On the Optimality of Sequential Test with Multiple Sensors

Vaibhav Katewa and Vijay Gupta

Abstract— **We study the problem of sequential detection for binary hypothesis testing using multiple sensors. We consider a randomized sensor selection strategy in which one sensor can be active at any given time step. We obtain an optimal sequential test using dynamic programming and show that it corresponds to the Sequential Probability Ratio Test (SPRT) when the sensor selection process is stationary. Further, we prove that Wald-Wolfowitz theorem holds true for sequential test with multiple sensors.**

I. INTRODUCTION

With the advance in communication and processing technology, networks with distributed components are being proposed widely. The different components of such distributed systems work in cooperation to achieve a particular objective. The class of hypothesis testing problems is one such application [5] where multiple sensors collect information from the environment. This collective information is used by the system to infer the true hypothesis. Using multiple sensors for detection may be desired as different sensors may provide different "views of the world". For example, a wireless sensor network can perform an object detection using the collective information from multiple sensors distributed over a geographical area. Presence of multiple sensors increases reliability, survivability and coverage of the system.

The use of multiple sensors introduces new challenges like sensor management, sensor scheduling, sensor selection and sensor fusion. Many works have focused on sensor management. The problem of sensor management for linear estimation has been widely considered (see, e.g. [1], [3], [4] and the references therein). In [2], [5], the authors study a decentralized detection problem where the fusion center implements an optimal policy to combine information from different sensors.

In the present work, we use a sequential detection rule to perform a simple binary hypothesis testing using multiple sensors. Sequential tests are the class of tests where the number of observations are not fixed, but can vary from one experiment to another. More accurate decisions can be made by taking large number of observations. However, there is a cost associated with taking observations and sequential tests resolve the trade-off between accuracy and experiment cost. A well known example is the Sequential Probability Ratio Test (SPRT) [8] developed by Wald. SPRT is optimal in the sense that for a given error performance, it requires minimal number of observations on an average [9]. We extend the conventional SPRT to the case when multiple sensors are used for taking observations.

Department of Electrical Engineering, University of Notre Dame, Notre Dame, IN 46556. <vkatewa@nd.edu, vijay.gupta.21@nd.edu>

There seem to be two main approaches in the literature to solve sequential detection problems. One approach is to formulate the sequential test as a solution to Bayesian optimization problem [10], [11]. Dynamic programming is applied to minimize the Bayesian cost. The dynamic programming introduces complexity in the solution, but the solution is optimal. SPRT is an example of tests obtained through this approach. Another approach is to analyze the sequential detection problem in asymptotic regime, where the cost of taking observations is very small [11], [6]. Although this approach results in tractable solutions, the asymptotically optimal results may not be optimal outside the asymptotic regime. Further, there may be situations in which there are limits on the average number of observations from the sensors. In such a scenario, the asymptotic approach cannot be applied. Therefore, we take the dynamic programming approach to solve the sequential detection problem.

Many works ([3], [14], [13] and references therein) have focused on sensor selection and management for detection purposes. To manage the sensors, we choose a probabilistic sensor selection strategy in which one sensor is chosen randomly among multiple sensors at each time step. Further, for tractability, the sensor selection probability distribution is assumed to be stationary throughout the experiment. The sensors have different costs of taking an observation that are constant throughout the experiment.

We formulate a Bayesian optimization problem and obtain SPRT as an optimal solution to the problem. The closest work to ours seems to be [7], where the authors present a sensor selection strategy for multiple hypothesis testing. However, they perform an asymptotic analysis and obtain SPRT as a asymptotically optimal sequential test. An important contribution of our work is that we show that SPRT is optimal outside the asymptotic regime also using dynamic programing.

II. PROBLEM FORMULATION

We now formally define the problem setup. The notations used here are motived from [10].

Notation: Random variables are denoted by uppercase letters (eg. Z) and their realizations are denoted by lowercase letters (eg. z). Let $Z_i^j = (Z_i, Z_{i+1}, \cdots, Z_j)$ for $0 \leq i \leq j$ denote a sequence of random variables and $z_i^j = (z_i, z_{i+1}, \cdots, z_j)$ denote a realization of this random sequence. Let $\mathbb{E}_Z[.]$ and $\mathbb{E}_Z[.|y]$ denotes unconditional expectation w.r.t. Z and conditional expectation w.r.t. Z given $Y = y$, respectively. Further, $\mathbb{R}^+$ denotes the non-negative real line.

Hypothesis Space: Let $\Phi = \{\theta_0, \theta_1\}$ denote the binary hypothesis space or a set of all possible states of nature. Let Θ denote the random binary hypothesis and θ denote a realization of the random hypothesis. Let hypothesis $\{\theta_0, \theta_1\}$ be denoted by $\{H_0, H_1\}$. Let the *a-priori* probabilities of H_0 and H_1 be $1 - \pi$ and π, respectively.

Action Space: Let $\mathcal{A}$ denote the action space or the set of all possible decisions available when the test is stopped. Let $a \in \mathcal{A}$ denote a general state of nature. For the given problem, $\mathcal{A} = \Phi = \{\theta_0, \theta_1\}$.

