Performance Analysis of Stochastic Network Coverage with Limited Mobility

Chris Y. T. Ma, David K. Y. Yau, Nung Kwan Yip
Purdue University
West Lafayette, IN, USA

Nageswara S. V. Rao
Oak Ridge National Lab
TN, USA

Jiming Chen
Zhejiang University
Hangzhou, China

Abstract

We analyze the ability of a stochastic coverage algorithm to achieve both accurate threat-based coverage and effective information capture. When mobile sensors are used to cover the region over time, the goal of threat-based coverage *is to allocate the sensors' coverage time between the subregions in proportion to their threat levels. We show that, in contrast to prior results on mobile coverage for maximizing simple event capture, limiting mobility by strategically* pausing *the sensor is important for threat-based coverage of physical world monitoring. Besides being energy efficient, pausing has two desirable effects. First, it can improve the accuracy of the threat-based coverage, in particular, the accuracy increases monotonically with a pause time parameter, and a large enough parameter will ensure* exact *matching of the sensor's coverage profile with the region's threat profile. Second, diverse natural phenomena require a non-negligible sensing time to overcome statistical uncertainties posed by the random nature of the phenomena. Suitable pausing allows a subregion to be observed long enough for reliable results.*

1. Introduction

Monitoring the physical world has important applications. One example is the protection of the country against chemical, biological, nuclear, radiological, and explosive (CBNRE) threats in homeland security. In such an application, the observed environmental parameters are mostly random in nature. For instance, a common characteristic of radiation detection is that a sequence of sensor readings over a significant time interval is needed to produce high-confidence detection, due to the inherent probabilistic nature of the radiation source, the presence of background radiation, and noise/interference in the surveillance area. This gives rise to a *temporal dimension* of the sensing problem, namely the need to collect and process radiation counts over suitable time intervals for reliable detection.

Research was supported in part by Office of Naval Research under grant no. DE-AC05-00OR22725, by U.S. National Science Foundation under grant no. DMS-0707926, by NSFC-Guangdong under Union Project grant no. U0735003, by NSFC under grant nos. 60604029 and 60736021, by 111 Projects under grant no. B07031, and by Zhejiang University State Key Lab of Industrial Control Technology under Open Funding grant no. ICT0901.

Under the federal SensorNet initiative, sensor-cyber networks have been deployed at the Port of Memphis [4] and the Washington DC area to protect the areas' population against the exposure to chemical and radiation risks, respectively. Both deployments are in populated areas and their primary objective is the detection of the threat. Lessons learned from these deployments highlight the importance of the following requirements:

Management of resource constraints. The sensors are expensive both in terms of procurement and management costs, and typically have only a modest sensing range. Thus, it is essential to optimize the sensor placements, and develop methods to further improve the coverage by exploiting sensor mobility.

Importance of people protection. The project report [4] states that the main decision of where to place a limited amount of sensing resources is based on the impact on the area's population, since "human effects represent the true consequences of failure" to detect a harmful agent and subsequently evacuate the affected population. Hence, at each step of the placement algorithm [4], a search procedure is used to determine the location of the next sensor that will maximize the marginal gain in protection for the residents.

Need for uncertainty reduction. The physical world being sensed is inherently noisy and probabilistic, and the measurements may be imprecise or have errors. The measurements for radiation detection are probabilistic with a high variance [7]. While chemical sensor measurements exhibit lower variances, the measurement errors are not negligible. Thus, there is a need to remove statistical outliers or cancel out their effects in the measurement process. Longer measurement periods lead to higher detection confidence, but if the sensors are to be stationary, the coverage extent becomes limited.

In this paper, we focus on the performance of mobility-based coverage that utilizes a limited number of sensors to defend against CBNRE threats under the above requirements. We will illustrate our problem for the case of radiation detection. The major goals of our sensor coverage algorithm are the following:

G1. We seek to allocate constrained sensing resources to different subregions in proportion to their threat levels, where a subregion's threat level is defined as its number of residents exposed to undetected radioactivity.

G2. In the resource allocation, both the amount and the

978-1-4244-5113-5/09/$25.00

quality of the information collected are relevant. It is desirable to capture a larger fraction of the interesting events. It is also important to maximize the confidence about the captured events.

G3. While constrained resources will require sensors to move between the subregions requiring coverage, we should limit the amount of the movement to achieve a low cost and practical solution.

The main **contribution** of this paper is to analyze the costs and benefits of mobility for threat-based sensing with a temporal dimension. We analyze a stochastic mobile coverage algorithm called WRW-aLP, and compare its performance with best-case static coverage and the mobile coverage algorithm in [1] via simulation experiments. Our analytical and simulation results argue for a limited form of mobility in which the sensor moves between points of interest (PoIs) for expanded area coverage, but also pauses strategically at the PoIs for improved performance. We show that, while mobility is useful, intermittent pausing has three main benefits. First, it reduces deployment costs such as energy use. Second, we prove that the accuracy of threat-based coverage increases monotonically with the pause time parameter of WRW-aLP, in particular, a large enough pause time parameter will ensure *exact* matching of the sensor's coverage profile with the surveillance region's threat profile. The analytical results hold in spite of *side effects* of coverage as the sensor moves from one PoI to another. Third, whereas a faster sensor always increases the fraction of events captured, as established by the results in [1], [5], the sensing *uncertainty* about each captured event also increases when the temporal dimension is present. In this case, suitable pausing at the PoIs can increase the utility of the information captured.

2. Related Work

Our work complements existing work on mobile coverage for simple event capture [1], [5]. The goal of threat-based mobile coverage is similar to that in [4] for static coverage in terms of people protection. Proportional sharing of coverage time has also been studied in [11]. Our work differs from theirs in three respects: (i) They study deterministic coverage algorithms, whereas our algorithm is stochastic; (ii) there are no side effects of coverage in their system model, whereas analyzing the side effects is one of our primary concerns; and (iii) they do not study the issues of sensor coordination.

Radiation detection has been studied in [2]. The relevance of the temporal dimension in radiation detection and other detection tasks has been documented in [8]. Our goal is to understand the impact of the temporal dimension on mobile sensor coverage.

We study the use of *limited mobility* to obtain the benefits of mobile coverage at low cost. Prior work has limited mobility by either (1) using a hybrid network architecture in which only a fraction of the nodes move during deployment [9], [10], or (2) moving nodes only when the network is first formed or when there are significant changes such as node failures [12]. In the first approach, the nodes that do move do so continuously, making it harder to manufacture these nodes or ensure their continuous operation under limited energy. In the second approach, the limited movement can improve the system's robustness against changes. However, it still requires that in the stable state, there are sufficient sensor resources to cover the whole surveillance region.

