PERGAMON Journal of the Franklin Institute 336 (1999) 301–322

Journal *of* The Franklin Institute

Sequential decision fusion: theory and applications[1]

Venugopal V. Veeravalli*

School of Electrical Engineering, Cornell University, Ithaca, NY 14853, USA

Abstract

A decision fusion problem is considered in which each one of a set of sensors receives a sequence of observations about the state of the system. The observations are quantized at each time step and sent to the fusion center where a *binary* decision is to be made at a stopping time. Applications of this problem setting are discussed and techniques for finding optimal solutions are presented. © 1998 The Franklin Institute. Published by Elsevier Science Ltd.

Keywords: Sequential analysis; Decentralized decision-making; Season fusion

1. Introduction

Statistical decision-making problems can be classified according to whether the information is *centralized* or *decentralized*. In a centralized setting, all information is available at a central processor where decisions are made according to some optimality criterion. Hence, even if more than one unit is involved in collecting information, the units can be considered together as a single decision maker. In the decentralized setting, the units which collect information (*sensors*) are also involved in decision making on some local level. However, all units are generally working towards a common objective (they are playing a cooperative game, according to a game theoretic definition [1, 2]).

Various sensor configurations are possible for decentralized decision-making [3]—the most common is the *decision fusion* configuration where a summary message

[1]Appeared in part in the *Proceedings of the Workshop on Foundations Information/Decision Fusion with Engineering Applications*, August 1996, Washington DC. This research was supported by the Office of Naval Research under Grant N00014-97-0823, through Cornell University.

*Tel.: 607-255-4847; fax: 607-255-9072; e-mail: venu@ee.cornell.edu

0016-0032/98/$ – See front matter
PII: S0016-0032(98)00023-4

(belonging to a finite alphabet) is sent from each sensor to a central processor, called the *fusion center*, where a final decision is made. The optimization to find the best decision policy needs to be done jointly over sensor and fusion decision rules.

Decision-making problems may also be classified according to whether the decisions are *static* or *dynamic*. In static frameworks, each decision maker makes only one decision. Static centralized decision-making problems are very well understood today and a well-developed theory exists for their analysis [4, 5]. Static decision fusion problems have also been researched over the past 15 years in various contexts, and are useful in many applications, including radar detection, surveillance systems, and human decision-making organizations [3, 6]. However, there is a large body of problems where the static setting is not useful. These problems call for more dynamic scenarios, where information is updated at certain intervals and new decisions are constantly being made in order to optimize system performance over time. Such dynamic problems are also known as sequential decision problems, a class of problems first introduced by Wald [7]. Centralized sequential detection and estimation problems have been the focus of the mature field of sequential analysis [8, 9]. Research on *decentralized sequential decision making* problems, on the other hand, is still in its infancy.

Our goal in this paper is to provide an introduction to a special class decentralized sequential decision making problems, namely *sequential decision fusion* problems.[2] Here each one of a set of sensors receives a sequence of observations. The observations are quantized at each time step and sent to the fusion center where a decision is to be made at a stopping time. The optimization of the decision rules needs to be done over time, as well as over the sensors and fusion center. For simplicity, we restrict our attention to problems where the final decision space is binary.

We consider a system with N sensors $S^{(1)}, \ldots, S^{(N)}$ as shown in Fig. 1. At time $k \in \{1, 2, \ldots\}$, an observation $X_k^{(l)}$ is made at sensor $S^{(l)}$. Further, based on the information available at $S^{(l)}$ at time k, a message $U_k^{(l)}$ is formed and sent to the fusion center. We include the possibility of feedback from the fusion center. In particular, at time k the fusion center could possibly broadcast, to each sensor, all the sensor messages it received at time $k-1$. Thus, at time k, each sensor could, in the most general case, have access to all its observations up to time k and all the messages of all the other sensors up to time $k-1$. The fusion center has access to all the sensor messages up to time k.

Various information structures are possible for sequential decision fusion problems of the form shown in Fig. 1.

Case A: *System with neither feedback from the fusion center nor local memory*. Here the sensor message $U_k^{(l)}$ is constrained to depend only on $X_k^{(l)}$:

$$U_k^{(l)} = \phi_k^{(l)}(X_k^{(l)})$$

[2] Other structures are possible for decentralized sequential detection [10, 11].

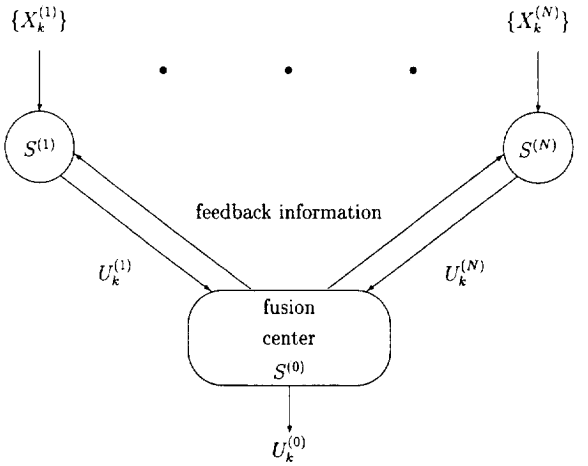

Fig. 1. General setting for sequential decision fusion.

Case B: *System with no feedback but full local memory*:

$$U_k^{(l)} = \phi_k^{(l)}(X_1^{(l)}, \ldots, X_k^{(l)})$$

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

$$U_k^{(l)} = \phi_k^{(l)}(X_k^{(l)}; U_1^{(l)}, \ldots, U_{k-1}^{(l)})$$

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[3]:

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

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 Case E, the past (one-step delayed) information at the fusion center and at each of the sensors is the same, and is nested at successive stages. This, together with the fact that the final decision depends only on the sensor messages (and through them on the observations), implies that the information structure for this case is *quasiclassical*. It is known that stochastic control or team problems with such an information structure are tractable via dynamic programming arguments [12–15]. As we explain in Sections (2.2) and (2.3), the same is true for sequential decision fusion problems.

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

There are two types of binary decision-making problems that have the above fusion structure: decentralized *sequential hypothesis testing* and decentralized *change detection*. We consider both of these problems and clearly delineate their common features and differences.

In decentralized sequential hypothesis testing, all of the observations at the sensors come from one of two hypotheses H_0 or H_1. The statistics of the observations are different under the two hypotheses, and the problem is to determine the true hypothesis with a certain accuracy as quickly as possible. Such a problem arises in many applications including clinical trials, radar detection, and channel selection in cellular communication systems.

In the channel selection application, the goal might be to determine whether a particular communications channel is "good" or "bad" based on interference measurements. Such a decision can be useful, for example, in dynamic channel allocation [16], or for borrowing voice channels in Cellular Digital Packet Data (CDPD) [17] applications. Depending on the system, these channel measurements can be taken both at the base station and at mobiles in the cell [18]. The base station plays the role of fusion center in these systems, while the mobiles act as sensors. The base station instructs the mobiles to take interference measurements at some starting time, and mobiles send back quantized versions[4] of these measurements (which can be regarded as local decisions) to the base station. Then, based on the measurements, the base station makes a decision about the channel at some stopping time.

In decentralized change detection, the statistics of the sensor observations change abruptly at some time, and the goal is to detect the change as soon as possible while containing the false alarm probability. Applications include machine monitoring, surveillance systems, and failure detection in large networks.

As an example consider the problem of link failure detection in a communication network. Information about the failure is available through measurements taken at several nodes (sensors) in the network [19, 20]. A network manager (fusion center) must detect the failure as soon as possible based on these measurements. The sensors, being geographically separated from the fusion center, are constrained to send *finite alphabet* messages.

