TOPICS IN DECENTRALIZED DETECTION

BY

VENUGOPAL VEERAVALLI

B. Tech., Indian Institute of Technology, Bombay, 1985
M.S., Carnegie Mellon University, Pittsburgh, 1987

THESIS

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Electrical Engineering
in the Graduate College of the
University of Illinois at Urbana-Champaign, 1992

Urbana, Illinois

TOPICS IN DECENTRALIZED DETECTION

Venugopal Veeravalli, Ph.D.
Department of Electrical and Computer Engineering
University of Illinois at Urbana-Champaign, 1992
H. Vincent Poor and Tamer Başar, Advisor

In this thesis we obtain several new results in the areas of decentralized sequential detection and robust decentralized detection.

In the area of decentralized sequential detection, we first consider the case in which each sensor performs a sequential test on its observations and arrives at a local decision about the true hypothesis; subsequently, the local decisions of all of the sensors are used for a common purpose. Here we assume that decision errors at the sensors are penalized through a common cost function and that each time step taken by the detectors as a team is assigned a positive cost. We show that optimal sensor decision functions can be found in the class of *generalized sequential probability ratio tests* with monotonically convergent thresholds. We present a technique for obtaining optimal thresholds.

We also consider the case in which each sensor sends a sequence of summary messages to a fusion center in which a sequential test is carried out to determine the true hypothesis. Here we assume that decision errors at the fusion center are penalized through a cost function and that each time step taken to arrive at the final decision costs a positive amount. We show that the problem is tractable when the information structure in the system is *quasiclassical*. In particular, we show that an optimal fusion center policy has a simple structure resembling a sequential probability ratio test and that a stationary set of monotone likelihood ratio tests is optimal at the sensors. Finally, we compute the optimal decision functions for some representative examples.

In the area of robust decentralized detection, we consider the case in which the sensor distributions are assumed to belong to known uncertainty classes. We show for a broad class of such decentralized detection problems that a set of least favorable distributions exists for minimax robust testing between the hypotheses. We thus establish that minimax robust tests are obtained as solutions to simple decentralized detection problems in which the sensor distributions are specified to be the least favorable distributions.

ACKNOWLEDGEMENTS

First and foremost, I would like to thank my advisors, Professor T. Başar and Professor H. V. Poor, for their guidance and support throughout my graduate work at the University of Illinois. I would also like to thank the other members of my thesis committee, Professors B. Hajek, S. Meyn and K. Jenkins, for many helpful and stimulating discussions. Thanks are also due to my fellow graduate students and the staff at the Coordinated Science Laboratory for their friendship and help. Last but not least, I would like to thank Starla for everything, including her help in getting this thesis together.

TABLE OF CONTENTS

CHAPTER 1
INTRODUCTION

Statistical decision-making is a generic term referring to scenarios in which one or more individuals (called decision makers) are confronted with the task of deciding between a number of alternatives (finite or infinite) in an uncertain environment, so as to satisfy a given objective or a set of objectives. The uncertain environment (sometimes called the state of nature) generally has a complete probabilistic description that is known to all decision makers, and the decisions are based on the measurements acquired (through sensors) on the unknown state of nature. Such decision problems arise in the fields of communications, control, and image and signal processing. If the number of alternatives is finite, they are called *detection*, or *hypothesis testing*, problems.

Detection problems admit two types of classifications, according to whether the information is *centralized* or *decentralized*, and subsequently to whether the decisions are *static* or *dynamic*. A paradigm for a centralized detection problem is one in which all of the information received by the sensors is sent to a central processor where a decision is made according to a given criterion. Hence, even if more than one unit is involved in the detection problem, in the centralized setting they can be viewed as a single decision maker. In the decentralized setting, however, only a summary message from each sensor is sent to the central processor, which therefore does not receive all of the available information. In static frameworks, each decision maker makes only one decision, whereas in dynamic scenarios, decision makers update their decisions as new information becomes available. Such dynamic problems are also known as sequential detection problems — a class of problems first introduced by Wald [1].

Centralized detection problems are well-understood today and a very well-developed theory exists for their analysis [2]. Decentralized detection, on the other hand, is a relatively nascent field. A theoretical framework for decentralized detection is still in the process of being developed; the goal of this thesis is to contribute to this development, for both static and dynamic problems.

1.1 Motivation

There are two main reasons why a decentralized setting, such as that introduced in [3], may be preferable to the centralized setting. First of all, the decentralized scheme offers the possibility for drastic reductions in the bandwidth requirements for communication between the sensors and the central processor. Hence a natural application for the decentralized scheme is a situation in which the sensors are far-removed from the central processor, and the bandwidth available for communication is limited. Second, the decentralized setting allows for distributed or shared information processing, thus reducing the burden on the central processor. An application where it would be very desirable to have distributed processing can be found in the context of fault detection in large scale systems such as power systems, surveillance systems and VLSI circuits. In these systems, the sensors could monitor different parts of the large scale system and send messages to a central unit, which would make the final decision about the existence of a fault. Another application for decentralized detection—one which has received a fair amount of attention recently—is in human decision-making organizations, where subordinates play the role of the sensors and a supervisor takes on the role of the fusion center (see, for example, [4, 5]).

It may seem that solutions to decentralized detection problems could be obtained readily by applying classical detection theory. However, each decision maker[1] in the decentralized setting receives only partial information. Therefore, decentralized detection problems fall in the general class of *team decision* problems, which are usually very difficult to solve, and in many cases are known to be intractable [6]. However, definite progress can be made in the particular case of decentralized detection problems. One of the main reasons for the tractability of these problems is that the search for optimal decision functions can often be restricted to classes that admit finite-dimensional parametrizations.

Decentralized detection theory is also very closely related to quantization theory [7, 8, 9, 10]. As in decentralized detection theory, the goal in quantization theory is to find an optimal

[1]Each sensor in the decentralized system can be regarded as a local decision maker even though the sensor decisions may not be directly related to decisions about the hypothesis.

way to quantize (summarize) the observations for decision making. However, there are some differences in emphasis. In contrast with decentralized detection theory, quantization theory typically assumes a single source of information and a large quantizer alphabet. Also, as pointed out in [11], the probability-of-error performance criterion has been avoided in the study of quantization problems for reasons of mathematical tractability, whereas it is the most commonly used performance criterion in decentralized detection theory.

1.2 Bayesian Framework for Decentralized Detection

The basic framework for centralized detection is the Bayesian framework [2]. Hence, as an introduction to decentralized detection, we first consider briefly a Bayesian framework for the fusion configuration discussed earlier.

The basic structure for decentralized detection is the one in which there are N sensors, $S_1, \ldots, S_N$, and one *fusion center*, as shown in Figure 1.1. (Figures appear at the end of chapters.) The hypothesis is denoted by a random variable H which is assumed to take on values H_0 and H_1, with a priori probabilities ν and $1-\nu$, respectively. Sensor S_i receives an observation X_i, which is a random variable that takes values on a measurable space $(\mathcal{X}_i, \mathcal{F}_i)$. It is assumed that the joint probability distribution function of $(X_1, \ldots, X_N)$ conditioned on each hypothesis is known. The sensor S_i, upon receiving observation X_i, evaluates a message $u_i = \phi_i(X_i)$ where $u_i \in \{1, \ldots, D_i\}$. The mapping $\phi : \mathcal{X}_i \mapsto \{1, \ldots, D_i\}$ is referred to as a local decision function (rule). The fusion center then makes a final binary-valued decision δ based on the information it receives from the sensors, i.e., $\delta = \gamma(u_1, \ldots, u_N)$, where $\gamma : \{1, \ldots, D_i\} \times \cdots \times \{1, \ldots, D_N\} \mapsto \{0, 1\}$ is the fusion decision function. The collection $(\phi_1, \ldots, \phi_N, \gamma)$ is referred to as a *strategy*.

The decision functions at the sensors and the fusion center are chosen to meet certain optimality criteria. A criterion that is commonly used is the *Bayesian* cost criterion in which the objective is to minimize an expected cost. In a general Bayesian formulation, a cost function $W : \{0, 1\} \times \{1, \ldots, D_1\} \times \cdots \times \{1, \ldots, D_N\} \times \{H_0, H_1\} \mapsto \mathbb{R}$ is given, with

$W(\delta, u_1, \ldots, u_N, H)$ being the cost associated with the sensor and fusion center decisions when the hypothesis is H. The Bayesian optimization problem is then to minimize the expected value of $W(\delta, u_1, \ldots, u_N, H)$ over all admissible strategies.

An example of a Bayesian cost criterion is the *probability of error criterion*. Here the cost function $W(\delta, u_1, \ldots, u_N, H_j)$ equals 1 when $\delta \neq j$, and equals 0 otherwise. The objective then is to minimize the probability of decision error at the fusion center.

The centralized Bayesian detection problem is well-known to be tractable and the solution is a simple likelihood ratio test. However, even the most basic decentralized version of this problem, described above, is not tractable unless we make certain assumptions. The most important of these is the conditional independence assumption stated below.

Assumption 1.1 *The observations received by the various sensors are conditionally independent given each hypothesis.*

The importance of the Assumption 1.1 is elucidated by the following computational complexity result which was proved in [12]. Consider the minimum probability of error problem described above with $N = 2$ and $D_1 = D_2 = 2$, where the sets $\mathcal{X}_i$, $i = 1, 2$ are finite. Let K be a rational number. Then the problem of verifying the existence of a strategy for which the expected Bayesian cost is less than or equal to K is NP-complete.

We now present some results for the Bayesian decentralized detection problem. Tenney and Sandell [3] were the first to consider this problem, and they explored the case in which $N = 2$ and $D_1 = D_2 = 2$. Under Assumption 1.1, they have shown that optimal decision functions can be found in the class of likelihood ratio tests (LRTs). This means that the search for optimal decision functions can be done over the finite-dimensional space of threshold values, thus making the optimization problem tractable. The result was extended to the general case of arbitrary N and D_i in [11]. Such a result is not true without Assumption 1.1 as was shown by a counterexample in [13]. However, in the conditionally dependent case, we can restrict the decision functions to the class of likelihood ratio tests and determine the best decision functions in this class (see for example, [14], [15], [16]).

In large-scale systems, symmetry assumptions are often made to simplify the analysis of these systems. A natural symmetry assumption for decentralized detection is the following:

Assumption 1.2 *The sensor observations are independent and identically distributed, conditioned on each hypothesis.*

Under Assumption 1.2 and for a cost function that is symmetric in $(u_1, \ldots, u_N)$, it may at first seem optimal for all sensors to use the same likelihood ratio test. This is not true, however, as illustrated by a counterexample in [13]. Despite the sub-optimality of identical decision functions under Assumption 1.2, a number of papers have made this simplifying restriction (see for example, [17], [18], [19]). With this restriction, the sensor decisions $(u_1, \ldots, u_N)$ are independent and identically distributed. It is clear that for the special case of binary sensor decisions, the number K, defined as the cardinality of the the set $\{i|u_i = 1\}$, becomes a sufficient statistic for the fusion center. Thus, an optimal fusion rule has the form $\gamma(u_1, \ldots, u_N) = 1$ if and only if $K \leq k$, where k is a threshold value. Such decision rules are often referred to as "k-out-of-N" rules.

1.3 Variations from the Basic Formulation and Existing Results

Bayesian optimization for the fusion configuration was the focus of many of the early papers in this field. But over the last decade a number of variations have been considered, as discussed below:

1.3.1 Neyman-Pearson problems

Decentralized detection problems with the Neyman-Pearson optimality criterion are usually more involved than their Bayesian counterparts. As in centralized Neyman-Pearson detection problems, we need to introduce a randomization factor in the decision functions.

There are two ways to randomize the decision functions: joint randomization, where the decision functions are randomized together, and independent randomization, where the decision functions are randomized independently. The class of jointly randomized strategies

is larger than the class of independently randomized strategies, but joint randomization may not be feasible in many applications.

In the class of jointly randomized strategies, it is quite straightforward to show that optimal strategies for the Neyman-Pearson problem can be obtained by randomizing between two deterministic LRT strategies [11]. In the class of independently randomized strategies, however, it is not so easy to establish the optimality of randomized LRT strategies. Several papers have been written on this topic (see for example, [20], [21], [22], [23]). However, as pointed out in [11], the optimization arguments made in these papers are incorrect. A correct argument was given in [24] for the case in which the sensors make binary decisions. This argument was extended to the general case of nonbinary sensor decisions in [42].

1.3.2 Other sensor configurations

The case in which the fusion center is completely absent has been investigated in [3]. This case can be handled by a Bayesian analysis of the type described in Section 1.2 by having a cost function W that does not depend on δ. In the special case in which $D_i = 2$, for $i = 1, \ldots, N$, each of the sensor decisions can be regarded as decisions about the true hypothesis, with these decisions being coupled through a common cost function.

Tree configurations with the base of the tree making the final decision have been studied in [25] where the optimality of likelihood ratio tests was established under Assumption 1.1. The optimality conditions were elaborated in [18] where special cases were studied in detail. A tandem sensor configuration has been studied in [26]. Here the sensor at the end of the tandem makes the final decision. This configuration generated some interest because in the two-sensor case, it is at least as good as the fusion (parallel) configuration. This is not true if the number of sensors is greater than two. In fact, the following asymptotic result holds: Consider a tandem of N sensors which receive independent and identically distributed observations, and where the sensors relay binary decisions. If the likelihood ratio of the observation is bounded away from zero and infinity, the asymptotic error probability

(as $N \to \infty$) is bounded away from zero [27]. This is in sharp contrast with the exponential decrease to zero of the error probability for the parallel configuration [13].

An important feature of this asymptotic result for the parallel configuration is that it is asymptotically optimal for all of the sensors to use the same decision rule [13]. This fact is intuitively appealing, and is sometimes used as a justification for using identical sensor decision functions when the number of sensors is large and Assumption 1.2 holds.

1.3.3 Decentralized sequential detection

The detection problem described in Section 1.2 and the variations discussed thus far are static problems in which the sensors receive either a single observation or a single block of observations about the hypothesis. In a dynamic setting (which is a generalization of the static setting), each sensor receives an entire sequence of observations and the detection system has the option to stop at any time and make a final decision, or to continue taking observations. There are two main categories of problems. In one case, each sensor sends a sequence of summary messages to a fusion center where a sequential test is carried out to determine the true hypothesis. In the other case, each sensor performs a sequential test on its observations and arrives at a final local decision. If a fusion center is present, the local decision is sent to the fusion center, which then makes the final decision about the hypothesis. If there is no fusion center, the local decisions are used for a common purpose at a site possibly remote to all the sensors.

Decentralized sequential detection problems are considerably more complex than their static counterparts, and very few positive results have been obtained previously. These problems will be the focus of Chapters 2 and 3 of this thesis. Since relevant previous results will be discussed in the introductory sections of these chapters, a description of these is omitted here.

1.3.4 Robust decentralized detection

The design of optimal decision rules for decentralized detection problems is based on the assumption that the probability distributions of the sensor observations (under each hypothesis) are known. In many applications, however, the distributions of the sensor observations are only specified as belonging to classes which are referred to as *uncertainty classes*. The problem here is to design decision rules that are robust with respect to uncertainties in the distributions. A common approach for such a design is the minimax approach where the goal is to minimize the worst-case performance over the uncertainty classes.

This problem will be the focus of Chapter 4 of this thesis. Very little work has been done previously on this problem, and relevant results will be discussed in Chapter 4.

1.4 Contribution of This Thesis

As evidenced in the above summary of existing results, simple static decentralized problems are well-understood and most tractable problems have been resolved. There has not been, however, significant progress in the related fields of decentralized sequential detection and robust decentralized detection. We attempt to address this situation by solving problems in these areas which were either previously unsolved or which had not been resolved adequately. In doing so, we hope to provide theoretical frameworks for analyzing the remaining open problems in these fields, some of which are identified in Chapter 5.

1.5 Outline

The remainder of the thesis is organized as follows. In Chapter 2, we study decentralized sequential detection problems in which the sensors perform sequential tests. The focus of Chapter 3 is on decentralized sequential detection problems in which a fusion center performs the sequential test. Next, we consider minimax robust decentralized detection in Chapter 4. Finally, the conclusions of the thesis are presented in Chapter 5, where we also identify several problems that still remain open in the field, and provide some partial solutions and ideas for future work.

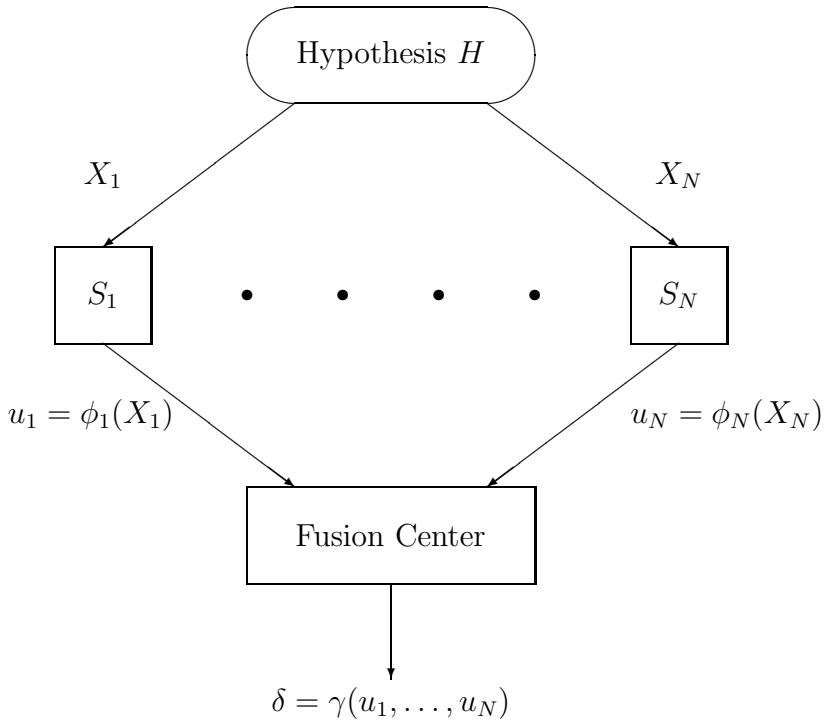

Figure 1.1: Basic structure for decentralized detection.

CHAPTER 2

DECENTRALIZED SEQUENTIAL DETECTION—PART I: SENSORS PERFORMING SEQUENTIAL TESTS

2.1 Introduction

In centralized sequential binary hypothesis testing, the detector is required to determine the true hypothesis based on a sequence of received observations. This decision problem can be posed in a Bayesian framework as follows: The hypothesis H is assumed to take on the two values, H_0 and H_1, with known prior probabilities ν and $1-\nu$, respectively. A positive cost c is associated with each observation (time step) taken by the detector. The detector stops receiving additional measurements at time τ, which is assumed to be a stopping time for the sigma field sequence generated by the observations, and makes a final decision δ based on the observations up to time τ. Decision errors are penalized through a decision cost function $W(\delta; H)$. The stopping rule together with the final decision rule represent the decision policy of the detector. The total expected cost (risk) for a given decision policy is given by $E\{c\tau + W(\delta; H)\}$. The centralized Bayesian sequential detection problem, which is sometimes referred to as the *Wald problem*, is to find a decision policy leading to minimum total expected cost. The solution to this problem for the case when the observations are i.i.d., conditioned on each hypothesis, is the well-known sequential probability ratio test (SPRT) [28].

In *decentralized* sequential hypothesis testing, each one of a set of sensors receives a sequence of observations about the hypothesis. As we mentioned in Section 1.3.3, there are two possible settings here. In this chapter, we consider the setting in which each sensor in the decentralized system performs a sequential test and arrives at a local decision. It is assumed that the sensor decisions are used for a common goal, possibly at some site remote to all of the sensors.