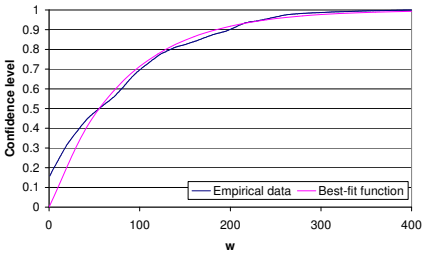

Figure 1. Utility of measurement as a function of w: Empirical characterization and least-square fit function.

3. Neyman-Pearson Test & Sensing Utility

We consider the detection of a point radiation source of strength A, in counts per minute (CPM), such that an ideal detector without background radiation located at a distance d from the source will register a count of c in a one second time interval. We know that c is Poisson distributed with parameter A/d^2 [3]. With background radiation, a detector may register a random count even when there is no identifiable source present. Hence, we use the Neyman-Pearson test [8] to ensure that a count is due to a radiation source, and not due to random fluctuations of the background radiation, which can be modeled as a point source of strength B. The hypothesis testing works as follows [8]: $\mathbf{H_0}$: c is Poisson distributed with parameter B; $\mathbf{H_1}$: c is Poisson distributed with parameter $B + A/d^2$. We can then formulate the Neyman-Pearson test with false alarm probability α by computing a threshold τ such that if $Pr(c|H_1)/Pr(c|H_0) > \tau$, then H_1 is chosen, otherwise, H_0 is chosen. The value of τ that yields the desired α can be computed using Lagrangian method [8].

To demonstrate the probabilistic nature of the sensing process, we used the RFTrax to measure a close-by radiation source. To be safe for experimentation, the source had an extremely low intensity, and was placed at a small distance of 14.6 cm from the sensor. We collected CPM readings from the sensor at two second intervals, for a total of more than 4000 seconds. Let x_i be the ith sensor reading, $i = 0, \ldots$. The true reading, $\bar{x}$, of the source can be taken as the mean over a large number of readings, i.e., $\bar{x} = \frac{1}{n}\sum_{k=0}^{n} x_k$, for a sufficiently large n. We compute a sequence of moving averages, of window size w, over the sequence $\{x_i\}$.

The jth moving average is given by $\bar{x}_j = \frac{1}{w}\sum_{k=j}^{j+w} x_k$. For each w, we measure the fraction of the $\bar{x}_j$'s that satisfy $|\bar{x}_j - \bar{x}| \leq \epsilon$, for a small ϵ. This fraction of the averages that deviate little from $\bar{x}$ is then the *confidence* of the measurement given the measurement window size w. A plot of the confidence against w is given in Fig. 1. Interpreting the confidence of a sensing result as its *utility*, we can also use Fig. 1 as a *utility function* of the sensing against the sensing time.

4. System Model

We assume that a 2D rectangular geographical region is under surveillance. The region is partitioned into a vector of n disjoint cells each of dimension $S \times S$, where S is in distance units. Each cell, say i, has a population size of s_i, where s_i is a non-negative integer. We define a cell to be a *point of interest* (PoI) if a radioactive source may appear at the center of the cell. The presence duration of a radioactive source at a PoI is called an *event*. An event is dynamic in that it appears and disappears according to given arrival/departure processes, and events do not overlap in time at a PoI. For simplicity, we assume that the event dynamics at the PoIs are identical and Poisson distributed, but the PoIs vary in importance in terms of their population sizes. We then define the *threat profile* of the surveillance region as an n-element row vector Φ such that if cell i is a PoI, $\Phi_i = s_i/\Sigma_j s_j$ is the threat level of i. If cell i is not a PoI, $\Phi_i = 0$.

We further assume that each sensor has a sensing range of $S/\sqrt{2}$, so that a sensor within PoI i can always sense a radioactive source in i towards the targeted false alarm rate α in Section 3. The sensor can move within the surveillance region, subjected to possible accessibility constraints (e.g., a sea area is inaccessible to a land sensor). Consider a deployment of m sensors up to real time T, of which T_O^j time is spent as *travel time overhead* by sensor j, meaning that sensor j is not within range of any PoI for $T_O^j \leq T$ time during the deployment, and T_R^j time is wasted by sensor j as it is within range of any PoI that is already monitored by another sensor. Denote $T_D = m \times T - \sum_j (T_O^j + T_R^j)$ as the *duty time* of all sensors up to time T. Define $C = \{C_i\}$, the *global coverage profile*, as a row vector such that C_i is the total amount of coverage time received by cell i from any sensor up to time T. *We seek to achieve threat-based coverage, i.e., if cell i is a PoI, C_i is targeted to be $\Phi_i \times T_D$.* To track the threat-based coverage goal, we define an n-element row vector $U = \{U_i\}$ such that $U_i = \Phi_i \times T_D - C_i$ is the *undercoverage time* of cell i. Note that U_i can be negative, in which case i is *over-covered*.

An event can be detected if its occurrence is within range of a sensor during the event's lifetime. However, as discussed in Section 3, a single reading while useful may not be reliable, and a longer sensing duration is needed to improve the sensing utility. We say that an event is *captured* if during its lifetime, it is within range of at least one sensor for one reading by the Neyman-Pearson method. Otherwise, the event is uncaptured, and the event is taken to have a zero utility of measurement.

5. Coverage Algorithm and Performance

5.1. Performance metrics

Before presenting the coverage algorithm, we address how a candidate algorithm should be evaluated. The evaluation consists of two sets of metrics. The first set quantifies three general properties of the mobility algorithm:

Matching. The matching between the achieved global coverage profile and the given threat profile at time t is quantified by the percentage deviation of the former from the latter, which is equal to $\frac{100}{2} \times \sum_i |\Phi(i) - \frac{C_i}{t}|$.

Unfairness. This is the average exposure time of the PoIs, i.e., the average duration of the continuous time interval over which the PoIs are not covered. The unfairness measure is given by $\sum_i \epsilon(i) \times \Phi(i)$, where $\epsilon(i)$ is the average exposure time of PoI i.

Effective coverage. The effective coverage at time t is given by $\frac{T_D}{t \times m}$, where m is the number of sensors. A higher effective coverage means that the sensors spend more of their time doing useful work in terms of monitoring events without duplicated efforts. Two factors may reduce the effective coverage of an algorithm: (i) the travel overhead of a sensor during which it is not covering any PoI but is traveling between the PoIs, and (ii) the redundancy of coverage when multiple sensors cover the same PoI at the same time.

