This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2008 proceedings.

Optimal Sleep–Wake Scheduling for Quickest Intrusion Detection using Sensor Networks

K. Premkumar and Anurag Kumar
Dept. of Electrical Communication Engineering (ECE)
Indian Institute of Science, Bangalore 560 012, INDIA
Email: {kprem, anurag}@ece.iisc.ernet.in

***Abstract*—We consider the problem of *quickest detection of an intrusion* using a sensor network, keeping only a minimal number of sensors *active*. By using a minimal number of sensor devices, we ensure that the energy expenditure for *sensing, computation* and *communication* is minimized (and the *lifetime* of the network is maximized). We model the intrusion detection (or change detection) problem as a *Markov decision process* (MDP). Based on the theory of MDP, we develop the following closed loop *sleep/wake* scheduling algorithms:**

1) **optimal control of M_{k+1}, the number of sensors in the *wake* state in time slot $k+1$,**
2) **optimal control of q_{k+1}, the probability of a sensor in the *wake* state in time slot $k+1$,**

and an open loop *sleep/wake* scheduling algorithm which

3) **computes q, the optimal probability of a sensor in the *wake* state (which does not vary with time),**

based on the sensor observations obtained until time slot k.

Our results show that an optimum closed loop control on M_{k+1} significantly decreases the cost compared to keeping any number of sensors active all the time. Also, among the three algorithms described, we observe that the total cost is minimum for the optimum control on M_{k+1} and is maximum for the optimum open loop control on q.

Keywords: Bayesian change detection, intrusion detection, quickest change detection with observation cost

I. Introduction

Sensor networks are *application specific* networks that comprise a large number of tiny, energy limited, low–powered smart sensor devices. A sensor can be in one of the two states, the *sleep* state or the *wake* state. In the *wake* state, the sensor makes measurements, performs some computation and then communicates information to the fusion center. The fusion center acts as a *decision maker* and may even act as a *controller*. The sensors being energy limited, impose stringent requirements on the energy–efficiency of the algorithms employed in the fusion center. In intrusion detection applications, the intrusions are typically rare events and hence sensor nodes spend a majority of their time in the *pre–intrusion* period which reduces the lifetime and the utility of the sensor network. Thus, it is essential to have an energy–efficient *sleep/wake* scheduling algorithm.

We are interested in quickest detection of intrusion which apparently requires the sensors to be in the *wake* state all the time. We thus have conflicting objectives (i) minimizing the energy consumption of the network (or equivalently, maximizing the lifetime of the network) and (ii) minimizing the detection delay.

Related literature: The problem of energy–efficient quickest change detection has been studied in various contexts. Appadwedula *et al.* [1] study a binary hypothesis testing problem. They minimize the probability of making an incorrect decision subject to energy constraints. Appadwedula *et al.* [1] and Rago *et al.* [6] study a *censoring* scheme, in which the uninformative observations are not sent to the fusion center. Note that [6] also considers a binary hypothesis testing problem. In [9], Wu *et al.* study *sleep/wake* scheduling for low duty cycle multi–hop sensor networks (employed for continuous monitoring applications) with synchronization errors. They maximize the lifetime of a sensor network that guarantees a data delivery performance. In [10], Zacharias and Sundaresan study a centralized detection problem based on physical layer fusion with power control at the sensors for energy efficiency. In our work, we achieve energy–efficiency by keeping only a *minimal* number of sensors *active*.

Our work differs from [1] and [6] in the following ways:

- We consider a sequential change–detection problem whereas [1] and [6] consider a simple binary hypothesis testing problem. [1] and [6] consider a static optimization problem whereas our work involves dynamic optimization.
- In [6], Rago *et al.* propose a *censoring* algorithm which makes a decision on whether to *transmit* an observation or *not* based on the *information content* of the observation and the communication–cost. Note that censoring runs at each of the sensors and does not take into account the sensing, sampling, and computation costs. [1] also uses the *censoring* scheme described above. In our work, we control the *sleep/wake* activity of the sensors thereby taking into account the energy required for sensing/computation. In our work, all sensor–observations are transmitted to the fusion center.

Summary of contributions: We summarize the main contributions of this paper below.

(i) We model the problem of quickest change (intrusion) detection by using a *minimal number of observations* in the Bayesian framework as a Markov decision process (MDP) that captures 1) the cost due to false alarm, 2) the cost due to the detection delay, and 3) the cost per observation per sensor in the network.
(ii) We derive the following closed loop scheduling algorithm

to control the number of sensors in the *wake* state, when there is feedback between the fusion center and the sensors.

At time slot k, the fusion center receives an observation vector, $\mathbf{X}_k^{(M_k)} := (X_k^{(1)}, X_k^{(2)}, \cdots, X_k^{(M_k)}) \in \mathbb{R}^{M_k}$ and computes the posterior probability of change, Π_k. The fusion center then chooses a control from the set of available controls, $\mathbf{A} = \left\{\mathsf{stop}, \bigcup_{m\in\{0,1,\cdots,n\}}(\mathsf{continue}, m)\right\}$. The control $u_k = (\mathsf{continue}, m)$ means that the detection process is continued at time slot $k+1$ with m sensors in the *wake* state. We observe that $\{\Pi_k\}$ forms a controlled Markov chain. Based on the theory of MDP, we derive the optimal policy μ^* which gives

1) **The stopping rule:**
Stop at time slot $\tau^* = \inf\ \{k \geq 0 : \Pi_k \geq \Gamma\}$, where $\Gamma \in [0,1]$ is a threshold, and
2) **The control policy for M_{k+1}:**
From the optimal policy, we can infer that there exists a map $M^* : [0,1] \to \mathbb{Z}_+$ such that the optimal number of sensors in the *wake* state in time slot $k+1$ is given by $M_{k+1} = M^*(\Pi_k)$.

We provide some *structural results* for the optimal policy.

(iii) We also derive another closed loop scheduling algorithm to control q_{k+1}, the probability that a sensor is in the *wake* state at time slot $k+1$.

At time slot k, the fusion center receives an observation vector $\mathbf{X}_k^{(M_k)}$ and computes Π_k, the posterior probability of change. The fusion center then chooses a control $u_k \in \mathbf{A} = \left\{\mathsf{stop}, \bigcup_{q\in[0,1]}(\mathsf{continue}, q)\right\}$. If $u_k = (0, q_{k+1})$ is the control chosen by the fusion center, then M_{k+1} is Bernoulli distributed with parameters (n, q_{k+1}). We observe that $\{(\Pi_k, M_k)\}$, $k \in \mathbb{N}$ forms a controlled Markov chain. Based on the theory of MDP, we derive the optimal policy μ^* which gives

1) **The stopping rule:**
Stop at time slot $\tau^* = \inf\ \{k \geq 0 : \Pi_k \geq \Gamma\}$, where $\Gamma \in [0,1]$ is a threshold, and
2) **The control policy for q_{k+1}:**
From the optimal policy, we can infer that there exists a map $q^* : [0,1] \to [0,1]$ such that the optimal probability that a sensor is in the *wake* state in time slot $k+1$ is given by $q_{k+1} = q^*(\Pi_k)$.

We show some structural results for the optimal policy.

(iv) We also derive the following open loop scheduling algorithm to control the *sleep/wake* activity of the sensors when the feedback between the fusion center and the sensors is not available. At time slot k, each sensor chooses to be in the *wake* state with probability q, independent of the state of other sensors. The fusion center receives a vector of observations, $\mathbf{X}_k^{(M_k)}$ ($M_k \sim$ Bernoulli(n, q)), and computes the posterior probability of change Π_k. The fusion center then decides whether to "stop" the decision process or "continue" sampling. We observe that $\{(\Pi_k, M_k)\}$ process forms a controlled Markov chain. Based on the theory of MDP, we obtain

1) **The stopping rule:**
Stop at time slot $\tau^* = \inf\ \{k \geq 0 : \Pi_k \geq \Gamma\}$, where $\Gamma \in [0,1]$ is a threshold.

Note that q is constant over time and we choose q that minimizes the Bayesian cost given in Eqn. 8.

Note: For all closed/open loop policies $M_{k+1}(\cdot) = 0, \forall k \geq \tau^*$.