We consider a Bayesian formulation of this problem with two hypotheses[1], and for simplicity of presentation, we study the case of two sensors. We denote the sensors by S_1 and S_2. Sensor S_i stops at time τ_i, and makes a decision u_i based on its observations up to time τ_i. The combined decision policy of the two sensors is denoted by $\gamma = (\gamma_1, \gamma_2)$, where $\gamma_i := (u_i, \tau_i)$ is the decision policy of sensor S_i.

Since the two decisions u_1 and u_2 are used for a common goal, it is natural to assume that decision errors are penalized through a common decision cost function $W(u_1, u_2; H)$. The choice of a time penalty is, however, not as unambiguous. If we are concerned with processing cost at the sensors, then we associate a positive cost c_i with each observation taken by sensor S_i. On the other hand, there may be situations in which we may wish to limit the time it takes for both decisions to be available at the remote site. In this case it may be more reasonable to associate a positive cost c with each time step taken by the sensors as a *team*.

Teneketzis and Ho [30] considered the situation in which a positive cost c_i is associated with each observation taken by sensor S_i. In this case, the total expected cost for a given combined decision policy γ is $E\{c_1\tau_1 + c_2\tau_2 + W(u_1, u_2; H)\}$. The Bayesian optimization problem is then to find the decision policy that minimizes this expected cost. A special case here is one in which the decision cost function is decoupled, i.e., $W(u_1, u_2; H) = W_1(u_1; H) + W_2(u_2; H)$. This is equivalent to the assumption that the sensor decisions are used for independent purposes. In this case, we have two decoupled Wald problems to solve, one at each of the sensors, and the solution is two independent SPRTs. Teneketzis and Ho showed in [30], using a rather involved argument, that even when there is coupling, optimal sensor decision policies can be found within the class of SPRTs. Their result can be derived immediately by recognizing that once the decision policy of sensor S_2 is fixed, sensor S_1 is faced with a classical Wald problem. This point was later clarified in [31], where a continuous time extension of this problem was solved.

[1]We will restrict our attention to binary hypothesis testing in this paper. Problems in sequential testing of multiple hypotheses are known to be very difficult and do not admit closed-form solutions even when the information is centralized [29].

In our analysis, we associate a positive cost c with each time step taken by the detectors as a team. The expected cost we wish to minimize over all admissible policies is then given by

$$E\{c\max(\tau_1, \tau_2) + W(u_1, u_2; H)\}.$$

The nonlinearity introduced by considering the maximum of the two stopping times makes this problem more difficult than the one solved in [30].

The rest of this chapter is organized as follows: In Section 2.2, we provide a more formal description of the problem we wish to solve. Then in Section 2.3, we focus on the structure of optimal solutions to this problem. In particular, we show that optimal solutions can be found in the class of generalized SPRTs (GSPRTs) with monotonically convergent thresholds. In Section 2.4, we address the problem of finding optimal GSPRT thresholds numerically. In Section 2.5, we present some numerical results for the case when the sensor observations are Gaussian under each hypothesis. We also compare the performance of optimal GSPRTs with the best performance that is obtained when the sensors are restricted to use SPRTs. Finally, in Section 2.6, we summarize the main points.

2.2 Mathematical Description

We begin with a formal description of the decentralized sequential detection problem we wish to analyze here.

1. The hypothesis is denoted by a binary random variable H which takes on values H_0 and H_1, with prior probabilities ν and $1-\nu$, respectively.

2. At time k, sensor S_i receives observation X_k^i, $i = 1, 2$. The sequences $\{X_k^1\}_{k=1}^{\infty}$ and $\{X_k^2\}_{k=1}^{\infty}$ are mutually independent i.i.d. sequences, conditioned on each hypothesis. The probability distributions of the sensor observations are assumed to have densities, and we denote the conditional density of X_k^i given H_j by f_j^i.

3. There is no communication between the sensors, i.e., the final decision at each sensor is based only on its own observations.

4. Let $\mathcal{X}_k^i = \sigma(X_j^i,\ j = 1, 2, ..., k)$. The decision policy γ_i for sensor S_i involves the selection of a termination time τ_i, and a binary valued decision u_i. For an admissible policy, τ_i is a $\{\mathcal{X}_k^i, k = 1, 2, ..\}$-stopping time, and u_i is measurable $\mathcal{X}_{\tau_i}^i$. The set of admissible policies is denoted by Γ_i.

5. If u_i denotes the final decision at sensor S_i, then the decision cost $W(u_1, u_2; H)$ satisfies the following inequalities for $u_2 = 0$ and $u_2 = 1$:

$$W(0, u_2; H_1) \geq W(1, u_2; H_1),$$

$$W(1, u_2; H_0) \geq W(1, u_2; H_1),$$

$$W(1, u_2; H_0) \geq W(0, u_2; H_0),$$

$$W(0, u_2; H_1) \geq W(0, u_2; H_0).$$

Similar inequalities hold for u_1, i.e., at most one error is not more costly than at least one error. Also, each unit of time taken by the sensors as a team costs a positive amount c.

The problem that we wish to solve is the following:

Problem P2.1

$$\min_{\{\gamma_i \in \Gamma_i\}_{i=1,2}} \quad E\{c \max(\tau_1, \tau_2) + W(u_1, u_2; H)\}.$$

□

2.3 The Structure of Optimal Solutions

In this section we study the common structure of all person-by-person optimal (p.b.p.o.) decision policies[2]. This structure would obviously be valid for globally optimal (g.o.) decision policies as well, since every g.o. decision policy is also p.b.p.o.

If γ_2 is fixed, possibly at the optimum, then u_2 and τ_2 have fixed distributions conditioned on each hypothesis. At sensor S_1, we are faced with the following optimization problem:

$$\min_{\{\gamma_1 \in \Gamma_1\}} \quad E\{c \max(\tau_1, \tau_2) + W(u_1, u_2; H)\}. \tag{2.1}$$

[2]A set of policies is said to be person-by-person optimal if it is not possible to improve the corresponding team performance by unilaterally changing any one of the policies. Clearly, globally optimal decision policies are also person-by-person optimal.

This can be posed as an infinite-horizon dynamic programming (DP) problem [32]. A sufficient statistic for this is given by

$$p_k = P(H = H_0 | \mathcal{X}_k^1).$$

A recursion for p_k is easily obtained by using Bayes' rule,

$$p_{k+1} = \frac{p_k f_0(X_{k+1}^1)}{p_k f_0(X_{k+1}^1) + (1-p_k) f_1(X_{k+1}^1)}, \qquad p_0 = \nu,$$

where $f_j(.)$ is the probability density of X_{k+1}^1 conditioned on H_j, $j = 0, 1$. Note that the conditional density of X_{k+1}^1 given $\mathcal{X}_k^1$, which we denote by $f(p_k; .)$, is given by

$$f(p_k; x) = p_k f_0(x) + (1-p_k) f_1(x).$$

We wish to solve the optimization problem of (2.1) using dynamic programming (DP). To this end, we first restrict the stopping time τ_1 to a finite interval, say $[0, T]$. The finite-horizon DP equations are derived as follows. The minimum expected cost-to-go at time k is a function of the sufficient statistic p_k, which we denote by $J_k^T(p_k)$. It is easily seen that

$$J_T^T(p_T) = \min\{G_0 p_T + K_0, G_1 p_T + K_1\}$$

where

$$K_i = \sum_{j=0}^{1} P_1(u_2 = j)\, W(i, j; H_1), \qquad i = 0, 1,$$

$$G_i = \sum_{j=0}^{1} P_0(u_2 = j)\, W(i, j; H_0) - K_i, \quad i = 0, 1,$$

and where P_j denotes the probability measure conditioned on H_j.

For $0 \le k \le T-1$, a standard DP argument yields the following recursion:

$$J_k^T(p_k) = \min\{G_0 p_k + K_0, G_1 p_k + K_1, c\, p_k\, P_0(\tau_2 \le k) + c\,(1-p_k)\, P_1(\tau_2 \le k) + \Lambda_k^T(p_k)\}, \quad (2.2)$$

where

$$\begin{aligned} \Lambda_k^T(p_k) &= E_{X_{k+1}^1 | \mathcal{X}_k^1} \left\{ J_{k+1}^T \left(\frac{p_k f_0(X_{k+1}^1)}{f(p_k; X_{k+1}^1)} \right) \right\} \\ &= \int J_{k+1}^T \left(\frac{p_k f_0(x)}{f(p_k; x)} \right) f(p_k; x) dx. \end{aligned} \quad (2.3)$$

In (2.2), the term $G_0 p_k + K_0$ represents the cost (conditioned on $\mathcal{X}_k^1$) of stopping at time k and choosing H_0, the term $G_1 p_k + K_1$ represents the cost of stopping at time k and choosing H_1, and the last term represents the cost of continuing at time k. Note that sensor S_1 is penalized for taking an additional step at time k only if sensor S_2 has stopped before time k.

The lemmas below present some useful properties of the functions J_k^T and Λ_k^T.

Lemma 2.1 *The functions $J_k^T(p)$ and $\Lambda_k^T(p)$ are nonnegative concave functions of p, for $p \in [0,1]$.* □

Lemma 2.2 *The functions $J_k^T(p)$, and $\Lambda_k^T(p)$ are monotonically nondecreasing in k, that is, for each $p \in [0,1]$,*

$$J_k^T(p) \leq J_{k+1}^T(p), \quad 0 \leq k \leq T-1,$$

$$\Lambda_k^T(p) \leq \Lambda_{k+1}^T(p), \quad 0 \leq k \leq T-2.$$

□

Lemma 2.3 *The functions $\Lambda_k^T(p)$ satisfy the following properties:*

$$\Lambda_k^T(0) = \min\{K_0, K_1\} = K_1,$$

$$\Lambda_k^T(1) = \min\{K_0 + G_0, K_1 + G_1\} = K_0 + G_0.$$

□

The above lemmas are easily proven by simple induction arguments. We can use these lemmas to derive the structure of finite-horizon optimal solutions as we did in [33]. Here we focus on the infinite-horizon case.

2.3.1 Infinite-horizon optimization

In order to solve the problem P2.1, we need to remove the restriction that τ_1 belongs to a finite interval, by letting $T \to \infty$. By an argument similar to the one in Section 3.3 of [30], we can establish that for each k, the following limit is well-defined:

$$\lim_{T\to\infty, T>k} J_k^T(p) = \inf_{T>k} J_k^T(p) =: J_k(p).$$

The function $J_k(p)$ is the infinite-horizon cost-to-go at time k. Unlike the infinite-horizon solution in [30], this limit need not be independent of k . In fact, if we let $T \to \infty$ in Lemma 2.2, we see that the following monotonicity holds in the limit:

$$J_k(p) \leq J_{k+1}(p), \quad \forall k.$$

Also, it is clear that $J_k(p)$ is bounded above by $\min\{G_0 p + K_0, G_1 p + K_1\}$ for all k. Hence, the limit

$$\lim_{k\to\infty} J_k(p) = \sup_k J_k(p) =: J(p)$$

is also well-defined, and satisfies the Bellman equation [32]

$$J(p) = \min\left\{G_0 p + K_0, G_1 p + K_1, c + \int J\left(\frac{p f_0(x)}{f(p;x)}\right) f(p;x)dx\right\}. \tag{2.4}$$

Teneketzis and Ho [30] obtain exactly the same Bellman equation in the context of the decentralized Wald problem with linear time penalty, where they also show that the equation has a unique solution (see Lemma 3.3 of [30]).

Now, by the Dominated Convergence Theorem the following limits are well-defined:

$$\Lambda_k(p) := \lim_{T\to\infty} \Lambda_k^T(p) = \int J_{k+1}\left(\frac{p f_0(x)}{f(p;x)}\right) f(p;x)dx,$$

and

$$\Lambda_J(p) := \lim_{k\to\infty} \Lambda_k(p) = \int J\left(\frac{p f_0(x)}{f(p;x)}\right) f(p;x)dx.$$

Hence the infinite-horizon cost-to-go function satisfies the recursion

$$J_k(p) = \min\{G_0 p + K_0, G_1 p + K_1, c\, p\, P_0(\tau_2 \leq k) + c\,(1-p)\, P_1(\tau_2 \leq k) + \Lambda_k(p)\}. \tag{2.5}$$

Taking limits as $T \to \infty$ in Lemmas 2.1-2.3, we obtain the following result:

Lemma 2.4 *The functions $\Lambda_k(p)$ are concave and satisfy*

$$\Lambda_k(p) \leq \Lambda_{k+1}(p), \quad \forall p \in [0,1],$$

$$\Lambda_k(0) = K_1, \qquad \Lambda_k(1) = K_0 + G_0.$$

□

It follows from Lemma 2.4 that provided the condition

$$c + \Lambda_J(\frac{K_0 - K_1}{G_1 - G_0}) \leq \frac{G_1 K_0 - G_0 K_1}{G_1 - G_0} \tag{2.6}$$

holds, we have the following result (see Section 6.3 of [32] for a similar analysis).

Theorem 2.1 *For fixed $\gamma_2 \in \Gamma_2$, let condition (2.6) hold. Then, an optimal infinite-horizon policy at sensor S_1 is of the form*

$$\begin{aligned} &\textit{accept } H_0 \quad &&\textit{if} \quad p_k \geq a_k, \\ &\textit{accept } H_1 \quad &&\textit{if} \quad p_k \leq b_k, \\ &\textit{continue} \quad &&\textit{if} \quad b_k < p_k < a_k, \end{aligned}$$

where the scalars a_k, b_k, $k = 0, 1, 2...$, are obtained from the relations

$$G_1 b_k + K_1 = c\, b_k\, P_0(\tau_2 \leq k) + c\,(1 - b_k)\, P_1(\tau_2 \leq k) + \Lambda_k(b_k),$$
$$G_0 a_k + K_0 = c\, a_k\, P_0(\tau_2 \leq k) + c\,(1 - a_k)\, P_1(\tau_2 \leq k) + \Lambda_k(a_k).$$

Furthermore, $\{a_k\}_{k=1}^{\infty}$ is a nonincreasing sequence converging to a and $\{b_k\}_{k=1}^{\infty}$ is a nondecreasing sequence converging to b, where a and b satisfy

$$c + \Lambda_J(b) = G_1 b + K_1,$$
$$c + \Lambda_J(a) = G_0 a + K_0.$$

□

Remark 2.1 If condition (2.6) does not hold, then the sequences a_k and b_k are both identically equal to $(K_0 - K_1)/(G_1 - G_0)$ for all k larger than some positive integer m, i.e., termination is guaranteed by time m. Hence, condition (2.6) does not bring in any loss of generality. □

For any fixed $\gamma_2 \in \Gamma_2$, Theorem 2.1 gives us the structure of any optimal infinite-horizon policy at sensor S_1. A similar structure is optimal at sensor S_2 for any fixed $\gamma_1 \in \Gamma_1$. Hence, every p.b.p.o. decision policy (at either of the sensors) has the structure given in Theorem

2.1. The existence of p.b.p.o. solutions can be established using sequential compactness arguments[3] as in [30]. However, unlike the result of [30], optimal sensor decision policies can be found not in the class of SPRTs, but rather in the class of generalized SPRTs (GSPRTs), which as shown above in Theorem 2.1 have monotonically convergent thresholds.

Remark 2.2 At this point it should be noted that the structure of p.b.p.o. decision policies remains the same (as specified in Theorem 2.1) even when the number of sensors is N ($N > 2$). To see this, we fix the decision policies of all of the sensors except sensor S_i. Then, we use a DP argument similar to the one used in establishing Theorem 2.1 to find an optimal policy at S_1. The structure of the optimal policy at S_1 is identical to the one in Theorem 2.1, with modified definitions for G_j and K_j and with $P_j(\tau_2 \leq k)$ replaced by $\prod_{l=2}^{N} P_j(\tau_l \leq k)$, $j = 0, 1$.

2.4 Threshold Computation

We now address the problem of finding optimal GSPRT thresholds numerically. Since the thresholds are known to be monotonically convergent, we could parametrize them as functions of time involving only a few parameters, and then optimize the expected cost over these parameters. This procedure would be facilitated if we could find good approximations for the error probabilities as well as for $E \max(\tau_1, \tau_2)$ in terms of the parameters. The usual Wald approximations, used in [30], cannot be used here; it is well known in sequential analysis that such approximations for time-varying threshold tests are very difficult to obtain [34].

An alternative to the above technique for finding optimal thresholds is the following recursive algorithm, that is motivated by the sequential compactness argument of the previous section (see footnote 3):

[3] An outline of the existence proof is the following: Start with any fixed policy $\gamma_2^{(0)}$ at S_2, and find an optimal policy at S_1, say $\gamma_1^{(1)}$. Then fix the policy of S_1 at $\gamma_1^{(1)}$ and find an optimal policy at S_2, say $\gamma_2^{(1)}$. Continue in this fashion, alternately optimizing at S_1 and S_2 to generate sequences of policies $\{\gamma_1^{(i)}, i = 1, 2, \ldots\}$ and $\{\gamma_2^{(i)}, i = 0, 1, \ldots\}$. These sequences must have convergent subsequences by the sequential compactness of the policy spaces [30]. The policies to which these subsequences converge define a p.b.p.o. solution.

1. Fix the decision policy of S_1 (an SPRT policy would be a reasonable starting point).

2. Run a simulation to obtain the probability distributions of τ_1 and error probabilities at S_1.

3. Use the result of step 2 in a DP recursion at S_2 (with a sufficiently large horizon) to obtain the thresholds at S_2 as described in Theorem 2.1.

4. Run a simulation to obtain the probability distributions of τ_2 and error probabilities at S_2.

5. Use the result of step 4 in a DP recursion at S_1 to obtain a new set of thresholds at S_1 as described in Theorem 2.1.

6. Stop if the policies at S_1 and S_2 have converged. Otherwise, go back to step 2.

If the above algorithm converges, it must converge to a p.b.p.o. solution of problem P2.1. One of these p.b.p.o. solutions is a g.o. solution to P2.1, if a g.o. solution exists.

2.4.1 Optimal SPRT policies

The simplicity of the SPRT structure makes it a good candidate sequential test even when it may not be an optimal test. Hence it is of interest to optimize the expected cost of problem P2.1 over decision policies which use SPRTs at the sensors. However, even if we restrict ourselves to using SPRTs, finding optimal thresholds numerically is difficult because an approximation for $E \max\{\tau_1, \tau_2\}$ is required for this purpose. We have derived one such approximation using characteristic functions, which we describe in the following.

An SPRT policy at sensor S_i has the following form:

$$\begin{aligned} &\text{accept } H_0 \quad \text{if} \quad p_k^i \geq a^i, \\ &\text{accept } H_1 \quad \text{if} \quad p_k^i \leq b^i, \\ &\text{continue} \quad \text{if} \quad b^i < p_k^i < a^i, \end{aligned}$$

where p_k^i denotes the a posteriori probability of H_0 given the observations up to time k at sensor S_i. The thresholds (a^i, b^i) are related to the thresholds (A_i, B_i) of the SPRTs written

in terms of the likelihood ratio [1] in the following way:

$$A_i = \frac{\nu\,(1-a^i)}{(1-\nu)\,a^i}, \qquad B_i = \frac{\nu\,(1-b^i)}{(1-\nu)\,b^i}. \tag{2.7}$$

Now let the error probabilities at S_i under H_0 and H_1 be denoted, respectively, by α_i and β_i. Then Wald's approximations [1] give us the following approximate expressions for α_i and β_i:

$$\alpha_i \approx \frac{1-A_i}{B_i-A_i}, \qquad \beta_i \approx \frac{A_iB_i-A_i}{B_i-A_i}. \tag{2.8}$$

We can also use renewal theory approximations for the error probabilities, which are known to be more accurate than Wald's approximations when the error probabilities are small [35]. With γ_i as defined in Theorem 3.1 of [35], we have the following approximations:

$$\alpha_i \approx \gamma_i/B_i, \qquad \beta_i \approx \gamma_i A_i \tag{2.9}$$

Using (2.7) and (2.8) or (2.9), we obtain an approximate expression for the expected decision cost $E\{W(u_1,u_2;H)\}$ in terms of the thresholds (a^i,b^i).