Sensor Selection Space: Let K be the total number of sensors available for taking measurements. Let $S_i^j = (S_i, S_{i+1}, \cdots, S_j)$ for $0 < i \leq j$ denote a random sensor selection sequence, where S_j denotes the random sensor selected at time step j. The sample space of S_j is the set $\mathcal{S} = \{1, 2, \cdots, K\}$ for all $j > 0$. At each time step, a single sensor is chosen with the stationary probability mass distribution $p = (p_1, p_2, \cdots, p_K)$ with $\sum_{i=1}^{K} p_i = 1$, where p_i is the probability of selecting sensor i at any time step.

Observation Space: Let $X_i^j = (X_i, X_{i+1}, \cdots, X_j)$ for $0 < i \leq j$ denote a random observation sequence. The sample space of X_j is denoted by $\mathcal{X}$ for all $j > 0$. At each time step, the observation can be taken by any one sensor with the probability distribution define above. Each observation is assumed to be i.i.d. conditioned on the hypothesis and selected sensor. Let $f_i^j(x_k)$ denote the probability density of X_k under hypothesis H_i and sensor j for $i = 0, 1$ and $j = 1, 2, \cdots, K$.

Loss Function: Let $L(\theta, a) : \Phi \times \mathcal{A} \to \mathbb{R}^+$ denote the finite, non-negative loss function when the true state of nature is θ and the terminal decision is a. For simplicity, we assume a finite non-negative loss function

$$L(\theta_0, \theta_0) = 0, \qquad L(\theta_1, \theta_1) = 0,$$
$$L(\theta_0, \theta_1) = w_{01}, \qquad L(\theta_1, \theta_0) = w_{10}.$$

Observation Cost: Let the cost of measurement from sensor j be d_j for $1 \leq j \leq K$. Each d_j is constant, finite, and non-negative. Let $C_i^j = (C_i, C_{i+1}, \cdots, C_j)$ for $0 < i \leq j$ denote a random observation cost sequence. The random variable C_j is distributed on the sample space $\mathcal{C} = \{d_1, d_2, \cdots, d_K\}$ with the distribution p for all $j > 0$. Let the mean observation cost be denoted by $\bar{d} = \sum_{i=1}^{K} p_i d_i$.

Stopping Rule: Let the stopping rule be denoted by a sequence of functions $\varphi = (\varphi_0, \varphi_1(x_1, s_1), \varphi_2(x_1^2, s_1^2), \cdots)$ where $\varphi_j : \mathcal{X}^j \times \mathcal{S}^j \to \{0, 1\}$. Thus, at time step j, the rule φ_j decides to stop or continue based on the observations collected and the sensors used till that time step. If $\varphi_j(x_1^j, s_1^j) = 1$, then the test is stopped, otherwise we continue to take one more observation. The stopping rule can be defined alternatively by a sequence of functions $\psi = (\psi_0, \psi_1(x_1, s_1), \psi_2(x_1^2, s_1^2), \cdots)$, where $\psi_j : \mathcal{X}^j \times \mathcal{S}^j \to \{0, 1\}$ is the stopping rule at time step j, given that the test

is not stopped till time step $j - 1$. Formally,

$$\psi_j(x_1^j, s_1^j) = (1 - \varphi_0)(1 - \varphi_1(x_1, s_1))(1 - \varphi_2(x_1^2, s_1^2)) \tag{1}$$
$$\cdots (1 - \varphi_{j-1}(x_1^{j-1}, s_1^{j-1}))\varphi_j(x_1^j, s_1^j).$$

Stopping Time: The stopping time N can be define as

$$N = \min\{k : \varphi_k(x_1^k, s_1^k) = 1\} = \{k : \psi_k(x_1^k, s_1^k) = 1.\}$$

Terminal Decision Rule: When the test is stopped by applying the stopping rule, a decision between the available hypothesis is made by the terminal decision rule(TDR) based on the available observations and the sensors used. It is denoted by a sequence of functions $\delta = (\delta_0, \delta_1(x_1, s_1), \delta_2(x_1^2, s_1^2), \cdots)$ where $\delta_j : \mathcal{X}^j \times \mathcal{S}^j \to \mathcal{A}$. The sequential detection rule is defined by the pair (φ, δ).

Risk Function: As discussed before, there exists a trade-off between the detection accuracy and the observation cost as we take more observations. To incorporate this trade-off, we define the risk function as

$$R(\pi, (\varphi, \delta)) = \sum_{j=0}^{\infty} \mathbb{E}_{\Theta, X_1^j, S_1^j} \left[\psi_j(X_1^j, S_1^j) \right. \tag{2}$$
$$\left. \left(L(\Theta, \delta_j(X_1^j, S_1^j)) + \sum_{i=1}^{j} C_i(S_i) \right) \right].$$

Problem Statement: The goal is to find a sequential detection rule (φ, δ) with the structure defined above, that minimizes the overall risk $R(\pi, (\varphi, \delta))$ of the experiment. This solution is referred to optimal/Bayes w.r.t. the risk R.

III. OPTIMAL SOLUTION VIA DYNAMIC PROGRAMMING

In this section, we solve the above optimization problem through dynamic programming and derive an optimal sequential test. We first proceed by finding the optimal terminal decision rule.

