the other sensor has already transmitted, the decision is made based on the local observations as well as the likelihood ratio transmitted by the other sensor and the time-stamp of that transmission. Let $\{\nabla_t\}$ be the random process defined through:

$$\nabla_t = \mathbf{1}_{\{t \geq F\}}.$$

Then, τ_1 is a stopping time with respect to the filtration

$$\mathcal{F}_t^{\left(x_t,\ \nabla_t,\ \nabla_t \cdot \lambda_{e,F},\ \nabla_t \cdot F \right)}.$$

Similarly, τ_2 is a stopping time with respect to the filtration

$$\mathcal{F}_t^{\left(y_t,\ \nabla_t,\ \nabla_t \cdot \lambda_{e,F},\ \nabla_t \cdot F \right)}.$$

5.2.3 Detection performance

Given sampling policies $\mathcal{SP}_1$, $\mathcal{SP}_2$ at the two sensors, the detection performance is measured through a linear combination of the expected time taken by the supervisor to arrive at a decision $(\hat{H})$ and the probability of error in that decision.

$$J\left(\mathcal{SP}_1, \mathcal{SP}_2\right) = \mathbb{E}\left[cS\right] + \mathbb{P}\left[\hat{H} \neq H\right],$$

where c is a positive constant.

The best decision by the supervisor at time S can be no better than that of the sensor that transmits second. Using the line of argument of Lemma 1 in chapter 4 (page 160) of [21], it can be shown that the optimal decision at the latter sensor is dependent only on its likelihood ratio at the time S. In fact, once the first packet has been transmitted, the optimal policy of the latter sensor is to implement a SPRT-like sampling policy to determine its transmit time and decision.

Before we seek good event-triggered sampling strategies, we will note down the performance provided by deterministic ones.

Deterministic sampling

When the transmit times are to be chosen deterministically, the best strategy is to choose the same time for both sensors. This is because the optimal decision is based on the following likelihood ratio at time S:

$$\lambda_{Sup} = \exp\left\{\frac{1}{\sigma_1^2}\left(x_{\tau_1} - \frac{\tau_1}{2}\right)\right\} \times \exp\left\{\frac{1}{\sigma_2^2}\left(y_{\tau_2} - \frac{\tau_2}{2}\right)\right\}.$$

The above product is the same as what we get when each sensor ignores the other's transmission and merely relays its likelihood ratio to the supervisor. The supervisor just has to form the product of the two likelihood ratios and uses the threshold rule:

$$\hat{H} = \begin{cases} 1 & \text{if } \lambda_{Sup} \geq \frac{1-\pi}{\pi}, \\ \\ 0 & \text{if } \lambda_{Sup} \geq \frac{1-\pi}{\pi}. \end{cases}$$

Intuitively, by increasing the earlier transmit time to equal the other, we can only improve the SNR of the earlier likelihood ratio because it is now based on longer measurement. So, the optimal deterministic strategy consists of using the same deterministic sampling policy $\mathcal{DP}(T)$ which when applied to sensor # i produces:

$$\tau_i = \Delta.$$

Thus the expected decision time is Δ. The resulting detection error is:

$$\frac{1-\pi}{2}\left\{1 - \operatorname{erf}\left(\frac{\Delta\left(\sigma_1^2+\sigma_2^2\right) + 2\sigma_1^2\sigma_2^2\ln\left(\frac{1-\pi}{\pi}\right)}{2\sqrt{2\Delta\sigma_1^2\sigma_2^2\left(\sigma_1^2+\sigma_2^2\right)}}\right)\right\} + \frac{\pi}{2}\left\{1 - \operatorname{erf}\left(\frac{\Delta\left(\sigma_1^2+\sigma_2^2\right) - 2\sigma_1^2\sigma_2^2\ln\left(\frac{1-\pi}{\pi}\right)}{2\sqrt{2\Delta\sigma_1^2\sigma_2^2\left(\sigma_1^2+\sigma_2^2\right)}}\right)\right\}.$$

The expression simplifies when we specialize to the case $\pi = \frac{1}{2}$. We will use this specialization in the next section to gain a lot of mileage in our calculations. The case of unequal priors is not any more difficult conceptually.

5.3 Threshold sampling policies

By a threshold strategy, we mean a set of sampling policies where both sampling times are chosen as threshold-crossing times for likelihood ratio process. Let us introduce some notation which will be useful in describing the performances of the threshold sampling policies.

5.3.1 Useful notation

Assume for a moment that sensor #1 acts in isolation without any information from sensor #2. Given $\lambda_1 > 1$, let the time for sampling the sensor's likelihood ratio be the first exit time from the interval $(\frac{1}{\lambda_1}, \lambda_1)$. Assume also that its likelihood ratio process at time zero is not 1 but an arbitrary λ_0 such that:

$$\frac{1}{\lambda_1} < \lambda_0 < \lambda_1.$$

Then following the calculations of [21] the expected value of the sampling time and the expected probability of error at that sampling time are:

$$\begin{aligned}\mathbb{E}\left[\tau(\lambda_0)\right] &= \frac{1}{2}\left[\sigma_1^2 \ln\left(\lambda_0\right) - \frac{4\sigma_1^2\lambda_1\lambda_0 \ln\lambda_1}{\lambda_1^2 - 1} + \frac{2\sigma_1^2(\lambda_1^2+1)\ln\lambda_1}{\lambda_1^2-1}\right] \\ &\quad + \frac{1}{2}\left[-\sigma_1^2 \ln\left(\lambda_0\right) - \frac{4\sigma_1^2\lambda_1 \ln\lambda_1}{(\lambda_1^2 - 1)\lambda_0} + \frac{2\sigma_1^2(\lambda_1^2+1)\ln\lambda_1}{\lambda_1^2-1}\right] \triangleq TTS_{\lambda_1}\left(\lambda_0\right),\end{aligned} \tag{5.10}$$

$$\mathbb{E}\left[\tau(\lambda_0)\right] = \frac{1}{2}\left[\frac{\lambda_1\lambda_0 - 1}{\lambda_1^2 - 1}\right] + \frac{1}{2}\left[\frac{\lambda_1 - \lambda_0}{(\lambda_1^2-1)\lambda_0}\right] \triangleq PErr_{\lambda_1}\left(\lambda_0\right). \tag{5.11}$$

We will describe four different threshold sampling policies starting with the simplest generalization of the deterministic sampling policy $\mathcal{DP}$.

5.3.2 Synchronous threshold sampling

For $i \in \{1, 2\}$, sensor # i is given two threshold values: α_i, β_i such that

$$0 \leq \alpha_i \leq 1 \leq \beta_i. \tag{5.12}$$

In the synchronized threshold sampling policy($\mathcal{STP}$), the two transmissions are synchronized forcibly. The transmissions are triggered by the sensor # i whose likelihood ratio $\lambda_{i,t}$ exceeds the interval (α_i, β_i) first. Immediately after, the other sensor is forced to transmit its likelihood for hypothesis 1 as well. In formal terms,

$$\tau_1 = \min\left\{\tau_2^+, \inf\left\{s|\lambda_{1,s} \notin (\alpha_1, \beta_1)\right\}\right\},$$
$$\tau_2 = \min\left\{\tau_1^+, \inf\left\{s|\lambda_{2,s} \notin (\alpha_2, \beta_2)\right\}\right\}.$$

Let $\tilde{\lambda}_{i,t}$ denote the likelihood ratio at sensor # i based solely upon the observations process at that sensor and ignoring any packets transmitted by the other sensor. This quantity obeys an SDE identical to equation(5.7) over the entire time horizon: $[0, \infty)$. Define the first exit times $\tilde{\tau}_1$, $\tilde{\tau}_2$ through:

$$\tilde{\tau}_1 = \inf\left\{s|\tilde{\lambda}_{1,s} \notin (\alpha_1, \beta_1)\right\},$$
$$\tilde{\tau}_2 = \inf\left\{s|\tilde{\lambda}_{2,s} \notin (\alpha_2, \beta_2)\right\}.$$

The common sampling time $\tau_{\mathcal{STP}}$ is nothing else than $\min\{\tilde{\tau}_1, \tilde{\tau}_2\}$. The decision is based on the product of the two likelihood ratios. The supervisor computes

$$\lambda_{Sup} = \exp\left\{\frac{1}{\sigma_1^2}\left(x_{\tau_1} - \frac{\tau_1}{2}\right)\right\} \times \exp\left\{\frac{1}{\sigma_2^2}\left(y_{\tau_2} - \frac{\tau_2}{2}\right)\right\},$$

and uses the threshold rule:

$$\hat{H} = \begin{cases} 1 & \text{if } \lambda_{Sup} \geq \frac{1-\pi}{\pi} \quad = 1, \\ 0 & \text{if } \lambda_{Sup} \geq \frac{1-\pi}{\pi} \quad 1. \end{cases}$$

Note that since we have assumed the prior probabilities to be equal, we can choose for $i \in \{1, 2\}$:

$$\beta_i = \frac{1}{\alpha_i}, \tag{5.13}$$

Performance computation

The performance of the scheme is computed through solving some PDEs. The expected sampling time is give through:

$$\mathbb{E}\left[\tau_{\mathcal{STP}}\right] = \frac{1}{2}f(1,1) + \frac{1}{2}g(1,1),$$

where,

$$\begin{aligned}
-1 &= \frac{x^2}{2\sigma_1^2}\frac{\partial^2 f}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 f}{\partial y^2},\\
-1 &= \frac{x}{\sigma_1^2}\frac{\partial g}{\partial x} + \frac{y}{\sigma_2^2}\frac{\partial g}{\partial y} + \frac{x^2}{2\sigma_1^2}\frac{\partial^2 g}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 g}{\partial y^2},
\end{aligned}$$

with both the PDEs satisfying the boundary conditions:

$$f(x,y) = g(x,y) = 0, \quad \text{on the boundary of } \left[\alpha_1, \frac{1}{\alpha_1}\right] \times \left[\alpha_2, \frac{1}{\alpha_2}\right].$$

Similarly, the probability of error is computed as:

$$\mathbb{P}\left[\hat{H} \neq H\right] = \frac{1}{2}f(1,1) + \frac{1}{2}\left[1 - g(1,1)\right],$$

where,

$$\begin{aligned}
0 &= \frac{x^2}{2\sigma_1^2}\frac{\partial^2 f}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 f}{\partial y^2},\\
0 &= \frac{x}{\sigma_1^2}\frac{\partial g}{\partial x} + \frac{y}{\sigma_2^2}\frac{\partial g}{\partial y} + \frac{x^2}{2\sigma_1^2}\frac{\partial^2 g}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 g}{\partial y^2},
\end{aligned}$$

with both the PDEs satisfying the boundary conditions:

$$f(x,y) = g(x,y) = \mathbf{1}_{\{xy>1\}}, \quad \text{on the boundary of } \left[\alpha_1, \frac{1}{\alpha_1}\right] \times \left[\alpha_2, \frac{1}{\alpha_2}\right].$$