The remainder of this paper is organized as follows. In Section 2, we study decentralized sequential hypothesis testing. In Section 3, we discuss the decentralized change detection problem. In Section 4, we provide examples of both types of sequential decision fusion problems. In Section 5, we give our conclusions and discuss directions for future research.

2. Decentralized sequential hypothesis testing

For this problem, the statistics of sensor observations are different under the two hypotheses. A decision about the hypothesis is made at the fusion center at a stopping

[4]Given the limited bandwidth available for control information, the measurements will typically be quantized to a small numbe. of bits.

time. The goal is to design the fusion stopping rule, fusion decision rule and the sensor message functions so as to strike an optimum tradeoff between decision accuracy and delay.

2.1. Problem formulation

We denote the hypothesis by a binary random variable H, and assume that H takes on values H_0 and H_1, with prior probabilities ν and $1-\nu$, respectively. Further, the observations at each sensor $S^{(l)}$ are independent with common p.d.f. $f_j^{(l)}$ if H_j is the true hypothesis, $j=0,1$. We also assume that the observations are independent from sensor to sensor. The sensor message $U_k^{(l)}$ takes on a value $d^{(l)}$ in the finite set $\{0, \ldots, D^{(l)}-1\}$.

The fusion center policy ψ consists of selecting a stopping time τ (that is measurable with respect to the sequence of sigma algebras generated by $\{I_k\}$), and a final decision $\delta \in \{0,1\}$ based on the information up to time τ. Decision errors are assumed to be penalized through a cost function $C(\delta, H)$. Typically, the cost function C is of the form: $C(0,H_0)=C(1,H_1)=0$, and $C(0,H_1)=L_0$, $C(1,H_0)=L_1$, where L_0 and L_1 are positive.

To each choice of fusion center policy and sensor quantizer functions, there corresponds an expected *decision cost* $E[C(\delta,H)]$ and an expected *delay* $E[\tau]$. Smaller decision cost can be obtained at the expense of larger delay and hence a tradeoff must be made. The optimum tradeoff problem may be posed in two ways:

Variational formulation: Minimize the expected delay $E[\tau]$ over all admissible choices of fusion center policy ψ and sensor quantizers $\{\phi_k^{(l)}\}$, subject to $E[C(\delta,H)] \leqslant \alpha$, where α is a control parameter.

Bayesian formulation (*P*1): Minimize the total Bayes cost $E[c\tau + C(\delta,H)]$ over all admissible choices of ψ and $\{\phi_k^{(l)}\}$, where the constant $c>0$ may be interpreted as the cost of each time step taken for decision making.

It is not difficult to establish that the solution to the variational problem for a given α is obtained as a solution to the Bayesian problem for some value c_α [21]. Hence, the optimum tradeoff curve for both formulations may be obtained by solving the Bayesian problem for various values of c.

2.2. Information structures and tractability

The tractability of problem (P1) depends crucially on the information structure in the system. For the information structures of cases B and D, the problem does not allow for any simplification even under independence assumptions on the sensor observations [22, 23]. Optimal decision rules in these cases can only be found by exhaustive search, which is in general intractable.

For the information structures of cases A and C, it is easily shown that optimal sensor functions are likelihood ratio quantizers [22, 24] if the observations are independent from sensor to sensor. 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. Further simplification would be possible in the i.i.d. observation case, if stationary sensor functions were optimal. However, such a result has not yet been established. The difficulty has its roots in the nonclassical[5] nature of the information in the system in these cases, which precludes the use of powerful dynamic optimization techniques as discussed in the following.

2.3. Dynamic programming and optimal solutions

Sequential decision fusion problems can be viewed as dynamic stochastic optimization problems, where an expected cost is to be minimized over the trajectory of a dynamical system. The goal is to find a set of (decision) functions, sometimes called a *policy*, that results in minimum cost. A powerful technique for obtaining optimal solutions for such problems is dynamic programming [12, 13]. This technique is applicable when the underlying system is a discrete-time dynamical system, the random variables have a conditional independence structure, and the cost to be minimized is additive over time. The idea behind this approach is to first find the optimal decision functions for the last time step in the trajectory and then work backwards, at each step minimizing the cost-to-go to the end of the trajectory. The added complexity in the case of sequential decision fusion problems is that at each time step, the optimization has to be done over all the decision makers in the system.

The system under consideration for decentralized sequential hypothesis testing can be considered to be a discrete-time dynamical system with state $\mathscr{S}_k$ that evolves as follows. $\mathscr{S}_0$ is the true hypothesis H. At time $k > 0$, the state $\mathscr{S}_k$ is either H or a termination state denoted by $\mathscr{T}$, i.e.,

$$\mathscr{S}_{k+1} = \begin{cases} \mathscr{S}_k & \text{if fusion center decision } U_k^{(0)} \text{ is continue} \\ \mathscr{T} & \text{if } U_k^{(0)} \text{ is terminate and decide } \delta = 0 \text{ or } 1 \end{cases}$$

We do not observe the state directly, but get noisy observations at each of the sensors. The observations $\{X_k^{(l)}\}$ at each sensor $S^{(l)}$ are independent with common pdf $f_j^{(l)}$ if $H = H_j$, $j = 0, 1$.

The cost for the Bayesian problem (P1) of Section (2.1) can be written as an additive function over time by defining the cost of time step k to be

$$g_k(\mathscr{S}_k, U_k^{(0)}) = \begin{cases} c & \text{if } U_k^{(0)} = \text{continue} \\ L_0 \mathbb{1}_{\{H = H_1\}} & \text{if } U_k^{(0)} = \text{terminate and decide 0} \\ L_1 \mathbb{1}_{\{H = H_0\}} & \text{if } U_k^{(0)} = \text{terminate and decide 1} \end{cases}$$

where $\mathbb{1}_{\{\cdot\}}$ is the indicator function. The total Bayes' risk (cost) is then given by

$$E[c\tau + C(\delta; H)] = \sum_k g_k(\mathscr{S}_k, U_k^{(0)})$$

[5] We 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.

and this cost needs to be minimized over all choices of sensor decision functions $\{\phi_k^{(l)}\}_{k,l}$ and fusion center policy ψ. Note that the cost is an explicit function of the fusion center decisions and only indirectly a function of the sensor decisions. The dynamic programming argument for this problem could proceed as follows. The information available for decision making at the fusion center at time k is

$$I_k = \{U_{[1,k]}^{(1)}, \ldots, U_{[1,k]}^{(N)}\} \tag{1}$$

Let the minimum cost-to-go at time k be denoted by $J_k(I_k)$. Then we can write the backward recursion:

$$J_{k-1}(I_{k-1}) = \min_{U_k^{(0)}, U_k^{(1)}, \ldots, U_k^{(N)}} E_{X_k^{(1)}, \ldots, X_k^{(N)} | I_{k-1}} \{J_k(I_{k-1}; U_k^{(0)}, \ldots, U_k^{(N)}) + g_k(\mathscr{S}_k, U_k^{(0)})\}$$

The above minimization would result in a sensor decision rule at sensor $S^{(l)}$ and at time k that would automatically depend on $X_k^{(l)}$ and I_{k-1}. But such a decision rule is not implementable unless the decision rule at sensor $S^{(l)}$ is forced to depend on exactly $X_k^{(l)}$ and I_{k-1}. This happens only when we have the *quasiclassical* information structure discussed in Case E of Section 3. Within this structure, optimal solutions can be obtained using dynamic programming arguments as we show in Theorem 1 below. Now, although we present the optimal solution under Case E only, the form of the optimal solution for this case may be useful in deriving good *ad hoc* solutions for the other information structures.