An approximation for $E\max\{\tau_1,\tau_2\}$ is not obtained as easily, since the basic Wald approximations are only for the first moments of τ_1 and τ_2, and we need the entire distributions to compute this expectation. Fortunately, we could obtain an expression for this expectation in terms of characteristic functions as given below.

$$\begin{aligned} E\{\max(\tau_1,\tau_2)\} &= E\{\tau_2\} + \frac{1}{4\pi}\int_0^\pi \operatorname{cosec}^2(\omega/2)\mathrm{Re}(\phi_2(\omega)(1-\phi_1(\omega)))d\omega \\ &+ \frac{E\{\tau_1\}}{2\pi}\int_0^\pi \left(\mathrm{Re}(\phi_2(\omega)) + \mathrm{Im}(\phi_1(\omega))\cot(\omega/2)\right)d\omega, \end{aligned}$$

where $\phi_i(\omega) = E\{\exp(-\imath\omega\tau_i)\}, i=1,2$ and $\imath$ here is $\sqrt{-1}$. The conditional expectations of τ_1 and τ_2 are given by the standard Wald approximations,

$$\begin{aligned} E\{\tau_i|H_0\} &\approx -2v(\log(A_i)(1-\alpha_i)+\log(B_i)\alpha_i) \quad i=1,2, \\ E\{\tau_i|H_1\} &\approx 2v(\beta_i\log(A_i)+\log(B_i)(1-\beta_i)) \quad i=1,2. \end{aligned}$$

The conditional characteristic functions under H_1 and H_0 can be approximated using Wald's fundamental identity (for details, see [1]). If we define, $t_{10}(\omega)$, $t_{20}(\omega)$, $t_{11}(\omega)$ and $t_{21}(\omega)$, by

$$t_{10}(\omega) = 0.5\left(1 + \sqrt{1 - 8v\omega\imath}\right),$$

$$t_{20}(\omega) = 0.5\left(1 - \sqrt{1 - 8v\omega\imath}\right),$$

$$t_{11}(\omega) = 0.5\left(-1 + \sqrt{1 - 8v\omega\imath}\right),$$

$$t_{21}(\omega) = 0.5\left(-1 - \sqrt{1 - 8v\omega\imath}\right),$$

then we can obtain the following approximations for the conditional characteristic functions under H_0:

$$E\{\exp(-\imath\omega\tau_i)|H_0\} \approx \frac{{A_i}^{t_{20}(\omega)} - {A_i}^{t_{10}(\omega)} + {B_i}^{t_{20}(\omega)} - {B_i}^{t_{10}(\omega)}}{{B_i}^{t_{10}(\omega)}{A_i}^{t_{20}(\omega)} - {A_i}^{t_{10}(\omega)}{B_i}^{t_{20}(\omega)}}. \tag{2.10}$$

Similar expressions hold for $E\{\exp(-\imath\omega\tau_i)|H_1\}, i = 1, 2$, with 0 replaced by 1 in (2.10).

All the approximations given so far can be put together to yield an approximate expression for the total expected cost in terms of the thresholds (a^1, b^1, a^2, b^2). This expression is then minimized over $[0, 1]^4$ to obtain the best SPRT thresholds for problem P2.1.

2.5 Numerical Results

For the numerical results presented in this section, we assume that the observations $\{X_k^1\}_{k=1}^{\infty}$ and $\{X_k^2\}_{k=1}^{\infty}$ are mutually independent i.i.d. Gaussian sequences with mean 0 and variance v under H_0, and mean 1 and variance v under H_1. We also assume that the decision cost is of the form

$$W(u_1, u_2; H) = \begin{cases} 0 & \text{if } u_1 = u_2 = H \\ 1 & \text{if } u_1 \neq u_2 \\ k_e & \text{if } u_1 = u_2 \neq H, 1 < k_e < \infty. \end{cases}$$

In Table 2.1, we present optimization results for the best SPRT thresholds at the sensors. The optimization was done using the approximate expected cost expression derived in the

previous section. Renewal theory approximations were used for the error probabilities. Optimal thresholds and the corresponding expected cost are listed. We have also listed the expected cost for these SPRT policies obtained by Monte-Carlo simulations.

We obtained optimal GSPRT policies by using the recursive algorithm described in the previous section. The algorithm was initialized by using an SPRT policy at S_1. We experimented with a variety of starting policies. A finite horizon of 100 was used for the DP recursions. This was considered to be a reasonable choice for the horizon because in the simulations of the SPRT policies the stopping time at either sensor never exceeded 50. The resulting GSPRT thresholds at the end of 10 iterations are shown in Figure 2.1 for a representative case. The policies at the two sensors converged to the same policy in all cases. Also, the sup-norm difference between the threshold vectors at the 9th and 10th iterates was less than 10^{-3} in all cases. The various choices of starting policies that we experimented with converged to the same GSPRT policy (i.e., the resulting threshold vectors differed in sup norm by less than 10^{-3}) in 10 iterations. Table 2.2 lists the expected cost of the GSPRT policies obtained from the DP recursions as well as by Monte-Carlo simulations for various cases. The expected cost for the corresponding best SPRT policies are repeated in this table for comparison. We note that, as expected, the GSPRT policies perform consistently better than the SPRT policies, and the improvement in performance is about 15-20%. In a practical application, a trade-off must be made between the simplicity of the SPRT policy and the performance gain obtainable with the GSPRT policy.

2.6 Summary

In this chapter, we formulated an extension of the Wald problem to the decentralized case. We used a DP argument to show that optimal sensor decision functions can be found in the class of GSPRTs with monotonically convergent thresholds. We presented some numerical results which illustrate a proposed technique to obtain optimal GSPRT thresholds. We also compared the performance of the GSPRT policies with that of the best SPRT policies.

Table 2.1: Optimization results for the best SPRT policies for the case in which $c = 0.01$, $v = 1.0$, and $k_e = 4.0$.

	SPRT thresholds				Expected Cost	
ν	$1 - a^1$	$1 - a^2$	b^1	b^2	Optimization	Simulation
0.1	4.87×10^{-3}	4.87×10^{-3}	5.14×10^{-3}	5.14×10^{-3}	4.96×10^{-2}	5.90×10^{-2}
0.2	4.25×10^{-3}	4.25×10^{-3}	4.36×10^{-3}	4.36×10^{-3}	5.59×10^{-2}	7.16×10^{-2}
0.3	4.04×10^{-3}	4.04×10^{-3}	4.03×10^{-3}	4.03×10^{-3}	5.98×10^{-2}	7.83×10^{-2}
0.4	3.93×10^{-3}	3.93×10^{-3}	3.91×10^{-3}	3.91×10^{-3}	6.17×10^{-2}	8.15×10^{-2}
0.5	3.88×10^{-3}	3.88×10^{-3}	3.88×10^{-3}	3.88×10^{-3}	6.23×10^{-2}	8.26×10^{-2}

Table 2.2: Comparison of the performance of GSPRT and SPRT policies for the case in which $c = 0.01$, $v = 1.0$ and $k_e = 4.0$.

	Expected Cost			
	SPRT Policy		GSPRT Policy	
ν	Optimization	Simulation	DP Recursion	Simulation
0.1	4.96×10^{-2}	5.90×10^{-2}	4.80×10^{-2}	4.91×10^{-2}
0.2	5.59×10^{-2}	7.16×10^{-2}	5.16×10^{-2}	5.99×10^{-2}
0.3	5.98×10^{-2}	7.83×10^{-2}	5.34×10^{-2}	6.56×10^{-2}
0.4	6.17×10^{-2}	8.15×10^{-2}	5.45×10^{-2}	6.86×10^{-2}
0.5	6.23×10^{-2}	8.26×10^{-2}	5.53×10^{-2}	6.97×10^{-2}

(a)

(b)

Figure 2.1: Optimal GSPRT thresholds for the case in which $c = 0.01$, $v = 1.0$ and $k_e = 4.0$: (a) $1 - a_k$ (b) b_k.

CHAPTER 3

DECENTRALIZED SEQUENTIAL DETECTION—PART II: FUSION CENTER PERFORMING THE SEQUENTIAL TEST

3.1 Introduction

In this chapter, we consider the decentralized sequential hypothesis testing problem in which each sensor sends a sequence of summary messages (local decisions) to a fusion center where a sequential test is carried out to determine the true hypothesis.

Let there be N sensors $S_1, \ldots, S_N$ in the system. At time $k \in \{1, 2, \ldots\}$, sensor S_l observes a random variable X_k^l, and forms a summary message u_k^l of the information it has up to time k. In a general setting, we allow a two-way communication between the sensors and the fusion center as shown in Figure 3.1. In particular, the fusion center could relay past decisions from the other sensors. This means that at time k, each sensor has access to all of its own observations up to time k, and the decisions of all of the other sensors up to time $k-1$.

We now introduce a Bayesian framework for this problem. The two hypotheses H_0 and H_1 are assumed to have known prior probabilities. Also, the conditional joint distributions of the sensor observations under each hypothesis are assumed to be known. A positive cost c is associated with each time step taken for decision making. The fusion center stops receiving additional information at a stopping time τ and makes a final decision δ based on the observations up to time τ. Decision errors are penalized through a decision cost function $W(\delta; H)$. The Bayesian optimization problem is then the minimization of $E\{c\tau + W(\delta; H)\}$ over all admissible decision policies at the fusion center and over all possible choices of local decision functions of the sensors.

Throughout this chapter we shall make the following assumption:

Assumption 3.1 *The sensor observations are independent, in time as well as from sensor to sensor, conditioned on each hypothesis.*

We will also have occasion to use the following extension to Assumption 3.1, especially when we consider infinite-horizon problems.

Assumption 3.2 *The sensor observation sequences are independent (from sensor to sensor) i.i.d. sequences, conditioned on each hypothesis.*

Once the decision rules of the sensors are fixed, the fusion center is faced with a classical sequential detection problem, and hence an optimal decision policy for the fusion center can be found in the class of GSPRTs [36]. Namely, at time k, the fusion center forms a likelihood-ratio L_k (as a function of all of the information it has accumulated) and compares it to two thresholds a_k and b_k. If $L_k \leq a_k$, then H_0 is chosen; if $L_k \geq b_k$ then H_1 is chosen; if $a_k < L_k < b_k$ then the decision is deferred.

Let us now consider the sensor decision functions. Several different cases can be considered depending on the information the sensor decisions are allowed to depend on.

Case A. *System with neither feedback from the fusion center nor local memory*
Here u_k^l is constrained to depend only on X_k^l, i.e.,

$$u_k^l = \phi_k^l(X_k^l).$$

This case was considered in [36], where it was easily shown that person-by-person optimal (p.b.p.o.)[1] sensor decision functions are likelihood ratio tests. The optimal thresholds satisfy a set of coupled equations, which are however almost impossible to solve numerically even if we restrict our attention to relatively short time horizons. Under Assumption 3.2, it may seem that for this case, stationary sensor decision functions are optimal and that an SPRT is optimal at the fusion center. Typically such "stationarity" results are established using dynamic programming (DP) arguments [32]. Unfortunately, dynamic programming cannot be used here because of the nonclassical[2] nature of the information in the system [37, 38], thus leaving this as an open problem.

[1]A set of decision functions is said to be person-by-person optimal if it is not possible to improve the corresponding team performance by unilaterally changing any one of the decision functions. Clearly, globally optimal decision functions are also person-by-person optimal.

[2]We refer to an information structure as nonclassical if, roughly speaking, all of the decision makers in the system do not have the same dynamic information about the past.

Case B. *System with no feedback, but full local memory*

$$u_k^l = \phi_k^l(X_1^l, \ldots, X_k^l).$$

Hashemi and Rhodes [39] considered this case with a finite horizon and argued incorrectly that p.b.p.o. sensor decision functions are likelihood ratio tests (a counterexample can be found in [36] which predates [39]). We point out this mistake in [40], where we also argue that likelihood ratio tests are indeed optimal if we restrict u_k^l to depend on X_k^l and $(u_1^l, \ldots, u_{k-1}^l)$, as given below in Case C.

Case C. *System with no feedback, and local memory restricted to past decisions*

$$u_k^l = \phi_k^l(X_k^l, u_1^l, \ldots, u_{k-1}^l).$$

Here p.b.p.o. sensor decision functions are likelihood ratio tests with thresholds depending on the past decision information. But just as in Cases A and B, we have a nonclassical information pattern and dynamic programming arguments cannot be used.

Case D. *System with full feedback and full local memory*

Here u_k^l is allowed to depend on all of the information that sensor S_l has access to in the setting of Figure 1, i.e.[3],

$$u_k^l = \phi_k^l(X_{[1,k]}^l; u_{[1,k-1]}^1, \ldots, u_{[1,k-1]}^N).$$

Then, as in Case B, likelihood ratio tests are not optimal. Furthermore, we still have a nonclassical information pattern.

Case E. *System with full feedback, but local memory restricted to past decisions*

$$u_k^l = \phi_k^l(X_k^l; u_{[1,k-1]}^1, \ldots, u_{[1,k-1]}^N).$$

For this system, the past (one-step delayed) information at the fusion center and each of the sensors is the same, and is nested at successive stages. This, together with the fact that the cost function depends only on the local decisions (and through them on the observations),

[3]We use the notation $[a, b]$ to represent the set of all time indices between a and b, inclusive.

implies that the information structure for this case is *quasiclassical*. It is well known that stochastic control or team problems with such an information structure are tractable via DP arguments [37, 38].

In the remainder of this chapter we study Case E in detail. As we will show, definite progress can be made in the analysis of this case. In Section 3.2, we provide a formal mathematical description of the problem. In Section 3.3, we provide a useful characterization of sensor decision functions. In Section 3.4, we consider a finite-horizon version of the problem and establish the optimality of likelihood-ratio tests at the sensors. Then, in Section 3.5, we study the infinite horizon optimization problem and show that stationary decision functions are optimal at the sensors and that an optimal fusion center policy can be characterized by two thresholds. In Section 3.6, we provide some numerical results. Finally, in Section 3.7, we summarize the main points.

3.2 Mathematical Description

We begin with a formal description of the decentralized sequential detection problem we wish to analyze here.

1. The hypothesis is denoted by a binary random variable H which takes on values H_0 and H_1, with prior probabilities ν and $1-\nu$, respectively.

2. There are N sensors in the system. The observation sequence received by sensor S_l is denoted by $\{X_k^l\}_{k=1}^{\infty}$, where k denotes the time index. Each observation at sensor S_l comes from a set $\mathcal{X}_l$. The sequences $\{X_k^1\}_{k=1}^{\infty}$, $\{X_k^2\}_{k=1}^{\infty}$, $\ldots$, $\{X_k^N\}_{k=1}^{\infty}$ are *independent*, i.i.d. sequences, when conditioned on each hypothesis. Let $P_{l|H_j}$ be the probability measure on $\mathcal{X}_l$ that describes the conditional distribution of X_1^l given H_j.

3. At time k, sensor S_l sends, to the fusion center, a local decision u_k^l which takes values in the finite set $\{1, \ldots, D_l\}$. Past decision information from all of the sensors is available at each sensor for local decision making. We denote the past decision information at time k by

I_{k-1}, which is given by

$$I_{k-1} = \{u^1_{[1,k-1]}, u^2_{[1,k-1]}, \ldots, u^N_{[1,k-1]}\},$$

with the understanding that I_0 is the null set. Now, let

$$\mathcal{D} = \{1, \ldots, D_1\} \times \{1, \ldots, D_2\} \times \cdots \times \{1, \ldots, D_N\}.$$

Then the *local decision function* (LDF) at sensor S_l at time k is a measurable mapping from $\mathcal{X}_l \times \mathcal{D}^{k-1}$ to $\{1, \ldots, D_l\}$. We denote this mapping by φ^l_k. The local decision u^l_k is then given by

$$u^l_k = \varphi^l_k(X^l_k; I_{k-1}).$$

But for a particular realization i_{k-1} of I_{k-1}, the LDF φ^l_k can be considered to be a mapping from $\mathcal{X}_l$ to $\{1, \ldots, D_l\}$, which we denote by $\varphi^l_{k,i_{k-1}}$, i.e.,$\varphi^l_k(\,\cdot\,; i_{k-1}) \equiv \varphi^l_{k,i_{k-1}}(\,\cdot\,)$. The set of all LDFs at time k is represented by the vector

$$\boldsymbol{\varphi}_k = \left(\varphi^1_k, \ldots, \varphi^N_k\right).$$

4. The fusion center performs a sequential test based on the information it receives from the sensors. That is, the policy γ of the fusion center consists of selecting a stopping time τ and a final decision $\delta \in \{0, 1\}$ based on the information up to time τ.

5. Decision errors are penalized through a cost function $W(\delta, H)$. For most of the analysis, we will assume that the cost function W is of the form: $W(0, H_0) = W(1, H_1) = 0$, and $W(0, H_1) = L_0,\ W(1, H_0) = L_1$, where L_0 and L_1 are positive. Also, each time step taken for decision making is assumed to cost a positive amount c.

The total expected cost resulting from the sequential procedure described above is $E\{c\tau + W(\delta, H)\}$. The problem that we wish to solve can now be stated as follows:

Problem P3.1:

Minimize $E\{c\tau + W(\delta, H)\}$ over all admissible decision policies at the fusion center and over all possible choices of local decision functions at each of the sensors. □

3.3 Local Decision Functions

The decision function $\varphi^l_{k,i_{k-1}}$ defined in Section 3.2 is a mapping from $\mathcal{X}_l$ to $\{1,\ldots,D_l\}$. Let Φ^l denote the set of all mappings from $\mathcal{X}_l$ to $\{1,\ldots,D_l\}$. We will refer to these mappings as *decision functions* in the sequel. Now, consider a representative element $\phi^l \in \Phi^l$, and let X^l denote the "generic" random variable in the i.i.d. sequence $\{X^l_k\}_{k=1}^{\infty}$. Then, we define the following:

$$\begin{aligned} q^j_{\phi^l}(d_l) &:= \text{Prob}(\phi^l(X^l) = d_l | H_j), \quad d_l = 1,\ldots,D_l, \quad j = 0,1; \\ q^j_{\phi^l} &:= \left(q^j_{\phi^l}(1),\ldots,q^j_{\phi^l}(D_l)\right), \quad j = 0,1; \\ q_{\phi^l} &:= \left(q^0_{\phi^l}, q^1_{\phi^l}\right). \end{aligned}$$

The vector q_{ϕ^l} describes the conditional distributions of $\phi^l(X^l)$, conditioned on each of the hypotheses. Let $Q^l := \left\{q_{\phi^l} \mid \phi^l \in \Phi^l\right\}$. We state the following result which was proved in [41] in the context of optimal likelihood ratio quantizers.

