

Decentralized Sequential Detection with a Fusion Center Performing the Sequential Test

Venugopal V. Veeravalli, *Student Member, IEEE*, Tamer Başar, *Fellow, IEEE*,
and H. Vincent Poor, *Fellow, IEEE*

Abstract—A decentralized sequential detection problem is considered in which each one of a set of sensors receives a sequence of observations about the hypothesis. Each sensor sends a sequence of summary messages to the fusion center where a sequential test is carried out to determine the true hypothesis. A Bayesian framework for this problem is introduced, and for the case when the information structure in the system is *quasi-classical*, it is shown that the problem is tractable. A detailed analysis of this case is presented along with some numerical results.

Index Terms—Decentralized detection, sequential analysis, dynamic programming.

I. INTRODUCTION

WITH THE INCREASING INTEREST in decentralized detection problems in recent years, extensions of various centralized detection problems to decentralized cases have been formulated and studied [1]. In particular, there has been considerable interest in the solution to decentralized detection problems of a sequential nature [2]–[6]. In decentralized sequential hypotheses testing, each one of a set of sensors receives a sequence of observations about the hypothesis. Two distinct formulations are possible. In one case, first each sensor performs a sequential test on its observations and arrives at a final local decision; subsequently the local decisions are used for a common purpose at a site possibly remote to all the sensors. In the other case, each sensor sends a sequence of summary messages to the fusion center, where a sequential test is carried out to determine the true hypothesis.

In this paper, we study the latter case. More formally, let there be N sensors S_1, \dots, S_N in the system. At time $k \in \{1, 2, \dots\}$, sensor S_l observes a random variable X_k^l , and forms a summary message u_k^l of the information available for

Manuscript received November 1, 1991; revised June 15, 1992. This work was supported in part by the Joint Services Electronics Program under Grant N00014-90-J-1270, through the University of Illinois, and in part by the Office of Naval Research under Grant N00014-89-J-1321, through Princeton University. This work was presented in part at the 1992 American Control Conference, Chicago, IL, June 24–26, 1992.

V. V. Veeravalli was with the Department of Electrical and Computer Engineering, and the Coordinated Science Laboratory, University of Illinois, Urbana, IL. He is now with the Division of Applied Sciences, Harvard University, Cambridge, MA 02139.

T. Başar is with the Department of Electrical and Computer Engineering, and the Coordinated Science Laboratory, University of Illinois, Urbana, IL 61801.

H. V. Poor is with the Electrical Engineering Department, Princeton University, Princeton, NJ 08544.

IEEE Log Number 9203865.

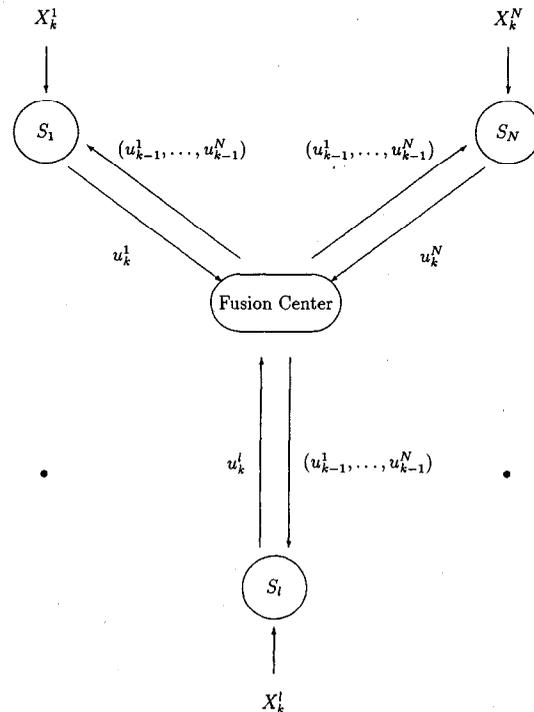


Fig. 1. General setting for decentralized sequential detection with a fusion center performing the sequential test.

decision at time k . In a general setting, we allow a two-way communication between the sensors and the fusion center as shown in Fig. 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 its observations up to time k and all the decisions of all the other sensors up to time $k-1$.

We now introduce a Bayesian framework for this sequential hypothesis testing 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 then is 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 at each of the sensors.

Throughout this paper we shall make the following assumption.

Assumption 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 1, especially when we consider infinite-horizon problems.

