2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Stochastic Surveillance Strategies for Spatial Quickest Detection

Vaibhav Srivastava Francesco Bullo

Abstract— **We present stochastic vehicle routing policies for detection of any number of anomalies in a set of regions of interest. The autonomous vehicle collects information from a set of regions and sends it to a fusion center. The vehicle follows a randomized region selection policy at each iteration. Using the collected information, the fusion center runs an ensemble of cumulative sum (CUSUM) algorithms in order to detect the presence of an anomaly in any region. We first determine optimal stationary policies that result in quickest detection of all anomalies. We then study an adaptive policy that assigns higher selection probability to a region with higher likelihood of an anomaly. We provide a comparative study of these policies.**

I. INTRODUCTION

Recent years have witnessed a surge in the application of autonomous agents like unmanned autonomous vehicles (UAVs) in various activities such as surveillance and information collection. In view of the recent Icelandic ash problem and the oil spill in the gulf of Mexico, quickest detection of anomalies is of considerable importance. Similar situation occurs in the case of wild fires. The extreme uncertainties in such situations call for surveillance strategies for quickest detection of anomalies. Generally, a limited number of UAVs are deployed to survey a large number of regions, and it becomes important that the UAVs collect the most pertinent information. Such scenarios motivate the characterization of the information and routing policies that result in quick information aggregation.

We study a spatial quickest detection problem: the simultaneous detection of anomalies at different regions. We adopt the Cumulative Sum (CUSUM) algorithm. In many situations the information collected by UAVs is sent to a human operator, and she decides on the presence of any anomaly. Recent advances in cognitive psychology [1], [13] show that human performance in decision making tasks, such as the two-alternative forced choice task, is well modeled by sequential statistical methods like CUSUM and the sequential probability ration test (SPRT). Roughly speaking, studying CUSUM algorithm for quickest detection may be appropriate even in situations where a human operator is making the decision.

Routing policies for UAVs have witnessed a lot of attention in the control literature. A survey on dynamic vehicle routing policies for servicing tasks is presented in [2]. Klein et al [9] present a vehicle routing policy for optimal localization of an acoustic source. They study the tradeoff between the Fisher information and the travel time. Quintero et al [12] minimized the error covariance to determine the optimal coordination policies for Dubin's vehicles. Zhang et al [23] study the estimation of environmental plumes with mobile sensors. They minimize the uncertainty of the estimate of the ensemble Kalman filter to determine the optimal trajectories for a swarm of mobile sensors. Gupta et al [7] determine trajectories for mobile sensors that minimize the error covariance of the Kalman filter estimate. Castañón [4] poses the search problem as a dynamic hypothesis test, and determines the optimal routing policy that maximizes the probability of detection of a target. Chung et al [5] study the probabilistic search problem in a decision theoretic framework. They present various search policies including sequential hypothesis tests. There has been a significant interest in design of policies for mixed human-automata systems. Savla et al [16] utilize models of human cognition to develop vehicle routing strategies for human in the loop systems. Certain optimal information aggregation strategies for human-automata system have been developed in [22], [21].

The problem of surveillance has received considerable attention recently. Grace et al [6] study stochastic surveillance strategies for multiple agents. They focus on local rules that minimize the computation time and communication. Pasqualetti et al [10] study the problem of optimal cooperative surveillance with multiple agents. They optimize the time gap between any two visits to the same region, and the time necessary to inform every agent about an event occurred in the environment. Smith et al [18] consider the surveillance of multiple regions with changing features and determine policies that minimize the maximum change in features between the observations. A persistent monitoring task where robots move on a given closed path has been considered in [19], and a speed controller has been designed that minimizes the penalty due to no surveillance of a region. Srivastava et al [20] present a stochastic surveillance strategy based on the Markov chain Monte Carlo method. Sak et al [15] present multi-agent routing and partitioning strategies for surveillance of regions with known intruder models. Hespanha et at [8] studied multi-agent probabilistic pursuit evasion game with the policy that, at each instant, directs pursuers to a location that maximizes the probability of finding an evader at that instant.

We study the optimal routing policies for a UAV performing surveillance. We consider a UAV which surveys a set of regions, collects the information and sends it to a fusion center. The fusion center, runs parallel CUSUM algorithms (one for each region) with the information collected and

This work has been supported in part by AFOSR MURI Award-FA9550-07-1-0528, by NSF Award CPS-1035917 and by ARO Award W911NF-11-1-0092.