Outline of the paper: The rest of the paper is organized as follows. In Section II, we formulate the *sleep/wake* scheduling problem for quickest change detection. In Section III, we solve the optimal *sleep/wake* scheduling problem that minimizes the detection delay when there is a feedback from the controller to the sensors. In Section IV, we discuss an optimal open loop *sleep/wake* scheduler that minimizes the detection delay. Finally, we summarize the paper in Section V.

II. Problem Formulation

In this section, we describe the *quickest intrusion detection problem with a minimal number of observations* and set up the model. We consider a sensor network comprising n (acoustic or vibration or magnetic or a combination of these) sensors deployed in a region $\mathcal{A}$ for an intrusion detection application. The sensors are collocated, i.e., the region $\mathcal{A}$ is covered by the *sensing coverage* of each of the sensors. An intrusion happens at a random time. The problem is to detect the intrusion as early as possible *using a minimal number of observations* subject to a false alarm constraint.

We consider a discrete time system and the basic unit of time is one slot. We assume that the sensor network is time synchronized. An event ("intruder" in a security system) happens at a random time T. The distribution of T (the time slot at which the intrusion/change/event happens) is given by

$$P\{T = k\} = \begin{cases} \pi_0 & \text{if } k \leq 0, \\ (1-\pi_0)(1-p)^{k-1}p & \text{if } k > 0. \end{cases}$$

where $0 < p \leq 1$ and $0 \leq \pi_0 \leq 1$ represents the probability that the change ("intrusion") happened even before the observations are made ($k \leq 0$). We model the intrusion by a change in the probability law of the sensor observations at a random time T. Note that the observations are obtained sequentially starting from time slot $k = 1$ onwards. We say that the state of nature, S_k is 0 before the occurrence of the event (i.e., $S_k = 0$ for $k < T$) and 1 after the occurrence of the event (i.e., $S_k = 1$ for $k \geq T$). Before the event takes place, i.e., for $1 \leq k < T$, sensor i observes $X_k^{(i)} \sim f_0^{(i)}(.)$ and after the event takes place, i.e., for $k \geq T$, sensor i observes $X_k^{(i)} \sim f_1^{(i)}(.)$ (because the sensors are collocated), where $f_0^{(i)}$ and $f_1^{(i)}$ are probability density functions (pdfs). Note that $f_0^{(i)} \neq f_1^{(i)}$ for all i. Conditioned on the state of the nature, the observations are independent across sensors and across time (the event and the observation models are essentially the same as in [8]). The observations are transmitted to a fusion center (the communication between the sensors and the fusion center is assumed to be error–free as the channel noise could be thought of as captured in the sensors itself) through parallel communication channels, [5]. At each time slot k, the fusion center receives a vector of observation $\mathbf{X}_k^{(M_k)} = \left(X_k^{(1)}, X_k^{(2)}, \cdots, X_k^{(M_k)}\right)$. At time slot k, based on the observations so far, $\left\{\mathbf{X}_t^{(M_t)}, t = 1, 2, \cdots, k\right\}$, π_0, the distribution of T, $f_0^{(i)}(\cdot)$s and $f_1^{(i)}(\cdot)$s, the fusion center

1) makes a decision on whether to raise an alarm or to continue sampling, and
2) controls the number of sensors in the *wake* state at time slot $k+1$.

The costs involved here are i) λ_s, the cost due to $(\mathsf{sampling} + \mathsf{computation} + \mathsf{communication})$ per observation per sensor, ii) λ_f, the cost of false alarm, and iii) the detection delay where the detection delay is defined as the delay between the occurrence of the event and the detection. Note that λ_s can also be considered as energy consumption per observation per sensor. We are interested in finding the quickest detection procedure with a minimal number of observations that minimizes the detection delay subject to a false alarm constraint, $\mathsf{PFA} \leq \alpha$, where PFA is the probability of false alarm. Note that the optimal detection procedure should give i) the optimal stopping time τ^* and ii) the optimal control on M_k, the number of sensors in the *wake* state at time slot k. We thus have a constrained optimization problem,

$$\begin{aligned} &\min_{\tau, M_1, M_2, \cdots, M_\tau} \mathbb{E}\left[(\tau - T)^+ + \lambda_s \sum_{k=1}^{\tau} M_k\right] \quad (1) \\ &\text{subject to} \qquad \mathsf{PFA} \leq \alpha \end{aligned}$$

where τ is a stopping time. Note that the above problem could be considered as a *quickest change detection* problem with energy constraint and λ_s as a *Lagrange multiplier* that relaxes the energy constraint.

Let Π_k be the posterior probability of the event happening at or before time slot k. Π_k is given by

$$\begin{aligned} \Pi_k &:= \mathbb{E}\left[\mathbf{1}_{\{T \leq k\}} \,\middle|\, \mathbf{X}_1^{(M_1)} \cdots \mathbf{X}_k^{(M_k)}\right] \quad (2) \\ \text{and hence, } \mathbb{E}\Pi_k &= \mathbb{E}\mathbb{E}\left[\mathbf{1}_{\{T \leq k\}} \,\middle|\, \mathbf{X}_1^{(M_1)} \cdots \mathbf{X}_k^{(M_k)}\right] \\ &= \mathbb{E}\mathbf{1}_{\{T \leq k\}} \quad (3) \end{aligned}$$

The Lagrangian relaxation of the problem defined in Eqn.1 is

$$\begin{aligned} R(\tau) &= \mathbb{E}\left[\lambda_f \mathbf{1}_{\{\tau < T\}} + (\tau - T)\mathbf{1}_{\{\tau \geq T\}} + \lambda_s \sum_{k=1}^{\tau} M_k\right] \\ &= \mathbb{E}\left[\lambda_f \mathbf{1}_{\{\tau < T\}} + \sum_{k=0}^{\tau-1} \mathbf{1}_{\{T \leq k\}} + \lambda_s \sum_{k=1}^{\tau} M_k\right] \\ &= \mathbb{E}\left[\lambda_f (1 - \Pi_\tau) + \sum_{k=0}^{\tau-1} \Pi_k + \lambda_s \sum_{k=1}^{\tau} M_k\right] \\ &= \mathbb{E}\left[\lambda_f (1 - \Pi_\tau) + \sum_{k=0}^{\tau-1} (\Pi_k + \lambda_s M_{k+1})\right] \quad (4) \end{aligned}$$

Refer Eqn. 3 and [7] for the justification of the steps in Eqn. 4. The Lagrange multiplier λ_f is chosen such that the false alarm constraint is satisfied with equality, i.e., $\mathsf{PFA} = \alpha$ (refer [7]). Thus, the constrained optimization problem defined in Eqn. 1 can be viewed as

$$\begin{aligned} &\min_{\tau, M_1, M_2, \cdots, M_\tau} R(\tau) \\ &= \min_{\tau, M_1, M_2, \cdots, M_\tau} \mathbb{E}\left[\lambda_f (1 - \Pi_\tau) + \sum_{k=0}^{\tau-1} (\Pi_k + \lambda_s M_{k+1})\right] \quad (5) \end{aligned}$$

We consider the following possibilities for the problem defined in Eqn. 5.