Assumption 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 generalized sequential probability ratio tests (GSPRT's) [5]. Namely, at time k , the fusion center forms a likelihood ratio L_k (as a function of all the information it has accumulated) and compares it to two thresholds α_k and β_k , $\alpha_k < \beta_k$. If $L_k \leq \alpha_k$, then H_0 is chosen. If $L_k \geq \beta_k$, then H_1 is chosen. If $\alpha_k < L_k < \beta_k$, 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 [5], 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 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 [7]. Unfortunately, dynamic programming cannot be used here because of the nonclassical² nature of the information in the system [8], [9], thus leaving this as an open problem.

Case B) System with no Feedback, but Full Local Memory:

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

Hashemi and Rhodes [6] 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 [5] which predates [6]). We point out this mistake in

¹A 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.

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

[10], where we also argue that likelihood ratio tests are indeed optimal if we restrict u_k^l to depend on X_k^l and $(u_1^l, \dots, u_{k-1}^l)$, as given below in Case C.

Case C) System with no Feedback, and Local Memory Restricted to Past Decisions:

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

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 the information that sensor S_l has access to in the setting of Fig. 1, i.e.,³

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

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:

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

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), implies that the information structure for this case is *quasi-classical*. It is well-known that stochastic control or team problems with such an information structure are tractable via DP arguments [8], [9].

In the remainder of this paper, we study Case E in detail. As we will show, definite progress can be made in the analysis of this case. In Section II, we provide a formal mathematical description of the problem. In Section III, we provide a useful characterization of sensor decision functions. In Section IV, we consider a finite-horizon version of the problem and establish the optimality of likelihood ratio tests at the sensors. Then, in Section V, 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 has a simple SPRT-like structure. In Section VI, we provide some numerical results. Finally, in Section VII, we include some concluding remarks.

II. 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 that 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 X_l . The sequences

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

$\{X_k^1\}_{k=1}^\infty, \{X_k^2\}_{k=1}^\infty, \dots, \{X_k^N\}_{k=1}^\infty$ 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_k^l given H_j .

- 3) At time k , sensor S_l sends, to the fusion center, a local decision u_k^l that takes values in the finite set $\{1, \dots, D_l\}$. Past decision information for *all* 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

$$I_{k-1} = \{u_{[1, k-1]}^1, u_{[1, k-1]}^2, \dots, u_{[1, k-1]}^N\},$$

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

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

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 $\{1, \dots, D_l\}$. We denote this mapping by ϕ_k^l . The local decision u_k^l is then given by

$$u_k^l = \phi_k^l(X_k^l; I_{k-1}).$$

But for a particular realization i_{k-1} , the LDF ϕ_k^l can be considered to be a mapping from \mathcal{X}_l to $\{1, \dots, D_l\}$, which we denote by $\phi_{k, i_{k-1}}^l$, i.e., $\phi_k^l(\cdot; i_{k-1}) \equiv \phi_{k, i_{k-1}}^l(\cdot)$. The set of all LDF's at time k is represented by the vector

$$\phi_k = (\phi_k^1, \dots, \phi_k^N).$$

- 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 decisionmaking is assumed to cost a positive amount c .

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

Problem (P1): 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.

III. LOCAL DECISION FUNCTIONS

The decision function $\phi_{k, i_{k-1}}^l$ defined in Section II is a mapping from \mathcal{X}_l to $\{1, \dots, D_l\}$. Let Φ^l denote the set of all mappings from \mathcal{X}_l to $\{1, \dots, D_l\}$. 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 $\{X_k^l\}_{k=1}^\infty$. Then, we

define the following:

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

The vector q_{ϕ^l} describes the conditional distributions of $\phi^l(X^l)$, conditioned on each of the hypotheses. Let $Q^l := \{q_{\phi^l} | \phi^l \in \Phi^l\}$. We state the following result which was proved in [11] in the context of optimal likelihood ratio quantizers.

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

To utilize this result in our framework, we concatenate the mappings ϕ^l , $l = 1, \dots, N$, into the vector $\phi = (\phi^1, \dots, \phi^N)$, and define

$$\mathbf{q}_\phi := (q_{\phi^1}, \dots, q_{\phi^N}).$$

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

Now, since $\phi_{k, i_{k-1}}^l \in \Phi^l$, $l = 1, \dots, N$, the vector $\mathbf{q}_{\phi_{k, i_{k-1}}}$ is well defined and describes the joint distribution of the observation vector $\mathbf{u}_k = (u_k^1, \dots, u_k^N)$, conditioned on each hypothesis and on the event that $I_{k-1} = i_{k-1}$. Note that two LDF's ϕ_k^l and $\hat{\phi}_k^l$ are equivalent, i.e., their use results in the same expected cost for the sequential test by the fusion center, if

$$q_{\phi_{k, i_{k-1}}^l} = q_{\hat{\phi}_{k, i_{k-1}}^l} \quad \text{a.e.}$$

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

IV. FINITE-HORIZON OPTIMIZATION

Before we address the solution of the infinite-horizon optimization problem (P1), 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 the LDF's 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 $\mathcal{X}_{[1, T]}$ denote the set of all observations up to time T , i.e., $\{X_{[1, T]}^1, \dots, X_{[1, T]}^N\}$. Then, the finite-horizon optimization problem can then be stated as follows.

