2017 IEEE 56th Annual Conference on Decision and Control (CDC)
December 12-15, 2017, Melbourne, Australia

Collaborative Sequential Detection of Gaussian Models from Observed Data and The Value of Information Exchanged

Aneesh Raghavan and John S Baras

Abstract— **We consider the problem of detecting which Gaussian model generates an observed time series data. We consider as possible generative models two linear systems driven by white Gaussian noise with Gaussian initial conditions. We also consider two collaborating observers. The observers observe a function of the state of the systems. Using these observations, the aim is to find which one of the two Gaussian models has generated the observations. For each observer we formulate a sequential hypothesis testing problem. Each observer computes its own likelihood ratio based on its own observations. Using the likelihood ratio, each observer performs sequential probability ratio test (SPRT) to arrive at its decision on the hypothesis. Taking into account the random and asymmetric stopping times of the two observers, we present a consensus algorithm which guarantees asymptotic convergence to the true hypothesis. The consensus algorithm involves exchange of information, i.e., the decision of the observers. Through simulations, the "value" of the information exchanged, probability of error and average time to consensus are computed.**

I. Introduction

Hypothesis testing and changepoint problems arise in various branches of engineering including quality control, detection and tracking of targets in war scenarios, detection of signals in seismology, econometrics, speech segmentation etc. Some recent applications are structural health monitoring of bridges, wind turbines, aircrafts, video scene analysis and sequential steganography [1]. Sequential analysis is a principal tool in addressing these problems. A sequential method is characterized by a stopping rule and a decision rule. These methods have been extensively studied in the literature when there is a single observer collecting all observations. In this paper we focus on a problem where there are multiple detectors collecting observations and work collaboratively to identify the true hypothesis.

The authors in [2] consider the problem where two detectors making independent observations need to decide which one of two hypotheses is true. The decisions of the two detectors are coupled through a common cost function. They prove that the optimal decisions are characterized by thresholds which are coupled and whose computation requires the solution of two coupled sets of dynamic programming equations. In [3] an information theoretic approach is presented to the distributed detection problem. They consider an entropy based cost function which maximizes the information transferred from the input to the output. They derive optimal decision and fusion rules with and without a fusion center. In [4] a decentralized sequential detection problem is considered. In their formulation, they consider a set of sensors making independent observations which need to decide as to which of the two hypotheses is true. The decision errors by the sensors are penalized through a common cost function. Each observation collected by the sensors as a team is assigned a positive cost. Optimal sensor decision rules are characterized through generalized sequential probability ratio tests (GSPRTs) and a technique for finding optimal thresholds is presented. In [5] the problem of noisy Bayesian active learning is addressed. They consider a hypothesis testing problem with observations corrupted by independent noise. Their objective is to find the true hypothesis using as few observations as possible by choosing the observations in an adaptive and strategic manner. They propose a sampling strategy which is based on collecting observations which maximize the Extrinsic Jensen - Shannon divergence at each step. In our previous work [6] we considered the problem of detecting Markov chain models from observed data. We used fixed block (given T observations) binary hypothesis testing and consensus to solve the problem.

In this paper we consider two Gaussian models and two observers. Under the alternate hypothesis, each observer observes a different function of the state of the first Gaussian model. Under the null hypothesis, each observer observes a different function of the state of the second Gaussian model. Thus each observer has its own sequence of observations. Given two sequences of observations (one for each observer), the objective is to find if the sequences were generated under the alternate hypothesis or under the null hypothesis. For each observer we formulate a sequential hypothesis testing problem which is solved using SPRT. We present a detection -estimation separation lemma which is useful in finding the likelihood ratio which is used in the SPRT. Based on the result of the SPRT, the observers could stop taking observations and arrive at the decision at the same time or at different times. We present a consensus algorithm which takes into account the various scenarios. Only the decisions made by the observers are exchanged in arriving at consensus. To understand the benefit of the one bit communication and its use by the two observers, the notion of value of information is introduced and discussed. Value of information, probability of error and average time to consensus have been calculated through Monte Carlo simulations. It should be noted that the two key differences of the formulation in this paper from the previous works mentioned here are: (i) each observer has its individual cost function (ii) the observations are not i.i.d.

Research supported by ARO grant W911NF-15-1-0646 and by DARPA through ARO grant W911NF-14-1-0384.

The authors are with the Department of Electrical and Computer Engineering and The Institute for Systems Research, University of Maryland, College Park, USA. Email: raghava@umd.edu, baras@umd.edu

978-1-5090-2872-6/17/$31.00 ©2017 IEEE

In the next section we discuss the problem formulation. In section [III], we first discuss the SPRT, followed by the detection -estimation separation lemma and then the consensus algorithm. In section [IV] we present the simulation results. In the last section we provide the conclusion and discuss future work.

