

Adaptive Sensor Selection in Sequential Hypothesis Testing

Vaibhav Srivastava

Kurt Plarre

Francesco Bullo

Abstract—We consider the problem of sensor selection for time-optimal detection of a hypothesis. We consider a group of sensors transmitting their observations to a fusion center. The fusion center considers the output of only one randomly chosen sensor at the time, and performs a sequential hypothesis test. We study sequential probability ratio test with randomized sensor selection strategy. We present optimal open loop sensor selection policies for three distinct cost functions. We utilize these policies to develop an adaptive sensor selection policy. We rigorously characterize the performance of the adaptive policy and show that it asymptotically achieves the performance of the globally optimal policy.

I. INTRODUCTION

The advent of cheap sensors has resulted in their extensive deployment. This makes plethora of information available for decision making. Such a data overload diminishes the efficiency of the operator processing this information and is often the root cause for missing critical information [10]. This calls for development of policies that help the operator focus her attention to the most pertinent information.

In large scale sensor networks, it may not be energy efficient to have all the sensors activated all the time. An energy efficient technique is to activate the most pertinent sensors and censor the remaining sensors. This needs a characterization of the most pertinent sensors.

In this paper, we consider the problem of quickest decision making and the Sequential Probability Ratio Test (SPRT). Recent advances in cognitive psychology [2] show that human performance in decision making tasks, such as the "two-alternative forced choice task," is well modeled by a drift diffusion process, i.e., by the continuous-time version of SPRT. Roughly speaking, modeling decision making as an SPRT process may be appropriate even for situations in which a human is making the decision, e.g., an operator processing the feeds of a camera network [3].

Recent years have witnessed a significant interest in the problem of sensor selection for optimal detection and estimation. Tay et al [13] discuss the problem of censoring sensors for decentralized binary detection. They assess the quality of sensor data by the Neyman-Pearson and a Bayesian binary hypothesis test and decide on which sensors should transmit their observation at that time instant. Gupta et al [6] focus on stochastic sensor selection and minimize the error covariance of a process estimation problem. Isler et al [7] propose

geometric sensor selection schemes for error minimization in target detection. Debouk et al [5] formulate a Markovian decision problem to ascertain some property in a dynamical system, and choose sensors to minimize the associated cost. Williams et al [15] use an approximate dynamic program over a rolling time horizon to pick a sensor-set that optimizes the information-communication trade-off. Zhao et al [16] study information driven sensor collaboration. They study information gain and cost trade-off in target tracking through directed diffusion routing. Wang et al [14] design entropy-based sensor selection algorithms for target localization. Joshi et al [8] present a convex optimization-based heuristic to select multiple sensors for optimal parameter estimation. Bajović et al [1] discuss sensor selection problems for Neyman-Pearson binary hypothesis testing in wireless sensor networks. Castañón [4] study an iterative search problem as a hypothesis testing problem over a fixed horizon.

In this paper, we analyze the problem of time-optimal sequential decision making in the presence of multiple switching sensors and determine a sensor selection strategy to achieve the same. We consider a sensor network where all sensors are connected to a fusion center. The fusion center, at each instant, receives information from only one sensor and implements the SPRT with the gathered information. The sensors may be heterogeneous (e.g., a camera sensor, a sonar sensor, a radar sensor, etc), hence, the time needed to collect, transmit, and process data may differ significantly for these sensors. We extend our previous work [12] where we develop a version of the SPRT algorithm in which the sensor is randomly switched at each iteration, and determine optimal open loop sensor selection policies that minimize the expected decision time. Here, we develop closed loop sensor selection policies where we learn the true state of the nature and adapt our sensor selection policy accordingly. We rigorously analyze this closed loop policy and show that it is asymptotically optimal. The major contributions of this paper are:

- i) We present a version of the SPRT where the sensor is selected randomly at each observation.
- ii) We present optimal open loop policies for this version of SPRT for three distinct cost functions.
- iii) We develop an adaptive sensor selection policy and rigorously characterize its performance.
- iv) We provide asymptotic bounds on performance of SPRT with multiple sensors under any sensor selection policy.
- v) We show that our adaptive sensor selection policy is asymptotically optimal.

The remainder of the paper is organized in following

This work has been supported in part by AFOSR MURI Award-FA9550-07-1-0528 and by NSF Award CPS-1035917.

Vaibhav Srivastava and Francesco Bullo are with the Center for Control, Dynamical Systems, and Computation, University of California, Santa Barbara, Santa Barbara, CA 93106, USA, {vaibhav,bullo}@engineering.ucsb.edu

Kurt Plarre is with the Department of Computer Science, University of Memphis, Memphis, TN 38152, USA, kplarre@memphis.edu

way. Some preliminaries are presented in Section II. In Section III, we present the problem setup. We present the SPRT procedure with randomized sensor selection strategy in Section IV. In Section V, we present open loop policies for time-optimal sensor selection. An adaptive sensor selection policy is presented in Section VI. We elucidate the results obtained through some numerical examples in Section VII. Our concluding remarks are in Section VIII.