2.4. Optimal solution to Bayesian formulation for Case E

For clarity of presentation, we now describe the optimum detector structure for Case E for a simple two sensor problem. This solution contains all of the essential ingredients of the general solution and the generalization should be clear. The details of the general solution are given in [23].

Suppose there are two sensors ($N = 2$) each sending one bit of information ($D = 2$) at each time step to the fusion center. As in Eq. (1), we define I_k as

$$I_k := \{U_{[1,k]}^{(1)}, U_{[1,k]}^{(2)}\} \tag{2}$$

Note that I_k is the information available for decision making at the fusion center at time k, and I_{k-1} is the past decision information available for sensor decision making at time k.

The solution to (P1) is obtained using dynamic programming (DP) [13] arguments given in the proof of Theorem 1. A sufficient statistic for the DP recursions is shown to be the *a posteriori* probability of the hypothesis H_0, given I_k, i.e.,

$$p_k = P(H = H_0 \mid I_k) \tag{3}$$

This one-dimensional sufficient statistic is all that the sensors and fusion center need to store at any given time k, and it can be easily updated using the recursion given below in Eq. (13). The complete solution to (P1) is stated below in the following theorem.

Theorem 1. (i) *At each time k, the sensors use likelihood ratio tests with thresholds which are functions of* p_{k-1}, *i.e.*,

$$U_k^{(l)} = \begin{cases} 1 & \text{if } L^{(l)}(X_k^{(l)}) > \lambda^{(l)}(p_{k-1}) \\ 0 & \text{otherwise} \end{cases} \tag{4}$$

where $L^{(l)}(X_k^{(l)}) = f_1^{(l)}(X_k^{(l)})/f_0^{(l)}(X_k^{(l)})$. *Note that the optimum sensor policy is stationary in time.*

(ii) *The optimum fusion center policy has the form:*

$$\begin{aligned} U_k^{(0)} &= \text{terminate and decide } 0 \quad \text{if } p_k \geqslant a \\ U_k^{(0)} &= \text{terminate and decide } 1 \quad \text{if } p_k \leqslant b \\ U_k^{(0)} &= \text{continue taking observations} \quad \text{if } b < p_k < a, \end{aligned} \tag{5}$$

The optimal sensor and fusion center thresholds are derived as follows. Let $J(p)$ *be the unique solution to the fixed point equation*

$$J(p) = \min\{L_1 p, L_0(1-p), c + \min_{\lambda^{(1)}, \lambda^{(2)}} W_J(\lambda^{(1)}, \lambda^{(2)}; p)\} \tag{6}$$

where

$$W_J(\lambda^{(1)}, \lambda^{(2)}; p) = \sum_{d^{(1)}, d^{(2)}} J\left(\frac{g(d^{(1)}, d^{(2)}; \lambda^{(1)}, \lambda^{(2)}; p)}{h(d^{(1)}, d^{(2)}; \lambda^{(1)}, \lambda^{(2)}; p)}\right) h(d^{(1)}, d^{(2)}; \lambda^{(1)}, \lambda^{(2)}; p) \tag{7}$$

with

$$g = p \prod_{l=1}^{2} [P_0\{L^{(l)}(X^{(l)}) > \lambda^{(l)}\}]^{d^{(l)}} [P_0\{L^{(l)}(X^{(l)}) \leqslant \lambda^{(l)}\}]^{1-d^{(l)}}, \tag{8}$$

and

$$h = g + (1-p) \prod_{l=1}^{2} [P_1\{L^{(l)}(X^{(l)}) > \lambda^{(l)}\}]^{d^{(l)}} [P_1\{L^{(l)}(X^{(l)}) \leqslant \lambda^{(l)}\}]^{1-d^{(l)}}. \tag{9}$$

Then the optimal sensor thresholds are given by

$$\{\lambda_{\text{opt}}^{(1)}(p), \lambda_{\text{opt}}^{(2)}(p)\} = \arg \min_{\lambda^{(1)}, \lambda^{(2)}} W_J(\lambda^{(1)}, \lambda^{(2)}; p), \tag{10}$$

and b_{opt} and a_{opt} *are obtained from the equations*

$$L_0(1 - b_{\text{opt}}) = c + W_J(\lambda_{\text{opt}}^{(1)}(b_{\text{opt}}), \lambda_{\text{opt}}^{(2)}(b_{\text{opt}}); b_{\text{opt}}) \tag{11}$$

and

$$L_1 a_{\text{opt}} = c + W_J(\lambda_{\text{opt}}^{(1)}(a_{\text{opt}}), \lambda_{\text{opt}}^{(2)}(a_{\text{opt}}); a_{\text{opt}}) \tag{12}$$

Finally, the recursion for p_k is given by

$$p_{k+1} = \frac{g(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \lambda_{\text{opt}}^{(1)}(p_k), \lambda_{\text{opt}}^{(2)}(p_k); p_k)}{h(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \lambda_{\text{opt}}^{(1)}(p_k), \lambda_{\text{opt}}^{(2)}(p_k); p_k)}, \quad p_0 = v \tag{13}$$

where v is the a priori probability of hypothesis H_0.

Proof. See the Appendix.

The fixed point of Eq. (6) can be obtained by successive approximation using the starting function $\eta(p) = \min\{L_1 p, L_0(1-p)\}$, as the following result shows.

Theorem 2. *Let $S \subset C[0,1]$ be the set of all non-negative concave functions on $[0,1]$ that are upper bounded (in sup norm) by the function*

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

Now define the mapping $\mathscr{T} : S \mapsto S$ by

$$\mathscr{T}G(p) = \min\left\{L_1 p, L_0(1-p), c + \min_{\lambda^{(1)}, \lambda^{(2)}} W_G(\lambda^{(1)}, \lambda^{(2)}; p)\right\}, \quad \text{for } G \in S$$

Then $J(p)$ is the unique fixed point of the mapping $\mathscr{T}$. Furthermore, $\mathscr{T}^n\eta$ converges monotonically to J as $n \to \infty$.

Proof. See the Appendix.

An illustrative example of the solution procedure is given in Section 4.

3. Decentralized change detection

The problem of detecting an abrupt change in a system based on stochastic observations of the system has a variety of applications including failure detection, quality control engineering, and channel monitoring for wireless communication systems. The *centralized* version of this problem—where all the information about the change is available at a single location—is well-understood and has been solved under a variety of criteria since the seminal work by Page [25]. (See, for example, [21, 26, 27].) However, there are situations where the information available for decision-making is decentralized.

A decentralized formulation of the change detection problem was considered in [28, 29] with each of the sensors performing a CUSUM type test. *Ad hoc* schemes for fault detection with multiple observers were considered in [20, 30]. A formulation of the decentralized change detection problem with a fusion center making the final decision about the change was given by us in [31]. Here the statistics of the sensor

observations change abruptly at some time, and a decision about the change is made at the fusion center at a stopping time. The goal is to design the fusion stopping rule and the sensor message functions so as to minimize the time taken to detect the change while containing the false alarm probability at an acceptable level.

Arguments similar to those given in Sections 2.2 and 2.3 show that only the information structure of Case E leads to a problem that is tractable via dynamic programming.

