86-09

A CONTINUOUS–TIME DISTRIBUTED VERSION OF WALD'S SEQUENTIAL HYPOTHESIS TESTING PROBLEM

by

Anthony LaVigna[1], Armand M. Makowski[2], and John S. Baras[3]

Electrical Engineering Department
and
Systems Research Center
University of Maryland
College Park, Maryland 20742

ABSTRACT

This paper discusses a distributed version of Wald's sequential hypothesis testing problem in the continuous time framework. For sake of concreteness, two decision-makers equipped with their own sensors, are faced with the following hypothesis testing problem : Decide between hypothesis H_0 and H_1, where

$$\begin{aligned} \text{Under } H_1: \quad & dX_t^i = \mu_i\, dt + \sigma_i\, dW_t^i, \quad i = 1,2 \\ \text{Under } H_0: \quad & dX_t^i = \sigma_i\, dW_t^i, \quad i = 1,2, \end{aligned}$$

with $\mu_i \neq 0$ and $\sigma_i \neq 0$, $i = 1,2$, non-random; here the noises $\{W_t^1,\ t \geq 0\}$ and $\{W_t^2,\ t \geq 0\}$ are independent Brownian motions.

Data is observed continuously and at each instant in time, each decision-maker can either declare one of the hypotheses to be true or continue collecting data. In either case, they base their individual decisions on the data collected by their own sensors up to that time; they do not communicate with each other and so do not share information. The decisions are selected to minimize a *joint* cost function with two components, the first one capturing the cost for *collecting* data, and the second assessing the cost for *incorrect* decisions. This is the simplest problem of its type, for the coupling between the two decision–makers occurs only through the cost structure. This problem was considered first in discrete–time by Teneketzis [6] who showed that the person-by-person optimal strategy was of threshold type for each sensor. Here a similar result is derived by *simple* and *direct* arguments based on well–known facts for the single detector problem. Moreover, *explicit* formulae are derived for this joint cost function when the detector policies are of threshold type, owing to the fact that at the decision times, the likelihood functionals assume *one* of two threshold values owing to the *continuity* of the paths of Brownian motion. This is in sharp contrast with the *overshoot* phenomena that leads in the discrete–time situation to the celebrated Wald approximations. These explicit formulae not only vividly display the cost interaction taking place between the two sensors but readily allow for a reduction of the original problem to a mathematical programming problem in four variables over a simple constraint set.

---

[1] The work of this author was supported by an ONR Fellowship.

[2] The work of this author was supported partially through ONR Grant N00014-84-K-0614, partially through NSF Grant ECS-83-51836 and partially through a Grant from GM Laboratories.

[3] The work of this author was supported partially through NSF Grant NSFD CDR -85-00108 and partially through ONR Grant N00014-83-K-0731

1. INTRODUCTION:

Consider the following *sequential testing* of *two* simple hypotheses H_0 and H_1 with *two* decision-makers: Decision-maker i is equipped with its own sensor and observes the increments dX_t^i of the stochastic process $X^i \equiv \{X_t^i, t \geq 0\}$ with values in $I\!R$. Under each hypothesis, the observed data is the output of a stochastic dynamical system, i.e.,

$$\text{Under } H_1: \qquad dX_t^i = \mu_i\, dt + \sigma_i\, dW_t^i, \quad i = 1, 2$$
$$\text{Under } H_0: \qquad dX_t^i = \qquad\quad \sigma_i\, dW_t^i, \quad i = 1, 2,$$

where $\mu_i \neq 0$ and $\sigma_i \neq 0$, $i = 1, 2$, are non-random. Here the noises $\{W_t^1, t \geq 0\}$ and $\{W_t^2, t \geq 0\}$ are independent Brownian motions. As usual, hypotheses H_1 and H_0 correspond to the "signal in noise" and "noise only" situations extensively treated in the literature on signal detection.