II. PRELIMINARIES

A. Kullback-Leibler divergence

Given two probability mass functions $f_1 : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$ and $f_2 : \mathcal{S} \rightarrow \mathbb{R}_{\geq 0}$, where \mathcal{S} is some countable set, the Kullback-Leibler divergence $\mathcal{D} : \mathcal{L}^1 \times \mathcal{L}^1 \rightarrow \mathbb{R} \cup \{+\infty\}$ is defined by

$$\mathcal{D}(f_1, f_2) = \mathbb{E}_{f_1} \left[\log \frac{f_1(X)}{f_2(X)} \right] = \sum_{x \in \text{supp}(f_1)} f_1(x) \log \frac{f_1(x)}{f_2(x)},$$

where $\text{supp}(f_1)$ is the support of f_1 . It is known that $0 \leq \mathcal{D}(f_1, f_2) \leq +\infty$, that the lower bound is achieved if and only if $f_1 = f_2$, almost surely, and that the upper bound is achieved if and only if the support of f_2 is a strict subset of the support of f_1 . Note that equivalent statements can be given for probability distribution functions.

B. Sequential Probability Ratio Test

Consider two hypotheses H_0 and H_1 with their conditional probability distribution functions $f^0(y) := f(y|H_0)$ and $f^1(y) := f(y|H_1)$. Given repeated measurements $\{y_1, y_2, \dots\}$, the SPRT algorithm collects evidence about the hypotheses and compares the integrated evidence to two thresholds η_0 and η_1 in order to decide on a hypothesis. The SPRT procedure is presented in Algorithm 1.

Algorithm 1 Sequential Probability Ratio Test

- 1: at time $\tau \in \mathbb{N}$, collect sample y_τ
 - 2: compute the log likelihood ratio $\lambda_\tau := \log \frac{f^1(y_\tau)}{f^0(y_\tau)}$
 - 3: integrate evidence up to current time $\Lambda_\tau := \sum_{t=1}^{\tau} \lambda_t$
% decide only if a threshold is crossed
 - 4: **if** $\Lambda_\tau > \eta_1$, **then** accept H_1
 - 5: **if** $\Lambda_\tau < \eta_0$, **then** accept H_0
 - 6: **else** ($\Lambda_\tau \in]\eta_0, \eta_1[$) continue sampling (step 1:)
-

Given the probability of false alarm $\mathbb{P}(H_1|H_0) = \alpha_0$ and probability of missed detection $\mathbb{P}(H_0|H_1) = \alpha_1$, the Wald's thresholds η_0 and η_1 are defined by

$$\eta_0 = \log \frac{\alpha_1}{1 - \alpha_0}, \quad \text{and} \quad \eta_1 = \log \frac{1 - \alpha_1}{\alpha_0}. \quad (1)$$

Let N_d denote the number of samples required for decision using SPRT. Its expected value is approximately [11] given by

$$\begin{aligned} \mathbb{E}[N_d|H_0] &\cong -\frac{(1 - \alpha_0)\eta_0 + \alpha_0\eta_1}{\mathcal{D}(f^0, f^1)}, \quad \text{and} \\ \mathbb{E}[N_d|H_1] &\cong \frac{\alpha_1\eta_0 + (1 - \alpha_1)\eta_1}{\mathcal{D}(f^1, f^0)}. \end{aligned} \quad (2)$$

The approximations in equation (2) are referred to as the Wald's approximations [11]. It is known that these expressions are accurate for large thresholds and small error probabilities.

III. PROBLEM SETUP

We consider a group of n agents (e.g., robots, sensors, or cameras), which take measurements and transmit them to a fusion center. We generically call these agents "sensors." We identify the fusion center with a person supervising the agents, and call it the "supervisor." The goal of the supervisor is to decide, based on the measurements it receives, which one of two alternative hypotheses or "states of nature" is correct. To do so, the supervisor implements the SPRT algorithm with the collected observations. Given pre-specified accuracy thresholds, the supervisor aims to make a decision in minimum time.




Fig. 1. The agents A transmit their observation to the supervisor S , one at the time. The supervisor performs SPRT to decide on the underlying hypothesis.

We assume that there are more agents than hypotheses (i.e., $n > 2$), and that only one sensor can transmit to the supervisor at each (discrete) time instant. Equivalently, the supervisor can process data from only one of the n agents at each time. Thus, at each time, the supervisor must decide which sensor should transmit its measurement. We aim to determine the optimal sensor(s) that the supervisor must observe in order to minimize the decision time.

We adopt the following notation. Let $\{H_0, H_1\}$ denote the two hypotheses. The time required by sensor $s \in \{1, \dots, n\}$ to collect, process and transmit its measurement is a random variable $T_s \in \mathbb{R}_{>0}$, with finite first and second moments. We denote the expected value of processing time of sensor s by \bar{T}_s . Let $s_t \in \{1, \dots, n\}$ indicate which sensor transmits its measurement at time $t \in \mathbb{N}$. The measurement of sensor s at time t is $y(t, s)$. For the sake of convenience, we denote $y(t, s_t)$ by y_t . For $k \in \{0, 1\}$, let $f_s^k(y)$ denote the probability that the measurement at sensor s is y conditioned on the hypothesis H_k , and α_k denote the desired bound on the probability of incorrect decision conditioned on hypothesis H_k . We make the following standard assumption:

Conditionally-independent observations: Conditioned on hypothesis H_k , the measurement $y(t, s)$ is independent of $y(\bar{t}, \bar{s})$, for $(t, s) \neq (\bar{t}, \bar{s})$.

