Math. Control Signals Systems (1994) 7: 292–305

Mathematics of Control, Signals, and Systems

Decentralized Sequential Detection with Sensors Performing Sequential Tests*

Venugopal V. Veeravalli,†‡ Tamer Başar,† and H. Vincent Poor§

Abstract. A decentralized sequential detection problem is considered where a set of sensors making independent observations must decide which of the given two hypotheses is true. Decision errors are penalized through a common cost function, and each time step taken by the sensors as a team is assigned a positive cost. It is shown that optimal sensor decision functions can be found in the class of *generalized sequential probability ratio tests* (GSPRTs) with monotonically convergent thresholds. A technique is presented for obtaining the optimal thresholds. The performance of the optimal policy is compared with that of a policy which uses SPRTs at each of the sensors.

Key words. Decentralized detection, Sequential analysis, Dynamic programming, Distributed decision making, Optimal stopping rules, Stochastic teams.

1. Introduction

Static decentralized detection problems are well understood today and most tractable problems in this area have been solved (for a comprehensive report see [T3]). There has also been considerable interest in the related field of decentralized sequential detection (see, for example, [TH], [LMB], [AV], [T2], [HR], and [VBP2]); however, this area has not witnessed significant progress, and several interesting problems still remain open. The goal of this paper is to address one of these problems.

As an introduction to sequential hypothesis testing, let us first consider the case when the information is centralized and the number of hypotheses is two. Here 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 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

* Date received: August 3, 1992. This research was supported in part by the Joint Services Electronics Program under Grant N00014-90-J-1270, through the University of Illinois.
† Department of Electrical and Computer Engineering, and the Coordinated Science Laboratory, University of Illinois, Urbana, Illinois 61801, U.S.A.

‡ Current address: Department of Electrical Engineering, Rice University, Houston, Texas, U.S.A.

§ Electrical Engineering Department, Princeton University, Princeton, New Jersey, U.S.A.

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 in which the observations are independent and identically distributed (i.i.d.), conditioned on each hypothesis, is the well-known *sequential probability ratio test* (SPRT) [L].

In *decentralized* sequential detection, each one of a set of sensors receives a sequence of observations about the hypothesis. Two distinct formulations are possible. 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 first 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 this paper we study the latter case. 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 where 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 [TH] considered the situation where 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 arises when 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 the case when 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 [TH] that

[1] We 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 [T1].

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 [LMB], where a continuous time extension of this problem was solved.

In our analysis in this paper, 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 [TH]. An earlier version of this paper was presented in [VBP1].

The rest of this paper is organized as follows. In Section 2 we provide a more formal description of the problem. Then in Section 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 4 we address the problem of finding optimal GSPRT thresholds numerically. In Section 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 using SPRTs. Finally, Section 6 gives some concluding remarks.

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 q and $1-q$, 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 $\mathscr{X}_k^i = \sigma(X_j^i, j = 1, 2, \ldots, k)$. 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 $\{\mathscr{X}_k^i, k = 1, 2, \ldots\}$ stopping time, and u_i is measurable $\mathscr{X}_{\tau_i}^i$. 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 (P1).

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

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:

$$\min_{\{\gamma_1 \in \Gamma_1\}} E\{c \max(\tau_1, \tau_2) + W(u_1, u_2; H)\}. \tag{1}$$

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

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

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

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

where $f_j(\cdot)$ 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 $\mathscr{X}_k^1$, which we denote by $f(p_k; \cdot)$, is given by

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

In order to solve the problem of (1), we first restrict the stopping time τ_1 to a finite horizon, 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

[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.