Problem (P2): Minimize

$$E_{\mathcal{X}_{[1, T]}, H} G_\gamma(u_{[1, T]}^1, \dots, u_{[1, T]}^N; H)$$

over all possible choices of γ and all possible choices of ϕ_k^l , $l = 1, \dots, N$, $k = 1, \dots, T$.

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

A. The Structure of Optimal LDF's

We can characterize p.b.p.o. LDF's as follows. We first fix l , $1 \leq l \leq N$, and k , $1 \leq k \leq T$. Then, we fix the policy γ and all the LDF's in the set $\{\varphi_m^j, j = 1, \dots, N, m = 1, \dots, T\}$, 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

$$R_{\varphi_k^l} = E_{u_{[1, k-1]}^1, \dots, u_{[1, k-1]}^N, X_k^l, H} \cdot \{E_{X_{[k+1, T]}^1, X_{[k, T]}^1, \dots, X_{[k, T]}^{l-1}, X_{[k, T]}^{l+1}, \dots, X_{[k, T]}^N | H} \cdot \{G_\gamma(\varphi_k^l(X_k^l; I_{k-1}), I_{k-1}, u_k^1, \dots, u_k^{l-1}, u_k^{l+1}, \dots, u_k^N, u_{[k+1, T]}^1, \dots, u_{[k+1, T]}^N; H)\}\}.$$

In the previous inner expectation, we do not need to condition on observations up to time $k-1$ because of the conditional independence assumption stated in Section II. 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 previous equation is a function of I_{k-1} , H , and $\varphi_k^l(X_k^l; I_{k-1})$, say $K(\varphi_k^l(X_k^l; I_{k-1}), I_{k-1}; H)$. Therefore,

$$R_{\varphi_k^l} = E_{I_{k-1}, X_k^l, H} \{K(\varphi_k^l(X_k^l; I_{k-1}), I_{k-1}; H)\} \\ = E_{I_{k-1}, X_k^l} \{\text{Prob}(H = H_0 | I_{k-1}, X_k^l) \cdot K(\varphi_k^l(X_k^l; I_{k-1}), I_{k-1}; H_0) \\ + \text{Prob}(H = H_1 | I_{k-1}, X_k^l) \cdot K(\varphi_k^l(X_k^l; I_{k-1}), I_{k-1}; H_1)\}.$$

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

$$\hat{\varphi}_k^l(X_k^l; I_{k-1}) \\ = \arg \min_{d_l \in \{1, \dots, D_l\}} \{\text{Prob}(H = H_0 | I_{k-1}, X_k^l) \cdot K(d_l, I_{k-1}; H_0) + \text{Prob}(H = H_1 | I_{k-1}, X_k^l) \cdot K(d_l, I_{k-1}; H_1)\} \text{ a.e.} \quad (1)$$

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

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

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

Definition 1:

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

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

where $I_1 = [0, \lambda_1]$, $I_{D_l} = [\lambda_{D_l-1}, \infty]$, and $I_{d_l} = [\lambda_{d_l-1}, \lambda_{d_l}]$, $d_l = 2, \dots, D_l - 1$.

- b) A decision function $\phi^l: \mathcal{X}_l \mapsto \{1, \dots, D_l\}$ is called a *likelihood ratio test* (LRT) if there exists a permutation mapping $\Sigma: \{1, \dots, D_l\} \mapsto \{1, \dots, D_l\}$ such that the composite function $\Sigma \circ \phi^l$ is a monotone likelihood ratio test.

Proposition 2: Person-by-person optimal local decision functions (when they exist) can be found in the class of LRT's, 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 (1). Using Bayes rule, we have

$$\frac{\text{Prob}(H = H_1 | I_{k-1}, X_k^l)}{\text{Prob}(H = H_0 | I_{k-1}, X_k^l)} = L_l(X_k^l) \frac{(1 - p_{k-1})}{p_{k-1}}, \text{ w.p. 1, (2)}$$

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

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

From (2), it follows that $\hat{\varphi}_k^l$ satisfies the equation at the bottom of the page. From this, it should be clear that there exists a solution for $\hat{\varphi}_k^l$ in the class of LRT's with thresholds that depend on I_{k-1} . \square

B. 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 (P2) is the posterior probability defined earlier, i.e.,

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

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, $g: \mathcal{D} \times \mathcal{Q} \times [0, 1] \mapsto [0, 1]$ and $f: \mathcal{D} \times \mathcal{Q} \times [0, 1] \mapsto [0, 1]$, as follows.