We adopt a *randomized strategy* in which the supervisor chooses a sensor randomly at each time instant; the probability to choose sensor s is stationary and given by q_s , for $s \in \{1, \dots, n\}$. Also, the supervisor uses the data collected from the randomized sensors to execute the SPRT. We study our proposed randomized strategy under the following assumptions about the sensors.

Distinct sensors: There are no two sensors with identical conditioned probability distribution $f_s^k(y)$ and processing time \bar{T}_s . (If there are such sensors, we club them together in a single node, and distribute the probability assigned to that node equally among them.)

Finitely-informative sensors: Each sensor $s \in \{1, \dots, n\}$ has the following property: for hypotheses $k, j \in \{0, 1\}$, $k \neq j$

- i) the support of f_s^k is equal to the support of f_s^j ,
- ii) $f_s^k \neq f_s^j$ almost surely in f_s^k , and
- iii) conditioned on hypothesis H_k , the first and second moment of $\log(f_s^k(Y)/f_s^j(Y))$ are finite.

Remark 1: The finitely-informative sensors assumption is equivalently restated as follows: each sensor $s \in \{1, \dots, n\}$ satisfies $0 < \mathcal{D}(f_s^0, f_s^1), \mathcal{D}(f_s^1, f_s^0) < +\infty$. \square

IV. SPRT WITH RANDOMIZED SENSOR SELECTION

We call the SPRT with the data collected from n sensors while observing only one randomly chosen sensor at a time as the SPRT with randomized sensor selection. The sensor to be observed at each time is sampled from a stationary distribution, and the probability of choosing sensor s is q_s . The SPRT with randomized sensor selection is defined identically to the Algorithm 1, where the first two instructions (steps 1: and 2:) are replaced by:

- 1: at time $\tau \in \mathbb{N}$, select a random sensor s_τ according to the probability vector q and collect a sample y_{s_τ}
- 2: compute the log likelihood ratio $\lambda_\tau := \log \frac{f_{s_\tau}^1(y_\tau)}{f_{s_\tau}^0(y_\tau)}$

Theorem 2 (SPRT with randomized sensor selection):

Assume finitely-informative sensors $\{1, \dots, n\}$, and independent observations conditioned on hypotheses H_k , $k \in \{0, 1\}$. For the SPRT with randomized sensor selection the following statements hold:

- i) Conditioned on a hypothesis, the sample size for decision N_d is finite almost surely.
- ii) The sample size for decision N_d satisfies

$$\begin{aligned} \mathbb{E}[N_d|H_0] &\cong -\frac{(1-\alpha_0)\eta_0 + \alpha_0\eta_1}{\sum_{s=1}^n q_s \mathcal{D}(f_s^0, f_s^1)}, \quad \text{and} \\ \mathbb{E}[N_d|H_1] &\cong \frac{\alpha_1\eta_0 + (1-\alpha_1)\eta_1}{\sum_{s=1}^n q_s \mathcal{D}(f_s^1, f_s^0)}. \end{aligned} \quad (3)$$

- iii) The decision time T_d satisfies

$$\mathbb{E}[T_d|H_k] \cong \mathbb{E}[N_d|H_k] \sum_{s=1}^n q_s \bar{T}_s. \quad (4)$$

Proof: Conditioned on hypothesis H_k , the observation y_t is sampled from the distribution f_s^k with probability q_s , $s \in \{1, \dots, n\}$. Therefore, the pairs $\{(s_t, y_t)\}_{t \in \mathbb{N}}$ are i.i.d. Further, $\mathbb{E}[\lambda_t|H_0] = -\sum_{s=1}^n q_s \mathcal{D}(f_s^0, f_s^1)$, and $\mathbb{E}[\lambda_t|H_1] = \sum_{s=1}^n q_s \mathcal{D}(f_s^1, f_s^0)$. The remaining proof of the first two statements follow similar to the proof of SPRT in [11].

Let $\{T_1, \dots, T_{N_d}\}$ be the sequence of the processing times at each iteration. Note that the decision time $T_d = \sum_{\ell=1}^{N_d} T_\ell$. It follows from Wald's identity [9] that $\mathbb{E}[T_d|H_k] = \mathbb{E}[N_d|H_k]\mathbb{E}[T]$. This completes the proof. \blacksquare

V. STATIONARY SENSOR SELECTION

In this section we consider sensor selection problems with the aim to minimize the expected decision time of the SPRT with randomized sensor selection. As exemplified in Theorem 2, the problem features multiple conditioned

decision times and, therefore, multiple distinct cost functions are of interest. In Scenario I below, we aim to minimize the decision time conditioned upon one specific hypothesis being true; in Scenarios II and III we will consider worst-case and average decision times. In all three scenarios the decision variables take values in the probability simplex. We state the results without proofs. The interested reader may refer to [12] for detailed proofs.