Further to the general performance properties, the second set of metrics quantifies the information capture of definite stochastic events:

Normalized utility of the events captured. This is the sum of utilities of all the events captured within a given time interval, normalized by the total number of events that appear during the time interval. A higher normalized utility shows that the sensors can collect a larger total amount or fraction of the interesting information.

Confidence about each captured event. A coverage algorithm may achieve a large aggregate utility by capturing many events, but with low confidence for each one. This metric further characterizes the algorithm performance by quantifying the utility of the measurements given the cumulative sensing time of each event.

5.2. Stochastic threat-based coverage

We present a stochastic mobile coverage algorithm based on a *weighted* random waypoint (WRW) design. (Further details of the algorithm can be found in [6].) In

the algorithm, the sensor moves in a sequence of trips. Each trip starts at the center of some PoI cell, say i, and ends at the center of another PoI cell, say j. The ending point, called a *waypoint*, of one trip becomes the starting point of the next trip, and so on. The speed of the sensor during a trip is fixed to be v.

Suppose cell i is the current position of the sensor. A cell $j, j \neq i$, is chosen to be the next waypoint with probability Φ_j. Such weighting of the random waypoint selection by the threat profile is a first step towards achieving threat-based coverage. For simplicity, we assume that every cell is reachable from every other cell. The exact path connecting i and j must not cross an inaccessible area, but its determination is otherwise flexible, e.g., it can be specified as the direct straight line from i to j if the line does not violate inaccessibility constraints.

The basic algorithm is simple, but its coverage profile fails to accurately match the threat profile, because it fails to consider the *intermediate* cells covered between the source and destination. For example, consider a surveillance region with a few high threat hotspots. In moving between the hotspots to give them adequate coverage, the sensor will also visit frequently all the cells on the paths between the hotspots, thus over-covering the intermediate cells. To solve this important problem of *side-effect coverage*, the basic algorithm is augmented with the features:

Maximum trip length. The distance of one trip is not allowed to exceed a parameter L (in distance units). Hence, when we choose the next waypoint, we restrict the candidate cells to be within the disc of radius L and centered at the current cell. Limiting the trip length forces the algorithm to consider more possible routes to go between any two hotspots, thus reducing the possibility of "warming up" the intermediate cells.

Adaptivity to prior coverage. Because of the probabilistic nature of the algorithm, the correlations between the cells visited, and the finite speed of the sensor, the algorithm's actual coverage at any time may deviate from the given threat profile. To correct the deviation, in selecting the next PoI to visit, the algorithm adapts to the current actual coverage and selects a more severely undercovered PoI with higher probability. The precise definition of such adaptation is given in Section 6.

Random pause time. If the sensor is at an undercovered cell, one way to correct the undercoverage is for the sensor to stay in the cell for some pause time T_p. The time T_p is drawn randomly from a distribution determined by a pause time parameter denoted by P (in time units). Specifically, at the end of the tth trip at destination cell j, $T_p \sim \mathbf{U}(0, \Omega_t(j))$, where

$$\Omega_t(j) = \frac{P \times \max(U_j, 0)}{\Sigma_{i \in \mathcal{C}} \max(U_i, 0)},$$

$\{U_j\}$ is the vector of undercoverage times defined in Section 4, and $\mathcal{C}$ is the set of cells that are candidates as the next waypoint.

Remarks: (A) Each of the above mentioned features can be enabled independently. We denote by WRW-$feat$ a particular extension of the WRW algorithm, where $feat$ is a combination of the enabled features represented by the letters a, L, and P for adaptivity, maximum trip length, and random pause time, respectively.

(B) The range of the pause time is controlled by P. In the special case that $P = 0$, the algorithm does not pause and performs as a continuous movement algorithm. In general, the pause time is expected to be larger when the undercoverage is higher. After the pause time, the selection of the next waypoint that defines the next trip occurs as in the previous description. The pause time attempts to correct the undercoverage in an extremely efficient way – with zero movement overhead and no possibility of inadvertently changing the coverage of other cells. *An important objective of this paper is to determine analytically and experimentally the impacts of such pausing on the accuracy of threat-based coverage and the utility of the sensing results.*

(C) We specify that sensors use the global coverage profile to determine the undercoverage of PoIs. We assume that there exists a wireless infrastructure for the sensors to exchange their local coverage information. If such an infrastructure is unavailable, the sensors can use their local coverage history to approximate the global one in computing the undercoverage, making the algorithm fully distributed in nature. On the other hand, a deterministic mobile coverage algorithm for maximizing the number of events captured at given PoIs has been proposed by Bisnik, Abouzeid, and Isler [1]. We call this the BAI algorithm. It is a *centralized* approach since a pre-defined path is first computed and all sensors follow the same path. The BAI algorithm is designed for simple event capture, without consideration for the temporal dimension or the threat-based coverage. Part of our simulation results illustrate the performance difference between these two algorithms. □

5.3. Multiple-sensor coordination

If there are multiple sensors, they may coordinate their operation for better performance. In [6], it is shown that *when the sensor density is low*, coordination may be unnecessary in that employing independent sensors may already lead to linear performance improvements in terms of the first two general performance metrics in Section 5.1 (i.e., matching and unfairness). When the number of available sensors increases for the same surveillance area, the sensor density increases and the redundancy of coverage problem may become more significant. Moreover, good coordination between sensors may allow to minimize the travel overhead, leading to superlinear performance improvements in terms of effective coverage. In this paper, we consider the performance of the three coordination approaches in [6] in

the case that the sensor density is high:

No coordination (NC): The sensors are deployed independently each according to the WRW-aLP algorithm. They use their local coverage history to approximate the global coverage profile. The approach is simple to implement but does not aim to reduce coverage redundancy.

Knowledge of global coverage profile (GK): A sensor knows the coverage profiles of all the other sensors, and the sensors use the global coverage profile to determine the undercoverage of the PoIs. The approach does not eliminate the redundancy of coverage since the sensors do not explicitly exchange their current locations and try to avoid each other in their movements. However, the global coverage profile allows the sensors to compensate for prior matching inaccuracies due to redundant coverage. Hence, GK gives better *matching* of the global coverage with the given threat profile. When the sensor density is high, however, the coverage redundancy may still cause significant reductions of the sensors' effective coverage for information capture.

Static map division (MD): The sensors are assigned to cover non-overlapping subregions of the surveillance area. Each sensor follows the WRW-aLP algorithm within its assigned subregion. The matching and unfairness of this approach rely heavily on the topology of the PoIs, the distribution of threats among the PoIs, and the manner in which the surveillance area is divided among the sensors. Particularly, a sensor in one assigned subregion will not be able to compensate for any undercoverage of PoIs in another subregion. Although the redundancy of coverage problem is eliminated in MD, the sensors' effective coverage is still not 100% due to travel overhead between the PoIs. Nevertheless, with proper division of the surveillance area, this travel overhead can be greatly reduced, and the effective coverage of MD may be significantly better than either NC or GK, thereby improving the information capture.