5.3.3 Tandem Threshold sampling

We now consider the set of strategies where the sensors transmit in a fixed sequence with the inter sampling times being chosen through threshold crossings. First the sensor slated to go early samples according to a threshold crossing time and passes its sample of the likelihood ratio to the other supervisor (and the other sensor). The sensor transmitting second modifies its likelihood ratio using the packet from its neighbour and samples next based on a different set of thresholds. Without loss of generality, let us assume that the first sensor has observations with higher SNR:

$$\sigma_1^2 \leq \sigma_2^2.$$

Then, we have two possible orders for the sequence of transmissions. In one (name it $\mathcal{CTANDP}$), we let the coarser sensor (the one with higher σ^2) sample first and pass its likelihood ratio to the other sensor. In the other (call it $\mathcal{FTANDP}$), let the finer sensor sample first. As in the previous section (eqns.(5.12,5.13)) let the sensors $\#1, 2$ choose α_1, α_2 respectively.

Calculations of peformances

Without any loss of generality, assume that $\sigma_2^2 \leq \sigma_1^2$. Under $\mathcal{CTANDP}$, we have:

$$
\begin{aligned}
e &= 1, \quad l \ = 2, \\
F &= \inf\left\{ t \geq 0 \,\middle|\, \lambda_{1,t} \notin \left(\alpha_2, \frac{1}{\alpha_2}\right) \right\} \\
S &= \inf\left\{ t \geq F \,\middle|\, \lambda_{1,t} \notin \left(\alpha_1, \frac{1}{\alpha_1}\right) \right\}
\end{aligned}
$$

The the performance of $\mathcal{CTANDP}$ is obtained by solving PDEs boundary value problems which are more complicated the ones we have met before. We have to make use of the notation in eqns. (5.10,5.11). The expected sampling time is given by:

$$
\mathbb{E}\left[S_{\mathcal{CTANDP}}\right] = \frac{1}{2} f(1,1) + \frac{1}{2} g(1,1),
$$

where,

$$
\begin{aligned}
-1 &= \frac{x^2}{2\sigma_1^2}\frac{\partial^2 f}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 f}{\partial y^2}, \\
-1 &= \frac{x}{\sigma_1^2}\frac{\partial g}{\partial x} + \frac{y}{\sigma_2^2}\frac{\partial g}{\partial y} + \frac{x^2}{2\sigma_1^2}\frac{\partial^2 g}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 g}{\partial y^2},
\end{aligned}
$$

with both the PDEs satisfying the boundary conditions:

$$
f(x,y) = g(x,y) = TTS_{\frac{1}{\alpha_1}}(xy)\mathbf{1}_{\left\{\alpha_1 < xy < \frac{1}{\alpha_1}\right\}},
$$
$$
\text{on the boundary of } (0, +\infty) \times \left[\alpha_2, \frac{1}{\alpha_2}\right].
$$

Similarly, the probability of error is given by:

$$
\mathbb{P}\left[\hat{H} \neq H\right] = \frac{1}{2} f(1,1) + \frac{1}{2}\left[1 - g(1,1)\right],
$$

where,

$$0 = \frac{x^2}{2\sigma_1^2}\frac{\partial^2 f}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 f}{\partial y^2},$$

$$0 = \frac{x}{\sigma_1^2}\frac{\partial g}{\partial x} + \frac{y}{\sigma_2^2}\frac{\partial g}{\partial y} + \frac{x^2}{2\sigma_1^2}\frac{\partial^2 g}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 g}{\partial y^2},$$

with both the PDEs satisfying the boundary conditions:

$$f(x,y) = g(x,y) = PErr_{\frac{1}{\alpha_1}}(xy)\mathbf{1}_{\left\{\alpha_1 < xy < \frac{1}{\alpha_1}\right\}} + \mathbf{1}_{\left\{xy>1, xy \notin \left(\alpha_1, \frac{1}{\alpha_1}\right)\right\}},$$

$$\text{on the boundary of}\,(0,+\infty) \times \left[\alpha_2, \frac{1}{\alpha_2}\right].$$

The performance of $\mathcal{FTANDP}$ is computed by the same calculation as above but with the roles of the two sensors reversed.

5.3.4 Optimal Threshold Sampling

In the optimal threshold sampling scheme, we have no control over the order of sampling. Lets name this asynchronous sampling scheme $\mathcal{ASP}$. Each sensor has to choose two sets of thresholds. $(\alpha_i, \frac{1}{\alpha_i})$ is for determining the first sampling time and $(\mu_i, \frac{1}{\mu_i})$ is for determining the second sampling time. Define the first exit times $\tilde{\tau}_1$, $\tilde{\tau}_2$ through:

$$\tilde{\tau}_1 = \inf\left\{s \,\middle|\, \tilde{\lambda}_{1,s} \notin (\alpha_1, \frac{1}{\alpha_1})\right\},$$

$$\tilde{\tau}_2 = \inf\left\{s \,\middle|\, \tilde{\lambda}_{2,s} \notin (\alpha_2, \frac{1}{\alpha_2})\right\}.$$

We have:

$$F = \min\{\tilde{\tau}_1, \tilde{\tau}_2\},$$

$$e = \arg\min\{i \in \{1,2\} | \, \tilde{\tau}_i\},$$

$$S = \inf\left\{t \geq F \,\middle|\, \lambda_{l,t} \notin \left(\mu_l, \frac{1}{\mu_l}\right)\right\}.$$

Then the performance of $\mathcal{ASP}$ can be computed exactly like for the tandem policies.