II. Problem Formulation

A. System Model

Let $(\Omega, \mathscr{F}, \mathbb{P})$ be a probability space. Two systems are considered whose dynamics are described as follows : Dynamics of the state of system 1 is described by a linear Gaussian model as follows :

$$X_{k+1}^1 = A^1 X_k^1 + B^1 W_k^1, \; \forall \, k \geq 0,$$

where W_k^1 is white noise process with zero mean and covariance $R_1 \delta_{kk'}$. X_0^1 is assumed to be Gaussian random variable with zero mean and variance Σ_1. The dynamics of the state of system 2 is also described by a linear Gaussian model as follows :

$$X_{k+1}^2 = A^2 X_k^2 + B^2 W_k^2, \; \forall \, k \geq 0,$$

where W_k^2 is white noise process with zero mean and covariance $R_2 \delta_{kk'}$. X_0^2 is assumed to be Gaussian random variable with zero mean and variance Σ_2. We assume X_k^1 and X_k^2 belong to $\mathbb{R}^{N_s}$ for all k. H (signifying the hypothesis) is a Bernoulli random variable such that

$$\mathbb{P}(H = 1) = p_1, \; \mathbb{P}(H = 0) = p_0 = 1 - p_1.$$

Consider Observer 1. Under the alternate hypothesis, it observes a function of the state of system 1 and is described as follows :

$$Y_k^1 = C^1 X_k^1 + V_k^1, \; \forall \, k \geq 0,$$

where V_k^1 is white noise process with zero mean and covariance $Q_1 \delta_{kk'}$. Under the null hypothesis, it observes a function of the state of system 2 and is described as follows:

$$Y_k^2 = C^2 X_k^2 + V_k^2, \; \forall \, k \geq 0,$$

where V_k^2 is white noise process with zero mean and covariance $Q_2 \delta_{kk'}$. Similarly, Observer 2, under the alternate hypothesis, observes a function of the state of system 1 (different from the function observed by Observer 1) and is described as follows :

$$Z_k^1 = D^1 X_k^1 + U_k^1, \; \forall \, k \geq 0,$$

where U_k^1 is white noise process with zero mean and covariance $S_1 \delta_{kk'}$. Under the null hypothesis, it observes a function of the state of system 2 (different from the function observed by Observer 1) which is described as :

$$Z_k^2 = D^2 X_k^2 + U_k^2, \; \forall \, k \geq 0,$$

where U_k^2 is white noise process with zero mean and covariance $S_2 \delta_{kk'}$. Thus, the dynamics of the observations at Observer 1 can be compactly written as :

$$Y_k = [(C^1 X_k^1 + V_k^1) H + (C^2 X_k^2 + V_k^2)(1 - H)],$$

and the dynamics of the observations at Observer 2 can be compactly written as :

$$Z_k = [(D^1 X_k^1 + U_k^1) H + (D^2 X_k^2 + U_k^2)(1 - H)].$$

It is assumed that $\{W_k^1\}_{k \geq 0}$, $\{W_k^2\}_{k \geq 0}$, $\{V_k^1\}_{k \geq 0}$, $\{V_k^2\}_{k \geq 0}$, $\{U_k^1\}_{k \geq 0}$, $\{U_k^2\}_{k \geq 0}$, X_0^1, X_0^2 and H are independent. The dimension of Y_k is assumed to be M_1, while the dimension of Z_k is assumed to be M_2. Let $\mathscr{Y}_n^k$ denote the complete σ algebra generated by $\{Y_n, ..., Y_k\}$. Let $\mathscr{Z}_n^k$ denote the complete σ algebra generated by $\{Z_n, ..., Z_k\}$. A $\mathscr{Y}_n^k$ stopping time is a random time $\tau : \Omega \to \{n, n+1, ..., \infty\}$ such that $\{\omega \in \Omega \, : \, \tau(\omega) \leq k\} \in \mathscr{Y}_n^k$. The sigma algebra associated with a $\mathscr{Y}_n^k$ stopping time τ is defined as: $\mathscr{F}_\tau = \{A \in \mathscr{Y}_n^\infty \, : \, A \cap \{\tau \leq k\} \in \mathscr{Y}_n^k \; \forall \;\; k\}$. Let $\{\mathbb{S}_n^1 \; , \; n \geq 0\}$ denote the set of all possible $\mathscr{Y}_n^k$ stopping time τ such that $\mathbb{P}(\tau < \infty) = 1$. Also, let $\{\mathbb{S}_n^2 \; , \; n \geq 0\}$ denote the set of all possible $\mathscr{Z}_n^k$ stopping time τ such that $\mathbb{P}(\tau < \infty) = 1$.