Proposition 3.1 *The set Q^l is a compact subset of $[0,1]^{2D_l}$, for $l = 1,\ldots,N$.* □

To utilize this result in our framework, we concatenate the mappings ϕ^l, $l = 1,\ldots N$, into the vector $\boldsymbol{\phi} = \left(\phi^1,\ldots,\phi^N\right)$, and define

$$\boldsymbol{q_\phi} := \left(q_{\phi^1},\ldots,q_{\phi^N}\right).$$

Then $\boldsymbol{q_\phi}$ belongs to the set $\boldsymbol{Q} = Q^1 \times \cdots \times Q^N$. By Proposition 3.1, $\boldsymbol{Q}$ is a compact set. Now suppose that $J : [0,1]^{2D_1\times\cdots\times 2D_N} \mapsto I\!R$ is a continuous function, and that the cost of using the decision function vector $\boldsymbol{\phi}$ is given by $J(\boldsymbol{q_\phi})$. Then by the Weierstrass theorem, Proposition 3.1 implies the existence of a decision function vector (say $\hat{\boldsymbol{\phi}}$) that minimizes J over the set $\boldsymbol{Q}$.

Now, since $\phi^l_{k,i_{k-1}} \in \Phi^l$, $l = 1,\ldots N$, the vector $\boldsymbol{q}_{\boldsymbol{\phi}_{k,i_{k-1}}}$ is well-defined and describes the joint distribution of the observation vector $\boldsymbol{u}_k = (u^1_k,\ldots,u^N_k)$, conditioned on each hypothesis

and on the event that $I_{k-1} = i_{k-1}$. Note that two LDFs φ_k^l and $\tilde{\varphi}_k^l$ are equivalent, i.e., their use results in the same expected cost for the sequential test by the fusion center, if

$$q_{\varphi_{k,I_{k-1}}^l} = q_{\tilde{\varphi}_{k,I_{k-1}}^l} \quad \text{a.e.}$$

That is, the LDFs for our problem are completely characterized by their corresponding conditional distribution vectors.

3.4 Finite-horizon Optimization

Before we address the solution of the infinite-horizon optimization problem P3.1, we study a finite-horizon version of it in which the stopping time τ is restricted to a finite interval, say $[0, T]$. In this case, the cost of the sequential procedure is a function of I_T (which in turn depends on all of the LDFs up to time T), as well as the decision policy γ of the fusion center and the hypothesis. We denote this cost by $G_\gamma(I_T, H)$. Let $\boldsymbol{X}_{[1,T]}$ denote the set of all observations up to time T, i.e., $\{X_{[1,T]}^1, \ldots, X_{[1,T]}^N\}$. Then, the finite-horizon optimization problem can then be stated as follows:

Problem P3.2:

Minimize

$$E_{\boldsymbol{X}_{[1,T]},\, H}\, G_\gamma\left(u_{[1,T]}^1, \ldots, u_{[1,T]}^N; H\right)$$

over all possible choices of γ and all possible choices of LDFs $\varphi_k^l, l = 1, \ldots, N,\ k = 1, \ldots, T$.

□

Now, before we consider globally optimal solutions to this problem, we first study the common structure of all p.b.p.o. LDFs. This common structure would obviously be valid for globally optimal LDFs as well.

3.4.1 The structure of optimal LDFs

We can characterize p.b.p.o. LDFs as follows. We first fix l, $1 \le l \le N$, and k, $1 \le k \le T$. Then we fix the policy γ and all the LDFs in the set $\{\varphi_m^j,\ j = 1, \ldots, N,\ m = 1, \ldots, T\}$,

except φ_k^l. The expected cost we wish to minimize is then a function of φ_k^l alone, say $R_{\varphi_k^l}$. The expectation needed for $R_{\varphi_k^l}$ can be computed in two steps as given below:

$$\begin{aligned} R_{\varphi_k^l} &= E_{u^1_{[1,k-1]},\ldots,u^N_{[1,k-1]},\, X_k^l,\, H}\,\{ \\ &\quad E_{X^l_{[k+1,T]},X^1_{[k,T]},\ldots,X^{l-1}_{[k,T]},X^{l+1}_{[k,T]},\ldots,X^N_{[k,T]}\,\mid H}\,\{ \\ &\quad G_\gamma\left(\varphi_k^l(X_k^l;I_{k-1}), I_{k-1}, u_k^1,\ldots,u_k^{l-1},u_k^{l+1},\ldots,u_k^N, u^1_{[k+1,T]},\ldots,u^N_{[k+1,T]}; H\right)\ \}\}\,. \end{aligned}$$

In the inner expectation above, we do not need to condition on observations up to time $k-1$ because of the conditional independence assumption stated in Section 3.2. Also, the outer expectation is taken with respect to the local decisions, since the local decision functions up to time $k-1$ are fixed. The inner expectation in the equation above is a function of I_{k-1}, H and $\varphi_k^l(X_k^l;I_{k-1})$, say $K\left(\varphi_k^l(X_k^l;I_{k-1}), I_{k-1}; H\right)$. Therefore,

$$\begin{aligned} R_{\varphi_k^l} &= E_{I_{k-1},\, X_k^l,\, H}\left\{K\left(\varphi_k^l(X_k^l;I_{k-1}), I_{k-1}; H\right)\right\} \\ &= E_{I_{k-1},\, X_k^l}\Big\{\mathrm{Prob}(H=H_0|I_{k-1},X_k^l)\,K\left(\varphi_k^l(X_k^l;I_{k-1}), I_{k-1}; H_0\right) \\ &\qquad +\mathrm{Prob}(H=H_1|I_{k-1},X_k^l)\,K\left(\varphi_k^l(X_k^l;I_{k-1}), I_{k-1}; H_1\right)\Big\}\,. \end{aligned}$$

Minimizing $R_{\varphi_k^l}$ with respect to φ_k^l is equivalent to minimizing the quantity inside the expectation almost everywhere. Hence, every p.b.p.o. solution $\hat{\varphi}_k^l$ for the LDF of sensor S_l at time k (when it *exists*) satisfies the equation

$$\begin{aligned} \hat{\varphi}_k^l(X_k^l;I_{k-1}) = \arg\min_{d_l\in\{1,\ldots,D_l\}} \quad & \Big\{\mathrm{Prob}(H=H_0|I_{k-1},X_k^l)\,K\left(d_l, I_{k-1}; H_0\right) \\ & +\mathrm{Prob}(H=H_1|I_{k-1},X_k^l)\;K\left(d_l, I_{k-1}; H_1\right)\Big\} \quad \text{a.e. (3.1)} \end{aligned}$$

Our goal in this section is to show that p.b.p.o. local decision functions (when they exist) can be found within a structured class of decision functions admitting a finite-dimensional parametrization. To this end, we first define $L_l : \mathcal{X}_l \mapsto [0,\infty]$ as the *likelihood ratio* of $P_{i|H_1}$ with respect to $P_{i|H_0}$. In particular, if $\mathcal{X}_l$ is a Euclidean space and if the conditional probability density function of X_k^l given H_j is f_j^l, then L_l is given by

$$L_l(X_k^l) = \frac{f_1^l(X_k^l)}{f_0^l(X_k^l)} \quad \text{w.p. 1.}$$

We now define a class of decision functions, based on this likelihood ratio, that can be parametrized by a set of thresholds[4].

Definition 3.4 (a) A decision function $\varphi^l : \mathcal{X}_l \mapsto \{1, \ldots, D_l\}$ is called a *monotone likelihood ratio test* (MLRT) if there exist thresholds $\lambda_1, \ldots, \lambda_{D_l-1}$ satisfying $0 \leq \lambda_1 \leq \lambda_2 \leq \ldots \leq \lambda_{D_l-1} \leq \infty$ such that

$$\varphi^l(x) = d_l \text{ only if } L_l(x) \in I_{d_l}, \quad d_l = 1, \ldots, D_l,$$

where $I_1 = [0, \lambda_1]$, $I_{D_l} = [\lambda_{D_l-1}, \infty]$, and $I_{d_l} = [\lambda_{d_l-1}, \lambda_{d_l}]$, $d_l = 2, \ldots, D_l - 1$.
(b) A decision function $\varphi^l : \mathcal{X}_l \mapsto \{1, \ldots , D_l\}$ is called a *likelihood ratio test* (LRT) if there exists a permutation mapping $\Sigma : \{1, \ldots, D_l\} \mapsto \{1, \ldots, D_l\}$ such that $\Sigma \circ \varphi^l$ is a monotone likelihood ratio test.

Proposition 3.2 *Person-by-person optimal local decision functions (when they exist) can be found in the class of LRTs, with thresholds that depend on the past decision information.*

Proof: We know that a p.b.p.o. solution $\hat{\varphi}_k^l$ for the LDF of sensor S_l at time k satisfies (3.1). Using Bayes rule, we have

$$\frac{\text{Prob}(H = H_1 | I_{k-1}, X_k^l)}{\text{Prob}(H = H_0 | I_{k-1}, X_k^l)} = L_l(X_k^l) \frac{(1 - p_{k-1})}{p_{k-1}}, \quad \text{w.p. } 1, \tag{3.2}$$

where p_k denotes the posteriori probability of H_0 given the decision information up to time k, i.e.,

$$p_k = \text{Prob}(H = H_0 | I_k), \quad k = 0, \ldots T.$$

From (3.2) it follows that $\hat{\varphi}_k^l$ satisfies

$$\hat{\varphi}_k^l(X_k^l; I_{k-1}) = \begin{cases} \arg\min_{d_l \in \{1,\ldots,D_l\}} \Big\{ (1 - p_{k-1}) K\left(d_l, I_{k-1}; H_1\right) L_l(X_k^l) \\ \qquad\qquad + p_{k-1} K\left(d_l, I_{k-1}; H_0\right) \Big\} & \text{if } L_l(X_k^l) < \infty \\ \arg\min_{d_l \in \{1,\ldots,D_l\}} \; K\left(d_l, I_{k-1}, H_1\right) & \text{if } L_l(X_k^l) = \infty. \end{cases}$$

[4]Similar definitions can be found in [42].

From this it should be clear that (see Figure 3.2) there exists a solution for $\hat{\varphi}_k^l$ in the class of LRTs with thresholds that depend on I_{k-1}. □

3.4.2 A sufficient statistic for DP

As we mentioned earlier, the information structure in the system under consideration is of a quasiclassical nature. Hence, we would expect that a sufficient statistic for a dynamic programming (DP) solution to problem P3.2 is the posteriori probability defined earlier, i.e.,

$$p_k = \text{Prob}(H = H_0 | I_k), \quad k = 0, \ldots, T.$$

Using the independence assumptions, a recursion for p_k can be obtained quite readily. Before we proceed to write the recursion equations, we introduce two functions, $g : \mathcal{D} \times \boldsymbol{Q} \times [0,1] \mapsto [0,1]$ and $f : \mathcal{D} \times \boldsymbol{Q} \times [0,1] \mapsto [0,1]$, as follows:
For $\boldsymbol{d} = (d_1, \ldots, d_N) \in \mathcal{D}$, $\boldsymbol{\phi} \in \Phi^1 \times \cdots \times \Phi^N$, and $p \in [0,1]$,

$$g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p) := p\, q^0_{\phi^1}(d_1) \cdots q^0_{\phi^N}(d_N),$$

and

$$f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p) := g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p) + (1-p)\, q^1_{\phi^1}(d_1) \cdots q^1_{\phi^N}(d_N).$$

Note that $f(\,\cdot\,; \boldsymbol{q}_{\boldsymbol{\varphi}_{k,I_{k-1}}}; p_{k-1})$ is the joint conditional distribution of $\boldsymbol{u}_k = (u_k^1, \ldots, u_k^N)$, given I_{k-1}.
We now give the recursion for p_k. For $k = 0, \ldots, T$,

$$\begin{aligned}
p_{k+1} &= \text{Prob}(H = H_0 | I_{k+1}) \\
&= \text{Prob}(H = H_0 | u_{k+1}^1, \ldots, u_{k+1}^N,\ I_k) \\
&= \frac{\text{Prob}(H = H_0 | I_k)\, p(u_{k+1}^1, \ldots, u_{k+1}^N | H_0,\ I_k)}{p(u_{k+1}^1, \ldots, u_{k+1}^N | I_k)} \\
&= \frac{g(\boldsymbol{u}_{k+1}; \boldsymbol{q}_{\boldsymbol{\varphi}_{k+1,I_k}}; p_k)}{f(\boldsymbol{u}_{k+1}; \boldsymbol{q}_{\boldsymbol{\varphi}_{k+1,I_k}}; p_k)},
\end{aligned} \tag{3.3}$$

with $p_0 = \nu$. However, we would not find this recursion useful unless we can show that the RHS of (3.3) depends on I_k only through p_k, i.e., that p_k is indeed a sufficient statistic for P3.2.

3.4.3 Finite-horizon DP

The DP equations for problem P3.2 are derived as follows. The (minimum) expected cost-to-go at time k, $0 \leq k \leq T$, is a function of the information available to the fusion center at time k, i.e., I_k, which we denote by $\tilde{J}_K^T(I_k)$. It is easily seen that

$$\tilde{J}_T^T(I_T) = \min\{(1-p_T)L_0, p_T L_1\},$$

where the first (respectively, second) term in the above minimum is the conditional expected cost of choosing H_0 (respectively, H_1), given I_T.

For $0 \leq k \leq T$, we have the following backward recursion:

$$\tilde{J}_k^T(I_k) = \min\left\{(1-p_k)L_0, p_k L_1, c + \inf_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}} E\left\{\tilde{J}_{k+1}^T(I_{k+1})|I_k\right\}\right\},$$

where the third term above is the minimum expected cost of continuing conditioned on I_k.

Proposition 3.3 *(i) For each k, $0 \leq k \leq T$, the function $\tilde{J}_k^T(I_k)$ can be written as a function of only p_k, say $J_k^T(p_k)$.*

(ii) For each k, $0 \leq k \leq T-1$, the search for optimal LDFs at time $k+1$ can be restricted to a class of decision functions that depend only on p_k.

Proof: Clearly, $\tilde{J}_T^T(I_T)$ is a function of only p_T, say $J_T^T(p_T)$. We now make the following induction argument. For any k, $0 \leq k \leq T-1$, suppose that $\tilde{J}_{k+1}^T(I_{k+1})$ is a function of only p_{k+1}, say $J_{k+1}^T(p_{k+1})$. Now,

$$\begin{aligned}\tilde{J}_k^T(I_k) \;=\; & \min\{(1-p_k)L_0, p_k L_1, \\ & c + \inf_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}} \sum_{\boldsymbol{d} \in \mathcal{D}} J_{k+1}^T\left(\frac{g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)}{f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)}\right) f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)\Bigg\}.\end{aligned}$$

In this equation, it is clear that the infimizing conditional distribution vector $\boldsymbol{q}_{\boldsymbol{\phi}^\star}$ is a function of only p_k. Both parts of the proposition follow from this fact. □

The main consequence of Proposition 3.3 is that we do not lose optimality if we restrict the local decision u_k^l to be a function of only X_k^l and p_{k-1}. From here on, we impose this restriction. Then, by a possible abuse of notation,

$$u_k^l = \phi_k^l(X_k^l; p_k). \tag{3.4}$$

For fixed $p \in [0,1]$, the mapping $\phi_k^l(\,\cdot\,; p)$ belongs to the set Φ^l defined earlier. We denote this decision function by $\phi_{k,p}^l$, i.e.,

$$\phi_k^l(\,\cdot\,; p) \equiv \phi_{k,p}^l(\,\cdot\,).$$

With the LDFs defined as in (3.4), we obtain the following useful recursion for p_k. For $k = 0, \ldots T-1$, we have

$$p_{k+1} = \frac{g(\boldsymbol{u}_{k+1}; \boldsymbol{q}_{\boldsymbol{\phi}_{k+1,p_k}}; p_k)}{f(\boldsymbol{u}_{k+1}; \boldsymbol{q}_{\boldsymbol{\phi}_{k+1,p_k}}; p_k)} \tag{3.5}$$

with $p_0 = \nu$. All of the decision makers in the system need to retain only the sufficient statistic p_k, which they can easily update using (3.5).

For completeness, we rewrite the finite-horizon DP equations in terms of the redefined LDFs:

$$J_T^T(p_T) = \min\{(1-p_T)\, L_0,\, p_T L_1\}, \tag{3.6}$$

and for $k = 0, \ldots, T-1$,

$$J_k^T(p_k) = \min\left\{(1-p_k)\, L_0,\, p_k L_1, c + A_k^T(p_k)\right\}, \tag{3.7}$$

where

$$A_k^T(p_k) := \inf_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}} \sum_{\boldsymbol{d} \in \mathcal{D}} J_{k+1}^T\left(\frac{g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)}{f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)}\right)\, f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k). \tag{3.8}$$

3.4.4 Finite-horizon policy of the fusion center

Our goal in this section is to use the DP Equations (3.6) - (3.8) to find the structure of an optimal finite-horizon policy of the fusion center. To this end, we first present some useful properties of the functions J_k^T and the functions A_k^T in the following lemmas.

Lemma 3.1 *The functions $J_k^T(p)$ and $A_k^T(p)$ are nonnegative concave functions of p, for $p \in [0,1]$.* □

Lemma 3.2 *The functions $J_k^T(p)$ and $A_k^T(p)$ are monotonically nondecreasing in k, that is, for each $p \in [0,1]$,*

$$J_k^T(p) \leq J_{k+1}^T(p), \quad 0 \leq k \leq T-1.$$

$$A_k^T(p) \leq A_{k+1}^T(p), \quad 0 \leq k \leq T-2.$$

□

Lemma 3.3 *The functions $A_k^T(p)$ satisfy the following property:*

$$A_k^T(0) = A_k^T(1) = 0.$$

□

Lemmas 3.2 and 3.3 are easily proven by simple induction arguments. The proof of Lemma 3.1 is not as straightforward and is given in Section 3.8.

If we now assume that the condition given below holds,

$$A_{T-1}^T\left(\frac{L_0}{L_1+L_0}\right) \leq \frac{L_0 L_1}{L_1+L_0}, \tag{3.9}$$

then Lemmas 3.1-3.3 give us the following threshold property of an optimal finite-horizon fusion center policy (see Figure 3.3; also see Section 3.5 of [32] for a similar analysis).

Theorem 3.1 *Let condition (3.9) hold. Then an optimal finite-horizon fusion center policy has the form*

$$\begin{aligned} &\textit{accept } H_0 \quad &&\textit{if} \quad p_k \geq a_k^T, \\ &\textit{accept } H_1 \quad &&\textit{if} \quad p_k \leq b_k^T, \\ &\textit{continue} \quad &&\textit{if} \quad b_k^T < p_k < a_k^T, \end{aligned}$$

where the scalars a_k^T, b_k^T, $k = 0, 1, ...T-1$, *are determined from the relations*

$$\begin{aligned} L_0\left(1 - b_k^T\right) &= c + A_k^T(b_k^T), \\ L_1 a_k^T &= c + A_k^T(a_k^T). \end{aligned}$$

Furthermore, $\{a_k^T\}_{k=0}^{T-1}$ *is a nonincreasing sequence and* $\{b_k^T\}_{k=0}^{T-1}$ *is a nondecreasing sequence.*

□

Remark 3.1 If condition (3.9) does not hold, then the thresholds a_k^T and b_k^T of Theorem 3.1 are both identically equal to $L_0/(L_0 + L_1)$ for all k greater than some m, $1 \leq m < T$, which essentially reduces the finite horizon to m. Hence, condition (3.9) does not impose any restrictions on the problem parameters. □

3.4.5 Optimal finite-horizon LDFs

The DP equations (3.6)-(3.8) can also be used to find optimal LDFs stagewise, starting from time T and going backwards. The concavity of the cost-to-go function J_{k+1}^T implies that the function

$$\sum_{\boldsymbol{d} \in \mathcal{D}} J_{k+1}^T\left(\frac{g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)}{f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)}\right) f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)$$

is continuous in $\boldsymbol{q}_{\boldsymbol{\phi}}$. By Proposition 3.1, this fact implies the *existence* of optimal LDFs at time $k+1$.