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

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

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

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:

$$J_k^T(p_k) = \min\{G_0 p_k + K_0, G_1 p_k + K_1, c p_k P_0(\tau_2 \le k) + c(1 - p_k) P_1(\tau_2 \le k) + \Lambda_k^T(p_k)\}, \tag{2}$$

where

$$\Lambda_k^T(p_k) = E_{X_{k+1}^1 | \mathscr{X}_k^1} \left\{ J_{k+1}^T \left(\frac{p_k f_0(X_{k+1}^1)}{f(p_k; X_{k+1}^1)} \right) \right\} = \int J_{k+1}^T \left(\frac{p_k f_0(x)}{f(p_k; x)} \right) f(p_k; x)\, dx. \tag{3}$$

In (2) the term $G_0 p_k + K_0$ represents the cost (conditioned on $\mathscr{X}_k^1$) 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 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. *The functions $J_k^T(p)$ and $\Lambda_k^T(p)$ are monotonically nondecreasing in k, that is, for each $p \in [0, 1]$,*

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

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

Lemma 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. Let us now assume that the following condition holds:

$$c p^* P_0(\tau_2 \le T - 1) + c(1 - p^*) P_1(\tau_2 \le T - 1) + \Lambda_{T-1}^T(p^*) \le \frac{G_1 K_0 - G_0 K_1}{G_1 - G_0}, \tag{4}$$

with p^* defined by

$$p^* = \frac{K_0 - K_1}{G_1 - G_0}.$$

Then, using Lemmas 1–3, we can prove the following threshold property of an optimal finite-horizon policy at S_1 (see Section 3.5 of [B] for a similar analysis).

Theorem 1. *For fixed $\gamma_2 \in \Gamma_2$, let condition* (4) *hold. Then an optimal finite-horizon policy at sensor S_1 is of the form*

$$\begin{array}{lll} \textit{accept } H_0 & \textit{if} & p_k \geq a_k^T, \\ \textit{accept } H_1 & \textit{if} & p_k \leq b_k^T, \\ \textit{continue taking observations} & \textit{if} & b_k^T < p_k < a_k^T, \end{array}$$

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

$$G_1 b_k^T + K_1 = c b_k^T P_0(\tau_2 \leq k) + c(1 - b_k^T) P_1(\tau_2 \leq k) + \Lambda_k^T(b_k^T),$$

$$G_0 a_k^T + K_0 = c a_k^T P_0(\tau_2 \leq k) + c(1 - a_k^T) P_1(\tau_2 \leq k) + \Lambda_k^T(a_k^T).$$

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 condition (4) does not hold, then the thresholds a_k^T and b_k^T of Theorem 1 are both identically equal to $(K_0 - K_1)/(G_1 - G_0)$ for all k greater than some m, $1 \leq m < T$, which essentially reduces the finite horizon to m. Hence, condition (4) does not impose any restrictions on the problem parameters.

3.1. *Infinite-Horizon Optimization*

In order to solve the infinite-horizon problem of (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 [TH], 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 [TH], this limit need not be independent of k. In fact, if we let $T \to \infty$ in Lemma 2, we see that the following monotonicity holds in the limit:

$$J_k(p) \leq J_{k+1}(p), \qquad \text{for all } k.$$

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 [B]

$$J(p) = \min\left\{G_0 p + K_0,\, G_1 p + K_1,\, c + \int J\left(\frac{pf_0(x)}{f(p;x)}\right) f(p;x)\,dx\right\}. \tag{5}$$

Teneketzis and Ho 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 [TH]).

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

$$\Lambda_k(p) := \lim_{T\to\infty} \Lambda_k^T(p) = \int J_{k+1}\left(\frac{pf_0(x)}{f(p;x)}\right) f(p;x)\,dx$$

and

$$\Lambda_J(p) := \lim_{k\to\infty} \Lambda_k(p) = \int J\left(\frac{pf_0(x)}{f(p;x)}\right) f(p;x)\,dx.$$

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

$$J_k(p) = \min\{G_0 p + K_0,\, G_1 p + K_1,\, cpP_0(\tau_2 \le k) + c(1-p)P_1(\tau_2 \le k) + \Lambda_k(p)\}. \tag{6}$$

Taking limits as $T \to \infty$ in Lemmas 1–3, we obtain the following result.

