Stochastic Surveillance Strategies for Spatial Quickest Detection

Vaibhav Srivastava Fabio Pasqualetti Francesco Bullo

Abstract—**We design persistent surveillance strategies for the quickest detection of anomalies taking place in an environment of interest. From a set of predefined regions in the environment, a team of autonomous vehicles collects noisy observations, which a control center processes. The overall objective is to minimize detection delay while maintaining the false alarm rate below a desired threshold. We present joint (i) anomaly detection algorithms for the control center and (ii) vehicle routing policies. For the control center, we propose parallel cumulative sum (CUSUM) algorithms (one for each region) to detect anomalies from noisy observations. For the vehicles, we propose a stochastic routing policy, in which the regions to be visited are chosen according to a probability vector. We study stationary routing policy (the probability vector is constant) as well as adaptive routing policies (the probability vector varies in time as a function of the likelihood of regional anomalies). In the context of stationary policies, we design a performance metric and minimize it to design an efficient stationary routing policy. Our adaptive policy improves upon the stationary counterpart by adaptively increasing the selection probability of regions with high likelihood of anomaly. Finally, we show the effectiveness of the proposed algorithms through numerical simulations and a persistent surveillance experiment.**

Index Terms—**vehicle routing, statistical decision making, quickest detection, persistent surveillance, patrolling, security, motion planning.**

I. Introduction

Recent years have witnessed a surge in the application of autonomous agents in various activities such as surveillance and information collection. In view of the recent Icelandic ash problem, the oil spill in the gulf of Mexico, and recurring wild fires, surveillance strategies resulting in the quickest detection of anomalies are of considerable importance. Due to extreme sensor and modeling uncertainties in these situations, robust anomaly detection methods need to be employed. Generally, a limited number of vehicles are deployed to survey a large number of regions, and it is fundamental that the vehicles collect the information that is most effective to minimize the detection delay of anomalies. In this paper we design surveillance strategies that result in quick detection of anomalies.

A preliminary version of this work (Srivastava and Bullo, 2011) was presented at IEEE Conference on Decision and Control and European Control Conference, 2011. In addition to the ideas in (Srivastava and Bullo, 2011), this paper contains a rigorous analysis of the single vehicle surveillance, the multiple vehicle surveillance, extensive numerical illustrations, and a persistent surveillance experiment.

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

A reliable detection of anomalies can be achieved by collecting observations sequentially until the evidence suggesting an anomaly reaches a substantial level. Various sequential algorithms for the detection of anomalies have been presented in (Basseville and Nikiforov, 1993). Furthermore, it is known that a human being typically performs well in detecting and identifying anomalies from observations. Recent advances in cognitive psychology (Bogacz et al., 2006; Ratnam et al., 2003), show that human performance in decision making tasks is well modeled by sequential statistical procedures such as the CUSUM algorithm. Inspired by the above human decision making models, in this work we adopt sequential statistical tests for anomaly detection.

Our setup and approach. We consider an environment comprising of potentially disjoint regions of interest, and we employ a team of autonomous vehicles for the persistent surveillance of these regions. In particular, the vehicles visit the regions, collect information, and send it to a control center. We study a spatial quickest detection problem with multiple vehicles, that is, the simultaneous quickest detection of anomalies at spatially distributed regions when the observations for anomaly detection are collected by autonomous vehicles. For this problem, we let the control center run parallel CUSUM algorithms (one for each region) with the collected information. The control center then decides on the presence of anomalies in the regions. Finally, we design vehicle routing policies to collect observations at different regions. Our vehicle routing policies aim to minimize the anomaly detection time at the control center.

Related work. Vehicle routing policies have witnessed a lot of attention in the robotics and controls literature. A survey on dynamic vehicle routing policies for servicing tasks is presented in (Bullo et al., 2011). Recently, the routing for information aggregation has been of particular interest. Klein et al. (2010) present a vehicle routing policy for optimal localization of an acoustic source. They consider a set of spatially distributed sensors and optimize the trade-off between the travel time required to collect a sensor observation and the information contained in the observation. They characterize the information in an observation by the volume of the Cramer-Rao ellipsoid associated with an optimal estimator. Hollinger et al. (2011b) study routing for an AUV to collect data from an underwater sensor network. They developed approximation algorithms for variants of the traveling salesperson problem to determine efficient policies that maximize the information collected while minimizing the travel time. Gupta et al. (2006) study the estimation in a linear dynamical system with the observations collected by a set of mobile sensors. They determine stochastic trajectories for mobile sensors that minimize the error covariance of the Kalman filter estimate. Zhang et al. (2011) study the estimation of environmental plumes with mobile sensors. They minimize the uncertainty of the estimate of the ensemble Kalman filter to determine optimal trajectories for a swarm of mobile sensors.

There has been some interest in decision theoretic information aggregation and vehicle routing as well. Castañón (1995) 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 and Burdick (2012) study the probabilistic search problem in a decision theoretic framework. They minimize the search decision time in a Bayesian setting. Certain optimal information aggregation strategies for sequential hypothesis testing have been developed in (Srivastava et al., 2011a,b). Hollinger et al. (2011a) study an active classification problem in which an autonomous vehicle classifies an object based on multiple views. They formulate the problem in an active Bayesian learning framework and apply it to underwater detection. The persistent surveillance problem in this paper also concerns with decision-theoretic information aggregation and vehicle routing. In contrast to the aforementioned works that focus on classification or search problems, our focus is on quickest detection of anomalies.

The problem of surveillance has received considerable attention recently. Preliminary results on this topic have been presented in (Chevaleyre, 2004; Elmaliach et al., 2008; Kingston et al., 2008). Pasqualetti et al. (2012b) 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 and Rus (2010) 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 vehicles move on a given closed path has been considered in (Pasqualetti et al., 2012a; Smith et al., 2012), and a speed controller has been designed to minimize the time lag between visits of regions.

Stochastic surveillance and pursuit-evasion problems have also fetched significant attention. In an earlier work, Hespanha et al. (1999) 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. Grace and Baillieul (2005) formulate the surveillance problem as a random walk on a hypergraph and parametrically vary the local transition probabilities over time in order to achieve an accelerated convergence to a desired steady state distribution. Sak et al. (2008) present partitioning and routing strategies for surveillance of regions for different intruder models. Srivastava et al. (2009) present a stochastic surveillance problem in centralized and decentralized frameworks. They use a Markov chain Monte Carlo method and a message passing based auction algorithm to achieve the desired surveillance criterion. They also show that the deterministic strategies fail to satisfy the surveillance criterion under general conditions. In this paper, we focus on stochastic surveillance policies. In contrast to aforementioned works on stochastic surveillance that assume a surveillance criterion is known, this work concerns the design of the surveillance criterion. The policies designed in this paper direct a vehicle with high probability to a region with high probability of being anomalous, a feature akin to the heuristic policy for the pursuer in (Hespanha et al., 1999). On the other hand, with respect to (Hespanha et al., 1999), our policy takes into account environmental factors, e.g., travel times and detection difficulty, and it satisfies an optimality criterion.

Paper contributions. The main contributions of this work are fivefold. First, we formulate the stochastic surveillance problem for spatial quickest detection of anomalies (Section II). We propose the ensemble CUSUM algorithm for a control center to detect concurrent anomalies at different regions from collected observations (Section III). For the ensemble CUSUM algorithm we characterize lower bounds for the expected detection delay and for the average (expected) detection delay at each region. Our bounds take into account the processing times for collecting observations, the prior probability of anomalies at each region, and the anomaly detection difficulty at each region.

Second, for the case of stationary routing policies, we provide bounds on the expected delay in detection of anomalies at each region (Section IV). In particular, we take into account both the processing times for collecting observations and the travel times between regions. For the single vehicle case, we explicitly characterize the expected number of observations necessary to detect an anomaly at a region, and the corresponding expected detection delay. For the multiple vehicles case, we characterize lower bounds for the expected detection delay and the average detection delay at the regions. As a complementary result, we show that the expected detection delay for a single vehicle is, in general, a non-convex function. However, we provide probabilistic guarantees that it admits a unique global minimum.

Third, we design stationary vehicle routing policies to collect observations from different regions (Section IV). For the single vehicle case, we design an efficient stationary policy by minimizing an upper bound for the average detection delay at the regions. For the multiple vehicles case, we first partition the regions among the vehicles, and then we let each vehicle survey the assigned regions by using the routing policy as in the single vehicle case. In both cases we characterize the performance of our policies in terms of expected detection delay and average (expected) detection delay.

Fourth, we describe our adaptive ensemble CUSUM algorithm, in which the routing policy is adapted according to the learned likelihood of anomalies in the regions (Section V). We derive an analytic bound for the performance of our adaptive policy. Finally, our numerical results show that our adaptive policy outperforms the stationary counterpart.

Fifth and finally, we report the results of extensive numerical simulations and a persistent surveillance experiment (Sections VI and VII). Besides confirming our theoretical findings, these practical results show that our algorithm are robust against realistic noise models, and sensors and motion uncertainties.

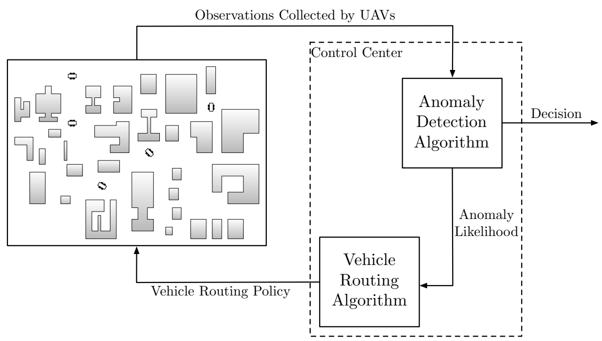

Fig. 1. Persistent Surveillance Setup. A set of n regions is surveyed by $m < n$ vehicles. Each vehicle visits the regions according to some policy and collects evidence from the visited region. The collected evidence is sent to an anomaly detection algorithm. The anomaly detection algorithm processes the collected evidence and decides on the presence of an anomaly. It also provides the likelihood of an anomaly being present, which in turn is used by the vehicle routing algorithm. The anomaly detection algorithm and vehicle routing algorithm constitute the control center, which can be implemented on-board of a vehicle.

II. Problem Setup

We consider the *persistent surveillance* of a set of n disjoint regions with a team of $m < n$ identical[1] autonomous vehicles capable of sensing, communicating, and moving from one region to another. In persistent surveillance, the vehicles visit the regions according to some routing policy, collect evidence (sensor observation), and send it to a control center. The control center runs an anomaly detection algorithm with the evidence collected by the vehicles to determine the likelihood of an anomaly being present at some region (the control center declares an anomaly if substantial evidence is present). Finally, the control center utilizes the likelihood of an anomaly at each region to determine a vehicle routing policy. The objective of the control center is to detect an anomaly at any region in minimum time subject to a desired bound on the expected time between any two subsequent false alarms. Notice that the time required to detect an anomaly depends on the anomaly detection algorithm and the time vehicles take to travel the regions. Thus, the control center needs to minimize the anomaly detection time jointly over anomaly detection policies and vehicle routing policies. Our problem setup is shown in Fig. 1.

We adopt the standard motion planning notation in (LaValle, 2006). We denote the k-th region by $\mathcal{R}_k, k \in \{1, \dots, n\}$, and r-th vehicle by $\mathcal{U}_r, r \in \{1, \dots, m\}$. Let the likelihood of an anomaly at region $\mathcal{R}_k$ be $\pi_k \in (0, 1)$. We study the persistent surveillance problem under the following assumptions.

Regarding the vehicles, we do not assume any specific dynamics and we assume that:

(A1) each vehicle takes time d_{ij} to travel from region $\mathcal{R}_i$ to region $\mathcal{R}_j$, $i, j \in \{1, \dots, n\}$;

(A2) the sensors on each vehicle take a random time T_k to collect an informative observation[2] from region $\mathcal{R}_k, k \in \{1, \dots, n\}$.

Regarding the observations, we assume that:

(A3) the observation collected by a vehicle from region $\mathcal{R}_k$ is sampled from probability density functions $f_k^0 : \mathbb{R} \to \mathbb{R}_{\geq 0}$ and $f_k^1 : \mathbb{R} \to \mathbb{R}_{\geq 0}$, respectively, in the presence and in the absence of anomalies;

(A4) for each $k \in \{1, \dots, n\}$, probability density functions f_k^1 and f_k^0 are non-identical with some non-zero probability, and the two distributions have the same support;

(A5) conditioned on the presence or absence of anomalies, the observations in each region are mutually independent; and

(A6) observations in different regions are also mutually independent.

Regarding the anomaly detection algorithm at the control center, we employ the cumulative sum (CUSUM) algorithm (see below) for anomaly detection at each region. In particular, we run n parallel CUSUM algorithms (one for each region) and declare an anomaly being present at a region as soon as a substantial evidence is present. We refer to such parallel CUSUM algorithms by *ensemble CUSUM algorithm*.