3.1. Problem formulation

The observations at each sensor $S^{(l)}$ are independent, have a common p.d.f. $f_0^{(l)}$ before the disruption, and common p.d.f. $f_1^{(l)}$ from the time of disruption. Further the observations are independent from sensor to sensor. Following the Bayesian analysis given in [21], we assume that the prior distribution of the change time Γ is geometric and is given by

$$P(\Gamma = 0) = \nu \quad \text{and} \quad P(\Gamma = i | \Gamma > 0) = \rho(1-\rho)^{(i-1)}$$

Based on the information received from the sensors, the fusion center must make a decision about the change.

The fusion center policy ψ consists of selecting a stopping time τ at which it is decided that the disruption has occurred. To each choice of fusion center policy and sensor quantizers, there corresponds a probability of *false alarm* $P(\tau < \Gamma)$ and an expected *delay* $E[\tau - \Gamma | \tau \geqslant \Gamma]$ of detecting the disruption under the condition that the alarm signal was correctly given. Here again, the optimum tradeoff problem can be posed in two ways [21].

Variational formulation: Minimize the expected delay $E[\tau - \Gamma | \tau \geqslant \Gamma]$ over all admissible choices of fusion center policy ψ and sensor quantizers $\{\phi_k^{(l)}\}$, subject to $P(\tau < \Gamma) \leqslant \alpha$, where α is a control parameter.

Bayesian formulation (P2): Minimize the total Bayes cost (risk)

$$\begin{aligned} R(c) &= P(\tau < \Gamma) + cE[\tau - \Gamma | \tau \geqslant \Gamma] P(\tau \geqslant \Gamma) \\ &= E[\mathbf{1}_{\{\tau < \Gamma\}} + c(\tau - \Gamma)\mathbf{1}_{\tau \geqslant \Gamma\}}] \\ &= P(\Gamma > \tau) + cE\left[\sum_{k=1}^{\tau-1} P(\Gamma \leqslant k)\right] \end{aligned}$$

over all admissible choices of ψ and $\{\phi_k^{(l)}\}$, where the constant $c > 0$ may be interpreted as the cost of each unit of delay.

3.2. Optimal solution to Bayesian formulation for Case E

The solution to (P2) is obtained using techniques very similar to those used for (P1). A sufficient statistic in this case is the *a posteriori* probability of the change having happened before time k given I_k, i.e.,

$$p_k = P(\Gamma \leqslant k | I_k) \tag{14}$$

where I_k is as defined in Eq. (2). This one-dimensional sufficient statistic is all that the sensors and fusion center need to store at any given time k, and it can be easily updated using the recursion given below in Eq. (23).

Again, to simplify the presentation, we present the solution to (P2) for the special case of two sensors ($N = 2$) each sending one bit of information ($D = 2$) at each time step to the fusion center. The general solution is given in [32].

Theorem 3. (i) *At each time k, the sensors use likelihood ratio tests with thresholds which are functions of p_{k-1}, i.e.,*

$$U_k^{(l)} = \begin{cases} 1 & \text{if } L^{(l)}(X_k^{(l)}) > \lambda^{(l)}(p_{k-1}) \\ 0 & \text{otherwise} \end{cases} \tag{15}$$

where $L^{(l)}(X_k^{(l)}) = f_1^{(l)}(X_k^{(l)})/f_0^{(l)}(X_k^{(l)})$. Note that the optimum sensor policy is stationary *in time.*

(ii) *The optimal fusion center policy has the form*

$$\begin{aligned} U_k^{(0)} &= \text{terminate and decide in favor of change} && \text{if } p_k \geqslant a \\ U_k^{(0)} &= \text{continue taking observations} && \text{if } p_k < a \end{aligned} \tag{16}$$

The optimal sensor and fusion center thresholds are derived as follows. Let $J(p)$ be the unique solution to the fixed point equation

$$J(p) = \min\{(1-p), cp + \min_{\lambda^{(1)}, \lambda^{(2)}} W_J(\lambda^{(1)}, \lambda^{(2)}; p)\} \tag{17}$$

where

$$W_J(\lambda^{(1)}, \lambda^{(2)}; p) = \sum_{d^{(1)}, d^{(2)}} J\left(\frac{g(d^{(1)}, d^{(2)}; \lambda^{(1)}, \lambda^{(2)}; p)}{h(d^{(1)}, d^{(2)}; \lambda^{(1)}, \lambda^{(2)}; p)}\right) h(d^{(1)}, d^{(2)}; \lambda^{(1)}, \lambda^{(2)}; p) \tag{18}$$

with

$$g = [p + (1-p)\rho] \prod_{l=1}^{2} [P_1\{L^{(l)}(X^{(l)}) > \lambda^{(l)}\}]^{d^{(l)}} [P_1\{L^{(l)}(X^{(l)}) \leqslant \lambda^{(l)}\}]^{1-d^{(l)}} \tag{19}$$

and

$$h = g + (1-\rho)(1-p) \prod_{l=1}^{2} [P_0\{L^{(l)}(X^{(l)}) > \lambda^{(l)}\}]^{d^{(l)}} [P_0\{L^{(l)}(X^{(l)}) \leqslant \lambda^{(l)}\}]^{1-d^{(l)}} \tag{20}$$

Then the optimal sensor thresholds are given by

$$\{\lambda_{\text{opt}}^{(1)}(p), \lambda_{\text{opt}}^{(2)}(p)\} = \arg\min_{\lambda^{(1)}, \lambda^{(2)}} W_J(\lambda^{(1)}, \lambda^{(2)}; p) \tag{21}$$

and a_{opt} *satisfies the equation*

$$(1 - a_{\text{opt}}) = ca_{\text{opt}} + W_J(\lambda^{(1)}_{\text{opt}}(a_{\text{opt}}), \lambda^{(2)}_{\text{opt}}(a_{\text{opt}}); a_{\text{opt}}) \tag{22}$$

Finally, the recursion for p_k *is given by*

$$p_{k+1} = \frac{g(U^{(1)}_{k+1}, U^{(2)}_{k+1}; \lambda^{(1)}_{\text{opt}}(p_k), \lambda^{(2)}_{\text{opt}}(p_k); p_k)}{h(U^{(1)}_{k+1}, U^{(2)}_{k+1}; \lambda^{(1)}_{\text{opt}}(p_k), \lambda^{(2)}_{\text{opt}}(p_k); p_k)}, \quad p_0 = v \tag{23}$$

where v *is the* a priori *probability of the event* $\{\Gamma = 0\}$.

Proof. See the Appendix.

The complete solution to (P2) is obtained by solving for the fixed point $J(p)$ of Eq. (17) by successive approximation using the starting function $\eta(p) = 1 - p$, as stated in the following theorem whose proof is very similar to that of Theorem 2.

Theorem 4. *Let* $S \subset C[0, 1]$ *be the set of all non-negative concave functions on* $[0, 1]$ *that are bounded (in sup norm) by the function* $\eta(p) = 1 - p$. *We define the mapping* $\mathcal{T} : S \mapsto S$ *by*

$$\mathcal{T}G(p) = \min\left\{(1 - p), cp + \min_{\lambda^{(1)}, \lambda^{(2)}} W_G(\lambda^{(1)}, \lambda^{(2)}; p)\right\}, \quad \text{for } G \in S.$$

Then $J(p)$ *is the unique fixed point of the mapping* $\mathcal{T}$. *Furthermore,* $\mathcal{T}^n\eta$ *converges monotonically to* J *as* $n \to \infty$.

4. Examples

4.1. Two sensor example of decentralized sequential hypothesis testing

Suppose there are two sensors each sending one bit of information to the fusion center, i.e., $N = D = 2$. Furthermore, we assume that the observations at both sensors 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.