6. Analytical Results

6.1. Markov model

We model the WRW-aLP algorithm by a Markov Chain for a single WRW-aLP sensor. Without loss of generality, we assume that the PoI cells of the surveillance region are denoted as cells $i = 1, \ldots, n$ of the Markov model, and use cell 0 in the model to represent the collection of non-PoI cells in the surveillance region. The state of the Markov Chain is denoted by $X = (i, U)$, where i is the location of the sensor (i.e. the PoI the sensor is at), and U is the vector of undercoverage times as defined in Section 4.

Besides the notations in Section 4, it is convenient to introduce the following quantities. T_{ij} specifies the actual travel time of the sensor for a movement from i to j. In addition, the numbers $Z_{ij}(k)$ specify how such a trip increases the coverage time of each PoI k *during* the travel from i to j. We use $Z_{ij}(0)$ to denote the travel time overhead of the trip from i to j, i.e., the period during which the sensor is not within the range of any PoI. Hence, $\sum_{0 \le k \le n} Z_{ij}(k) = T_{ij}$. Finally, we use $D_{ij} = T_{ij} - Z_{ij}(0)$ to denote the duty time portion of the actual travel time.

To admit more general models, we allow $i = j$, i.e., the sensor can decide to remain at its current location in a transition. In this case, T_{ii} is set to be zero.

Notice that all the quantities $\{T_{ij}\}$, $\{Z_{ij}(k)\}$ and $\{D_{ij}\}$ are deterministic properties of the surveillance region and the paths between pairs of PoIs. Thus they can be pre-computed by the coverage algorithm.

A transition from state $X = (i, U^X)$ to state $Y = (j, U^Y)$ in the Markov Chain is labeled by (i, j), a trip of the sensor from cell i to cell j. As the goal of the WRW-aLP algorithm is to reduce the undercoverage, the sensor tends to go to PoIs with large undercoverage. In a stochastic algorithm, the transition probability is chosen according to:

(1) Whether j is within a given distance from i. If so, we say that j is *eligible*. We indicate the condition by the indicator function $I(i, j)$, i.e., $I(i, j) = 1$ if j is within the allowed distance of i, otherwise $I(i, j) = 0$.

(2) The relative under-coverage times between the eligible PoIs. The probability that j is chosen as the next PoI is given by

$$Pr_{ij} = \frac{I(i,j) \times W_j U_j^+}{\sum_{1 \le k \le n} I(i,k) \times W_k U_k^+} \qquad (1)$$

where $U_k^+ = \max(U_k, 0)$ and W_k are the weights given as some positive numbers, e.g., $W_k = \Phi_k$.

In the degenerate cases that $I(i, k) = 0$ or $U_k^+ = 0$ for all k, we can simply choose the next PoI randomly, according to some uniform distribution. The precise choices of the W_k's and what to do in the degenerate cases will not affect the qualitative results.

In WRW-aLP, once the sensor has moved from i to j, it will stay there for a pause time, $T_P(U_j)$, which can depend on j's current undercoverage. The value of this pause time is part of the decision of the transition. Hence a movement from i to j increases the sensor's duty time by $D_{ij} + T_P(U_j)$. In addition, for each cell, its coverage time will be increased by $Z_{ij}(k)$ for $k \neq j$ and $T_P(U_k)$ for $k = j$. Using the above, the change of the undercoverage $U_k^Y - U_k^X$ (for $k = 1, \ldots n$) can be computed explicitly.

Note that the transition probability from X to Y given by (1) depends only on the current state i and U^X. Hence it models a Markov Chain process.

6.2. Performance analysis

This section analyzes the WRW-aLP Markov Chain model for the single sensor. One of the main behaviors of any Markov Chain concerns its long time behavior and ergodicity and stationarity properties. It is natural

to ask these questions to our WRW-aLP model. Many of the standard results are expected to be true. *However, due to the side effects of coverage—the sensor can pass by some unintended PoIs due to $Z_{ij}(k) > 0$ for $k \neq j$—it is not immediately clear whether the long time behavior can be related to the desired threat-based coverage, or the threat profile can be achieved.*

For example, without the side effects, i.e., $Z_{ij}(k) = 0$, to achieve the threat profile Φ, we can simply construct a Markov Chain with Φ as its stationary distribution. Such a chain can be realized by appropriate Monte-Carlo simulations. However, with non-zero Z_{ij}'s, the analysis needs to be modified with some care. Our main results state that the desired threat profile can *always* be achieved, provided that the pause time is *long enough* compared with the Z_{ij}'s. They are proved through the concept of *Lyaponouv function*, which appears in the study of the stability properties of dynamical systems.

To simplify the analysis, we set all the pause times $T_P(U_j)$'s equal to some fixed constant T_P. Furthermore, $I(i,j) = 1$ for all i and j, i.e., any PoI is accessible within one transition. Such an assumption can certainly be relaxed. In general, the pause time can even be random.

Consider the following quantity:

$$V(X) = \sum_{i=1}^{n} \alpha_i (U_i^X)^2 \tag{2}$$

where α_i's are some weights to be determined. Clearly, $V = 0$ refers to perfect matching with the threat profile. Hence the goal is, in the long run, to make V as *small* as possible. The following two results show that "on average," V decreases as the mobile algorithm continues. To better illustrate the idea, we first consider the case of no side effects and treat the latter as a *perturbation*.

Theorem 1: Let $Z_{ij}(k) = 0$ for $i, j, k = 1, \ldots n$. There exists a positive number $\lambda > 0$ such that if $V(X) \geq \lambda$, then

$$E_X V(Y) \leq V(X) - 1, \tag{3}$$

where X and Y refer to the current and next states, respectively, and $E_X(\cdot)$ denotes the expectation given the state X. (The choice of the weights α_i's will be specified in the proof of this Theorem.)