Lemma 4. *The functions $\Lambda_k(p)$ are concave and satisfy*

$$\Lambda_k(p) \le \Lambda_{k+1}(p), \qquad \text{for all} \quad p \in [0,1],$$
$$\Lambda_k(0) = K_1, \qquad \Lambda_k(1) = K_0 + G_0.$$

It follows from Lemma 4 that provided the condition

$$c + \Lambda_J\left(\frac{K_0 - K_1}{G_1 - G_0}\right) \le \frac{G_1 K_0 - G_0 K_1}{G_1 - G_0} \tag{7}$$

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

Theorem 2. *For fixed $\gamma_2 \in \Gamma_2$, let condition (7) hold. Then an optimal infinite-horizon policy at sensor S_1 is of the form*

$$\begin{array}{lll} \textit{accept } H_0 & \textit{if} & p_k \ge a_k, \\ \textit{accept } H_1 & \textit{if} & p_k \le b_k, \\ \textit{continue} & \textit{if} & b_k < p_k < a_k, \end{array}$$

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

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

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. If condition (7) 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 (7) does not entail any loss of generality.

For any fixed $\gamma_2 \in \Gamma_2$, Theorem 2 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. The existence of p.b.p.o. solutions can be established using sequential compactness arguments[3] as in [TH]. However, unlike the result of [TH], optimal sensor decision policies can be found not in the class of SPRTs, but rather in the class of GSPRTs, which as shown above in Theorem 2 have monotonically convergent thresholds.

Remark 3. 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) even when the number of sensors is N ($N > 2$). To see this, we fix the decision policies of all the sensors except sensor S_i. Then we use a DP argument similar to the one used in establishing Theorem 2 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, with modified definitions for G_j and K_j and with $P_j(\tau_2 \le k)$ replaced by $\prod_{l=2}^{N} P_j(\tau_l \le k)$, $j = 0, 1$.

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 [TH], cannot be used here; it is well known in sequential analysis that such approximations for time-varying threshold tests are very difficult to obtain [S].

An alternative to the above technique for finding optimal thresholds is the fol-

[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 [TH]. The policies to which these subsequences converge define a p.b.p.o. solution.

lowing recursive algorithm, that is motivated by the sequential compactness argument of the previous section (see footnote 3):

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.
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.
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 (P1). One of these p.b.p.o. solutions is a g.o. solution to (P1), if a g.o. solution exists.

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 (P1) 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 below.

An SPRT policy at sensor S_i has the following form:

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

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 [W1] in the following way:

$$A_i = \frac{q(1-a^i)}{(1-q)a^i}, \qquad B_i = \frac{q(1-b^i)}{(1-q)b^i}. \tag{8}$$

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 [W1] give us the following approximate expressions for α_i and β_i:

$$\alpha_i \approx \frac{1-A_i}{B_i-A_i}, \qquad \beta_i \approx \frac{A_iB_i-A_i}{B_i-A_i}. \tag{9}$$

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 [W2]. With γ_i as defined in Theorem 3.1 of [W2], we have the following approximations:

$$\alpha_i \approx \frac{\gamma_i}{B_i}, \qquad \beta_i \approx \gamma_i A_i. \tag{10}$$

Using (8) and (9) or (10), 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. If f_{τ_i} denotes the probability mass function of τ_i, then

$$E\{\max(\tau_1, \tau_2)\} = \sum_{n=0}^{\infty}\left(\sum_{m=0}^{n} nf_{\tau_1}(n)f_{\tau_2}(m) + \sum_{m=n+1}^{\infty} mf_{\tau_1}(n)f_{\tau_2}(m)\right). \tag{11}$$

It is straightforward to show that the above expectation can also be written in terms of characteristic functions as given below:

$$\begin{aligned} E\{\max(\tau_1, \tau_2)\} = E\{\tau_2\} &+ \frac{1}{4\pi}\int_0^{\pi} \operatorname{cosec}^2\left(\frac{\omega}{2}\right)\operatorname{Re}(\varphi_2(\omega)(1 - \varphi_1(\omega)))\, d\omega \\ &+ \frac{E\{\tau_1\}}{2\pi}\int_0^{\pi}\left(\operatorname{Re}(\varphi_2(\omega)) + \operatorname{Im}(\varphi_1(\omega))\cot\left(\frac{\omega}{2}\right)\right) d\omega, \end{aligned} \tag{12}$$