Since the observations are Gaussian with the same variance, the likelihood ratio is monotonic and hence the sensor thresholds may be considered to be thresholds on the observations. That is,

$$U^{(l)}_k = \begin{cases} 1 & \text{if } X^{(l)}_k \geqslant \lambda^{(l)}(p_{k-1}) \\ 0 & \text{if } X^{(l)}_k < \lambda^{(l)}(p_{k-1}) \end{cases} \tag{24}$$

The functions given in Theorem 1 simplify to

$$g = p \prod_{l=1}^{2} [P_0\{X > \lambda^{(l)}\}]^{d^{(l)}} [P_0\{X \leqslant \lambda^{(l)}\}]^{1-d^{(l)}}$$

and

$$h = g + (1-p)\prod_{l=1}^{2}[P_1\{X > \lambda^{(l)}\}]^{d^{(l)}}[P_1\{X \leqslant \lambda^{(l)}\}]^{1-d^{(l)}}$$

Optimal sensor thresholds (as functions of p) are obtained by minimizing $W_J(\lambda^{(1)}, \lambda^{(2)}, p)$ over $(\lambda^{(1)}, \lambda^{(2)}) \in \mathbb{R}^2$ at each stage of the successive approximation procedure. Typical results are shown in Fig. 2. The optimal thresholds are identical (functions of p) at the two sensors. It should be noted that the sensor threshold is discontinuous and non-monotonic (the spikes at the end points are attributed to quantization and finite precision). This might be surprising at first since all of the functions involved in the optimization are smooth, 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 $\mathrm{sgn}(x)$.

Optimal fusion center thresholds a_{opt} and b_{opt} are obtained from $J(p)$, using Eqs (11) and (12).

4.2. *Two sensor example of decentralized change detection*

Suppose there are two sensors each sending one bit of information to the fusion center, i.e., $N = D = 2$. Furthermore, we assume that the observations at both sensors

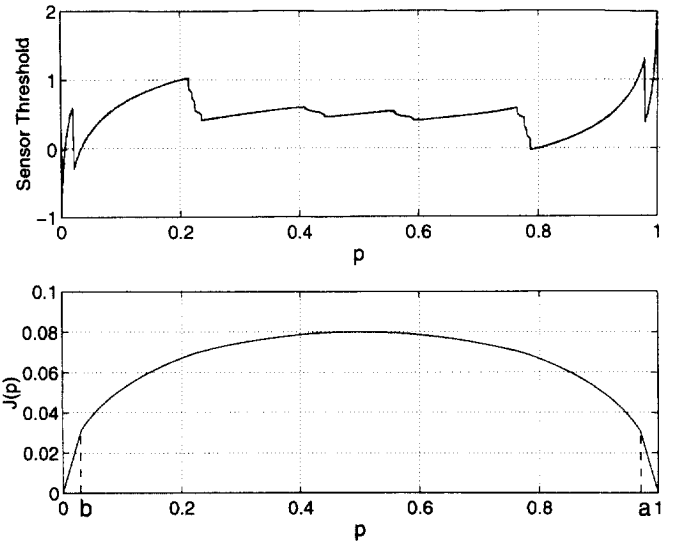

Fig. 2. Decentralized sequential detection with two sensors making binary decisions. Results for Gaussian sensor observations with mean 0 and variance 1 under H_0, and mean 1 and variance 1 under H_1. The parameter values are $c = 0.01$, and $L_0 = L_1 = 1.0$. Results were obtained using 250 iterations of the successive approximation procedure with 5000 points on the p-axis. The optimal thresholds are identical (functions of p) at the two sensors. The fusion center thresholds a and b are marked on the plot showing $J(p)$.

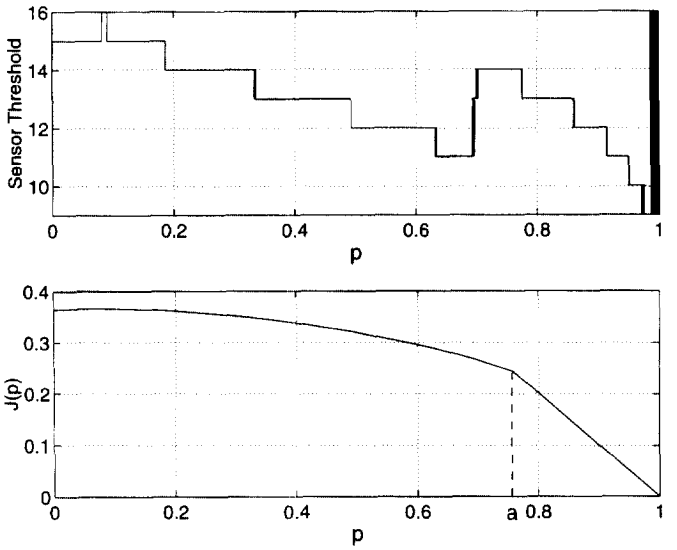

Fig. 3. Decentralized change detection with two sensors making binary decisions. Results for Poisson sensor observations with mean 10 before the disruption, and mean 14 after the disruption. The parameter values are $c = 0.1$, and $\rho = 0.05$. Results were obtained using 250 iterations of the successive approximation procedure with 5000 points on the p-axis. The optimal thresholds are identical (functions of p) at the two sensors. The fusion center threshold a is marked on the plot showing $J(p)$.

are i.i.d. Poisson random variable with mean μ_0 before the disruption and mean μ_1 after the disruption.

Here again the likelihood ratio is monotonic, and the sensor decision rules characterized by thresholds $\lambda^{(1)}$ and $\lambda^{(2)}$ on the observations. That is, $U_k^{(l)}$ has the same form as in Eqs. (24), but with

$$g = [p + (1-p)\rho] \prod_{l=1}^{2} [P_1\{X > \lambda^{(l)}\}]^{d^{(l)}} [P_1\{X \leqslant \lambda^{(l)}\}]^{1-d^{(l)}}$$

and

$$h = g + (1-\rho)(1-p) \prod_{l=1}^{2} [P_0\{X > \lambda^{(l)}\}]^{d^{(l)}} [P_0\{X \leqslant \lambda^{(l)}\}]^{1-d^{(l)}}.$$

Typical optimization results are shown in Fig. 3. Here again it is interesting to note that the optimal threshold is not a monotonic function. In particular, the optimal sensor threshold progressively favors the "change" hypothesis as p increases. But right before the value of p at which the fusion center would decide to stop, the threshold takes on a large value (i.e., favors the "no change" hypothesis). This behavior is again very similar to that observed in "bang–bang" optimal control.

The optimal fusion center threshold a_{opt} is obtained from $J(p)$, using Eq. (22).

5. Conclusions

The distinguishing feature of the decision fusion problems discussed in this paper is the dynamic nature of the information processing and decision-making. We demonstrated that these problems are tractable under a quasiclassical information structure (Case E).

It is of great interest to design decision-making strategies for other information structures, particularly those described in Cases A and B of Section 1. These other information structures may be more prevalent in practical applications. The insight we gain from deriving the optimal strategy for Case E and the form of the optimal solution might prove useful in deriving good suboptimal solutions for other information structures. Some preliminary work in this direction is presented in [33].

It is also of interest to develop a better understanding of the non-monotonic, discontinuous behavior of the optimal sensor threshold functions. It might be possible, for example, to find simple parametric models for these functions that would facilitate implementation of the optimal solution.

Finally, although the analysis of this paper allows us to derive optimal solutions to sequential decision fusion problems, it does not allow for a performance analysis of these optimal tests. The expected stopping times and expected decision costs of the optimal tests may be obtained via simulations. Nevertheless, it would be useful to derive analytic expressions or approximations for the expected stopping times and expected decision costs.