Proof: Let the current and next sensor locations be i and j. Then the undercoverage of PoI k will be increased by $\Phi_k T_p$ for $k \neq j$ and $(\Phi_j - 1)T_P$ if $k = j$. Hence,

$$V(Y) = \sum_{k \neq j} \alpha_k \left(U_k^X + \Phi_k T_p\right)^2 + \alpha_j \left(U_j^X + (\Phi_j - 1)T_P\right)^2.$$

Then, the expectation $E_X(V(Y))$ equals

$$\sum_{j=1}^{n} \left[\sum_{k \neq j} \alpha_k \left(U_k^X + \Phi_k T_p\right)^2 + \alpha_j \left(U_j^X + (\Phi_j - 1)T_P\right)^2 \right] Pr_{ij}.$$

Now consider

$$\begin{aligned} & E_X V(Y) - V(X) \\ = & \sum_{j=1}^{n} \Big[\sum_{k \neq j} \alpha_k \Big(U_k + \Phi_k T_p\Big)^2 + \alpha_j \left(U_j + (\Phi_j - 1)T_P\right)^2 \\ & - \sum_{k=1}^{n} \alpha_k (U_k)^2 \Big] Pr_{ij} \quad \text{(note: } \textstyle\sum_j Pr_{ij} = 1\text{)} \\ = & Z_*^{-1} T_P \Big[2 \times I + II \Big], \end{aligned}$$

where $Z_* = \sum_{k=1}^{n} W_k U_k^+$. The quantities I and II are defined and analyzed as follows.

$$\begin{aligned} I &= \sum_{j=1}^{n} \Big(\sum_{k=1}^{n} \alpha_k \Phi_k U_k - \alpha_j U_j \Big) W_j U_j^+ \\ &\leq \Big(\sum_{k=1}^{n} \alpha_k \Phi_k U_k^+ \Big) \Big(\sum_{j=1}^{n} W_j U_j^+ \Big) - \Big(\sum_{k=1}^{n} \alpha_k W_k (U_k^+)^2 \Big), \end{aligned}$$

where we have used $U \leq U^+$ and $UU^+ = (U^+)^2$. Using the Cauchy-Schwartz inequality, we have

$$\begin{aligned} \Big(\sum_{k=1}^{n} \alpha_k \Phi_k U_k^+ \Big) &\leq \Big(\sum_{k=1}^{n} \Phi_k \Big)^{1/2} \Big(\sum_{k=1}^{n} \alpha_k^2 \Phi_k (U_k^+)^2 \Big)^{1/2} \\ \Big(\sum_{k=1}^{n} W_k U_j^+ \Big) &\leq \Big(\sum_{k=1}^{n} \frac{W_k}{\alpha_k} \Big)^{1/2} \Big(\sum_{k=1}^{n} \alpha_k W_k (U_k^+)^2 \Big)^{1/2}. \end{aligned}$$

By choosing $\Phi_k \alpha_k^2 = \alpha_k W_k$, i.e., $\alpha_k = W_k \Phi_k^{-1}$, then

$$\Big(\sum_{k=1}^{n} \frac{W_k}{\alpha_k} \Big)^{1/2} = \Big(\sum_{k=1}^{n} \Phi_k \Big)^{1/2} = 1,$$

which gives $I \leq 0$. On the other hand, equalities in the above Cauchy-Schwartz inequalities hold if and only if the vector $(U_k^+)_{k=1}^n$ is a multiple of $(\alpha_k \Phi_k)_{k=1}^n$ or $(W_k)_{k=1}^n$. This is impossible as $\sum_{k=1}^{n} U_k = 0$. We can thus infer the existence of a $\mu > 0$ such that

$$I \leq -\mu \sum_{k=1}^{n} \alpha_k (U_k^+)^2. \tag{4}$$

(See also Remark (B) on Page 7.)

Next, the quantity II equals:

$$II = T_P \sum_{j=1}^{n} \Big[\sum_{k \neq j}^{n} \alpha_k \Phi_k^2 + \alpha_j (\Phi_j - 1)^2 \Big] W_j U_j^+.$$

Note that it involves only *linear* terms of U_j^+. Hence, there exists a $C > 0$ such that $II \leq C \sqrt{\sum_{k=1}^{n} (U_k^+)^2}$.

So if $\sqrt{\sum_{k=1}^{n} (U_k^+)^2} \geq \lambda$ for some large number λ, we have

$$2 \times I + II \leq -\frac{\mu}{2} \sum_{k=1}^{n} \alpha_k (U_k^+)^2.$$

After taking into account of the prefactor Z_*^{-1} and redefining the value of the constant μ, we get the desired result:

$$E_X V(Y) - V(X) \leq -\mu \sqrt{V(X)} < -1.$$

In the above, we have used the fact that for $\sum_k U_k = 0$, there exists a $C > 0$ such that

$$\sum (U_k^+)^2 \leq \sum (U_k)^2 \leq C \sum (U_k^+)^2.$$

□

The next result incorporates the presence of side effects, i.e., $Z_{ij} > 0$. We use the same notation as in the previous Theorem.

Corollary 1: There exist $\lambda, T_P^* > 0$ such that for $T_P > T_P^*$ and $V(X) > \lambda$, then

$$E_X V(Y) \leq V(X) - 1.$$

Proof: We will only outline the proof as it is very similar to Theorem 1. Due to the Z_{ij}'s, the function $V(Y)$ evaluated at the new state is now given as:

$$\sum_{k \neq j} \alpha_k \left(U_k^X + \Phi_k T_p - Z_{ij}(k)\right)^2 + \alpha_j \left(U_j^X + (\Phi_j - 1)T_P\right)^2.$$

Hence, the quantity $E_X V(Y) - V(X)$ equals

$$\sum_{j=1}^{n} \Big[\sum_{k \neq j} \alpha_k \Big(U_k + \Phi_k T_p - Z_{ij}(k) \Big)^2$$

$$+ \quad \alpha_j \left(U_j + (\Phi_j - 1)T_P\right)^2 - \sum_{k=1}^{n} \alpha_k (U_k)^2 \Big] Pr_{ij}.$$

We proceed as before. The quantity I now has an extra term given by:

$$-\sum_{j=1}^{n} \left[\sum_{k \neq j} \alpha_k U_k \frac{Z_{ij}(k)}{T_P} \right] W_j U_j^+.$$

Note that if $\frac{Z_{ij}(k)}{T_P} \ll 1$, this extra term can be bounded by $c \sum_{k=1}^{n} (U_k^+)^2$ for some *small* constant c so that it can be absorbed by the $-\mu V(X)$ appearing in Theorem 1.

For II, it becomes

$$T_P \sum_{j=1}^{n} \Big[\sum_{k \neq j}^{n} \alpha_k \Big(\Phi_k - \frac{Z_{ij}(k)}{T_P} \Big)^2 + (\Phi_j - 1)^2 \Big] W_j U_j^+.$$

Again, it can be considered in exactly the same way as before because only linear terms of U_k^+'s are involved.

