Quickest Detection over Robotic Roadmaps

Pushkarini Agharkar, Francesco Bullo

Abstract—**We study the problem of quickest detection of anomalies in an environment under extreme uncertainties in sensor measurements. The robotic roadmap corresponding to the environment can be represented as a graph with an arbitrary topology. We analyze the Ensemble CUSUM Algorithm for this surveillance problem. We quantify the delay in detection of anomalies using the Ensemble CUSUM Algorithm and also frame an optimization problem to minimize this detection delay. We then provide an upper bound on the optimal detection delay and frame a convex optimization problem to minimize this upper bound. We also propose an efficient policy which achieves this upper bound and which can be computed by solving a semidefinite program. We illustrate the efficacy of the Ensemble CUSUM Algorithm using numerical simulations. We observe that the efficient policy outperforms policies based on other well-known Markov chains. This trend is more noticeable for higher levels of uncertainties and noise in sensor measurements.**

I. Introduction

The topic of surveillance in environments using autonomous agents has received considerable attention lately. Specific examples include the monitoring of chemical and oil spills [5], the detection of forest fires [9], search and rescue missions and topological mapping [21], the periodic patrolling of an environment [6], [20]. Other applications include building and infrastructure maintenance [10], minimizing emergency vehicle response times [2] and robotic warehouse management [23]. In this paper, we consider surveillance strategies under extreme uncertainties in sensor measurements.

A. Related Work

Theoretical analysis of the surveillance problem was conducted in [4] and a survey of various surveillance scenarios and the corresponding approaches was presented in [1]. Surveillance strategies that minimize the refresh time, i.e., time period between subsequent visits to regions have been proposed in [12],[16] and [17]. In [12], authors propose optimal algorithms which minimize the refresh time for chain and tree graphs and constant factor algorithm for cyclic graphs. Authors in [17] consider the problem of minimizing specific weighted sums of refresh times and design non-intersecting tours on graphs for this surveillance criterion. In [16], the authors design speed controllers on closed paths to minimize the refresh time for a given set of points of interest in the environment.

The surveillance policies proposed in [12],[16] and [17] are deterministic in nature. Stochastic surveillance strategies assume importance in scenarios where the intruders can move or hide to avoid detection and as a result, the movement of the surveillance vehicle is required to be non-deterministic. The computation of an optimal deterministic strategy is also tedious. This is because the length of the tour which is repeatedly executed in the course of the deterministic strategy may be arbitrarily long. This is important especially in large environments and in situations where the strategy has to be re-evaluated to incorporate new information. We hence devote our attention to stochastic strategies.

Several authors have used Markov chain based approaches to design stochastic strategies for various surveillance tasks. Authors in [18] use the Metropolis-Hastings algorithm to achieve specified frequency of visits to regions of the environment. In [7], authors design random walk strategies on hypergraphs and parametrically vary the local transition probabilities over time in order to achieve fast convergence to a desired visit frequency distribution. In [19], authors use the fastest mixing Markov chain for quickest detection of anomalies. Authors in [13] use a Markov chain with minimum mean first passage time in order to detect intruders in unknown locations in the environment. Authors in [14] consider different intruder models and present routing strategies for surveillance in scenarios corresponding to them.

The framework for the surveillance problem studied in this paper was introduced in [19]. In the setup of this problem, the surveillance vehicle conducts surveillance of an environment which can be represented as a graph. It takes observations from different regions of the environment in order to detect the presence of anomalies in the regions as quickly as possible after their occurrence. The authors in [19] proposed a Markov chain based routing policy termed the *Ensemble CUSUM Algorithm* to determine the movement of the surveillance vehicle across regions. The Markov chains that they considered were required to have transition matrices with identical columns.

B. Contributions

We revisit the surveillance problem studied in [19]. The authors in [19] considered an all to all graph topology more suitable for environments with aerial vehicles. We extend their setup to graphs with arbitrary topologies, which we broadly refer to as robotic roadmaps. Further, we only keep the assumption of irreducibility on the Markov chains and look at a wider class of Markov chains than the one considered in [19]. We determine an expression for the average detection delay in the generalized setting and find that it depends on the first passage times of the Markov chain corresponding to the routing policy. We then frame an optimization problem to find the Markov chain corresponding to the *optimal policy* which minimizes the detection delay. We also provide an upper bound on the minimum detection delay and frame an optimization

This work was supported in part by AFOSR grant FA9550-15-1-0138.

[1] Pushkarini Agharkar and Francesco Bullo are with the Mechanical Engineering Department and the Center for Control, Dynamical Systems and Computation, University of California at Santa Barbara, CA 93106-5070, USA
agharkar@umail.ucsb.edu, bullo@enigneering.ucsb.edu

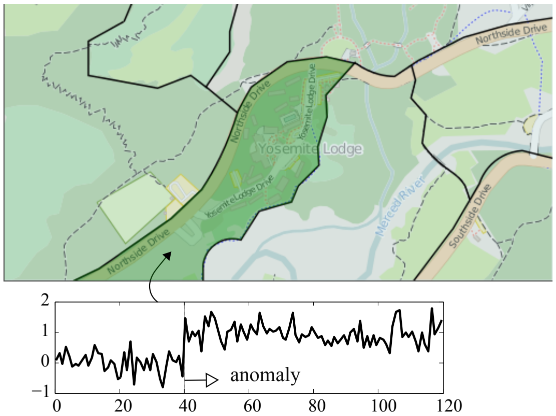

Fig. 1. The environment is an area separated into seven regions of interest. Observations made in the highlighted region change after an anomaly occurs. The aim of the surveillance vehicle is to detect this change as soon as possible.