The expected sampling time is given by:

$$\mathbb{E}\left[S_{\mathcal{ASP}}\right] = \frac{1}{2}f(1,1) + \frac{1}{2}g(1,1),$$

where,

$$-1 = \frac{x^2}{2\sigma_1^2}\frac{\partial^2 f}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 f}{\partial y^2},$$

$$-1 = \frac{x}{\sigma_1^2}\frac{\partial g}{\partial x} + \frac{y}{\sigma_2^2}\frac{\partial g}{\partial y} + \frac{x^2}{2\sigma_1^2}\frac{\partial^2 g}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 g}{\partial y^2},$$

with both the PDEs satisfying the boundary conditions:

$$f(x,y) = g(x,y) = \left\{TTS_{\frac{1}{\mu_2}}(xy)\mathbf{1}_{\left\{xy\in\left(\mu_2,\frac{1}{\mu_2}\right), x\notin\left(\alpha_1,\frac{1}{\alpha_1}\right)\right\}}\right\}$$
$$+\left\{TTS_{\frac{1}{\mu_1}}(xy)\mathbf{1}_{\left\{xy\in\left(\mu_1,\frac{1}{\mu_1}\right), y\notin\left(\alpha_2,\frac{1}{\alpha_2}\right)\right\}}\right\},$$
$$\text{on the boundary of } \left[\alpha_1, \frac{1}{\alpha_1}\right] \times \left[\alpha_2, \frac{1}{\alpha_2}\right].$$

Similarly, the probability of error is given by:

$$\mathbb{P}\left[\hat{H} \neq H\right] = \frac{1}{2}f(1,1) + \frac{1}{2}\left[1 - g(1,1)\right],$$

where,

$$0 = \frac{x^2}{2\sigma_1^2}\frac{\partial^2 f}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 f}{\partial y^2},$$

$$0 = \frac{x}{\sigma_1^2}\frac{\partial g}{\partial x} + \frac{y}{\sigma_2^2}\frac{\partial g}{\partial y} + \frac{x^2}{2\sigma_1^2}\frac{\partial^2 g}{\partial x^2} + \frac{y^2}{2\sigma_2^2}\frac{\partial^2 g}{\partial y^2},$$

with both the PDEs satisfying the boundary conditions:

$$f(x,y) = g(x,y) = \left\{ PErr_{\frac{1}{\mu_2}}(xy)\mathbf{1}_{\left\{xy\in\left(\mu_2,\frac{1}{\mu_2}\right),x\notin\left(\alpha_1,\frac{1}{\alpha_1}\right)\right\}} \right\}$$
$$+ \left\{ \mathbf{1}_{\left\{xy>1,xy\notin\left(\mu_2,\frac{1}{\mu_2}\right),x\notin\left(\alpha_1,\frac{1}{\alpha_1}\right)\right\}} \right\}$$
$$+ \left\{ PErr_{\frac{1}{\mu_1}}(xy)\mathbf{1}_{\left\{xy\in\left(\mu_1,\frac{1}{\mu_1}\right),y\notin\left(\alpha_2,\frac{1}{\alpha_2}\right)\right\}} \right\}$$
$$+ \left\{ \mathbf{1}_{\left\{xy>1,xy\notin\left(\mu_1,\frac{1}{\mu_1}\right),y\notin\left(\alpha_2,\frac{1}{\alpha_2}\right)\right\}} \right\},$$
$$\text{on the boundary of} \left[\alpha_1, \frac{1}{\alpha_1}\right] \times \left[\alpha_2, \frac{1}{\alpha_2}\right].$$

$$f(x,y) = g(x,y) = PErr_{\frac{1}{\alpha_1}}(xy)\mathbf{1}_{\left\{\alpha_1<xy<\frac{1}{\alpha_1}\right\}} + \mathbf{1}_{\left\{xy>1,xy\notin\left(\alpha_1,\frac{1}{\alpha_1}\right)\right\}},$$
$$\text{on the boundary of}\,(0,+\infty) \times \left[\alpha_2, \frac{1}{\alpha_2}\right].$$

We should note the none of the threshold sampling policies are person-by-person optimal in the parlance of team theory [45, 46].

5.4 Relative Performances and Conclusion

Here, we present graphs (fig:5.2) detailing the relative performances of the various sampling schemes.