⁴Similar definitions can be found in [11].

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

For $\mathbf{d} = (d_1, \dots, d_N) \in \mathcal{D}$, $\phi \in \Phi^1 \times \dots \times \Phi^N$, and $p \in [0, 1]$,

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

and

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

Note that $f(\cdot; \mathbf{q}_{\phi_{k, I_{k-1}}}; p_{k-1})$ is the joint conditional distribution of $\mathbf{u}_k = (u_k^1, \dots, u_k^N)$, given I_{k-1} .

We now give the recursion for p_k . For $k = 0, \dots, T$,

$$\begin{aligned} p_{k+1} &= \text{Prob}(H = H_0 | I_{k+1}) \\ &= \text{Prob}(H = H_0 | u_{k+1}^1, \dots, u_{k+1}^N, I_k) \\ &= \frac{\text{Prob}(H = H_0 | I_k) p(u_{k+1}^1, \dots, u_{k+1}^N | H_0, I_k)}{p(u_{k+1}^1, \dots, u_{k+1}^N | I_k)} \\ &= \frac{g(\mathbf{u}_{k+1}; \mathbf{q}_{\phi_{k+1, I_k}}; p_k)}{f(\mathbf{u}_{k+1}; \mathbf{q}_{\phi_{k+1, I_k}}; p_k)}, \end{aligned} \quad (3)$$

with $p_0 = \nu$. However, we will not find this recursion useful unless we can show that the RHS of (3) depends on I_k only through p_k , i.e., that p_k is indeed a sufficient statistic for (P2).

C. Finite-Horizon DP

The DP equations for problem (P2) 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:

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

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

Proposition 3:

- 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 $J_k^T(p_k)$.
- For each k , $0 \leq k \leq T-1$, the search for optimal LDF's 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,

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

In this equation, it is clear that the conditional distribution vector \mathbf{q}_{ϕ^*} is a function of only p_k . Both parts of the proposition follow from this fact. \square

The main consequence of Proposition 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,

$$u_k^l = \phi_k^l(X_k^l; p_{k-1}). \quad (4)$$

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 LDF's defined as in (4), we obtain the following useful recursion for p_k . For $k = 0, \dots, T-1$, we have

$$p_{k+1} = \frac{g(\mathbf{u}_{k+1}; \mathbf{q}_{\phi_{k+1, p_k}}; p_k)}{f(\mathbf{u}_{k+1}; \mathbf{q}_{\phi_{k+1, p_k}}; p_k)} \quad (5)$$

with $p_0 = \nu$. All the decision makers in the system have to retain only the sufficient statistic p_k , which they can easily update using (5).

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

$$J_T^T(p_T) = \min \{(1-p_T)L_0, p_T L_1\}, \quad (6)$$

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

$$J_k^T(p_k) = \min \{(1-p_k)L_0, p_k L_1, c + A_k^T(p_k)\}, \quad (7)$$

where

$$A_k^T(p_k) := \inf_{\mathbf{q}_\phi \in \mathcal{Q}} \sum_{\mathbf{d} \in \mathcal{D}} J_{k+1}^T \left(\frac{g(\mathbf{d}; \mathbf{q}_\phi; p_k)}{f(\mathbf{d}; \mathbf{q}_\phi; p_k)} \right) f(\mathbf{d}; \mathbf{q}_\phi; p_k). \quad (8)$$

D. Finite-Horizon Policy of the Fusion Center

Our goal in this section is to use the DP equations (6)–(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 1: The functions $J_k^T(p)$ and $A_k^T(p)$ are nonnegative concave functions of p , for $p \in [0, 1]$.

Lemma 2: The functions $J_k^T(p)$ and $A_k^T(p)$ are monotonically nondecreasing in k , that is, for each $p \in [0, 1]$,

$$\begin{aligned} J_k^T(p) &\leq J_{k+1}^T(p), & 0 \leq k \leq T-1. \\ A_k^T(p) &\leq A_{k+1}^T(p), & 0 \leq k \leq T-2. \end{aligned}$$

Lemma 3: The functions $A_k^T(p)$ satisfy the following property:

$$A_k^T(0) = A_k^T(1) = 0. \quad \square$$

Lemmas 2 and 3 are easily proven by simple induction arguments. The proof of Lemma 1 is not as straightforward and is given in the Appendix section.