B. Sequential Hypothesis Testing Problem

We consider the two observer problem given by :

$$\begin{aligned}
&Under \;\; H \; = \; 1 \;\; : \;\; X_{k+1}^1 = A^1 X_k^1 + B^1 W_k^1, \\
&Under \;\; H \; = \; 0 \;\; : \;\; X_{k+1}^2 = A^2 X_k^2 + B^2 W_k^2, \\
&Observer \;\; O1 \;\; : \\
&\quad Y_k = \left[(C^1 X_k^1 + V_k^1) H + (C^2 X_k^2 + V_k^2)(1 - H) \right], \\
&Observer \;\; O2 \;\; : \\
&\quad Z_k = \left[(D^1 X_k^1 + U_k^1) H + (D^2 X_k^2 + U_k^2)(1 - H) \right].
\end{aligned}$$

We define the following collection of optimization problems for each observer. n denotes the starting time for the optimization problem. The objective of Observer 1 is to find $\tau_n^1 \in \mathbb{S}_n^1$ and $D_{\tau_n^1}^1 \in \{0, 1\}$ which is $\mathscr{F}_{\tau_n^1}$ measurable such that following cost is minimized:

$$J^1(\tau_n^1, D_{\tau_n^1}^1) = \mathbb{E}[\alpha^1 \tau_n^1 + H(1 - D_{\tau_n^1}^1) + (1 - H) D_{\tau_n^1}^1], \quad (1)$$

where $\alpha^1 > 0$. The objective of Observer 2 is to find $\tau_n^2 \in \mathbb{S}_n^2$ and $D_{\tau_n^2}^2 \in \{0, 1\}$ which is $\mathscr{F}_{\tau_n^2}$ measurable such that following cost is minimized:

$$J^2(\tau_n^2, D_{\tau_n^2}^2) = \mathbb{E}[\alpha^2 \tau_n^2 + H(1 - D_{\tau_n^2}^2) + (1 - H) D_{\tau_n^2}^2], \quad (2)$$

where $\alpha^2 > 0$.

C. Consensus

The optimal decisions (beliefs of the true hypothesis) by Observer 1 and Observer 2 are obtained (as result of the previous optimization problem) at random times. The objective is to design an algorithm so that the two observers arrive at consensus about their beliefs by only exchanging their decisions.

III. Solution

A. Sequential Probability Ratio Test

For solutions using the dynamic programming appraoch we refer to [1]. Both finite horizon and infinite horizon problems have been studied in detail. The main drawback of this approach is that it is not computable. The sequential probability ratio test (SPRT) is also very well studied in the literature [[7], [8] and [1]] and is often used as a tool in sequential analysis. In the following, we discuss the SPRT for observations which are not i.i.d. We use ideas and techniques which are similar to the instance where the SPRT is derived for i.i.d observations. Consider the optimization problem (II-B) for Observer 1 starting at time 0. Define :

$$\pi_0^1 = f(H = 1|Y_0 = y_0).$$

It follows that,

$$\pi_0^1 = \frac{f(Y_0 = y_0|H = 1) \times p_1}{f(Y_0 = y_0|H = 1) \times p_1 + f(Y_0 = y_0|H = 0) \times p_0},$$
$$f(Y_0 = y_0|H = 1) = \int_{\mathbb{R}^{N_s}} f_{V^1}(y_0 - C^1 x) f_{X_0^1}(x)dx,$$
$$f(Y_0 = y_0|H = 0) = \int_{\mathbb{R}^{N_s}} f_{V^2}(y_0 - C^2 x) f_{X_0^2}(x)dx.$$

Minimizing the cost function [1] of the optimization problem starting at time 0 is equivalent to minimizing:

$$J^1(\tau_0^1, D^1_{\tau_0^1}) = \mathbb{E}[\alpha^1 \tau_0^1] + \pi_0^1 \mathbb{P}(D^1_{\tau_0^1} = 0|H = 1) + (1 - \pi_0^1)\mathbb{P}(D^1_{\tau_0^1} = 1|H = 0)$$

Define:

$$\mathbb{V}_1^1(\pi) = \inf_{\{\tau_0^1 \in \mathbb{S}_0^1 : \tau_0^1(\omega) \geq 1 \ \forall \omega \in \Omega\}, \{D^1_{\tau_0^1} \in \{0,1\}\}} \mathbb{E}\left[\alpha^1 \tau_0^1\right] + \pi \left[\mathbb{P}(D^1_{\tau_0^1} = 0|H = 1)\right] + (1 - \pi)\left[\mathbb{P}(D^1_{\tau_0^1} = 1|H = 0)\right].$$

For every $\tau_0^1 \in \mathbb{S}_0^1$, and $D^1_{\tau_0^1} \in \{0, 1\}$, $\mathbb{E}\left[\alpha^1 \tau_0^1\right] + \pi \left[\mathbb{P}(D^1_{\tau_0^1} = 0|H = 1)\right] + (1 - \pi)\left[\mathbb{P}(D^1_{\tau_0^1} = 1|H = 0)\right]$ is affine function of π. Hence $\mathbb{V}_1^1(\pi)$ is continuous and concave in π. The posterior cost incurred at time 0 is $\min((1 - \pi_0^1), (\pi_0^1))$. Let $\phi_0(\pi) = 1 - \pi$ and $\varphi_0(\pi) = \pi$. Let $\pi_U^* = \{0 < \pi < 1 : \mathbb{V}_1^1(\pi) = \phi_0(\pi)\}$ and $\pi_L^* = \{0 < \pi < 1 : \mathbb{V}_1^1(\pi) = \varphi_0(\pi)\}$. By concavity of $\mathbb{V}_1^1(\pi)$, it follows that if $\pi_0 \leq \pi_L^*$, it is optimal to stop with $D_0^1 = 0$. If $\pi_0 \geq \pi_U^*$, it is optimal to stop with $D_0^1 = 1$. Else the optimal strategy is to collect the next observation. At time k, let