Even though we have a sequential detection problem in infinite horizon, the optimal sampling schemes are not threshold-triggered. The sharing of data packets brings a time dependence on the cost structure. Loosely speaking, from the point of view a sensor in this team, the transmission by the other sensor is an event which affects the overall cost through the time of its occurrence. That is why the asynchronous sampling scheme $\mathcal{ASP}$ does not even possess the person-by-person

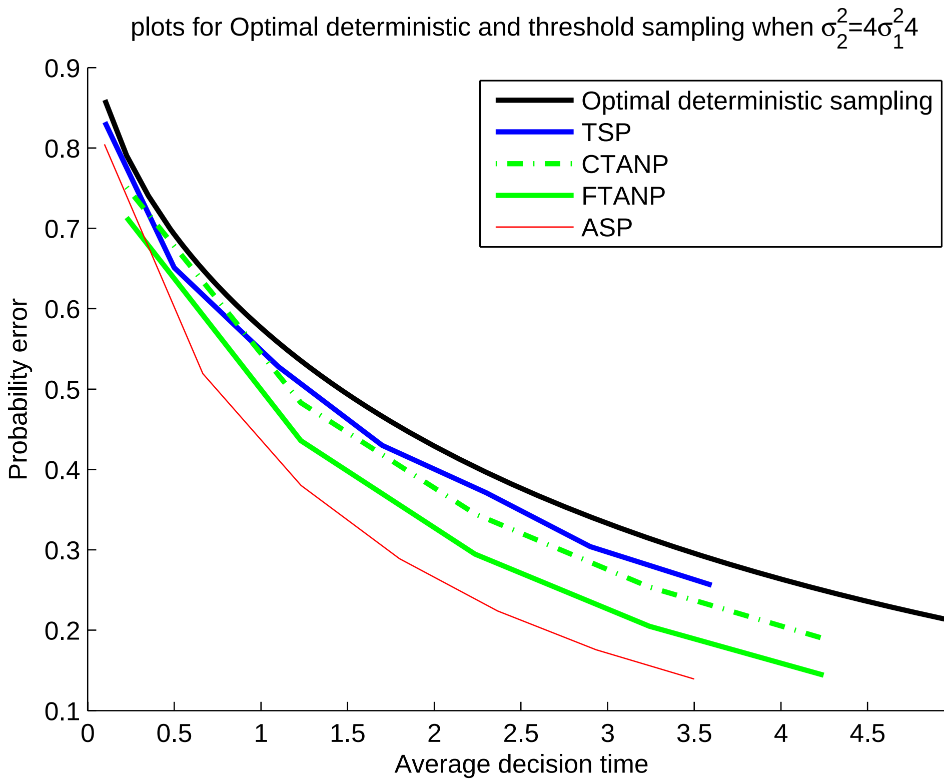

Figure 5.2: Plots of Probability of error vs. expected time taken to decide when the priors are equal. The asynchronous scheme performs well but the other threshold schemes are not clearly distinguishable.

optimality property it does in the context of [45, 46].

Chapter 6

Conclusions

This thesis provides solutions to some key design problems for event-triggered control. The design of good event-triggered sampling schemes answers the question: *What events should be reported and acted on for good control performance ?* On the other hand it also extends the work of Lazar and Tóth [24] on deterministic event-triggered sampling. It also provides a generalization of the works of Åström and Bernhardsson [1] and of Kushner [2].

6.1 Finite horizon estimation

In chapter 2 we have furnished methods to obtain good sampling policies for the finite horizon filtering problem. When the signal to be kept track of is a Wiener process, 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 performance.

We will report elsewhere on the solution to the case when the sensor has access only to noisy observations of the signal instead of perfect observations. The approach leads us to also consider some simple multi-sensor sampling and filtering problems which can be solved in the same way.

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 chapter.

Another set of unanswered questions involve 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 is decreased by using sampling designs that achieve merely good filtering performance.

6.2 Estimation on an infinite horizon

In chapter 3 we solved the repeated sampling problem over an infinite horizon when the signal has a scalar linear model. Åström and Bernhardsson [1], treat an average cost repeated resetting problem similar to the one we have solved. In fact, it is the same problem couched in different terms. The estimation error signal we have in our problem corresponds to a state process in [1], which needs to reset to zero by means of a resetting impulse control. The sampling rate of our problem corresponds to the rate at which the zero-resetting impulse control is invoked. By establishing the optimality of the symmetric threshold-triggered sampling policy, we have proved the optimality of the so-called Lebesgue sampling scheme discussed in [1] and shown to be superior to a periodic sampling scheme.

The problem we have solved has important extensions which need to be solved for direct applicability to practical designs. One extension involves finding the optimal or near-optimal strategies when the samples generated by the sensor get lost during transmission to the supervisor. Another direction involves non-Gaussian signals.

The estimation problem with multiple sensors which measure different noisy versions of the signal and 'listen' to each other's samples is also important for many sensor networks.

Problems of optimal repeated sampling with an average cost have been considered in [33, 34]. The survey [35] on optimal stochastic control contains some pointers to the literature on such average cost problems. The works [23, 36] discuss optimal single stopping problems with constraints like we have investigated in this chapter.

6.3 Average cost Control

In chapter 4 we have solved an average cost feedback control problem with reliable delivery of samples. We need to find ways of obtaining the optimal set of levels for event-triggered sampling. We need to see how the performances of the various sampling schemes compare when the signal is nonlinear. Extension to the case of a vector signal is non-trivial.

On the other hand, using multiple sensors for estimating a scalar state signal leads to a tractable analysis of level-triggered sampling. We could sample when the local conditional mean for the state at a sensor freshly crosses levels. The performance of such a sampling scheme can be analyzed with or without mutual listening of sensor samples. In principle, incorporating packet losses is possible but not transmission delays. This of course adds to the computational burden.