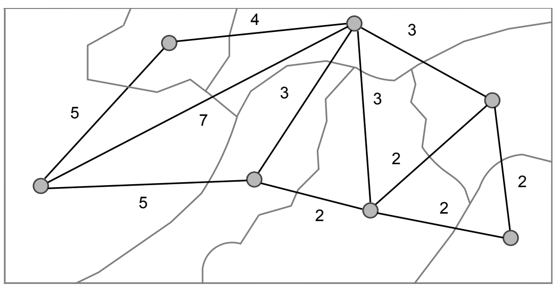

Fig. 2. The robotic roadmap corresponding to the environment can be represented by a graph. The edge weights of the graph represent travel times between neighboring regions.

problem to minimize the upper bound. We prove that the upper bound optimization problem is convex and provide a semidefinite program (SDP) formulation to solve it and obtain the corresponding *efficient policy*. Using an illustrative example, we validate our expression for the detection delay and also surmise that the efficient policy provides a detection delay close to that of the optimal policy.

C. Organization

In Section II, we describe the setup of the surveillance problem, formally define the *quickest detection task* and state the *Ensemble CUSUM Algorithm* to address the task. In Section III, we review results on the CUSUM algorithm and the mean first passage time of Markov chain random walks on graphs which will be used to analyze the Ensemble CUSUM Algorithm. In Section IV, we analyze the performance of the Ensemble CUSUM Algorithm and provide an upper bound on its performance. In Section V, we present numerical simulations which validate our findings. Conclusions of the paper are summarized in Sec. VI

II. Problem Setup

We first describe our model for the environment and the mathematical model used for simulating the presence of anomalies in the environment.

A. Environment

The environment in the problem setup of this paper can be modeled as a graph $G = (V, E)$ with node set $V := \{1, \ldots, n\}$ and edge set $E \subseteq V \times V$. The nodes in the graph correspond to the regions in the environment and the edges correspond to the interconnections between them. The time taken to travel from region i to the neighboring region j is d_{ij} and travel time matrix $D = [d_{ij}] \in \mathbb{R}^{n\times n}$ with the property that $d_{ij} \geq 0$ if $(i, j) \in E$ and $d_{ij} = 0$ otherwise.

The level of importance w_i is assigned to region i and $w = [w_i] \in \mathbb{R}^{n\times 1}$ is the referred to as the *priority vector*. Without loss of generality, $w^T \mathbb{1}_n = 1$. The environment can thus be described by the 4-tuple: $\mathcal{E} = \langle V, E, D, w\rangle$. An example of the environment and the graph corresponding to it is shown in Fig. 1 and Fig. 2 respectively.

B. Observations in Environment

When the surveillance vehicle visits a region in the environment, it makes an observation about the region. Based on all the observations made in the region up to that point, it predicts the presence of anomalies in the region.

Let the set of observations made by the surveillance vehicle at the region k be $\{y_{k,1}, y_{k,2}, \ldots\}$. If an anomaly occurs in the region at some iteration ν, then the observations $\{y_{k,1}, \ldots, y_{k,\nu}\}$ are i.i.d. with probability density function f_k^0 and the observations $\{y_{k,\nu}, y_{k,\nu+1}, \ldots\}$ are i.i.d. with probability density function f_k^1. We use the notation $\mathcal{D}(f_k^1, f_k^0)$ to denote the Kullback-Leibler divergence of f_k^0 from f_k^1 and also denote $\mathcal{D}_k := \mathcal{D}(f_k^1, f_k^0)$ for convenience. We now describe the *spatial quickest detection task* and quantify it.

C. Quickest Detection of Anomalies

The surveillance vehicle adopts a policy described by the tuple $\mathcal{P} = \langle P, q\rangle$. It moves in the environment $\mathcal{E} = \langle V, E, D, w\rangle$ according to a Markov chain with stationary distribution $q = [q_i] \in \mathbb{R}^{n\times 1}$ and transition matrix $P = [p_{ij}] \in \mathbb{R}^{n\times n}$. If $(i, j) \in E$, then $p_{ij} \geq 0$ and $p_{ij} = 0$ otherwise.

The aim of the vehicle is to detect anomalies in a region based on observations made in that region in least amount of time possible. More specifically, using a routing policy $\mathcal{P}$, it is required to minimize the average detection delay defined below.

Definition 1 (Average Detection Delay): Let the vehicle service the environment $\mathcal{E} = \langle V, E, D, w\rangle$ using policy $\mathcal{P}$ for $k \in \{1, \ldots, n\}$ and let $\delta_k(\mathcal{P})$ be the delay in detecting an anomaly at region k. Then the task of the vehicle is to minimize the *average detection delay* $\delta_{\text{avg}}(\mathcal{P})$ given by

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

D. Ensemble CUSUM Algorithm

The surveillance vehicle visits the regions in $\mathcal{E}$ according to a realization of the Markov chain with stationary distribution q and transition matrix P. When the vehicle is in a particular region of the environment, it runs a local version

Algorithm 1: Ensemble CUSUM Algorithm

Given: Policy $\mathcal{P} = \langle P, q \rangle$, threshold η, initial state x.
Set: $\Lambda_{k,0} = 0$ for $k \in \{1, \ldots, n\}$, local variable $\tau = 1$ for all regions.

1 Make observation $y_{k,\tau}$ at region x;
2 $\Lambda_{x,\tau} = \sup\left(\Lambda_{x,\tau-1} + \log \frac{f_x^1(y_{x,\tau})}{f_x^0(y_{x,\tau})}, 0\right)$;
3 **if** $\Lambda_{x,\tau} > \eta$ **then**
4 Declare an anomaly at region x;
5 Set $\Lambda_{x,\tau} = 0$;
6 **end**
7 Set $\tau \leftarrow \tau + 1$ for x;
8 Select $x \leftarrow z$ with probability $P(x, z)$;
9 Repeat from step 1.

of the *CUSUM algorithm*. We refer to the n parallel CUSUM algorithms by *Ensemble CUSUM Algorithm* (Algorithm 1). We wish to find the surveillance policy $\mathcal{P} = \langle P, q \rangle$ for the environment $\mathcal{E} = \langle V, E, D, w \rangle$ which minimizes the average detection delay $\delta_{\text{avg}}(\mathcal{P})$ defined in equation (1) in the previous section.

Remark 2: (Service times): The service times required for conducting surveillance in different regions are not modeled in the problem setup. However, they can be incorporated in a straightforward manner. If $v \in \mathbb{R}^{n \times 1}$ is the constant vector of service times, they can be accounted for by modifying the travel time matrix to $\bar{D} := D + \mathbb{1}_n v^T$.