Combining the above statements about I and II, we have the desired conclusions. □

Remarks: (A) The above two results show that the positive quantity V decreases in expectation if it has a large value to start with. Hence on average, *large* values of V will be *reduced*. Combining with the theory of martingales, we can show that

$$\lim_{T \longrightarrow +\infty} \frac{1}{T} V(X(T)) = 0,$$

i.e., the overall undercoverage has at most *sub-linear* growth. This leads to the result that, *in the long run, the threat profile can be achieved exactly*.

(B) The number μ in Eqn (4) can be estimated. In fact,

$$\mu = 1 - \max_j \sqrt{\sum_{k \neq j} \Phi_k^2} \ (> 0).$$

Combining with the actual pre-computed values of the Z_{ij}'s, we can also estimate λ and most importantly, T_P^*. This can provide *practical* guidance of how long the pause time should be in order to ensure bounded global undercoverage. □

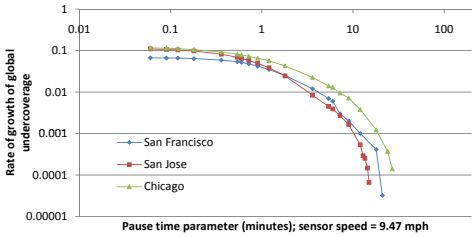

Figure 3. Rate of growth of global undercoverage time.

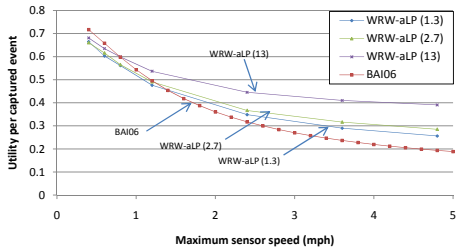

Figure 4. Per captured event utilities for San Francisco map.

7. Simulation Results (Residential Maps)

(A) Single sensor. We evaluate the performance of the WRW-aLP algorithm in Section 5 and the BAI algorithm [1]. We show the results for a real-life topology, namely a residential region in San Francisco. The map of the region is shown in Fig. 2(a). The region is of size 2000 feet by 2000 feet. It is divided into 8×8 cells, 51 of them are PoIs. The threat level of a cell is taken to be the number of residents in that cell, estimated from the LandScanTM 2004 database of population data. The sensing utility function is the concave function given in Section 3. We use the dynamic events defined in Section 4. The event durations are exponentially distributed with mean 13 minutes. Their interarrival times at a cell are exponentially distributed with mean value of 1 hour.

Figs. 2(c)–(f) show that progressively adding the a, L, and P features to WRW achieves actual coverage that increasingly match the threat profile in Fig. 2(b). WRW-aLP achieves the threat profile exactly when $P = 30$ minutes (Fig. 2(f)), verifying Remark (A) on Page 7. This important property of exact threat-based coverage is further investigated in Fig. 3 for WRW-aLP. The figure shows a log-log plot of the growth rate of the global undercoverage time (as a percentage of the sensor duty time) against the pause time parameter P, for San Francisco and also two other cities Chicago and San Jose. Note that as P increases, the growth rate decreases. The log-log scale of the plot amplifies the effect of *zero* rate of growth as P becomes larger than some *critical value*.

Fig. 4 shows the per-captured event utilities achieved for San Francisco by BAI and WRW-aLP (with varying P given in minutes) for different maximum speeds of the sensor. The results show that (1) pausing, which is possible with WRW-aLP and increases with P, can

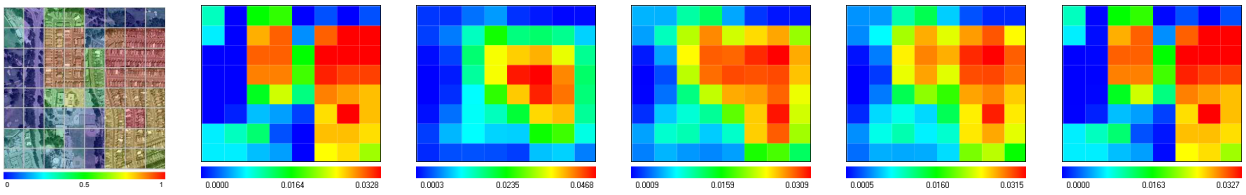

(a) SF Map (b) Threat Profile (c) WRW (d) WRW-a (e) WRW-aL (f) WRW-aLP

Figure 2. (a) Map of residential area in San Francisco. (b) Threat profile of the residential area. (c)–(f) Actual coverage achieved by progressive variants of WRW.

P (min)	Maximum sensor speed v (mph)						
	0.4	0.6	0.8	1.2	2.4	3.6	4.8
1.3	0.983	0.975	0.967	0.950	0.900	0.852	0.807
2.7	0.967	0.950	0.933	0.900	0.807	0.725	0.658
13	0.837	0.765	0.702	0.601	0.422	0.323	0.260

Table 1. Fraction of time sensor moves under WRW-aLP, for different maximum sensor speeds v and pause time parameters P.

significantly increase the utilities of the sensing results; and (2) the utilities do not increase, but rather decrease consistently, as the sensor speed increases. Table 1 verifies that the amount of sensor movement required by WRW-aLP is practically small.

(B) Multiple sensors. This set of experiments illustrates the performance of NC, GK, and MD coordination (Section 5.3) under different sensor densities. Figs. 5 and 6 show the pause fraction, unfairness, effective coverage, and the percentage deviation from the threat profile of WRW-aLP for San Francisco. The number of sensors is varied to be 2, 4, 7, 10 and 20. The pause fraction, unfairness, and effective coverage results for NC are similar to those for GK, and are omitted due to space constraints. Figs. 5(a) and 5(d) show that for GK and MD, the sensors move for a small fraction of the time only, verifying the *limited mobility* property of WRW-aLP. Notice from Fig. 5(b) that for GK, the unfairness is roughly reduced by half when we double the number of sensors. The same improvement applies for NC, showing that sensor coordination will likely not further improve the unfairness beyond independent deployment of multiple sensors. The conclusion (for unfairness) is similar to that in [6] for the low sensor density case. Figs. 6(a) and 6(b) show that in contrast to unfairness, the steady-state matching does *not* improve as we use more sensors for NC and GK. This is because having more sensors will exacerbate the coverage redundancy, thus hurting the matching metric. However, for the high sensor speed shown in Fig. 6, GK achieves a 60% improvement over NC in terms of the percentage deviation (of the global coverage profile) from the threat profile when the sensor density is high. It is because a sensor under GK can exploit the actual global coverage to better compensate for matching inaccuracies occurring from prior redundant coverage.