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

978-1-61284-799-3/11/$26.00 ©2011 IEEE

decides on the presence of any anomaly. This setup also models the situations where the UAV surveys a region, collects evidence, and sends that to a fusion center where an operator processes it to detect any anomaly in any region. For a given stochastic routing policy, we determine the expected time the CUSUM algorithms at different regions take to detect any anomaly. We minimize the expected detection time over the policy space and thus, obtain the policy for quickest detection of any anomaly. The main contributions of this work are as follows:

i) We present a novel *ensemble* CUSUM algorithm to detect more than one anomaly simultaneously.
ii) We incorporate the travel time of the UAV to determine the decision times for the ensemble CUSUM algorithm.
iii) We show that the minimization of the detection delay in ensemble CUSUM algorithm is a non-convex problem, and provide probabilistic guarantees that it achieves a unique minimum.
iv) We present an adaptive vehicle routing policy for quickest detection of any number of anomalies.

The remainder of the paper is organized in the following way. We present some preliminaries in Section II. The problem set up is presented in Section III. We develop the ensemble CUSUM algorithm in Section IV. The optimal vehicle routing policies are derived in Section V. We elucidate on the ideas in the paper through some numerical examples in Section VI. Our conclusions are in Section VII.

II. Preliminaries

A. Kullback-Leibler divergence

Given two probability mass functions $f_1 : \mathcal{S} \to \mathbb{R}_{\geq 0}$ and $f_2 : \mathcal{S} \to \mathbb{R}_{\geq 0}$, where $\mathcal{S}$ is some countable set, the Kullback-Leibler divergence $\mathcal{D} : \mathcal{L}^1 \times \mathcal{L}^1 \to \mathbb{R} \cup \{+\infty\}$ 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 $\mathcal{L}^1$ is the set of integrable functions and $\text{supp}(f_1)$ 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 density functions.

Definition 1 (Finitely informative distribution): Given two probability distribution functions $f^0, f^1 : \mathbb{R} \to \mathbb{R}_{>0}$. f^1 is said to be finitely informative with respect to f^0, if $\mathcal{D}(f^1, f^0) \in \mathbb{R}_{>0}$. This implies that the two distributions are non-identical with some non-zero probability and there exists no sample that absolutely demarcates one distribution from other.

B. Cumulative sum algorithm

Given a set of observations $\{y_1, y_2, \ldots\}$ such that $\{y_1, \ldots, y_{\nu-1}\}$ are i.i.d. with probability density function f^0 and $\{y_\nu, y_{\nu+1}, \ldots\}$ are i.i.d. with probability density function f^1 with ν unknown. Let $\delta \geq \nu$ be the iteration at which the change is detected. The non-Bayesian quickest detection problem [11], [17] is posed as following

$$\begin{aligned} \text{minimize} \quad & \sup_{\nu \geq 1} \mathbb{E}_\nu[\delta - \nu + 1 | \delta \geq \nu] \\ \text{subject to} \quad & \mathbb{E}_{f^0}[\delta] \geq \gamma, \end{aligned} \tag{1}$$

where $\mathbb{E}_\nu[\cdot]$ represents expected value with respect to distribution of observation at iteration ν, $\gamma \in \mathbb{R}_{>0}$ is a large constant and is called *false alarm rate.*

The solution to the problem (1) is the cumulative sum (CUSUM) algorithm [11], which is described in Algorithm 1.

Algorithm 1 Cumulative Sum (CUSUM) Algorithm

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 $\Lambda_\tau := (\Lambda_{\tau-1} + \lambda_\tau)^+$
% decide only if the threshold is crossed
4: **if** $\Lambda_\tau > \eta$, **then** declare change is detected
5: **else** $(\Lambda_\tau \in [0, \eta[)$ continue sampling (step `1:`)

For a given threshold η, the false alarm rate and the worst expected detection delay for CUSUM algorithm are

$$\begin{aligned} \mathbb{E}_{f^0}(\delta) &\cong \frac{|e^\eta - \eta - 1|}{\mathcal{D}(f^0, f^1)}, \text{ and} \\ \mathbb{E}_{f^1}(\delta) &\cong \frac{|e^{-\eta} + \eta - 1|}{\mathcal{D}(f^1, f^0)}. \end{aligned} \tag{2}$$

The approximations in equation (2) are referred to as the Wald's approximations [17], and are known to be accurate for large values of threshold η. The case when $\eta \to +\infty$ is called the asymptotic case.

III. Problem setup

We consider surveillance of a set of disjoint regions $\mathcal{R} = \{1, \ldots, n\}$ by a UAV. The UAV, at each iteration, goes to a randomly selected region, collects information and sends it to a fusion center. We assume that the UAV moves with unit speed. We identify the fusion center with a human operator. The fusion center decides on the presence of anomalies in any region. To do so, the fusion center runs n parallel CUSUM algorithms (one for each region) with the collected observations. The objective of the fusion center is to detect all the anomalies in minimum time.