If we now assume that

$$A_{T-1}^T \left(\frac{L_0}{L_1 + L_0} \right) \leq \frac{L_0 L_1}{L_1 + L_0}, \quad (9)$$

holds, then Lemmas 1–3 give us the following threshold property of an optimal finite-horizon fusion center policy (see Section 3.5 of [7] for a similar analysis).

Theorem 1: Let (9) hold. Then an optimal finite-horizon fusion center policy has the form

$$\begin{array}{ll} \text{accept } H_0, & \text{if } p_k \geq a_k^T, \\ \text{accept } H_1, & \text{if } p_k \leq b_k^T, \\ \text{continue,} & \text{if } b_k^T < p_k < a_k^T, \end{array}$$

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

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

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 1: If (9) does not hold, then the thresholds a_k^T and b_k^T of Theorem 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 (9) does not impose any restrictions on the problem parameters.

E. Optimal Finite-Horizon LDF's

The DP equations (6)–(8) can also be used to find optimal LDF's stagewise, starting from time T and going backwards. The concavity of the cost-to-go function J_{k+1}^T implies that the function

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

is continuous in \mathbf{q}_ϕ . By Proposition 1, this fact implies the *existence* of optimal LDF's at time $k+1$.

We showed earlier (Proposition 2) that the search for globally optimal LDF's at time $k+1$ can be restricted to the set of LRT's with thresholds depending on I_k . Propositions 1 and 3 further show that globally optimal LDF's at time $k+1$ can be *found* in a class of LRT's 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 LDF's at time $k+1$ can be found in the smaller class of MLRT's with thresholds that depend on p_k .

Now, suppose an MLRT $\hat{\phi}_{k,p_k}^l$ is characterized by the thresholds $\lambda_1(p_k), \dots, \lambda_{D_l-1}(p_k)$. Then,

$$q_{\hat{\phi}_{k,p_k}^l}^j(d_l) = \text{Prob}(L(X_k^l) \in [\lambda_{d_l-1}(p_k), \lambda_{d_l}(p_k)] | H_j),$$

with the understanding that $\lambda_0(p_k) = 0$ and $\lambda_{D_l}(p_k) = \infty$. Hence, the minimization to obtain $A_k^T(p_k)$ in (8) can be done over $|\mathcal{D}|$ thresholds.

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

$$\mathcal{Q}_M = \{\mathbf{q}_\phi \in \mathcal{Q}: \phi \text{ is a vector of MLRT's}\},$$

then $A_k^T(p_k)$ can be written as

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

V. INFINITE-HORIZON OPTIMIZATION

In order to solve the original optimization problem (P1), we need to remove the restriction that τ belongs to a finite interval, by letting $T \rightarrow \infty$. Toward this end, we first note the inequality

$$J_k^{T+1}(p) \leq J_k^T(p),$$

which holds because the set of stopping times increases with T . Furthermore, by leaving out the third term in (7), we obtain

$$0 \leq J_k^T(p) \leq \eta(p), \quad \text{for all } T, \quad \text{and for all } k \leq T,$$

where

$$\eta(p) = \min\{L_1 p, L_0(1-p)\}. \quad (10)$$

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

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

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 :

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

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

$$J(p) = \min\{L_1 p, L_0(1-p), c + A_J(p)\}. \quad (11)$$

We note that the optimum cost for problem (P1) is $J(\nu)$.

A. 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 (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 $T \rightarrow \infty$ in Lemmas 1–3.

Lemma 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 4, it is clear that, provided the condition

$$J\left(\frac{L_0}{L_1 + L_0}\right) < \frac{L_1 L_0}{L_1 + L_0} \quad (12)$$

holds, an optimal policy of the supervisor will have the threshold property given in Theorem 2 (see [7, Section 6.3] for a similar analysis).

Theorem 2: Let condition (12) hold. Then an optimal fusion center policy for problem (P1) has the form

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

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 2: It should be noted that if (12) does not hold, then it would be optimal for the fusion center to ignore all the data it receives from the sensors, and base its decision solely on the value of the prior probability ν . Hence, (12) does not bring any loss of generality to the result of Theorem 2.

B. 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 (10). For $G \in S$, we define

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

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

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

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

$$\begin{aligned} TG(p) &= \min \{ L_1 p, L_0(1 - p), c \\ &\quad + \min_{\mathbf{q}_\phi \in \mathcal{Q}_M} W_G(\mathbf{q}_\phi; p) \}, \quad \text{for } G \in S. \end{aligned}$$

Theorem 3: The infinite-horizon cost-to-go function J is the unique fixed point of the mapping T .