Remark 3: (Knowledge of probability density functions): The probability density functions in the absence and presence of anomalies are assumed to be known to the surveillance vehicle. In a scenario where the probability density functions are not known, the CUSUM algorithm can be replaced by the minimax robust quickest change detection algorithm [22] and the results presented in this paper can be extended to apply to that scenario as well.

III. Preliminaries

We will now state some preliminary results which will be used in analysing the Ensemble CUSUM Algorithm. We will start by reviewing some performance guarantees on the CUSUM algorithm.

A. CUSUM Algorithm

The CUSUM algorithm is designed for quick prediction of anomalies while at the same time, avoiding making *false alarms* [15], [11]. In the CUSUM algorithm, at each iteration $\tau \in \mathbb{N}$ made in region k, (i) observation $y_{k,\tau}$ is collected, (ii) the statistic

$$\Lambda_{k,\tau} = \sup\left(\Lambda_{k,\tau-1} + \log \frac{f_k^1(y_{k,\tau})}{f_k^0(y_{k,\tau})}, 0\right)$$

with $\Lambda_{k,0} = 0$ is computed and (iii) a change is declared if $\Lambda_{k,\tau} > \eta$. Let O_k be the observation at which an anomaly is declared at region k. For a given threshold η, the expectation of O_k conditioned on the presence of an anomaly, i.e., the worst expected number of observations of the CUSUM algorithm [11] is

$$\mathbb{E}_{f_k^1}(O_k) \approx \frac{e^{-\eta} + \eta - 1}{\mathcal{D}(f_k^1, f_k^0)} = \frac{\bar{\eta}}{\mathcal{D}_k}, \tag{2}$$

where $\bar{\eta} = e^{-\eta} + \eta - 1$, and the expectation of O_k conditioned on the absence of an anomaly, i.e., the false alarm rate for CUSUM algorithm is

$$\mathbb{E}_{f_k^0}(O_k) \approx \frac{e^{\eta} - \eta - 1}{\mathcal{D}(f_k^0, f_k^1)} \tag{3}$$

The approximations in equations (2,3) are referred to as the Walds approximations [15]. For large values of the threshold η, these approximations are known to be accurate. We also set

$$s_k := \frac{\bar{\eta}}{\mathcal{D}_k}, \tag{4}$$

and $s = [s_k] \in \mathbb{R}^{n \times 1}$, referring to it as the *vector of CUSUM samples*. Given $\bar{\eta}$ and $\mathcal{D}_k$, the constant s_k is the expected number of visits to region k required to detect an anomaly in that region.

The expression for the average detection delay of the Ensemble CUSUM Algorithm (Algorithm 1) depends on the property of Markov chains called the first passage time. So we will define the first passage times of a graph and review some of its properties.

B. First Passage Time

Simply put, the first passage time from node i to node j of a graph is the expectation of the time T_{ij} taken for a Markov chain to start from i and visit j for the first time. This is taking into account the time taken to traverse edges in the graph specified by the travel time matrix D of the graph. Given a realization $X_1, X_2, X_3, \ldots$ of the Markov chain, the mathematical definition of the random variable T_{ij} is as

$$T_{ij} = \min\Bigg\{\sum_{n=0}^{k-1} d_{X_n, X_{n+1}}, \text{ for } k \geq 1 \mid X_k = j \\ \text{given that } X_0 = i\Bigg\}.$$

The expectation of T_{ij}, i.e., $\mathbb{E}[T_{ij}] =: n_{ij}$ is the first passage time from node i to node j. The matrix $N = [n_{ij}]$ is referred to as the first passage time matrix. The following results provide the governing equation for N, as well as a special weighted sum of its entries. The proofs for these results can be found in [13].

Theorem 4: (Governing equations for first passage times): For an irreducible Markov chain with transition matrix P, stationary distribution q and travel time matrix D,

(i) The first passage time matrix $N = [n_{ij}]$ satisfies the following equation:

$$n_{ij} = p_{ij} d_{ij} + \sum_{k \neq j} p_{ik}(n_{kj} + d_{ik}),$$

or in matrix form,

$$(I - P)N = (P \circ D)\mathbb{1}_n \mathbb{1}_n^T - P N_{\text{d}}.$$

(ii) The diagonal elements of N satisfy the following equation:

$$n_{ii} = q^T (P \circ D)\mathbb{1}_n q_i.$$

where $P \circ D$ is the element-wise product between P and D.

A specific weighted sum of the first passage time matrix N, called the *mean first passage time* is now defined. It is also referred to as the *weighted Kemeny constant* in [13].

Theorem 5: (Mean first passage time): The mean first passage time of an irreducible Markov chain with transition matrix P, stationary distribution q, travel time matrix D and first passage time matrix N is

$$q^T N q = q^T (P \circ D)\mathbb{1}_n \left(\sum_{i=2}^{n} \frac{1}{1-\lambda_i(P)} \right), \qquad (5)$$

where $\{\lambda_1(P), \ldots, \lambda_n(P)\}$ are the eigenvalues of P with $\lambda_1(P) = 1$.

IV. Performance of the Ensemble CUSUM Algorithm

We are now ready to state our main results on the average detection delay $\delta_{\text{avg}}(\mathcal{P})$ of the Ensemble CUSUM Algorithm (Algorithm 1).

Theorem 6: (Performance of the Ensemble CUSUM Algorithm): For a single vehicle conducting surveillance of the environment $\mathcal{E} = \langle V, E, D, w\rangle$ according to the Ensemble CUSUM Algorithm (Algorithm 1) using the policy $\mathcal{P} = \langle P, q\rangle$,

(i) the expected detection delay $\mathbb{E}[\delta_k(\mathcal{P})]$ at region k satisfies

$$\mathbb{E}[\delta_k(\mathcal{P})] = \sum_{i=1}^{n} q_i n_{ik} + (s_k - 1)\, n_{kk}, \qquad (6)$$

(ii) the average detection delay $\delta_{\text{avg}}(\mathcal{P})$ over the entire environment satisfies

$$\delta_{\text{avg}}(\mathcal{P}) = \sum_{k=1}^{n} w_k \left(\sum_{i=1}^{n} q_i n_{ik} + (s_k - 1)\, n_{kk} \right), \qquad (7)$$

where $N = [n_{ij}] \in \mathbb{R}^{n\times n}$ is the first passage time matrix for the irreducible Markov chain with transition matrix $P \in \mathbb{R}^{n\times n}$ and stationary distribution $q \in \mathbb{R}^{n\times 1}$ and the constant $s \in \mathbb{R}^{n\times 1}$ is the vector of CUSUM samples.