We adopt a *randomized* region selection policy. The UAV at each iteration picks a region $\ell \in \mathcal{R}$, with a stationary probability q_ℓ, and surveys it. The region selection process is, in fact, Markovian, but we are interested in the case where large number of observations are required. Since, the UAV can move from one region to another, the underlying Markov chain is irreducible. Thus, in this case the region selection process can be approximated by a stationary randomized process with the distribution corresponding to the stationary distribution of the Markov chain.

Let π_ℓ^1 be the prior probability of the anomaly being present at region $\ell \in \mathcal{R}$. We assume that the distance between two regions $i, j \in \mathcal{R}$ is $d_{ij} \in \mathbb{R}_{\geq 0}$. Let $f_\ell^1, f_\ell^0 : \mathbb{R} \to \mathbb{R}_{\geq 0}$

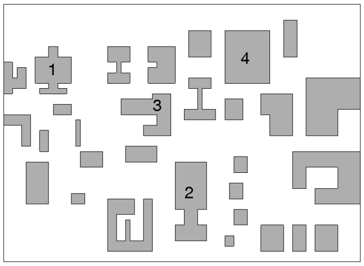

Fig. 1. A model of the UCSB campus map. A UAV surveys the regions shown. At each iteration the UAV surveys a region randomly and send the observation to a fusion center. The objective of the fusion center is to quickly detect any anomaly is any of the regions.

be the probability distribution functions at region ℓ in the presence and absence of an anomaly, respectively. Let $T_\ell \in \mathbb{R}_{>0}, \ell \in \mathcal{R}$ be the expected time that the UAV take to collect, process and send the information at each iteration at region ℓ. We study the quickest detection problem under following assumptions.

Finitely informative distributions: We assume that the conditional probability density functions f_ℓ^0, f_ℓ^1 at each region $\ell \in \mathcal{R}$ are finitely informative with respect to each other.

Conditionally independent observations: We assume that conditioned on the presence or absence of any anomaly in a region, the observations in that region are mutually independent. We further assume that the observations in different regions are also mutually independent.

Finite moments: We assume that at each region $\ell \in \mathcal{R}$, the statistic $\lambda = \log(f_\ell^1(Y)/f_\ell^0(Y))$ has finite first and second moments with respect to f_ℓ^0 and f_ℓ^1.

Notation: We denote the probability simplex in $\mathbb{R}^n$ by Δ_{n-1}. We denote the region selection probability vector by $\boldsymbol{q} = [\, q_1 \ \dots \ q_n \,]^T$.

IV. Randomized ensemble CUSUM

We assume that the fusion center runs n parallel CUSUM algorithms, one for each region. At each iteration, the UAV moves to a randomly chosen region k, and collects the observation. This observation is used to update the CUSUM statistic for the region k only. We refer to such parallel CUSUM algorithms by *randomized ensemble CUSUM*. We assume that the probability to choose region k is stationary and equal to q_k. We further assume that the threshold for each CUSUM algorithm is the same and is equal to $\eta \in \mathbb{R}_{\geq 0}$. The randomized ensemble CUSUM procedure is formally presented in Algorithm 2.

Theorem 2 (Randomized ensemble CUSUM): For the randomized ensemble CUSUM where the region $\ell \in \mathcal{R}$ is selected at each iteration with a stationary probability q_ℓ, the following statements hold:

i) The worst expected sample size for detecting the change at region ℓ satisfies

$$\mathbb{E}_{f_\ell^1}[\delta^\ell] \cong \frac{|e^{-\eta} + \eta - 1|}{q_\ell \mathcal{D}(f_\ell^1, f_\ell^0)}.$$

Algorithm 2 Randomized Ensemble CUSUM Algorithm

1: at time $\tau \in \mathbb{N}$, sample a region ℓ from the distribution q
2: collect sample y_τ from region ℓ
3: update the CUSUM statistic at each region

$$\Lambda_\tau^j = \begin{cases} \left(\Lambda_{\tau-1}^\ell + \log \frac{f_\ell^1(y_\tau)}{f_\ell^0(y_\tau)}\right)^+, & \text{if } j = \ell, \\ \Lambda_{\tau-1}^j, & \text{if } j \in \mathcal{R} \setminus \{\ell\}. \end{cases}$$