*Remark 1 (**Knowledge of distributions**):* For the ease of presentation, we assume that the probability density functions in presence and absence of an anomaly are known. In general, only the probability density function in absence of any anomaly may be known, or both the probability density functions may be unknown. In the first case, the CUSUM algorithm can be replaced by the weighted CUSUM algorithm or the Generalized Likelihood Ratio (GLR) algorithm (Basseville and Nikiforov, 1993), while in the second case, it can be replaced by the robust minimax quickest change detection algorithm (Unnikrishnan et al., 2011). The ideas presented in this paper extend to these cases in a straightforward way. A related example is in Section VI. $\square$

*Remark 2 (**Independence of observations**):* For the ease of presentation, we assume that the observations collected from each region are independent conditioned on the presence and absence of anomalies. In general, the observations may be dependent and the dependence can be captured through an appropriate hidden Markov model. If the observations can be modeled as a hidden Markov model, then the CUSUM like algorithm in (Chen and Willett, 2000) can be used instead of the standard CUSUM algorithm. The analysis presented in this paper holds in this case as well but in an asymptotic sense, i.e., in the limit when a large number of observations are needed for anomaly detection.

We also assumed that the observations collected from different regions are mutually independent. Although the ideas in this paper also work when the observations at different regions are dependent, the performance can be improved with a slight modification in the procedure presented here (see Remark 4). In this case the algorithm performance improves because each observation is now informative about more than one region. $\square$

Regarding the vehicle routing policy, we propose the *randomized routing policy*, and the *adaptive routing policy*. In the randomized routing policy, each vehicle (i) selects a region from a stationary distribution, (ii) visits that region,

[1] The vehicle routing policies designed in this paper also work for non-identical vehicles. We make this assumption for the convenience of analysis.

[2] An informative observation may require the acquisition of several observations from different locations at the same region. In this case the processing time equals the total time required to collect all these observations.

(iii) collects an evidence, and (iv) transmits this evidence to the control center and iterates this process endlessly. In the randomized routing policy, the evidence collected by the vehicles is not utilized to modify their routing policy. In other words, there is no feedback from the anomaly detection algorithm to the vehicle routing algorithm. In the adaptive routing policy, instead, the evidence collected by the vehicles is used to modify the routing policy, and thus, the loop between the vehicle routing algorithm and the anomaly detection algorithm is closed. The adaptive routing policy follows the same steps as in the randomized routing policy, with the exception that the distribution in step (i) is no longer stationary and is adapted based on the collected evidence.

For brevity of notation, we will refer to the joint anomaly detection and vehicle routing policy comprising of the ensemble CUSUM algorithm and the randomized routing policy by *randomized ensemble CUSUM algorithm.* We will show that the randomized ensemble CUSUM algorithm provides a solution that is within a factor of optimality. Similarly, we refer to the joint anomaly detection and vehicle routing policy comprising of the ensemble CUSUM algorithm and adaptive routing policy by *adaptive ensemble CUSUM algorithm.* We will show that adaptive ensemble CUSUM algorithm makes the vehicles visit anomalous regions with high probability, and thus it improves upon the performance of the randomized ensemble CUSUM algorithm. The following standard definition (Cover and Thomas, 1991) will be used in the remaining sections.

*Definition 1 (**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

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

where $\mathcal{L}^1$ is the set of integrable functions, $\mathbb{E}_{f_1}[\cdot]$ represents expected value with respect to f_1, X is a random variable sampled from f_1, and $\operatorname{supp}(f_1)$ is the support of f_1. □

It is known that (i) $0 \leq \mathcal{D}(f_1, f_2) \leq +\infty$, (ii) the lower bound is achieved if and only if $f_1 = f_2$ almost everywhere, and (iii) the upper bound is achieved if and only if the support of f_2 is a strict subset of the support of f_1. Observe that Assumption (A4) on the observations is equivalent to $\mathcal{D}(f_k^1, f_k^0) > 0$, for each $k \in \{1, \dots, n\}$.

We now introduce some notations that will be used throughout the paper. We denote the probability simplex in $\mathbb{R}^n$ by Δ_{n-1}, and the space of vehicle routing policies by Ω. For the processing time T_k, we let $\bar{T}_k$ denote its expected value. Consider m realizations of the processing time T_k, we denote the expected value of the minimum of these m realized values by $\bar{T}_k^{m\text{-smlst}}$. Note that $\bar{T}_k^{1\text{-smlst}} = \bar{T}_k$. We also define $\bar{T}_{\max} = \max\{\bar{T}_k \mid k \in \{1, \dots, n\}\}$ and $\bar{T}_{\min} = \min\{\bar{T}_k \mid k \in \{1, \dots, n\}\}$. We denote the Kullback-Leibler divergence between the probability density functions f_k^1 and f_k^0 by $\mathcal{D}_k$. Finally, $\mathcal{D}_{\max} = \max\{\mathcal{D}_k \mid k \in \{1, \dots, n\}\}$ and $\mathcal{D}_{\min} = \min\{\mathcal{D}_k \mid k \in \{1, \dots, n\}\}$. For the convenience of the reader, we have enlisted the notation in Table I.

TABLE I
LIST OF SYMBOLS

n	number of regions
m	number of robots
$\mathcal{R}_k, k \in \{1, \dots, n\}$	k-th region
$\mathcal{U}_r, r \in \{1, \dots, m\}$	r-th vehicle
$\pi_k, k \in \{1, \dots, n\}$	prior probability of anomaly at $\mathcal{R}_k$
w_k	$\pi_k / (\sum_{j=1}^n \pi_j)$
$d_{ij}, i, j \in \{1, \dots, n\}$	travel time between $\mathcal{R}_i$ and $\mathcal{R}_j$
$T_k, k \in \{1, \dots, n\}$	processing time at $\mathcal{R}_k$
$\bar{T}_k, k \in \{1, \dots, n\}$	expected processing time at $\mathcal{R}_k$
$\bar{T}_{\max}$	$\max\{\bar{T}_k \mid k \in \{1, \dots, n\}\}$
$\bar{T}_{\min}$	$\min\{\bar{T}_k \mid k \in \{1, \dots, n\}\}$
$\bar{T}_k^{m\text{-smlst}}$	$\mathbb{E}[\min\{T_k^{(1)}, \dots, T_k^{(m)}\}]$, where $T_k^{(1)}, \dots, T_k^{(m)}$ are m realizations of T_k
$\bar{T}_{\min}^{m\text{-smlst}}$	$\min\{\bar{T}_k^{m\text{-smlst}} \mid k \in \{1, \dots, n\}\}$
Ξ	set of sets of m arbitrary regions
$\bar{T}_{\text{one}}$	$\min\{\mathbb{E}[\min\{t_1^\xi, \dots, t_m^\xi\}] \mid \xi \in \Xi\}$, where t_i^ξ's are the processing times at regions in ξ
$f_k^0, k \in \{1, \dots, n\}$	pdf in absence of anomaly at $\mathcal{R}_k$
$f_k^1, k \in \{1, \dots, n\}$	pdf in presence of anomaly at $\mathcal{R}_k$
$\mathcal{D}_k, k \in \{1, \dots, n\}$	K-L divergence between f_k^1 and f_k^0
$\mathcal{D}_{\max}$	$\max\{\mathcal{D}_k \mid k \in \{1, \dots, n\}\}$
$\mathcal{D}_{\min}$	$\min\{\mathcal{D}_k \mid k \in \{1, \dots, n\}\}$
Ω	space of vehicle routing policies
$N_k : \Omega \to \mathbb{N} \cup \{+\infty\}$	observations required for detection at $\mathcal{R}_k$
$\delta_k : \Omega \to \mathbb{R}_{>0} \cup \{+\infty\}$	detection delay at $\mathcal{R}_k$
$\delta_{\text{avg}} : \Omega \to \mathbb{R}_{>0} \cup \{+\infty\}$	$\sum_{k=1}^n w_k \mathbb{E}[\delta_k(\omega)]$
$\delta_{\text{upper}} : \Omega \to \mathbb{R}_{>0} \cup \{+\infty\}$	upper bound to δ_{avg}
$\delta_k^{m\text{-min}}$	$\inf\{\mathbb{E}[\delta_k(\omega)] \mid \omega \in \Omega\}$
$\delta_{\text{avg}}^{m\text{-min}}$	$\inf\{\delta_{\text{avg}}(\omega) \mid \omega \in \Omega\}$
Λ_τ^j	CUSUM statistic at $\mathcal{R}_j$ at τ-th iteration
η	CUSUM threshold
$\bar{\eta}$	$e^{-\eta} + \eta - 1$
Δ_{n-1}	probability simplex in $\mathbb{R}^n$
$\boldsymbol{q} \in \Delta_{n-1}$	single vehicle randomized routing policy
$\boldsymbol{q}^* \in \Delta_{n-1}$	optimal $\boldsymbol{q}$
$\boldsymbol{q}^\dagger \in \Delta_{n-1}$	efficient $\boldsymbol{q}$
$\vec{\boldsymbol{q}}_m \in \Delta_{n-1}^m$	m vehicle randomized routing policy
$\vec{\boldsymbol{q}}_{\text{part}} \in \Delta_{n-1}^m$	$\vec{\boldsymbol{q}}_m$ with region partitioning
$\boldsymbol{a} \in \Delta_{n-1}$	single vehicle adaptive routing policy
$\boldsymbol{a}_{\text{part}} \in \Delta_{n-1}^m$	m vehicle adaptive routing policy with region partitioning

*Remark 3 (**Randomized routing policy**):* The randomized routing policy samples regions to visit from a stationary distribution; this assumes that each region can be visited from another region in a single hop. While this is the case for aerial vehicles, it may not be true for ground vehicles. In the latter case, the motion from one region to another can be modeled as a Markov chain. The transition probabilities of this Markov chain can be designed to achieve a desired stationary distribution. This can optimally be done, for instance, by picking the fastest mixing Markov chain proposed in (Boyd et al., 2004) or heuristically by using the standard Metropolis-Hastings algorithm (Wasserman, 2004). Related examples are presented in Section VI and VII. It should be noted that

under the randomized routing policy, the desired stationary distribution of the Markov chain is fixed, and the Markov chain converges to this distribution exponentially. Thus, the policy designed using Markov chain is arbitrarily close to the desired policy. However, in the case of adaptive routing policy, the desired stationary distribution keeps on changing, and the performance of the Markov chain based policy depends on rate of convergence of the Markov chain and the rate of change of desired stationary distribution. □

III. Spatial Quickest Detection

In this section we propose the ensemble CUSUM algorithm for the simultaneous quickest detection of anomalies in spatially distributed regions. We start by recalling the standard quickest change detection problem. Then we describe and characterize the ensemble CUSUM algorithm.

A. Quickest change detection

Consider a set of observations $\{y_1, y_2, \dots\}$, where, for some ν, the observations $\{y_1, \dots, y_{\nu-1}\}$ are i.i.d. with probability density function f^0, and $\{y_\nu, y_{\nu+1}, \dots\}$ are i.i.d. with probability density function f^1. The objective of the quickest change detection is to detect the change in the underlying distribution in minimum number of observations subject to a desired lower bound on the number of samples between two false alarms. Let $N \geq \nu$ be the observation at which the change is detected. The non-Bayesian quickest detection problem (Poor and Hadjiliadis, 2008; Siegmund, 1985), is posed as

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

where $\mathbb{E}_\nu[\cdot]$ represents expected value with respect to the observations distribution at iteration ν and $\gamma > 0$ is a small constant called *false alarm rate.*

An algorithmic solution to the minimization problem (1) is the cumulative sum (CUSUM) algorithm (Poor and Hadjiliadis, 2008), in which, at each iteration $\tau \in \mathbb{N}$, (i) an observation y_τ is collected, (ii) the statistic $\Lambda_\tau = \big(\Lambda_{\tau-1} + \log \frac{f_k^1(y_\tau)}{f_k^0(y_\tau)}\big)^+$ with $\Lambda_0 = 0$ is computed, and (iii) a change is declared if $\Lambda_\tau > \eta$. For a given threshold η, the false alarm rate and the worst expected number of observations for CUSUM algorithm are

$$\mathbb{E}_{f^0}(N) \approx \frac{e^{\eta} - \eta - 1}{\mathcal{D}(f^0, f^1)} \text{ and } \mathbb{E}_{f^1}(N) \approx \frac{e^{-\eta} + \eta - 1}{\mathcal{D}(f^1, f^0)}. \tag{2}$$

The approximations in equation (2) are referred to as the Wald's approximations (Siegmund, 1985), and are known to be accurate for large values of the threshold η. In the following, we assume that the chosen threshold is large enough and the expressions in equation (2) are exact. Let $u > 0$ be the uniform time duration between two iterations of the CUSUM algorithm. The *expected detection delay* δ, i.e., the expected time required to detect an anomaly after its appearance, satisfies $\mathbb{E}_{f^1}[\delta] = u\mathbb{E}_{f^1}(N)$.

Algorithm 1: *Ensemble CUSUM Algorithm*

Input : threshold η, pdfs $f_k^0, f_k^1, k \in \{1, \dots, n\}$;
Output : decision on presence of an anomaly ;

at time τ receive observation y_τ for region $\mathcal{R}_k$;
update the CUSUM statistic at each region:

$$\Lambda_\tau^j = \begin{cases} \Big(\Lambda_{\tau-1}^k + \log \frac{f_k^1(y_\tau)}{f_k^0(y_\tau)}\Big)^+, & \text{if } j = k; \\ \Lambda_{\tau-1}^j, & \text{if } j \in \{1, \dots, n\} \setminus \{k\}; \end{cases}$$

if $\Lambda_\tau^k > \eta$ **then** change detected at region $\mathcal{R}_k$;
else wait for next observations and iterate.