1) **Closed loop control on** M_{k+1}: At time slot k, the fusion center makes a decision on M_{k+1}, the number of sensors in the *wake* state in time slot $k+1$, based on the information available (at the fusion center) up to time slot k. The decision is then fed back to the sensors via a feedback channel. Thus, the problem becomes

$$\begin{aligned} \tau^*, M_1^*, M_2^*, \cdots, M_{\tau^*}^* = \arg \min_{\tau, M_{k+1}, k=0,1,2,\cdots,\tau-1} \\ \mathbb{E}\left[\lambda_f (1 - \Pi_\tau) + \sum_{k=0}^{\tau-1} (\Pi_k + \lambda_s M_{k+1})\right] \quad (6) \end{aligned}$$

2) **Closed loop control on** q_{k+1}: At time slot k, the fusion center makes a decision on q_{k+1}, the probability that a sensor is in the *wake* state at time slot $k+1$. q_{k+1} is then broadcast via a feedback channel to the sensors. Thus, it is easy to see that the number of sensors in the *wake* state M_{k+1}, at time slot $k+1$, is Bernoulli distributed with parameters (n, q_{k+1}). Also note that $\mathbb{E}M_{k+1} = nq_{k+1}$. Thus, the problem defined in Eqn. 5 becomes

$$\begin{aligned} \tau^*, q_1^*, q_2^*, \cdots, q_{\tau^*}^* = \arg \min_{\tau, q_{k+1}, k=0,1,2,\cdots,\tau-1} \\ \mathbb{E}\left[\lambda_f (1 - \Pi_\tau) + \sum_{k=0}^{\tau-1} (\Pi_k + \lambda_s n q_{k+1})\right] \quad (7) \end{aligned}$$

3) **Open loop control on** q: At time slot k, each sensor node is in the *wake* state with probability q. Note that M_k, the number of sensors in the *wake* state at time slot k is Bernoulli distributed with parameters (n, q). Also note that $\{M_k\}$ process is i.i.d. over time and that $\mathbb{E}M_k = nq$. Note that the probability q is constant over time. Thus, the problem defined in Eqn. 5 becomes

$$\begin{aligned} \tau^* = \arg \min_{\tau} \\ \mathbb{E}\left[\lambda_f (1 - \Pi_\tau) + \sum_{k=0}^{\tau-1} (\Pi_k + \lambda_s n q)\right] \quad (8) \end{aligned}$$

Here, q is chosen such that it minimizes the above cost.

Note that the first two scenarios require a feedback channel between the fusion center and the sensors whereas the last scenario does not require a feedback channel.

In Section III, we formulate the optimization problem defined in Eqns. 6 and 7 in the framework of MDP and study the optimal policy. We formulate the open loop control optimization problem defined in Eqn. 8 in the MDP framework and obtain optimal policy in Section IV.

III. Quickest Change Detection with Feedback

In this section, we study the *sleep/wake* scheduling problem when there is a feedback channel from the controller to the sensors.

At time slot k, the fusion center receives $\mathbf{X}_k^{(M_k)}$ and computes Π_k. Recall that $\Pi_k = \mathsf{P}\left\{T \leq k \,\middle|\, \mathbf{X}_1^{(M_1)}, \cdots, \mathbf{X}_k^{(M_k)}\right\}$ is the posterior probability of the event having occurred at or before time slot k. For the change detection problem, a sufficient statistic [4] for the sensor observations up to time slot k, is given by $\{\Pi_0, \Pi_1, \cdots, \Pi_k\}$. When an *alarm* is raised the system enters into an absorbing state 'a'. Thus, the state space of the $\{\Pi_k\}$ process is $\mathbf{S} = \{[0,1] \cup \{\mathsf{a}\}\}$ (we always work with the Borel space, $(\mathbf{S}, \mathscr{B}(\mathbf{S}))$). Note that Π_k is also called the *information state* of the system.

In the rest of the section, we explain the closed loop control techniques used to obtain the *sleep/wake* scheduling algorithms.

A. Control on M_{k+1}, the number of sensors in the wake state at time slot $k+1$

The fusion center on receiving $\mathbf{X}_K^{(M_k)}$ computes Π_k, the posterior probability of the event. It then makes a decision D_k on "*stop*"ping or to "*continue*" sampling. Note that $D_k = 0$ means that the sampling (and the detection process) is continued and $D_k = 1$ means that an alarm is raised. If $D_k = 0$, the controller chooses $M_{k+1} = m$, the number of sensors to be in the *wake* state at time slot $k+1$. Thus the set of controls at time slot k is given by $\mathbf{A} = \left\{\mathsf{stop}, \bigcup_{m \in \{0,1,\cdots,n\}} (\mathsf{continue}, m)\right\} = \{1, (0,0), (0,1), \cdots, (0,n)\}$.

When the control $u_k = (0,m)$ is chosen, the information state at time slot $k+1$, Π_{k+1} can be computed recursively using the map $\Phi : \mathbf{S} \times \mathbf{A} \to \mathbf{S}$ as

$$\Pi_{k+1} := \Phi\Big(\Pi_k, (0,m)\Big) = \frac{\widetilde{\Pi}_k \phi_1\Big(\mathbf{X}_{k+1}^{(m)}\Big)}{\phi_2\Big(\mathbf{X}_{k+1}^{(m)}; \widetilde{\Pi}_k\Big)} \tag{9}$$

where

$$\begin{aligned} \widetilde{\Pi}_k &:= \Pi_k + (1 - \Pi_k)p \qquad (10)\\ \mathbf{X}_k^{(m)} &:= (X_k^{(1)}, X_k^{(2)}, \cdots, X_k^{(m)}),\\ \phi_1(\mathbf{X}_k^{(m)}) &:= \prod_{i=1}^{m} f_1^{(i)}(X_k^{(i)}),\\ \phi_2(\mathbf{X}_k^{(m)}; \widetilde{\Pi}) &:= \widetilde{\Pi} \prod_{i=1}^{m} f_1^{(i)}(X_k^{(i)}) + (1 - \widetilde{\Pi}) \prod_{i=1}^{m} f_0^{(i)}(X_k^{(i)}) \end{aligned}$$

and when the control $u_k = 1$ is chosen, Π_{k+1} is given as

$$\Pi_{k+1} := \Phi(\Pi_k, 1) := \mathsf{a} \quad \text{w.p.1} \tag{11}$$

Note: $\widetilde{\Pi}_k = \mathbb{E}\Pi_{k+1}$ before $\mathbf{X}_{k+1}^{(\cdot)}$ is observed and $\{\Pi_k\}$ forms a controlled Markov chain. From Eqn. 6, we see that when the (state, action) pair is (π, u), the single stage cost function is

$$c(\pi, u) = \begin{cases} \lambda_f (1 - \pi), & \pi \in [0,1], u = 1\\ \pi + \lambda_s m, & \pi \in [0,1], u = (0,m)\\ 0, & \pi = \mathsf{a} \end{cases}$$

Therefore, the optimization problem defined in Eqn. 6 can be stated as

$$J^*(\pi_0) = \min_{u_0, u_1, u_2, \cdots} \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, u_k) \,\middle|\, \Pi_0 = \pi_0\right]$$

We use the theory of MDP to solve the above optimization problem. From [2], it is clear that we can find a stationary optimal policy which solves the above optimization problem. Let $\mu : \mathbf{S} \to \mathbf{A}$ be a stationary policy and $\mu^* : \mathbf{S} \to \mathbf{A}$ be the optimal stationary policy. Then, we have

$$\begin{aligned} J_\mu(\pi_0) &= \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right]\\ \text{and } J^*(\pi_0) &= \min_{\mu} \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right]\\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu^*(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right] \end{aligned}$$

Note that the optimal stationary policy μ^* is independent of the initial state $\Pi_0 = \pi_0$ [2]. By Bellman's equation, we have

$$J^*(\pi) = \min_{u \in \mathbf{A}} \left\{ c(\pi, u) + \mathbb{E}_{\phi_2(\mathbf{x}^{(m)}; \tilde{\pi})}\left[J^*\Big(\Phi(\pi, u)\Big)\right]\right\}.$$

Note: The notation $\mathbb{E}_{\phi_2(\mathbf{x}^{(m)}; \tilde{\pi})}[\cdot]$ means that the expectation is taken with respect to the pdf $\phi_2(\mathbf{x}^{(m)}; \tilde{\pi})$[1]. Since $J^*(\mathsf{a}) = 0$, the above equation can be written as

$$J^*(\pi) = \min\left\{\lambda_f \cdot (1 - \pi), \pi + A_{J^*}(\tilde{\pi})\right\}. \tag{12}$$

where the function $A_{J^*} : [0,1] \to \mathbb{R}_+$ is defined as

$$A_{J^*}(\tilde{\pi}) = \min_{0 \leq m \leq n} \left\{\lambda_s m + \mathbb{E}\left[J^*\left(\frac{\tilde{\pi} \cdot \phi_1(\mathbf{x}^{(m)})}{\phi_2(\mathbf{x}^{(m)}; \tilde{\pi})}\right)\right]\right\} \tag{13}$$

Note that J^* gives the cost defined in Eqn. 6. Thus, the optimal policy μ^* that achieves J^* gives τ^* and $M_k^*,\ k = 1, 2, \cdots, \tau^*$. The existence of the optimal policy is shown in the following theorem.