We showed earlier (in Section 3.4.1) that the search for globally optimal LDFs at time $k+1$ can be restricted to the set of LRTs with thresholds depending on I_k. Propositions 3.1 and 3.3 further show that globally optimal LDFs at time $k+1$ can be *found* in a class of LRTs with thresholds depending only on p_k. Now, suppose $\hat{\phi}_{k+1}^l$ is a globally optimal LDF for sensor l at time $k+1$. Then we can replace $\hat{\phi}_k^l$ by $\tilde{\phi}_k^l = \Sigma_l \circ \hat{\phi}_{k,p_k}^l$, where Σ_l is a permutation mapping that makes $\hat{\phi}_{k,p_k}^l$ a *monotone likelihood ratio test* (MLRT), without changing the value of $E\{J_{k+1}^T(p_{k+1}) \,|\, p_k\}$. Hence, globally optimal LDFs at time $k+1$ can be found in the smaller class of MLRTs with thresholds that depend on p_k.

Now suppose an MLRT $\tilde{\phi}^l_{k,p_k}$ is characterized by the thresholds $\lambda_1(p_k), \ldots, \lambda_{D_l-1}(p_k)$. Then

$$q^j_{\tilde{\phi}^l_{k,p_k}}(d_l) = \text{Prob}\left(L(X^l_k) \in [\lambda_{d_l-1}(p_k), \lambda_{d_l}(p_k)] \mid H_j\right),$$

with the understanding that $\lambda_0(p_k) = 0$ and $\lambda_{D_l}(p_k) = \infty$. Hence, the minimization to obtain $A^T_k(p_k)$ in (3.8) can be done over $|\mathcal{D}|$ thresholds.

Finally, if we define the set $\boldsymbol{Q}_M$ by

$$\boldsymbol{Q}_M = \left\{ \boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q} \,:\, \boldsymbol{\phi} \text{ is a vector of MLRTs } \right\},$$

then $A^T_k(p_k)$ can be written as

$$A^T_k(p_k) := \min_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}_M} \sum_{\boldsymbol{d} \in \mathcal{D}} J^T_{k+1}\left(\frac{g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)}{f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k)}\right) f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p_k).$$

3.5 Infinite-horizon Optimization

In order to solve the original optimization problem P3.1, we need to remove the restriction that τ belongs to a finite interval, by letting $T \to \infty$. Toward this end, we first note the inequality

$$J^{T+1}_k(p) \leq J^T_k(p),$$

which holds because the set of stopping times increases with T. Furthermore, by leaving out the third term in (3.7), we obtain

$$0 \leq J^T_k(p) \leq \eta(p), \quad \forall T, \quad \text{and } \forall k \leq T,$$

where

$$\eta(p) = \min\{L_1 p, L_0(1-p)\}. \tag{3.10}$$

The fact that J^T_k is bounded below implies that, for each finite k, the following limit

$$\lim_{T\to\infty, T>k} J^T_k(p) = \inf_{T>k} J^T_k(p) =: J^\infty_k(p)$$

is well-defined. Also, due to the i.i.d. nature of the observations, a time-shift argument easily shows that

$$J_k^\infty(p) = J_{k+1}^\infty(p),$$

for all k, and we can denote the common value by $J(p)$, which we will refer to as the infinite-horizon cost-to-go function.

Now, by the Dominated Convergence Theorem, the following limit is well-defined for all k:

$$\lim_{T\to\infty} A_k^T(p) = \min_{\boldsymbol{q}_{\boldsymbol{\phi}}\in\boldsymbol{Q}_M} \sum_{\boldsymbol{d}\in\mathcal{D}} J\left(\frac{g(\boldsymbol{d};\boldsymbol{q}_{\boldsymbol{\phi}};p)}{f(\boldsymbol{d};\boldsymbol{q}_{\boldsymbol{\phi}};p)}\right) f(\boldsymbol{d};\boldsymbol{q}_{\boldsymbol{\phi}};p).$$

This limit, which is independent of k, is denoted by $A_J(p)$. It follows that the infinite-horizon cost-to-go function $J(p)$ satisfies the Bellman equation

$$J(p) = \min\left\{L_1 p, L_0(1-p), c + A_J(p)\right\}. \tag{3.11}$$

We note that the optimum cost for problem P3.1 is $J(\nu)$.

3.5.1 The structure of an optimal fusion center policy

If we compute the infinite-horizon cost-to-go function $J(p)$, $p \in [0,1]$, then an optimal policy of the fusion center can be obtained from the RHS of (3.11). However, it is possible to obtain the qualitative structure of an optimal fusion center policy without actually computing $J(p)$. To this end, we state the following result, whose proof follows by taking limits as $T \to \infty$ in Lemmas 3.1-3.3.

Lemma 3.4 *The functions $J(p)$ and $A_J(p)$ are nonnegative concave functions of p, $p \in [0,1]$. Furthermore, they satisfy the end-point conditions*

$$J(0) = J(1) = A_J(0) = A_J(1) = 0.$$

□

From Lemma 3.4 it is clear that, provided the condition

$$J\left(\frac{L_0}{L_1+L_0}\right) < \frac{L_1 L_0}{L_1+L_0} \tag{3.12}$$

holds, an optimal policy of the supervisor will have the threshold property given in the theorem below (see Figure 3.4; also see Section 6.3 of [32] for a similar analysis).

Theorem 3.2 *Let condition (3.12) hold. Then an optimal fusion center policy for problem P3.1 has the form*

$$\begin{array}{lll} \textit{accept } H_0 & \textit{if} & p_k \geq a \\ \textit{accept } H_1 & \textit{if} & p_k \leq b \\ \textit{continue taking observations} & \textit{if} & b < p_k < a, \end{array}$$

where the thresholds a and b are determined from the relations

$$\begin{aligned} L_0(1-b) &= c + A_J(b), \\ L_1 a &= c + A_J(a). \end{aligned}$$

□

Remark 3.2 It should be noted that if condition (3.12) does not hold, then it would be optimal for the fusion center to ignore all of the data it receives from the sensors, and base its decision solely on the value of the prior probability ν. Hence condition (3.12) does not bring any loss of generality to the result of Theorem 3.2 above. □

3.5.2 Uniqueness of $J(p)$ and its consequences

Let $S \subset C[0,1]$ be the set of all concave functions on $[0,1]$ that are bounded (in sup norm) by the function $\eta(p)$, $p \in [0,1]$, defined in (3.10). For $G \in S$, we define

$$W_G(\boldsymbol{q}_{\boldsymbol{\phi}}; p) := \sum_{\boldsymbol{d} \in \mathcal{D}} G\left(\frac{g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p)}{f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p)}\right) \; f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p).$$

It is clear that the infinite-horizon cost-to-go function J belongs to the set S. Furthermore, the Bellman equation that J satisfies can be written as

$$J(p) = \min\left\{ L_1 p, L_0(1-p), c + \min_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}_M} W_J(\boldsymbol{q}_{\boldsymbol{\phi}}; p) \right\}.$$

Then, we define the mapping $\mathcal{T} : S \mapsto S$ by

$$\mathcal{T}\, G(p) = \min \left\{ L_1 p, L_0(1-p), c + \min_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}_M} W_G(\boldsymbol{q}_{\boldsymbol{\phi}}; p) \right\}, \quad \text{for } G \in S.$$

Theorem 3.3 *The infinite-horizon cost-to-go function J is the unique fixed point of the mapping $\mathcal{T}$.*

Proof: Let G be any fixed point of $\mathcal{T}$, and let $\boldsymbol{\phi}_p^\star$ be such that

$$\boldsymbol{q}_{\boldsymbol{\phi}_p^\star} = \arg \min_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}_M} W_G(\boldsymbol{q}_{\boldsymbol{\phi}}; p).$$

Fix $p_0 = \nu \in [0,1]$, and let $p_1, p_2, \ldots$, be defined recursively by

$$p_{k+1} = \frac{g(\boldsymbol{u}_{k+1}; \boldsymbol{q}_{\boldsymbol{\phi}^\star_{p_k}}; p_k)}{f(\boldsymbol{u}_{k+1}; \boldsymbol{q}_{\boldsymbol{\phi}^\star_{p_k}}; p_k)}.$$

Now define a stopping time $\mathcal{N}$ and a decision rule $\delta_{\mathcal{N}}$ as follows:

$$\mathcal{N} = \min \left\{ k \geq 0 \;\middle|\; \eta(p_k) \leq c + W_G(\boldsymbol{q}_{\boldsymbol{\phi}^\star_{p_k}}; p_k) \right\},$$

and

$$\delta_{\mathcal{N}} = \begin{cases} 1 & \text{if } L_1 p_{\mathcal{N}} \leq L_0(1-p_{\mathcal{N}}) \\ 0 & \text{if } L_1 p_{\mathcal{N}} > L_0(1-p_{\mathcal{N}}). \end{cases}$$

From the definition of $\mathcal{N}$ and the fact that G is a fixed point of $\mathcal{T}$, we obtain the following relations:

$$\begin{aligned} G(\nu) &= c + E\{G(p_1)\} \\ G(p_1) &= c + E\{G(p_2)|I_1\} \\ &\vdots \\ G(p_{\mathcal{N}-1}) &= c + E\{G(p_{\mathcal{N}})|I_{\mathcal{N}-1}\} \\ G(p_{\mathcal{N}}) &= \eta(p_{\mathcal{N}}). \end{aligned}$$

Substituting backwards and taking expectations, we obtain

$$G(\nu) = E\left\{c\mathcal{N} + W(\delta_{\mathcal{N}}, H)\right\} \geq J(\nu),$$

where the last inequality follows from the definition of J.

To show the reverse inequality, we first note that for each $p \in [0,1]$,

$$G(p) \leq \eta(p) = J_T^T(p), \quad \forall T.$$

Now fix T, and suppose that for some $m < T-1$, $J_{m+1}^T \geq G(p)$. Then

$$\begin{aligned} J_m^T(p) &= \min\left\{\eta(p), c + \min_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}_M} W_{J_{m+1}^T}(\boldsymbol{q}_{\boldsymbol{\phi}}; p)\right\} \\ &\geq \min\left\{\eta(p), c + \min_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}_M} W_G(\boldsymbol{q}_{\boldsymbol{\phi}}; p)\right\} \\ &= G(p). \end{aligned}$$

By induction, it follows that for each $p \in [0,1]$,

$$J_k^T(p) \geq G(p), \quad \forall T, \text{ and } \forall k \leq T.$$

Fixing k and taking the limit as $T \to \infty$ in the above equation, we obtain

$$J(p) \geq G(p). \qquad \square$$

The first important consequence of Theorem 3.3 is that $J(p)$ can be obtained by successive approximation. We can show, using an induction argument, that

$$\mathcal{T}^{n+1}\eta(p) \leq \mathcal{T}^n\eta(p), \text{ for each } p \in [0,1].$$

This means that $\mathcal{T}^n\eta$ converges monotonically to J as $n \to \infty$.

3.5.3 Optimal infinite-horizon LDFs

Theorem 3.3 also implies that a stationary vector of LDFs is optimal for the infinite-horizon problem P3.1, as the following argument shows. Let $\boldsymbol{\phi}_p^\star$ be such that

$$\boldsymbol{q}_{\boldsymbol{\phi}_p^\star} = \arg \min_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}_M} W_J(\boldsymbol{q}_{\boldsymbol{\phi}}; p),$$

where $J(p)$ is the infinite-horizon cost-to-go function for problem P3.1. Then, in the problem setting for P3.1, we restrict ourselves to the singleton vector of LDFs $\boldsymbol{\phi}^\star = (\phi^{1\star}, \ldots, \phi^{N\star})$, where $\phi^{l\star} : \mathcal{X}^l \times [0,1] \mapsto \{1, \ldots, D_l\}$ is such that

$$\phi^{l\star}(\,\cdot\,; p) \equiv \phi_p^{l\star}(\,\cdot\,).$$

In other words, for each l, $l = 1, \ldots, N$,

$$u_k^l = \phi^{l\star}(X_l^k; p_{k-1}), \quad \forall\, k.$$

We denote the optimization problem with this restriction by P3.1′. We can solve P3.1′ in a manner parallel to the way we solved P3.1, i.e., by first solving the corresponding finite-horizon problem and then extending this solution to the infinite-horizon case. The Bellman equation for the infinite-horizon cost-to-go function $J'(p)$ for problem P3.1′ satisfies

$$J'(p) = \min\left\{L_1 p, L_0(1-p), c + W_{J'}(\boldsymbol{q}_{\boldsymbol{\phi}_p^\star}; p)\right\}.$$

By Theorem 3.3, it follows that $J(p) = J'(p)$, $\forall p \in [0,1]$, which implies the optimality of the stationary vector of LDFs $\boldsymbol{\phi}^\star$ for problem P3.1.

3.6 Numerical Results

For all of the examples presented in this section we assume that the local decisions are binary. For these examples, it is convenient to write the LDFs in terms of the log-likelihood ratio. In this section, the function $L(\cdot)$ represents the log-likelihood ratio of the observations. We consider three cases in increasing order of complexity.

Case A. *Single Sensor*

Here the LDF is characterized by a single threshold λ. Hence, for each $G \in S$, W_G is a function of only λ and p. Let X denote the generic random variable in the set of i.i.d. observations that the system receives. Then

$$W_G(\lambda, p) = \sum_{d=1}^{2} G\left(\frac{g(d,\lambda,p)}{f(d,\lambda,p)}\right)\, f(d,\lambda,p),$$

where

$$\begin{aligned} g(d,\lambda,p) &= p\,[P_0(L(X)>\lambda)]^{d-1}\,[P_0(L(X)\le\lambda)]^{2-d}, \\ f(d,\lambda,p) &= g(d,\lambda,p)+(1-p)\,[P_1(L(X)>\lambda)]^{d-1}\,[P_1(L(X)\le\lambda)]^{2-d}. \end{aligned}$$

An optimal threshold (as a function of p) is obtained by minimizing $W_G(\lambda,p)$ over $\lambda\in\mathbb{R}$. It is easy to see that

$$\lim_{\lambda\to\infty} W_G(\lambda,p)=\lim_{\lambda\to-\infty} W_G(\lambda,p)=G(p).$$

Also, by the concavity of G, for fixed $p\in[0,1]$

$$W_G(\lambda,p)\le G(p),\ \forall\lambda.$$

In addition it is easy to show that, for fixed p, $W_G(\lambda,p)$ has bounded left- and right-hand derivatives for every $\lambda\in\mathbb{R}$. This means that the minimizing threshold can be found to within a desired accuracy by a systematic search procedure [43].

Example 3.1 The observations that the sensor receives are i.i.d. Gaussian random variables with mean 0 and variance v under H_0 and mean 1 and variance v under H_1. In this case $L(X)$ is $N(-1/2v,1/v)$ under H_0 and $N(1/2v,1/v)$ under H_1. An optimal stationary LDF threshold $\lambda^\star(p)$ and the infinite-horizon cost-to-go function are obtained by successive approximation. As indicated earlier, we start the iteration with $\eta(p)$ and repeatedly apply the transformation $\mathcal{T}$, and stop at iteration n if $\mathcal{T}^n\eta$ is sufficiently close to $\mathcal{T}^{n+1}\eta$.

Numerical experimentation suggests that $W_G(\lambda,p)$ is unimodal in λ, for all $G\in S$. We have hence used a *golden section* search procedure [43] to obtain an optimal threshold at each stage of the successive approximation. Representative results are shown in Figure 3.5. A hundred iterations were run, and the norm difference between the 99th and 100th iterates was less than 10^{-4}. The figure indicates the values of the optimal fusion center thresholds a and b. The optimal local decision threshold as a function of p is also plotted.

It is interesting to observe that $\lambda^\star(p)$ is a discontinuous function in both cases (the spikes around the points of discontinuity and at the end points are attributed to quantization and finite precision). This might be surprising at first, but such behavior is commonly observed in control systems where "bang-bang" control is optimal. For example, if we

consider $f(u,x) = -ux$, and we wish to minimize f over $u \in [-1,1]$ for each fixed x, then the minimizing u as a function of x is $\mathrm{sgn}(x)$.

Case B *Two Identical Sensors*

Here, in addition to Assumption 3.2, we assume that the observations received by the two sensors are identically distributed conditioned on each hypothesis. The vector of LDFs is characterized by two thresholds λ_1 and λ_2, with λ_i being the threshold at sensor S_i. Hence W_G is a function of λ_1, λ_2 and p, and is given by

$$W_G(\lambda_1,\lambda_2,p) = \sum_{d_1=1}^{2}\sum_{d_2=1}^{2} G\left(\frac{g(d_1,d_2,\lambda_1,\lambda_2,p)}{f(d_1,d_2,\lambda_1,\lambda_2,p)}\right)\ f(d_1,d_2,\lambda_1,\lambda_2,p)$$

where

$$\begin{aligned} g(d_1,d_2,\lambda_1,\lambda_2,p) &= \prod_{l=1}^{2} p\ [P_0(L(X)>\lambda_l)]^{d_l-1}\ [P_0(L(X)\le\lambda_l)]^{2-d_l}\,, \\ f(d_1,d_2,\lambda_1,\lambda_2,p) &= g(d_1,d_2,\lambda_1,\lambda_2,p) \\ &\quad + \prod_{l=1}^{2}(1-p)\ [P_1(L(X)>\lambda_l)]^{d_l-1}\ [P_1(L(X)\le\lambda_l)]^{2-d_l}\,. \end{aligned}$$

Optimal thresholds (as functions of p) are obtained by minimizing $W_G(\lambda_1,\lambda_2,p)$ over $(\lambda_1,\lambda_2) \in I\!R^2$.

Example 3.2 The observations received by the system are i.i.d. Gaussian random variables with mean 0 and variance v under H_0 and mean 1 and variance v under H_1. In this case $L(X)$ is $N(-1/2v, 1/v)$ under H_0 and $N(1/2v, 1/v)$ under H_1.

Here also, numerical experimentation suggests that for each $G \in S$, $W_G(\lambda_1,\lambda_2,p)$ is unimodal on the set $\{(\lambda_1,\lambda_2) : (\lambda_1,\lambda_2) \in I\!R^2\}$. The unimodality would imply that the search for optimal thresholds can be restricted to the set $\{(\lambda_1,\lambda_2) : \lambda_1 = \lambda_2\}$. This is confirmed in the optimization results (see Figure 3.6) where the optimal thresholds $\lambda_1^\star(p)$ and $\lambda_2^\star(p)$ are seen to be identical functions of p. A hundred iterations were used to obtain these results, and the norm difference between the 99th and 100th iterates was less than 10^{-4}. We note that the same results are obtained if we set $\lambda_1 = \lambda_2 = \lambda$ and optimize W_G over the single threshold λ.

Case C. *Two Nonidentical Sensors*

This case is similar to Case B above except that functions f and g are given by

$$\begin{aligned} g(d_1, d_2, \lambda_1, \lambda_2, p) &= \prod_{l=1}^{2} p\ [P_0(L(X_l) > \lambda_l)]^{d_l - 1}\ [P_0(L(X_l) \leq \lambda_l)]^{2-d_l}\,, \\ f(d_1, d_2, \lambda_1, \lambda_2, p) &= g(d_1, d_2, \lambda_1, \lambda_2, p) \\ &\quad + \prod_{l=1}^{2} (1-p)\ [P_1(L(X_l) > \lambda_l)]^{d_l - 1}\ [P_1(L(X_l) \leq \lambda_l)]^{2-d_l}\,, \end{aligned}$$

where X_l denotes the generic random variable in the i.i.d. sequence of observations received by sensor S_l.