$$\pi_k^1 = f(H = 1|\{Y_m = y_m\}_{m=0}^{m=k}).$$

Define :

$$\mathbb{V}_{k+1}^1(\pi) = \inf_{\{\tau_0^1 \in \mathbb{S}_0^1 : \tau_0^1(\omega) \geq k+1 \ \forall \omega \in \Omega\}, \{D^1_{\tau_0^1} \in \{0,1\}\}} \mathbb{E}\left[\alpha^1 \tau_0^1\right] + \pi \left[\mathbb{P}(D^1_{\tau_0^1} = 0|H = 1)\right] + (1 - \pi)\left[\mathbb{P}(D^1_{\tau_0^1} = 1|H = 0)\right].$$

The posterior cost incurred at time k is $\alpha^1 k + \min((1 - \pi_k^1), (\pi_k^1))$. Let $\pi_U^k = \{0 < \pi < 1 : \mathbb{V}_{k+1}^1(\pi) = \alpha^1 k + 1 - \pi\}$ and $\pi_L^k = \{0 < \pi < 1 : \mathbb{V}_{k+1}^1(\pi) = \alpha^1 k + \pi\}$. By same arguments as before, if $\pi_k \leq \pi_L^k$, it is optimal to stop with $D_k^1 = 0$. Else if $\pi_k \geq \pi_U^k$, it is optimal to stop with $D_k^1 = 1$. Else the optimal strategy is to collect the next observation. Hence threshold policies are optimal. We define the *Likelihood Ratio* (LLR) at time k (denoted by λ_k^1) as follows :

$$\lambda_k^1 = \frac{f(Y_k = y_k, Y_{k-1} = y_{k-1}, ..., Y_0 = y_0|H = 1)}{f(Y_k = y_k, Y_{k-1} = y_{k-1}, ..., Y_0 = y_0|H = 0)} = \frac{f(Y_k^1 = y_k, Y_{k-1}^1 = y_{k-1}, ..., Y_0^1 = y_0)}{f(Y_k^2 = y_k, Y_{k-1}^2 = y_{k-1}, ..., Y_0^2 = y_0)}.$$

From the above definition and definition of π_k^1, it follows, that

$$\pi_k^1 = \frac{p_1 \lambda_k^1}{p_0 + p_1 \lambda_k^1}$$
$$\Rightarrow \pi_k^1 \geq \pi_U^k \Leftrightarrow \lambda_k^1 \geq \frac{p_0 \pi_U^k}{p_1(1 - \pi_U^k})$$
$$\pi_k^1 \leq \pi_L^k \Leftrightarrow \lambda_k^1 \leq \frac{p_0 \pi_L^k}{p_1(1 - \pi_L^k)}$$

Hence, it suffices to compute the LLR and its associated thresholds. It remains to find the thresholds. Instead of finding the optimal thresholds, we find one pair of thresholds which is used at every k to achieve a desired level of performance. We denote the lower threshold associated with LLR by $\mathscr{A}$ and the upper threshold by $\mathscr{B}$. To find the pair $(\mathscr{A}, \mathscr{B})$, we use *Wald's approximation*.

Lemma 3.1: (*Wald's approximation*) Let β_d denote the desired probability of false alarm ($\mathbb{P}(D^1_{\tau_0^1} = 1|H = 0)$) and γ_d denote the desired probability of miss detection ($\mathbb{P}(D^1_{\tau_0^1} = 0|H = 1)$) to be achieved . Then the thresholds associated with LLR can be approximated as :

$$\mathscr{A} = \frac{\gamma_d}{1 - \beta_d} \ , \ \mathscr{B} = \frac{1 - \gamma_d}{\beta_d}. \tag{3}$$

Further, if $\beta_d = \gamma_d$, then the actual probabilities of false alarm (β_a) and miss detection (γ_a) are bounded above :

$$\beta_a \leq \beta_d + O(\beta_d^2), \ \gamma_a \leq \gamma_d + O(\gamma_d^2).$$

For proof, we refer to [8].

Thus, given desired probabilities of false alarm and miss detection, the thresholds associated with LLR can be computed. The test can be defined as :

$$\begin{aligned}
&\text{SPRT}(\mathscr{A}, \mathscr{B}) : \\
&\lambda_k^1 \geq \mathscr{B} \Rightarrow \tau_0^1 = k, D^1_{\tau_0^1} = 1, \\
&\mathscr{A} < \lambda_k^1 < \mathscr{B} \Rightarrow \text{collect next observation}, \\
&\lambda_k^1 \leq \mathscr{A} \Rightarrow \tau_0^1 = k, D^1_{\tau_0^1} = 0,
\end{aligned}$$

B. Detection Estimation Separation Lemma

To calculate the LLR, the joint distribution of the observations under either hypothesis needs to be found. The calculation of the joint distribution can be be simplified by invoking the following lemma. The general detection estimation separation theorem was studied in [9].