Theorem 1: As the state space is a Borel space, the action space is compact, the transition kernel is strongly continuous, and the single stage and terminal cost functions are bounded continuous functions, an optimal policy exists.

Proof: See Chap. 3, Vol. II of [2] and [3]. ■

We now prove some properties of the *minimum* total cost function J^*.

Theorem 2: The total cost function J^* is concave.

Proof: See Appendix I. ■

Theorem 3: The optimal stopping rule for the quickest change-detection problem with a dynamic control on M_{k+1}, the number of sensors in the *wake* state at time slot $k+1$ is a threshold based rule, where the threshold is on posterior probability of change.

Proof: See Appendix II. ■

[1] Unless explicitly stated, all the expectations are taken with respect to the pdf $\phi_2(\mathbf{x}^{(m)}; \tilde{\pi})$.

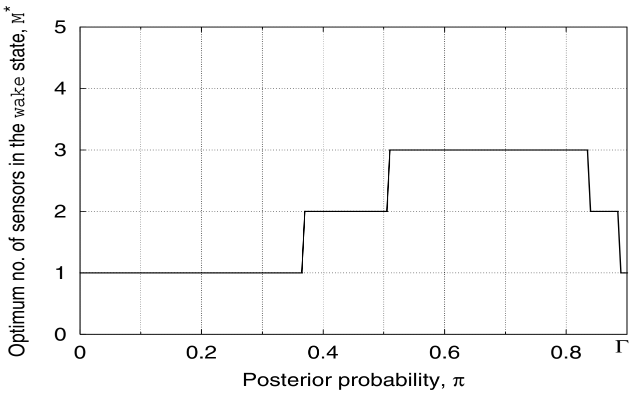

Fig. 1. Optimum number of sensors in the *wake* state M^* for $n = 10$ sensors, $\lambda_f = 100.0$, $\lambda_s = 0.5$, $f_0 \sim \mathcal{N}(0,1)$ and $f_1 \sim \mathcal{N}(1,1)$. Note that $\Gamma = 0.9$ corresponds to the threshold.

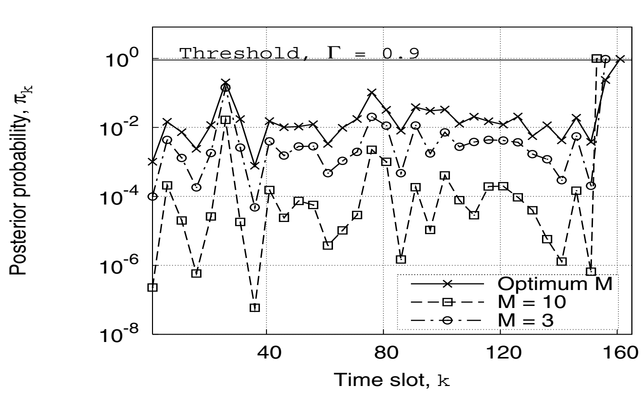

Fig. 2. A sample run of *event detection* with $n = 10$ sensors, $\lambda_f = 100.0$, $\lambda_s = 0.5$, $f_0 \sim \mathcal{N}(0,1)$ and $f_1 \sim \mathcal{N}(1,1)$.

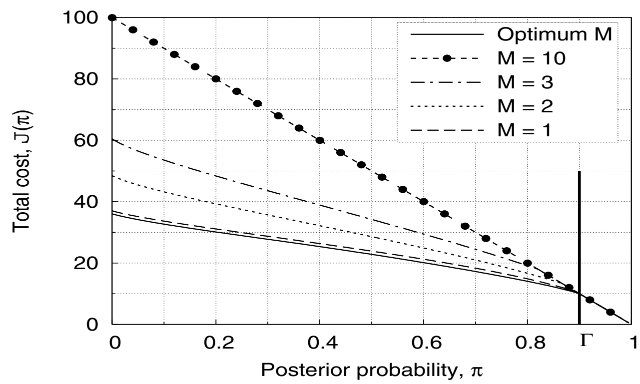

Fig. 3. Total cost $J(\pi)$ for $n = 10$ sensors, $\lambda_f = 100.0$, $\lambda_s = 0.5$, $f_0 \sim \mathcal{N}(0,1)$ and $f_1 \sim \mathcal{N}(1,1)$. Note that the threshold corresponding to $M = 1$ is 0.895, for $M = 2$ is 0.870, for $M = 3$ is 0.825, and for M^* is $\Gamma = 0.9$.

Note that Theorem 3 addresses only the *stopping time* part of the optimal policy μ^*. We now explore the structure of the optimal closed loop control policy for $M^* : [0,1] \to \mathbb{Z}_+$, the optimal number of sensors in the *wake* state in the *next* time slot. We choose $M_{k+1} = M^*_{k+1} = M^*(\Pi_k)$. We define the functions $B^{(m)}_{J^*} : [0,1] \to \mathbb{R}_+$ and $A^{(m)}_{J^*} : [0,1] \to \mathbb{R}_+$ as

$$\begin{aligned} B^{(m)}_{J^*}(\tilde{\pi}) &:= \mathbb{E}_{\phi_2(\mathbf{X}^{(m)};\tilde{\pi})}\left[J^*\left(\frac{\tilde{\pi}\cdot\phi_1(\mathbf{X}^{(m)})}{\phi_2(\mathbf{X}^{(m)};\tilde{\pi})}\right)\right], \\ \text{and } A^{(m)}_{J^*}(\tilde{\pi}) &:= \lambda_s m + B^{(m)}_{J^*}(\tilde{\pi}). \end{aligned}$$

Theorem 4: For any $\tilde{\pi} \in [0,1]$, the cost-to-go function $B^{(m)}_{J^*}(\tilde{\pi})$ monotonically decreases with m.

Proof: See Appendix III. ■

Note that $B^{(m)}_{J^*}(\tilde{\pi})$ can be thought of as the cost–to–go function for sensors used in the next slot. Thus, Theorem 4 states that the cost–to–go reduces as the number of observations increases. The intuition behind this is that $B^{(m)}_{J^*}(\tilde{\pi})$ represents the *uncertainty* about the *event* left at the end of the current time slot and more sensors resolve the uncertainty in a better way.

Numerical Results: We consider the following scenario: the change–time $T \sim geometric(0.01)$, $\pi_0 = 0$, $f_0^{(i)} \sim \mathcal{N}(0,1)$ and $f_1^{(i)} \sim \mathcal{N}(1,1)$. We set the cost per observation per sensor, λ_s to 0.5 and the cost of false alarm, λ_f to 100.0 (this sets PFA to 0.04). We consider $n = 10$ sensors. We compute M^* (from the optimal policy μ^* given by Eqn.12) by the *value iteration* algorithm [2],[3] and plot in Fig. 1. We note that in any time slot, it is not economical to use more than 3 sensors (though we have 10 sensors). Also, from Fig. 1, it is clear that M^* increases monotonically for $\pi < 0.6$ and then decreases monotonically for $\pi \geq 0.6$. Also note that, the region $\pi \in [0.5,\ 0.82]$ requires many sensors for optimal detection whereas the region $[0.0,\ 0.3] \cup [0.9,\ 1.0]$ requires the least number of sensors. This is due to the fact that *uncertainty* is more in the region $\pi \in [0.5,\ 0.82]$ whereas it is less in the region $[0.0,\ 0.3] \cup [0.9,\ 1.0]$.