B. Ensemble CUSUM algorithm

We run n parallel CUSUM algorithms (one for each region), and update the CUSUM statistic for region $\mathcal{R}_k$ only if an observation is received from region $\mathcal{R}_k$. We refer to such parallel CUSUM algorithms by *ensemble CUSUM algorithm* (Algorithm 1). Notice that an iteration of this algorithm is initiated by the collection of an observation.

We are particularly interested in the performance of the ensemble CUSUM algorithm when the observations are collected by autonomous vehicles. In this case, the performance of the ensemble CUSUM algorithm is a function of the vehicle routing policy. For the ensemble CUSUM algorithm with autonomous vehicles collecting observation, let the number of iterations (collection of observations) required to detect an anomaly at region $\mathcal{R}_k$ be $N_k : \Omega \to \mathbb{N} \cup \{+\infty\}$, and let the detection delay, i.e., the time required to detect an anomaly, at region $\mathcal{R}_k$ be $\delta_k : \Omega \to \mathbb{R}_{>0} \cup \{+\infty\}$, for each $k \in \{1, \dots, n\}$, where Ω is the space of vehicle routing policies. We also define average detection delay as follows:

*Definition 2 (**Average detection delay**):* For any vector of weights $(w_1, \dots, w_n) \in \Delta_{n-1}$, define the average detection delay $\delta_{\text{avg}} : \Omega \to \mathbb{R}_{>0} \cup \{+\infty\}$ for the ensemble CUSUM algorithm with autonomous vehicles collecting observations by

$$\delta_{\text{avg}}(\omega) = \sum_{k=1}^{n} w_k \mathbb{E}[\delta_k(\omega)]. \tag{3}$$

For the ensemble CUSUM algorithm with m vehicles collecting observation, define $\delta_k^{m\text{-min}}$ and $\delta_{\text{avg}}^{m\text{-min}}$ by

$$\begin{aligned} \delta_k^{m\text{-min}} &= \inf\{\mathbb{E}[\delta_k(\omega)] \mid \omega \in \Omega\}, \text{ and} \\ \delta_{\text{avg}}^{m\text{-min}} &= \inf\{\delta_{\text{avg}}(\omega) \mid \omega \in \Omega\}, \end{aligned}$$

respectively. Note that $\delta_k^{m\text{-min}}$ and $\delta_{\text{avg}}^{m\text{-min}}$ are lower bounds for the expected detection delay and average detection delay at region $\mathcal{R}_k$, respectively, independently of the routing policy. Let $\bar{\eta} = e^{-\eta} + \eta - 1$. We now state lower bounds on the performance of the ensemble CUSUM algorithm with autonomous vehicles collecting observations.

*Lemma 3 (**Global lower bound**):* The following statements hold for the ensemble CUSUM algorithm with m vehicles collecting information:

(i) the lower bound $\delta_k^{m\text{-min}}$ for the expected detection delay at region $\mathcal{R}_k$ satisfies

$$\delta_k^{m\text{-min}} \geq \frac{\bar{\eta}\ \bar{T}_k^{m\text{-smlst}}}{m\mathcal{D}_k};$$

(ii) the lower bound $\delta_{\text{avg}}^{m\text{-min}}$ for the average detection delay satisfies

$$\delta_{\text{avg}}^{m\text{-min}} \geq \frac{\bar{\eta}\ \bar{T}_{\min}^{m\text{-smlst}}}{m\mathcal{D}_{\max}},$$

where $\bar{T}_{\min}^{m\text{-smlst}} = \min\{\bar{T}_k^{m\text{-smlst}} \mid k \in \{1,\dots,n\}\}$.

Proof: We start by establishing the first statement. We note that a lower bound on the expected detection delay at region $\mathcal{R}_k$ is obtained if all the vehicles always stay at region $\mathcal{R}_k$. Since, each observation is collected from region $\mathcal{R}_k$, the number of iterations of the ensemble CUSUM algorithm required to detect an anomaly at region $\mathcal{R}_k$ satisfies $\mathbb{E}[N_k] = \bar{\eta}/\mathcal{D}_k$. Let $T_k^r(b)$ be realized value of the processing time of vehicle $\mathcal{U}_r$ at its b-th observation. It follows that $T_k^{m\text{-smlst}}(b) = \min\{T_k^r(b) \mid r \in \{1,\dots,m\}\}$ is a lower bound on the processing time of each vehicle for its b-th observation. Further, $T_k^{m\text{-smlst}}(b)$ is identically distributed for each b and $\mathbb{E}[T_k^{m\text{-smlst}}(b)] = \bar{T}_k^{m\text{-smlst}}$. Consider a modified stochastic process where the realized processing time of each vehicle for its bth observation in $T_k^{m\text{-smlst}}(b)$. Indeed, such a stochastic process underestimates the time required to collect each observation and, hence, provides a lower bound to the expected detection delay. Therefore, the detection delay satisfies the following bound

$$\delta_k(\omega) \geq \sum_{b=1}^{\lceil N_k/m \rceil} T_k^{m\text{-smlst}}(b), \text{ for each } \omega \in \Omega.$$

It follows from Wald's identity (Resnick, 1999), that

$$\mathbb{E}[\delta_k(\omega)] \geq \bar{T}_k^{m\text{-smlst}}\mathbb{E}[\lceil N_k/m \rceil] \geq \bar{T}_k^{m\text{-smlst}}\mathbb{E}[N_k]/m.$$

This proves the first statement.

The second statement follows from Definition 2 and the first statement. ■

*Remark 4 (**Dependence across regions**):* We assumed that the observations collected from different regions are mutually independent. If the observations from different regions are dependent, then, at each iteration, instead of updating only one CUSUM statistic, the CUSUM statistic at each region should be updated with an appropriate marginal distribution. □

IV. Randomized Ensemble CUSUM Algorithm

We now study the persistent surveillance problem under randomized ensemble CUSUM algorithm. First, we derive an exact expression for the expected detection delay for the randomized ensemble CUSUM algorithm with a single vehicle, and use the derived expressions to develop an efficient stationary policy for a single vehicle. Second, we develop a lower bound on the expected detection delay for the randomized ensemble CUSUM algorithm with multiple vehicles, and develop a generic partitioning policy that (i) constructs a complete and disjoint m-partition of the regions, (ii) allocates one partition each to a vehicle, and (iii) lets each vehicle survey its assigned region with some single vehicle policy. Finally, we show that the partitioning policy where each vehicle implements the efficient stationary policy in its regions is within a factor of an optimal policy.

A. Analysis for single vehicle

Consider the randomized ensemble CUSUM algorithm with a single vehicle. Let $q_k \in [0,1]$ denote the probability to select region $\mathcal{R}_k$, and let $\boldsymbol{q} = (q_1,\dots,q_n) \in \Delta_{n-1}$. Let the threshold for the CUSUM algorithm at each region be uniform and equal to $\eta > 0$. We note that for the randomized ensemble CUSUM algorithm with a single vehicle the space of vehicle routing policies is $\Omega = \Delta_{n-1}$.

*Theorem 4 (**Single vehicle randomized ensemble CUSUM**):* For the randomized ensemble CUSUM algorithm with a single vehicle and stationary routing policy $\boldsymbol{q} \in \Delta_{n-1}$, the following statements hold:

(i) the number of observations $N_k(\boldsymbol{q})$ required to detect a change at region $\mathcal{R}_k$ satisfies

$$\mathbb{E}_{f_k^1}[N_k(\boldsymbol{q})] = \frac{\bar{\eta}}{q_k\mathcal{D}_k};$$

(ii) the detection delay $\delta_k(\boldsymbol{q})$ at region $\mathcal{R}_k$ satisfies

$$\mathbb{E}_{f_k^1}[\delta_k(\boldsymbol{q})] = \Big(\sum_{i=1}^n q_i\bar{T}_i + \sum_{i=1}^n\sum_{j=1}^n q_iq_jd_{ij}\Big)\mathbb{E}_{f_k^1}[N_k(\boldsymbol{q})].$$

Proof: Let $\tau \in \{1,\dots,N_k\}$ be the iterations at which the vehicle collects and sends information about the regions, where N_k denotes the iteration at which an anomaly is detected at region $\mathcal{R}_k$. Let the log likelihood ratio at region $\mathcal{R}_k$ at iteration τ be λ_τ^k. We have

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

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

$$\mathbb{E}_{f_k^1}[\lambda_\tau^k] = q_k\mathcal{D}_k.$$

The remaining proof of the first statement follows similar to the proof for CUSUM in (Siegmund, 1985).

To prove the second statement, note that the information aggregation time T^{agr} comprises of the processing time and the travel time. At an iteration the vehicle is at region $\mathcal{R}_i$ with probability q_i and picks region $\mathcal{R}_j$ with probability q_j. Additionally, the vehicle travels between the two regions in d_{ij} units of time. Thus, the average travel time at each iteration is

$$\mathbb{E}[T_{\text{travel}}] = \sum_{i=1}^n\sum_{j=1}^n q_iq_jd_{ij}.$$

Hence, the expected information aggregation time at each iteration is

$$\mathbb{E}[T^{\text{agr}}] = \mathbb{E}[T_{\text{travel}} + T_{\text{process}}] = \sum_{i=1}^n\sum_{j=1}^n q_iq_jd_{ij} + \sum_{i=1}^n q_i\bar{T}_i.$$

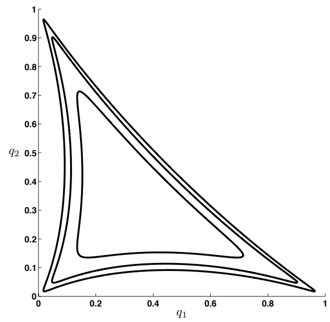

Fig. 2. Level-sets of the objective function in problem (5). It can be seen that the level sets are not convex.

Let $\{T_\tau^{\text{agr}}\}_{\tau\in\{1,\dots,N_k\}}$, be the information aggregation times at each iteration. We have that $\delta_k = \sum_{\tau=1}^{N_k} T_\tau^{\text{agr}}$, and it follows from Wald's identity (Resnick, 1999), that

$$\mathbb{E}[\delta_k] = \mathbb{E}[T^{\text{agr}}]\mathbb{E}[N_k].$$

This completes the proof of the statement. ■

B. Design for single vehicle

Our objective is to design a stationary policy that simultaneously minimizes the detection delay at each region, that is, to design a stationary policy that minimizes each term in $(\delta_1(\boldsymbol{q}),\dots,\delta_n(\boldsymbol{q}))$ simultaneously. For this multiple-objective optimization problem, we construct a single aggregate objective function as the average detection delay. After incorporating the expressions for the expected detection delays derived in Theorem 4, the average detection delay becomes

$$\delta_{\text{avg}}(\boldsymbol{q}) = \Big(\sum_{k=1}^{n}\frac{w_k\bar{\eta}}{q_k\mathcal{D}_k}\Big)\Big(\sum_{i=1}^{n} q_i T_i + \sum_{i=1}^{n}\sum_{j=1}^{n} q_i q_j d_{ij}\Big), \quad (4)$$

where $w_k = \pi_k/(\sum_{i=1}^n \pi_i)$ is the weight on the expected detection delay at region $\mathcal{R}_k$ and π_k is the prior probability of an anomaly being present at region $\mathcal{R}_k$. Our objective is to solve the average detection delay minimization problem:

$$\underset{\boldsymbol{q}\in\Delta_{n-1}}{\text{minimize}} \quad \delta_{\text{avg}}(\boldsymbol{q}). \quad (5)$$

In general, the objective function δ_{avg} is non-convex. For instance, let $n=3$, and consider the level sets of δ_{avg} on the two dimensional probability simplex (Fig. 2). It can be seen that the level sets are non-convex, yet there exists a unique critical point and it corresponds to a minimum. We now state the following conjecture about the average detection delay:

*Conjecture 5 (**Single vehicle optimal stationary policy**):* For the randomized ensemble CUSUM algorithm with a single vehicle, the average detection delay function δ_{avg} has a unique critical point at which the minimum of δ_{avg} is achieved. □.

In the Appendix we provide probabilistic guarantees that, for a particular stochastic model of the parameters in δ_{avg}, with at least confidence level 99.99% and probability at least 99%, the optimization problem (5) has a unique critical point at which the minimum is achieved. Such a minimum can be computed via standard gradient-descent methods (Boyd and Vandenberghe, 2004).

We now construct an upper bound for the expected detection delay. We will show that minimization of this upper bound yields a policy that is within a factor of an optimal policy. From equation (4), we define the upper bound $\delta_{\text{upper}}:\Delta_{n-1}\to\mathbb{R}_{>0}\cup\{+\infty\}$ as

$$\delta_{\text{avg}}(\boldsymbol{q}) \le \delta_{\text{upper}}(\boldsymbol{q}) = \Big(\sum_{k=1}^{n}\frac{w_k\bar{\eta}}{q_k\mathcal{D}_k}\Big)(\bar{T}_{\max}+d_{\max}),$$

where $d_{\max} = \max\{d_{ij} \mid i,j\in\{1,\dots,n\}\}$.