Theorem 3.1: Let $\hat{\delta}_j(x_1^j, s_1^j)$ denote the *Bayes rule* (optimal rule that minimizes the Bayes risk) for the binary hypothesis testing problem w.r.t. hypothesis space Φ, the action space $\mathcal{A}$, the observation space $\mathcal{X}$, the sensor space $\mathcal{S}$ and the loss function L based on observations x_1^j and selected sensors s_1^j. Then, for any fixed stopping rule φ, the risk $R(\pi, (\varphi, \delta))$ is minimized by $\hat{\delta} = (\hat{\delta}_0, \hat{\delta}_1(x_1, s_1), \hat{\delta}_2(x_1^2, s_1^2), \cdots)$.

Proof: See [10](section 7.2, theorem 1) for details. ∎

We observe that the optimal terminal decision rule is a fixed sample size rule and is independent of the stopping rule. Now we find the optimal stopping rule. We truncate the test at some finite time step J in the sense that at most J observations are available for the test. Mathematically, it can be stated as

$$\varphi_J^J(x_1^J) = 1 \qquad \text{or} \qquad \sum_{j=0}^{J} \psi_j^J(x_1^j) = 1,$$

where the superscript indicates that the test is truncated at time step J. We use dynamic programming to find the

optimal truncated stopping rule. Then, we let $J \to \infty$ and obtain the optimal non-truncated stopping rule.

We minimize the truncated risk using backward recursion which is described as follows. Suppose we have taken J measurements. Then, no more measurements can be taken and we incur a risk by using the Bayes terminal decision rule based on (x_1^J, s_1^J). If we reach $J-1$ step, then we continue if the current risk incurred based on (x_1^{J-1}, s_1^{J-1}) is more than the expected risk incurred given (x_1^{J-1}, s_1^{J-1}), of taking one more measurement and then stopping; otherwise we stop at $J-1$ steps. Based on this comparison, we have found the optimal φ_{J-1}^J. Now suppose that using this recursion, we have found $(\varphi_{i+1}^J, \cdots, \varphi_J^J)$. If we have taken i measurements, then we continue if the current risk incurred based on (x_1^i, s_1^i) is more than the expected risk given (x_1^i, s_1^i), of taking one more observation and then using the stopping rule $(\varphi_{i+1}^J, \cdots, \varphi_J^J)$ from there on. In this way, we can find φ_i^J recursively. We now present the analysis formally.

The posteriori probability that θ_1 is the true state of nature given (x_1^j, s_1^j) is denoted by $\pi_{x_1^j, s_1^j}$ and is given by

$$\pi_{x_1^j, s_1^j} = \frac{\pi \prod\limits_{i=1}^{j} f_1^{s_i}(x_i)}{\pi \prod\limits_{i=1}^{j} f_1^{s_i}(x_i) + (1-\pi) \prod\limits_{i=1}^{j} f_0^{s_i}(x_i)}. \quad (3)$$

As stated in Theorem 3.1, the optimal terminal decision rule when the test is stopped is a fixed sample size Bayes rule. For binary hypothesis testing, it can be stated as [10]

$$\hat{\delta}_j(x_1^j, s_1^j) = \begin{cases} \theta_0 & \text{if} \quad \pi_{x_1^j, s_1^j} < \dfrac{w_{01}}{w_{01}+w_{10}}, \\ any & \text{if} \quad \pi_{x_1^j, s_1^j} = \dfrac{w_{01}}{w_{01}+w_{10}}, \\ \theta_1 & \text{otherwise.} \end{cases} \quad (4)$$

Further, let the conditional minimum risk of stopping at time step j and using the Bayes decision rule in (4) based on (x_1^j, s_1^j) be denoted by $U_j(x_1^j, s_1^j; \pi)$. It can be defined as

$$U_j(x_1^j, s_1^j; \pi) = \begin{cases} w_{10} \pi_{x_1^j, s_1^j} + \sum\limits_{i=1}^{j} c_i(s_i) \\ \qquad \text{if} \quad \pi_{x_1^j, s_1^j} < \dfrac{w_{01}}{w_{01}+w_{10}}, \\ w_{01}(1 - \pi_{x_1^j, s_1^j}) + \sum\limits_{i=1}^{j} c_i(s_i) \quad \text{otherwise.} \end{cases} \quad (5)$$

Note that the minimum risk depends explicitly on the a-priori probability π. Further, let $V_j^J(x_1^j, s_1^j; \pi)$ denote the conditional minimum risk based on (x_1^j, s_1^j) and using the optimal stopping rules $(\hat{\varphi}_j^J, \hat{\varphi}_{j+1}^J, \cdots, \hat{\varphi}_J^J)$ for a test truncated at time step J. We use dynamic programming to compute the stopping rule and the conditional minimum risk recursively.

If we have obtained J measurements, then we stop taking new measurements and clearly $V_J^J(x_1^J, s_1^J; \pi) = U_J(x_1^J, s_1^J; \pi)$. If we are at step $J-1$, then we can either stop or continue. If we stop, we incur a risk $U_{J-1}(x_1^{J-1}, s_1^{J-1}; \pi)$. If we continue, the expected risk will be $\mathbb{E}_{X_J, S_J}[U_J(X_1^J, S_1^J; \pi)|(x_1^{J-1}, s_1^{J-1})]$. Thus, comparing

these two quantities, the stopping rule and minimum conditional risk can be calculated. Continuing the recursion, if we have taken $J-2$ measurements, then $\hat{\varphi}_{J-2}^J$ and V_{J-2}^J can be obtained by comparing $U_{J-2}(x_1^{J-2}, s_1^{J-2}; \pi)$ with $\mathbb{E}_{X_{J-1}, S_{J-1}}[V_{J-1}^J(X_1^{J-1}, S_1^{J-1}; \pi)|x_1^{J-2}, s_1^{J-2}]$. In general, the dynamic programming recursion can be described as