In our numerical experiment, the event occurs at $T = 152$. When $M_{k+1}(\pi) = M^*(\pi)$ (taken from Fig. 1), we see that the detection happens at $\tau_{M^*} = 161$. When $M_{k+1}(\pi) = 10$ sensors (no *sleep* scheduling), we find the detection epoch to be $\tau_{10} = 153$. When $M_{k+1} = 3$ sensors (we chose 3 because $M^* \leq 3$), the stopping happens at $\tau_3 = 156$. From the above stopping times, it is clear that the detection delay does not vary significantly in the above three cases. We plot the trajectory of a sample path of Π_k versus the time slot k in Fig. 2. A close look at the Fig. 2 shows that all the three strategies follow almost the same trend. Also, we see from Fig. 2, that the π_k trajectory corresponding to $M_{k+1}(\pi) = 10$ (and $M_{k+1}(\pi) = 3$) gives more reliable information about the event than the π_k trajectory corresponding to $M_{k+1}(\pi) = M^*$. We also plot the total cost function $J(\pi)$ for the above cases in Fig. 3. Though the detection delays do not vary much, the total cost varies significantly. This is because the event happens at time slot $T = 152$. In the case of $M_{k+1} = M^*$, it is clear from Figs. 1 and 2 that only one sensor is used for the first 158 time slots. This reduces the cost by 10 times compared to the case of $M_{k+1} = 10$ (in this sample path) and about 3 times compared to the case of $M_{k+1} = 3$ (in this sample path). We note from Fig. 3, that it is better to keep 3 sensors active all the time than keeping 10 sensors active all the time. Note that $M_{k+1} = 1$ achieves a small cost but a large detection delay.

Define the differential cost $d : \{1, 2, \cdots, n\} \to \mathbb{R}_+$ as

$$d(m;\pi) = B_{J^*}^{(m-1)}(\tilde{\pi}) - B_{J^*}^{(m)}(\tilde{\pi}) \quad (14)$$

Recall that $\tilde{\pi} = \pi + (1-\pi)p$. Note that d is bounded, continuous, and $d(\cdot\,;1) = 0$. We are interested in $d(m;\pi)$ for $\pi \in [0,\ \Gamma)$. In Fig. 4, we plot the differential cost function $d(m;\cdot)$ given in Eqn. 14 for $m = 1, 2$ and 3. We observe that $d(m;\pi)$ monotonically decreases in m, for each $\pi \in [0, \Gamma)$ (i.e., $d(1;\pi) \geq d(2;\pi) \geq d(3;\pi)$). We observe this monotonicity property for different sets of experiments for the case when f_0 and f_1 belong to the Gaussian class of distributions. We hypothesize this monotonicity property of d and state the following theorem which gives a *structure* for M^*, the optimal number of sensors in the *wake* state.

Theorem 5: If for each $\pi \in [0, \Gamma)$, $d(m;\pi)$ decreases monotonically in m, then the optimal number of sensors in the *wake* state, $M^* : [0,1] \to \{0, 1, \cdots, n\}$ is given by

$$M^*(\pi) = \max\left\{m : d(m;\pi) \geq \lambda_s\right\}$$

Proof: Eqn. 13 and the monotone property of $d(m;.)$ proves the theorem. ∎

The above theorem is evident from Fig. 4.

B. Control on q_{k+1}, the probability that a sensor is in the wake *state at time slot $k+1$*

Here, we consider the scenario, where at time slot k, the fusion center (controller) upon receiving $\mathbf{X}_k^{(M_k)}$ makes a decision D_k on *stop*ping ($D_k = 1$) or to *continue* sampling ($D_k = 0$). If $D_k = 0$, the fusion center computes the optimal $q_{k+1} \in [0,1]$, the probability of a sensor node being *awake* at slot $k+1$. Thus the set of controls at time slot k is given by $\mathbf{A} = \left\{\mathsf{stop}, \cup_{q\in[0,1]}(\mathsf{continue}, q)\right\} = \left\{1, \cup_{q\in[0,1]}(0,q)\right\}$.

When the control $u_k = (0, q_{k+1})$ is chosen, M_{k+1}, the number of sensors in the *wake* state at time slot $k+1$ is *Bernoulli* distributed with parameters (n, q_{k+1}). Let $\gamma_m(q_{k+1})$ be the probability that m sensors are in the *wake* state at time slot $k+1$. $\gamma_m(q_{k+1})$ is given by

$$\gamma_m(q_{k+1}) = \binom{n}{m} q_{k+1}^m (1-q_{k+1})^{n-m} \quad (15)$$

The information state at time slot $k+1$ can be computed recursively as

$$\Pi_{k+1} := \Phi\Big(\Pi_k, (0, M_{k+1})\Big) = \frac{\widetilde{\Pi}_k \phi_1\Big(\mathbf{X}_{k+1}^{(M_{k+1})}\Big)}{\phi_2\Big(\mathbf{X}_{k+1}^{(M_{k+1})}; \widetilde{\Pi}_k\Big)} \quad (16)$$

where $\widetilde{\Pi}_k$, $\mathbf{X}_k^{(m)}$, $\phi_1(\mathbf{X}_k^{(m)})$, $\phi_2(\mathbf{X}_k^{(m)}; \widetilde{\Pi})$ are the same as those defined in Eqn. 10. When the control $u_k = 1$ is chosen, Π_{k+1} is given as

$$\Pi_{k+1} := \Phi(\Pi_k, 1) := \mathsf{a} \quad \text{w.p.1} \quad (17)$$

From Eqns. 16, 17, it is clear that the $\{(\Pi_k, M_k)\}$ process is a controlled Markov chain, the state space of the chain being

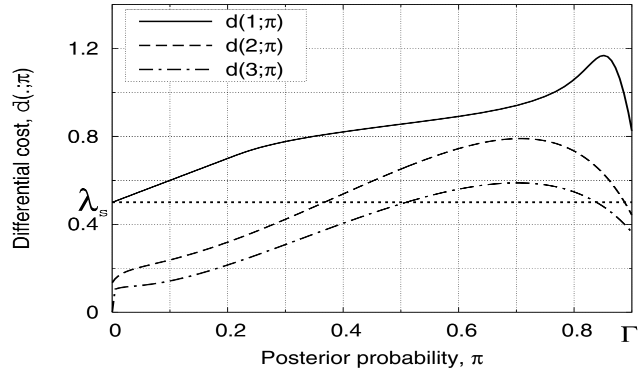

Fig. 4. Differential costs $d(\cdot;\pi)$, for $n = 10$ sensors, $\lambda_f = 100.0$, $\lambda_s = 0.5$, $f_0 \sim \mathcal{N}(0,1)$ and $f_1 \sim \mathcal{N}(1,1)$.

$\mathbf{S} \times \{0, 1, \cdots, n\}$. From Eqn. 7, it is clear that the single stage cost when the information state is π and the action is u, is

$$c(\pi, u) = \begin{cases} \lambda_f(1-\pi), & \pi \in [0,1], u = 1 \\ \pi + \lambda_s nq, & \pi \in [0,1], u = (0,q) \\ 0, & \pi = \mathsf{a} \end{cases}$$

Therefore, the optimization problem defined in Eqn. 7 can be stated as

$$J^*(\pi_0) = \min_{u_0, u_1, u_2, \cdots} \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, u_k) \,\middle|\, \Pi_0 = \pi_0\right]$$

We use the theory of MDP to solve the above optimization problem. From [2], it is clear that we can find a stationary optimal policy which solves the above optimization problem. Recall that the state space of the MDP is $\mathbf{S} \times \{0, 1, \cdots, n\}$. Since, given Π_k and u_k, (Π_{k+1}, M_{k+1}) is independent of M_k, it is easy to see that the optimal stationary policy does not depend on M_k. Let $\mu : \mathbf{S} \to \mathbf{A}$ be a stationary policy and $\mu^* : \mathbf{S} \to \mathbf{A}$ be the optimal stationary policy. Hence, we have

$$\begin{aligned} J_\mu(\pi_0) &= \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right] \\ \text{and } J^*(\pi_0) &= \min_{\mu} \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu^*(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right] \end{aligned}$$

Note that the optimal stationary policy μ^* is independent of the initial state $\Pi_0 = \pi_0$ [2]. By Bellman's equation, we have

$$J^*(\pi) = \min_{u\in\mathbf{A}} \left\{ c(\pi, u) + \sum_{m=0}^{n} \gamma_m(q) \mathbb{E}\Big[J^*\Big(\Phi(\pi, u)\Big)\Big]\right\}$$