Proof: Let $\tau \in \{1, \ldots, O_k\}$ be the iterations at which the vehicle visits region region k and sends information about it to the control center. Let O_k be the iteration at which an anomaly is detected in region k. The observation made at region k at the τ-th iteration in that region is denoted by $y_{k,\tau}$. Let the log likelihood ratio calculated by the local CUSUM algorithm for that iteration be $\epsilon_{k,\tau}$. Then,

$$\epsilon_{k,\tau} = \log \frac{f_k^1(y_{k,\tau})}{f_k^0(y_{k,\tau})}.$$

Conditioned on the presence of an anomaly, $\{\epsilon_{k,\tau}\}_{\tau_k\in\mathbb{N}}$ are i.i.d. and $\mathbb{E}_{f_k^1}[\epsilon_{k,\tau}] = \mathcal{D}_k$. Then, referring to result summarized in equation (2), $\mathbb{E}_{f_k^1}[O_k] = \bar{\eta}/\mathcal{D}_k$. Thus, the expected time it takes for the Ensemble CUSUM Algorithm to make the O_k-th observation at region k is essentially the expected detection delay $\delta_k(\mathcal{P})$ at region k. We will now devote our attention to computing the expectation of $\delta_k(\mathcal{P})$.

Let t_k^0 be the time at which the vehicle starts the CUSUM algorithm. Let $\{t_k^1, t_k^2, \ldots, t_k^{O_k}\}$ be the time instant at which it leaves region k, having serviced it, and $\Delta t_k^i = t_k^{i+1} - t_k^i$ for $i = \{0, 1, 2, \ldots\}$. Then, the detection delay $\delta_k(P) = t_k^{O_k} = \sum_{i=0}^{O_k} \Delta t_k^i$. The expectation of $\delta_k(P)$ can be computed:

$$\begin{aligned}
\mathbb{E}[\delta_k(P)] &= \mathbb{E}\left[\sum_{i=0}^{O_k-1} \Delta t_k^i\right] \\
&= \mathbb{E}[\Delta t_k^0] + \mathbb{E}\left[\sum_{i=1}^{O_k-1} \Delta t_k^i\right] \\
&= \mathbb{E}[\Delta t_k^0] + (\mathbb{E}[O_k] - 1)\mathbb{E}[\Delta t_k^i] \qquad (8)\\
&= \sum_{i=1}^{n} q_i n_{ik} + \left(\frac{\bar{\eta}}{\mathcal{D}_k} - 1\right) n_{kk}. \qquad (9)
\end{aligned}$$

Equation (8) comes from the application of Wald's identity. Notice that $\mathbb{E}[\Delta t_k^0]$ is the expected time to start from any node and visit node k for the first time and given $i > 0$, $\mathbb{E}[\Delta t_k^i]$ is the expected time taken to return to node k. Recollect that n_{ij} is the expected time for the vehicle to start from node i to visits node j for the first time. Hence, we can conclude that $\mathbb{E}[\Delta t_k^0] = \sum_{i=1}^{n} q_i n_{ik}$ and $\mathbb{E}[\Delta t_k^i] = n_{kk}$ for $i > 0$ to obtain equation (9). Using the definition of s from equation (4), the first result follows. Next, using the definition of $\delta_{\text{avg}}(\mathcal{P})$ from equation (1), the second result follows. ■

Thus, the average detection delay depends on the first passage times between nodes of the graph representing the environment $\mathcal{E}$. We now present a modified expression for $\delta_{\text{avg}}(\mathcal{P})$, removing the dependence of the first passage times, in the following theorem. The proof of the theorem is postponed to Appendix B.

Theorem 7: (Average detection delay): For a single vehicle conducting surveillance of the environment $\mathcal{E} = \langle V, E, D, w\rangle$ according to the Ensemble CUSUM Algorithm (Algorithm 1) using the policy $\mathcal{P} = \langle P, q\rangle$,

$$\begin{aligned}
\delta_{\text{avg}}(\mathcal{P}) = \beta \mathbb{1}_n^T [((I-P) + (P\circ D)\mathbb{1}_n q^T)^{-1} \circ I](r\cdot w) \\
+ (\beta - 1) + \beta(s - \mathbb{1}_n)^T (r \cdot w), \quad (10)
\end{aligned}$$

where $r \in \mathbb{R}^{n\times 1}$ with $r \cdot q = \mathbb{1}_n$, $I \in \mathbb{R}^{n\times n}$ is the identity matrix, $\beta = q^T (P \circ D)\mathbb{1}_n$ and the constant $s \in \mathbb{R}^{n\times 1}$ is the vector of CUSUM samples.

In our setup, the environment can have an arbitrary graph topology and the routing policy can also take an arbitrary form adhering to the restrictions imposed by the graph topology. A specific simplification, where the environment is an all to all graph and where the transition matrix for the routing policy has the form $P = \mathbb{1}_n q^T$ is explored in [19]. While they provide algorithms to optimize the stationary distribution q in the simplified setup, we consider the more generalized problem. Specifically, our goal is to find policy $\mathcal{P} = \langle P, q\rangle$ for the vehicle such that $\delta_{\text{avg}}(\mathcal{P})$ is minimized. This can be framed as the following optimization problem:

Problem 1: (Minimizing the average detection delay): Given the environment $\mathcal{E} = \langle V, E, D, w\rangle$ and the constant vector of CUSUM samples $s \in \mathbb{R}^{n\times 1}$, determine the stationary distribution $q = [q_i] \in \mathbb{R}^{n\times 1}$ and transition probabilities $P = [p_{ij}] \in \mathbb{R}^{n\times n}$ solving:

$$\begin{aligned}
\text{minimize} \quad & \beta \mathbb{1}_n^T[((I-P)+(P\circ D)\mathbb{1}_n q^T s)^{-1}\circ I](r\cdot w) \\
& +(\beta-1)+\beta(s-\mathbb{1}_n)^T(r\cdot w) \\
\text{subject to} \quad & P\mathbb{1}_n=\mathbb{1}_n, \\
& 0\le p_{ij}\le 1, \text{ for each } (i,j)\in E \\
& p_{ij}=0, \text{ for each } (i,j)\notin E \\
& q^T P = q^T, \text{ for each } (i,j)\in E \\
& q^T\mathbb{1}_n = 1, \\
& q_i\ge 0,\ r_i=1/q_i, \text{ for each } i\in\{1,\dots,n\} \\
& P \text{ irreducible}, \\
& \beta = q^T(P\circ D)\mathbb{1}_n.
\end{aligned} \tag{11}$$

The above optimization problem contains the constraint that the transition matrix be irreducible. Since it is hard to enforce the irreducibility constraint during each step of an iterative optimization algorithm, our approach is to relax the irreducibility constraint and verify that the final solution satisfies the constraint. A Markov chain that is not irreducible contains multiple communicating classes, making the first passage time between at least one pair of regions infinite. Since the average detection delay depends on the first passage times of the chain, the outcome where the final solution is a reducible chain would drive up the cost function of the optimization problem, making such an outcome highly unlikely. Because of this reason, the relaxation of the irreducibility constraint works very well in practice.

Let $\mathcal{P}^*$ be the solution to Problem 1 and let $\delta^*_{\text{avg}} := \delta_{\text{avg}}(\mathcal{P}^*)$. The cost function of this optimization problem is not a convex function of P and q. Moreover, one of the constraints is also nonlinear. We now devote some attention to determining an upper bound on δ^*_{avg}, and frame an optimization problem to minimize it. We start with evaluating policies of the form $\mathcal{P}_w = \langle P_w, w\rangle$, i.e., where the Markov chain corresponding to the policy has stationary distribution equal to the priority vector w. We leverage the result known on the weighted sum of the first passage times from Theorem 5 to simplify expressions for the detection delay of the Ensemble CUSUM algorithm in this case.

Theorem 8: (Upper bound on average detection delay): For a single vehicle conducting surveillance of the environment $\mathcal{E} = \langle V, E, D, w\rangle$ according to the Ensemble CUSUM Algorithm (Algorithm 1) using the policy $\mathcal{P}_w = \langle P_w, w\rangle$,

$$\delta_{\text{avg}}(\mathcal{P}_w) = \left(w^T(P_w\circ D)\mathbb{1}_n\right)\left(\sum_{i=2}^n \frac{1}{1-\lambda_i(P_w)} + (s-\mathbb{1}_n)^T\mathbb{1}_n\right)$$

for all $P_w \in S_w$, where S_w is the set of transition matrices corresponding to irreducible Markov chains with stationary distribution w, $\{\lambda_1(P_w),\dots,\lambda_n(P_w)\}$ are the eigenvalues of P_w with $\lambda_1(P_w)=1$, and the constant $s\in\mathbb{R}^{n\times 1}$ is the vector of CUSUM samples.

Proof: We start with the expression for $\delta_{\text{avg}}(\mathcal{P})$ obtained in Theorem 6. In matrix form, equation (7) can be rewritten as

$$\delta_{\text{avg}}(\mathcal{P}) = q^T N w + (s-\mathbb{1}_n)^T N_d w, \tag{12}$$

where $\mathcal{P} = \langle P, q\rangle$. Setting the variable q to w, and using the result from Theorem 5, as well as the result from Theorem. 4 (ii), equation (12) can be simplified:

$$\begin{aligned}
&\delta_{\text{avg}}(\mathcal{P}_w) \\
&\quad = w^T N w + (s-\mathbb{1}_n)^T N_d w \\
&\quad = \left(w^T(P_w\circ D)\mathbb{1}_n\right)\left(\sum_{i=2}^n \frac{1}{1-\lambda_i(P_w)} + (s-\mathbb{1}_n)^T\mathbb{1}_n\right).
\end{aligned}$$

■

We can make the upper bound obtained on the optimal detection delay tighter by choosing $P_{\text{ub}} \in S_w$ which minimizes the average detection delay. The following optimization problem can be framed to find the matrix P_{ub}.

Problem 2: (Minimizing the upper bound on optimal average detection delay): Given the environment $\mathcal{E} = \langle V, E, D, w\rangle$, the vector of CUSUM samples $s\in\mathbb{R}^{n\times 1}$ and the stationary distribution $w\in\mathbb{R}^{n\times 1}$, determine the transition probabilities $P = [p_{ij}]\in\mathbb{R}^{n\times n}$ solving:

$$\begin{aligned}
\text{minimize} \quad & \left(w^T(P\circ D)\mathbb{1}_n\right)\left(\sum_{i=2}^n \frac{1}{1-\lambda_i(P)} + (s-\mathbb{1}_n)^T\mathbb{1}_n\right) \\
\text{subject to} \quad & P\mathbb{1}_n=\mathbb{1}_n, \\
& 0\le p_{ij}\le 1, \text{ for each } (i,j)\in E \\
& p_{ij}=0, \text{ for each } (i,j)\notin E \\
& w_i p_{ij} = w_j p_{ji}, \text{ for each } (i,j)\in E.
\end{aligned} \tag{13}$$

Note that the above optimization problem also involves the restriction of non-reversibility on the transition matrix P as denoted by the last equality.

Theorem 9: (Convexity of Optimization Problem 2): Let S_w be the set of transition matrices associated with irreducible non-reversible Markov chains on graph $G = (V, E)$ and having the stationary distribution w. Then, the Optimization Problem 2 is convex.