% decide only if the threshold is crossed
4: **if** $\Lambda_\tau^\ell > \eta$, **then** declare change detected at region ℓ
5: go to step `1:`

ii) The worst expected detection delay at region ℓ, T_δ^ℓ satisfies

$$\mathbb{E}_{f_\ell^1}[T_\delta^\ell] = \Big(\sum_{i \in \mathcal{R}} q_i T_i + \sum_{i \in \mathcal{R}} \sum_{j \in \mathcal{R}} q_i q_j d_{ij}\Big) \mathbb{E}_{f_\ell^1}[\delta^\ell].$$

Proof: Let the log likelihood ratio at region ℓ at iteration τ be λ_τ^ℓ. Note that

$$\lambda_\tau^\ell = \begin{cases} \log \frac{f_\ell^1(y_\tau)}{f_\ell^0(y_\tau)}, & \text{with probability } q_\ell, \\ 0, & \text{with probability } 1 - q_\ell. \end{cases}$$

Therefore, conditioned on the presence on an anomaly, $\{\lambda_\tau^\ell\}_{\tau \in \mathbb{N}}$ are i.i.d., and

$$\mathbb{E}_{f_\ell^1}[\lambda_\tau^\ell] = q_\ell \mathcal{D}(f_\ell^1, f_\ell^0).$$

Now, the remaining proof of the first statement is similar to the proof for CUSUM in [17].

To prove the second statement, we note that the information aggregation time at each iteration T^{agr} comprises of the processing time and the travel time. At an iteration the UAV is at region i with probability q_i and picks region j with probability q_j, and UAV travels between the two regions in d_{ij} units of time. Thus, the average travel time at each iteration is

$$\mathbb{E}[T_{\text{travel}}] = \sum_{i \in \mathcal{R}} \sum_{j \in \mathcal{R}} q_i q_j d_{ij}.$$

Hence, the expected information aggregation time at each iteration is

$$\mathbb{E}[T^{\text{agr}}] = \mathbb{E}[T_{\text{travel}} + T_{\text{process}}] = \sum_{i \in \mathcal{R}} \sum_{j \in \mathcal{R}} q_i q_j d_{ij} + \sum_{i \in \mathcal{R}} q_i T_i.$$

Let $\{T_\tau^{\text{agr}}\}_{\tau \in \{1, \dots, \delta^\ell\}}$, be the information aggregation times at each iteration. Thus, $T_\delta^\ell = \sum_{\tau=1}^{\delta^\ell} T_\tau^{\text{agr}}$, and it follows from Wald's identity [14] that

$$\mathbb{E}[T_\delta^\ell] = \mathbb{E}[T_\tau^{\text{agr}}] \mathbb{E}[\delta^\ell].$$

This completes the proof of the statement. ■

V. Optimal Vehicle routing

A. Stationary policy

We intend to determine a vehicle routing policy that minimizes the detection time of any anomaly present at any region. As exemplified in Theorem 2, this problem requires multiple detection delays to be minimized together.

We minimize the weighted sum of the expected detection delay at each region. We pick the weights to be the likelihood of the presence of an anomaly at each region. Formally, we pick the weight for region $\ell \in \mathcal{R}$ as $w_\ell = \pi_\ell^1/(\sum_{j\in\mathcal{R}} \pi_j^1)$.

Before we state the optimization problem to determine the optimal stationary policy, we introduce some notations. Let $\boldsymbol{v} \in \mathbb{R}^n_{>0}$ be the vector with entries $v_\ell = w_\ell(|e^{-\eta} + \eta - 1|)/\mathcal{D}(f_\ell^1, f_\ell^0)$, for each $\ell \in \mathcal{R}$. We denote the array of processing times by $\boldsymbol{T}$ and the matrix of distances between regions by D.

We define a single aggregate objective function $g : \Delta_{n-1} \to \mathbb{R}_{>0} \cup \{+\infty\}$ as the weighted sum of decision times at each region, i.e.,

$$g(\boldsymbol{q}) = \Big(\sum_{\ell\in\mathcal{R}} \frac{v_\ell}{q_\ell}\Big)\Big(\sum_{i\in\mathcal{R}} q_i T_i + \sum_{i\in\mathcal{R}}\sum_{j\in\mathcal{R}} q_i q_j d_{ij}\Big),$$

where $\boldsymbol{q}$ is the vector of the region selection probabilities.

We now pose the following detection delay minimization problem:

$$\underset{\boldsymbol{q}\in\Delta_{n-1}}{\text{minimize}} \quad g(\boldsymbol{q}) \tag{3}$$

Level sets of g on the two dimensional probability simplex are shown in Figure 2. It can be seen that for the example considered, the level sets are not convex, but there exists a unique minimum.