Data is observed continuously starting at an initial time which is taken for convenience to be zero. At each time $t > 0$, each decision maker can either declare one of the hypotheses to be true or continue collecting data. In either case, the decision makers base their individual decisions on the data collected by their own sensors up to that time; they do not communicate with each other and so do not share information. These decisions are made in some "optimal" fashion in that they minimize a *joint* cost function with two components, the first one capturing the cost for *collecting* data and the second assessing the cost for *incorrect* decisions. This is the simplest problem of its type, for the coupling between the two decision–makers occurs only through the cost structure. Such a problem was considered first in discrete–time by Teneketzis [6] who showed that the person-by-person optimal strategy was of threshold type for each decision-maker. Here a similar result is derived by *simple* and *direct* arguments based on well–known facts for the single detector problem.

The advantages of the continuous–time formulation over the discrete–time one lie in the following features: *Explicit* formulae can be derived for the joint cost function when the detector policies are of threshold type. This follows from the fact that at the decision times, the likelihood functionals assume *one* of two threshold values owing to the *continuity* of the paths of Brownian motion. This is in sharp contrast with the *overshoot* phenomena that leads in the discrete–time situation to the celebrated Wald approximations. These explicit formulae not only vividly display the cost interaction taking place between the two sensors but readily allow for a reduction of the original problem to a mathematical programming problem in four variables over a simple constraint set.

The paper is organized as follows: The precise problem formulation is given in Section 2, the single–detector problem is presented in Section 3, and the results described in Section 4. The technical discussion is contained in Sections 5 and 6, each section being devoted to the proof of a main result. A useful technical lemma is given in the appendix for completeness.

2 PROBLEM FORMULATION:

In this paper, a Bayesian formulation is adopted for the problem loosely described in the introduction. Let $(\Omega, \mathcal{F}, P)$ be a probability triple carrying a $\{0,1\}$-valued random variable (RV) H and two $\mathbb{R}$-valued processes $\{W_t^1,\, t \geq 0\}$ and $\{W_t^2,\, t \geq 0\}$. Throughout the discussion, it is assumed that *under* P,

(A1): The RV H and the processes $\{W_t^1,\, t \geq 0\}$ and $\{W_t^2,\, t \geq 0\}$ are *mutually* independent,
(A2): Each process $\{W_t^i,\, t \geq 0\}$, $i = 1,2$, is a *standard Brownian motion*, and
(A3): For some λ , $0 < \lambda < 1$,

$$P[H = 1] = \lambda = 1 - P[H = 0]. \tag{2.1}$$

The $\mathbb{R}$-valued processes $\{X_t^1,\, t \geq 0\}$ and $\{X_t^2,\, t \geq 0\}$ are defined as outputs of the dynamical system

$$\begin{aligned} dX_t^i &= H\,\mu_i\,dt + \sigma_i\,dW_t^i, \\ X_0^i &= 0 \end{aligned} \qquad i = 1,2, \tag{2.2}$$

or equivalently,

$$X_t^i = H\,\mu_i\,t + \sigma_i\,W_t^i, \quad t \geq 0, \quad i = 1,2, \tag{2.3}$$

where $\mu_i \neq 0$ and $\sigma_i \neq 0$ are *non*-random.

To fix the notation, let $\{\mathcal{F}_t^{W^i},\, t \geq 0\}$ and $\{\mathcal{F}_t^{X^i}, t \geq 0\}$ be the (P-completion of) the filtrations generated by the processes $\{W_t^i,\, t \geq 0\}$ and $\{X_t^i,\, t \geq 0\}$, respectively. Let

$$\mathcal{G}_t^i := \mathcal{F}_t^{X^i} \vee \sigma(H), \qquad t \geq 0 \tag{2.4}$$

with

$$\mathcal{G}^i := \bigvee_{t \geq 0} \mathcal{G}_t^i. \tag{2.5}$$

The decision policy for decision-maker i , $i = 1,2$, involves the selection of a *termination* time τ_i and of a binary valued *decision* δ_i, and an *admissible* decision policy for decision-maker i is thus any pair $\gamma_i = (\tau_i, \delta_i)$ of RV's, where τ_i is an $\mathcal{F}_t^{X^i}$-stopping time and δ_i is an $\mathcal{F}_{\tau_i}^{X^i}$-measurable $\{0,1\}$-valued RV. Denote by Γ_i the collection of all admissible decision policies for decision-maker i and let $\Gamma := \Gamma_1 \times \Gamma_2$.