Proof: From Theorem 8, the cost function $f(P_w)$ of the Optimization Problem 2 can be written down as:

$$\begin{aligned}
f(P_w) = &\left(w^T(P_w\circ D)\mathbb{1}_n\right)\left(\sum_{i=2}^n \frac{1}{1-\lambda_i(P_w)}\right) \\
&+ \left(w^T(P_w\circ D)\mathbb{1}_n\right)(s-\mathbb{1}_n)^T\mathbb{1}_n.
\end{aligned} \tag{14}$$

The first term in equation (14) is the mean first passage time of the Markov chain as defined in Theorem 5. The mean first passage time is a convex function over the set S_w (refer to [13] for the proof). Moreoever, the second term in equation (14) is an affine function over the set S_w. Since the positive weighted sum of convex and affine functions is convex, the function $P_w \mapsto f(P_w)$ is convex over the set S_w. The set S_w is also convex and the constraints of the problem are affine. Hence, the optimization problem is convex. ■

The Optimization Problem 2 can be written as a semidefinite program. In order to do this, the expression for the detection

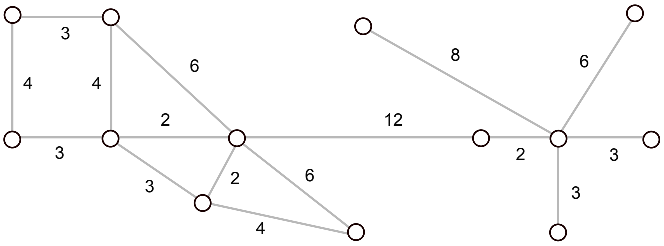

Fig. 3. The environment consists of thirteen regions of interest. A long passage connects a subgraph containing seven of the regions to another star subgraph consisting of the remaining six regions.

delay is rewritten in terms of the trace of a matrix as

$$\begin{aligned}\delta_{\text{avg}}(\mathcal{P}_w) &= \left(w^T(P_w \circ D)\mathbb{1}_n\right) \\ &\quad \times \left(\sum_{i=2}^{n} \frac{1}{1-\lambda_i(P_w)} + (s-\mathbb{1}_n)^T\mathbb{1}_n\right) \\ &= \left(w^T(P_w \circ D)\mathbb{1}_n\right) \\ &\quad \times \operatorname{Tr}\left((I - W^{1/2}P_w W^{-1/2} + w_c w_c^T)^{-1}\right) \\ &\quad + \left(w^T(P_w \circ D)\mathbb{1}_n\right)(s-\mathbb{1}_n)^T\mathbb{1}_n,\end{aligned}$$

where $W = \operatorname{diag}[w]$ and the column vector $w_c = (\sqrt{w_1}, \ldots, \sqrt{w_n})^T$. The first equation comes from Theorem 8 and the first part of the second equation is because of a relation between the trace of a function of P_w and its eigenvalues [13]. Using this form for $\delta_{\text{avg}}(\mathcal{P}_w)$, we can now formulate an SDP as shown below.

Problem 3: (Minimizing the upper bound on the optimal average detection delay (SDP)): Given the environment $\mathcal{E} = \langle V, E, D, w\rangle$ and vector of CUSUM samples $s \in \mathbb{R}^{n\times 1}$, with $W = \operatorname{diag}[w]$ and $w_c = (\sqrt{w_1}, \ldots, \sqrt{w_n})^T$, determine $Y = [y_{ij}] \in \mathbb{R}^{n\times n}$, $X \in \mathbb{R}^{n\times n}$, $t \in \mathbb{R}$ and $u \in \mathbb{R}$ solving:

$$\begin{aligned}&\text{minimize} \quad \operatorname{Tr}(X) + u(s^T\mathbb{1}_n) \\ &\text{subject to} \\ &\qquad \begin{bmatrix} t(I + w_c w_c^T) - W^{1/2}YW^{-1/2} & I \\ I & X\end{bmatrix} > 0 \\ &\qquad \begin{bmatrix} t & 1 \\ 1 & u\end{bmatrix} > 0 \\ &\qquad \sum_{j=1}^{n} y_{ij} = t, \text{ for each } i \in \{1,\ldots,n\} \\ &\qquad w_i y_{ij} = w_j y_{ji}, \text{ for each } (i,j) \in E \\ &\qquad 0 \le y_{ij} \le t, \text{ for each } (i,j) \in E \\ &\qquad y_{ij} = 0, \text{ for each } (i,j) \notin E \\ &\qquad w^T(Y \circ W)\mathbb{1}_n = 1 \\ &\qquad t \ge 0.\end{aligned}$$

Then, the transition matrix P_w is given by $P_w = Y/t$.s

Let P_{ub} be the solution to the Optimization Problem 2. We refer to the policy $\mathcal{P}_{\text{ub}} = \langle P_{\text{ub}}, w\rangle$ as the *efficient policy* for convenience.

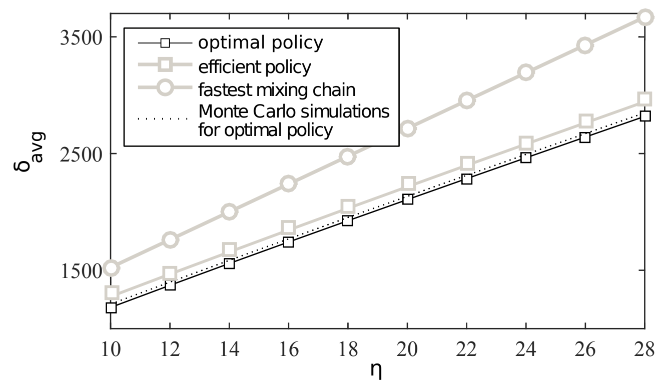

Fig. 4. Variation of the average detection delay using the optimal policy δ^*_{avg} (black squares), the efficient policy δ_{ub} (grey squares) and the policy based on the fastest mixing non-reversible Markov chain with a uniform stationary distribution (grey circles) with respect to the threshold η of the CUSUM algorithm. Expected detection delay for the optimal policy using Monte Carlo Simulations (dashed lines).