Lemma 3.2: Consider Observer 1 with observations $\{Y_m = y_m\}_{m=0}^{m=k}$. Then,

$$\lambda_k^1 = \frac{\prod_{j=1}^{j=k} f_{\Gamma_j^1}(y_j - C^1 A^1 \hat{x}_{j-1}^1) f_{Y_0^1}(y_0)}{\prod_{j=1}^{j=k} f_{\Gamma_j^2}(y_j - C^2 A^2 \hat{x}_{j-1}^2) f_{Y_0^2}(y_0)},$$

where, for $i = 1, 2,\ k \geq 1$,

$$\begin{aligned}
\hat{x}_k^i &= A^i \hat{x}_{k-1}^i + \mathbb{K}_k^i \eta_k^i, \\
\eta_k^i &= y_k - C^i A^i \hat{x}_{k-1}^i, \\
M_k^i &= A^i P_{k-1}^i A^{i^T} + B^i R_i B^{i^T}, \\
\mathbb{K}_k^i &= M_k^i C^{i^T} \left[C^i M_k^i C^{i^T} + Q_i \right]^{-1}, \\
P_k^i &= (I - \mathbb{K}_k^i C^i) M_k^i, \\
\hat{x}_0^i &= \Sigma_i C^{i^T} \left[C^i \Sigma_i C^{i^T} + Q_i \right]^{-1} \times y_0, \\
P_0^i &= \Sigma_i - \Sigma_i C^{i^T} \left[C^i \Sigma_i C^{i^T} + Q_i \right]^{-1} C^i \Sigma_i, \\
f_{\Gamma_k^i} &= \mathscr{N}(0, C^i M_k^i C^{i^T} + Q_i), \\
\Phi_k^i &= A^i \Phi_{k-1}^i A^{i^T} + B^i R_i B^{i^T}, \Phi_0^i = \Sigma_i; \\
\Psi_k^i &= C^i \Phi_k^i C^{i^T} + Q_i, \Psi_0^i = C^i \Sigma_i C^{i^T} + Q_i \\
f_{Y_k^i} &= \mathscr{N}(0, \Psi_k^i).
\end{aligned}$$

Proof: Using the theory of Kalman filters, it follows that the observation equations for Observer 1 under either hypothesis can be equivalently written as :

$$H = 1 : \begin{cases} Y_k^1 = C^1 A^1 \hat{X}_{k-1}^1 + \Gamma_k^1, \\ \hat{X}_{k-1}^1 = A^1 \hat{X}_{k-1}^1 + \mathbb{K}_k^1 \Gamma_k^1. \end{cases}$$

$$H = 0 : \begin{cases} Y_k^2 = C^2 A^2 \hat{X}_{k-1}^2 + \Gamma_k^2, \\ \hat{X}_{k-1}^2 = A^2 \hat{X}_{k-1}^2 + \mathbb{K}_k^2 \Gamma_k^2. \end{cases}$$

where $\mathbb{K}_k^i$ follows the recursions mentioned in the statement of the lemma and Γ_k^i are the innovation processes. Hence Γ_k^i is independent of the past observations $\{Y_m^i\}_{m=0}^{m=k-1}$. Using the definition of λ_k^1,

$$\lambda_k^1 = \frac{\prod_{j=1}^{j=k} f(Y_j^1 = y_j | Y_{j-1}^1 = y_{j-1}, ..., Y_0^1 = y_0) f_{Y_0^1}(y_0)}{\prod_{j=1}^{j=k} f(Y_j^2 = y_j | Y_{j-1}^2 = y_{j-1}, ..., Y_0^2 = y_0) f_{Y_0^2}(y_0)}.$$

The numerator of the R.H.S can be further simplified as:

$$\prod_{j=1}^{j=k} f(C^1 A^1 \hat{X}_{j-1}^1 + \Gamma_j^1 = y_j | \{Y_m^1 = y_m\}_{m=0}^{m=j-1}) f_{Y_0^1}(y_0)$$
$$= \prod_{j=1}^{j=k} f(\Gamma_j^1 = y_j - C^1 A^1 \hat{x}_{j-1}^1 | \{Y_m^1 = y_m\}_{m=0}^{m=j-1}) f_{Y_0^1}(y_0).$$

A similar simplification for the denominator can also be obtained. Since $\{\Gamma_k^i\}_{k \geq 1}$ are the innovation processes, the result of the lemma follows. ∎

C. Consensus Algorithm

Each observer arrives at its decision about the true hypothesis based on its own observations at random times. We now present the algorithm used by the observers to arrive at a consensus. We first mention the pseudo code for SPRT [Algorithm 1]. The consensus algorithm is described in detail in Algorithm 2. The summary of the consensus algorithm is as follows : The observers start taking observations at $k = 0$ with the objective of achieving certain probability of error. At each time instant they collect their observations and