Example 3.3 The observations received by sensor S_1 are i.i.d. Gaussian random variables with mean 0 and variance v under H_0 and mean 1/2 and variance v under H_1. The observations received by sensor S_2 are i.i.d. Gaussian random variables with mean 0 and variance v under H_0 and mean 1 and variance v under H_1. In this case $L(X_1)$ is $N(-1/8v, 1/4v)$ under H_0 and $N(1/8v, 1/4v)$ under H_1, and $L(X_2)$ is $N(-1/2v, 1/v)$ under H_0 and $N(1/2v, 1/v)$ under H_1.

Here again, numerical experimentation suggests that for each $G \in S$, $W_G(\lambda_1, \lambda_2, p)$ is unimodal on the set $\{(\lambda_1, \lambda_2) : (\lambda_1, \lambda_2) \in I\!R^2\}$. Optimal thresholds at each iteration were hence found by a two-dimensional golden section search procedure. Representative results are shown in Figure 3.7. A hundred iterations were run, and, as before, the norm difference between the 99th and 100th iterates was less than 10^{-4}.

3.7 Summary

As we demonstrated in the preceding sections, the information pattern that we assumed for our analysis (Case E of Section 3.1) gave rise to a very tractable problem. Our main results are:

(i) At each stage k, it is optimal for each sensor to quantize its current observation using a likelihood ratio test whose thresholds are determined by the past decision information I_{k-1}.

(ii) The optimal thresholds at the sensors at stage k depend on the I_{k-1} only through the one-dimensional sufficient statistic p_{k-1}. Furthermore, the sufficient statistic can be updated using a simple recursion.

(iii) An optimal policy for the fusion center is a sequential test based on p_k, with fixed boundaries (a and b) in the infinite-horizon case. Also for the infinite-horizon problem, a stationary set of decision functions is optimal at the sensors. That is, the optimal MLRT for each sensor is a time-invariant function of the current observation and the sufficient statistic of the past decision information. This reduces the complexity of the design considerably.

3.8 Proof of Lemma 3.1

The assertion is true for $k = T$ since $J_T^T(p)$ is the minimum of two affine functions of p. Now suppose $J_{m+1}^T(p)$ is concave in p, $p \in [0, 1]$. This is possible if, and only if, there exists a collection of affine functions $\{\lambda_z p + \mu_z : z \in Z\}$, for some index set Z, such that [44]

$$J_{m+1}^T(p) = \inf_{z \in Z} \{\lambda_z p + \mu_z\}.$$

Then,

$$\begin{aligned} A_m^T(p) &= \inf_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}} \sum_{\boldsymbol{d} \in \mathcal{D}} \inf_{z \in Z} \left\{ \lambda_z g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p) + \mu_z f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p) \right\} \\ &= \inf_{\boldsymbol{q}_{\boldsymbol{\phi}} \in \boldsymbol{Q}} \inf_{z \in Z} \sum_{\boldsymbol{d} \in \mathcal{D}} \left\{ \lambda_z g(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p) + \mu_z f(\boldsymbol{d}; \boldsymbol{q}_{\boldsymbol{\phi}}; p) \right\}. \end{aligned}$$

Hence, $A_m^T(p)$ is concave in p, because each term in the above infimum is affine in p. This further implies that $J_m^T(p)$ is concave in p, which completes the proof.

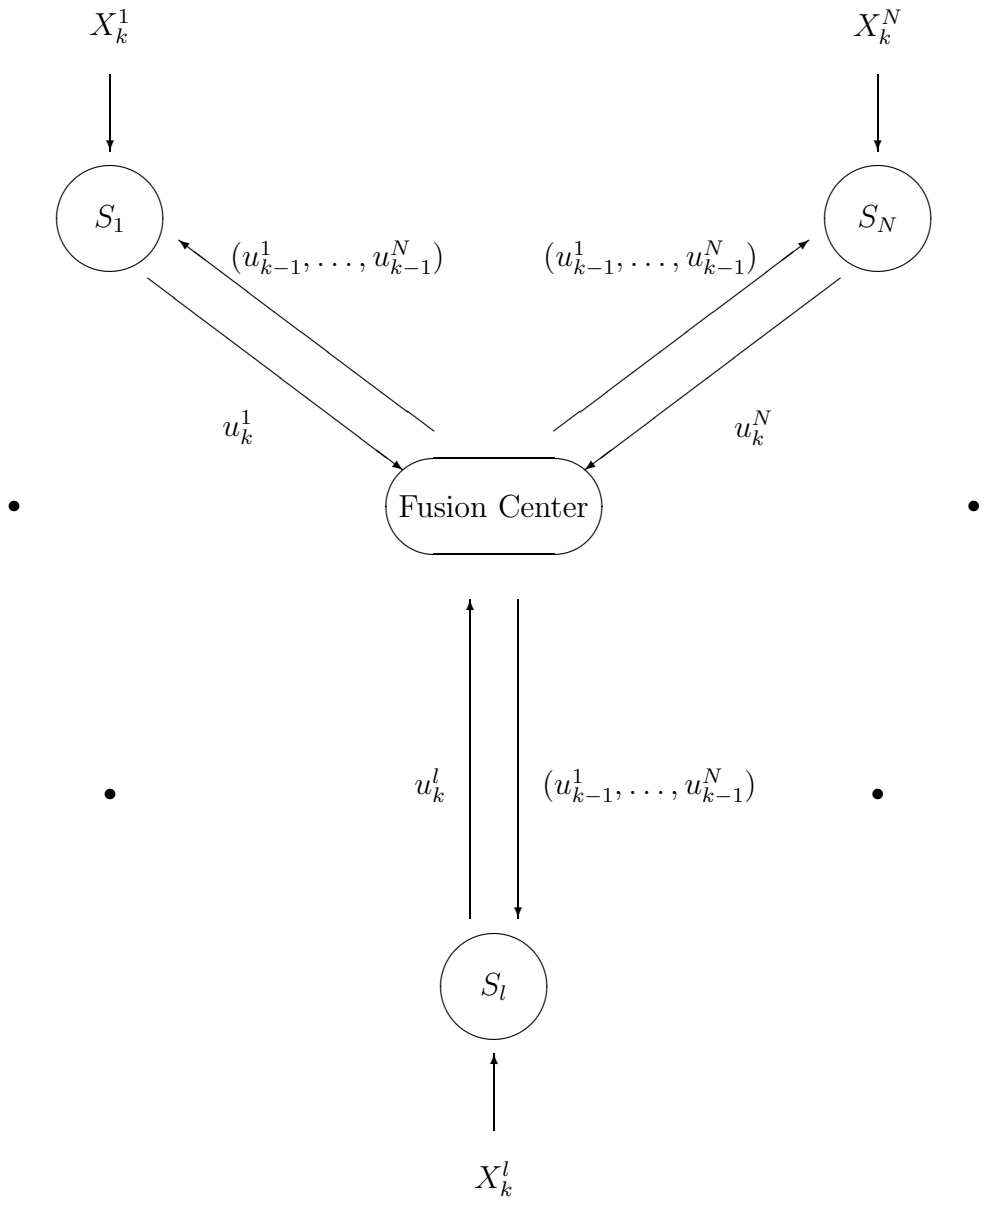

Figure 3.1: General setting for decentralized sequential detection with a fusion center performing the sequential test.

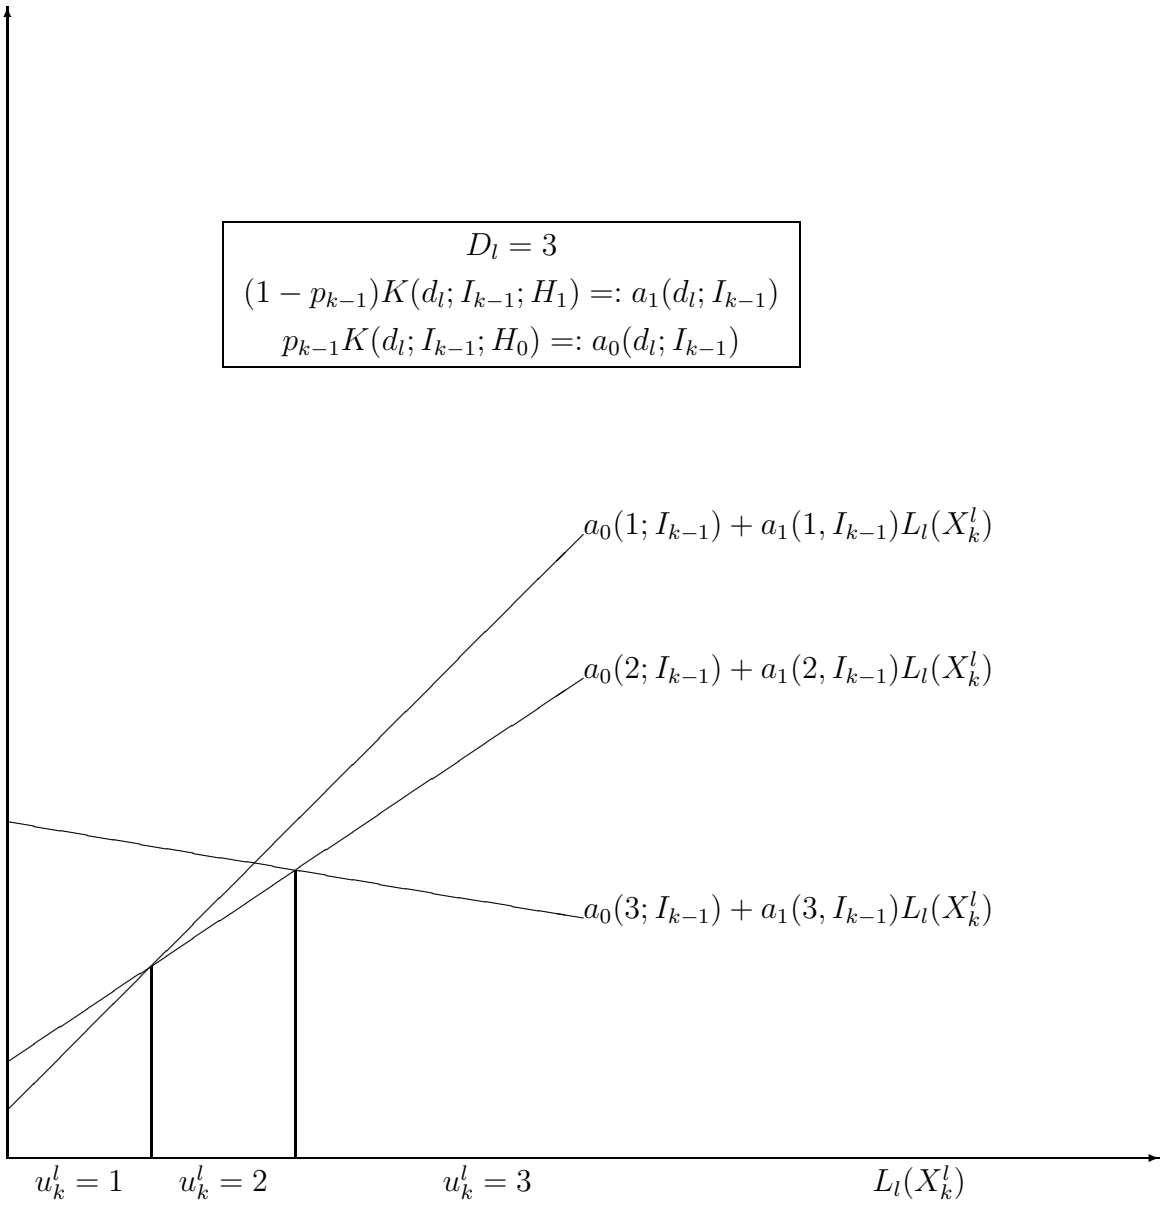

Figure 3.2: Illustration for Proposition 3.2.

Figure 3.3: Illustration for Theorem 3.1.

Figure 3.4: Illustration for Theorem 3.2.

(a)

(b)

Figure 3.5: Results for the single sensor case with $c = 0.01$, $v = 1.0$, and $L_0 = L_1 = 1.0$: (a) Infinite-horizon cost-to-go function (b) Optimal stationary LDF threshold.

(a)

(b)

Figure 3.6: Results for the case of two identical sensors with $c = 0.01$, $v = 1.0$, and $L_0 = L_1 = 1.0$: (a) Infinite-horizon cost-to-go function (b) Optimal stationary LDF thresholds.

(a)

(b)

Figure 3.7: Results for the case of two nonidentical sensors with $c = 0.01$, $v = 1.0$, and $L_0 = L_1 = 1.0$: (a) Infinite-horizon cost-to-go function (b) Optimal stationary LDF thresholds.

CHAPTER 4

MINIMAX ROBUST DECENTRALIZED DETECTION

4.1 Introduction

The design of optimal decision rules in detection problems requires the knowledge of the conditional probability distributions of the observations, given each hypothesis. In many applications, however, the probability distributions are not specified completely. In these cases, the probability distributions are usually specified to belong to classes (sets) of distributions, often termed as *uncertainty classes*. One way to design decision rules when the probability distributions are given to belong to uncertainty classes is the *minimax* approach, where the goal is to minimize the worst-case performance over the uncertainty classes. The decision rules thus obtained are said to be robust to the uncertainties in the probability distributions.

Minimax robust detection problems with two hypotheses[1] and with centralized information have been the subject of numerous papers (for an excellent survey of results in this area, see [46]). The solutions to these problems invariably involve identifying a pair of *least favorable distributions* (LFDs), and subsequently designing a simple hypothesis test between the LFDs.

An extension of the minimax robust detection problem to a decentralized setting with two sensors and without a fusion center was considered by Geraniotis [47]. The problem was formulated in a Bayesian framework with the observations at each of the sensors belonging to uncertainty classes generated by alternating capacities of order two. The binary sensor decisions about the hypothesis were assumed to be coupled through a common cost function. For a specific choice of cost structure, it was shown in [47] that the task of finding LFDs at the sensors can be decoupled into two independent tasks, one at each of the sensors. This implies that the LFDs for the decentralized problem are the same as those for two independent centralized detection problems at the sensors.

[1]Minimax robust detection problems with more than two hypotheses are known to be difficult and do not admit closed-form solutions [45].

Minimax robust decentralized detection with a fusion center has also been studied. In the only existing analyses of this problem ([48] and [49]), the authors restricted their study to a Bayesian formulation and to binary sensor decisions. They further limited the scope of their study by only considering the following special cases: (i) the case of identical sensors using identical decision rules, (ii) the asymptotic case of a large number of sensors and (iii) the asymptotic case of large observation block lengths.

In this chapter, we attempt to find a more comprehensive solution to robust decentralized detection problems. We study both the cases with and without a fusion center. For the case when a fusion center is present, we give a solution to the minimax robust detection problem for the general case of finite number of sensors, finite observation block length, and non-binary sensor decisions. This solution covers all of the block detection cases considered in [48] and [49]. Furthermore, our analysis is not restricted to Bayesian detection. For the case when no fusion center is present, we extend the work in [47] to more than two sensors and more general cost functions. We also give sufficient conditions for the decoupling of the minimax robust detection problem.

The remainder of this chapter is organized as follows: In Section 4.2, we give a detailed introduction to robust centralized detection. The purpose of this introduction is twofold: first, we believe that we have provided a framework whereby most of the previous results in robust centralized detection are unified; second, the results here are used explicitly in the solution to the decentralized problems in the subsequent sections. In Section 4.3, we consider decentralized detection problems where a fusion center is present, and in Section 4.4 we consider the case where the fusion center is absent. Finally, in Section 4.5, we summarize the main points.

4.2 Robust Centralized Detection

We begin with a description of a minimax robust detection problem which was first introduced by Huber [50]. The basic setup is as follows: Let $(\mathcal{X}, \mathcal{F})$ be a measurable space,

and let P_0 and P_1 be distinct probability measures on it. Let X be an observation taking values in $\mathcal{X}$, and let the distribution of X be P_0 (respectively, P_1) under H_0 (respectively, H_1). A decision δ about the true hypothesis is to be made based on X, i.e., $\delta = \phi(X)$. The objective here is to construct a hypothesis test between H_0 and H_1, when P_0 and P_1 are not specified completely. The approach taken by Huber was to first define classes of allowable distributions (or uncertainty classes) under H_0 and H_1, and then solve a minimax test between this pair of classes. If we denote the uncertainty class under H_j by $\mathcal{P}_j$, then the minimax robust versions of Bayesian, minimax and Neyman-Pearson formulations of the hypothesis test between H_0 and H_1 are given respectively by

$$\text{(a)} \qquad \min_{\phi} \left[\nu \sup_{P_0 \in \mathcal{P}_0} P_F(\phi, P_0) + (1-\nu) \sup_{P_1 \in \mathcal{P}_1} P_M(\phi, P_1) \right],$$

$$\text{(b)} \qquad \min_{\phi} \max \left\{ \sup_{P_0 \in \mathcal{P}_0} P_F(\phi, P_0), \sup_{P_1 \in \mathcal{P}_1} P_M(\phi, P_1) \right\},$$

$$\text{(c)} \qquad \min_{\phi} \sup_{P_1 \in \mathcal{P}_1} P_M(\phi, P_1) \text{ subject to } \sup_{P_0 \in \mathcal{P}_0} P_F(\phi, P_0) \leq \alpha,$$

where $P_M(\phi, \mathcal{P}_1) = P_1(\delta = 0)$ and $P_F(\phi, \mathcal{P}_0) = P_0(\delta = 1)$.

The classes considered in [50] are neighborhood classes containing, under each hypothesis, a nominal distribution and distributions in its vicinity. The two types of neighborhood classes studied in [50] are the ϵ-contamination and the total variation. For each case, Huber showed [50] that a pair of LFDs can be found for the minimax robust detection problems described above. He also gave a characterization of a least favorable pair in terms of the parameters of the uncertainty neighborhoods, and showed that the corresponding minimax robust tests are "censored" versions of the nominal likelihood ratio tests.

Huber and Strassen [51] have shown in a later paper that pairs of LFDs can be found for the cases when the neighborhood classes can be described in terms of alternating capacities of order 2. When the observation set is *compact*, several uncertainty models such as ϵ-contaminated neighborhoods, total variation neighborhoods, band-classes and p-point classes are special cases of this model with different choices of capacity.