*Theorem 6 (**Single vehicle efficient stationary policy**):* The following statements hold for the randomized ensemble CUSUM algorithm with single vehicle:

(i) the upper bound on the expected detection delay satisfies

$$\min_{\boldsymbol{q}\in\Delta_{n-1}} \delta_{\text{upper}}(\boldsymbol{q}) = \Big(\sum_{k=1}^{n}\sqrt{\frac{w_k}{\mathcal{D}_k}}\Big)^2 \bar{\eta}(\bar{T}_{\max}+d_{\max}),$$

and the minimum is achieved at $\boldsymbol{q}^\dagger$ defined by

$$q_k^\dagger = \frac{\sqrt{w_k/\mathcal{D}_k}}{\sum_{j=1}^n\sqrt{w_j/\mathcal{D}_j}},\ k\in\{1,\dots,n\};$$

(ii) the average detection delay satisfies the following lower bound

$$\delta_{\text{avg}}(\boldsymbol{q}) \ge \Big(\sum_{k=1}^{n}\sqrt{\frac{w_k}{\mathcal{D}_k}}\Big)^2 \bar{\eta}\,\bar{T}_{\min},$$

for all $\boldsymbol{q}\in\Delta_{n-1}$;

(iii) the stationary policy $\boldsymbol{q}^\dagger$ is within a factor of optimal, that is

$$\frac{\delta_{\text{avg}}(\boldsymbol{q}^\dagger)}{\delta_{\text{avg}}(\boldsymbol{q}^*)} \le \frac{\bar{T}_{\max}+d_{\max}}{\bar{T}_{\min}}, \text{ and}$$
$$\frac{\delta_{\text{avg}}(\boldsymbol{q}^\dagger)}{\delta_{\text{avg}}^{\text{1-min}}} \le n\frac{\bar{T}_{\max}+d_{\max}}{\bar{T}_{\min}}\frac{\mathcal{D}_{\max}}{\mathcal{D}_{\min}},$$

where $\boldsymbol{q}^*$ is an optimal stationary policy;

(iv) the expected detection delay at region $\mathcal{R}_k$ under policy $\boldsymbol{q}^\dagger$ satisfy

$$\frac{\mathbb{E}[\delta_k(\boldsymbol{q}^\dagger)]}{\delta_k^{\text{1-min}}} \le \frac{(\bar{T}_{\max}+d_{\max})}{\bar{T}_k}\sqrt{\frac{n\mathcal{D}_k}{w_k\mathcal{D}_{\min}}}.$$

Proof: We start by establishing the first statement. It follows from the stationarity conditions on the Lagrangian that the minimizer $\boldsymbol{q}^\dagger$ of δ_{upper} satisfy $q_k^\dagger \propto \sqrt{w_k\bar{\eta}/\mathcal{D}_k}$, for each $k\in\{1,\dots,n\}$. Incorporating this fact into $\sum_{k=1}^n q_k^\dagger = 1$ yields the expression for $q_k^\dagger$. The expression for $\delta_{\text{upper}}(\boldsymbol{q}^\dagger)$ can be verified by substituting the expression for $\boldsymbol{q}^\dagger$ into δ_{upper}.

To prove the second statement, we construct a lower bound $\delta_{\text{lower}}:\Delta_{n-1}\to\mathbb{R}_{>0}\cup\{+\infty\}$ to the average detection delay δ_{avg} defined by $\delta_{\text{lower}}(\boldsymbol{q}) = \sum_{k=1}^n w_k\bar{\eta}\bar{T}_{\min}/\mathcal{D}_k q_k$. It can be verified that δ_{lower} also achieves its minimum at $\boldsymbol{q}^\dagger$, and

$$\delta_{\text{lower}}(\boldsymbol{q}^\dagger) = \Big(\sum_{k=1}^{n}\sqrt{\frac{w_k}{\mathcal{D}_k}}\Big)^2 \bar{\eta}\,\bar{T}_{\min}.$$

We note that

$$\delta_{\text{lower}}(\boldsymbol{q}^\dagger) \le \delta_{\text{lower}}(\boldsymbol{q}^*) \le \delta_{\text{avg}}(\boldsymbol{q}^*) \le \delta_{\text{avg}}(\boldsymbol{q}), \forall \boldsymbol{q}\in\Delta_{n-1}.$$

Thus, the second statement follows.

To prove the first part of the third statement, we note that

$$\delta_{\text{lower}}(\boldsymbol{q}^\dagger) \leq \delta_{\text{avg}}(\boldsymbol{q}^*) \leq \delta_{\text{avg}}(\boldsymbol{q}^\dagger) \leq \delta_{\text{upper}}(\boldsymbol{q}^\dagger).$$

Therefore, the policy $\boldsymbol{q}^\dagger$ is within $\delta_{\text{upper}}(\boldsymbol{q}^\dagger)/\delta_{\text{lower}}(\boldsymbol{q}^\dagger) = (T_{\max} + d_{\max})/T_{\min}$ factor of optimal stationary policy.

To prove the second part of the third statement, we note

$$\begin{aligned}\frac{\delta_{\text{avg}}(\boldsymbol{q}^\dagger)}{\delta_{\text{avg}}^{\text{1-min}}} &\leq \frac{\mathcal{D}_{\max}(\bar{T}_{\max} + d_{\max})}{\mathcal{D}_{\min}\bar{T}_{\min}}(\sqrt{w_1} + \ldots + \sqrt{w_n})^2 \\ &\leq n\frac{\bar{T}_{\max} + d_{\max}}{\bar{T}_{\min}}\frac{\mathcal{D}_{\max}}{\mathcal{D}_{\min}},\end{aligned}$$

where the last inequality follows from the fact: $\max\{\sqrt{w_1} + \ldots + \sqrt{w_n} \mid w_1 + \ldots + w_n = 1\} = \sqrt{n}$.

To establish the last statement, we note that

$$\begin{aligned}\frac{\mathbb{E}[\delta_k(\boldsymbol{q}^\dagger)]}{\delta_k^{\text{1-min}}} &\leq \frac{(\bar{T}_{\max} + d_{\max})}{q_k^\dagger \bar{T}_k} \\ &\leq \frac{(\bar{T}_{\max} + d_{\max})}{\bar{T}_k}\sqrt{\frac{\mathcal{D}_k}{w_k\mathcal{D}_{\min}}}(\sqrt{w_1} + \ldots + \sqrt{w_n}) \\ &\leq \frac{(\bar{T}_{\max} + d_{\max})}{\bar{T}_k}\sqrt{\frac{n\mathcal{D}_k}{w_k\mathcal{D}_{\min}}}.\end{aligned}$$

This concludes the proof of the theorem. ■

In the following, we would refer to $\boldsymbol{q}^\dagger$ as the *single vehicle efficient stationary policy*.

*Remark 5 (**Efficient stationary policy**):* As opposed to the average detection delay δ_{avg}, the upper bound δ_{upper} does not depend upon any travel time information. Then, our efficient policy does not take this information into account, and it may not be optimal. Instead, an optimal policy allocates higher visiting probabilities to regions that are located more centrally in the environment. We resort to the efficient policy because (i) if the problem (5) does not admit a unique minimum, then the optimal policy can not be computed efficiently; and (ii) the efficient policy has an intuitive, tractable, and closed form expression. □

C. Analysis for multiple vehicles

We now consider the randomized ensemble CUSUM with $m > 1$ vehicles. In this setting the vehicles operate in an asynchronous fashion. This asynchronicity, which did not occur in the single vehicle case, is due to (i) different travel times between two different pair of regions, and (ii) different realized value of processing time at each iteration. Such an asynchronous operation makes the time durations between two subsequent iterations non-identically distributed and makes it difficult to obtain closed form expressions for the expected detection delay at each region.

Motivated by the above discussion, we determine a lower bound on the expected detection delay for the randomized ensemble CUSUM algorithm with multiple vehicles. Let $\boldsymbol{q}^r = (q_1^r, \ldots, q_n^r) \in \Delta_{n-1}$ denote the *stationary policy* for vehicle $\mathcal{U}_r$, i.e., the vector of probabilities of selecting different regions for vehicle $\mathcal{U}_r$, and let $\vec{\boldsymbol{q}}_m = (\boldsymbol{q}^1, \ldots, \boldsymbol{q}^m) \in \Delta_{n-1}^m$. We note that for the randomized ensemble CUSUM algorithm with m vehicles the space of vehicle routing policies is $\Omega = \Delta_{n-1}^m$. We construct a lower bound on the processing times at different regions for different vehicles in the following way. Let Ξ be the set of all the sets with cardinality m in which each entry is an arbitrarily chosen region; equivalently, $\Xi = \{\mathcal{R}_1, \ldots, \mathcal{R}_n\}^m$. Let a realization of the processing times at the regions in a set $\xi \in \Xi$ be $t_1^\xi, \ldots, t_m^\xi$. We now define a lower bound $\bar{T}_{\text{one}}$ to the expected value of the minimum of the processing times at m arbitrary regions as $\bar{T}_{\text{one}} = \min\{\mathbb{E}[\min\{t_1^\xi, \ldots, t_m^\xi\}] \mid \xi \in \Xi\}$.

Algorithm 2: *Partitioning Algorithm*

Input : vehicles $\{\mathcal{U}_1, \ldots, \mathcal{U}_m\}$, regions $\mathcal{R} = \{\mathcal{R}_1, \ldots, \mathcal{R}_n\}$, a single vehicle routing policy;
Require : $n > m$;
Output : a m-partition of the regions ;

1 partition $\mathcal{R}$ into m arbitrary subsets $\{\mathcal{S}^r\}_{r\in\{1,\ldots,m\}}$ with cardinalities $n_r \leq \lceil n/m \rceil, r \in \{1, \ldots, m\}$;
2 allocate vehicle $\mathcal{U}_r$ to subset $\mathcal{S}^r$, for each $r \in \{1, \ldots, m\}$;
3 implement the single vehicle efficient stationary policy in each subset.

*Theorem 7 (**Multi-vehicle randomized ensemble CUSUM**):* For the randomized ensemble CUSUM algorithm with m vehicles and stationary region selection policies $\boldsymbol{q}^r$, $r \in \{1, \ldots, m\}$, the detection delay δ^k at region $\mathcal{R}_k$ satisfies:

$$\mathbb{E}_{f_k^1}[\delta_k(\vec{\boldsymbol{q}}_m)] \geq \frac{\bar{\eta}\,\bar{T}_{\text{one}}}{\sum_{r=1}^m q_k^r \mathcal{D}_k}.$$

Proof: We construct a modified stochastic process to determine a lower bound on the expected detection delay. For the randomized ensemble CUSUM algorithm with multiple vehicles, let t_r^b be the the processing time for the vehicle $\mathcal{U}_r$ during its b-th visit to any region. We assume that the sampling time for each vehicle at its b-th visit in the modified process is $\min\{t_1^b, \ldots, t_m^b\}$. Therefore, the sampling time for the modified process is the same at each region. Further, it is identically distributed for each visit and has expected value greater than or equal to $\bar{T}_{\text{one}}$. We further assume that the distances between the regions are zero. Such a process underestimates the processing and travel time required to collect each observation in the randomized ensemble CUSUM algorithm. Hence, the expected detection delay for this process provides a lower bound to the expected detection delay for randomized ensemble CUSUM algorithm. Further, for this process the vehicles operate synchronously and the expected value of the likelihood ratio at region k at each iteration is $\sum_{r=1}^m q_k^r \mathcal{D}(f_k^1, f_k^0)$. The remainder of the proof follows similar to the proof for single vehicle case in Theorem 4. ■

D. Design for multiple vehicles

We now design an efficient stationary policy for randomized ensemble CUSUM algorithm with multiple vehicles. We propose an algorithm that partitions the set of regions into m subsets, allocates one vehicle to each subset, and implements our single vehicle efficient stationary policy in each subset. This procedure is formally defined in Algorithm 2.

Let the subset of regions allocated to vehicle $\mathcal{U}_r$ be $\mathcal{S}^r, r \in \{1, \ldots, m\}$. We will denote the elements of subset $\mathcal{S}^r$ by $\mathcal{S}_i^r, i \in \{1, \ldots, n_r\}$. Let $\vec{\boldsymbol{q}}_{\text{part}}^\dagger \in \Delta_{n-1}^m$ be a stationary routing policy under the partitioning algorithm that implements single vehicle efficient stationary policy in each partition. We define the weights in equation (3) by $w_k = \pi_k^1/\sum_{j=1}^n \pi_j^1$, where π_k^1 is the prior probability of anomaly at region $\mathcal{R}_k$. Let $w_{\min} = \min\{w_1, \dots, w_n\}$ and $w_{\max} = \max\{w_1, \dots, w_n\}$. We now analyze the performance of the partitioning algorithm and show that it is within a factor of optimal.