Since $J^*(\mathsf{a}) = 0$, the above equation can be written as

$$J^*(\pi) = \min\left\{\lambda_f \cdot (1-\pi), \pi + A_{J^*}(\tilde{\pi})\right\}. \quad (18)$$

where the function $A_{J^*} : [0,1] \to \mathbb{R}_+$ is defined as

$$A_{J^*}(\tilde{\pi}) = \min_{q\in[0,1]} \left\{\lambda_s nq + \sum_{m=0}^{n} \gamma_m(q) \mathbb{E}\left[J^*\left(\frac{\tilde{\pi}\cdot\phi_1(\mathbf{x}^{(m)})}{\phi_2(\mathbf{x}^{(m)}; \tilde{\pi})}\right)\right]\right\}.$$

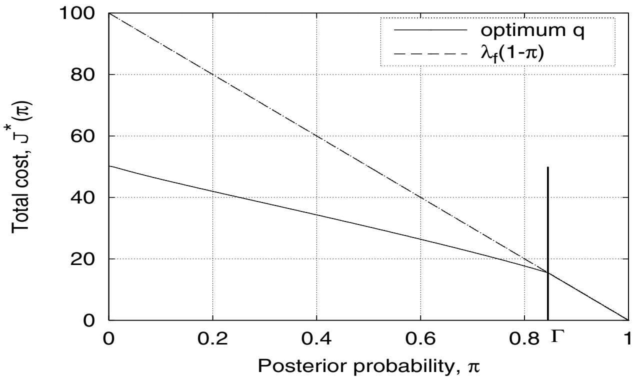

Fig. 5. Total cost $J^*(\pi)$ for $n = 10$ sensors, $\lambda_f = 100.0$, $\lambda_s = 0.5$, $f_0 \sim \mathcal{N}(0, 1)$ and $f_1 \sim \mathcal{N}(1, 1)$.

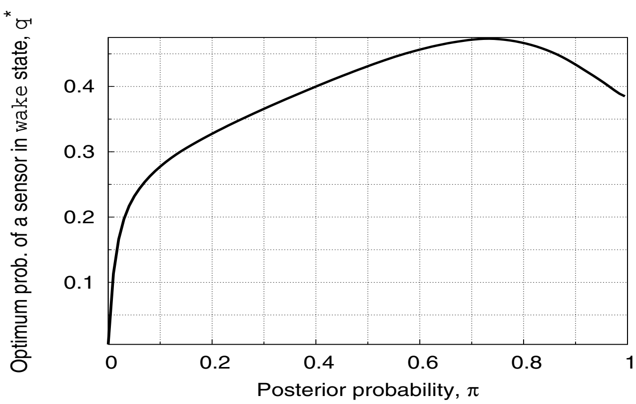

Fig. 6. Optimum probability of a sensor in the *wake* state, $q^*(\pi)$ for $n = 10$ sensors, $\lambda_f = 100.0$, $\lambda_s = 0.5$, $f_0 \sim \mathcal{N}(0, 1)$ and $f_1 \sim \mathcal{N}(1, 1)$.

Note that J^* gives the cost defined in Eqn. 7. Thus, the optimal policy μ^* that achieves J^* gives τ^* and $q_k^*,\ k = 1, 2, \cdots, \tau^*$. The existence of the optimal policy and its structure is shown in the following theorems.

Theorem 6: As the state space is a Borel space, the action space is compact, the transition kernel is strongly continuous, and the single stage and terminal cost functions are bounded continuous functions, an optimal policy exists.

Proof: See Chap. 3, Vol. II of [2] and [3]. ■

Theorem 7: The total cost function J^* is concave.

Proof: Follows from Appendix I. ■

Theorem 8: The optimal stopping rule for the quickest change-detection problem with a dynamic control on q_{k+1}, the probability that a sensor is in the *wake* state at time slot $k+1$ is a threshold based rule, where the threshold is on posterior probability of change.

Proof: See Appendix II. ■

Numerical Results: We consider the same scenario as in the case of control on M_{k+1}. We plot the total cost $J^*(\pi)$ in Fig. 5. We also plot the optimum probability of a sensor in the *wake* state, $q^*(\pi)$ in Fig. 5. **We observe that $q^*(\pi)$ is concave in π.** This agrees well with the intuition for the control on M_{k+1}.

IV. Quickest Change Detection without Feedback

In this section, we study the *sleep/wake* scheduling problem defined in Eqn. 8. Open loop control is applicable to the systems in which there is no feedback channel from the fusion center (controller) to the sensors. Here, at any time slot k, a sensor chooses to be in the *wake* state with probability q independent of other sensors. Hence, $\{M_k\}$, the number of sensors in the *wake* state at time slot k is i.i.d. *Bernoulli distributed* with parameters (n, q). Let γ_m be the probability that m sensors are in the *wake* state. γ_m is given by

$$\gamma_m = \binom{n}{m} q^m (1-q)^{n-m} \tag{19}$$

We can choose q that minimizes the Bayesian cost given by Eqn. 8.

At time slot k, the fusion center receives a vector of observation $\mathbf{X}_k^{(M_k)}$ and computes Π_k. Recall that $\{\Pi_0, \Pi_1, \cdots, \Pi_k\}$ is a sufficient statistic for the sensor observations up to time slot k [4]. In the open loop scenario, the state space is $\mathbf{S} = \big\{[0,1] \cup \{\mathsf{a}\}\big\}$ (we always work with the Borel space, $(\mathbf{S}, \mathscr{B}(\mathbf{S}))$). The set of actions is given by $\mathbf{A} = \{\mathsf{stop}, \mathsf{continue}\} = \{1, 0\}$ where '1' represents *stop* and '0' represents *continue*.

Note that given a control u_k, Π_{k+1} can be recursively computed in the same way as shown in Eqns. 16, 17. Thus, $\{(\Pi_k, M_k)\},\ k \in \mathbb{Z}_+$ is a controlled Markov chain. From Eqn. 8, it is clear that when the information state is π, and the action is u, then the single stage cost $c(\pi, u)$ is given by

$$c(\pi, u) = \begin{cases} \lambda_f(1-\pi), & \pi \in [0,1], u = 1 \\ \pi + \lambda_s n q, & \pi \in [0,1], u = 0 \\ 0, & \pi = \mathsf{a} \end{cases}$$

Therefore, the optimization problem defined in Eqn. 8 can be stated as

$$J^*(\pi_0) = \min_{u_0, u_1, u_2, \cdots} \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, u_k) \,\middle|\, \Pi_0 = \pi_0\right]$$

We use the theory of MDP to solve the above optimization problem. From [2], it is clear that we can find a stationary optimal policy which solves the above optimization problem. Note that the state space of the MDP is $\mathbf{S} \times \{0, 1, \cdots, n\}$. Since, $\{M_k\}$ is an i.i.d. process, given Π_k and u_k, (Π_{k+1}, M_{k+1}) is independent of M_k. Hence, the optimal stationary policy does not depend on M_k. Let $\mu : \mathbf{S} \to \mathbf{A}$ be a stationary policy and $\mu^* : \mathbf{S} \to \mathbf{A}$ be the optimal stationary policy. Hence, we have

$$\begin{aligned} J_\mu(\pi_0) &= \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right] \\ \text{and } J^*(\pi_0) &= \min_{\mu} \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right] \\ &= \mathbb{E}\left[\sum_{k=0}^{\infty} c(\Pi_k, \mu^*(\Pi_k)) \,\middle|\, \Pi_0 = \pi_0\right] \end{aligned}$$

By Bellman's equation, we have

$$J^*(\pi) = \min_{u \in \mathbf{A}} \left\{ c(\pi, u) + \sum_{m=0}^{n} \gamma_m \mathbb{E}\Big[J^*\Big(\Phi(\pi, u)\Big)\Big] \right\}$$

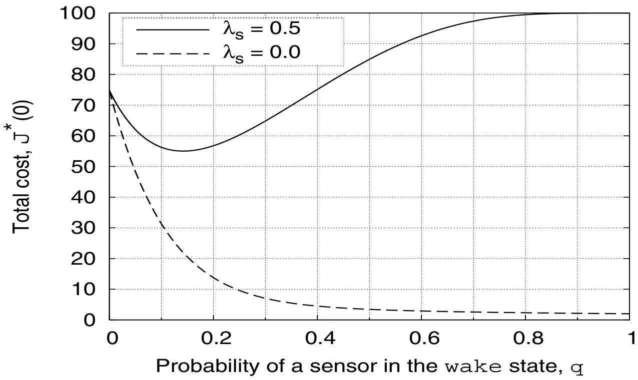

Fig. 7. Total cost $J^*(0)$ for $n = 10$ sensors, $\lambda_f = 100.0$, $f_0 \sim \mathcal{N}(0,1)$ and $f_1 \sim \mathcal{N}(1,1)$.