Let k_1 and k_2 be positive constants and let C be a mapping $\{0,1\} \times \{0,1\} \times \{0,1\} \to \mathbb{R}^+$. For every pair $\gamma = (\gamma_1, \gamma_2)$ in Γ, define the cost function $J(\gamma)$ to be

$$J(\gamma) = E[\, k_1\tau_1 + k_2\tau_2 + C(\delta_1, \delta_2; H)\,]. \tag{2.6}$$

The problem $(\mathcal{P})$ investigated in this paper is then

(P): *Minimize* $J_\lambda(\gamma)$ *over* Γ.

An admissible policy γ^* in Γ is said to solve problem (P) if

$$J(\gamma^*) \leq J(\gamma). \qquad \textit{for all } \gamma \textit{ in } \Gamma. \tag{2.7}$$

The purpose of this paper is to identify the ***structure*** of an optimal policy.

3. THE SINGLE–DETECTOR PROBLEM

Consider the following single-detector sequential hypothesis testing problem defined say on some underlying probability triple $(\Omega, \mathcal{F}, P)$. A decision-maker observes the increments of a stochastic process $\{X_t,\ t \geq 0\}$ with values in $\mathbb{R}$. Under each hypothesis, the observed data include a signal which is the output of a stochastic dynamical system, i.e.,

$$dX_t = \mu\, H\, dt + \sigma\, dW_t \tag{3.1}$$

where as before $\mu \neq 0$ and $\sigma \neq 0$ are *non*-random, and $\{W_t,\ t \geq 0\}$ is a standard Brownian motion.

Data is observed continuously starting at an initial time which is taken for convenience to be zero. At each time $t > 0$, the decision maker can either declare one of the hypotheses to be true or continue collecting data. Decisions are made so as to optimize a cost function with two components, the first one capturing the cost for collecting data, and the second one assessing the cost for incorrect decisions.

More formally, let $\{\mathcal{F}_t^X,\ t \geq 0\}$ be the (P–completion of the) filtration generated by the process $\{X_t,\ t \geq 0\}$. An *admissible policy* is any pair $\gamma = (\tau, \delta)$ of RV's where τ is an $\mathcal{F}_t^X$-stopping time and δ is a $\{0,1\}$-valued $\mathcal{F}_t^X$ -measurable RV. The collection of all admissible policies is denotes by Γ_s. The cost corresponding to any γ in Γ_s is given by

$$\tilde{J}(\gamma; k, c) = E[\, k\tau + c\, 1[\delta \neq H]\,] \tag{3.2}$$

The problem faced by the single detector is the to find γ^* in Γ_s such that

$$\tilde{J}(\gamma^*; k, c) \leq \tilde{J}(\gamma; k, c) \tag{3.3}$$

for all other γ in Γ_s.

Let $\{\pi_t,\ t \geq 0\}$ be the $\mathcal{F}_t^X$-adapted process defined by

$$\pi_t := P[\, H = 1 \mid \mathcal{F}_t^X\,], \quad t \geq 0. \tag{3.4}$$

Note that π_t is the posteriori probability of hypothesis H_1 given the observations $\mathcal{F}_t^X$.

Theorem 3.1. *The admissible decision policy* $\gamma^* = (\tau^*, \delta^*)$ *given by*

$$\tau^* = \inf\{t \geq 0 \mid \pi_t \neq (A^*, B^*)\}, \tag{3.5a}$$

$$\delta^* = \begin{cases} 1, & \pi_{\tau^*} \geq B^*, \\ 0, & \pi_{\tau^*} \leq A^*, \end{cases} \tag{3.5b}$$

is optimal in the sense of (3.3). The constants A^* *and* B^* *are the unique solutions to the transcendental equations*

$$2c = \frac{2k\sigma^2}{\mu^2}(\Psi'(A^*) - \Psi'(B^*)) \tag{3.6}$$

$$c(1-B^*) = cA^* + (B^* - A^*)(c - \frac{2k\sigma^2}{\mu^2}\Psi'(A^*)) + k(\Psi(B^*) - \Psi(A^*)) \tag{3.7}$$

where $\Psi(x) = (1-2x)\log(x/(1-x))$.