*Theorem 8 (**Performance of the partitioning policy**):* For the partitioning algorithm with m vehicles and n regions that implements the single vehicle efficient stationary policy in each partition, the following statements hold:

(i) the average detection delay under partitioning policy satisfies the following upper bound

$$\delta_{\text{avg}}(\vec{\boldsymbol{q}}_{\text{part}}^\dagger) \leq m \left\lceil \frac{n}{m} \right\rceil^2 \frac{w_{\max}\bar{\eta}(\bar{T}_{\max} + d_{\max})}{\mathcal{D}_{\min}};$$

(ii) the average detection delay satisfies the following lower bound

$$\delta_{\text{avg}}(\vec{\boldsymbol{q}}_m) \geq \Big(\sum_{k=1}^n \sqrt{\frac{w_k}{\mathcal{D}_k}} \Big)^2 \frac{\bar{\eta}\bar{T}_{\text{one}}}{m},$$

for any $\vec{\boldsymbol{q}}_m \in \Delta_{n-1}^m$;

(iii) the stationary policy $\vec{\boldsymbol{q}}_{\text{part}}^\dagger$ is within a factor of optimal, and

$$\frac{\delta_{\text{avg}}(\vec{\boldsymbol{q}}_{\text{part}}^\dagger)}{\delta_{\text{avg}}(\vec{\boldsymbol{q}}_m^*)} \leq \frac{4w_{\max}}{w_{\min}} \frac{(\bar{T}_{\max} + d_{\max})}{\bar{T}_{\text{one}}} \frac{\mathcal{D}_{\max}}{\mathcal{D}_{\min}}, \text{ and}$$
$$\frac{\delta_{\text{avg}}(\vec{\boldsymbol{q}}_{\text{part}}^\dagger)}{\delta_{\text{avg}}^{m\text{-min}}} \leq m^2 \left\lceil \frac{n}{m} \right\rceil \frac{(\bar{T}_{\max} + d_{\max})}{\bar{T}_{\min}^{m\text{-smlst}}} \frac{\mathcal{D}_{\max}}{\mathcal{D}_{\min}},$$

where $\vec{\boldsymbol{q}}_m^*$ is optimal stationary policy;

(iv) the expected detection delay at region $\mathcal{R}_k$ under the stationary policy $\vec{\boldsymbol{q}}_{\text{part}}^\dagger$ satisfies

$$\frac{\mathbb{E}[\delta_k(\vec{\boldsymbol{q}}_{\text{part}}^\dagger)]}{\delta_k^{m\text{-min}}} \leq \frac{m(\bar{T}_{\max} + d_{\max})}{\bar{T}_k^{m\text{-smlst}}} \sqrt{\left\lceil \frac{n}{m} \right\rceil \frac{\mathcal{D}_k}{w_k \mathcal{D}_{\min}}}.$$

Proof: We start by establishing the first statement. We note that under the partitioning policy, the maximum number of regions a vehicle serves is $\lceil n/m \rceil$. It follows from Theorem 6 that for vehicle $\mathcal{U}_r$ and the associated partition $\mathcal{S}^r$, the average detection delay is upper bounded by

$$\begin{aligned}\delta_{\text{avg}}(\boldsymbol{q}_{\text{part}}^r) &\leq \Big(\sum_{i=1}^{n_r} \sqrt{\frac{w_i}{\mathcal{D}_i}} \Big)^2 \bar{\eta}(\bar{T}_{\max} + d_{\max}) \\ &\leq \left\lceil \frac{n}{m} \right\rceil^2 \frac{\bar{\eta} w_{\max}(\bar{T}_{\max} + d_{\max})}{\mathcal{D}_{\min}}.\end{aligned}$$

Therefore, the overall average detection delay satisfies $\delta_{\text{avg}}(\vec{\boldsymbol{q}}_{\text{part}}^\dagger) \leq m\delta_{\text{avg}}(\boldsymbol{q}_{\text{part}}^r)$. This establishes the first statement.

To prove the second statement, we utilize the lower bounds obtained in Theorem 7 and construct a lower bound to the average detection delay $\delta_{\text{lower}}^m : \Delta_{n-1}^m \to \mathbb{R}_{>0} \cup \{+\infty\}$ defined by $\delta_{\text{lower}}^m(\vec{\boldsymbol{q}}_m) = \sum_{k=1}^n (v_k \bar{T}_{\text{one}}/\sum_{r=1}^m q_k^r)$. It can be verified that

$$\min_{\vec{\boldsymbol{q}}_m \in \Delta_{n-1}^m} \delta_{\text{lower}}^m(\vec{\boldsymbol{q}}_m) = \Big(\sum_{k=1}^n \sqrt{\frac{w_k}{\mathcal{D}_k}} \Big)^2 \frac{\bar{\eta}\bar{T}_{\text{one}}}{m}.$$

We now establish the first part of the third statement. Note that

$$\begin{aligned}\frac{\delta_{\text{avg}}(\vec{\boldsymbol{q}}_{\text{part}}^\dagger)}{\delta_{\text{avg}}(\vec{\boldsymbol{q}}_m^*)} &\leq \frac{\lceil n/m \rceil^2}{(n/m)^2} \frac{w_{\max}}{w_{\min}} \frac{(\bar{T}_{\max} + d_{\max})}{\bar{T}_{\text{one}}} \frac{\mathcal{D}_{\max}}{\mathcal{D}_{\min}} \\ &\leq \frac{4w_{\max}}{w_{\min}} \frac{(\bar{T}_{\max} + d_{\max})}{\bar{T}_{\text{one}}} \frac{\mathcal{D}_{\max}}{\mathcal{D}_{\min}},\end{aligned}$$

where the last inequality follows from the fact that $(\lceil n/m \rceil)/(n/m) \leq 2$.

The remainder of the proof follows similar to the proof of Theorem 6. ■

Algorithm 3: *Single Vehicle Adaptive Ensemble CUSUM*

Input : parameters η, $\mathcal{D}_k$, pdfs f_k^0, f_k^1, for each $k \in \{1, \dots, n\}$;
Output : decision on anomaly at each region ;

set $\Lambda_0^j = 0$, for all $j \in \{1, \dots, n\}$, and $\tau = 1$;
while true **do**
  set new prior $\pi_k^1 = e^{\Lambda_\tau^k}/(1 + e^{\Lambda_\tau^k})$, for each $k \in \{1, \dots, n\}$
  set $q_k = \dfrac{\sqrt{\pi_k^1/\mathcal{D}_k}}{\sum_{j=1}^n \sqrt{\pi_j^1/\mathcal{D}_j}}$, for each $k \in \{1, \dots, n\}$;
  sample a region from probability distribution $(q_1, \dots, q_n)$;
  collect sample y_τ from region k;
  update the CUSUM statistic at each region

$$\Lambda_\tau^j = \begin{cases} \Big(\Lambda_{\tau-1}^k + \log \frac{f_k^1(y_\tau)}{f_k^0(y_\tau)} \Big)^+, & \text{if } j = k; \\ \Lambda_{\tau-1}^j, & \text{if } j \in \{1, \dots, n\} \setminus \{k\}; \end{cases}$$

  if $\Lambda_\tau^k > \eta$ **then**
    anomaly detected at region $\mathcal{R}_k$;
    set $\Lambda_\tau^k = 0$;
  set $\tau = \tau + 1$;

V. Adaptive ensemble CUSUM Algorithm

The stationary vehicle routing policy does not utilize the real-time information regarding the likelihood of anomalies at the regions. We now develop an adaptive policy that incorporates the anomaly likelihood information provided by the anomaly detection algorithm. We consider the CUSUM statistic at a region as a measure of the likelihood of an anomaly at that region, and utilize it at each iteration to design new prior probability of an anomaly for each region. At each iteration, we adapt the efficient stationary policy using this new prior probability. This procedure results in higher probability of visiting an anomalous region and, consequently, it improves the performance of our efficient stationary policy. In Section VI we provide numerical evidence showing that the adaptive ensemble CUSUM algorithm improves the performance of randomized ensemble CUSUM algorithm.

Our adaptive ensemble CUSUM algorithm is formally presented in Algorithm 3 for the single vehicle case. For the case of multiple vehicles we resort to the partitioning Algorithm 2 that implements the single vehicle adaptive ensemble CUSUM Algorithm 3 in each partition. Let us denote the adaptive routing policy for a single vehicle by $\boldsymbol{a}$ and the policy obtained from the partitioning algorithm that implements single vehicle adaptive routing policy in each partition by $\boldsymbol{a}_{\text{part}}$. We now analyze the performance of the adaptive ensemble CUSUM algorithm. Since, the probability to visit any region varies with time in the adaptive ensemble CUSUM algorithm, we need to determine the number of iterations between two consecutive visit to a region, i.e., the number of iterations for the recurrence of the region. We first derive a bound on the expected number of samples to be drawn from a time-varying probability vector for the recurrence of a particular state.

*Lemma 9 (**Mean observations for region recurrence**):* Consider a sequence $\{x_\tau\}_{\tau\in\mathbb{N}}$, where x_τ is sampled from a probability vector $\boldsymbol{p}^\tau \in \Delta_{n-1}$. If the kth entry of $\boldsymbol{p}^\tau$ satisfy $p_k^\tau \in (\alpha_k, \beta_k)$, for each $\tau \in \mathbb{N}$ and some $\alpha_k, \beta_k \in (0,1)$, then the number of iterations I_k for the recurrence of state k satisfy $\mathbb{E}[I_k] \le \beta_k/\alpha_k^2$.

Proof: The terms of the sequence $\{x_\tau\}_{\tau\in\mathbb{N}}$ are statistically independent. Further, the probability mass function p^τ is arbitrary. Therefore, the bound on the expected iterations for the first occurrence of state k is also a bound on the subsequent recurrence of state k. The expected number of iterations for first occurrence of region k are

$$\mathbb{E}[I_k] = \sum_{i\in\mathbb{N}} i p_k^i \prod_{j=1}^{i-1}(1-p_k^j) \le \beta_k \sum_{i\in\mathbb{N}} i(1-\alpha_k)^{i-1} = \beta_k/\alpha_k^2.$$

This establishes the statement. ■

We utilize this upper bound on the expected number of iterations for recurrence of a region to derive performance metrics for the adaptive ensemble CUSUM algorithm. We now derive an upper bound on the expected detection delay at each region for adaptive ensemble CUSUM algorithm. We derive these bounds for the expected evolution of the CUSUM statistic at each region.

*Theorem 10 (**Adaptive ensemble CUSUM algorithm**):* Consider the expected evolution of the CUSUM statistic at each region. For the partitioning algorithm that implements single vehicle adaptive ensemble CUSUM algorithm (Algorithm 3) in each subset of the partition, the following statement holds:

$$\mathbb{E}[\delta_k(\boldsymbol{a}_{\text{part}})] \le \Big(\frac{\bar\eta}{\mathcal{D}_k} + \frac{2(\lceil n/m\rceil - 1)e^{\eta/2}\sqrt{\mathcal{D}_k}(1-e^{-\bar\eta/2})}{\sqrt{\mathcal{D}_{\min}}(1-e^{-\mathcal{D}_k/2})} + \frac{(\lceil n/m\rceil-1)^2 e^{\eta}\mathcal{D}_k(1-e^{-\bar\eta})}{\mathcal{D}_{\min}(1-e^{-\mathcal{D}_k})}\Big)(\bar T_{\max} + d_{\max}).$$

Proof: We start by deriving expression for a single vehicle. Let the number of iterations between the $(j-1)$th and jth visit to region $\mathcal{R}_k$ be I_j^k.

Let the observation during the jth visit to region $\mathcal{R}_k$ be y_j and the CUSUM statistic at region $\mathcal{R}_k$ after the visit be C_j^k. It follows that the probability to visit region $\mathcal{R}_k$ between $(j-1)$th and jth visit is greater than

$$p_k^{j-1} = \frac{e^{C_{j-1}^k/2}/\sqrt{\mathcal{D}_k}}{e^{C_{j-1}^k/2}/\sqrt{\mathcal{D}_k} + (n-1)e^{\eta/2}/\sqrt{\mathcal{D}_{\min}}}.$$

Therefore, it follows from Lemma 9 that

$$\mathbb{E}[I_j^k] \le (1 + (n-1)e^{(\eta - C_{j-1}^k)/2}\sqrt{\mathcal{D}_k/\mathcal{D}_{\min}})^2.$$

Note that $C_j^k = \max\{0, C_{j-1}^k + \log(f_k^1(y_j)/f_k^0(y_j))\}$. Since, maximum of two convex function is a convex function, it follows from Jensen inequality (Resnick, 1999), that

$$\mathbb{E}[C_j^k] \ge \max\{0, \mathbb{E}[C_{j-1}^k] + \mathcal{D}_k\} \ge \mathbb{E}[C_{j-1}^k] + \mathcal{D}_k.$$

Therefore, $\mathbb{E}[C_j^k] \ge j\mathcal{D}_k$ and for expected evolution of the CUSUM statistics

$$\mathbb{E}[I_j^k] \le (1 + (n-1)e^{(\eta-(j-1)\mathcal{D}_k)/2}\sqrt{\mathcal{D}_k/\mathcal{D}_{\min}})^2.$$

Therefore, the total number of iterations N_k required to collect N_k^{obs} observations at region $\mathcal{R}_k$ satisfy

$$\begin{aligned}\mathbb{E}[N_k(\boldsymbol{a})|N_k^{\text{obs}}] &= \sum_{j=1}^{N_k^{\text{obs}}}(1+(n-1)e^{(\eta-(j-1)\mathcal{D}_k)/2}\sqrt{\mathcal{D}_k/\mathcal{D}_{\min}})^2\\ &= N_k^{\text{obs}} + \frac{2(n-1)e^{\eta/2}\sqrt{\mathcal{D}_k}(1-e^{-\mathcal{D}_k N_k^{\text{obs}}/2})}{\sqrt{\mathcal{D}_{\min}}(1-e^{-\mathcal{D}_k/2})}\\ &\quad + \frac{(n-1)^2 e^{\eta}\mathcal{D}_k(1-e^{-\mathcal{D}_k N_k^{\text{obs}}})}{\mathcal{D}_{\min}(1-e^{-\mathcal{D}_k})}.\end{aligned}$$

Note that the number of observations N_k^{obs} required at region $\mathcal{R}_k$ satisfy $\mathbb{E}[N_k^{\text{obs}}] = \bar\eta/\mathcal{D}_k$. It follows from Jensen's inequality that

$$\mathbb{E}[N_k(\boldsymbol{a})] \le \frac{\bar\eta}{\mathcal{D}_k} + \frac{2(n-1)e^{\eta/2}\sqrt{\mathcal{D}_k}(1-e^{-\bar\eta/2})}{\sqrt{\mathcal{D}_{\min}}(1-e^{-\mathcal{D}_k/2})} + \frac{(n-1)^2 e^{\eta}\mathcal{D}_k(1-e^{-\bar\eta})}{\mathcal{D}_{\min}(1-e^{-\mathcal{D}_k})}.$$

Since the expected time required to collect each evidence is smaller $\bar T_{\max} + d_{\max}$, it follows that

$$\mathbb{E}[\delta_k(\boldsymbol{a})] \le (\bar T_{\max} + d_{\max})\mathbb{E}[N_k(\boldsymbol{a})].$$

The expression for the partitioning policy that implement single vehicle adaptive routing policy in each partition follow by substituting $\lceil n/m\rceil$ in the above expressions. This completes the proof of the theorem. ■

*Remark 6 (**Performance bound**):* The bound derived in Theorem 10 is very conservative. Indeed, it assumes the CUSUM statistic at each region to be fixed at its maximum value η, except for the region in consideration. This is practically never the case. In fact, if at some iteration the CUSUM statistic is close to η, then it is highly likely that the vehicle visits that region at the next iteration, so that the updated statistic crosses the threshold η and resets to zero. □

VI. Numerical Results

We now elucidate on the concepts developed in this paper through some numerical examples. We first validate the expressions for expected detection delay obtained in Section IV.