Proof: Let G be any fixed point of T , and let ϕ_p^* be such that

$$\mathbf{q}_{\phi_p^*} = \arg \min_{\mathbf{q}_\phi \in \mathcal{Q}_M} W_G(\mathbf{q}_\phi; p).$$

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

$$p_{k+1} = \frac{g(\mathbf{u}_{k+1}; \mathbf{q}_{\phi_{p_k}^*}; p_k)}{f(\mathbf{u}_{k+1}; \mathbf{q}_{\phi_{p_k}^*}; p_k)}.$$

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

$$\mathcal{N} = \min \{ k \geq 0 \mid \eta(p_k) \leq c + W_G(\mathbf{q}_{\phi_{p_k}^*}; p_k) \},$$

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 T , we obtain the following relations:

$$\begin{aligned} G(\nu) &= c + E\{G(p_1)\} \\ G(p_1) &= c + E\{G(p_2) \mid I_1\} \end{aligned}$$

$$\begin{aligned} G(p_{\mathcal{N}-1}) &= c + E\{G(p_{\mathcal{N}}) \mid I_{\mathcal{N}-1}\} \\ G(p_{\mathcal{N}}) &= \eta(p_{\mathcal{N}}). \end{aligned}$$

Substituting backwards and taking expectations, we obtain

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

where the last inequality follows from the definition of J .

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

$$G(p) \leq \eta(p) = J_T^T(p), \quad \text{for all } T.$$

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

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

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

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

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

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

□

The first important consequence of Theorem 3 is that $J(p)$ can be obtained by successive approximation. We can show, using an induction argument, that

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

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

C. Optimal Infinite-Horizon LDF's

Theorem 3 also implies that a stationary vector of LDF's is optimal for the infinite-horizon problem (P1), as the following argument shows. Let ϕ_p^* be such that

$$\mathbf{q}_{\phi_p^*} = \arg \min_{\mathbf{q}_{\phi} \in \mathcal{Q}_M} W_J(\mathbf{q}_{\phi}; p),$$

where $J(p)$ is the infinite-horizon cost-to-go function for problem (P1). Then, in the problem setting for (P1), we restrict ourselves to the singleton vector of LDF's $\phi^* = (\phi^{1*}, \dots, \phi^{N*})$, where $\phi^{l*}: \mathcal{X}^l \times [0, 1] \mapsto \{1, \dots, D_l\}$ is such that

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

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

$$u_k^l = \phi^{l*}(X_l^k; p_{k-1}), \quad \text{for all } k.$$

We denote the optimization problem with this restriction by (P1'). We can solve (P1') in a manner parallel to the way we solved (P1), 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 (P1') satisfies

$$J'(p) = \min \{L_1 p, L_0(1-p), c + W_{J'}(\mathbf{q}_{\phi_p^*}; p)\}.$$

By Theorem 3, it follows that $J(p) = J'(p)$, for all $p \in [0, 1]$, which implies the optimality of the stationary vector of LDF's ϕ^* for problem (P1).

VI. NUMERICAL RESULTS

For all the examples presented in this section, we assume that the local decisions are binary. For these examples, it is convenient to write the LDF's 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 1) Single Sensor: Here, the LDF is characterized by a single threshold λ . Hence, for each $G \in \mathcal{S}$, 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

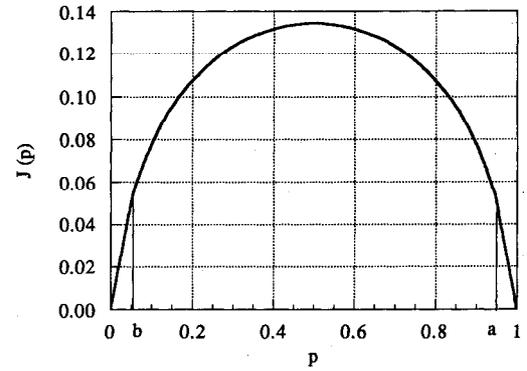
$$\begin{aligned} g(d, \lambda, p) &= p[P_0(L(X) > \lambda)]^{d-1} [P_0(L(X) \leq \lambda)]^{2-d}, \\ f(d, \lambda, p) &= g(d, \lambda, p) + (1-p) \\ &\quad \cdot [P_1(L(X) > \lambda)]^{d-1} [P_1(L(X) \leq \lambda)]^{2-d}. \end{aligned}$$

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

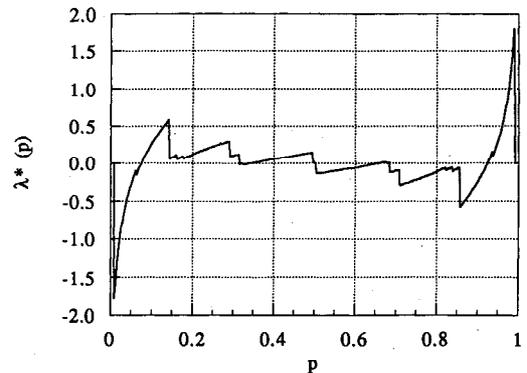
$$\lim_{\lambda \rightarrow \infty} W_G(\lambda, p) = \lim_{\lambda \rightarrow -\infty} W_G(\lambda, p) = G(p).$$

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

$$W_G(\lambda, p) \leq G(p), \quad \text{for all } \lambda.$$



(a)



(b)

Fig. 2. 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) Optimum stationary LDF threshold.