The proofs of existence of LFDs in [50] and [51] rely on the following property possessed by all of the pairs of uncertainty classes considered in [50] and [51]:

Definition 4.1 (Joint Stochastic Boundedness)*: A pair* $(\mathcal{P}_0, \mathcal{P}_1)$ *of classes of distributions defined on a measurable space* $(\mathcal{X}, \mathcal{F})$ *is said to be jointly stochastically bounded by* (Q_0, Q_1)*, if there exist distributions* $Q_0 \in \mathcal{P}_0$ *and* $Q_1 \in \mathcal{P}_1$ *such that for any* $(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1$ *and all* $t \geq 0$*,*

$$P_0(l_q(X) > t) \leq Q_0(l_q(X) > t), \qquad P_0(l_q(X) < t) \geq Q_0(l_q(X) < t)$$

and

$$P_1(l_q(X) < t) \leq Q_1(l_q(X) < t), \qquad P_1(l_q(X) > t) \geq Q_1(l_q(X) > t)$$

where l_q *is the likelihood ratio between* Q_1 *and* Q_0*.* □

Example 4.1 ϵ-contamination uncertainty classes:

$$\mathcal{P}_i := \left\{ Q \in \mathcal{H} \,|\, Q = (1-\epsilon_i)\tilde{P}_i + \epsilon_i M_i, \quad M_i \in \mathcal{H} \right\}$$

where $\mathcal{H}$ denotes the set of all probability measures on $\mathcal{X}$, and $\tilde{P}_0$ and $\tilde{P}_1$ are distinct probability measures with densities $\tilde{p}_0$ and $\tilde{p}_1$ with respect to some measure μ. It can be shown [50] that this pair $(\mathcal{P}_0, \mathcal{P}_1)$ is jointly stochastically dominated by the pair of distributions Q_0 and Q_1 which have densities with respect to μ given by

$$q_0(x) = \begin{cases} (1-\epsilon_0)\tilde{p}_0(x) & \text{if } l_p(x) < c_0 \\ (1-\epsilon_0)\tilde{p}_1(x)/c_0 & \text{if } l_p(x) \geq c_0 \end{cases}$$

and

$$q_1(x) = \begin{cases} (1-\epsilon_1)\tilde{p}_1(x) & \text{if } l_p(x) > c_1 \\ c_1(1-\epsilon_1)\tilde{p}_0(x) & \text{if } l_p(x) \leq c_1 \end{cases},$$

where c_0 and c_1 are such that the densities integrate to 1, and $l_p(x) = \tilde{p}_1(x)/\tilde{p}_0(x)$. □

It can be shown [50, 51] that the distributions Q_0 and Q_1 in Definition 4.1 are LFDs for minimax robust hypothesis testing between $\mathcal{P}_0$ and $\mathcal{P}_1$. That is, the solutions to (a), (b) and (c) are obtained as solutions to the following simple hypothesis testing problems:

(a′) $\min_\phi \ [\nu\, P_F(\phi, Q_0) + (1-\nu)\, P_M(\phi, Q_1)]$,

(b′) $\min_\phi \ \max\{P_F(\phi, Q_0), P_M(\phi, Q_1)\}$,

(c′) $\min_\phi P_M(\phi, Q_1)$ subject to $P_F(\phi, Q_0) \leq \alpha$.

In many applications, the observation X is a vector (block) of independent observations $(X_1, \ldots, X_n)$, where the observation X_i takes values in a measurable space $(\mathcal{X}_i, \mathcal{F}_i)$ and has a distribution which belongs to the class $\mathcal{P}_j^i$ when the hypothesis is H_j. In this chapter, the set $\mathcal{P}_j := \mathcal{P}_j^1 \times \cdots \times \mathcal{P}_j^n$ represents a class of distributions on $(\mathcal{X}, \mathcal{F})$ which are products of distributions in $\mathcal{P}_j^i, i = 1, \ldots, n$. To further clarify this point, let P_j^i denote a typical element of $\mathcal{P}_j^i$. Then $P_j := (P_j^1, \ldots, P_j^n) \in \mathcal{P}_j$ represents the product distribution $P_j^1 \times \cdots \times P_j^n$.

In the above context, we have the following result[2]:

Lemma 4.1 *For each i, $i = 1, \ldots, n$, let the pair $(\mathcal{P}_0^i, \mathcal{P}_1^i)$ be jointly stochastically bounded by (Q_0^i, Q_1^i). Then the pair $(\mathcal{P}_0, \mathcal{P}_1)$ is jointly stochastically bounded by (Q_0, Q_1).*

Proof: Let l_q denote the likelihood ratio between Q_1 and Q_0, and let l_q^i denote the likelihood ratio between Q_1^i and Q_0^i. Then for any $t \geq 0$, we have

$$\begin{aligned}\{l_q(X) > t\} &= \left\{ \prod_{i=1}^n l_q^i(X_i) \ > \ t \right\} \\ &= \bigcup_{t_1, \ldots, t_n \geq 0, \text{ rational} : \prod_{i=1}^n t_i \geq t} \left\{ \bigcap_{i=1}^n \ \left\{ l_q^i(X_i) > t_i \right\} \right\}.\end{aligned}$$

By the joint stochastic boundedness property of $(\mathcal{P}_0^i, \mathcal{P}_1^i)$, each event in the countable union above has a larger probability under P_0 than under Q_0. This proves the first condition required for the joint stochastic boundedness of $(\mathcal{P}_0, \mathcal{P}_1)$ by (Q_0, Q_1). The other conditions are proved similarly. □

Lemma 1 implies the following: Suppose we are given a minimax robust detection problem with a block of independent observations, and with uncertainty classes satisfying the joint stochastic boundedness property. Then, this problem can be reduced to a single observation problem without sacrificing the joint stochastic boundedness property.

[2]The reader may be familiar with other forms of Lemma 4.1. We believe that the one given here is the most general form.

The above description of centralized robust detection problems, we believe, unifies most of the results in this area. Hence the material of Section 4.2 should be of independent interest to the reader.

4.3 Robust Decentralized Detection with a Fusion Center

A description of the decentralized detection system considered in this section is given in Fig. 1.1. The hypothesis H takes values H_0 and H_1. There are N sensors and one fusion center. The sensor S_i receives an observation X_i which is assumed to take values on a measurable space $(\mathcal{X}_i, \mathcal{F}_i)$. By virtue of Lemma 4.1, X_i could represent a block of independent observations. The observations at the sensors are independent, and X_i has a distribution which belongs to the class $\mathcal{P}_j^i$ when the hypothesis is H_j. For each i, the pair of uncertainty classes $(\mathcal{P}_0^i, \mathcal{P}_1^i)$ is jointly stochastically bounded by (Q_0^i, Q_1^i) (see Definition 4.1). Let $\mathcal{P}_j := \mathcal{P}_j^1 \times \cdots \times \mathcal{P}_j^N$. Then $P_j = (P_j^1, \ldots, P_j^N) \in \mathcal{P}_j$ represents the product distribution $P_j^1 \times \cdots \times P_j^N$.

By Lemma 4.1, the pair $(\mathcal{P}_0, \mathcal{P}_1)$ is jointly stochastically dominated by (Q_0, Q_1). Hence, if all of the information received by the sensors is made available to the fusion center, then (Q_0, Q_1) are LFDs for robust hypothesis testing between H_0 and H_1. But in the decentralized setting, only a summary of the sensor observations is available at the fusion center. At sensor S_i, there is a decision function ϕ_i which maps the observation vector X_i to a local decision $u_i \in \{1, \ldots, D_i\}$, and the fusion center makes a final binary-valued decision δ based on the information it receives from the sensors, i.e., $\delta = \gamma(u_1, \ldots, u_N)$. For compactness of notation, we represent the set of local decision functions $(\phi_1, \ldots, \phi_N)$ by $\boldsymbol{\phi}$.

We consider here, in detail, a Bayesian formulation of the robust decentralized detection problem where the objective is to minimize the worst-case error probability at the fusion center. The minimax and Neyman-Pearson formulations are discussed briefly at the end of this section. The hypothesis is assumed to take on values H_0 and H_1, with prior probabilities ν and $1 - \nu$, respectively. Let $P_F(P_0, \boldsymbol{\phi}, \gamma) := P_0(\delta = 1)$ and $P_M(P_1, \boldsymbol{\phi}, \gamma) := P_1(\delta = 0)$.

Then the problem we wish to solve is the following:

Problem P4.1:

$$\inf_{\boldsymbol{\phi},\gamma} \quad \left[\nu \sup_{P_0 \in \mathcal{P}_0} P_F(P_0, \boldsymbol{\phi}, \gamma) + (1-\nu) \sup_{P_1 \in \mathcal{P}_1} P_M(P_1, \boldsymbol{\phi}, \gamma) \right]$$

In the following, we will establish that (Q_0, Q_1), which were LFDs for the centralized problem, are LFDs for Problem P4.1 as well. That is, the solution to P4.1 is obtained as the solution to the simple decentralized detection problem P4.1′ given below.

Problem P4.1′:

$$\inf_{\boldsymbol{\phi},\gamma} \quad [\, \nu\, P_F(Q_0, \boldsymbol{\phi}, \gamma) + (1-\nu)\, P_M(Q_1, \boldsymbol{\phi}, \gamma)\,]$$

Optimal decision rules for Problem P4.1′ are monotone likelihood ratio tests (MLRTs) [11] of the form:

$$\phi_i^R(X_i) = \begin{cases} 1 & \text{if } l_q^i(X_i) < \lambda_1^i \\ d & \text{if } \lambda_{d-1}^i \le l_q^i(X_i) < \lambda_d^i, d = 2, \ldots, D_i - 1 \\ D_i & \text{if } l_q^i(X_i) \ge \lambda_{D_i}^i \end{cases}$$

and

$$\gamma^R(u_1, \ldots, u_N) = \begin{cases} 1 & \text{if } l_q(u_1, \ldots, u_N) > t \\ 0 & \text{otherwise} \end{cases},$$

where

$$l_q(\nu_1, \ldots, \nu_N) = \prod_{i=1}^{N} \frac{Q_{1,u}^i(\nu_i)}{Q_{0,u}^i(\nu_i)}.$$

In the above expressions, l_q^i denotes the generalized likelihood ratio between Q_1^i and Q_0^i. Also, $Q_{j,u}^i$ denotes the distribution induced on u_i by Q_j^i.

Theorem 4.1 *Let $\boldsymbol{\phi}^R$, γ^R be any set of MLRTs based on Q_0 and Q_1. Then for all $(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1$, we have*

$$P_F(Q_0, \boldsymbol{\phi}^R, \gamma^R) \ge P_F(P_0, \boldsymbol{\phi}^R, \gamma^R),$$

$$P_M(Q_1, \boldsymbol{\phi}^R, \gamma^R) \ge P_M(P_1, , \boldsymbol{\phi}^R, \gamma^R).$$

Before we address the proof of this result, we consider its implications. Suppose $\boldsymbol{\phi}^\star$, $\gamma^\star$ constitute a solution to P4.1′. Then

$$\begin{aligned}
\nu \sup_{P_0 \in \mathcal{P}_0} P_F(P_0, \boldsymbol{\phi}^\star, \gamma^\star) \;+\; & (1-\nu) \sup_{P_1 \in \mathcal{P}_1} P_M(P_1, \boldsymbol{\phi}^\star, \gamma^\star) \\
&= \nu P_F(Q_0, \boldsymbol{\phi}^\star, \gamma^\star) + (1-\nu) P_M(Q_1, \boldsymbol{\phi}^\star, \gamma^\star) \\
&\leq \nu P_F(Q_0, \boldsymbol{\phi}, \gamma) + (1-\nu) P_M(Q_1, \boldsymbol{\phi}, \gamma) \\
&\leq \nu \sup_{P_0 \in \mathcal{P}_0} P_F(P_0, \boldsymbol{\phi}, \gamma) + (1-\nu) \sup_{P_1 \in \mathcal{P}_1} P_M(P_1, \boldsymbol{\phi}, \gamma), \qquad (4.1)
\end{aligned}$$

for any $\boldsymbol{\phi}$, γ. This means that (Q_0, Q_1) are LFDs for P4.1, and that $\boldsymbol{\phi}^\star$, $\gamma^\star$ solve P4.1.

We now state and prove the following lemma which will be used in the proof of Theorem 4.1.

Lemma 4.2 *For each i and for any $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ satisfying $0 \leq \lambda_1 \leq \lambda_2 \leq \lambda_3 \leq \lambda_4 < \infty$,*

$$\frac{Q_1^i(\lambda_1 \leq l_q^i(X_i) \leq \lambda_2)}{Q_0^i(\lambda_1 \leq l_q^i(X_i) \leq \lambda_2)} \leq \frac{Q_1^i(\lambda_3 \leq l_q^i(X_i) \leq \lambda_4)}{Q_0^i(\lambda_3 \leq l_q^i(X_i) \leq \lambda_4)}.$$

Proof: The LHS above can be written as

$$\frac{1}{Q_0^i(\lambda_1 \leq l_q^i(X_i) \leq \lambda_2)} \int_{\{\lambda_1 \leq l_q^i(x) \leq \lambda_2\}} l_q^i(x) dQ_0^i(x).$$

This means that the LHS is the average of $l_q^i(x)$ with respect to the distribution Q_0^i on the set $\{\lambda_1 \leq l_q^i(x) \leq \lambda_2\}$ and hence is in between λ_1 and λ_2. Similarly the RHS is between λ_3 and λ_4. The lemma follows. □

Proof of Theorem 4.1 For any $(P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1$,

$$P_F(P_0, \boldsymbol{\phi}^R, \gamma^R) = P_0\left(l_q(u_1, \ldots, u_N) > t\right).$$

First, we note that

$$\begin{aligned}
\{l_q(u_1, \ldots, u_N) > t\} &= \left\{ \prod_{i=1}^N \frac{Q_{1,u}^i(u_i)}{Q_{0,u}^i(u_i)} > t \right\} \\
&= \bigcup_{t_1, \ldots, t_N \geq 0, \text{ rational} : \prod_{i=1}^N t_i \geq t} \left\{ \bigcap_{i=1}^N \left\{ \frac{Q_{1,u}^i(u_i)}{Q_{0,u}^i(u_i)} > t_i \right\} \right\}. \qquad (4.2)
\end{aligned}$$

By Lemma 4.2, $Q^i_{1,u}(u_i)/Q^i_{0,u}(u_i)$ is a nondecreasing function of u_i and hence

$$\left\{\frac{Q^i_{1,u}(u_i)}{Q^i_{0,u}(u_i)} > t_i\right\} = \{l^i_q(X_i) > \lambda_{t_i}\},$$

for an appropriately chosen λ_{t_i}. Now, an application of the joint stochastic boundedness property of the pairs $(\mathcal{P}^i_0, \mathcal{P}^i_1)$ gives us that

$$\begin{aligned} Q_0\left(\bigcap_{i=1}^{N} \{l^i_q(X_i) > \lambda_{t_i}\}\right) &= \prod_{i=1}^{N} Q^i_0(l^i_q(X_i) > \lambda_{t_i}) \\ &\geq \prod_{i=1}^{N} P^i_0(l^i_q(X_i) > \lambda_{t_i}) \\ &= P_0\left(\bigcap_{i=1}^{N} \{l^i_q(X_i) > \lambda_{t_i}\}\right), \end{aligned}$$

i.e., each event in the countable union of (4.2) has a larger probability under Q_0 than under P_0. The first part of the theorem follows. An analogous argument can be used to establish the second part. □

The minimax and Neyman-Pearson versions of P4.1 can be stated as follows:

Problem P4.2:

$$\inf_{\phi_1,\ldots\phi_N,\gamma} \max\left[\sup_{P_0\in\mathcal{P}_0} P_F(P_0, \boldsymbol{\phi}, \gamma), \sup_{P_1\in\mathcal{P}_1} P_M(P_1, \boldsymbol{\phi}, \gamma)\right]$$

Problem P4.3:

$$\inf_{\phi_1,\ldots\phi_N,\gamma} \sup_{P_1\in\mathcal{P}_1} P_M(P_1, \boldsymbol{\phi}, \gamma) \text{ subject to } \sup_{P_0\in\mathcal{P}_0} P_F(P_0, \boldsymbol{\phi}, \gamma) \leq \alpha.$$

Just as in the simple versions of P4.2 and P4.3, we extend the class of allowable decision functions to include jointly randomized decision rules [11]. In the class of jointly randomized strategies, the simple versions of P4.2 and P4.3 have solutions that are randomized tests obtained by joint randomization between two deterministic MLRT strategies. Hence, an argument similar to the one given in (4.1) can be used to show that (Q_0, Q_1) are LFDs for P4.2 and P4.3. The corresponding robust tests are then obtained as solutions to the simple decentralized detection problems in which the sets $\mathcal{P}_0$ and $\mathcal{P}_1$ are replaced by the singletons Q_0 and Q_1, respectively.

4.4 Robust Decentralized Detection without a Fusion Center

The system under consideration here is similar to the one in Section 4.3, with the only differences being that the local decisions u_i are binary and that there is no fusion center. This setup is useful only in a Bayesian framework in which we assume that the local decisions are coupled through a common cost function $W(u_1, \ldots, u_N; H)$. The expected cost is a function of the conditional distributions at the sensors and the local decision functions. Thus the expected cost is given by

$$C(P_0, P_1, \boldsymbol{\phi}) = \nu\, E_{P_0}\{W(u_1, \ldots, u_N; H_0)\} \;+\; (1-\nu)\, E_{P_1}\{W(u_1, \ldots, u_N; H_1)\}.$$

The Bayes minimax robust detection problem at hand is then the following:

Problem P4.4:

$$\inf_{\boldsymbol{\phi}} \sup_{(P_0,P_1)\in\mathcal{P}_0\times\mathcal{P}_1} C(P_0, P_1, \boldsymbol{\phi}).$$

If the distributions of the observations are known, i.e., if the uncertainty classes $\mathcal{P}_j^i$ are singletons, then the optimal decision rules for P4.4 are binary likelihood ratio tests (LRTs) [11, 3]. Here we have assumed that for each i, the pair $(\mathcal{P}_0^i, \mathcal{P}_1^i)$ is jointly stochastically bounded (Q_0^i, Q_1^i). We showed in Section 4.3 that (Q_0, Q_1) are LFDs for the Bayesian decentralized detection problem P4.1. Although a similar result cannot be proved in its most generality for problem P4.4, we will show below that there are many interesting cases where (Q_0, Q_1) are indeed LFDs for P4.4.

The distributions (Q_0, Q_1) are LFDs for P4.4, if for any LRTs $\boldsymbol{\phi}^R$ based on Q_1 and Q_0, the following inequality holds:

$$C(Q_0, Q_1, \boldsymbol{\phi}^R) \geq C(P_0, P_1, \boldsymbol{\phi}^R), \quad \forall (P_0, P_1) \in \mathcal{P}_0 \times \mathcal{P}_1. \tag{4.3}$$

Now suppose that $\boldsymbol{\phi}^\star$ constitutes a solution to the following problem:

$$\inf_{\boldsymbol{\phi}} C(Q_0, Q_1, \boldsymbol{\phi}).$$

Then, by an argument similar to one following Theorem 4.1, we can show that $\boldsymbol{\phi}^\star$ solves P4.4.

Of course, any likelihood ratio test ϕ_i^R at sensor S_i based on Q_1^i and Q_0^i has the form

$$\phi_i^R(X_i) = \begin{cases} 1 & \text{if } l_q^i(X_i) \geq t_i \\ 0 & \text{if } l_q^i(X_i) < t_i \end{cases} .$$

We now consider some special cases.

4.4.1 The two-sensor case

Here the expected cost has the form:

$$\begin{aligned} C(P_0, P_1, \phi_1^R, \phi_2^R) &= \nu\, W(0,0;H_0)\left(1 - P_0^1(l_q^1(X_1) \geq t_1)\right)\left(1 - P_0^2(l_q^2(X_2) \geq t_2)\right) \\ &+ \nu\, W(0,1;H_0)\left(1 - P_0^1(l_q^1(X_1) \geq t_1)\right) P_0^2(l_q^2(X_2) \geq t_2) \\ &+ \nu\, W(1,0;H_0) P_0^1(l_q^1(X_1) \geq t_1)\left(1 - P_0^2(l_q^2(X_2) \geq t_2)\right) \\ &+ \nu\, W(1,1;H_0) P_0^1(l_q^1(X_1) \geq t_1) P_0^2(l_q^2(X_2) \geq t_2) \\ &+ \text{ similar terms in } W(i,j;H_1). \end{aligned}$$

If $W(0,0;H_0) = W(1,1;H_1) = 0$, then for (4.3) to hold, it is sufficient that the following conditions hold:

$$W(1,1,H_0) \geq W(0,1;H_0) + W(1,0;H_0)$$

and

$$W(0,0;H_1) \geq W(0,1;H_1) + W(1,0;H_1).$$

A special case of the above conditions is found in [47] where it is assumed that the cost function is of the form:

$$W(u_1,u_2;H) = \begin{cases} 0 & \text{for } u_1 = u_2 = H \\ e & \text{for } u_1 \neq u_2 \\ f > 2e & \text{for } u_1 = u_2 \neq H. \end{cases}$$

4.4.2 The case of fixed symmetric fusion rules

Bayesian decentralized detection problems with binary local decisions and fixed fusion rules such as the "AND" rule and the "OR" rule can be posed in the framework of this section with an appropriately chosen cost function W.