Proof: See ([5], Theorem 5, p.185).

It is further known [5] that for all admissible policies of the form (3.5) with constants (A, B) satisfying $0 < A \leq 1 \leq B < \infty$ and $A \neq B$, the relations

$$E[\tau \mid H = 1] = 2\frac{\sigma^2}{\mu^2} w(\beta, \alpha), \qquad E[\tau \mid H = 0] = 2\frac{\sigma^2}{\mu^2} w(\alpha, \beta) \tag{3.8}$$

hold with

$$w(x, y) = (1-x)\log(\frac{1-x}{y}) + x\log(\frac{x}{1-y}), \quad 0 < x, y < 1, \tag{3.9}$$

and

$$\alpha = \frac{(1-B)(\lambda - A)}{(B-A)(1-\lambda)}, \qquad \beta = \frac{A(B-\lambda)}{\lambda(B-A)}. \tag{3.10}$$

4. THE RESULTS:

Throughout the discussion, the decision cost satisfies the natural conditions

$$C(m, n; n) \geq C(n, n; n) \tag{4.1}$$

$$C(m, m; n) \geq C(n, m; n) \tag{4.2}$$

$$C(m, n; h) = C(n, m; h) \tag{4.3}$$

for all n , m , h in $\{0, 1\}$, with $m \neq n$. In other words, it always cost more to make an error than it does to make the correct decision, regardless of which decision is taken by the other decision-maker. As readily checked, there is no loss in generality in assuming zero cost for correct decisions, i.e.,

$$C(1, 1; 1) = C(0, 0; 0) = 0, \tag{4.4}$$

an assumption enforced from now on.

To state the results of this paper, it is convenient to introduce the $\mathcal{F}_t^{X^i}$-adapted processes $\{\pi_t^i, t \geq 0\}$ and $\{\phi_t^i, t \geq 0\}$, $i = 1, 2$, where

$$\pi_t^i := P[\, H = 1 \,|\, \mathcal{F}_t^{X^i}], \qquad t \geq 0, \tag{4.5}$$

and

$$\phi_t^i := \exp[\, \frac{\mu_i}{\sigma_i^2}(X_t^i - \frac{1}{2}\mu^i t)], \qquad t \geq 0. \tag{4.6}$$

For each decision-maker, a policy $\gamma_i = (\tau_i, \delta_i)$ in Γ_i, $i = 1, 2$, is said to be of *threshold* type if there exists constants A_i and B_i, with $0 \leq A_i \leq B_i \leq 1$, such that

$$\tau_i := \inf\{t \geq 0 : \quad \pi_t^i \notin (A_i, B_i)\} \tag{4.7}$$

and

$$\delta_i := \begin{cases} 1 & \text{if } \pi_{\tau_i}^i \geq B_i \\ 0 & \text{if } \pi_{\tau_i}^i \leq A_i. \end{cases} \tag{4.8}$$

As is well known [5, p. 181], the relation

$$\pi_t^i = \frac{\lambda \phi_t^i}{(1-\lambda) + \lambda \phi_t^i}, \qquad t \geq 0, \tag{4.9}$$

holds. Therefore, any threshold policy for decision-maker i with constants (A_i, B_i) takes the form

$$\tau_i := \inf\{t \geq 0 : \quad \phi_t^i \notin (a_i, b_i)\} \tag{4.10}$$

and

$$\delta_i := \begin{cases} 1 & \text{if } \phi_{\tau_i}^i \geq b_i \\ 0 & \text{if } \phi_{\tau_i}^i \leq a_i, \end{cases} \tag{4.11}$$

where the new thresholds a_i and b_i are given by

$$a_i = \left(\frac{1-\lambda}{\lambda}\right)\left(\frac{A_i}{1-A_i}\right), \qquad b_i = \left(\frac{1-\lambda}{\lambda}\right)\left(\frac{B_i}{1-B_i}\right). \tag{4.12}$$