A hybrid sampling scheme based on level-triggered and time triggered sampling lets us stabilize unstable plants using only a finite set. The scheme depends on deterministic sampling when the state is beyond the last level in the finite $\mathcal{L}$. This sort of scheme is needed in practice in order to combat unreliable delivery of packets. However, analyzing and designing such a scheme gets more computationally burdensome.

Extension of the level-triggered scheme to the case of vector state signal is somewhat tricky. On the one hand, levels could be replaced by concentric spherical shells of the form,

$$\left|\vec{x}\right| = l_i > 0.$$

Of course, one could use ellipsoidal shells or other non-symmetrical and non-concentric shells. But this would differ from a scheme which samples based on threshold crossings of the magnitude of the estimation error signal. The latter scheme would be optimal for the average cost filtering problem when the state signal is scalar and linear.

6.4 Sequential detection in teams

In chapter 5 even though we have a sequential detection problem in infinite horizon, the optimal sampling schemes are not threshold-triggered. The sharing of data packets brings a time dependence on the cost structure. Loosely speaking, from the point of view a sensor in this team, the transmission by the other sensor is an event which affects the overall cost through the time of its occurrence. That is why the asynchronous sampling scheme $\mathcal{ASP}$ does not even possess the person-by-person optimality property it does in the context of [45, 46].

On the other hand, the mere numerical comparison of performances of the various threshold sampling schemes is not enough to develop design principles. Perhaps a different formulation of the multi-agent problems is needed to obtain insights.

Taking an overall view of the thesis in the context of [24], we need to explore connections with the familiar Shannon sampling theorem for bandlimited functions.

Appendix A

Expectations of some variables at hitting time

The purpose of this appendix is to provide a partial derivation of the expressions for expected time of a diffusion process to exit a regular set in $\mathbb{R}^n$. We will also sketch the similar result for computing the expected values of other quantities at first exit.

Now, we recall some basic results from the book of Øksendal (Theorem 5.2.1 in page 66 of [32]). Suppose that the diffusion process x_t is governed by the SDE:

$$dx_t = f(x_t)dt + g(x_t)dW_t,$$

where, $x_t \in \mathbb{R}^n$ and W_t is a standard Brownian motion process in $\mathbb{R}^n$. Let $f(\cdot)$, $g(\cdot)$ be such that

$$|f(x) - f(y)| + |g(x) - g(y)| \leq K\, |x - y|\,, \quad \forall\; x, y \in \mathbb{R}^n,$$

for some positive K. This condition ensures that the SDE above has a strong solution. The Brownian motion process and the Ornstein-Uhlenbeck process that we use in this thesis clearly satisfy the above conditions. The two dimensional process defined by the vector of likelihood ratios in chapter 5 also satisfies these conditions.

Let D be a regular set in $\mathbb{R}^n$ (section 9.2 in [32]). In the level-triggered sampling problems in this thesis, D is either a finite interval in $\mathbb{R}$ (3.3.2, 5.3.1), a finite, axes parallel rectangle in $\mathbb{R}^2$ (5.3.2, 5.3.4) or a semi-infinite axes-parallel rectangle in $\mathbb{R}^2$ (5.3.3). All of these sets are regular.

Let τ be the first exit time of x_t from D:

$$D = \inf \{ t \geq 0 \, | x_t \notin D \} .$$

Let $\mathbb{E}_x [\cdot]$ denote the conditional expectation given that $x_0 = x$, $\forall \, x \in D$. Then, if h, ϕ are integrable scalar functions,

$$\mathbb{E}_x \left[\int_0^\tau h(x_s) \, ds \right] + \mathbb{E}_x \left[\phi(x_\tau) \, ds \right] = \psi(x) .$$

where, $\psi(\cdot)$ satisfies the PDE (chapter 9 of [32]):

$$\begin{aligned} & L\phi = -h, \quad \text{with the boundary conditions,} \\ & \lim_{\substack{x \to y \\ x \in D}} = \phi(y), \quad \forall y \in \partial D, \end{aligned}$$

where L is given by:

$$L = \sum_{i=1}^{n} f_i(x) \frac{\partial}{\partial x_i} + \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} g_i(x) g_j(x) \frac{\partial^2}{\partial x_i \partial x_j} .$$

A rigorous proof is available in chapter 9 of [32]. In what follows, we will sketch how the result comes about.

By the Ito rule, we have:

$$\begin{aligned} d\phi(x_t) &= \left\{ \sum_{i=1}^{n} \frac{\partial \psi}{\partial x_i} \right\} dx_{i,t} + \left\{ \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} g_i(x) g_j(x) \frac{\partial^2 \phi}{\partial x_i \partial x_j} \right\} dW_{i,t} dW_{j,t} \\ &= \left\{ \sum_{i=1}^{n} f_i(x) \frac{\partial \psi}{\partial x_i} + \frac{1}{2} \sum_{i=1}^{n} \sum_{j=1}^{n} g_i(x) g_j(x) \frac{\partial^2 \phi}{\partial x_i \partial x_j} \right\} dt + \left\{ \sum_{i=1}^{n} g_i(x) \frac{\partial \phi}{\partial x_i} \right\} dW_{i,t}. \end{aligned}$$