$$\hat{\varphi}_j^J(x_1^j, s_1^j) = \begin{cases} 1 & \text{if} \quad U_j(x_1^j, s_1^j; \pi) \\ & \leq \mathbb{E}_{X_{j+1}, S_{j+1}}[V_{j+1}^J(X_1^{j+1}, S_1^{j+1}; \pi)|x_1^j, s_1^j], \\ 0 & \text{otherwise,} \end{cases}$$
$$\qquad\qquad (6)$$

$$V_j^J(x_1^j, s_1^j; \pi) = \min\{U_j(x_1^j, s_1^j; \pi), \quad\quad (7)$$
$$\mathbb{E}_{X_{j+1}, S_{j+1}}[V_{j+1}^J(X_1^{j+1}, S_1^{j+1}; \pi)|x_1^j, s_1^j]\},$$

with $V_J^J(x_1^J, s_1^J; \pi) = U_J(x_1^J, s_1^J; \pi)$. Thus, $V_0^J(\pi)$ denotes the minimum risk of a sequential detection test that is truncated at time step J.

We have found an optimal sequential test for a truncated problem. Now, we intend to increase the truncation step J and obtain the result for an non-truncated problem. First, we show that the minimum unconditional risk sequence $\{V_0^J(\pi)\}_{J=0}^{\infty}$ converges to a finite value.

Lemma 3.2: $V_0^J(\pi)$ is a non increasing function of J

$$V_0^0(\pi) \geq V_0^1(\pi) \geq V_0^2(\pi) \geq V_0^3(\pi) \cdots. \quad (8)$$

Proof: The theorem can be proved using an induction on V_0^k and is omitted due to space constraints. ∎

Since $V_0^J(\pi) \geq 0$, the sequence in (8) converges to a finite value $V_0^\infty(\pi)$.

Theorem 3.3: If the loss function $L(\theta, a)$ is bounded, then $\lim\limits_{J \to \infty} V_0^J(\pi) = V_0^\infty(\pi)$.

Proof: The proof is similar as stated in [11] (theorem 5, chapter 7) and is omitted due to space constraints. ∎

We showed that truncated Bayes risk converges to a finite value. Thus, we know that a non-truncated Bayes test exists. Now, we state an important property of the Bayes risk.

Lemma 3.4: The functions $V_0^J(\pi)$, $V_0^\infty(\pi)$ and $R(\pi, (\hat{\varphi}^J, \hat{\delta}))$ are concave functions of $\pi \in [0, 1]$.

Proof: The proof follows from the fact that each of these functions is a infimum over a class of functions that are linear in π and the definition of concave functions. ∎

The sequential stopping rule in (6) has an abstract structure and is not implementable. Note that the pair (X_k, S_k) are i.i.d. give the hypothesis i with probability density $p_s f_i^s(x_k)$. In remaining section, we show how to use this conditional independence property to reduce the general test to a more tractable sequential test structure. We start by stating the independence property.

Lemma 3.5: If (X_1^n, S_1^n) are i.i.d. with density $p_s f_\theta^s(x)$ and π is the a-priori probability that θ is the true state of nature, then the distribution of (X_{j+1}^n, S_{j+1}^n) given $(X_1^j, S_1^j) = (x_1^j, s_1^j)$ is i.i.d. with density $p_s f_\theta^s(x)$ and $\pi_{x_1^j, s_1^j}$ being the a-posteriori probability that θ is the true state of nature.

Proof: The proof follows from Bayes theorem and is omitted due to space constraints. ∎

The lemma states that when we have obtained the measurements x_1^j from sensors s_1^j, the only thing that changes about

the distribution of future observations is that the *a-priori* hypothesis probability changes from π to $\pi_{x_1^j,s_1^j}$. It suggests that $\pi_{x_1^j,s_1^j}$ should be the sufficient statistics for the stopping rule. We have already shown in (4) that this is indeed a sufficient statistics for the terminal decision rule.

As shown in (6), the stopping rule depends upon the the risks U and V. To show that $\pi_{x_1^j,s_1^j}$ is a sufficient statistic for stopping rule, we show that $U_j(x_1^j,s_1^j;\pi)$ and $V_j^J(x_1^j,s_1^j;\pi)$ depend on (x_1^j,s_1^j) only through $\pi_{x_1^j,s_1^j}$ and a common cost term. Following lemma states this result and is a generalization of the corresponding lemma in [10].

Lemma 3.6: 1) $\pi_{x_1^j,s_1^j|(x_{j+1}^k,s_{j+1}^k)} = \pi_{x_1^k,s_1^k}$.

2) $U_j(x_1^j,s_1^j;\pi) = U_{j-1}(x_2^j,s_2^j;\pi_{x_1,s_1}) + c_1(s_1) =$
$\cdots = U_0(\pi_{x_1^j,s_1^j}) + \sum_{i=1}^{j} c_i(s_i)$.

3) $V_j^J(x_1^j,s_1^j;\pi) = V_{j-1}^{J-1}(x_2^j,s_2^j;\pi_{x_1,s_1}) + c_1(s_1) =$
$\cdots = V_0^{J-j}(\pi_{x_1^j,s_1^j}) + \sum_{i=1}^{j} c_i(s_i)$.