From now on, a threshold policy γ_i in Γ_i will be described in the form (4.10)–(4.11) and will thus be identified with two threshold constants a_i and b_i, $0 \leq a_i \leq 1 \leq b_i$, $a_i \neq b_i$. Let $\mathcal{T}_i$ be the collection of all threshold policies in Γ_i for decision-maker i, with the notation $\mathcal{T} = \mathcal{T}_1 \times \mathcal{T}_2$.

Theorem 4.1: *Under the foregoing assumptions,*

$$\inf_{\gamma \in \Gamma} J(\gamma) = \inf_{\gamma \in \mathcal{T}} J(\gamma), \tag{4.13}$$

i.e., only threshold policies need to be considered in solving problem $(\mathcal{P})$.

For simplicity of exposition, the discussion is carried out under the additional assumption that the decision cost has the form

$$C(d_1, d_2; h) = \begin{cases} 0 & \text{for } d_1 = d_2 = h \\ e & \text{for } d_1 \neq d_2 \\ f & \text{for } d_1 = d_2 \neq h. \end{cases} \tag{4.14}$$

The general case can be handled in a similar way but the calculations are more involved. Details are worked out in a lengthier version of this paper [3].

Let γ_1 and γ_2 be threshold policies with parameters (a_1, b_1) and (a_2, b_2), respectively, and pose

$$\alpha_i = P[\,\delta_i = 1\,|\,H = 0\,], \qquad \beta_i = P[\,\delta_i = 0\,|\,H = 1\,], \quad i = 1, 2, \tag{4.15}$$

the so-called error probabilities of the *first* and *second* kind, respectively, or in radar parlance, the *false alarm* and *miss probabilities.*

Theorem 4.2: *Under the foregoing assumptions, any pair of threshold policies γ_1 and γ_2 with parameters (a_1, b_1) and (a_2, b_2), respectively, incurs a cost $J(\gamma)$ given by the expression*

$$J(\gamma) = \tilde{J}(\gamma_1; k_1, e) + \tilde{J}(\gamma_2; k_2, e) + (f - 2e)[\lambda\beta_1\beta_2 + (1-\lambda)\alpha_1\alpha_2] \tag{4.16}$$

where $\tilde{J}(\gamma; k, c)$ is the function defined in (3.2).

Together Theorems 4.1 and 4.2 thus reduce the search for a solution γ^* to problem $(\mathcal{P})$ to a mathematical programming problem in the variables $(\alpha_1, \beta_1, \alpha_2, \beta_2)$ over a simple constraint set.

Theorem 4.3: *Under the foregoing assumptions, problem $(\mathcal{P})$ is solved by a pair $\gamma^* = (\gamma_1^*, \gamma_2^*)$ of threshold policies in $\mathcal{T}$.*

5. A DISCUSSION OF THEOREM 4.1.

The discussion given below hinges upon arguments very similar to the ones proposed by Teneketzis in the discrete-time case [6], as would be expected. Here too, the basic idea lies in switching attention from optimality for problem $(\mathcal{P})$ defined in (2.7) to *person-by-person* optimality, as understood in the literature on multi-agent decision-making. Under the strong independence assumptions made here, this procedure leads to a decoupling of the two decision-makers. Some of the details underlying this line of reasoning are presented in the next two technical lemmas.

For every admissible policy γ_2 in Γ_2, define the mapping $\widehat{C}_{\gamma_2}: \{0,1\} \times \{0,1\} \to I\!R$ by

$$\begin{aligned}\widehat{C}_{\gamma_2}(d_1, h) &:= E[\,C(d_1, \delta_2; h)\,|\,H = h\,] \\ &= C(d_1, 1; h)\,P[\delta_2 = 1\,|\,H = h] + C(d_1, 0; h)\,P[\delta_2 = 0\,|\,H = h]\end{aligned} \tag{5.1}$$

for all (d_1, h) in $\{0,1\} \times \{0,1\}$. The conditions (4.1)-(4.2) readily imply that this mapping satisfies the condition

$$\widehat{C}_{\gamma_2}(\delta_1, h) \geq \widehat{C}_{\gamma_2}(h, h). \tag{5.2}$$

for all (d_1, h) in $\{0,1\} \times \{0,1\}$.