Algorithm 1 SPRT

function SPRT($\lambda, \mathscr{A}, \mathscr{B}, n, \tau, D, k$) ▷ Where λ - LLR, $\mathscr{A}$,$\mathscr{B}$ are the thresholds, n denotes number of decisions, τ denotes stopping time, D denotes current decision and k denotes time
    $true \leftarrow 0$
    if $\lambda \geq \mathscr{B}$ **then**
        $n \leftarrow n + 1$, $\tau \leftarrow k$
        $Store\ k$, $D = 1$
        $\mathscr{A} \leftarrow \frac{1}{(\mathscr{B}+1) \times \nu - 1}$, $\mathscr{B} \leftarrow (\mathscr{B}+1) \times \nu - 1$
        $true \leftarrow 1$
    else if $\lambda \leq \mathscr{A}$ **then**
        $n \leftarrow n + 1$, $\tau \leftarrow k$
        $Store\ k$, $D = 0$
        $\mathscr{A} \leftarrow \frac{1}{(\mathscr{B}+1) \times \nu - 1}$, $\mathscr{B} \leftarrow (\mathscr{B}+1) \times \nu - 1$
        $true \leftarrow 1$
    return $[D, \mathscr{A}, \mathscr{B}, n, \tau, true]$

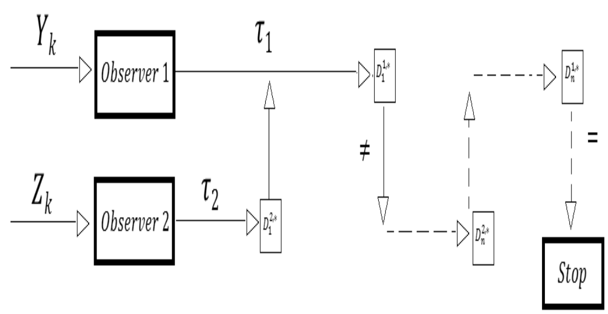

Fig. 1. Consensus Algorithm

update their LLR. Using the updated likelihood ratio they perform SPRT test. They could stop or continue collecting observations depending on the result of the test. If both the observers stop at the same time, then they exchange their decisions. If their decisions are the same, then they stop. If their decisions are different then they repeat SPRT test starting from next time instant with updated thresholds. If Observer 1 (Observer 2) stops first, it communicates its decision to Observer 2 (Observer 1). Observer 2 (Observer 1) continues with SPRT (with updated thresholds). When Observer 2 (Observer 1) stops, it checks its own decision with the decision obtained from Observer 1 (Observer 2). If the decisions are the same, then consensus has been achieved, else Observer 1 (Observer 2) starts performing SPRT again. When Observer 1 (Observer 2) starts performing SPRT again, note that it has not collected observations from $\tau_0^1 + 1$ to τ_0^2 (for Observer 2 it would be from $\tau_0^2 + 1$ to τ_0^1). Observer 1 updates its LLR as follows:

$$\lambda_{\tau_0^2+1}^1 = \frac{f_{Y_{\tau_0^2+1}^1}(y_{\tau_0^2+1}) \lambda_{\tau_0^1}^1}{f_{Y_{\tau_0^2+1}^2}(y_{\tau_0^2+1})},$$

$$\lambda_k^1 = \frac{f_{\Gamma_k^1}(y_k - C^1 A^1 \hat{x}_{k-1}^1) \lambda_{k-1}^1}{f_{\Gamma_k^2}(y_k - C^2 A^2 \hat{x}_{k-1}^2)}, k \geq \tau_0^2 + 2.$$

Algorithm 2 Consensus Algorithm

procedure CONSENSUS
  $D_f^1 \leftarrow -1$, $D_f^2 \leftarrow -2$, $true \leftarrow 0$
  $\tau^1 \leftarrow \infty$, $\tau^2 \leftarrow \infty$, $count \leftarrow 0$
  $n \leftarrow 0$, $m \leftarrow 0$, $\mu \leftarrow 3$, $\nu \leftarrow 2$
  $\mathscr{A}^j \leftarrow \dfrac{1}{\mu - 1}$, $\mathscr{B}^j \leftarrow \mu - 1$, $j = 1, 2$
  $State \leftarrow 1$, $i \leftarrow 0$,
  while $D_f^1 \neq D_f^2$ **do**
    $i \leftarrow i + 1$,
    if $State = 1$ **then**
      $Update\ \lambda_i^1$, λ_i^2
      $[D_f^1, \mathscr{A}^1, \mathscr{B}^1, n, \tau^1, true] \leftarrow SPRT(\lambda_i^1, \mathscr{A}^1, \mathscr{B}^1, n, D_f^1, \tau^1)$
      $[D_f^2, \mathscr{A}^2, \mathscr{B}^2, m, \tau^2, true] \leftarrow SPRT(\lambda_i^2, \mathscr{A}^2, \mathscr{B}^2, m, D_f^2, \tau^2)$
      if $\tau^1 = \tau^2$ **then**
        $State \leftarrow 1$
      else if $\tau^1 > \tau^2$ **then**
        $State \leftarrow 2$
      else if $\tau^1 < \tau^2$ **then**
        $State \leftarrow 3$
    else if $State = 2$ **then**
      if $count = 0$ **then**
        $\mathscr{A}^1 \leftarrow \dfrac{1}{\mu \times \nu - 1}$, $\mathscr{B}^1 \leftarrow \mu \times \nu - 1$
        $count \leftarrow 1$
      $Update\ \lambda_i^1$
      $[D_f^1, \mathscr{A}^1, \mathscr{B}^1, n, \tau^1, true] \leftarrow SPRT(\lambda_i^1, \mathscr{A}^1, \mathscr{B}^1, n, D_f^1, \tau^1)$
      if $true = 1$ **then**
        $State \leftarrow 3$
    else if $State = 3$ **then**
      if $count = 0$ **then**
        $\mathscr{A}^2 \leftarrow \dfrac{1}{\mu \times \nu - 1}$, $\mathscr{B}^2 \leftarrow \mu \times \nu - 1$
        $count \leftarrow 1$
      $Update\ \lambda_i^2$
      $[D_f^2, \mathscr{A}^2, \mathscr{B}^2, m, \tau^2, true] \leftarrow SPRT(\lambda_i^2, \mathscr{A}^2, \mathscr{B}^2, m, D_f^2, \tau^2)$
      if $true = 1$ **then**
        $State \leftarrow 2$