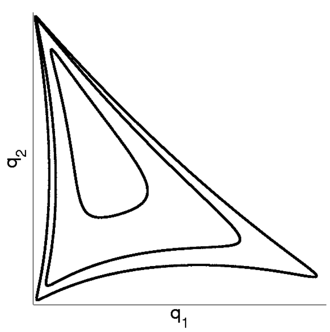

Fig. 2. Level-sets of the objective function in problem (3). It can be seen that the level sets are not convex.

Conjecture 3 (Uniqueness): The optimization problem (3) achieves a unique minimum.

We now provide some probabilistic guarantees for Conjecture 3. We assume that the parameter set $\{\boldsymbol{v}, \boldsymbol{T}, D, n\}$ in a given instance of optimization problem (3) is a realization of random variables sampled from some space Ω. For a given realization $\omega \in \Omega$, the associated optimization problem is:

$$\underset{\boldsymbol{q}\in\Delta_{n(\omega)-1}}{\text{minimize}} \quad g_\omega(\boldsymbol{q}), \tag{4}$$

where $g_\omega : \Delta_{n(\omega)-1} \to \mathbb{R}_{>0} \cup \{+\infty\}$ is defined by

$$g_\omega(\boldsymbol{q}) = \Big(\sum_{\ell\in\mathcal{R}(\omega)} \frac{v_\ell(\omega)}{q_\ell}\Big) \times \Big(\sum_{i\in\mathcal{R}(\omega)} q_i T_i(\omega) + \sum_{i\in\mathcal{R}(\omega)}\sum_{j\in\mathcal{R}(\omega)} q_i q_j d_{ij}(\omega)\Big),$$

where $\mathcal{R}(\omega) = \{1, \dots, n(\omega)\}$.

For such a realization ω, a local minimum of the optimization problem (4) can be found be substituting $q_{n(\omega)} = 1 - \sum_{j=1}^{n(\omega)-1} q_j$, and then running the gradient descent algorithm from some point $q_0 \in \Delta_{n(\omega)-1}$ on the resulting objective function. For a given ω, we sample N_1 points $q_0^r, r \in \{1, \dots, N_1\}$ in the simplex $\Delta_{n(\omega)-1}$ and, from each point, run the gradient descent algorithm to solve the optimization problem (4). Let $q_\omega^* : \Delta_{n(\omega)-1} \to \Delta_{n(\omega)-1}$ be the function that determines the outcome of the gradient descent algorithm, i.e., the gradient descent algorithm starting from point q_0^r converges to the point $q_\omega^*(q_0^r)$. Let $q_\omega^{\min} = q_\omega^*(\frac{1}{n(\omega)}\mathbf{1}_{n(\omega)})$. Define

$$\hat{\gamma}_\omega = \max\{\|q_\omega^*(q_0^r) - q_\omega^{\min}\| \mid r \in \{1, \dots, N_1\}\}.$$

It is known [3] that if $N_1 \geq \frac{1}{\epsilon_1}\log\frac{1}{\delta_1}$, then with at least probability $1 - \delta_1$

$$\mathbb{P}(\{q_0 \in \Delta_{n(\omega)-1} \mid \|q_\omega^*(q_0) - q_\omega^{\min}\| \leq \hat{\gamma}_\omega\}) \geq 1 - \epsilon_1.$$

Now consider realizations : $\{\omega_s \in \Omega \mid s \in \{1, \dots, N_2\}\}$. Define

$$\hat{\gamma} = \max\{\hat{\gamma}_{\omega_s} \mid s \in \{1, \dots, N_2\}\}.$$

If $N_2 \geq \frac{1}{\epsilon_2}\log\frac{1}{\delta_2}$, then at least with probability $1 - \delta_2$

$$\mathbb{P}(\{\omega \in \Omega \mid \hat{\gamma}_\omega \leq \hat{\gamma}\}) \geq 1 - \epsilon_2.$$

We sampled n uniformly from $\{3, \dots, 12\}$, n regions from the bivariate normal distribution with mean $\mathbf{0}_2$ and covariance $100\boldsymbol{I}_2$, T_ℓ, for each $\ell \in \mathcal{R}$, from the half normal distribution with mean 0 and variance 100, and v_ℓ, for each $\ell \in \mathcal{R}$, uniformly from $(0, 1)$. We normalized $\boldsymbol{v}$ to make it convex. The matrix D was chosen as the Euclidean distance matrix between the n sampled regions. We considered $N_2 = 92$ realizations of the $\{\boldsymbol{v}, \boldsymbol{T}, D, n\}$. For each realization, we solved the optimization problem (4) from 920 different initial points. The sample sizes were determined for $\epsilon_1 = 0.01$, $\delta_1 = 10^{-4}$, $\epsilon_2 = 0.05$, and $\delta_2 = 0.01$. The value of $\hat{\gamma}$ obtained was 10^{-4}.