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

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

The conditional characteristic functions under H_1 and H_0 can be approximated using Wald's fundamental identity (for details, see [W1]). If we define $t_{10}(\omega)$, $t_{20}(\omega)$, $t_{11}(\omega)$, and $t_{21}(\omega)$ by

$$\begin{aligned} t_{10}(\omega) &= 0.5(1 + \sqrt{1 - 8v\omega j}), \\ t_{20}(\omega) &= 0.5(1 - \sqrt{1 - 8v\omega j}), \\ t_{11}(\omega) &= 0.5(-1 + \sqrt{1 - 8v\omega j}), \\ t_{21}(\omega) &= 0.5(-1 - \sqrt{1 - 8v\omega j}), \end{aligned}$$

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

$$E\{\exp(-j\omega\tau_i)|H_0\} \approx \frac{A_i^{t_{20}(\omega)} - A_i^{t_{10}(\omega)} + B_i^{t_{20}(\omega)} - B_i^{t_{10}(\omega)}}{B_i^{t_{10}(\omega)}A_i^{t_{20}(\omega)} - A_i^{t_{10}(\omega)}B_i^{t_{20}(\omega)}}. \tag{13}$$

Similar expressions hold for $E\{\exp(-j\omega\tau_i)|H_1\}$, $i = 1, 2$, with 0 replaced by 1 in (13).

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 (P1).

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

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

In Table 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 Fig. 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 lists the expected cost of the

Table 1. Optimization results for the best SPRT policies for the case where $c = 0.01$, $v = 1.0$, and $k_e = 4.0$.

	SPRT thresholds				Expected cost	
q	$1 - a^1$	$1 - a^2$	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}