The filter updates are done as per Lemma 3.2. The Kalman filtering begins afresh, i.e., for $k \geq \tau_0^2 + 2$, the observations from $\tau_0^2 + 1$ to k are considered while filtering. The influence of the past information is considered in the LLR calculation. The LLR is calculated as the product of the LLR at τ_0^1 and ratio of the joint distribution of the observations from $\tau_0^2 + 1$ to k under $H = 1$ to that under $H = 0$. Observer 1 (Observer 2) performs SPRT based on the LLR computed and updated thresholds. When Observer 1 (Observer 2) stops it compares its decision to that of Observer 2 (Observer 1). If they are not equal then Observer 2 (Observer 1) starts SPRT at time $\tau^1_{\tau_0^2+1} + 1$ ($\tau^2_{\tau_0^1+1} + 1$). Hence, the observers alternatively collect observations and perform SPRT until consensus is achieved.

In algorithm 2, at the first iteration, if the observers stop at the same time, then $State = 1$. At the first iteration, if Observer 2 stops before Observer 1, then $State = 2$. Else if Observer 1 stops before Observer 2, then $State = 3$. After the first iteration, if the $State = 1$, the $State$ remains at 1 if the observers stop at the same time in further iterations as well. The first time, Observer 2 (Observer 1) stops before Observer 1 (Observer 2), the $State$ changes from 1 to 2 (3). Once the $State$ is equal to 2 or 3 it oscillates between these two states until the algorithm stops. It is also possible that the $State$ remains at 1 until consensus is achieved.

In figure 1, a simple scenario is depicted where Observer 2 arrives at its decision first and sends it to Observer 1. After Observer 1 has arrived at its decision, it compares its own decision to that of Observer 2. Since they are not equal, it communicates its decision to Observer 2 and Observer 2 starts collecting observations from the next time instant onwards. The algorithm is executed until consensus is achieved.

The thresholds are updated for each observer after every iteration. The lower threshold is monotonically decreasing with every iteration while the upper threshold is monotonically increasing. Thus, the consensus algorithm has been designed in such way that at the nth iteration, i.e., after both observers have made their final decisions n times, the probability of error is bounded above by $\dfrac{2}{\mu \times \nu^{n-1}}$ where μ and ν are greater than 1. Hence as n tends to ∞ the probability of error tends to zero.

IV. Simulation Results

The consensus algorithm described above involves information exchange between the observers. Through simulations we would like to understand if the exchange of the decision information has led to reduction in false alarms and miss detections. We consider the centralized decision of the observers as the decision of the observers after the first iteration. Hence each observer has its own centralized decision. A heuristic way to calculate the value of information for this specific problem would be to calculate the average reduction in detection error as :

α = Number of simulations in which consensus occurs to correct hypothesis after one iteration. We exclude the cases in which decision of both the observers is equal to the true hypothesis at the first iteration as exchange of decision information is not useful.

β = Number of simulations in which consensus occurs to wrong hypothesis, while the decision for either observers after the first iteration was equal to true hypothesis. In such cases the exchange of information is not useful, as the neither of the observer gain from the information exchange.

γ = Total number of bits communicated in all the simulations.