Therefore, the gradient descent algorithm for the optimization problem (3) starting from any feasible point yields the same solution with high probability. In other words, the optimization problem (3) achieves a unique minimum with high probability.

B. Adaptive policy

The observations collected by the UAV can be utilized to learn the state of each region. This information is not utilized in the stationary policy considered in the previous section. We wish to incorporate our learning about the environment through each observation into our policies, and we do it in an adaptive fashion in the following way. After each observation, the CUSUM statistic for each region gives the likelihood of an anomaly in that region. We utilize this likelihood to design the weights of the cost functions in optimization problem (3). This ensures that the region with high likelihood of an anomaly is surveyed with a higher probability. We formally present these ideas in Algorithm 3.

Algorithm 3 Adaptive vehicle routing policy

1: **Given:** $\mathcal{R} = \{1, \dots, n\}$ and distances $d_{ij}, i, j \in \mathcal{R}$
2: **Given:** η, T_i, π_i^1, f_i^0, f_i^1, for each $i \in \mathcal{R}$
3: set $\Lambda_0^j = 0$, for all $j \in \mathcal{R}$, and $\tau = 0$
4: set $w_i = e^{\Lambda_\tau^i} / (\sum_{j \in \mathcal{R}} e^{\Lambda_\tau^j})$, for each $i \in \mathcal{R}$
5: obtain solution $\boldsymbol{q}^*$ of optimization problem (3)
6: at time $\tau \in \mathbb{N}$, select a random region $\ell \in \mathcal{R}$ according to the probability distribution $\boldsymbol{q}^*$
7: collect sample y_τ from region ℓ
8: update the CUSUM statistic at each region

$$\Lambda_\tau^j = \begin{cases} \left(\Lambda_{\tau-1}^\ell + \log \frac{f_\ell^1(y_\tau)}{f_\ell^0(y_\tau)}\right)^+, & \text{if } j = \ell, \\ \Lambda_{\tau-1}^j, & \text{if } j \in \mathcal{R} \setminus \{\ell\}. \end{cases}$$

% detect change if the threshold is crossed
9: **if** $\Lambda_\tau^\ell > \eta$, **then** declare anomaly detected at region ℓ and set $\Lambda_\tau^\ell = 0$
10: continue to step 4:

VI. Numerical Results

We now elucidate on the ideas presented in the previous sections with some examples. In the first example, we compare the analytic expressions for the expected detection delay obtained in Theorem 2 with the empirical expected detection delay. In the second example, we elucidate on the adaptive vehicle routing policy. We consider two scenarios, first when no anomaly is present, and second, when multiple anomalies appear at different regions. We also present a comparison of the adaptive policy with the optimal stationary policy.

Example 4: Consider the surveillance of four regions $\mathcal{R} = \{1, \dots, 4\}$ shown in Figure 1 . If no anomaly is present, then the observation is a normal $\mathcal{N}(0, 1)$ random variable. Otherwise, the observation is a normal $\mathcal{N}(1, 1)$ random variable. The distance between the regions is given by the matrix

$$D = \begin{bmatrix} 0 & 25 & 9 & 17 \\ 25 & 0 & 18 & 28 \\ 9 & 18 & 0 & 14 \\ 17 & 28 & 14 & 0 \end{bmatrix} \text{ units.}$$

Assume that the UAV moves with a speed 5 units per second and the processing time at each region is 1 second. A comparison of the analytic and empirical expected detection delay is shown in Figure 3. It can be seen that the analytic expected detection delay provides a lower bound to the empirical expected detection delay and the gap between them is small. It can also be seen that the expected detection delay for the policy where one stays at the anomalous region all the time is minimum because it requires minimum expected sample size and zero travel time.

Example 5: For the same set of data as in Example 4, we now study the adaptive vehicle routing policy. First, we present the scenario when no anomaly is present. A sample evolution of the CUSUM statistics and the corresponding region selection policy is shown in Figure 4. It can be seen that if the threshold is small, there may be potential false alarms. The region selection probabilities grow with the likelihood of the anomaly. When the CUSUM statistic is zero for each region, then the region 3 is surveyed with highest probability as it is closest to all other regions.

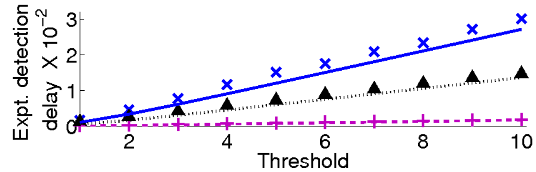

Fig. 3. Expected detection delay for randomized ensemble CUSUM in seconds. The anomaly appears at region 2. The solid blue line and blue $\times$, respectively, represent the analytic and empirical expected detection delay for the uniform region selection policy. The dotted black line and black triangles, respectively, represent the analytic and empirical expected detection delay for the region selection probabilities $\{0, 0.5, 0, 0.5\}$. The dashed magenta line and magenta $+$, respectively, represent the analytic and empirical expected detection delay for the policy when region 2 is surveyed all the time.