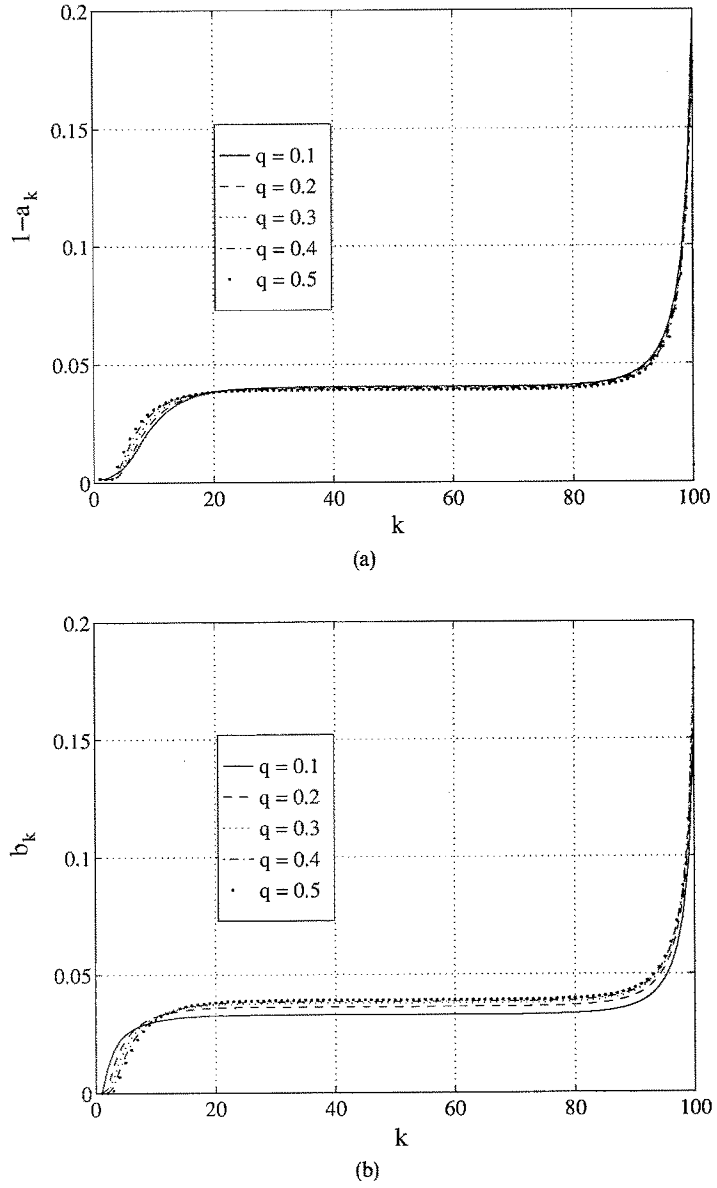

Fig. 1. Optimal GSPRT thresholds for the case where $c = 0.01$, $v = 1.0$, and $k_e = 4.0$.

GSPRT policies obtained from the DP recursions. The expected cost was also obtained by direct Monte-Carlo simulations of the optimal GSPRT policies. 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 tradeoff might be made between the

Table 2. Comparison of the performance of GSPRT and SPRT policies for the case where $c = 0.01$, $v = 1.0$, and $k_e = 4.0$.

	Expected cost			
	SPRT policy		GSPRT policy	
q	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}

simplicity of the SPRT policy and the performance gain obtainable with the GSPRT policy.

6. Conclusion

In this paper we formulated an extension of the Wald problem to the decentralized case. We used a dynamic programming 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.

The analysis contained in this paper can routinely be extended to the general case in which there are N ($N > 2$) sensors, without any conceptual difficulties (see Remark 3). Also, the case where the stopping time penalty is of the form $c_1\tau_1 + c_2\tau_2 + c\max(\tau_1, \tau_2)$ is 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.

References

[AV] M. M. Al-Ibrahim and P. K. Varshney, A simple multi-sensor sequential detection procedure, *Proceedings of the 27th IEEE Conference on Decision and Control*, Austin, TX, 1988, pp. 2479–2483.

[B] D. Bertsekas, *Dynamic Programming*, Prentice-Hall, Englewood Cliffs, NJ, 1987.

[HR] H. R. Hashemi and I. B. Rhodes, Decentralized sequential detection, *IEEE Trans. Inform. Theory*, **35** (1989), 509–520.

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

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

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

[T1] A. G. Tartakovskii. Sequential testing of many simple hypotheses with independent observations. *Problems Inform. Transmission*, **24** (1988), 299–309.

[TH] D. Teneketzis and Y. C. Ho, The decentralized Wald problem, *Inform. and Comput.*, **73** (1987), 23–44.

[T2] J. N. Tsitsiklis, On threshold rules in decentralized detection, *Proceedings of the 25th IEEE Conference on Decision and Control*, Athens, 1986, pp. 232–236.

[T3] J. N. Tsitsiklis, Decentralized detection, in *Advances in Statistical Signal Processing*, vol. 2 (H. V. Poor and J. B. Thomas, eds.), JAI Press, Greenwich, CT, 1993.

[VBP1] V. V. Veeravalli, T. Başar, and H. V. Poor, The decentralized Wald problem with a nonlinear penalty on stopping times, *Proceedings of the 25th Annual Conference on Information Sciences and Systems*, Johns Hopkins University, Baltimore, MD, 1991, pp. 277–282.

[VBP2] V. V. Veeravalli, T. Başar, and H. V. Poor, Decentralized sequential detection with a fusion center performing the sequential test, *IEEE Trans. Inform. Theory*, **39** (1993), 433–442.

[W1] A. Wald, *Sequential Analysis*, Wiley, New York, 1947.

[W2] M. Woodroofe, *Nonlinear Renewal Theory in Sequential Analysis*, SIAM, Philadelphia, PA, 1982.