Since $J^*(\mathsf{a}) = 0$, the above equation can be written as

$$J^*(\pi) \quad = \quad \min\left\{\lambda_f \cdot \big(1-\pi\big), \pi + A_{J^*}\big(\tilde{\pi}\big)\right\}. \qquad (20)$$

where the function $A_{J^*} : [0,1] \to \mathbb{R}_+$ is defined as

$$A_{J^*}(\tilde{\pi}) = \lambda_s n q + \sum_{m=0}^{n} \gamma_m \mathbb{E}\left[J^*\left(\frac{\tilde{\pi}\cdot\phi_1(\mathbf{X}^{(m)})}{\phi_2(\mathbf{X}^{(m)};\tilde{\pi})}\right)\right].$$

Note that J^* gives the cost defined in Eqn. 8. Thus, the optimal policy μ^* that achieves J^* gives τ^*.

Theorem 9: As the state space is a Borel space, the action space is compact, the transition kernel is strongly continuous, and the single stage and terminal cost functions are bounded continuous functions, an optimal policy exists.

Proof: See Chap. 3, Vol. II of [2] and [3]. ■

We now prove some properties of the optimal policy.

Theorem 10: The total cost function J^* is concave.

Proof: Follows from Appendix I. ■

Theorem 11: The optimal stopping rule for the quickest change-detection problem with the open loop control on the number of sensors in the *wake* state is a threshold based rule, where the threshold is on posterior probability of change.

Proof: See Appendix II. ■

Numerical Results: We consider a sensor network of $n = 10$ nodes. We consider the following models, $T \sim geometric(0.01)$, $\pi_0 = 0$, $f_0^{(i)} \sim \mathcal{N}(0,1)$ and $f_1^{(i)} \sim \mathcal{N}(1,1)$. We set λ_f to 100.0 (this sets the PFA to 0.04, approximately). This is the same scenario as before. We obtain $J^*(0)$ for various values of q and plotted in the Fig. 7. We obtain the plot for $\lambda_s = 0.5$ and for $\lambda_s = 0.0$. Note that when $\lambda_s > 0$, for low values of q, the detection delay cost dominates over the observation costs in $J^*(0)$ and for high values of q, the observation costs dominate over the detection delay cost. Thus, there is a trade–off between the detection delay cost and the observation costs as q varies. This is captured in the Fig. 7. Note that the Bayesian cost is optimal at $q = 0.15$. When $\lambda_s = 0$, as q increases the detection delay decreases. Hence, we see the monotonically decreasing trend for $\lambda_s = 0.0$.

V. Summary

In this paper, we formulated the problem of *sleep/wake* scheduling in a sensor network that minimizes the detection delay by optimal use of sensing/communication resources. Recall that we have set out to solve the problem in Eqn. 5. We have derived the optimal control for three approaches using the theory of MDP. We showed the existence of the optimal policy and obtained some structural results.

From Figs. 3, 5, and 7, we note that the total cost $J(\pi)$ is the least for optimal control on M_{k+1}. Also, we note that in the open loop control case, the least total cost $J^*(0) = 55$ is achieved when the attempt probability, q is 0.15 (this corresponds to an average duty cycle of 0.15). It is to be noted that this cost is larger than that achieved by the optimal closed loop policies ($J^*(0) = 50$ for the closed loop control on q_{k+1} and $J^*(0) = 38$ for the closed loop control on M_{k+1}). From Figs. 2 and 1, we see that when $M_{k+1}(\pi) = M^*(\pi)$, the switching of the sensors between *sleep* and *wake* states happen only in 2 slots out of 161 slots.

We formulated the problem and showed the optimal policy in the context of *centralized detection*. It is easy to extend our results to a *decentralized detection* setting, [8]. As a future work, we also plan to provide some simple heuristic policies.

Appendix - I

Proof of Theorem 2

We use the following Lemma to prove Theorem 2.

Lemma 1: If $f : [0,1] \to \mathbb{R}$ is concave, then the function $h : [0,1] \to \mathbb{R}$ defined by

$$h(y) = \mathbb{E}_{\phi_2(\mathbf{X}^{(m)};y)}\left[f\left(\frac{y\phi_1(\mathbf{x}^{(m)})}{y\phi_1(\mathbf{x}^{(m)}) + (1-y)\phi_0(\mathbf{x}^{(m)})}\right)\right]$$

is concave for any m, where $\phi_1(\mathbf{x}^{(m)})$ and $\phi_0(\mathbf{x}^{(m)})$ are pdfs on $\mathbf{X}^{(m)}$, $0 \le y \le 1$, and $\phi_2(\mathbf{x};y) = y\phi_1(\mathbf{x}) + (1-y)\phi_0(\mathbf{x})$.

Proof: Define the function $h_1 : [0,1] \to \mathbb{R}$ as

$$h_1(y;\mathbf{x}) := f\left(\frac{y\phi_1(\mathbf{x})}{y\phi_1(\mathbf{x}) + (1-y)\phi_0(\mathbf{x})}\right)\Big[y\phi_1(\mathbf{x}) + (1-y)\phi_0(\mathbf{x})\Big].$$

As $\mathsf{T} := \int \cdots \, d\mathbf{x}$ is a linear operator and $h(y) = \mathsf{T}h_1(y;\mathbf{x})$, it is sufficient to show that $h_1(y;\mathbf{x})$ is concave in y. If $f(y)$ is concave then

$$f(y) = \inf_{(a_i,b_i)\in I}\left\{a_i y + b_i\right\}$$

where $I = \{(a,b) \in \mathbb{R}^2 : ay + b \ge f(y), y \in [0,1]\}$. Hence,

$$\begin{aligned}
&h_1(y;\mathbf{x})\\
&= f\left(\frac{y\phi_1(\mathbf{x})}{y\phi_1(\mathbf{x}) + (1-y)\phi_0(\mathbf{x})}\right)\Big[y\phi_1(\mathbf{x}^{(m)}) + (1-y)\phi_0(\mathbf{x}^{(m)})\Big]\\
&= \inf_{(a_i,b_i)\in I}\left\{a_i\left(\frac{y\phi_1(\mathbf{x})}{y\phi_1(\mathbf{x}) + (1-y)\phi_0(\mathbf{x})}\right) + b_i\right\}\Big[y\phi_1(\mathbf{x}) + (1-y)\phi_0(\mathbf{x})\Big]\\
&= \inf_{(a_i,b_i)\in I}\left\{a_i y\phi_1(\mathbf{x}) + b_i\Big[y\phi_1(\mathbf{x}) + (1-y)\phi_0(\mathbf{x})\Big]\right\}\\
&= \inf_{(a_i,b_i)\in I}\left\{\Big((a_i+b_i)\phi_1(\mathbf{x}) - b_i\phi_0(\mathbf{x})\Big)y + b_i\phi_0(\mathbf{x})\right\}
\end{aligned}$$

This implies that $h_1(y;\mathbf{x}^{(m)})$ is concave in y. ■

In the *value iteration*, the finite K–horizon cost–to–go function, $J_K^K(\pi) = \lambda_f \cdot \big(1-\pi\big)$ is concave. Hence, by lemma 1, we see that the cost–to–go functions $J_{K-1}^K(\pi)$, $J_{K-2}^K(\pi)$, $\cdots, J_0^K(\pi)$ are concave. Hence for $0 \le \lambda \le 1$,

$$\begin{aligned}
J^*(\pi) &= \lim_{K\to\infty} J_0^K(\pi) \\
J^*(\lambda\pi_1 + (1-\lambda)\pi_2) &= \lim_{K\to\infty} J_0^K\Big(\lambda\pi_1 + (1-\lambda)\pi_2\Big) \\
&\ge \lim_{K\to\infty} \lambda J_0^K(\pi_1) + \lim_{K\to\infty}(1-\lambda)J_0^K(\pi_2) \\
&= \lambda J^*(\pi_1) + (1-\lambda)J^*(\pi_2)
\end{aligned}$$