Now,

$$\begin{aligned}
\mathbb{E}_x\left[\int_0^\tau h\left(x_s\right)ds\right] &= \mathbb{E}_x\left[\int_0^\tau -L\left(\phi(x)\right)ds\right] \\
&= \mathbb{E}_x\left[\int_0^\tau -L\left(\phi\right)ds\right] - \mathbb{E}_x\left[\int_0^\tau \left\{\sum_{i=1}^n g_i(x)\frac{\partial\phi(x)}{\partial x_i}\right\} dW_{i,s}\right] \\
&= \mathbb{E}_x\left[\int_0^\tau -d\phi(x_s)\right] \\
&= -\mathbb{E}_x\left[\phi\left(x_\tau\right) - \phi\left(x_0\right)\right] \\
&= -\mathbb{E}_x\left[\psi\left(x_\tau\right)\right] + \phi\left(x_0\right).
\end{aligned}$$

Rearranging the last equation above gives us:

$$\mathbb{E}_x\left[\int_0^\tau h\left(x_s\right)ds\right] + \mathbb{E}_x\left[\psi\left(x_\tau\right)\right] = \phi\left(x_0\right).$$

By choosing $h(\cdot) = 1$ and $\psi \equiv 0$, we can compute the expected first exit time $\mathbb{E}_x\left[\tau\right]$. By choosing $h(x) = x^2$ and $\psi \equiv 0$, we can compute the quantity:

$$\mathbb{E}_x\left[\int_0^\tau x_s^2 ds\right].$$

And by choosing $h \equiv 0$, and

$$\psi(x) = \mathbf{1}_{\{x\in C\}},$$

for some $C \subset \partial D$, we get the probability that the first exit occurs through C:

$$\mathbb{E}_x\left[\mathbf{1}_{\{x_\tau\in C\}}\right] = \mathbb{P}\left[x_\tau \in C\right].$$

BIBLIOGRAPHY

[1] Karl Johan Åström and Bo 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)*, pages 2011–2016. IEEE Control Systems Society, 2002.

[2] Harold J. Kushner. On the optimum timing of observations for linear control systems with unknown initial state. *IEEE Trans. Automatic Control*, AC-9:144–150, 1964.

[3] Toby Berger, Zhen Zhang, and Harish Viswanathan. The CEO problem. *IEEE Trans. Inform. Theory*, 42(3):887–902, 1996.

[4] Harish Viswanathan and Toby Berger. The quadratic Gaussian CEO problem. *IEEE Trans. Inform. Theory*, 43(5):1549–1559, 1997.

[5] Yasutada Oohama. The rate-distortion function for the quadratic Gaussian CEO problem. *IEEE Trans. Inform. Theory*, 44(3):1057–1070, 1998.

[6] Michael Gastpar, Bixio Rimoldi, and Martin Vetterli. To code, or not to code: lossy source-channel communication revisited. *IEEE Trans. Inform. Theory*, 49(5):1147–1158, 2003.

[7] Yao W. Hong and Anna Scaglione. Opportunistic large arrays: Cooperative transmission in wireless multi hop ad hoc networks to far distances. *IEEE transactions on Signal Processing*, 51(8):2082–2092, 2003.

[8] Bruce Hajek. Minimun mean hitting times of brownian motion with constrained drift. In *Proceedings of Conference on Stochastic processes and their applications*, 2000.

[9] Anna Hać. *Wireless Sensor Network Designs.* John Wiley, 2003.

[10] Wing Shing Wong and Roger W. Brockett. Systems with finite communication bandwidth constraints. I. State estimation problems. *IEEE Trans. Automat. Control*, 42(9):1294–1299, 1997.

[11] Wing Shing Wong and Roger W. Brockett. Systems with finite communication bandwidth constraints. II. Stabilization with limited information feedback. *IEEE Trans. Automat. Control*, 44(5):1049–1053, 1999.

[12] Roger W. Brockett and Daniel Liberzon. Quantized feedback stabilization of linear systems. *IEEE Trans. Automat. Control*, 45(7):1279–1289, 2000.

[13] Vivek S. Borkar, Sanjoy K. Mitter, and Sekhar Tatikonda. Optimal sequential vector quantization of Markov sources. *SIAM J. Control Optim.*, 40(1):135–148 (electronic), 2001.

[14] H. Vincent Poor. Fine quantization in signal detection and estimation. *IEEE Trans. Inform. Theory*, 34(5, part 1):960–972, 1988.

[15] Karl Johan Åström and Bo Bernhardsson. Systems with Lebesgue sampling. In *Directions in mathematical systems theory and optimization*, volume 286 of *Lecture Notes in Control and Inform. Sci.*, pages 1–13. Springer, Berlin, 2003.

[16] Dmitris Hristu-Varsakelis and Panganamala R. Kumar. Interrupt-based feedback control over a shared communication medium. In *Proceedings of the 41st IEEE conference on Decision and Control and European Control Conference (Las Vegas, 2002)*, pages 3223–3228. IEEE Control Systems Society, 2002.

[17] Vivek S. Borkar and Sanjoy K. Mitter. Lqg control with communication constraints. In *Communications, Computation, Control, and Signal Processing: a Tribute to Thomas Kailath*, pages 365–373. Kluwer Academic Publishers, Norwell, MA, 1997.

[18] Tunc Simsek, Rahul Jain, and Pravin Varaiya. Scalar estimation and control with noisy binary observations. *IEEE Trans. Automat. Control*, 49(9):1598–1603, 2004.

[19] Hiroaki Morimoto. Variational inequalities for combined control and stopping. *SIAM J. Control Optim.*, 42(2):686–708 (electronic), 2003.