$total$ = Total number of simulations.

$$Value\ of\ information = \frac{\frac{\alpha - \beta}{total}}{\frac{\gamma}{total}} = \frac{\alpha - \beta}{\gamma}.$$

Hence the value of information captures the fraction of the total bits communicated which led to correction of the centralized decision of exactly one of the observers. Hence it captures how much of the information exchange was useful in the decision making process. Probability of error is calculated as :
υ = Number of simulations in which consensus occurs to wrong hypothesis.

$$Probability\ of\ error = \frac{\upsilon}{total}.$$

Average time to consensus is calculated as :
ϱ = Sum of the time to consensus over all simulations.

$$Average\ time\ to\ consensus = \lceil \frac{\varrho}{total} \rceil.$$

The simulations were performed with two Gaussian models. The states for both models were considered to be 3-dimensional. The parameters defining the systems under either hypothesis were considered as follows :

$$A^1 = \begin{bmatrix} -0.5 & 0 & 0 \\ 0 & -0.25 & 0 \\ 0 & 0 & 0.6 \end{bmatrix},\ A^2 = \begin{bmatrix} 0.7 & 0 & 0 \\ 0 & -0.4 & 0 \\ 0 & 0 & 0.35 \end{bmatrix},$$

$B^1 = B^2 = \mathbb{I}_3,\ \Sigma_1 = \Sigma_2 = R_1 = R_2 = 3 \times \mathbb{I}_3$. Observer 1 was considered to have 3-dimensional observations. The other parameters which define the observer were chosen as:

$$C^1 = \begin{bmatrix} 1 & 3 & 5 \\ 2 & 4 & 0 \\ 7 & 0 & 0 \end{bmatrix},\ C^2 = \begin{bmatrix} 2 & 4 & 6 \\ 1 & 3 & 0 \\ 8 & 0 & 0 \end{bmatrix},$$

$Q_1 = Q_2 = \mathbb{I}_3$. Observer 2 was considered to have 2-dimensional observations. The other parameters which define the observer were chosen as :

$$D^1 = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix},\ D^2 = \begin{bmatrix} 1 & 2 & 1 \\ 0 & 1 & 3 \end{bmatrix},$$

$S_1 = \mathbb{I}_2$, $S_2 = 2 \times \mathbb{I}_2$. The number of simulations was varied from 10 to 10^4. The value of information, probability of error and average time to consensus were calculated in each case and have been tabulated [I],[II]. For the simulation setting mentioned, the value of information is 0.27, which means that approximately 27% of the bits communicated led to correction of the centralized decision of exactly one of the observers. From table II, on an average, each observer collects 7 observations. With the same simulation settings and with minimum of 7 observations, the probability of error of Observer 1 was found to be 0.137. For Observer 2, the probability of error was found to be 0.184. By collecting 7 samples on an average, in a complementary manner and by the information exchange, the observers achieve lower probability of error equal to 0.062.

Number of Simulations	Value of Information	Probability of Error
10	0.3333	0.00
100	0.3066	0.03
1000	0.2609	0.068
10000	0.2719	0.0616

TABLE I
VALUE OF INFORMATION AND PROBABILITY OF ERROR

Number of Simulations	Average Time to Consensus
10	10
100	13
1000	13
10000	12

TABLE II
AVERAGE TIME TO CONSENSUS

V. CONCLUSION AND FUTURE WORK

In this paper , two collaborating detectors perform sequential hypothesis testing based on observations generated by Gaussian models. The SPRT is used to solve the hypothesis testing problem. A consensus algorithm with monotonically changing thresholds is presented. The convergence of the algorithm is discussed. To understand the value of the one bit communication used to achieve consensus, simulations were performed. It was observed that there was a reduction in erroneous detection. For the simulation settings mentioned, on an average, 27% of the information exchange resulted in an improved performance; i.e., the centralized decision of one of the observers was wrong while the consensus decision was the true hypothesis.

The stopping time problem for each one of the observers could be studied using approximate dynamic programming methods. It would also be interesting to consider the problem in a framework where the observations used to make the decision are chosen strategically and not all observations are used.

REFERENCES

[1] A. Tartakovsky, I. Nikiforov, and M. Basseville, *Sequential analysis: Hypothesis testing and changepoint detection*. CRC Press, 2014.
[2] D. Teneketzis and Y.-C. Ho, "The decentralized Wald problem," *Information and Computation*, vol. 73, no. 1, pp. 23–44, 1987.
[3] I. Y. Hoballah and P. K. Varshney, "An information theoretic approach to the distributed detection problem," *IEEE Transactions on Information Theory*, vol. 35, no. 5, pp. 988–994, 1989.
[4] V. V. Veeravalli, T. Başar, and H. Vincent Poor, "Decentralized sequential detection with sensors performing sequential tests," *Mathematics of Control, Signals, and Systems (MCSS)*, vol. 7, no. 4, pp. 292–305, 1994.
[5] M. Naghshvar, T. Javidi, and K. Chaudhuri, "Noisy Bayesian active learning," in *Communication, Control, and Computing (Allerton), 2012 50th Annual Allerton Conference on*. IEEE, 2012, pp. 1626–1633.
[6] A. Raghavan and J. Baras, "Detection of Markov chain models from observed data using binary hypothesis testing and consensus," to appear in Proceedings of 2017 IFAC World Congress.
[7] A. Wald, "Sequential tests of statistical hypotheses," *The Annals of Mathematical Statistics*, vol. 16, no. 2, pp. 117–186, 1945.
[8] H. V. Poor, *An introduction to signal detection and estimation*. Springer Science & Business Media, 2013.
[9] T. Kailath, "A general likelihood-ratio formula for random signals in Gaussian noise," *IEEE Transactions on Information Theory*, vol. 15, no. 3, pp. 350–361, 1969.