*Example 1 (**Expected detection delay**):* Consider a set of 4 regions surveyed by a single vehicle. Let the location of the regions be $(10,0), (5,0), (0,5)$, and $(0,10)$, respectively. The vector of processing times at each region is $(1,2,3,4)$. Under the nominal conditions, the observations at each region are sampled from normal distributions $\mathcal{N}(0,1), \mathcal{N}(0,1.33), \mathcal{N}(0,1.67)$ and $\mathcal{N}(0,2)$, respectively, while under anomalous conditions, the observations

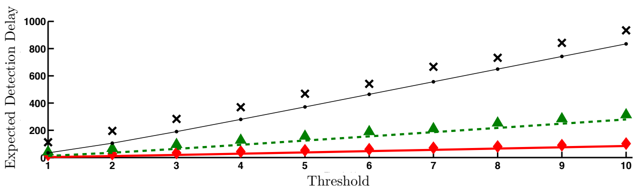

(a) Expected detection delay at region $\mathcal{R}_1$

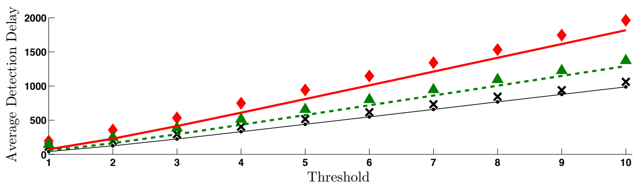

(b) Average detection delay

Fig. 3. Expected and average detection delay. Solid black line with dots and black $\times$, respectively, represent the theoretical expression and the value obtained by Monte-Carlo simulations under stationary policy $\boldsymbol{q} = [\ 0.2\ 0.25\ 0.25\ 0.3\ ]$. Dashed green line and green triangles, respectively, represent the theoretical expression and the value obtained by Monte-Carlo simulations under stationary policy $\boldsymbol{q} = [\ 0.5\ \ 0.2\ \ 0.2\ \ 0.1\ ]$. Solid red line and red diamonds, respectively, represent the theoretical expression and the value obtained by Monte-Carlo simulations under stationary policy $\boldsymbol{q} = [\ 0.85\ 0.05\ 0.05\ 0.05\ ]$.

are sampled from normal distributions with unit mean and same variance as in nominal case. Let the prior probability of anomaly at each region be 0.5. An anomaly appears at each region at time $50, 200, 350$, and 500, respectively. Assuming that the vehicle is holonomic and moves at unit speed, the expected detection delay at region $\mathcal{R}_1$ and the average detection delay are shown in Fig. 3. It can be seen that the theoretical expressions provide a lower bound to the expected detection delay obtained through Monte-Carlo simulations. This phenomenon is attributed to the Wald's approximation.□

We remarked in Section II that if each region cannot be reached from another region in a single hop, then a fastest mixing Markov chain (FMMC) with the desired stationary distribution can be constructed. Consider a set of regions modeled by the graph $\mathcal{G} = (V, \mathcal{E})$, where V is the set of nodes (each node corresponds to a region) and $\mathcal{E}$ is the set of edges representing the connectivity of the regions. The transition matrix of the FMMC $P \in \mathbb{R}^{n \times n}$ with a desired stationary distribution $\boldsymbol{q} \in \Delta_{n-1}$ can be determined by solving the following convex minimization problem (Boyd et al., 2004):

$$\begin{aligned} \text{minimize} \quad & \|Q^{1/2} P Q^{1/2} - \boldsymbol{q}_{\text{root}} \boldsymbol{q}_{\text{root}}^T\|_2 \\ \text{subject to} \quad & P\mathbf{1} = \mathbf{1} \\ & QP = P^T Q \\ & P_{ij} \geq 0, \ \text{for each } (i,j) \in \mathcal{E} \\ & P_{ij} = 0, \ \text{for each } (i,j) \notin \mathcal{E}, \end{aligned}$$

where Q is a diagonal matrix with diagonal $\boldsymbol{q}$, $\boldsymbol{q}_{\text{root}} = (\sqrt{q_1}, \dots, \sqrt{q_n})$, and $\mathbf{1}$ is the vector of all ones. We now demonstrate the effectiveness of FMMC in our setup.

*Example 2 (**Effectiveness of FMMC**):* Consider the same set of data as in Example 1. We study the expected and average detection delay for randomized ensemble CUSUM algorithm when the regions to visit are sampled from the FMMC. The expected and average detection delay for all-to-all connection topology, line connection topology and ring connection topology are shown in Fig. 4. It can be seen that the performance under all three topologies is remarkably close to each other. □

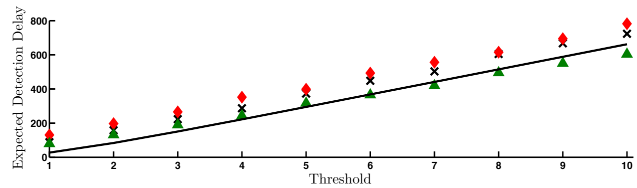

(a) Expected detection delay at region $\mathcal{R}_1$

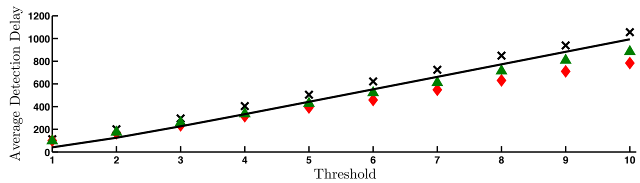

(b) Average detection delay

Fig. 4. Expected and average detection delay for uniform stationary policy. The solid black line represents the theoretical expression. The black $\times$, red diamonds, and green triangles, respectively, represent the values obtained by Monte-Carlo simulations for all-to-all, line, and ring connection topology. For the line and ring topologies, the region to visit at each iteration is sampled from the fastest mixing Markov chain with the desired stationary distribution.

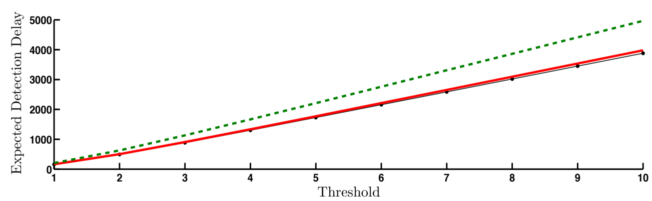

Fig. 5. Average detection delay for a single vehicle. The solid red line, the dashed green line, and the solid black line with dots represent efficient, uniform, and optimal stationary policies, respectively.

We now study the performance of the (numerically computed) optimal and our efficient stationary policies for the single vehicle randomized ensemble CUSUM algorithm.

*Example 3 (**Single vehicle optimal stationary policy**):* For the same set of data as in Example 1, we now study the performance of the uniform, the (numerically computed) optimal and our efficient stationary routing policies. A comparison is shown in Fig. 5. Notice that the performance of the optimal and efficient stationary policy is extremely close to each other. □

We now study the performance of the optimal, partitioning and uniform stationary policies for randomized ensemble CUSUM algorithm with multiple vehicles.

*Example 4 (**Multiple-vehicle optimal stationary policy**):* Consider a set of 6 regions surveyed by 3 vehicles. Let the regions be located at $(10, 0), (5, 0), (0, 5), (0, 10), (0, 0)$ and $(5, 5)$. Let the processing time at each region be unitary. Under nominal conditions, the observations at each region are sampled from normal distributions $\mathcal{N}(0, 1)$, $\mathcal{N}(0, 1.4)$, $\mathcal{N}(0, 1.8)$, $\mathcal{N}(0, 2.2)$, $\mathcal{N}(0, 2.6)$ and $\mathcal{N}(0, 3)$, respectively. Under anomalous conditions, the observations are sampled

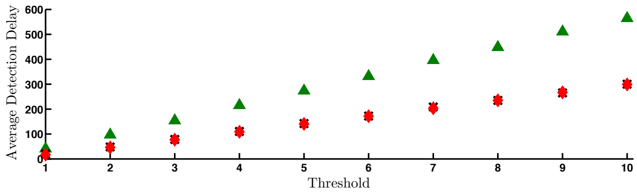

Fig. 6. Average detection delay for 3 vehicles surveying 6 regions. The green triangles represent the policy in which each vehicle surveys each region uniformly. The red diamonds and black $\times$ represent the partitioning policy in which each vehicle implements the single vehicle efficient stationary policy and the single vehicle optimal stationary policy, respectively.

from normal distributions with unit mean and same variance as in the nominal case. Let the prior probability of anomaly at each region be 0.5. An anomaly appears at each region at time $25, 35, 45, 55, 65$ and 75, respectively. Assuming that the vehicles are holonomic and moves at unitary speed, the average detection delay for the uniform stationary policy for each vehicle, the partitioning policy in which each vehicle implements single vehicle efficient stationary policy in each subset of the partition, and the partitioning policy in which each vehicle implements single vehicle optimal stationary policy in each subset of the partition is shown in Fig. 6. □

We now study the performance of the adaptive ensemble CUSUM algorithm, and we numerically show that it improves the performance of our stationary policy.

*Example 5 (**Adaptive ensemble CUSUM algorithm**):* Consider the same set of regions as in Example 1. Let the processing time at each region be unitary. The observations at each region are sampled from normal distributions $\mathcal{N}(0, \sigma^2)$ and $\mathcal{N}(1, \sigma^2)$, in nominal and anomalous conditions, respectively. Under the nominal conditions at each region and $\sigma^2 = 1$, a sample evolution of the adaptive ensemble CUSUM algorithm is shown in Fig. 7(a). The anomaly appears at regions $\mathcal{R}_2$, $\mathcal{R}_3$, and $\mathcal{R}_4$ at time 100, 300, and 500, respectively. Under these anomalous conditions and $\sigma^2 = 1$, a sample evolution of the adaptive ensemble CUSUM algorithm is shown in Fig. 7(b). It can be seen that the adaptive ensemble algorithm samples a region with high likelihood of anomaly with high probability, and, hence, it improves upon the performance of the stationary policy.