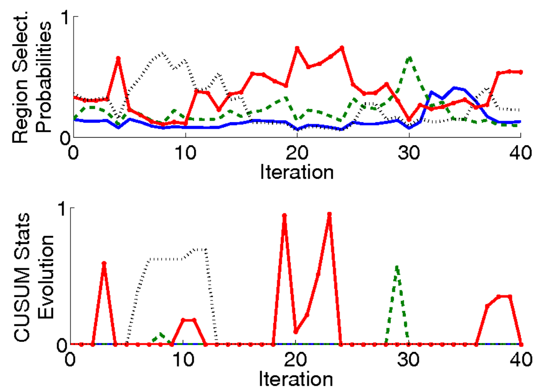

Fig. 4. Sample evolution of region selection probabilities and the CUSUM statistic at each region when no anomaly is present. Regions $\{1, 2, 3, 4\}$ are shown in solid blue, dashed green, solid red with dots, and dotted black lines, respectively. The region selection probability is a function of distance of the region from other regions and the likelihood of anomaly being present. As the likelihood of an anomaly being present at a region increases the probability to survey that region increases. When all CUSUM statistics are zero, the region 3 is surveyed with highest probability because it is closest to all other regions.

Second, we present the scenario when anomalies appear at region 2, 4, and 3 at iteration number 10, 20, 30, respectively. A sample evolution of the CUSUM statistics and the corresponding region selection policy is shown in Figure 5. Again, it can be seen that the region selection probability increases with the likelihood of the anomaly in that region, and if the likelihood of an anomaly is very high at a region, the probability to survey that region is very close to one.

We present a comparison of the optimal stationary region selection policy with the adaptive policy in Figure 6. It can be noticed that the optimal stationary policy performs better than the adaptive policy at low thresholds, but this regime is marred with high false-alarm rate. For higher thresholds the adaptive policy performs better than the optimal stationary policy. Notice that the improvement in the detection delay of region 3 is not significant, since region 3 is already surveyed

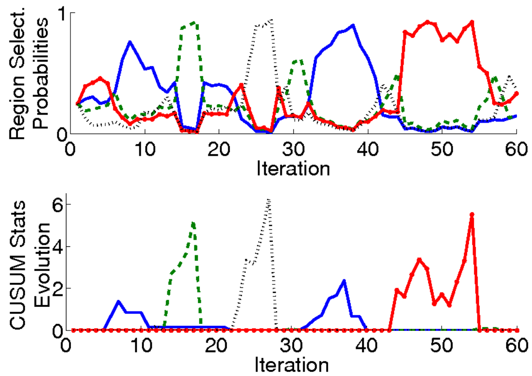

Fig. 5. Sample evolution of the region selection probabilities and CUSUM statistic at each region. Regions $\{1, 2, 3, 4\}$ are shown in solid blue, dashed green, solid red with dots, and dotted black lines, respectively. Anomalies appear at region 2, 3 and 4 at iteration $10, 60$ and 100, respectively. The region selection probability is a function of distance of the region from other regions and the likelihood of anomaly being present. Once an anomaly is detected, it is removed and the statistic is reset to zero.

with high probability under optimal stationary policy, but the improvement in the performance at region 2 and 4 is significant.

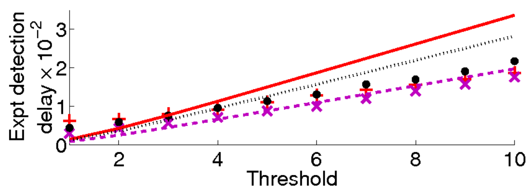

Fig. 6. Comparison of the adaptive policy with the optimal stationary surveillance policy. The anomalies appear at times 10 secs, 20 secs, and 30 secs at region 2, 4, and 3, respectively. The red $+$, the magenta $\times$ and the black dots represent the empirical expected detection delay for the adaptive horizon policy at region 2, 3 and 4, respectively. The solid red line, the magenta dashed line and the black dotted line represent the analytic expected detection delay for the optimal stationary policy at region 2, 3 and 4, respectively.