[20] Ioannis Karatzas and Ingrid-Mona Zamfirescu. Martingale approach to stochastic control with discretionary stopping. *Appl. Math. Optim.*, 53(2):163–184, 2006.

[21] A. N. Shiryaev. *Optimal stopping rules*. Springer-Verlag, 1978. translated from the Russian *Statisticheskii posledovatelnyi analiz*.

[22] J. Zabczyk. Introduction to the theory of optimal stopping. In *Stochastic control theory and stochastic differential systems (Proc. Workshop, Deutsch. Forschungsgemeinsch., Univ. Bonn, Bad Honnef, 1979)*, volume 16 of *Lecture Notes in Control and Information Sci.*, pages 227–250. Springer, Berlin, 1979.

[23] Hiroaki Morimoto. On average cost stopping time problems. *Probab. Theory Related Fields*, 90(4):469–490, 1991.

[24] Aurel A. Lazar and László T. Tóth. Perfect recovery and sensitivity analysis of time encoded bandlimited signals. *IEEE Trans. Circuits Syst. I Regul. Pap.*, 51(10):2060–2073, 2004.

[25] Abhay G. Bhatt, Amarjit. Budhiraja, and Rajeeva L. Karandikar. Markov property and ergodicity of the nonlinear filter. *SIAM J. Control Optim.*, 39(3):928–949 (electronic), 2000.

[26] Maben Rabi and John 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)*, pages 4163–4168. IEEE Control Systems Society, 2004.

[27] Ioannis Karatzas and Steven E. Shreve. *Methods of mathematical finance*, volume 39 of *Applications of Mathematics (New York)*. Springer-Verlag, New York, 1998.

[28] Roland Glowinski. *Numerical methods for nonlinear variational problems.* Springer Series in Computational Physics. Springer-Verlag, New York, 1984.

[29] Maben Rabi, John S. Baras, and George V. Moustakides. Efficient sampling for keeping track of a gaussian process. In *Proceedings of the 14th Mediterranean Conferences on Control and Automation (Ancona, 2006)*. IEEE Control Systems Society, 2006.

[30] Maben Rabi, George V. Moustakides, and John S. Baras. Multiple sampling for estimation on a finite horizon. In *forthcoming Proceedings of the 45rd IEEE conference on Decision and Control (San Diego, CA, 2006)*, pages xxxx–xxxx. IEEE Control Systems Society, 2006.

[31] Bruno Sinopoli, Luca Schenato, Massimo Franceschetti, Kameshwar Poolla, Michael I. Jordan, and Shankar S. Sastry. Kalman filtering with intermittent observations. *IEEE Trans. Automat. Control*, 49(9):1453–1464, 2004.

[32] Bernt Øksendal. *Stochastic differential equations.* Universitext. Springer-Verlag, Berlin, sixth edition, 2003. An introduction with applications.

[33] Dariusz Gątarek. Ergodic impulsive control of Feller processes with costly information. *Systems Control Lett.*, 15(3):247–257, 1990.

[34] Łukasz Stettner. On impulsive control with long run average cost criterion. *Studia Math.*, 76(3):279–298, 1983.

[35] Vivek S. Borkar. Controlled diffusion processes. *Probability Surveys*, 2:213–244 (electronic), 2005.

[36] Monique Pontier and Jacques Szpirglas. Optimal stopping with constraint. In *Analysis and optimization of systems, Part 2 (Nice, 1984)*, volume 63 of *Lecture Notes in Control and Inform. Sci.*, pages 82–91. Springer, Berlin, 1984.

[37] Ioannis Karatzas and Hui Wang. Utility maximization with discretionary stopping. *SIAM J. Control Optim.*, 39(1):306–329 (electronic), 2000.

[38] Michel Adès, Peter E. Caines, and Roland P. Malhamé. Stochastic optimal control under Poisson-distributed observations. *IEEE Trans. Automat. Control*, 45(1):3–13, 2000.

[39] Vivek S. Borkar and Pravin Varaiya. Finite chain approximation for a continuous stochastic control problem. *IEEE Trans. Automat. Control*, 26(2):466–470, 1981.

[40] Oneésimo Hernández-Lerma. *Adaptive Markov control processes*, volume 79 of *Applied Mathematical Sciences*. Springer-Verlag, New York, 1989.

[41] Dimitri P. Bertsekas. *Dynamic programming and optimal control. Vol. II*. Athena Scientific, Belmont, MA, second edition, 2001.

[42] Vivek S. Borkar. Controlled Markov chains and stochastic networks. *SIAM J. Control Optim.*, 21(4):652–666, 1983.

[43] John N. Tsitsiklis. Decentralized detection. In *Adv. Statist. Signal Processing*, volume 2 of *Lecture Notes in Control and Inform. Sci.*, pages 297–344. JAI press, Greenwich, CT, 1993.

[44] J.-F. Chamberland and Venugopal V. Veeravalli. Decentralized detection in sensor networks. *IEEE Transactions on Signal Processing*, 51(2):407–416, 2003.

[45] Demosthenis Teneketzis and Yu-Chi Ho. The decentralized wald problem. *Infomation and Computation*, 73(1):23–44, 1987.

[46] Anthony LaVigna, Armand M. Makowski, and John S. Baras. A continuous-time distributed version of Wald's sequential hypothesis testing problem. In *Analysis and optimization of systems (Antibes, 1986)*, volume 83 of *Lecture Notes in Control and Inform. Sci.*, pages 533–543. Springer, Berlin, 1986.