Appendix

Proof of Theorem 1. The general version of this proof is given in [23]. However, we have decided to include the proof given below for the special case of $N = D = 2$, because we believe it illustrates more lucidly the intricacies of the steps involved in establishing this result.

To begin we first restrict the fusion center stopping time τ to a finite horizon, say the interval $[0, T]$ (we will remove this restriction later). The minimization of the Bayes risk $E[c\tau + C(\delta, H)]$ over the finite horizon can be done recursively via dynamic programming.

Since the decision about the change is made at the fusion center, the (minimum) expected cost-to-go at time k is a function of the information available for decision making at the fusion center at time k, i.e., I_k. We denote the expected cost-to-go at time k by $\tilde{J}_k^T(I_k)$. It is easily seen that

$$\tilde{J}_T^T(I_T) = \min\{L_0 P(H = H_1|I_T), L_1 P(H = H_0|I_T)\} \tag{A.1}$$

and for $0 \leqslant k < T$,

$$\tilde{J}_k^T(I_k) = \min\Bigg\{L_0 P(H = H_1|I_k), L_1 P(H = H_0|I_k), c + \min_{\phi_{k+1}^{(1)}, \phi_{k+1}^{(2)}} E[\tilde{J}_{k+1}^T(I_{k+1})|I_k]\Bigg\} \tag{A.2}$$

with the understanding that I_0 is the empty set. The first (respectively, second) term in the outer minimum is the conditional expected cost of choosing H_0 (respectively, H_1), given I_k. The third term is the cost of continuing at time k. Note that the minimum expected cost for the finite horizon optimization problem is simply $\tilde{J}_0^T$.

We now establish that optimal sensor decision rules are likelihood ratio tests with thresholds that depend on I_{k-1}. From Eq. (26) it follows that optimal $\{\phi_k^{(1)}, \phi_k^{(2)}\}$ minimize

$$R_k = E[\tilde{J}_k^T(I_k) \mid I_{k-1}]$$

We fix $\phi_{k+1}^{(2)}$ and minimize R_k w.r.t. $\phi_{k+1}^{(1)}$ to find the structure of all person-by-person optimal solutions (including the globally optimal one) for $\phi_{k+1}^{(1)}$. Towards this end, we write

$$\begin{aligned} R_k &= E_{X_k^{(1)}, X_k^{(2)}, H \mid I_{k-1}} \tilde{J}_k^T(\phi_k^{(1)}(X_k^{(1)}; I_{k-1}), \phi_k^{(2)}(X_k^{(2)}; I_{k-1}), I_{k-1}) \\ &= E_{X_k^{(1)}, H \mid I_{k-1}} E_{X_k^{(2)} \mid H, I_{k-1}} \tilde{J}_k^T(\phi_k^{(1)}(X_k^{(1)}; I_{k-1}), \phi_k^{(2)}(X_k^{(2)}; I_{k-1}), I_{k-1}) \end{aligned}$$

where the conditioning with respect to $X_k^{(1)}$ is dropped in the inner expectation on the last line due to the independence of the sensor observations given H. The result of the inner expectation is a function of $\phi_k^{(1)}(X_k^{(1)}; I_{k-1})$, I_{k-1}, and H, say $K(\phi_k^{(1)}(X_k^{(1)}; I_{k-1}), I_{k-1}; H)$. Therefore,

$$\begin{aligned} R_k &= E_{X_k^{(1)}, H \mid I_{k-1}} K(\phi_k^{(1)}(X_k^{(1)}; I_{k-1}), I_{k-1}; H) \\ &= E_{X_k^{(1)} \mid I_{k-1}} \sum_{j=0}^{1} K(\phi_k^{(1)}(X_k^{(1)}; I_{k-1}), I_{k-1}; H_j) P(H = H_j \mid I_{k-1}; X_k^{(1)}) \end{aligned}$$

Minimizing R_k with respect to $\phi_k^{(1)}$ is equivalent to minimizing the quantity inside the above expectation almost everywhere. Thus, the optimal decision rule at $S^{(1)}$ at time k has the form

$$\hat{\phi}_k^{(1)}(X_k^{(1)}; I_{k-1}) = \begin{cases} 1 \text{ if } \dfrac{P(H = H_1 \mid I_{k-1}; X_k^{(1)})}{P(H = H_0 \mid I_{k-1}; X_k^{(1)})} > \dfrac{K(1, I_{k-1}, H_0) - K(0, I_{k-1}, H_0)}{K(0, I_{k-1}, H_1) - K(1, I_{k-1}, H_1)} \\ 0 \text{ otherwise} \end{cases}$$

An application of Bayes' rule gives

$$\hat{\phi}_k^{(1)}(X_k^{(1)}; I_{k-1}) = \begin{cases} 1 & \text{if } L^{(1)}(X_k^{(l)}) > \lambda_k^{(1)}(I_{k-1}) \\ 0 & \text{otherwise} \end{cases}$$

where $L^{(1)}(X_k^{(1)}) = f_1^{(1)}(X_k^{(1)}) / f_0^{(1)}(X_k^{(1)})$. Similar results hold for $\hat{\phi}_k^{(2)}$.

Next, we establish that the optimal sensor threshold $\lambda_k^{(l)}$ is a function of I_{k-1} only through the sufficient statistic $p_{k-1} = P(H = H_0 | I_{k-1})$. To this end we first see that by

Bayes' rule

$$\begin{aligned} p_{k+1} &= P(H = H_0 | I_{k+1}) \\ &= P(H = H_0 | U_{k+1}^{(1)}, U_{k+1}^{(2)}, I_k) \\ &= \frac{p_k f(U_{k+1}^{(1)}, U_{k+1}^{(2)} | H_0, I_k)}{f(U_{k+1}^{(1)}, U_{k+1}^{(2)} | I_k)} \end{aligned}$$

where $f(\cdot)$ denotes the joint p.m.f. of $U_{k+1}^{(1)}$ and $U_{k+1}^{(2)}$. Now,

$$\begin{aligned} p_k f(U_{k+1}^{(1)}, U_{k+1}^{(2)} | H_0, I_k) &= p_k \prod_{l=1}^{2} [P_0\{L^{(l)}(X_{k+1}^{(l)}) > \lambda_{k+1}^{(l)}(I_k)\}]^{U_{k+1}^{(l)}} \\ &\quad \times [P_0\{L^{(l)}(X_{k+1}^{(l)}) \leqslant \lambda_{k+1}^{(l)}(I_k)\}]^{1-U_{k+1}^{(l)}} \\ &= g(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \lambda_{k+1}^{(1)}(I_k), \lambda_{k+1}^{(2)}(I_k); p_k) \end{aligned}$$

where g is as defined in Eq. (8). Similarly,

$$\begin{aligned} f(U_{k+1}^{(1)}, U_{k+1}^{(2)} | I_k) &= p_k f(U_{k+1}^{(1)}, U_{k+1}^{(2)} | H_0, I_k) + (1 - p_k) f(U_{k+1}^{(1)}, U_{k+1}^{(2)} | H_1, I_k) \\ &= h(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \lambda_{k+1}^{(1)}(I_k), \lambda_{k+1}^{(2)}(I_k); p_k) \end{aligned}$$

where h is as defined in Eq. (9).
Thus,

$$p_{k+1} = \frac{g(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \lambda_{k+1}^{(1)}(I_k), \lambda_{k+1}^{(2)}(I_k); p_k)}{h(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \lambda_{k+1}^{(1)}(I_k), \lambda_{k+1}^{(2)}(I_k); p_k)}, \quad p_0 = v. \tag{A.3}$$

Note that the RHS of Eq. (A.3) depends on I_k (and not just p_k), since $\lambda_{k+1}^{(l)}$ depends on I_k in general. However, we will establish in the following lemma that optimal $\lambda_{k+1}^{(l)}$ depend on I_k only through p_k. Thus, Eq. (A.3) is a useful recursion for p_k if optimal sensor quantizers are used.