We now study the expected detection delay under adaptive ensemble CUSUM algorithm and compare it with the efficient stationary policy. The anomaly at each region appears at time $50, 200, 350$ and 500, respectively. The expected detection delay obtained by Monte-Carlo simulations for $\sigma^2 = 1$ and different thresholds is shown in Fig. 8(a). It can be seen that the adaptive policy improves the detection delay significantly over the efficient stationary policy for large thresholds. It should be noted that the detection delay minimization is most needed at large thresholds because the detection delay is already low at small thresholds. Furthermore, frequent false alarms are encountered at low thresholds and hence, low thresholds are not typically chosen. The expected detection delay obtained by Monte-Carlo simulations for different value of σ^2 and threshold $\eta = 5$ is shown in Fig. 8(b). Note that for a given value of σ^2, the Kullback-Leibler divergence between

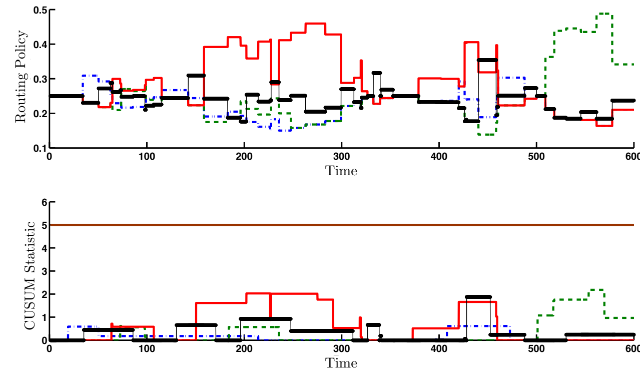

(a) CUSUM statistic and vehicle routing probabilities under nominal conditions

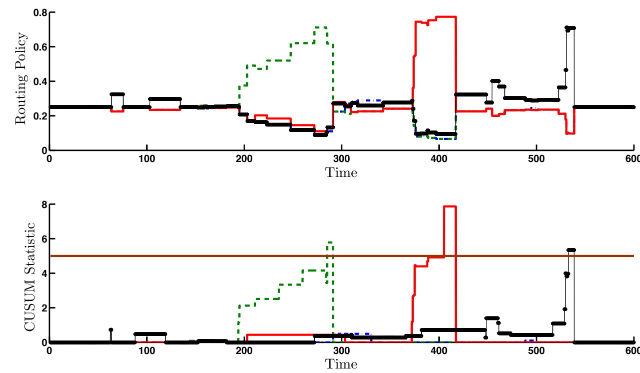

(b) CUSUM statistic and vehicle routing probabilities under anomalous conditions

Fig. 7. Sample evolution of the adaptive ensemble CUSUM algorithm. The dashed-dotted blue line, dashed green line, solid red line and solid black line with dots represent data from regions $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ and $\mathcal{R}_4$, respectively. The solid brown horizontal line represents the threshold. The vehicle routing probability is a function of the likelihood of anomaly at each region. As the likelihood of an anomaly being present at a region increases, also the probability to survey that region increases. Anomalies appear at region $\mathcal{R}_2$, $\mathcal{R}_3$ and $\mathcal{R}_4$ at times 100, 300 and 500, respectively. Once an anomaly is detected, it is removed and the statistic is reset to zero.

$\mathcal{N}(1, \sigma^2)$ and $\mathcal{N}(0, \sigma^2)$ is $1/2\sigma^2$. It can be seen that the adaptive policy improves the performance of the stationary policy for each value of noise. □

We now apply the adaptive ensemble CUSUM algorithm to a more general scenario where the anomalous distribution is not completely known. As remarked in Section II, in this case, the CUSUM algorithm should be replaced with the GLR algorithm. Given the nominal probability density function f_k^0 and the anomalous probability density function $f_k^1(\cdot|\theta)$ parameterized by $\theta \in \Theta \subseteq \mathbb{R}^\ell$, for some $\ell \in \mathbb{N}$, the GLR algorithm (Basseville and Nikiforov, 1993), works identically to the CUSUM algorithm, except that the CUSUM statistic is replaced by the statistic

$$\Lambda_\tau^k = \max_{t \in \{1,\dots,\tau\}} \sup_{\theta \in \Theta} \sum_{i=t}^{\tau} \log \frac{f_k^1(y_i|\theta)}{f_k^0(y_i)}.$$

*Example 6 (**Generalized Likelihood Ratio**):* For the same set of data as in Example 5, assume that there are three types of potential anomalies at each region. Since any combination of these anomalies can occur simultaneously, there are 7 potential distributions under anomalous conditions. We char-

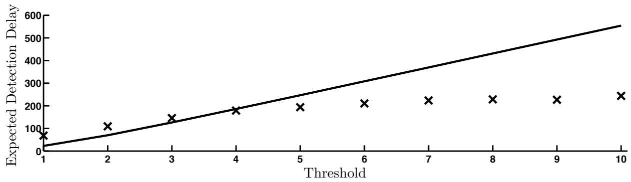

(a) Expected detection delay as a function of threshold

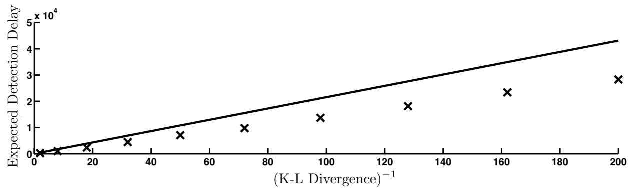

(b) Expected detection delay as a function of KL divergence

Fig. 8. Performance of the adaptive ensemble CUSUM algorithm. The solid black line represents the theoretical expected detection delay for the efficient stationary policy and the black $\times$ represent the expected detection delay for the adaptive ensemble CUSUM algorithm.

acterize these distributions as different hypothesis and assume that the observations under each hypothesis $h \in \{1, \dots, 8\}$ are sampled from a normal distribution with mean μ_h and covariances Σ_h. Let

$$\mu_1 = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \mu_2 = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \mu_3 = \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix}, \mu_4 = \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix},$$
$$\mu_5 = \begin{bmatrix} 1 \\ 1 \\ 0 \end{bmatrix}, \mu_6 = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \mu_7 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}, \mu_8 = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \text{ and}$$

$$\Sigma_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \Sigma_2 = \begin{bmatrix} 2 & 1 & 0 \\ 1 & \frac{3}{2} & 0 \\ 0 & 0 & 1 \end{bmatrix}, \Sigma_3 = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 1 & \frac{3}{2} \end{bmatrix}, \Sigma_4 = \begin{bmatrix} \frac{3}{2} & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 1 & 2 \end{bmatrix},$$
$$\Sigma_5 = \begin{bmatrix} 2 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 1 & 1 \end{bmatrix}, \Sigma_6 = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 2 & 1 \\ 0 & 1 & 2 \end{bmatrix}, \Sigma_7 = \begin{bmatrix} 2 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 2 \end{bmatrix}, \Sigma_8 = \begin{bmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{bmatrix}.$$

We picked region $\mathcal{R}_1$ as non-anomalous, while hypothesis 4, 6, and 8 were true at regions $\mathcal{R}_2, \mathcal{R}_3$, and $\mathcal{R}_4$, respectively. The Kullback-Leibler divergence at a region was chosen as the minimum of all possible Kullback-Leibler divergences at that region. A sample evolution of the adaptive ensemble CUSUM algorithm with GLR statistic replacing the CUSUM statistic is shown in Fig 9(a). It can be seen the performance is similar to the performance in Example 5. As an additional ramification of this algorithm, we also get the likelihood of each hypothesis at each region. It can be seen in Fig 9(b) that the true hypothesis at each region corresponds to the hypothesis with maximum likelihood. $\square$

VII. Experimental Results

We first detail our implementation of the algorithms using the Player/Stage robot control software package and the specifics of our robot hardware. We then present the results of the experiment.

Robot hardware: We use Erratic mobile robots from Videre Design shown in Fig. 10. The robot platform has a roughly square footprint (40cm $\times$ 37cm), with two differential drive wheels and a single rear caster. Each robot carries an onboard computer with a 1.8Ghz Core 2 Duo processor, 1 GB of memory, and 802.11g wireless communication. For navigation and localization, each robot is equipped with a Hokuyo URG-

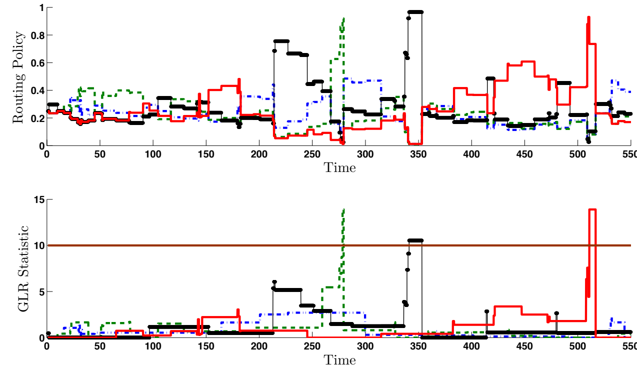

(a) GLR statistic under anomalous conditions

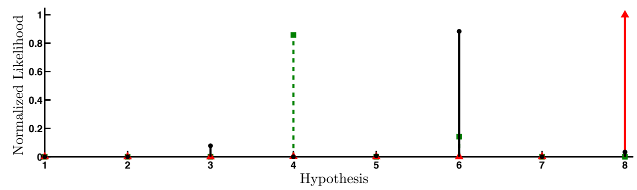

(b) Normalized likelihood of each hypothesis

Fig. 9. Sample evolution of the adaptive ensemble CUSUM algorithm with GLR statistic. The dashed-dotted blue line, dashed green line, solid red line and solid black line with dots represent data from regions $\mathcal{R}_1, \mathcal{R}_2, \mathcal{R}_3$ and $\mathcal{R}_4$, respectively. The solid brown horizontal line represents the threshold. The vehicle routing probability is a function of the likelihood of anomaly at each region. As the likelihood of an anomaly being present at a region increases, also the probability to survey that region increases. Anomalies appear at region $\mathcal{R}_2$, $\mathcal{R}_3$ and $\mathcal{R}_4$ at times 100, 300 and 500, respectively. Once an anomaly is detected, it is removed and the statistic is reset to zero. The true hypothesis at each region corresponds to the hypothesis with maximum likelihood

Fig. 10. Erratic mobile robot with URG-04LX laser rangefinder.

04LX laser rangefinder. The rangefinder scans 683 points over 240° at 10Hz with a range of 5.6 meters.

Localization: We use the *amcl* driver in Player which implements Adaptive Monte-Carlo Localization (Thrun et al., 2001). The physical robots are provided with a map of our lab with a 15cm resolution and told their starting pose within the map (Fig. 11). We set an initial pose standard deviation of 0.9m in position and 12° in orientation, and request updated localization based on 50 of the sensors range measurements for each change of 2cm in robot position or 2° in orientation. We use the most likely pose estimate by amcl as the location of the robot.

Navigation: Each robot uses the *snd* driver in Player for the Smooth Nearness Diagram navigation (Durham and Bullo, 2008). For the hardware, we set the robot radius parameter to 22cm, obstacle avoidance distance to 0.5m, and maximum speed to 0.2m/s. We let a robot achieve its target when it is within 10cm of the target.

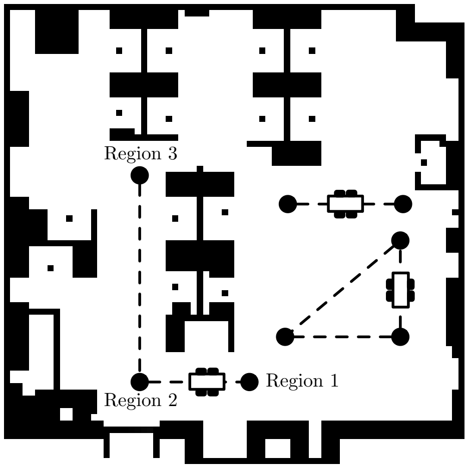

Fig. 11. This figure shows a map of our lab together with our surveillance configuration. Three erratic robots survey the selected 8 regions (black dots), which have been partitioned among the robots. Regions 1, 2, and 3 are also considered in Fig. 12, where we report the statistics of our detection algorithm.

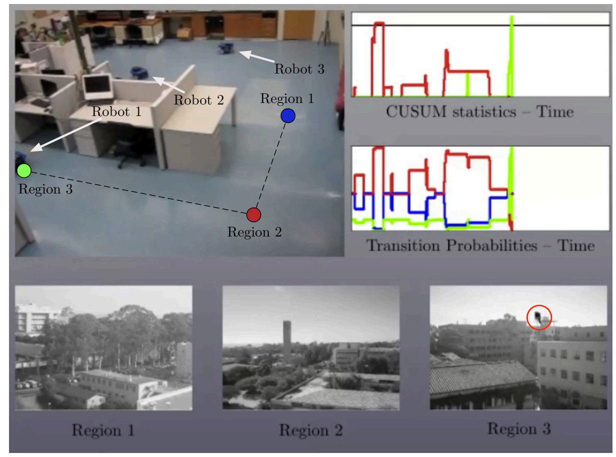

Fig. 12. A snapshot of our surveillance experiment, where three robots survey six locations in our lab (Fig. 11). In this figure we show the three regions assigned to the first robot. Each region correspond to a part of our campus, and observations are taken accordingly. Notice that Region 3 contains an anomaly (black smoke), and that the CUSUM statistics, which are updated upon collection of observations, reveal the anomaly (green peak). The transition probabilities are updated according to our adaptive ensemble CUSUM algorithm.

Experiment setup: For our experiment we employed our team of 3 Erratic robots to survey our laboratory. As in Fig. 11, a set of 8 important regions have been chosen and partitioned among the robots. Each robot surveys its assigned regions. In particular, each robot implements the single robot adaptive ensemble CUSUM algorithm in its regions. Notice that Robot 1 cannot travel from region 1 to region 3 in a single hop. Therefore, Robot 1 selects the regions according to a Markov chain with desired stationary distribution. This Markov chain was constructed using the Metropolis-Hastings algorithm. In particular, for a set of regions modeled as a graph $\mathcal{G} = (V, \mathcal{E})$, to achieve a desired stationary routing policy $\boldsymbol{q}$, the Metropolis-Hastings algorithm (Wasserman, 2004), picks the transition matrix P with entries:

$$P_{ij} = \begin{cases} 0, & \text{if } (i,j) \notin \mathcal{E}, \\ \min\left\{\frac{1}{d_i}, \frac{q_j}{q_i d_j}\right\} & \text{if } (i,j) \in \mathcal{E} \text{ and } i \neq j, \\ 1 - \sum_{k=1, k\neq i}^{n} P_{ik} & \text{if } (i,j) \in \mathcal{E} \text{ and } i = j, \end{cases}$$

where d_i is the number of regions that can be visited from region $\mathcal{R}_i$.