Before we pose the problem of optimal sensor selection, we introduce the following notation. We denote the probability simplex in \mathbb{R}^n by Δ_{n-1} , and the vertices of the probability simplex Δ_{n-1} by e_i , $i \in \{1, \dots, n\}$. We refer to the line joining any two vertices of the simplex as an *edge*. Finally, we define the conditional expected decision time $g^k : \Delta_{n-1} \rightarrow \mathbb{R}$, $k \in \{0, 1\}$, by $g^k(\mathbf{q}) = \mathbf{q} \cdot \mathbf{T} / \mathbf{q} \cdot \mathbf{I}^k$, where $\mathbf{T} \in \mathbb{R}_{>0}^n$ is the array of expected processing times, $\mathbf{I}^0 \in \mathbb{R}_{>0}^n$ is an array with elements $I_s^0 = -\mathcal{D}(f_s^0, f_s^1) / ((1-\alpha_0)\eta_0 + \alpha_0\eta_1)$, and $\mathbf{I}^1 \in \mathbb{R}_{>0}^n$ is an array with elements $I_s^1 = \mathcal{D}(f_s^1, f_s^0) / (\alpha_1\eta_0 + (1-\alpha_1)\eta_1)$.

A. Scenario I (Optimization of conditioned decision time):

We consider the case when the supervisor is trying to validate the true hypothesis. The corresponding optimization problem for a fixed $k \in \{0, 1\}$ is:

$$\underset{\mathbf{q} \in \Delta_{n-1}}{\text{minimize}} \quad g^k(\mathbf{q}). \quad (5)$$

The conditional expected decision time g^k is a quasi-linear function and achieves its minimum at some vertex of the feasible simplex. This is rigorously characterized in the following theorem:

Theorem 3 (Optimization of conditioned decision time):

The solution to the minimization problem (5) is $\mathbf{q}^* = \mathbf{e}_{s^*}$, where s^* is given by

$$s^* = \operatorname{argmin}\{\bar{T}_s / I_s^k \mid s \in \{1, \dots, n\}\},$$

and the minimum objective function is $\mathbb{E}[T_d^*|H_k] = \bar{T}_{s^*} / I_{s^*}^k$.

B. Scenario II (Optimization of the worst decision time):

We consider the multi-objective optimization problem of minimizing both conditional expected decision times. We construct single aggregate objective function by considering the maximum of the two objective functions. The associated optimization problem is:

$$\underset{\mathbf{q} \in \Delta_{n-1}}{\text{minimize}} \quad \max\{g^0(\mathbf{q}), g^1(\mathbf{q})\}. \quad (6)$$

The conditional expected decision times g^k are monotone along any line in the feasible simplex, and hence, the solution of the optimization problem (6) lies either at some vertex of the feasible simplex or at the intersection of the graphs of g^0 and g^1 . Further, the minimum along the the intersection of the graphs of g^0 and g^1 is achieved at some edge of the feasible simplex. To rigorously characterize these ideas, we introduce some notation. We assume that the sensors have been re-ordered such that the entries in $\mathbf{I}^0 - \mathbf{I}^1$ are in ascending order. We further assume that, for $\mathbf{I}^0 - \mathbf{I}^1$, the first m entries, $m < n$, are non positive, and the remaining entries are positive. For each $s \in \{1, \dots, m\}$, $r \in \{m+1, \dots, n\}$

and $i \in \{1, \dots, n\}$, we define the optimal point at an edge joining vertices e_r and e_s by

$$\tilde{q}_i^{sr} = \begin{cases} \frac{(I_r^0 - I_r^1)}{(I_r^0 - I_r^1) - (I_s^0 - I_s^1)}, & \text{if } i = s, \\ 1 - \tilde{q}_i^{sr}, & \text{if } i = r, \\ 0, & \text{otherwise.} \end{cases}$$

The optimal edge and the associated minimum value is:

$$(s^*, r^*) \in \underset{\substack{r \in \{m+1, \dots, n\} \\ s \in \{1, \dots, m\}}}{\operatorname{argmin}} \frac{(I_r^0 - I_r^1)\bar{T}_s - (I_s^0 - I_s^1)\bar{T}_r}{I_s^1 I_r^0 - I_s^0 I_r^1}, \quad \text{and}$$

$$g_{\text{two-sensors}}(s^*, r^*) = \frac{(I_{r^*}^0 - I_{r^*}^1)\bar{T}_{s^*} - (I_{s^*}^0 - I_{s^*}^1)\bar{T}_{r^*}}{I_{s^*}^1 I_{r^*}^0 - I_{s^*}^0 I_{r^*}^1}.$$

The optimal vertex and associated minimum value is:

$$w^* = \underset{w \in \{1, \dots, n\}}{\operatorname{argmin}} \max \left\{ \frac{\bar{T}_w}{I_w^0}, \frac{\bar{T}_w}{I_w^1} \right\}, \quad \text{and}$$

$$g_{\text{one-sensor}}(w^*) = \max \left\{ \frac{\bar{T}_{w^*}}{I_{w^*}^0}, \frac{\bar{T}_{w^*}}{I_{w^*}^1} \right\}.$$

Theorem 4 (Worst case optimization): For the optimization problem (6), an optimal probability vector is given by:

$$\mathbf{q}^* = \begin{cases} \mathbf{e}_{w^*}, & \text{if } g_{\text{one-sensor}}(w^*) \leq g_{\text{two-sensors}}(s^*, r^*), \\ \tilde{\mathbf{q}}^{s^* r^*}, & \text{if } g_{\text{one-sensor}}(w^*) > g_{\text{two-sensors}}(s^*, r^*), \end{cases}$$

and the minimum value of the objective function is:

$$\min \{g_{\text{one-sensor}}(w^*), g_{\text{two-sensors}}(s^*, r^*)\}.$$

C. Scenario III (Optimization of the average decision time):

For the multi-objective optimization problem of minimizing both conditional expected decision times, we formulate the single aggregate objective function as the average of these decision times. Let π_k be the prior probability of H_k being true, the resulting optimization problem is posed in the following way:

$$\underset{\mathbf{q} \in \Delta_{n-1}}{\operatorname{minimize}} \quad \pi_0 g^0(\mathbf{q}) + \pi_1 g^1(\mathbf{q}). \quad (7)$$