For the "AND" rule, a final decision in favor of H_1 is made whenever all of the local decisions are 1; otherwise, a decision in favor of H_0 is made. If the Bayesian criterion is to minimize the error probability at the fusion center, then the corresponding cost function $W(u_1, \ldots, u_N; H)$ has the form:

$$W(u_1, \ldots, u_N; H_1) = \begin{cases} 0 & \text{for } u_1 = u_2 = \cdots = u_N = 1 \\ 1 & \text{otherwise} \end{cases},$$

and

$$W(u_1, \ldots, u_N; H_0) = \begin{cases} 1 & \text{for } u_1 = u_2 = \cdots = u_N = 1 \\ 0 & \text{otherwise} \end{cases}.$$

In this case the expected cost is

$$C(P_0, P_1, \boldsymbol{\phi}^R) = \nu \prod_{i=1}^{N} P_0^i(l_q^i(X_i) \geq t_i) + (1-\nu)\left(1 - \prod_{i=1}^{N} P_1^i(l_q^i(X_i) \geq t_i)\right).$$

A straightforward application of the joint stochastic boundedness property shows that condition (4.3) holds. That is, (Q_0, Q_1) are LFDs for P4.4 in this case. A similar result holds for the "OR" fusion rule.

4.5 Summary

In this chapter, we studied decentralized detection problems in which the sensor distributions were not specified completely, i.e., the sensor distributions were assumed to belong to known uncertainty classes. We showed, for a broad class of problems, that a set of least favorable distributions exists for minimax robust testing between the hypotheses. These LFDs can be obtained by previously known techniques [46], and the corresponding minimax robust tests are solutions to simple decentralized detection problems for which the sensor distributions are specified to be the LFDs.

CHAPTER 5

CONCLUSIONS

The primary purpose of our work has been to explore problems in decentralized sequential detection and robust decentralized detection. In the first chapter, however, we have first provided a unified summary of existing results in the general area of decentralized detection in order to facilitate an understanding of the field as a whole and to place significance of the results of the subsequent chapters in proper perspective.

In the area of decentralized sequential detection, we first considered the case in which the sensors perform sequential tests (Chapter 2). We formulated a Bayesian problem with the assumptions that decision errors at the sensors are penalized through a common cost function, and that each time step taken by the detectors as a team is assigned a positive cost. We then showed that optimal sensor decision functions can be found in the class of GSPRTs with monotonically convergent thresholds and introduced a technique to obtain optimal GSPRT thresholds. We also compared the performance of the GSPRT policies with that of the best SPRT policies, and we noted that a trade-off must be made between the simplicity of the SPRT policy and the performance gain obtainable with the GSPRT policy.

The analysis contained in Chapter 2 can easily be extended to the general case in which there are N ($N > 2$) sensors, without any conceptual difficulties. Also, the case in which the stopping time penalty is of the form $c_1\tau_1 + c_2\tau_2 + c\max(\tau_1, \tau_2)$ is also easily handled through minor modifications. Here again it can be shown that optimal solutions can be found in the class of GSPRTs with monotonically convergent thresholds.

In Chapter 3, we explored the class of decentralized sequential detection problems where the sequential test is carried out at a fusion center. We introduced a Bayesian framework for this problem and showed that the problem is tractable when the information structure in the system is quasiclassical (Case E of Section 3.1). In particular, we showed that an optimal fusion center policy has a simple SPRT-like structure and that a stationary set of MLRTs is optimal at the sensors.

At this point it is important to note that, in general, an optimal scheme for the information pattern of Case D (of Section 3.1) could outperform an optimal scheme for the information pattern we adopted in Chapter 3. However, as we pointed out, Case D is highly intractable, and even if an optimal scheme were found, it would most likely require the sensors to retain all information from the past. One might wonder if some ad hoc scheme for Case D which uses only a one-dimensional sufficient statistic of the past information would perform better than an optimal scheme for our case. Performance analysis for the ad hoc scheme would still be difficult and we might have to resort to simulations. We have some simulation results for an obvious scheme for Case D with $N = 1$ (for Example 1 in Section 3.6) which show that this scheme performs consistently worse than an optimal scheme for Case E.

We were able to use dynamic programming arguments in our analysis to obtain optimality results because all decision makers in the system have the same information about the past. This is not true for the information patterns of Cases A, B, C and D, as discussed in Section 3.1. An interesting open problem for these cases would be to investigate if stationary LDFs are optimal under Assumption 3.2. Such a result would be useful especially for Cases A and C since we have already established the optimality of likelihood ratio tests. Also, if we do not allow the fusion center to send messages back to the sensors, then Case E reduces to Case C. Hence, any results for the infinite-horizon problem in Case C would tie in very well with the results presented in Chapter 3.

In Chapter 4, we studied decentralized detection problems in which the sensor distributions were not specified completely, i.e., the sensor distributions were assumed to belong to known uncertainty classes. We showed for a broad class of such detection problems that LFDs exist for minimax robust testing between the hypotheses. These LFDs can be obtained by previously known techniques, and the corresponding minimax robust tests are then obtained as solutions to simple decentralized detection problems in which the sensor distributions are specified to be the LFDs.

We note that the analysis presented in Chapter 4 was restricted to static or block detection schemes. Robustification of decentralized sequential detection schemes such as those presented in Chapters 2 and 3 remains an interesting problem for further research.

5.1 Some Open Research Topics

While the work done here has provided significant progress in the area of decentralized detection, there still remain many problems which need to be explored further:

1. Previously, decentralized detection by a tandem of sensors has been considered mainly in the context of attempting to determine necessary conditions for the error probability to go to zero asymptotically [27], [52]. We are unaware of any results on the rates of convergence of the error probabilities.

If we consider the special case in which the sensor observations are i.i.d. $N(-\mu, \sigma^2)$ under H_0 and i.i.d. $N(\mu, \sigma^2)$ under H_1, then an argument given in [27] shows that if we pick

$$u_k = \begin{cases} 1 & \text{if } X_k > \sqrt{2\sigma^2 \log k} \\ 0 & \text{if } X_k < -\sqrt{2\sigma^2 \log k} \\ u_{k-1} & \text{otherwise,} \end{cases}$$

then the asymptotic error probability is indeed zero. Now, if we let p_n denote the error probability for n stages, then we have some preliminary results which show that: (1) for any fixed positive integer m, $p_n(\log n)^m$ goes to zero as $n \to \infty$, and (2) p_n goes to zero at a subexponential rate as $n \to \infty$. So far we have not been able to identify the exact rate at which p_n goes to zero.

2. As we commented in Chapter 1, in the two-sensor case, the tandem configuration is always at least as good as the parallel configuration. This is because for $N = 2$, the detector which makes the decision in the tandem configuration has more information than the fusion center for the parallel configuration. This advantage over the fusion configuration can be maintained for $N > 2$ if in the tandem configuration we allow the decision at the k-th stage to depend on all past decisions rather than on only the most recent decision. Under

Assumption 1.1, it is clear that the asymptotic (minimum) error probability for the modified tandem configuration is zero and that the error probability goes to zero at an exponential rate. The open problem here is to find the optimum error exponent, and to see if this error exponent is achieved by a stationary sensor decision policy.

3. Consider the decentralized detection problem where the distributions of the observations are known but the prior probabilities of the hypotheses are unknown, i.e., ν is unknown. Suppose the goal is to find a strategy that minimizes the worst-case Bayesian cost (for $\nu \in [0, 1]$). Just as in the centralized version of this problem, we extend the class of allowable decision functions to include randomized decision functions. As we remarked in Section 1.3.1, there are two ways to randomize the decision functions: joint randomization and independent randomization. If we allow jointly randomized strategies, then it is quite straightforward to show that an optimal strategy (equalizer rule) can be obtained by randomizing between two deterministic LRT strategies. The open question here is whether equalizer rules exist in the smaller, more useful class of independently randomized strategies, a question akin to the existence of behavioral strategies in dynamic games [53].

REFERENCES

[1] A. Wald, *Sequential Analysis*. New York: Wiley, 1947.

[2] H. V. Poor, *An Introduction to Signal Detection and Estimation*. New York: Springer-Verlag, 1988.

[3] R. R. Tenney and N. R. Sandell, Jr., "Detection with distributed sensors," *IEEE Trans. Aerosp. Electron. Sys.*, vol. 17, pp. 501–510, July 1981.

[4] R. Mallubhatla, K. R. Pattipati, D. L. Kleinman, and Z. Tang, "A model for distributed team information processing under ambiguity," *IEEE Trans. Syst., Man, Cybern.*, vol. 21, no. 4, pp. 713–725, 1991.

[5] Z. Tang, K. R. Pattipati, and D. L. Kleinman, "An algorithm for determining the decision thresholds in a distributed detection problem," *IEEE Trans. Syst., Man, Cybern.*, vol. 21, pp. 231–237, Jan. 1991.

[6] C. H. Papadimitriou and J. N. Tsitsiklis, "Intractable problems in control theory," *SIAM J. Control Optim.*, vol. 24, pp. 639–654, 1986.

[7] H. V. Poor and J. B. Thomas, "Applications of Ali-Silvey distance measures in the design of generalized quantizers for binary decision systems," *IEEE Trans. Commun.*, vol. 25, pp. 893–900, Sept. 1977.

[8] T. J. Flynn and R. M. Gray, "Encoding of correlated observations," *IEEE Trans. Inform. Th.*, vol. IT-33, pp. 773–787, Nov. 1987.

[9] H. V. Poor, "Fine quantization in signal detection and estimation," *IEEE Trans. Inform. Th.*, vol. 34, pp. 960–972, Sept. 1988.

[10] G. R. Benitz and J. A. Bucklew, "Asymptotically optimal quantizers for i.i.d. data," *IEEE Trans. Inform. Th.*, vol. 35, pp. 316–325, Mar. 1989.

[11] J. N. Tsitsiklis, "Extremal properties of likelihood-ratio quantizers," in *Proc. 29th IEEE Conf. Dec. and Contr.*, (Honolulu, Hawaii), pp. 2680–2685, Dec. 1990.

[12] J. N. Tsitsiklis and M. Athans, "On the complexity of decentralized decision making and detection problems," *IEEE Trans. Automat. Contr.*, vol. 30, pp. 440–446, May 1985.

[13] J. N. Tsitsiklis, "Decentralized detection by a large number of sensors," *Math. Contr. Signals Syst.*, vol. 1, pp. 167–182, 1988.

[14] G. S. Lauer and N. R. Sandell, Jr., "Distributed detection of known signal in correlated noise," tech. rep., ALPHATECH, Burlington, MA, 1982.

[15] R. Vishwanathan and A. Ansari, "Distributed detection of a signal in generalized Gaussian noise," *IEEE Trans. Acoust., Speech, Signal Proc.*, vol. 37, pp. 775–778, May 1989.

[16] V. Aalo and R. Vishwanathan, "Asymptotic performance of a distributed detection system in correlated Gaussian noise," *IEEE Trans. Signal Proc.*, vol. 40, pp. 211–213, Jan. 1992.

[17] A. R. Reibman and L. W. Nolte, "Optimal detection performance of distributed sensor systems," *IEEE Trans. Aerosp. Electron. Sys.*, vol. 23, pp. 24–29, Jan. 1987.

[18] A. R. Reibman and L. W. Nolte, "Design and performance comparison of distributed detection networks," *IEEE Trans. Aerosp. Electron. Sys.*, vol. 23, pp. 789–797, Nov. 1987.

[19] I. Y. Hoballah and P. K. Varshney, "Distributed Bayesian signal detection," *IEEE Trans. Inform. Th.*, vol. 35, pp. 995–1000, Sept. 1989.

[20] R. Srinivasan, "Distributed radar detection theory," *Proc. IEEE*, vol. 133, pp. 55–60, Feb. 1986.

[21] I. Y. Hoballah and P. K. Varshney, "Neyman-Pearson detection with distributed sensors," in *Proc. 25th IEEE Conf. Decision & Contr.*, (Athens, Greece), pp. 237–241, Dec. 1986.

[22] S. C. A. Thomopoulos, R. Vishwanathan, and D. P. Bougoulias, "Optimal decision fusion in multiple sensor systems," *IEEE Trans. Aerosp. Electron. Sys.*, vol. 23, Nov. 1987.

[23] M. Barkat and P. K. Varshney, "Decentralized CFAR signal detection," *IEEE Trans. Acoust., Speech, Signal Proc.*, vol. AES-25, pp. 141–149, Mar. 1989.

[24] S. C. A. Thomopoulos, R. Vishwanathan, and D. P. Bougoulias, "Optimal distributed decision fusion," *IEEE Trans. Aerosp. Electron. Sys.*, vol. 25, Nov. 1989.

[25] L. K. Ekchian and R. R. Tenney, "Detection networks," in *Proc. 21st IEEE Conf. Decision Contr.*, (Orlando, FL), pp. 686–691, Dec. 1982.

[26] R. Vishwanathan, S. C. A. Thomopoulos, and R. Tumulri, "Optimal serial distributed decision fusion," *IEEE Trans. Aerosp. Electron. Sys.*, vol. 24, pp. 366–376, July 1988.

[27] T. M. Cover, "Hypothesis testing with finite statistics," *Ann. Math. Statist.*, vol. 40, no. 3, pp. 828–835, 1969.

[28] E. L. Lehmann, *Testing Statistical Hypotheses*. New York: Wiley, 1959.

[29] A. G. Tartakovskii, "Sequential testing of many simple hypotheses with independent observations," *Probl. Inform. Transm.*, vol. 24, pp. 299–309, October-December 1988. Translated from the Russian version.

[30] D. Teneketzis and Y. C. Ho, "The decentralized Wald problem," *Inform. and Comput.*, vol. 73, no. 1, pp. 23–44, 1987.

[31] A. LaVigna, A. M. Makowski, and J. S. Baras, "A continuous-time distributed version of wald's sequential hypothesis testing problem," in *Proceedings of the 7th International Conference on Analysis and Optimization of Systems* (A. Bensoussan and J. L. Lions, eds.), (Antibes, France), pp. 533–543, 1986.

[32] D. Bertsekas, *Dynamic Programming*. Upper Saddle River, NJ: Prentice-Hall, 1987.

[33] V. V. Veeravalli and H. V. Poor, "Quadratic detection of signals with drifting phase," *J. Acoust. Soc. Amer.*, vol. 89, pp. 811–819, Feb. 1991.

[34] D. Siegmund, *Sequential Analysis: Tests and Confidence Intervals*. New York: Springer-Verlag, 1985.

[35] M. Woodroofe, *Nonlinear Renewal Theory in Sequential Analysis*. Philadelphia, PA: SIAM, 1982.

[36] J. N. Tsitsiklis, "On threshold rules in decentralized detection," in *Proc. 25th IEEE Conf. Dec. and Contr.*, (Athens, Greece), pp. 232–236, Dec. 1986.

[37] T. Başar and J. B. Cruz, "Concepts and methods in multi-person coordination and control," in *Optimisation and Control of Dynamic Operational Research Models* (S. G. Tzafestas, ed.), pp. 351–394, Amsterdam: North-Holland Publishing Company, 1982.

[38] H. S. Witsenhausen, "Separation of estimation and control for discrete-time systems," *Proc. IEEE*, vol. 59, pp. 1557–1566, 1971.

[39] H. R. Hashemi and I. B. Rhodes, "Decentralized sequential detection," *IEEE Trans. Inform. Th.*, vol. 35, pp. 509–520, May 1989.

[40] V. V. Veeravalli and H. V. Poor, "Quadratic detection of Lorentzian signals," in *Proc. 24th Ann. Conf. Inform. Sci. Syst.*, (Princeton University, Princeton, NJ), Mar. 1990.

[41] J. N. Tsitsiklis, "Extremal properties of likelihood-ratio quantizers," *IEEE Trans. Commun.*, vol. COM-41, pp. 550–8, Apr. 1993.

[42] J. N. Tsitsiklis, "Decentralized detection," in *Advances in Statistical Signal Processing* (H. V. Poor and J. B. Thomas, eds.), vol. 2, Greenwich, CT: JAI Press, 1993.

[43] R. P. Brent, *Algorithms for Minimization Without Derivatives*. Princeton, NJ: Prentice-Hall, 1973.

[44] H. L. Royden, *Real Analysis*. New York: Macmillan, 1988.

[45] V. P. Kuznetsov, "Stable rules for discrimination of hypotheses," *Probl. Inform. Transm.*, vol. 18, pp. 41–51, Jan-March 1982.

[46] S. A. Kassam and H. V. Poor, "Robust techniques for signal processing: A survey," *Proc. IEEE*, vol. 73, pp. 433–481, Mar. 1985.

[47] E. Geraniotis, "Robust distributed discrete-time block and sequential detection," in *Proc. 1987 Conf. Inform. Sci. Syst.*, (Johns Hopkins Univ., Baltimore, MD), pp. 354–360, Mar. 1987.

[48] E. Geraniotis and Y. A. Chau, "On minimax robust data fusion," in *Proc. 1988 Conf. Inform. Sci. Syst.*, (Princeton University, Princeton, NJ), pp. 876–881, Mar. 1988.

[49] E. Geraniotis and Y. A. Chau, "Robust data fusion for multisensor detection systems," *IEEE Trans. Inform. Th.*, vol. IT-36, pp. 1265–1279, Nov. 1990.

[50] P. J. Huber, "A robust version of the probability ratio test," *Ann. Math. Statist.*, vol. 36, pp. 1753–1758, 1965.

[51] P. J. Huber and V. Strassen, "Minimax tests and the Neyman-Pearson lemma for capacities," *Ann. Statist.*, vol. 1, no. 2, pp. 251–263, 1973.

[52] J. D. Papastavrou and M. Athans, "Distributed detection by a large team of sensors in tandem," in *Proc. 29th Conf. Decision and Contr.*, (Honolulu, Hawaii), pp. 246–251, Dec. 1990.

[53] T. Başar and G. J. Olsder, *Dynamic Noncooperative Game Theory.* New York: Academic Press, 1982.

VITA

Venugopal Veeravalli was born on July 28, 1963. He received his B. Tech degree from the Indian Institute of Technology, Bombay, in 1985, and the M.S. degree from Carnegie Mellon University, Pittsburgh, Pennsylvania, in 1987, both in electrical engineering. From August 1985 to December 1987, he was a graduate research assistant at the Magnetics Technology Center at Carnegie Mellon, and since January 1988, has been a graduate research assistant at the Coordinated Science Laboratory, University of Illinois. In the spring semester of 1992, he taught a graduate course in detection and estimation in the Electrical Engineering Department, University of Illinois. Mr. Veeravalli was the recipient of the 1985 President of India Silver Medal for being the most outstanding student in electrical engineering at the Indian Institute of Technology, Bombay.