Lemma 5. (i) *The function $\tilde{J}_k^T(I_k)$ of Eq. (25) can be written as a function of only p_k, say $J_k^T(p_k)$.*

(ii) *Optimal sensor thresholds at time $k+1$ depends on I_k only through 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, $0 \leqslant k \leqslant T-1$, suppose that $\tilde{J}_{k+1}^T(I_{k+1})$ is a function of only p_{k+1}, say $J_{k+1}^T(p_{k+1})$. Then

$$\tilde{J}_k^T(I_k) = \min\{(1 - p_k)L_0, p_k L_1, c + \min_{\lambda_{k+1}^{(1)}, \lambda_{k+1}^{(2)}} E[J_{k+1}^T(p_{k+1}) | I_k]\}.$$

But for a given choice of $\{\lambda_{k+1}^{(l)}, \lambda_{k+1}^{(2)}\}$,

$$E[J_{k+1}^T(p_{k+1}) \mid I_k] = \sum_{d^{(1)}, d^{(2)}} J_{K+1}^T\left(\frac{g(d^{(1)}, d^{(2)}, \lambda_{k+1}^{(1)}, \lambda_{k+1}^{(2)}; p)}{h(d^{(1)}, d^{(2)}, \lambda_{k+1}^{(1)}, \lambda_{k+1}^{(2)}; p)}\right) \times h(d^{(1)}, d^{(2)}, \lambda_{k+1}^{(1)}, \lambda_{k+1}^{(2)}; p)$$

$$= W_{J_{k+1}^T}(\lambda_{k+1}^{(1)}, \lambda_{k+1}^{(2)}; p)$$

where $W_{J_{k-1}^T}$ is as defined in Eq. (7).

Thus, optimal thresholds $\lambda_{k+1}^{(l)}$ depend on I_k only through p_k. Both parts of the lemma follow from this fact.

In order to solve the original Bayesian optimization problem (P1), the restriction that τ belongs to a finite interval is removed by letting $T \to \infty$ in the finite horizon DP equations. As a first step, we note that $J_k^{T+1}(p) \leqslant J_k^T(p)$, because the set of stopping times increases with T. This together with the fact that J_k^T is bounded below by zero implies that for each k, the following limit is well defined:

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

Denote this limit by $J_k^\infty(p)$.

Exploiting the i.i.d. nature of the sensor observations, a time-shift argument easily shows that $J_{k+1}^{T+1}(p)$ equals $J_k^T(p)$. This implies that J_k^∞ is independent of k. We hence denote the limit $J_k^\infty(p)$ by $J(p)$, which we will refer to as the infinite-horizon cost-to-go function. Taking the limit as $T \to \infty$ in the finite horizon DP equations, we get that $J(p)$ satisfies the Bellman equation

$$J(p) = \min\left\{L_1 p, L_0(1-p), c + \min_{\lambda^{(1)}, \lambda^{(2)}} W_J(\lambda^{(1)}, \lambda^{(2)}; p)\right\} \tag{A.4}$$

As described in the statement of Theorem 1, the solution to problem (P1) is obtained from the solution to the above Bellman equation. First, the uniqueness of the solution (established in Theorem 2) implies that optimal sensor thresholds are stationary in time and are as given in Eq. (10). Second, the RHS of Eq. (A.4) implies that the optimal fusion center policy is

$U_k^{(0)}$ = terminate and decide 0 if $L_1 p_k = J(p_k)$

$U_k^{(0)}$ = terminate and decide 1 if $L_0(1-p_k) = J(p_k)$

$U_k^{(0)}$ = continue taking observations otherwise

Straightforward concavity arguments (see [23]) then lead to the threshold rule given in Eq. (5) with thresholds satisfying Eqs (11) and (12). This completes the proof. □

Proof of Theorem 2. Let G be any fixed point of $\mathcal{T}$, and let $\{\hat{\lambda}^{(1)}(p), \hat{\lambda}^{(2)}(p)\}$ be such that

$$\{\hat{\lambda}^{(1)}(p), \hat{\lambda}^{(2)}(p)\} = \arg\min_{\lambda^{(1)}, \lambda^{(2)}} W_G(\lambda^{(1)}, \lambda^{(2)}; p)$$

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

$$p_{k+1} = \frac{g(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \hat{\lambda}^{(1)}(p_k), \hat{\lambda}^{(2)}(p_k); p_k)}{h(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \hat{\lambda}^{(1)}(p_k), \hat{\lambda}^{(2)}(p_k); p_k)}$$

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

$$\mathcal{N} = \min\{k \geqslant 0 \,|\, \eta(p_k) \leqslant c + W_G(\hat{\lambda}^{(1)}(p_k), \hat{\lambda}^{(2)}(p_k); p_k)\}$$

and

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

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

$$G(v) = c + E\{G(p_1)\}$$

$$G(p_1) = c + E\{G(p_2)|I_1\}$$

$$\vdots$$

$$G(p_{\mathcal{N}-1}) = c + E\{G(p_{\mathcal{N}})|I_{\mathcal{N}-1}\}$$

$$G(p_{\mathcal{N}}) = \eta(p_{\mathcal{N}})$$

Substituting backwards and taking expectations, we obtain

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

where the last inequality follows from the definition of J.

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

$$G(p) \leqslant \eta(p) = J_T^T(p), \forall T$$

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

$$\begin{aligned} J_m^T(p) &= \min\{\eta(p), c + \min_{\lambda^{(1)}, \lambda^{(2)}} W_{J_{m+1}^T}(\lambda^{(1)}, \lambda^{(2)}; p)\} \\ &\geqslant \min\{\eta(p), c + \min_{\lambda^{(1)}, \lambda^{(2)}} W_G(\lambda^{(1)}, \lambda^{(2)}; p)\} \\ &= G(p) \end{aligned}$$

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

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

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

$$J(p) \geqslant G(p)$$

Thus uniqueness is established.

Now, a straightforward induction argument shows that

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

Since $\mathscr{T}^{n}\eta(p) > 0$, the above inequality implies that $\mathscr{T}^{n}\eta$ converges monotonically to a fixed point of $\mathscr{T}$. The uniqueness of the fixed point implies that $\mathscr{T}^{n}\eta$ converges to J. □

Proof of Theorem 3. The steps in the proof are nearly identical to those in the proof of Theorem 1. The key differences are the form of the DP equations, the definition of the sufficient statistic and its recursion equation. The finite horizon DP equations (A.1) and (A.2) are replaced by the equations

$$\tilde{J}_T^T(I_T) = P(\Gamma > T|I_T), \tag{A.5}$$

and for $0 \leqslant k \leqslant T$,

$$\tilde{J}_k^T(I_k) = \min\left\{P(\Gamma > k|I_k), cP(\Gamma \leqslant k|I_k) + \min_{\phi_{k+1}^{(1)}, \phi_{k+1}^{(2)}} E[\tilde{J}_{k+1}^T(I_{k+1})\,|\,I_k]\right\} \tag{A.6}$$

where the first term in the outer minimum is the cost of stopping at time k and deciding that a change has taken place, and the second term is the cost continuing at time k.