For a generic set of sensors, the Jacobian of the average expected decision time does not vanish in the interior of the feasible simplex, and the optimal solution lies at some edge of the simplex. The global optimal can be computed by comparing the optimal values at each edge. We rigorously characterize the solution in the following theorems:

Theorem 5 (Optimal number of sensors): For $n > 2$, if each 3×3 submatrix of the matrix $[\mathbf{T} \quad -\mathbf{I}^0 \quad -\mathbf{I}^1] \in \mathbb{R}^{n \times 3}$ is full rank, then the following statements hold:

- i) every solution of the optimization problem (7) lies on an edge of the probability simplex Δ_{n-1} ; and
- ii) every time-optimal policy requires at most two sensors to be observed.

Theorem 6 (Optimization on an edge): Given two vertices e_s and e_r , $s \neq r$, of the probability simplex Δ_{n-1} , then for the problem (7), the following statements hold:

- i) if $g^0(e_s) < g^0(e_r)$, and $g^1(e_s) < g^1(e_r)$, then the minima, along the edge joining e_s and e_r , lies at e_s , and optimal value is given by $\pi_0 g^0(e_s) + \pi_1 g^1(e_s)$; and

- ii) if $g^0(e_s) < g^0(e_r)$, and $g^1(e_s) > g^1(e_r)$, then the minima, along the edge joining e_s and e_r , lies at the point $\mathbf{q}^* = (1 - \nu^*)\mathbf{e}_s + \nu^*\mathbf{e}_r$, where

$$\nu^* = \min\left(1, \left(\frac{1}{1 + \mu}\right)^+\right), \quad \text{and}$$

$$\mu = \frac{I_r^0 \sqrt{\pi_1 (\bar{T}_s I_r^1 - \bar{T}_r I_s^1)} - I_r^1 \sqrt{\pi_0 (\bar{T}_r I_s^0 - \bar{T}_s I_r^0)}}{I_s^1 \sqrt{\pi_0 (\bar{T}_r I_s^0 - \bar{T}_s I_r^0)} - I_s^0 \sqrt{\pi_1 (\bar{T}_s I_r^1 - \bar{T}_r I_s^1)}},$$

and the optimal value is given by

$$\begin{aligned} \pi_0 g^0(\mathbf{q}^*) + \pi_1 g^1(\mathbf{q}^*) &= \\ \begin{cases} \pi_0 g^0(e_s) + \pi_1 g^1(e_s), & \text{if } \nu^* = 0, \\ \pi_0 g^0(e_r) + \pi_1 g^1(e_r), & \text{if } \nu^* = 1, \\ \left(\sqrt{\frac{\pi_0 (\bar{T}_s I_r^1 - \bar{T}_r I_s^1)}{I_s^1 I_r^0 - I_r^1 I_s^0}} + \sqrt{\frac{\pi_1 (\bar{T}_r I_s^0 - \bar{T}_s I_r^0)}{I_s^0 I_r^1 - I_r^0 I_s^1}} \right)^2, & \text{otherwise.} \end{cases} \end{aligned}$$

VI. ADAPTIVE SENSOR SELECTION

In the previous sections, we considered the open loop policies, i.e., we did not incorporate the information available from the observations to modify our policy. We now present an adaptive policy where we learn the likelihood of the true hypothesis from each observation, and utilize it to modify our sensor selection policy at each iteration. The SPRT with adaptive sensor selection policy is presented in Algorithm 2. We now present a formal analysis of this policy.

Algorithm 2 SPRT with adaptive sensor selection policy

- 1: set $(\pi_0, \pi_1) = (0.5, 0.5)$
 - 2: solve problem (7) to obtain \mathbf{q}^*
 - 3: at time $\tau \in \mathbb{N}$, select a random sensor s_τ according to the probability vector \mathbf{q}^* and collect a sample y_{s_τ}
 - 4: compute the log likelihood ratio $\lambda_\tau := \log \frac{f_{s_\tau}^1(y_\tau)}{f_{s_\tau}^0(y_\tau)}$
 - 5: integrate evidence up to current time $\Lambda_\tau := \sum_{t=1}^{\tau} \lambda_t$
 - 6: compute $\pi_0 = 1/(1 + \exp(\Lambda_\tau))$, $\pi_1 = 1 - \pi_0$
% decide only if a threshold is crossed
 - 7: **if** $\Lambda_\tau > \eta_1$, **then** accept H_1
 - 8: **if** $\Lambda_\tau < \eta_0$, **then** accept H_0
 - 9: **else** $(\Lambda_\tau \in]\eta_0, \eta_1[)$ continue sampling (step 2:)
-

We introduce the following notation. As stated in Theorem 3, conditioned on a hypothesis, optimal policy is to pick one sensor deterministically. Let s_k^* be the optimal sensor conditioned of hypothesis H_k .