Proof:
1) Follows using the definition of *a-posteriori* probability in (3).
2) This part follows directly by using the definition of $U_j(x_1^j,s_1^j;\pi)$ in (5) and the result of part 1.
3) The result can be proved using induction on the variable $k = J - j$. For details, see [10](section 7.5, lemma 2)

∎

The following lemma presents a recursive expression to calculate V_0^J.

Lemma 3.7:

$$V_0^{J+1}(\pi) = \min\{U_0(\pi), \bar{d} + \mathbb{E}_{X_1,S_1}[V_0^J(\pi_{X_1,S_1})]\}.$$

Proof: The lemma can be proved using the definition of V_j^J in (7) and part 3 of lemma 3.6.

$$V_0^{J+1}(\pi) = \min\{U_0(\pi), \mathbb{E}_{X_1,S_1}[V_1^{J+1}(X_1,S_1;\pi)]\}$$
$$= \min\{U_0(\pi), \mathbb{E}_{X_1,S_1}[V_0^J(\pi_{X_1,S_1}) + C_1(S_1)]\}$$
$$= \min\{U_0(\pi), \bar{d} + \mathbb{E}_{X_1,S_1}[V_0^J(\pi_{X_1,S_1})]\}.$$

∎

We now use these results to present an alternate form of stopping rule that explicitly depends on the *a-posteriori* probability $\pi_{x_1^j,s_1^j}$.

Let Ω^J denote the set of all prior distributions π for which the test is stopped without taking any observations. Using lemma 3.7, we have

$$\Omega^J \triangleq \{\pi : V_0^J(\pi) = U_0(\pi)\}. \qquad (9)$$

Theorem 3.8: The optimum stopping rule $\hat{\varphi}^J = (\hat{\varphi}_0^J, \hat{\varphi}_1^J(x_1,s_1), \hat{\varphi}_2^J(x_1^2,s_1^2), \cdots, \hat{\varphi}_J^J(x_1^J,s_1^J))$ for a problem truncated at J is given by

$$\hat{\varphi}_j^J(x_1^j,s_1^j) = \begin{cases} 1 & \text{if} \quad \pi_{x_1^j,s_1^j} \in \Omega^{J-j} \\ 0 & \text{otherwise.} \end{cases} \qquad (10)$$

Proof: From (6) and (7), the test is stopped at j if

$$V_j^J(x_1^j,s_1^j;\pi) = U_j(x_1^j,s_1^j;\pi).$$

From lemma 3.6 the above condition is equivalent to

$$V_0^{J-j}(\pi_{x_1^j,s_1^j}) = U_0(\pi_{x_1^j,s_1^j}),$$

and the theorem follows from definition (9). ∎

Thus, we see that the truncated stopping rule depends only on $\pi_{x_1^j,s_1^j}$. Now we generalize this result to obtain a non-truncated stopping rule. Since $V_0^J(\pi)$ are non-decreasing (lemma 3.2), the sets Ω^J satisfy

$$\Omega^0 \supset \Omega^1 \supset \Omega^2 \supset \cdots.$$

Further, since $\lim_{J \to \infty} V_0^J(\pi) = V_0^\infty(\pi)$, from definition (9) we have

$$\lim_{J \to \infty} \Omega^J = \Omega^\infty = \{\pi : V_0^\infty(\pi) = U_0(\pi)\}. \qquad (11)$$

Also, lemma 3.7 can be generalized as

$$V_0^\infty(\pi) = \min\{U_0(\pi), \bar{d} + \mathbb{E}_{X_1,S_1}[V_0^\infty(\pi_{X_1,S_1})]\}. \qquad (12)$$

Equation (12) represents the fundamental equation of dynamic programming for the sequential detection problem for the randomized sensor selection.

As a result, the non-truncated stopping rule becomes

$$\hat{\varphi}_j(x_1^j,s_1^j) = \begin{cases} 1 & \text{if} \quad \pi_{x_1^j,s_1^j} \in \Omega^\infty \\ 0 & \text{otherwise.} \end{cases} \qquad (13)$$

As we expected, the stopping rule depends only on whether the posteriori probability $\pi_{x_1^j,s_1^j}$ is in the fixed set Ω^∞ or not. For a rigorous treatment of the limiting stopping rule, see [10].

Although we have presented the optimal stopping rule explicitly in terms of $\pi_{x_1^j,s_1^j}$, we only have an abstract characterization of the set Ω^∞ in (11). To implement the sequential test, we need to explicitly find this set of prior distributions. In the next section, we show that the set Ω^∞ can be characterized by two thresholds and the optimal Bayesian sequential test corresponds to an SPRT.

IV. OPTIMAL SEQUENTIAL TEST AS AN SPRT

In this section, we show that the optimal sequential test developed in the previous section can be reduced to SPRT which can be easily implemented. We proceed by finding the set Ω^∞ explicitly. Let the term

$$W(\pi) \triangleq \bar{d} + \mathbb{E}_{X_1,S_1}[V_0^\infty(\pi_{X_1,S_1})]$$

represent the minimum risk over all the sequential tests that take at least one observation. Using this definition and (12), the set Ω^∞ can be equivalently defined as

$$\Omega^\infty = \{\pi : U_0(\pi) \leq W(\pi)\}.$$

To characterize this set, we need to first characterize the functions $U_0(\pi)$ and $W(\pi)$. The function $U_0(\pi)$ can be easily obtained through (5) and consists of two linear parts

$$U_0(\pi) = \begin{cases} w_{10}\pi & \text{if} \quad \pi < \dfrac{w_{01}}{w_{01} + w_{10}}, \\ w_{01}(1 - \pi) & \text{otherwise.} \end{cases} \qquad (14)$$

Next, we proceed by stating some properties of $W(\pi)$.