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 [12].

Example 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 $\lambda^*(p)$ 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 $G \in \mathcal{S}$. We have hence used a *golden section* search procedure [12] to obtain an optimal threshold at each stage of the successive approximation. Representative results are shown in Fig. 2. 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 $\lambda^*(p)$ 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 $\text{sgn}(x)$.

Case 2) Two Identical Sensors: Here, in addition to Assumption 2, we assume that the observations received by the two sensors are identically distributed conditioned on each hypothesis. The vector of LDF’s 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

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

where

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

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

Example 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 $G \in \mathcal{S}$, $W_G(\lambda_1, \lambda_2, p)$ is unimodal on the set $\{(\lambda_1, \lambda_2): (\lambda_1, \lambda_2) \in \mathbb{R}^2\}$. 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 Fig. 3) where the optimal thresholds $\lambda_1^*(p)$ and $\lambda_2^*(p)$ are seen to be identical functions of p . The thresholds at each iteration were found by a two-dimensional golden section search procedure. 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 3) Two Nonidentical Sensors: This case is similar to Case 2 except that functions f and g are given by

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

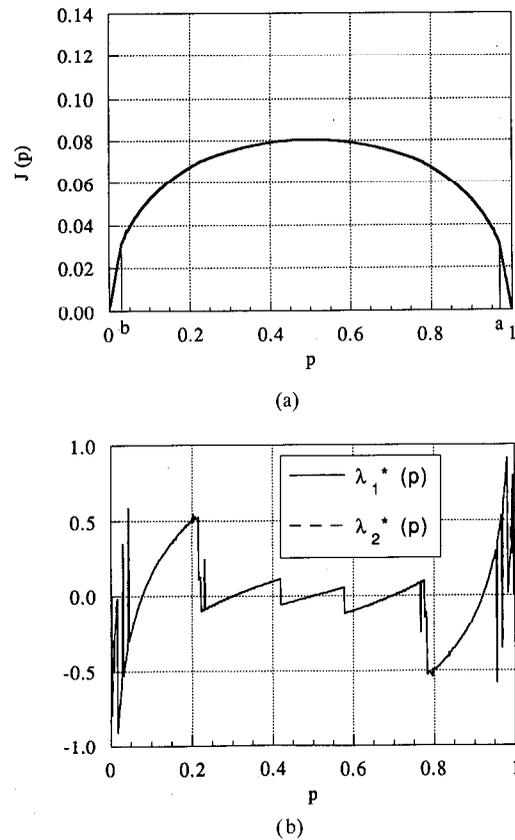


Fig. 3. Results for the case of two identical sensors with $c = 0.01$, $v = 1.0$, and $L_0 = L_1 = 1.0$. (a) Infinite-horizon cost-to-go function. (b) Optimum stationary LDF thresholds.

$$\begin{aligned} &\cdot [P_1(L(X_l) > \lambda_l)]^{d_l-1} \\ &\cdot [P_1(L(X_l) \leq \lambda_l)]^{2-d_l}, \end{aligned}$$

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

Example 3: 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 $G \in \mathcal{S}$, $W_G(\lambda_1, \lambda_2, p)$ is unimodal on the set $\{(\lambda_1, \lambda_2): (\lambda_1, \lambda_2) \in \mathbb{R}^2\}$. Optimal thresholds at each iteration were hence found by a two-dimensional golden section search procedure. Representative results are shown in Fig. 4. A hundred iterations were run, and, as before, the norm difference between the 99th and 100th iterates was less than 10^{-4} .