It follows that $J^*(\pi)$ is concave. ■

APPENDIX - II

Proof of Theorems 3, 8 and 11

Define the maps $C : [0,1] \to \mathbb{R}_+$ and $H : [0,1] \to \mathbb{R}_+$, as

$$\begin{aligned}
C(\pi) &:= \lambda_f \cdot \big(1-\pi\big) \\
H(\pi) &:= \pi + A_{J^*}(\tilde{\pi})
\end{aligned}$$

Note that $C(1) = 0$, $H(1) = 1$, $C(0) = \lambda_f$ and $H(0) = A_{J^*}(p)$. In Theorem 3, we have

$$\begin{aligned}
& A_{J^*}(p) \\
= & \min_{0\le m\le n}\left\{\lambda_s m + \mathbb{E}_{\phi_2(\mathbf{X}^{(m)};p)}\left[J^*\left(\frac{p\cdot\phi_1(\mathbf{X}^{(m)})}{\phi_2(\mathbf{X}^{(m)};p)}\right)\right]\right\} \\
\le & \min_{0\le m\le n}\left\{\lambda_s m + J^*\left(\mathbb{E}_{\phi_2(\mathbf{X}^{(m)};p)}\left[\frac{p\cdot\phi_1(\mathbf{X}^{(m)})}{\phi_2(\mathbf{X}^{(m)};p)}\right]\right)\right\} \\
= & \min_{0\le m\le n}\left\{\lambda_s m + J^*\left(p\right)\right\} \\
= & J^*\left(p\right) \\
\le & \lambda_f\cdot\big(1-p\big)
\end{aligned}$$

and in Theorems 8 and 11, we have

$$\begin{aligned}
& A_{J^*}(p) \\
= & \sum_{m=0}^{n}\gamma_m\mathbb{E}_{\phi_2(\mathbf{X}^{(m)};p)}\left[J^*\left(\frac{p\cdot\phi_1(\mathbf{X}^{(m)})}{\phi_2(\mathbf{X}^{(m)};p)}\right)\right] \\
\le & \sum_{m=0}^{n}\gamma_m J^*\left(\mathbb{E}_{\phi_2(\mathbf{X}^{(m)};p)}\left[\frac{p\cdot\phi_1(\mathbf{X}^{(m)})}{\phi_2(\mathbf{X}^{(m)};p)}\right]\right) \\
= & \sum_{m=0}^{n}\gamma_m J^*\left(p\right) \\
= & J^*\left(p\right) \\
\le & \lambda_f\cdot\big(1-p\big).
\end{aligned}$$

The inequality in the second step is justified using Jensen's inequality and the inequality in the last step follows from the definition of J^*.

Note that $H(1) - C(1) > 0$ and $H(0) - C(0) < 0$. As the function $H(\pi) - C(\pi)$ is concave, by the *intermediate value theorem*, there exists $\Gamma \in [0,1]$ such that $H(\Gamma) = C(\Gamma)$. This Γ is unique as $H(\pi) = C(\pi)$ for at most two values of π. If in the interval $[0,1]$, there are two distinct values of π for which $H(\pi) = C(\pi)$, then the signs of $H(0)-C(0)$ and $H(1)-C(1)$ should be the same. Hence, the optimal stopping rule is given by

$$\tau^* = \inf\left\{k : \Pi_k \ge \Gamma\right\}$$

where the threshold Γ is given by

$$\Gamma + A_{J^*}(\Gamma) = \lambda_f \cdot \big(1-\Gamma\big) \qquad ■$$

APPENDIX - III

Proof of Theorem 4

Define

$$\begin{aligned}
\phi_j(\mathbf{x}^{(m)}) &:= \prod_{i=1}^{m} f_j(x^{(i)}),\ j = 0,1. \\
\mathbf{x}^{(l)} &:= (x^{(1)}, x^{(2)}, \cdots, x^{(m)}, x^{(m+1)}, \cdots, x^{(l)}) \\
\mathbf{u} &:= (x^{(1)}, x^{(2)}, \cdots, x^{(m)}) \\
\mathbf{v} &:= (x^{(m+1)}, x^{(m+2)}, \cdots, x^{(l)}) \\
\hat{\pi} &:= \frac{\tilde{\pi}\phi_1(\mathbf{u})}{\tilde{\pi}\phi_1(\mathbf{u}) + (1-\tilde{\pi})\phi_0(\mathbf{u})}
\end{aligned}$$

Note that

$$\begin{aligned}
& B_{J^*}^{(l)}(\tilde{\pi}) \\
= & \int_{\mathbb{R}^l} J^*\left(\frac{\tilde{\pi}\cdot\phi_1(\mathbf{x}^{(l)})}{\phi_2(\mathbf{x}^{(l)};\tilde{\pi})}\right)\Big[\phi_2(\mathbf{x}^{(l)};\tilde{\pi})\Big]\ d\mathbf{x}^{(l)} \\
= & \int_{\mathbb{R}^m}\int_{\mathbb{R}^{l-m}} J^*\left(\frac{\hat{\pi}\phi_1(\mathbf{v})}{\phi_2(\mathbf{v};\hat{\pi})}\right)\phi_2(\mathbf{v};\hat{\pi})\ d\mathbf{v}\phi_2(\mathbf{u};\tilde{\pi})\ d\mathbf{u} \\
\le & \int_{\mathbb{R}^m} J^*\left(\int_{\mathbb{R}^{l-m}}\frac{\hat{\pi}\phi_1(\mathbf{v})}{\phi_2(\mathbf{v};\hat{\pi})}\Big[\phi_2(\mathbf{v};\hat{\pi})\Big]d\mathbf{v}\right)\phi_2(\mathbf{u};\tilde{\pi})\ d\mathbf{u} \\
= & \int_{\mathbb{R}^m} J^*\left(\hat{\pi}\right)\phi_2(\mathbf{u};\tilde{\pi})d\mathbf{u} \\
= & B_{J^*}^{(m)}(\tilde{\pi})
\end{aligned}$$

As J^* is concave, the inequality in the second line follows from Jensen's inequality. Hence proved. ■

REFERENCES

[1] S. Appadwedula, V. V. Veeravalli, and D. L. Jones, “Energy-efficient detection in sensor networks,” *IEEE J. Sel. Areas Commun.*, vol. 23, no. 4, pp. 693–702, April 2005.

[2] D. Bertsekas, *Dynamic Programming and Optimal Control, vol. I and II*. Belmont, Massachusetts: Athena Scientific, 1995.

[3] O. Hernández-Lerma and J. B. Lasserre, *Discrete–Time Markov Control Processes:Basic Optimality Criteria*. New York: Springer–Verlag, 1996.

[4] E. L. Lehmann and G. Casella, *Theory of Point Estimation*. New York: Springer–Verlag, 1998.

[5] V. K. Prasanthi and Anurag Kumar, “Optimizing delay in sequential change detection on ad hoc wireless sensor networks,” in *IEEE SECON*, VA, USA, Sep. 2006.

[6] C. Rago, P. Willet, and Y. Bar-Shalom, “Censoring sensors: A low communication-rate scheme for distributed detection,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 32, no. 2, pp. 554–568, April 1996.

[7] A. N. Shiryaev, *Optimal Stopping Rules*. New York: Springer, 1978.

[8] V. V. Veeravalli, “Decentralized quickest change detection,” *IEEE Trans. Inf. Theory*, vol. 47, no. 4, pp. 1657–1665, May 2001.

[9] Y. Wu, S. Fahmy, and N. B. Shroff, “Energy efficient sleep/wake scheduling for multi-hop sensor networks: Non-convexity and approximation algorithm,” in *Infocom*, Alaska, USA, May 2007.

[10] L. Zacharias and R. Sundaresan, “Decentralized sequential change detection using physical layer fusion,” in *IEEE ISIT*, Nice, France, Jun. 2007.