Lemma 4.1: The function $W(\pi)$ is continuous and concave in $[0,1]$ and $W(0) = W(1) = \bar{d}$.

Proof: The proof of concavity of $W(\pi)$ is similar to that of lemma 3.4. Continuity of $W(\pi)$ follows from its concavity. Further, from (12) and (14) we have $0 \leq V_0^\infty(\pi) \leq U_0(\pi) \leq w_{10}\pi$. Using this relation, we obtain

$$0 \leq \mathbb{E}_{X_1,S_1}[V_0^\infty(\pi_{X_1,S_1})] = W(\pi) - \bar{d}$$
$$\leq w_{10}\mathbb{E}_{X_1,S_1}[\pi_{X_1,S_1}] = w_{01}\pi.$$

Thus, as $\pi \to 0$, $W(\pi) \to \bar{d}$. By symmetry, same arguments can be made for $\pi \to 1$. ∎

Since, $W(\pi)$ is concave, continuous and $W(0) > U_0(0)$, the equation $W(\pi) = w_{10}\pi$ has at most one solution in the interval $[0, \frac{w_{01}}{w_{01}+w_{10}}]$. Denote π_L as the solution of this equation if it exists. It a solution does not exist, then define π_L as $\frac{w_{01}}{w_{01}+w_{10}}$. Similarly, π_U is defined as the solution of the equation $W(\pi) = w_{01}(1-\pi)$ in the interval $[\frac{w_{01}}{w_{01}+w_{10}}, 1]$ if it exists, otherwise it is defined as $\frac{w_{01}}{w_{01}+w_{10}}$. From the preceding arguments, it follows that

$$0 < \pi_L \leq \frac{w_{01}}{w_{01}+w_{10}} \leq \pi_U < 1, \tag{15}$$

and the set Ω^∞ can be specified by the two thresholds as

$$\Omega^\infty = [0, \pi_L] \cup [\pi_U, 1]. \tag{16}$$

A. Sequential Probability Ratio Test

We define the standard Sequential Probability Ratio Test (SPRT)[8] as a particular class of sequential decision tests. Let the likelihood ratio based on measurements x_1^j and sensor s_1^j be defined as

$$L_j(x_1^j, s_1^j) \triangleq \prod_{i=1}^{j} \frac{p_{s_i} f_1^{s_i}(x_i)}{p_{s_i} f_0^{s_i}(x_i)} = \prod_{i=1}^{j} \frac{f_1^{s_i}(x_i)}{f_0^{s_i}(x_i)}. \tag{17}$$

Then, SPRT with thresholds A and B, denoted by $SPRT(A,B)$ with $0 < A \leq 1 \leq B < \infty$ is defined by the stopping rule φ and terminal decision rule δ

$$\varphi_j(x_1^j, s_1^j) = \begin{cases} 0 & \text{if} \quad A < L_j(x_1^j, s_1^j) < B, \\ 1 & \text{otherwise.} \end{cases} \tag{18}$$

$$\delta_j(x_1^j, s_1^j) = \begin{cases} \theta_1 & \text{if} \quad L_j(x_1^j, s_1^j) \geq B, \\ \theta_0 & \text{if} \quad L_j(x_1^j, s_1^j) \leq A. \end{cases} \tag{19}$$

Thus, the test continues till the likelihood ratio stays between the two thresholds. When the test is stopped, H_1 is accepted if $L_j \geq B$ and H_0 is accepted if $L_j \leq A$. We now show the equivalence between the optimal sequential Bayes test and SPRT.

Theorem 4.2: [10] The optimal sequential detection rule $(\hat{\varphi}, \hat{\delta})$ as given in (13) and (4) for $\pi_L \leq \pi \leq \pi_U$ is equivalent to $SPRT(A,B)$, where

$$A = \frac{(1-\pi)\pi_L}{\pi(1-\pi_L)} \quad \text{and} \quad B = \frac{(1-\pi)\pi_U}{\pi(1-\pi_U)}. \tag{20}$$

Proof: If $\pi_L \leq \pi \leq \pi_U$, then (A, B) satisfy $0 < A \leq 1 \leq B < \infty$. Using (16), the stopping rule in (13) can be written as

$$\hat{\varphi}_j(x_1^j, s_1^j) = \begin{cases} 0 & \text{if} \quad \pi_L < \pi_{x_1^j, s_1^j} < \pi_U, \\ 1 & \text{otherwise.} \end{cases}$$

Further, (3) can be rewritten as

$$\pi_{x_1^j, s_1^j} = \frac{\pi L_j(x_1^j, s_1^j)}{\pi L_j(x_1^j, s_1^j) + (1-\pi)}.$$

Using the above relations and (20), it is easy to see that the relation $\pi_L < \pi_{x_1^j, s_1^j} < \pi_U$ is equivalent to $A < L_j(x_1^j, s_1^j) < B$. Thus the stopping rules in (13) and (18) are identical.