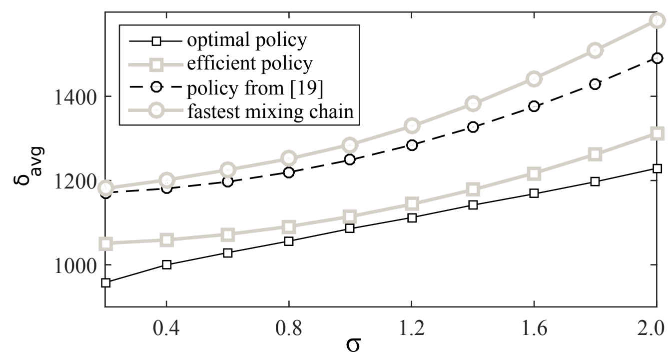

Fig. 5. Average detection delay using the optimal policy δ^*_{avg} (black squares) and the efficient policy δ_{ub} (grey squares) are compared with the average detection delay obtained using the policy based on the fastest mixing non-reversible Markov chain (grey circles) with a uniform stationary distribution and the policy from [19] (black circles) for various levels of noise in sensor measurements made in one of the regions.

V. Numerical Simulations

We now study the spatial quickest detection task for a specific environment. In particular, we are interested in examining the efficiency of the upper bound service policy $\mathcal{P}_{\text{ub}} = \langle P_{\text{ub}}, w\rangle$ (termed the efficient policy) compared to the optimal policy $\mathcal{P}^*$ which minimizes the average detection delay. We also compare these two policies to some candidate policies (namely the policy based on the fastest mixing non-reversible Markov chain [3] and a policy proposed in [19]).

Environment and modeling of anomalies: The environment modeled as a graph (Fig. 3) is an area separated into thirteen regions of interest. The edge weights of this graph represent the travel times between neighboring regions. All regions in the environment have equal priority, so that $w = \mathbb{1}_n$, and the service time required to make an observation in each of the regions is one time unit. The probability density functions of the observations made in the environment in the absence and presence of anomalies are normal distributions $f_k^0 = \mathcal{N}(0,1)$ and $f_k^1 = \mathcal{N}(1,1)$ respectively for $k \in \{1,\ldots,n\}$.

Computation of service policies: The Optimization Problem 1 to determine the optimal policy $\mathcal{P}^*$ is non-convex with nonlinear constraints. We solve it using the sqp algorithm in Matlab and verify that the solutions obtained are at least local minima. This is done by ensuring that the solutions satisfies 1. the regularity condition and 2. the Karush-Kuhn-Tucker (KKT) conditions necessary for the solution to optimal. On the other hand, the Optimization Problem 2 to compute the Markov chain corresponding to the upper bound service policy $\mathcal{P}_{\text{ub}}$ is convex and can be written as a semidefinite program. It is solved using CVX, a Matlab-based package for convex programs [8]. The fastest mixing non-reversible Markov chain is also computed by solving a semidefinite program in CVX. The policy proposed in [19] is stated as follows: $\mathcal{P}^\dagger = \langle P^\dagger, q^\dagger \rangle$ where

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

and $P^\dagger$ is the fastest mixing non-reversible Markov chain with stationary distribution $q^\dagger$.

Validation of theoretical expressions: We start with comparing the theoretical expression for the average detection delay δ_{avg} in the environment obtained in Theorem 7 (black squares) to the expected detection delay computed through Monte-Carlo simulations (dotted lines) in Fig 4 for the optimal policy. For these simulations, the initial position of the surveillance vehicle is sampled from the stationary distribution corresponding to the Markov chain it uses. No anomalies are present in the beginning of the simulation. The gap between the theoretical and the numerically obtained values is attributed to Wald's approximation introduced in equation (2).

Comparison of performances of service policies: We first compare variation in the performance of various service policies with respect to different thresholds η of the CUSUM algorithm in Fig. 4(a). The average detection delay δ_{ub} obtained using the efficient policy $\mathcal{P}_{\text{ub}}$ (grey squares) is close to the optimal average detection delay δ^*_{avg} (black squares) for all values of the threshold η. The gap observed between the optimal solution and the upper bound can be attributed to two factors: freedom to choose any stationary distribution as well as relaxation of the nonreversibility constraint for computing the optimal solution.

In comparison, the performance of the fastest mixing non-reversible Markov chain with stationary distribution $w = \mathbb{1}_n$ (grey circles) is much poorer. This is expected since the efficient policy is guaranteed to perform better in comparison to the fastest mixing non-reversible Markov chain with the same stationary distribution.

Next, we study the effect of the variation in the probability density functions of observations on the performance of the service policies. We consider a scenario where the region in the star subgraph farthest away from the center, is affected by noisy observations. While the probability distribution functions for observations in all the other regions remain same, they are different for this region: $f^0 = \mathcal{N}(0,\sigma)$ and $f^1 = \mathcal{N}(1,\sigma)$ and average detection delays of the various policies considered in the paper are compared for different values of σ in Fig. 5. For this simulation, $\eta = 10$ is chosen. The performance of the efficient policy (grey squares) is close to the optimal performance (black squares) for a wide range of σ in this case. In comparison, the performances of the policy based on the fastest mixing non-reversible chain with stationary distribution $w = \mathbb{1}_n$ (grey circles) and the policy from [19] (black circles) are much poorer.

VI. Conclusion

We studied the problem of quickest detection of anomalies based on sensor observations in environments with arbitrary graph topologies. We analyzed the Ensemble CUSUM Algorithm for this surveillance task and provided guarantees on its performance. We framed an optimization problem to compute the optimal policy for the Ensemble CUSUM Algorithm. We proposed an efficient policy which can be computed by solving a convex optimization problem. Through numerical simulations, we compared the performance of the optimal policy to the efficient policy. The detection delays guaranteed by the efficient policy were much smaller compared to alternative policies considered, especially for higher levels of uncertainties in sensor observations.

The policies introduced in this work can be modified further. For instance, the state of the CUSUM filters can be utilized to re-evaluate the routing strategy and achieve faster anomaly detection. An example of such a modification is the Adaptive Ensemble CUSUM Algorithm introduced in [19]. The service time at each region can also be optimized to obtain smaller detection delays. In situations where unpredictability is not required, deterministic policies which may potentially outperform stochastic policies can be further explored.