VII. Conclusions and Future Directions

We studied surveillance of multiple regions with a UAV and presented stochastic surveillance strategies that detect multiple anomalies in a set of regions in minimum time. We adopted randomized region selection strategy for the UAV. A randomized strategy ensures that an intruder does not preempt its detection during a future visit by the UAV. We studied the information collection and travel time trade-off. An adaptive policy was presented and it was numerically demonstrated that the adaptive policy performs better than the optimal stationary policy.

An immediate extension of this work is to consider surveillance with multiple UAVs. A region partitioning approach can be utilized to deal with multiple UAVs. Other potential extensions include learning the regions and studying the exploration-exploitation trade-off in stochastic surveillance.

References

[1] 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.

[2] F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems. *Proceedings of the IEEE*, May 2011. To appear.

[3] G. C. Calafiore, F. Dabbene, and R. Tempo. Research on probabilistic methods for control system design. *Automatica*, 47(7):1279–1293, 2011.

[4] D. A. Castañón. Optimal search strategies in dynamic hypothesis testing. *IEEE Transactions on Systems, Man & Cybernetics*, 25(7):1130–1138, 1995.

[5] T. H. Chung and J. W. Burdick. A decision-making framework for control strategies in probabilistic search. In *IEEE Int. Conf. on Robotics and Automation*, pages 4386–4393, Roma, Italy, April 2007.

[6] J. Grace and J. Baillieul. Stochastic strategies for autonomous robotic surveillance. In *IEEE Conf. on Decision and Control and European Control Conference*, pages 2200–2205, Seville, Spain, December 2005.

[7] 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.

[8] J. P. Hespanha, H. J. Kim, and S. S. Sastry. Multiple-agent probabilistic pursuit-evasion games. In *IEEE Conf. on Decision and Control*, pages 2432–2437, Phoenix, AZ, USA, December 1999.

[9] D. J. Klein, J. Schweikl, J. T. Isaacs, and J. P. Hespanha. On UAV routing protocols for sparse sensor data exfiltration. In *American Control Conference*, pages 6494–6500, Baltimore, MD, USA, June 2010.

[10] F. Pasqualetti, A. Franchi, and F. Bullo. On cooperative patrolling: Optimal trajectories, complexity analysis and approximation algorithms. *IEEE Transactions on Robotics*, January 2011. Submitted.

[11] H. V. Poor and O. Hadjiliadis. *Quickest Detection*. Cambridge University Press, 2008.

[12] S. A. P. Quintero, F. Papi, D. J. Klein, L. Chisci, and J. P. Hespanha. Optimal UAV Coordination for Target Tracking using Dynamic Programming. In *IEEE Conf. on Decision and Control*, pages 4541–4546, Atlanta, GA, USA, December 2010.

[13] R. Ratnam, J. Goense, and M. E. Nelson. Change-point detection in neuronal spike train activity. *Neurocomputing*, 52:849–855, 2003.

[14] S. I. Resnick. *A Probability Path*. Birkhäuser, 1999.

[15] T. Sak, J. Wainer, and S. Goldenstein. Probabilistic multiagent patrolling. In *Brazilian Symposium on Artificial Intelligence*, Advances in Artificial Intelligence, pages 124–133, Salvador, Brazil, 2008. Springer.

[16] K. Savla, T. Temple, and E. Frazzoli. Human-in-the-loop vehicle routing policies for dynamic environments. In *IEEE Conf. on Decision and Control*, pages 1145–1150, Cancún, México, December 2008.

[17] D. Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. Springer, 1985.

[18] S. L. Smith and D. Rus. Multi-robot monitoring in dynamic environments with guaranteed currency of observations. In *IEEE Conf. on Decision and Control*, pages 514–521, Atlanta, GA, USA, December 2010.

[19] S. L. Smith, M. Schwager, and D. Rus. Persistent robotic tasks: Monitoring and sweeping in changing environments. *IEEE Transactions on Robotics*, February 2011. To appear.

[20] K. Srivastava, D. M. Stipanovic, and M. W. Spong. On a stochastic robotic surveillance problem. In *IEEE Conf. on Decision and Control*, pages 8567–8574, Shanghai, China, December 2009.

[21] V. Srivastava, K. Plarre, and F. Bullo. Adaptive sensor selection in sequential hypothesis testing. In *IEEE Conf. on Decision and Control and European Control Conference*, Orlando, FL, USA, December 2011. to appear.

[22] V. Srivastava, K. Plarre, and F. Bullo. Randomized sensor selection in sequential hypothesis testing. *IEEE Transactions on Signal Processing*, 59(5):2342–2354, 2011.

[23] D. Zhang, C. Colburn, and T. Bewley. Estimation and adaptive observation of environmental plumes. In *American Control Conference*, pages 4821–4286, San Francisco, CA, USA, June 2011.