Further, using the relation in (15), it is easy to see that the terminal decision rules in (4) and (19) are identical. ∎

The equivalence of the optimal sequential detection rule $(\hat{\varphi}, \hat{\delta})$ and $SPRT(A,B)$ allows a simple tractable form of the test in terms of the thresholds. The thresholds $\pi_L(A)$ and $\pi_U(B)$ depend upon the w_{01}, w_{10}, the sensor costs $d = (d_1, \cdots, d_K)$ and the sensor selection probabilities $p = (p_1, \cdots, p_K)$. However, the closed form expression of π_L and π_U in terms of these quantities is not tractable. To overcome this problem, we observe that the standard Wald's approximations that approximate the thresholds in terms of the error probabilities, also hold true for the randomized sensor selection sequential test. Thus, we have

$$A \approx \frac{\alpha_1}{1-\alpha_0} \quad \text{and} \quad B \approx \frac{1-\alpha_1}{\alpha_0}. \tag{21}$$

where, $\alpha_0 = P(\text{accept } H_1 | H_0 \text{ is true})$ denote the probability of false detection and $\alpha_1 = P(\text{accept } H_0 | H_1 \text{ is true})$ denote the probability of miss. It has been shown that for large number of observations (low error probabilities), the above approximations are good.

We now show that the standard Wald-Wolfowitz inequality [9], which establishes the optimality of SPRT also holds true for the proposed randomized sensor selection sequential test. We extend the ideas presented in [10] to establish this result. We rewrite the conditional risks as

$$R(\theta_j, (\hat{\varphi}, \hat{\delta})) = \mathbb{E}\left[L(\theta_j, \hat{\delta}_N(X_1^N, S_1^N))\right]$$
$$+ \mathbb{E}\left[\sum_{i=1}^{N} C_i(S_i)|(H_j, \hat{\varphi})\right], \quad j = 0, 1$$

where N is the random stopping time. Since $C_i(S_i)$ is an i.i.d. random sequence and N depends on $C_1^j(S_1^j)$ and is independent of $C_{j+1}^\infty(S_{j+1}^\infty)$, we have [12]

$$\mathbb{E}\left[\sum_{i=1}^{N} C_i(S_i)|(H_j, \hat{\varphi})\right] = \mathbb{E}[C_i(S_i)]\mathbb{E}[N|(H_j, \hat{\varphi})]$$
$$= \bar{d}\,\mathbb{E}[N|(H_j, \hat{\varphi})] \qquad j = 0, 1.$$

Thus, the conditional risks can be rewritten as

$$R(\theta_0, (\hat{\varphi}, \hat{\delta})) = w_{01}\alpha_0(\hat{\varphi}, \hat{\delta}) + \bar{d}\,\mathbb{E}[N|(H_0, \hat{\varphi})], \tag{22}$$
$$R(\theta_1, (\hat{\varphi}, \hat{\delta})) = w_{10}\alpha_1(\hat{\varphi}, \hat{\delta}) + \bar{d}\,\mathbb{E}[N|(H_1, \hat{\varphi})]. \tag{23}$$

Now we state the Wald-Wolfowitz theorem.

Theorem 4.3: Let $(\hat{\varphi}, \hat{\delta})$ be SPRT(A, B) and (φ, δ) be any other sequential rule with random sensor selection for which $\alpha_0(\varphi, \delta) \leq \alpha_0(\hat{\varphi}, \hat{\delta})$ and $\alpha_1(\varphi, \delta) \leq \alpha_1(\hat{\varphi}, \hat{\delta})$. Then,

$$\mathbb{E}[N|(H_0, \varphi)] \geq \mathbb{E}[N|(H_0, \hat{\varphi})], \qquad \text{and}$$
$$\mathbb{E}[N|(H_1, \varphi)] \geq \mathbb{E}[N|(H_1, \hat{\varphi})].$$

This result is intuitive and follows from the Bayesian optimality of the $(\hat{\varphi}, \hat{\delta})$ which is equivalent to SPRT(A, B) as shown in theorem 4.2. To formally prove the statement, we denote the explicit dependence of the thresholds and $W(\pi)$ on $w = (w_{01}, w_{10}), p, d$ by $\pi_L(w, p, d)$, $\pi_U(w, p, d)$ and $W(\pi; w, p, d)$. To prove the theorem, we first prove the following properties of the thresholds. These results are the same as presented in [10].

Lemma 4.4: For fixed w and p, $\pi_L(w, p, d)$ and $\pi_U(w, p, d)$ are continuous functions of d, and $\pi_L(w, p, d) \rightarrow 0$ and $\pi_U(w, p, d) \rightarrow 1$ as $d \rightarrow 0$.

Proof: The proof is similar to that presented in [10](section 7.6, lemma 3) and is omitted due to space constraints. ∎

Using the properties in the above lemma, we state the relation between the thresholds and the test parameters π, w, p, d.

Lemma 4.5: For a given $\epsilon > 0$ and $0 < A \leq 1 \leq B < \infty$, there exist

1) π, w, p, d with $0 < \pi < \epsilon$ such that (20) holds true,
2) π', w, p, d with $1 - \epsilon < \pi' < 1$ s.t. (20) holds true.

Proof: See [10](section 7.6, lemma 4) for details. ∎

We now state the proof of theorem 4.3.