Observations (in the form of pictures) are collected by a robot each time a region is visited. In order to have a more realistic experiment, we map each location in our lab to a region in our campus. Then, each time a robot visit a region in our lab, a picture of a certain region in our campus is selected as observation (see Fig. 12). Pictures have been collected prior to the experiment.

Finally, in order to demonstrate the effectiveness of our anomaly detection algorithm, some pictures from regions 2 and 3 have been manually modified to contain an anomalous pattern; see Fig. 13. Anomalous pictures are collected by Robot 1 at some pre-specified time instants (the detection algorithm, however, does not make use of this information).

Probability density function estimation: In order to implement our adaptive ensemble CUSUM algorithm, the probability density functions of the observations at the regions in presence and absence of an anomaly need to be estimated.

Fig. 13. This figure shows sample pictures from Region 2 and Region 3, both with and without the anomaly to be detected.

For this task, we first collect sample images, and we register them in order to align their coordinates (Radke et al., 2005). We then select a reference image, and compute the difference between the sample pictures and the reference image. Then, we obtain a coarse representation of each difference image by dividing the image into blocks. For each difference image, we create a vector containing the mean value of each block, and we compute the mean and standard deviation of these vectors. Finally, we fit a normal distribution to represent the collected nominal data. In order to obtain a probability density distribution of the images with anomalies, we manually modify the nominal images, and we repeat the same procedure as in the nominal case.

Experiment results: The results of our experiment are illustrated in Fig. 12, Fig. 13, and in the multimedia extension available at `http://www.ijrr.org`. From the CUSUM statistics we note that the anomalies in Region 2 and Region

3 are both detected: indeed both the red curve and the green curve pass the decision threshold. We also note that few observations are necessary to detect the anomaly. Since the robots successfully survey the given environment despite sensor and modeling uncertainties due to real hardware, we conclude that our modeling assumptions in Section II are not restrictive.

VIII. Conclusions

In this paper we studied a spatial quickest detection problem in which multiple vehicles surveil a set of regions to detect anomalies in minimum time. We developed a novel ensemble CUSUM algorithm to detect an anomaly in any of the regions. A stochastic vehicle routing policy was adopted in which the vehicle samples the next region to visit from a probability vector. In particular, we studied (i) stationary policy: the probability vector is a constant function of time; and (ii) adaptive policy: the probability vector is adapted with time based on the collected observations. We designed an efficient stationary policy that depends on the travel time of the vehicles, the processing time required to collect information at each region, and the anomaly detection difficulty at each region. In adaptive policy, we modified the efficient stationary policy at each iteration to ensure that the regions with high likelihood of anomaly are visited with high probability, and thus, improved upon the performance of the stationary policy. We also mentioned the methods that extend the ideas in this paper immediately to the scenario in which the distributions of the observations in presence and absence of anomaly are not completely known, but belong to some parametrized family, or to the scenario in which the observations collected from each region are not independent (e.g., in the case of dynamic anomalies).

There are several possible extensions of the ideas considered here. First, in the case of dependent observations at each region, the current method assumes known distributions in presence and absence of anomalies. An interesting direction is to design quickest detection strategies that are robust to the uncertainties in these distributions. Second, the anomalies considered in this paper are always contained in the same region. It would be of interest to consider anomalies that can move from one region to another. Third, the policy presented in this paper considers an arbitrary partition that satisfy some cardinality constraints. It is of interest to come up with *smarter* partitioning policies that take into consideration the travel times, and the difficulty of detection at each region. Last, to construct the fastest mixing Markov chain with desired stationary distribution, we relied on time-homogeneous Markov chains. A time varying Markov chain may achieve a faster convergence to the desired stationary distribution (Grace and Baillieul, 2005). This is also an interesting direction to be pursued.

Appendix

A. Probabilistic guarantee to the uniqueness of critical point

We now provide probabilistic guarantee for Conjecture 5. The average detection delay for a single vehicle under a stationary policy $\boldsymbol{q}$ is

$$\delta_{\text{avg}}(\boldsymbol{q}) = \Big(\sum_{i=1}^{n} \frac{v_i}{q_i}\Big)\Big(\sum_{i=1}^{n} q_i \bar{T}_i + \sum_{i=1}^{n}\sum_{j=1}^{n} q_i q_j d_{ij}\Big),$$

where $v_i = w_i \bar{\eta}/\mathcal{D}_i$ for each $i \in \{1,\dots,n\}$. A local minimum of δ_{avg} can be can be found by substituting $q_n = 1 - \sum_{j=1}^{n-1} q_j$, and then running the gradient descent algorithm from some initial point $\boldsymbol{q}_0 \in \Delta_{n-1}$ on the resulting objective function.

Let $\boldsymbol{v} = (v_1,\dots,v_n)$ and $\boldsymbol{T} = (\bar{T}_1,\dots,\bar{T}_n)$. We assume that the parameters $\{\boldsymbol{v},\boldsymbol{T},D,n\}$ in a given instance of optimization problem (5) and the chosen initial point $\boldsymbol{q}_0$ are realizations of random variables sampled from some space $\mathcal{K}$. For a given realization $\kappa \in \mathcal{K}$, let the realized value of the parameters be $\{\boldsymbol{v}(\kappa),\boldsymbol{T}(\kappa),D(\kappa),n(\kappa)\}$, and the chosen initial point be $\boldsymbol{q}_0(\kappa)$. The associated optimization problem is:

$$\underset{\boldsymbol{q}\in\Delta_{n(\kappa)-1}}{\text{minimize}} \quad \delta_{\text{avg}}(\boldsymbol{q}\,|\,\kappa), \tag{A-1}$$

where, for a given realization $\kappa \in \mathcal{K}$, $\delta_{\text{avg}}(\cdot\,|\,\kappa) : \Delta_{n(\kappa)-1} \to \mathbb{R}_{>0} \cup \{+\infty\}$ is defined by

$$\delta_{\text{avg}}(\boldsymbol{q}\,|\,\kappa) = \Big(\sum_{i=1}^{n(\kappa)} \frac{v_i(\kappa)}{q_i}\Big)\Big(\sum_{i=1}^{n(\kappa)} q_i \bar{T}_i(\kappa) + \sum_{i=1}^{n(\kappa)}\sum_{j=1}^{n(\kappa)} q_i q_j d_{ij}(\kappa)\Big).$$

For a given realization κ, let $gd(\cdot\,|\,\kappa) : \Delta_{n(\kappa)-1} \to \Delta_{n(\kappa)-1}$ be the function that determines the outcome of the gradient descent algorithm applied to the function obtained by substituting $q_{n(\kappa)} = 1 - \sum_{j=1}^{n(\kappa)-1} q_j$ in $\delta_{\text{avg}}(\boldsymbol{q}\,|\,\kappa)$. In other words, the gradient descent algorithm starting from point $q_0(\kappa)$ converges to the point $gd(q_0(\kappa)\,|\,\kappa)$. Consider N_1 realizations $\{\kappa_1,\dots,\kappa_{N_1}\} \in \mathcal{K}^{N_1}$. Let $\boldsymbol{q}_{\text{optimal}}(\kappa) = gd(\frac{1}{n(\kappa)}\mathbf{1}_{n(\kappa)}\,|\,\kappa)$, and define

$$\hat{\gamma} = \max\{\|gd(\boldsymbol{q}_0(\kappa_s)\,|\,\kappa_s) - \boldsymbol{q}_{\text{optimal}}(\kappa_s)\| \mid s \in \{1,\dots,N_1\}\}.$$

It is known (Calafiore et al., 2011) that if $N_1 \geq -(\log \nu_1)/\mu_1$, for some $\mu_1, \nu_1 \in\,]0,1[$, then, with at least confidence $1-\nu_1$, it holds

$$\mathbb{P}(\{\boldsymbol{q}_0(\kappa) \in \Delta_{n(\kappa)-1} \mid \|gd(\boldsymbol{q}_0(\kappa)\,|\,\kappa) - \boldsymbol{q}_{\text{optimal}}(\kappa)\| \leq \hat{\gamma}\}) \geq 1 - \mu_1,$$

for any realization $\kappa \in \mathcal{K}$.

We sample the following quantities: the value n as uniformly distributed in $\{3,\dots,12\}$; each coordinate of the n regions in two dimensional space from the normal distribution with mean 0 and variance 100; the value T_i, for each $i \in \{1,\dots,n\}$, from the half normal distribution with mean 0 and variance 100; and the value v_i, for each $i \in \{1,\dots,n\}$, uniformly from $]0,1[$. For a realized value of n, we chose $\boldsymbol{q}_0$ uniformly in Δ_{n-1}. Let the matrix D be the Euclidean distance matrix between the n sampled regions.

We considered $N_1 = 1000$ realizations of the parameters $\{\boldsymbol{v},\boldsymbol{T},D,n\}$ and initial value $\boldsymbol{q}_0$. The sample sizes were determined for $\mu_1 = 0.01$ and $\nu_1 = 10^{-4}$. The value of $\hat{\gamma}$ obtained was 10^{-4}. Consequently, the gradient descent algorithm for the optimization problem (5) starting from any feasible point yields the same solution with high probability. In other words, with at least confidence level 99.99% and probability at least 99%, the optimization problem (5) has a unique critical point at which the minimum is achieved.

REFERENCES

M. Basseville and I. V. Nikiforov. *Detection of Abrupt Changes: Theory and Application*. Prentice Hall, 1993.

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.

S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. ISBN 0521833787.

S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, 2004.

F. Bullo, E. Frazzoli, M. Pavone, K. Savla, and S. L. Smith. Dynamic vehicle routing for robotic systems. *Proceedings of the IEEE*, 99(9):1482–1504, 2011.

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

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

B. Chen and P. Willett. Detection of hidden Markov model transient signals. *IEEE Transactions on Aerospace and Electronic Systems*, 36(4):1253–1268, 2000.

Y. Chevaleyre. Theoretical analysis of the multi-agent patrolling problem. In *IEEE/WIC/ACM Int. Conf. on Intelligent Agent Technology*, pages 302–308, Beijing, China, September 2004.

T. H. Chung and J. W. Burdick. Analysis of search decision making using probabilistic search strategies. *IEEE Transactions on Robotics*, 28(1):132–144, 2012.

T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley, 1991.

J. W. Durham and F. Bullo. Smooth nearness-diagram navigation. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 690–695, Nice, France, September 2008.

Y. Elmaliach, A. Shiloni, and G. A. Kaminka. A realistic model of frequency-based multi-robot polyline patrolling. In *International Conference on Autonomous Agents*, pages 63–70, Estoril, Portugal, May 2008.

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.

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.

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.

G. A. Hollinger, U. Mitra, and G. S. Sukhatme. Active classification: Theory and application to underwater inspection. In *International Symposium on Robotics Research*, Flagstaff, AZ, USA, August 2011a.

G. A. Hollinger, U. Mitra, and G. S. Sukhatme. Autonomous data collection from underwater sensor networks using acoustic communication. In *IEEE/RSJ Int. Conf. on Intelligent Robots & Systems*, pages 3564–3570, San Francisco, CA, USA, September 2011b.

D. B. Kingston, R. W. Beard, and R. S. Holt. Decentralized perimeter surveillance using a team of UAVs. *IEEE Transactions on Robotics*, 24 (6):1394–1404, 2008.

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.

S. M. LaValle. *Planning Algorithms*. Cambridge University Press, 2006. ISBN 0521862051. Available at `http://planning.cs.uiuc.edu`.

F. Pasqualetti, J. W. Durham, and F. Bullo. Cooperative patrolling via weighted tours: Performance analysis and distributed algorithms. *IEEE Transactions on Robotics*, 28(5):1181–1188, 2012a.

F. Pasqualetti, A. Franchi, and F. Bullo. On cooperative patrolling: Optimal trajectories, complexity analysis and approximation algorithms. *IEEE Transactions on Robotics*, 28(3):592–606, 2012b.

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

R. J. Radke, S. Andra, O. Al-Kofahi, and B. Roysam. Image change detection algorithms: A systematic survey. *IEEE Transactions on Image Processing*, 14(3):294–307, 2005.

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

S. I. Resnick. *A Probability Path*. Birkhäuser, 1999.

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.

D. Siegmund. *Sequential Analysis: Tests and Confidence Intervals*. Springer, 1985. ISBN 978-0-387-96134-7.

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.

S. L. Smith, M. Schwager, and D. Rus. Persistent robotic tasks: Monitoring and sweeping in changing environments. *IEEE Transactions on Robotics*, 28(2):410–426, 2012.

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

V. Srivastava and F. Bullo. Stochastic surveillance strategies for spatial quickest detection. In *IEEE Conf. on Decision and Control and European Control Conference*, pages 83–88, Orlando, FL, USA, December 2011.

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

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*, pages 6284–6289, Orlando, FL, USA, December 2011b.

S. Thrun, D. Fox, W. Burgard, and F. Dellaert. Robust Monte Carlo localization for mobile robots. *Artificial Intelligence*, 128(1-2):99–141, 2001.

J. Unnikrishnan, V. V. Veeravalli, and S. P. Meyn. Minimax robust quickest change detection. *IEEE Transactions on Information Theory*, 57(3):1604–1614, 2011.

L. Wasserman. *All of Statistics: A Concise Course in Statistical Inference*. Springer, 2004.

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.