The sufficient statistic for (P2) is $p_k = P(\Gamma \leqslant k|I_k)$. The recursion for p_k is obtained as follows:

$$\begin{aligned} p_{k-1} &= P(\Gamma \leqslant k+1|I_{k+1}) \\ &= P(\Gamma \leqslant k+1|U_{k+1}^{(1)}, U_{k+1}^{(2)}, I_k) \\ &= \frac{P(\Gamma \leqslant k+1|I_k) f(U_{k+1}^{(1)}, U_{k+1}^{(2)}|\Gamma \leqslant k+1; I_k)}{f(U_{k-1}^{(1)}, U_{k+1}^{(2)}|I_k)} \end{aligned}$$

Now,

$$\begin{aligned} P(\Gamma \leqslant k+1|I_k) &= P(\Gamma \leqslant k|I_k) + P(\Gamma = k+1|I_k) \\ &= p_k + P(\Gamma = k+1|I_k; \Gamma \geqslant k+1)P(\Gamma \geqslant k+1|I_k) \\ &= p_k + \rho(1 - p_k) \end{aligned}$$

and

$$f(U_{k+1}^{(1)}, U_{k+1}^{(2)}|\Gamma \leqslant k+1; I_k) = \prod_{l=1}^{2} [P_1\{L^{(l)}(X_{k+1}^{(l)}) > \lambda_{k+1}^{(l)}(I_k)\}]^{U_{k+1}^{(l)}} \times [P_1\{L^{(l)}(X_{k+1}^{(l)}) \leqslant \lambda_{k+1}^{(l)}(I_k)\}]^{1-U_{k+1}^{(l)}}$$

and

$$f(U_{k+1}^{(1)}, U_{k+1}^{(2)}|I_k) = f(U_{k+1}^{(1)}, U_{k+1}^{(2)}|I_k; \Gamma \leqslant k+1)P(\Gamma \leqslant k+1|I_k) + f(U_{k+1}^{(1)}, U_{k+1}^{(2)}|I_k; \Gamma > k+1)P(\Gamma > k+1|I_k)$$

Thus

$$p_{k+1} = \frac{g(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \lambda_{k+1}^{(1)}(I_k), \lambda_{k+1}^{(2)}(I_k); p_k)}{h(U_{k+1}^{(1)}, U_{k+1}^{(2)}; \lambda_{k+1}^{(1)}(I_k), \lambda_{k+1}^{(2)}(I_k); p_k)}, \quad p_0 = v$$

where g and h are as defined in Eqs (19) and (20).

With the above modifications, Theorem 3 can be proved using a sequence of arguments that is identical to those given in the proof of Theorem 1. □

References

[1] T. Başar, G.J. Olsder, Dynamic Noncooperative Game Theory, Academic Press, New York, 1982.
[2] G. Owen, Game Theory, Academic Press, San Diego, 1995.
[3] J.N. Tsitsiklis, Decentralized detection, in: H.V. Poor, J.B. Thomas (Eds.), Advances in Statistical Signal Processing, vol. 2 — Signal Detection, JAI Press, Greenwich, CT, 1993.
[4] H.V. Poor, An Introduction to Signal Detection and Estimation, 2nd ed., Springer, New York, 1988.
[5] H.L. Van Trees, Detection, Estimation, and Modulation theory, Wiley, New York, 1968.
[6] P.K. Varshney, Distributed Detection and Data Fusion, Springer, New York, 1996.
[7] A. Wald, Sequential Analysis, Wiley, New York, 1947.
[8] D. Siegmund, Sequential Analysis: Tests and Confidence Intervals, Springer, New York, 1985.
[9] B.K. Ghosh, P.K. Sen, Handbook of Sequential Analysis, Marcel Dekker, New York, 1991.
[10] P.F. Swaszek, On the performance of serial networks in distributed detection, IEEE Trans. Aerosp. Electron. Systems 29 (1993) 254–260.
[11] P.F. Swaszek, P. Willett, Parley as an approach to distributed detection, IEEE Trans. Aerosp. Electron. Systems 31 (1995) 447–457.
[12] P. Whittle, Optimization Over Time, vols. 1 and 2, Wiley, New York 1983.
[13] D. Bertsekas, Dynamic Programming, Prentice-Hall, Princeton, NJ, 1987.
[14] T. Başar, J.B. Cruz, Concepts and methods in multi-person coordination and control, in: S.G. Tzafestas (Ed.), Optimisation and Control of Dynamic Operational Research Models, North-Holland, Amsterdam, 1982, pp. 351–394.
[15] Y.C. Ho, Team decision theory and information structures, Proc. IEEE, 68 (1980) 644–654.
[16] I. Katzela, M. Naghshineh, Channel assignment schemes for cellular mobile telecommunication systems: a comprehensive survey, IEEE Pers. Commun. Mag. 3 (3) (1996) 10–31.
[17] N.G. Badr, Cellular digital packet data CDPD Proc Int. Phoenix Conf. Comput. Commun., Scottsdale, AZ, March 1995, pp. 659–665.
[18] K. Pahlavan, A.H. Levesque, Wireless Information Networks, Wiley, New York, 1995.
[19] P. Papatoni-Kazakos, Algorithms for monitoring changes in quality of communication links, IEEE Trans. Commun. COM-27 (1979) 1261–1266.

[20] C. Wang, M. Schwartz, Fault detection with multiple observers, IEEE/ACM Trans. Networking 1 (1994) 48–55.
[21] A.N. Shiryayev, Optimal Stopping Rules, Springer, New York, 1978.
[22] J.N. Tsitsiklis, On threshold rules in decentralized detection, Proc. 25th IEEE Conf. on Decision and Control, Athens, Greece, December 1986, pp. 232–236.
[23] V.V. Veeravalli, T. Başar, H.V. Poor, Decentralized sequential detection with a fusion center performing the sequential test, IEEE Trans. Inform. Theory IT-39 (1993) 433–442.
[24] V.V. Veeravalli, Comments on 'Decentralized sequential detection', IEEE Trans. Inform. Theory IT-38 (1992) 1428–1429.
[25] E.S. Page, Continuous inspection schemes. Biometrika 41 (1954) 100–114.
[26] G. Lorden, Procedures for reacting to a change in distribution, Ann. Math. Statist. 42 (1971) 1897–1908.
[27] G.V. Moustakides, Optimal stopping times for detecting changes in distributions, Ann. Statist. 14 (4) (1986) 1379–1387.
[28] D. Teneketzis, P. Varaiya, The decentralized quickest detection problem, IEEE Trans. Automat. Control, AC-29 (1984) 641–644.
[29] R.W. Crow, S.C. Schwartz, Quickest detection for sequential decentralized decision systems, IEEE Trans. Aerosp. Electron. Systems 32 (1996) 267–283.
[30] Y. Gao, H.F. Durrant-Whyte, Multi-sensor fault detection and diagnosis using combined qualitative and quantitative techniques, Proc. IEEE Int. Conf. Multisensor Fusion and Integration for Intelligent Systems, Las Vegas, NV, 1994, pp. 43–50.
[31] V.V. Veeravalli, Decentralized quickest change detection, Proc. Int. Symp. Inform. Theory Whistler, BC, Canada, September 1995, p. 294.
[32] V.V. Veeravalli, Decentralized change detection, IEEE Trans. Inform. Theory (submitted).
[33] V.V. Veeravalli, Further results on decentralized change detection, Proc. Internat. Symp. Inform. Theory, Ulm, Germany, June–July 1997, p. 294.