Lemma 5.1: *Under the foregoing assumptions, the σ-fields $\mathcal{G}^1$ and $\mathcal{G}^2$ are conditionally independent given H.*

Proof: This readily follows from the assumptions (A1), (2.3) and the obvious σ-field identity

$$\mathcal{G}_t^i = \mathcal{F}_t^{X^i} \vee \sigma(H) = \mathcal{F}_t^{W^i} \vee \sigma(H), \quad t \geq 0. \qquad \square$$

Lemma 5.2. *With the notation (5.1), the relation*

$$E[C(\delta_1, \delta_2; H)] = E[\widehat{C}_{\gamma_2}(\delta_1, H)] \tag{5.3}$$

holds true for every pair $\gamma = (\gamma_1, \gamma_2)$ in Γ.

Proof: By standard properties of conditional expectations, the equalities

$$E[C(\delta_1, \delta_2; H)] = E[\ E[C(\delta_1, \delta_2; H \mid \mathcal{G}^1]\ ] \tag{5.4}$$

$$= E[\ E[C(d_1, \delta_2; h) \mid \mathcal{G}^1]_{d_1 = \delta_1, h = H}\ ] \tag{5.5}$$

are readily obtained. From Proposition A.1 of the appendix and from Lemma 5.1, it follows that

$$E[C(d_1, \delta_2; h) \mid \mathcal{G}^1] = E[C(d_1, \delta_2; h) \mid \sigma(H)] \tag{5.6}$$

for all (d_1, h) in $\{0, 1\} \times \{0, 1\}$. Substitution of this last relation into (5.5) and use of the definition (5.1) now yield the result. $\square$

Proof of Theorem 4.1: With $\gamma = (\gamma_1, \gamma_2)$ in Γ, Lemma 5.2 implies that

$$J(\gamma) = E[\,k_1 \tau_1 + k_2 \tau_2 + \widehat{C}_{\gamma_2}(\delta_1, H)\,]$$

$$= k_2\, E[\,\tau_2\,] + \widehat{J}_{\gamma_2}(\gamma_1), \tag{5.7}$$

where

$$\widehat{J}_{\gamma_2}(\gamma_1) = E[\,k_1 \tau_1 + \widehat{C}_{\gamma_2}(\delta_1, H)\,]. \tag{5.8}$$

The sequential hypothesis problem $(\mathcal{P}/\gamma_2)$ faced by decision maker 1, with γ_2 fixed, takes the form

$(\mathcal{P}/\gamma_2)$ *Minimize* $\widehat{J}_{\gamma_2}(\gamma_1)$ *over* Γ_1.

Under the assumed conditions, this problem is a single-detector sequential hypothesis testing problem with decision cost $\widehat{C}_{\gamma_2}$ of the type discussed in Section 3. Here $k = k_1$ and $c = c$. As pointed out in Theorem 3.1, the admissible policy γ_1^* in Γ_1 that solves problem $(\mathcal{P}/\gamma_2)$ is of *threshold* type ([5], Theorem 5, p. 185).

Therefore, for every pair γ in Γ, there exists a threshold policy $\widetilde{\gamma}_1$ in $\mathcal{T}_1$ such that

$$\widehat{J}_{\gamma_2}(\widetilde{\gamma}_1) \leq \widehat{J}_{\gamma_2}(\gamma_1), \tag{5.9}$$

or equivalently,

$$J(\widetilde{\gamma}_1, \gamma_2) \leq J(\gamma_1, \gamma_2) = J(\gamma). \tag{5.10}$$