MD is an explicit attempt to solve the coverage redundancy problem and reduce the travel overhead. Fig. 5(f) shows that as the sensor speed is increased under higher sensor density, MD can achieve an effective

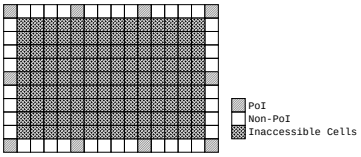

Figure 7. The ring topology.

coverage of more than 95%. Matching, however, may significantly worsen as there are more sensors. This is because as the sensor density increases, the size of a subregion assigned to a sensor becomes smaller. It then becomes more difficult to divide the surveillance area into subregions of adjacent cells having similar threat levels. The unfairness achieved by MD also worsens with higher sensor densities, as shown in Fig. 5(e). It is because under MD, the number of possible destination cells of a sensor is much reduced with more sensors. Hence, it is more likely for the sensor to stay at the current PoI. This results in higher unfairness as the sensor travels less, and has longer pause times as shown in Fig. 5(d).

In summary, NC and GK give better matching when the sensor speed or the sensor density is high, but they do so at the cost of significantly reduced effective coverage. On the other hand, MD may achieve a much better effective coverage, especially when the sensor density is high. This may lead to improved information capture globally, although higher-threat PoIs may no longer receive proportionally higher allocations of the constrained sensing resources.

8. Simulation Results (Ring Topology)

In this section, we discuss results for the ring topology in Fig. 7. The same topology is used in [1] to evaluate the BAI algorithm. The ring consists of a sequence of 50 cells, each of dimension 250 ft. $\times$ 250 ft., and 10 PoIs are uniformly placed on the ring. We use the same sensing utility function and the same type of dynamic events as in the previous section. The event durations and interarrival times are exponentially distributed with mean 800 s. For WRW-aLP, L is set to be 1500 ft.. The pause time parameter P is varied to control the amount of sensor movement. We use WRW-aLP(k) to denote the algorithm running with $P = k$ minutes.

To assess the impact of mobility, we also compare with *best-case* static coverage, in which each sensor's static position is chosen to give the best performance.

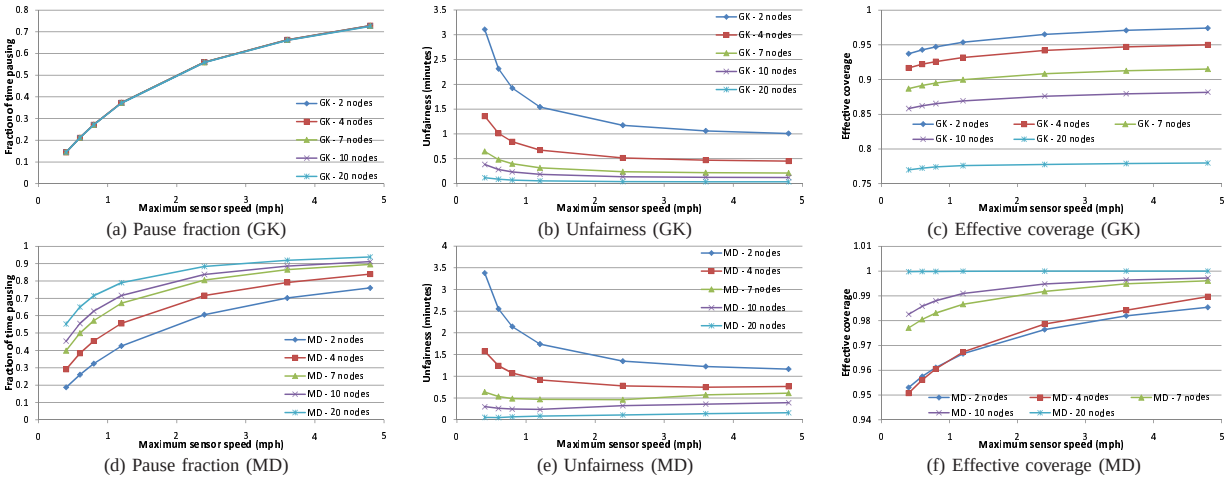

(a) Pause fraction (GK) (b) Unfairness (GK) (c) Effective coverage (GK)

(d) Pause fraction (MD) (e) Unfairness (MD) (f) Effective coverage (MD)

Figure 5. (a) Pause fraction, (b), unfairness, and (c) effective coverage of WRW-aLP for multiple sensors under GK in San Francisco. Corresponding results for MD are shown in (d), (e), and (f).

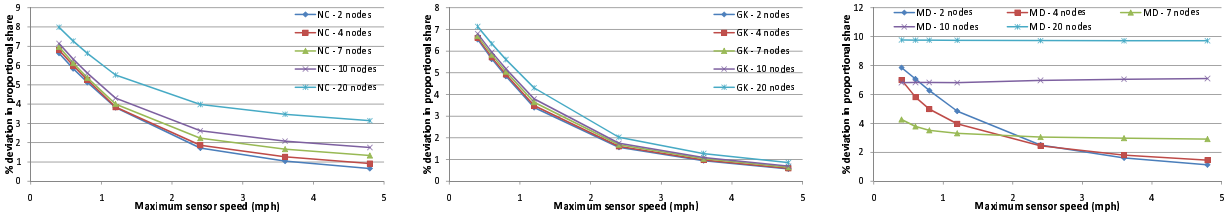

(a) NC (b) GK (c) MD

Figure 6. Percentage deviation of WRW-aLP global coverage profile from given threat profile, for multiple sensors under NC, GK, and MD in San Francisco.

We report results that are averages of at least 10 different runs. The standard deviations are very small (within 1% of the averages). *Hence, we do not report the standard deviations or the error bars.*

In this set of experiments, one sensor is used to cover the whole area. This corresponds to a severely constrained resource environment, in which the resource availability is only 10% (i.e., one sensor for 10 PoIs).

(A) Normalized Utility. We compare WRW-aLP (with different values of P), BAI, and static coverage by the normalized utility measure in Section 5.1. Fig. 8(a) plots the normalized utility achieved by the different algorithms as a function of the sensor's maximum speed. The following observations are in order:

(i) For static coverage, the sensor always stays at one PoI. Hence, it should be able to capture 10% of the events at their maximum utilities. From Fig. 8(a), however, notice that static coverage has a normalized utility of about 0.08, which is less than 0.1. This is because some of the events are short-lived, and do not last long enough for them to be captured at utility one.

(ii) From Fig. 8(a), notice that WRW-aLP(0) has similar performance as BAI. This is because $P = 0$ ensures that the sensor will continuously move between the PoIs, similar to the BAI algorithm. When P increases to one time unit, however, WRW-aLP(1.3) can perform significantly better than BAI. The results show that pausing at PoIs can improve the quality of the sensing by allowing the events to be measured for longer and therefore with higher confidence.