For $r, s \in \{1, \dots, n\}$, we define the following quantities

$$\eta_{rs}^{\max} = \max \left\{ \left(\frac{I_r^1}{I_r^0} \right)^2, \left(\frac{I_s^1}{I_s^0} \right)^2 \right\} \left(\frac{(\bar{T}_r I_s^0 - \bar{T}_s I_r^0)}{(\bar{T}_s I_r^1 - \bar{T}_r I_s^1)} \right)^+,$$

$$\eta_{rs}^{\max} = \max\{\eta_{rs}^{\max} \mid r, s \in \{1, \dots, n\}, r \neq s\},$$

$$\eta_s^{\text{one}} = ((g^0(e_{s_1^*}) - g^0(e_s))^+ / (g^1(e_s) - g^1(e_{s_1^*}))), \quad \text{and}$$

$$\eta_s^{\text{one}} = \max\{\eta_s^{\text{one}} \mid s \in \{1, \dots, n\} \setminus \{s_1^*\}\}.$$

Let $\mathcal{D}_s(0) := \mathcal{D}(f_s^0, f_s^1)$, and $\mathcal{D}_s(1) := \mathcal{D}(f_s^1, f_s^0)$. We now state some asymptotic properties of the adaptive sensor selection policy.

Theorem 7 (Adaptive sensor selection): Given finitely informative sensors $\{1, \dots, n\}$, the following statements hold for the SPRT with adaptive sensor selection policy:

- i) Conditioned on hypothesis H_k , the sample size required to cross a threshold $(-1)^{k+1}\eta$, $\eta \in \mathbb{R}_{\geq 0}$ is finite almost surely.
- ii) Conditioned on a hypothesis, the sample size for decision $N_d \rightarrow +\infty$ almost surely as $\alpha_0, \alpha_1 \rightarrow 0^+$.
- iii) Conditioned on hypothesis H_k , the sample size for decision N_d satisfies

$$\frac{N_d}{|\eta_k|} \rightarrow \frac{1}{\mathcal{D}_{s_k^*}(k)}, \text{ almost surely, as } \alpha_0, \alpha_1 \rightarrow 0^+.$$

- iv) The expected sample size satisfies

$$\frac{\mathbb{E}[N_d|H_k]}{|\eta_k|} \rightarrow \frac{1}{\mathcal{D}_{s_k^*}(k)}, \text{ as } \alpha_0, \alpha_1 \rightarrow 0^+.$$

- v) Conditioned on hypothesis H_k , the decision time T_d satisfies

$$\frac{T_d}{|\eta_k|} \rightarrow \frac{\bar{T}_{s_k^*}}{\mathcal{D}_{s_k^*}(k)}, \text{ almost surely, as } \alpha_0, \alpha_1 \rightarrow 0^+.$$

- vi) The expected decision time satisfies

$$\frac{\mathbb{E}[T_d|H_k]}{|\eta_k|} \rightarrow \frac{\bar{T}_{s_k^*}}{\mathcal{D}_{s_k^*}(k)}, \text{ as } \alpha_0, \alpha_1 \rightarrow 0^+.$$

Proof: The proof of the first two statements follow from Theorem 1 and Lemma 1 in [12]. Note that the observations in the adaptive policy are not identically distributed, but the proof of convergence requires only independence and sensors being finitely informative.

To prove the third statement, we observe from Theorem 6 that the optimal solution on the edge joining vertex e_s and e_r always occurs at a vertex if $\mu \leq 0$, which holds if $\pi_1/\pi_0 > \eta_{rs}^{\max}$. Therefore, the solution on any edge lies at a vertex if $\pi_1/\pi_0 > \eta^{\max}$. Without loss of generality assume that hypothesis H_1 is true. From Theorem 3, it follows that $g^1(e_{s_1^*}) \leq g^1(e_s)$ for all $s \in \{1, \dots, n\}$. It can be verified that $\pi_0 g^0(e_{s_1^*}) + \pi_1 g^1(e_{s_1^*}) \leq \pi_0 g^0(e_s) + \pi_1 g^1(e_s)$ if $\pi_1/\pi_0 > \eta_s^{\text{one}}$. Therefore, it follows that the adaptive sensor selection policy always picks sensor s_1^* if $\log(\pi_1/\pi_0) > \eta^*$, where $\eta^* = \log(\max\{\eta^{\max}, \eta^{\text{one}}\})$. Note that $\eta^* < +\infty$, and hence, from the first statement of the theorem, conditioned on hypothesis H_1 , the SPRT with adaptive sensor selection crosses η^* in finite time, almost surely. Let N^* be this time. From second statement of the theorem, it follows that $N_d \rightarrow +\infty$ almost surely as $\alpha_0, \alpha_1 \rightarrow 0^+$. We now observe that

$$\begin{aligned} \frac{1}{N_d} \sum_{t=1}^{N_d} \log \frac{f_{s_t}^1(y_t)}{f_{s_t}^0(y_t)} &= \frac{1}{N_d} \sum_{t=1}^{N^*} \log \frac{f_{s_t}^1(y_t)}{f_{s_t}^0(y_t)} \\ &+ \frac{N_d - N^*}{N_d} \frac{1}{N_d - N^*} \sum_{t=N^*+1}^{N_d} \log \frac{f_{s_1^*}^1(y_t)}{f_{s_1^*}^0(y_t)}. \end{aligned} \quad (8)$$