Proof: (*of Theorem 4.3*) Using lemma 4.5, we can find π, w and d such that (20) is satisfied and $\pi < \epsilon$. Since $(\hat{\varphi}, \hat{\delta})$ is equivalent to $SPRT(A, B)$ (from theorem 4.2) and is the optimal sequential rule, we have

$$0 \leq R(\pi, (\varphi, \delta)) - R(\pi, (\hat{\varphi}, \hat{\delta}))$$
$$= \pi w_{10}(\alpha_1(\varphi, \delta) - \alpha_1(\hat{\varphi}, \hat{\delta}))$$
$$+ (1 - \pi)w_{01}(\alpha_1(\varphi, \delta) - \alpha_1(\hat{\varphi}, \hat{\delta}))$$
$$+ \pi \bar{d}(\mathbb{E}[N|(H_1, \varphi)] - \mathbb{E}[N|(H_1, \hat{\varphi})])$$
$$+ (1 - \pi)\bar{d}(\mathbb{E}[N|(H_0, \varphi)] - \mathbb{E}[N|(H_0, \hat{\varphi})])$$
$$\leq \pi \bar{d}(\mathbb{E}[N|(H_1, \varphi)] - \mathbb{E}[N|(H_1, \hat{\varphi})])$$
$$+ (1 - \pi)\bar{d}(\mathbb{E}[N|(H_0, \varphi)] - \mathbb{E}[N|(H_0, \hat{\varphi})]).$$

Since the above statement is valid for π arbitrarily close to zero, we have $\mathbb{E}[N|(H_0, \varphi)] \geq \mathbb{E}[N|(H_0, \hat{\varphi})]$. Further, using symmetry and part 2 of lemma 4.5, we have $\mathbb{E}[N|(H_1, \varphi)] \geq \mathbb{E}[N|(H_1, \hat{\varphi})]$, thus completing the proof. ∎

Thus, the stationarity of the probabilistic sensor selection process results in SPRT being the optimal test, as it was the case with single sensor case. With multiple sensors, we have more degree of freedom to choose the sensors and improve the test performance.

V. CONCLUSION

We have obtained an optimal test for binary hypothesis testing using multiple sensors. When a single sensor is chosen randomly at each time step with a stationary distribution, the optimal test reduces to SPRT. Further, we prove that Wald-Wolfowitz theorem can be extended for the test involving multiple sensors. We plan to extend the study by characterizing the average stopping time and the experiment cost of the sequential test and optimizing it over the sensor selection distribution. Thus, we plan to obtain a sensor selection strategy that optimally balances the trade-off between the sensor costs and the sensor performance.

REFERENCES

[1] V. Gupta, T. H. Chung, B. Hassibi and R. M. Murray, "On a Stochastic Sensor Selection Algorithm with Applications in Sensor Scheduling and Sensor Coverage," Automatica, vol. 42, no. 2, pp. 251-260, 2006.
[2] V. V. Veeravalli, T. Basar and H. V. Poor, "Decentralized Sequential Detection with a Fusion Center Performing the Sequential Test," IEEE Transactions on Information Theory, vol. 39, no. 2, pp. 433-442, 1993.
[3] Y. Mo, R. Ambrosino, B. Sinopoli, "Sensor Selection Strategies for State Estimation in Energy Constrained Wireless Sensor Networks," Automatica, vol. 47, no. 7, pp. 1330-1338, 2011.
[4] S. Joshi and S. Boyd, "Sensor Selection via Convex Optimization," IEEE Transactions on Signal Processing, vol. 57, no. 2, pp. 451-462, 2009.
[5] Z. Chair and P. K. Varshney, "Optimal data fusion in multiple sensor detection systems," IEEE Trans. Aerosp. Electron. Syst., vol. AES-22, pp. 98-101, 1986.
[6] C. W. Baum and V. V. Veeravalli, "A Sequential Procedure for Multihypothesis Testing," IEEE Transactions on Information Theory, vol. 40, no. 6, pp. 1994-2007, 1994.
[7] V. Srivastava, K. Plarre and F. Bullo, "Randomized Sensor Selection in Sequential Hypothesis Testing," IEEE Transactions on Signal Processing, vol. 59, issue 5, pp. 2342 - 2354, 2011.
[8] A. Wald, "Sequential Analysis," Wiley, New York, 1947.
[9] A. Wald and J. Wolfowitz, "Optimum Character of the Sequential Probability Ratio Test," Annals of Mathematical Statistics, vol. 19, no. 3, pp. 326-339, 1948.
[10] T. S. Ferguson, "Mathematical Statistics: A Decision Theoretic Approach," Academic Press, New York, 1967.
[11] J. O. Berger, "Statistical Decision Theory and Bayesian Analysis," Springer, 2nd Edition, 1985.
[12] D. Siegmund, "Sequential Analysis," Springer-Verlag, New York, 1985.
[13] C. K. Sestok, M. R. Said and A. V. Oppenheim, "Randomized Data Selection in Detection with Applications to Distributed Signal Processing," Proceedings of IEEE, vol. 91, issue 8, pp. 1184 - 1198, 2003.
[14] Y. Wang, I. I. Hussein and R. S. Erwin, "Risk-based Sensor Management for Integrated Detection and Estimation," American Control Conference, pp. 3633 - 3638, 2011.