By symmetry, interchanging the role of γ_1 and γ_2, there exists a threshold policy $\tilde{\gamma}_2$ in $\mathcal{T}_2$ such that

$$J(\gamma_1, \tilde{\gamma}_2) \leq J(\gamma). \tag{5.11}$$

This last step, when applied with the pair $(\tilde{\gamma}_1, \gamma_2)$ (instead of (γ_1, γ_2)) yields the existence of a threshold policy $\tilde{\gamma}_2$ in $\mathcal{T}$ such that

$$J(\tilde{\gamma}_1, \tilde{\gamma}_2) \leq J(\tilde{\gamma}_1, \gamma_2). \tag{5.12}$$

This shows that to every admissible policy γ in Γ, there corresponds an admissible policy $\tilde{\gamma}$ in $\mathcal{T}$ such that

$$J(\tilde{\gamma}) \leq J(\gamma), \tag{5.13}$$

and the result of Theorem 4.1 is now immediate since $\mathcal{T} \subseteq \Gamma$. □

6. PROOF OF THEOREM 4.2:

Let γ be an admissible policy $\mathcal{T}$. By elementary calculations, the corresponding cost is given by

$$\begin{aligned}
J(\gamma) &= E\Big[\, k_1\tau_1 + k_2\tau_2 + C(\delta_1, \delta_2; H) \Big] &(6.1)\\
&= E\Big[\, k_1\tau_1 + k_2\tau_2 + C(1, \delta_2; H)\, P[\,\delta_1 = 1 \mid H\,] + C(0, \delta_2; H)\,(1 - P[\,\delta_1 = 1 \mid H\,]) \Big] &(6.2)\\
&= E\Big[\, k_1\tau_1 + k_2\tau_2 + C(1,1;H)\, P[\,\delta_1 = 1 \mid H\,] P[\,\delta_2 = 1 \mid H\,]\\
&\qquad + C(1,0;H)\, P[\,\delta_1 = 1 \mid H\,] \cdot P[\,\delta_2 = 0 \mid H\,]\\
&\qquad + C(0,1;H)\,(1 - P[\,\delta_1 = 1 \mid H\,]) P[\,\delta_2 = 1 \mid H\,]\\
&\qquad + C(0,0;H)\,(1 - P[\,\delta_1 = 1 \mid H\,])(1 - P[\,\delta_2 = 1\,]) \Big] &(6.3)\\
&= E\Big[\, k_1\tau_1 + k_2\tau_2 \Big] + \lambda\Big[C(1,1;1)(1-\beta_1)(1-\beta_2) + C(1,0;1)(1-\beta_1)\beta_2\\
&\qquad + C(0,0;1)\beta_1(1-\beta_2) + C(0,0;1)\beta_1\beta_2\Big]\\
&\qquad + (1-\lambda)\Big[C(1,1;0)\alpha_1\alpha_2 + C(1,0;0)\alpha_1(1-\alpha_2)\\
&\qquad + C(0,1;0)(1-\alpha_1)\alpha_2 + C(0,0;0)(1-\alpha_1)(1-\alpha_2)\Big] &(6.4)\\
&= E\Big[\, k_1\tau_1 + k_2\tau_2 \Big] + \lambda\Big[e(\beta_1 + \beta_2) + (k - 2e)\beta_1\beta_2\Big]\\
&\qquad + (1-\lambda)\Big[e(\alpha_1 + \alpha_2) + (k - 2e)\alpha_1\alpha_2\Big] &(6.5)
\end{aligned}$$

where α_i and β_i were defined in (4.16).

In order to compute the first expectation term in (6.5), it is convenient to observe that

$$E\Big[\tau_i\Big] = \lambda E\Big[\tau_i \mid H = 1\Big] + (1-\lambda)\, E\Big[\tau_i \mid H = 0\Big], \qquad i = 1, 2 \tag{6.6}$$

thus reducing the final step of the computation to the evaluation of the conditional expectations

$$E\Big[\tau_i \mid H = h\Big], \qquad i = 1, 2,$$

whose expressions are given in (3.8). The result of Theorem 4.2 now follows. □

7. CONCLUSIONS:

Based on Theorem 4.2 an explicit solution can be computed for the thresholds of each agent once the parameters of the problem are given (k_1, k_2, e, f). The mathematical programming problem described in Theorem 4.2 is readily seen to have a unique solution, and can be solved by a variety of methods, say a Newton-like scheme. By solving (6.5) for the constants $(\alpha_1, \beta_1, \alpha_2, \beta_2)$, the thresholds are easily determined via the relations (3.10) and (4.12).