Note that the observations are i.i.d. for $t \geq N^*+1$. Therefore, it follows from strong law of large numbers [9] that

$$\lim_{N_d \rightarrow +\infty} \frac{1}{N_d} \sum_{t=1}^{N_d} \log \frac{f_{s_t}^1(y_t)}{f_{s_t}^0(y_t)} = \mathbb{E} \left[\log \frac{f_{s_1^*}^1(y_t)}{f_{s_1^*}^0(y_t)} \middle| H_1 \right] = \mathcal{D}_{s_1^*}(1).$$

Since, N_d is sample size for decision, and $\alpha_0, \alpha_1 \rightarrow 0^+$, there is no false alarm or mis-detection, it follows

$$\sum_{t=1}^{N_d} \log \frac{f_{s_t}^1(y_t)}{f_{s_t}^0(y_t)} > \eta_1, \text{ and } \sum_{t=1}^{N_d-1} \log \frac{f_{s_t}^1(y_t)}{f_{s_t}^0(y_t)} < \eta_1.$$

We take \liminf and \limsup as $N_d \rightarrow +\infty$, respectively, on the two inequalities to obtain

$$\lim_{N_d \rightarrow +\infty} \sum_{t=1}^{N_d} \log \frac{f_{s_t}^1(y_t)}{f_{s_t}^0(y_t)} = \eta_1,$$

and thus, the third statement of the Theorem follows.

The proof of the remaining three statements is similar to the proof of Theorem 1 in [12] with a construction as in the proof of third statement. ■

Now we establish bounds on the performance of the adaptive policy. We denote the least expected decision time, conditioned on hypothesis H_k , among all the sensor selection policies by T_{\min}^k . Let T_{adp} be the decision time of the adaptive sensor selection policy. We now establish the following bounds on the performance:

Theorem 8 (Global bounds on performance): The following statements hold for the SPRT with adaptive sensor selection policy as $\alpha_0, \alpha_1 \rightarrow 0^+$

- i) Conditioned on a hypothesis H_k , the least expected decision time among all the sensor selection policies satisfies $T_{\min}^k \geq g^k(e_{s_k^*})$.
- ii) The expected decision time for the adaptive sensor selection policy, conditioned on hypothesis H_k , satisfies $\mathbb{E}[T_{\text{adp}}|H_k] \leq \min\{g_{\text{one-sensor}}(w^*), g_{\text{two-sensors}}(s^*, r^*)\}$.

Proof: We first establish the global lower bound. We first note that as $\alpha_0, \alpha_1 \rightarrow 0^+$, $-\eta_0, \eta_1 \rightarrow +\infty$. It is known that conditioned on a hypothesis, the decision time of the SPRT $T_d \rightarrow +\infty$ almost surely as $-\eta_0, \eta_1 \rightarrow 0^+$ [12]. Further, as $-\eta_0, \eta_1 \rightarrow +\infty$, the correct decision is made with probability one. Without loss of generality, suppose H_1 is true. Suppose that x_s is the number of times the sensor s is chosen. For a given expected decision time \bar{T}_d , the optimal policy will maximize the expected aggregate likelihood ratio. Therefore, the optimal sensor selection policy is the solution of the following optimization problem

$$\begin{aligned} &\text{maximize}_{x_s \in \mathbb{Z}_{\geq 0}} \sum_{s=1}^n x_s \mathcal{D}_s(1) \\ &\text{subject to} \quad \sum_{s=1}^n x_s \bar{T}_s \leq \bar{T}_d. \end{aligned} \quad (9)$$

The relaxation of the optimization problem (9) is a linear program. Therefore, the relaxed solution lies at one of the vertex of the feasible simplex and is given by:

$$x_s = \begin{cases} \bar{T}_d / \bar{T}_s, & \text{if } s = \text{argmax}_{r \in \{1, \dots, n\}} \mathcal{D}_r(1) / \bar{T}_r, \\ 0, & \text{otherwise.} \end{cases} \quad (10)$$

The non-zero component of the solution in equation (10) may not be an integer, but the fractional part is negligible as $\bar{T}_d \rightarrow +\infty$. This establishes the first statement.

We observe that $\pi_0 g^0(\mathbf{q}) + \pi_1 g^1(\mathbf{q}) \leq \max(g^0(\mathbf{q}), g^1(\mathbf{q}))$, and the second statement follows immediately. ■

VII. NUMERICAL EXAMPLES

We now elucidate on the results in previous sections with an example. We first demonstrate that the expressions for expected decision time in equation (4) provide a lower bound to the empirical expected decision time. We show that the

optimal policies for the worst case decision time and the average expected decision time in Section V may yield an expected decision time much larger than the optimal expected decision time conditioned on the true hypothesis. We show that the adaptive sensor selection policy works better than the open loop optimal policies for the worst case decision time and the average expected decision time.

Example 9: We consider a set of four sensors. Conditioned on hypothesis H_0 , the observations from each sensor is normally distributed with mean value $\{0, 0, 0, 0\}$ and variance $\{1, 4, 9, 16\}$, respectively. Conditioned on hypothesis H_1 , the observations from each sensor are normally distributed with mean value $\{1, 1, 1, 1\}$ and variance $\{9, 1, 16, 4\}$, respectively. The expected processing time of sensors are $\{4, 5, 10, 2\}$ secs, respectively.