VII. DISCUSSION

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

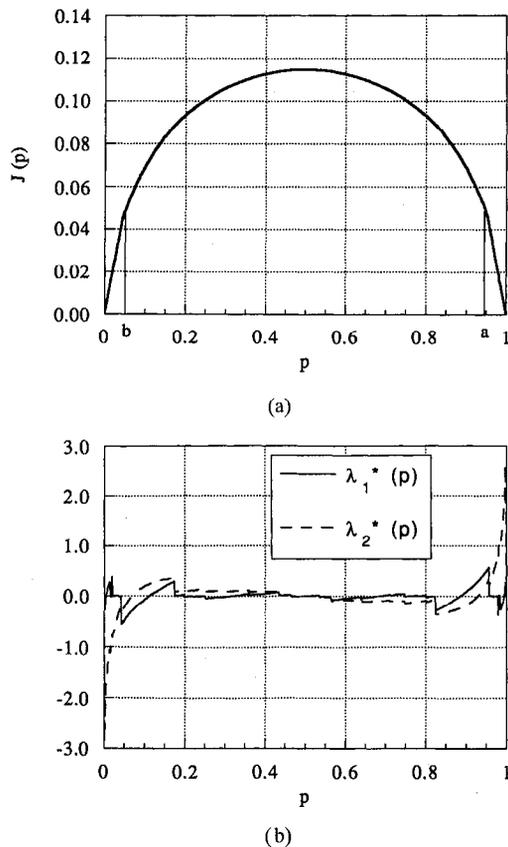


Fig. 4. Results for the case of two nonidentical sensors with $c = 0.01$, $v = 1.0$, and $L_0 = L_1 = 1.0$. (a) Infinite-horizon cost-to-go function. (b) Optimum stationary LDF thresholds.

- 1) 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} .
- 2) 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.
- 3) 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.

The reason we were able to use dynamic programming arguments in our analysis to obtain optimality results is that all the decisionmakers in the system have the same dynamic information about the past. This is not true for the information patterns of Cases A, B, C, and D, as discussed in Section I. An interesting open problem for these cases would be to investigate if stationary LDF's are optimal under Assumption 2. Such a result would be especially useful for Cases A and C since we have already established the optimality of likelihood

ratio tests. Also, if we do not allow feedback from the fusion center, 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 this paper.

VIII. APPENDIX

Proof of Lemma 4.1: 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 [13]

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

Then,

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

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

ACKNOWLEDGMENT

The authors would like to thank the reviewers for their helpful comments and suggestions.

REFERENCES

- [1] J. N. Tsitsiklis, "Decentralized detection," in *Advances in Statistical Signal Processing, Vol. 2—Signal Detection*, H. V. Poor and J. B. Thomas, Eds. Greenwich, CT: JAI Press, 1990.
- [2] D. Teneketzis and Y. C. Ho, "The decentralized Wald problem," *Inform. Comput.*, vol. 73, pp. 23–44, 1987.
- [3] A. LaVigna, A. M. Makowski, and J. S. Baras, "A continuous-time distributed version of Wald's sequential hypothesis testing problem," in *Proc. 7th Int. Conf. Anal. and Optimization of Syst.*, A. Bensoussan and J. L. Lions, Eds. Antibes, France, 1986, pp. 533–543.
- [4] M. M. Al-Ibrahim and P. K. Varshney, "A simple multi-sensor sequential detection procedure," in *Proc. 27th IEEE Conf. Decision Contr.*, Austin, TX, Dec. 1988, pp. 2479–2483.
- [5] J. N. Tsitsiklis, "On threshold rules in decentralized detection," in *Proc. 25th IEEE Conf. Decision Contr.*, Athens, Greece, 1986, pp. 232–236.
- [6] H. R. Hashemi and I. B. Rhodes, "Decentralized sequential detection," *IEEE Trans. Inform. Theory*, vol. 35, pp. 509–520, May 1989.
- [7] D. Bertsekas, *Dynamic Programming*. Englewood Cliffs, NJ: Prentice-Hall, 1987.
- [8] T. Başar and J. B. Cruz, "Concepts and methods in multi-person coordination and control," in *Optimization and Control of Dynamic Operational Research Models*, S. G. Tzafestas, Ed. Amsterdam: North-Holland, 1982, pp. 351–394.
- [9] H. S. Witsenhausen, "Separation of estimation and control for discrete-time systems," *Proc. IEEE*, vol. 59, pp. 1557–1566, 1971.
- [10] V. V. Veeravalli, "Comments on 'Decentralized sequential detection,'" *IEEE Trans. Inform. Theory*, vol. 38, pp. 1428–1429, July 1992.
- [11] J. N. Tsitsiklis, "Extremal properties of likelihood-ratio quantizers," in *Proc. 29th IEEE Conf. Decision Contr.*, Honolulu, HI, Dec. 1990, pp. 2680–2685.
- [12] R. P. Brent, *Algorithms for Minimization Without Derivatives*. Englewood Cliffs, NJ: Prentice-Hall, 1973.
- [13] H. L. Royden, *Real Analysis*. New York: Macmillan, 1988.