Appendix A
Computing the first passage time matrix of a Markov chain

The following result from [13] provides an expression for the first massage time matrix N of a Markov chain with transition matrix P in terms on the generalized inverse G of $(I-P)$.

Theorem 10: (First passage times): For a Markov chain on a graph with transition matrix P, stationary distribution q and travel time matrix D, the first passage time matrix N satisfies the following equations:

$$N_d = q^T(P \circ D)\mathbb{1}_n Q, \tag{15}$$

$$N = \beta(\mathbb{1}_n \mathbb{1}_n^T G_d + I - G)Q^{-1}. \tag{16}$$

where $G = ((I-P) + (P \circ D)\mathbb{1}_n q^T)^{-1}$, N_d and G_d are diagonal matrices with elements same as that of N and G respectively, $Q = \text{diag}[q]$ and $(P \circ D)$ is the element-wise product between P and D.

The following identity satisfied by the generalized inverse G of $(I-P)$ will also be used.

Lemma 11: For a Markov chain on a graph with transition matrix P, stationary distribution q and generalized inverse $G = ((I-P) + (P \circ D)\mathbb{1}_n q^T)^{-1}$, the following holds true:

$$q^T G = \frac{q^T}{q^T(P \circ D)\mathbb{1}_n}. \tag{17}$$

Proof: The proof can be found in Lemma 16 of [13]. ■

APPENDIX B
PROOF OF THEOREM 7

Proof: We start from the expression for $\delta_{\text{avg}}(\mathcal{P})$ obtained in Theorem 6. Using the definition of s stated in the statement of the theorem, and equation (7), the expression for $\delta_{\text{avg}}(\mathcal{P})$ can be written in matrix form as follows:

$$\delta_{\text{avg}}(\mathcal{P}) = q^T N w + (s - \mathbb{1}_n)^T N_d w. \tag{18}$$

We first work towards simplifying the first term in equation (18). The first passage time matrix satisfies equation (16). Using this and the identity from Lemma 11, and the assumption that $\mathbb{1}_n^T w = 1$, the term $q^T N w$ can be simplified:

$$\begin{aligned} q^T N w &= q^T \beta(\mathbb{1}_n \mathbb{1}_n^T G_d + I - G) Q^{-1} w \\ &= \beta(\mathbb{1}_n^T G_d + q^T(I - G)) Q^{-1} w \\ &= \beta(\mathbb{1}_n^T G_d + q^T - \frac{q^T}{\beta}) Q^{-1} w \\ &= \beta \mathbb{1}_n^T G_d Q^{-1} w + (\beta - 1) q^T Q^{-1} w \\ &= \beta \mathbb{1}_n^T G_d Q^{-1} w + (\beta - 1) \mathbb{1}_n^T w \\ &= \beta \mathbb{1}_n^T G_d Q^{-1} w + (\beta - 1). \end{aligned} \tag{19}$$

where $Q = \text{diag}[q]$. Looking at the first term in equation (19),

$$\begin{aligned} \mathbb{1}_n^T G_d Q^{-1} w &= \mathbb{1}_n^T [((I - P) + (P \circ D)\mathbb{1}_n q^T)^{-1} \circ Q^{-1}] w \\ &= \mathbb{1}_n^T [((I - P) + (P \circ D)\mathbb{1}_n q^T)^{-1} \circ I](r \cdot w). \end{aligned}$$

Substituting $N_d w = \beta Q^{-1} w = \beta(r \circ w)$ from equation (15) into the second term in equation (18), the result follows.

■

REFERENCES

[1] A. Almeida, G. Ramalho, H. Santana, P. Tedesco, T. Menezes, V. Corruble, and Y. Chevaleyre. Recent advances on multi-agent patrolling. In *Advances in Artificial Intelligence*, volume 3171 of *Lecture Notes in Computer Science*, pages 474–483. Springer, 2004.

[2] T. H. Blackwell and J. S. Kaufman. Response time effectiveness: Comparison of response time and survival in an urban emergency medical services system. *Academic Emergency Medicine*, 9(4):288–295, 2002.

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

[4] 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.

[5] J. Clark and R. Fierro. Mobile robotic sensors for perimeter detection and tracking. *ISA Transactions*, 46(1):3–13, 2007.

[6] 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.

[7] 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.

[8] M. Grant and S. Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. http://cvxr.com/cvx, October 2014.

[9] 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.

[10] J. López, D. Pérez, E. Paz, and A. Santana. WatchBot: A building maintenance and surveillance system based on autonomous robots. *Robotics and Autonomous Systems*, 61(12):1559–1571, 2013.

[11] G. Lorden. Procedures for reacting to a change in distribution. *The Annals of Mathematical Statistics*, 42:1897–1908, 1971.

[12] 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, 2012.

[13] R. Patel, P. Agharkar, and F. Bullo. Robotic surveillance and Markov chains with minimal weighted Kemeny constant. *IEEE Transactions on Automatic Control*, 2015. To appear.

[14] T. Sak, J. Wainer, and S. Goldenstein. Probabilistic multiagent patrolling. In G. Zaverucha and A. Loureiro da Costa, editors, *Brazilian Symposium on Artificial Intelligence*, volume 5249 of *Advances in Artificial Intelligence*, pages 124–133. Springer, 2008.

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

[16] 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.

[17] 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.

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

[19] V. Srivastava, F. Pasqualetti, and F. Bullo. Stochastic surveillance strategies for spatial quickest detection. *International Journal of Robotics Research*, 32(12):1438–1458, 2013.

[20] S. Susca, P. Agharkar, S. Martínez, and F. Bullo. Synchronization of beads on a ring by feedback control. *SIAM Journal on Control and Optimization*, 52(2):914–938, 2014.

[21] S. Thrun. Robotic mapping: A survey. In G. Lakemeyer and B. Nebel, editors, *Exploring Artificial Intelligence in the New Millennium*, pages 1–35. Morgan Kaufmann, 2002.

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

[23] P. R. Wurman, R. D'Andrea, and M. Mountz. Coordinating hundreds of cooperative, autonomous vehicles in warehouses. *AI Magazine*, 29(1):9–20, 2008.