REFERENCES:

[1] S. Karlin and H. M. Taylor, *A First Course in Stochastic Processes*, Academic Press, New York, 1975.

[2] A. LaVigna, *Real Time Sequential Detection for Diffusion Signals* , M.S. Thesis, Electrical Engineering Department, University of Maryland, College Park, 1986.

[3] A. LaVigna, *Distributed sequential hypothesis testing* , SRC Technical report, 1986, *Under preparation.*

[4] R. S. Lipster and A. N. Shiryayev, *Statistics of Random Processes II, Applications* (English Translation), Springer-Verlag, New York, 1978.

[5] A. N. Shiryayev, *Optimal Stopping Rules* (English Translation), Springer-Verlag, New York, 1977.

[6] D. Teneketzis, " The Decentralized Wald Problem." , *Proceedings of the Large-Scale Systems Symposium*, Virginia Beach, Virginia, October 1982.

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

APPENDIX

On some probability triple $(\Omega, \mathcal{F}, P)$, consider three σ-fields $\mathcal{F}_1$, $\mathcal{F}_2$ and $\mathcal{F}_3$, and X be an integrable $\mathbb{R}$-valued RV.

Proposition A.1: *If the σ-fields $\mathcal{F}_1$ and $\mathcal{F}_2$ are conditionally independent given the σ-field $\mathcal{F}_3$, then*

$$E[\, X \mid \mathcal{F}_2 \vee \mathcal{F}_3 \,] = E[\, X \mid \mathcal{F}_3 \,] \quad P-a.s. \tag{A1}$$

whenever X is an $\mathcal{F}_1$-measurable RV.

Proof: By virtue of the Monotone Class Theorem, it suffices to establish (A.1) for any *bounded* $\mathcal{F}_1$-measurable RV. To that end, let F_2 and F_3 be arbitrary elements of the σ-fields $\mathcal{F}_2$ and $\mathcal{F}_3$, respectively. By the very definition of conditional expectation,

$$E[\, 1_{F_2 \cap F_3}\, E[\, X \mid \mathcal{F}_2 \vee \mathcal{F}_3 \,]\,] = E[\, 1_{F_2} \times 1_{F_3} \,] \tag{A2}$$

$$= E[\, 1_{F_3}\, E[\, 1_{F_2} X \mid \mathcal{F}_3 \,]\,] \tag{A3}$$

$$= E[\, 1_{F_3}\, E[\, 1_{F_2} \mid \mathcal{F}_3 \,]\, E[\, X \mid \mathcal{F}_3]\,], \tag{A4}$$

where the last equality was obtained by making use of the conditional independence of the σ-fields $\mathcal{F}_1$ and $\mathcal{F}_2$ given $\mathcal{F}_3$. The $\mathcal{F}_3$-measurability of the RV's 1_{F_3} and $E[X \mid \mathcal{F}_3]$ now readily yields

$$E[1_{F_3} E[1_{F_2} \mid \mathcal{F}_3] E[X \mid \mathcal{F}_3]] = E[E[1_{F_3} E[X \mid \mathcal{F}_3] 1_{F_2} \mid \mathcal{F}_3]] \tag{A5}$$

$$= E[1_{F_2 \cap F_3} E[X \mid \mathcal{F}_3]] \tag{A6}$$

via the smoothing property for conditional expectations.

Consequently,

$$E[1_{F_2 \cap F_3} E[X \mid \mathcal{F}_2 \vee \mathcal{F}_3]] = E[1_{F_2 \cap F_3} E[X \mid \mathcal{F}_3]] \tag{A7}$$

and the result (A1) immediately follows by standard arguments. □