A comparison of the empirical and analytic conditional expected decision time for different sensor selection probabilities is shown in Figure 2. It can be seen that the analytic expected decision time is a lower bound to the empirical expected decision time. Conditioned on H_1 , the optimal policy for the average expected decision time performs very badly.

A sample evolution of the SPRT with adaptive sensor selection policy is shown in Figure 3. It can be seen that as the posterior probability of a particular hypothesis being true increases, the adaptive policy chooses the best sensor for that hypothesis with higher probability. In this example, sensor 1 is the best for hypothesis H_1 , while sensor 4 is the best for hypothesis H_0 .

The performance of the adaptive policy is shown in Figure 4. It can be seen that as the error probability decreases, the expected decision time of the adaptive policy is very close to the optimal conditional decision time.




Fig. 2. Empirical and analytic expected decision times conditioned on hypothesis H_1 . The solid blue line and blue + represent analytic and empirical expected decision time for the optimal policy for conditional decision time. The dotted black line and black triangles represent analytic and empirical expected decision time for the optimal policy for worst decision time. The dashed red line and red x represent analytic and empirical expected decision time for the optimal policy for average expected decision time.




Fig. 3. A sample evolution of the SPRT with adaptive sensor selection policy. In the figure on left, the solid blue line represents the posterior probability of H_1 being true and the dashed red line shows the posterior probability of H_0 being true. In the figure on right, solid blue line represents the selection probability of sensor 1, while the dashed red line represents the selection probability of sensor 4. All other sensors have zero selection probability.

VIII. CONCLUSIONS

In this paper, we considered a sequential decision making problem with switching sensors. We studied version of the SPRT algorithm where the sensor switches at each observation. We used this sequential procedure to decide reliably. We




Fig. 4. Performance of the adaptive sensor selection policy. The solid blue line represents the expected decision time for the adaptive sensor selection policy. The dotted black line and dashed red line represent the optimal worst case expected decision time and the optimal conditioned expected decision time, respectively.

studied the set of optimal sensors to be observed in order to decide in minimum time. We observed the trade off between the information carried by a sensor and its processing time. A randomized sensor selection strategy was adopted. Optimal open loop policies were developed for three performance metrics. An adaptive sensor selection policy was developed that learns the true hypothesis and accordingly adapts the sensor selection strategy. The asymptotic performance of the adaptive sensor selection policy was characterized. An asymptotic lower bound on the performance of SPRT with multiple sensor for any sensor selection strategy was developed. It was shown that the adaptive sensor selection policy achieves this lower bound asymptotically.

REFERENCES

- [1] D. Bajović, B. Sinopoli, and J. Xavier. Sensor selection for hypothesis testing in wireless sensor networks: a Kullback-Leibler based approach. In *IEEE Conf. on Decision and Control*, pages 1659–1664, Shanghai, China, December 2009.
- [2] R. Bogacz, E. Brown, J. Moehlis, P. Holmes, and J. D. Cohen. The physics of optimal decision making: A formal analysis of performance in two-alternative forced choice tasks. *Psychological Review*, 113(4):700–765, 2006.
- [3] W. M. Bulkeley. Chicago’s camera network is everywhere. *The Wall Street Journal*, Nov 17 2009.
- [4] D. A. Castañón. Optimal search strategies in dynamic hypothesis testing. *IEEE Transactions on Systems, Man & Cybernetics*, 25(7):1130–1138, 1995.
- [5] R. Debouk, S. Lafortune, and D. Teneketzis. On an optimization problem in sensor selection. *Discrete Event Dynamic Systems*, 12(4):417–445, 2002.
- [6] 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*, 42(2):251–260, 2006.
- [7] V. Isler and R. Bajcsy. The sensor selection problem for bounded uncertainty sensing models. *IEEE Transactions on Automation Sciences and Engineering*, 3(4):372–381, 2006.
- [8] S. Joshi and S. Boyd. Sensor selection via convex optimization. *IEEE Transactions on Signal Processing*, 57(2):451–462, 2009.
- [9] S. I. Resnick. *A Probability Path*. Birkhäuser, 1999.
- [10] T. Shanker and M. Richtel. In new military, data overload can be deadly. *The New York Times*, Jan 16, 2011.
- [11] D. Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. Springer, 1985.
- [12] V. Srivastava, K. Plarre, and F. Bullo. Randomized sensor selection in sequential hypothesis testing. *IEEE Transactions on Signal Processing*, 59(5):2342 – 2354, 2011.
- [13] W. P. Tay, J. N. Tsitsiklis, and M. Z. Win. Asymptotic performance of a censoring sensor network. *IEEE Transactions on Information Theory*, 53(11):4191–4209, 2007.
- [14] H. Wang, K. Yao, G. Pottie, and D. Estrin. Entropy-based sensor selection heuristic for target localization. In *Symposium on Information Processing of Sensor Networks*, pages 36–45, Berkeley, CA, April 2004.
- [15] J. L. Williams, J. W. Fisher, and A. S. Willsky. Approximate dynamic programming for communication-constrained sensor network management. *IEEE Transactions on Signal Processing*, 55(8):4300–4311, 2007.
- [16] F. Zhao, J. Shin, and J. Reich. Information-driven dynamic sensor collaboration for tracking applications. *IEEE Signal Processing Magazine*, 19(2):61–72, 2002.