(iii) Static coverage is extremely efficient. Hence, while it is inherently unfair, it might perform the best purely from a utility standpoint. Fig. 8(a), however, shows that WRW-aLP always outperforms static coverage when the maximum speed exceeds a modest value. This is partly due to the concavity of the utility function. When the utility function is concave, much of the utility is obtained during the initial period of observing a new event. This encourages the sensor to occasionally move from one PoI to another in order to catch more new events, as long as the moving speed is not too low to make the travel overhead too high.

The normalized utility as a function of P and the maximum sensor speed is shown in Fig. 8(b). Notice that the binary function is concave in both arguments, showing that standard techniques can be applied to compute, for example, the optimal P for maximizing performance.

(B) Utility per captured event. Fig. 8(c) compares the utility per captured event for the different algorithms as a function of the maximum sensor speed. Notice from Fig. 8(c) that WRW-aLP(p), for $p \geq 1$, achieves a significantly higher per-captured event utility than BAI or WRW-aLP(0) for the same maximum speed. Hence, whereas the previous results show that a positive pause time with WRW-aLP achieves a higher normalized utility than a continuous movement algorithm such as BAI or WRW-aLP(0), these results further show that they do

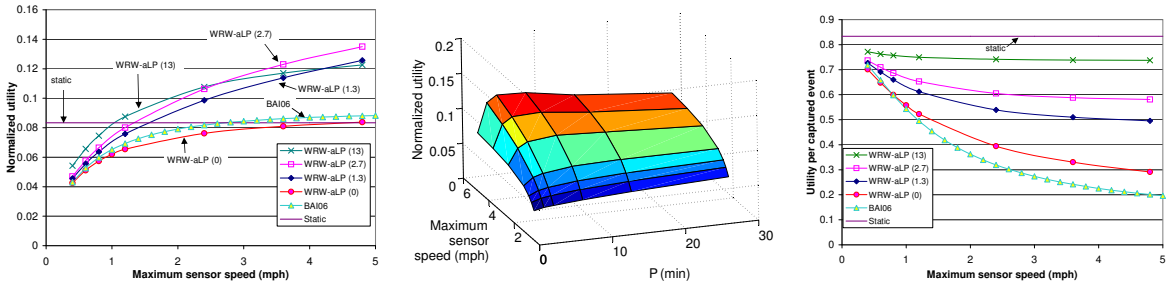

(a) (b) (c)

Figure 8. Normalized utility as a function of (a) maximum sensor speed, (b) different WRW-aLP parameters. (c) Utility per-captured event as a function of maximum sensor speed.

so not by capturing more events, but by improving the confidence of each captured event. For the continuous movement algorithms, notice also that the per-captured event utility drops significantly as the sensor speed increases. This shows that for sensing with a temporal dimension, too much mobility can be counterproductive.

Remark about multiple sensors: We omit detailed multiple-sensor results for the utility of captured information. We report, however, that when the number of sensors under NC doubles, the normalized utility roughly doubles also, while the improvement in terms of the utility per captured event is quite small. The observations indicate that the higher normalized utility achieved is mostly due to proportionately more events captured by the additional sensors. In addition, similar to the single-sensor case, WRW-aLP(p) for $p \geq 1$ consistently outperforms BAI in terms of both normalized utility and the utility per captured event. □

9. Conclusions

We have analyzed the costs/benefits of a stochastic mobile coverage algorithm for sensing tasks with a temporal dimension. Our results show that when sensor resources are constrained, a limited amount of sensor movement can provide threat-based fairness. Moreover, the quality of the sensing can be even better than best-case static coverage. Comparisons between WRW-aLP and BAI have provided insights on how mobility may impact our performance differently than simple event capture. We show that, in contrast to existing basic results on simple event capture, too much mobility may be counterproductive in our case, by reducing the accuracy of threat-based coverage and compromising the utility of the sensing results.

We have also studied how coordination between multiple sensors may impact performance when the sensor density is high so that reducing redundant coverage is an important goal. In terms of matching, GK and NC may perform significantly better than MD, and GK is in turn better than NC because the former can make use of the global coverage to better compensate for prior matching inaccuracies. MD has performance that is highly dependent on how the surveillance area is divided among the sensors. In general, however, MD may achieve much better effective coverage than either GK or NC. This is because under MD, coverage redundancy is eliminated and sensors have reduced travel overhead.

References

[1] N. Bisnik, A. Abouzeid, and V. Isler, "Stochastic event capture using mobile sensors subject to a quality metric," in *Proc. ACM MobiCom*, 2006.

[2] S. M. Brennan, A. M. Mielke, and D. C. Torney. *Radiation detection with distributed sensor networks. IEEE Computer*, August 2004.

[3] R. E. Lapp and H. L. Andrews. *Nuclear Radiation Physics*. Prentice-Hall, 1948.

[4] R. W. Lee and J. J. Kulesz, "A risk-based sensor placement methodology," Computational Sciences and Engineering Division, Oak Ridge National Lab, Tech. Rep., 2006.

[5] B. Liu, P. Brass, O. Dousse, P. Nain, and D. Towsley, "Mobility improves coverage of sensor networks," in *Proc. ACM MobiHoc*, 2005.

[6] C. Ma, D. K. Y. Yau, J. C. Chin, N. Rao, and M. Shankar, "Matching and fairness in threat-based mobile sensor coverage," *IEEE Trans. Mobile Computing*, to appear.

[7] A. Sundaresan, P. K. Varshney, and N. S. V. Rao. "Distributed detection of a nuclear radioactive source using fusion of correlated decisions," in *Fusion*, 2007.

[8] P. K. Varshney. *Distributed Detection and Data Fusion*. Springer-Verlag, 1997.

[9] Z. Vincze, D. Vass, R. Vida, A. Vidacs, and A. Telcs, "Adaptive sink mobility in event-driven multi-hop wireless sensor networks," in *Proc. InterSense*, 2006.

[10] Y.-C. Wang, C.-C. Hu, and Y.-C. Tseng, "Efficient deployment algorithms for ensuring coverage and connectivity of wireless sensor networks," in *Proc. IEEE WICON*, 2005.

[11] D. K. Y. Yau, N. K. Yip, C. Y. T. Ma, N. S. V. Rao, and M. Shankar, "Quality of Monitoring of Stochastic Events by Periodic and Proportional-Share Scheduling of Sensor Coverage," in *Proc. ACM CoNext*, 2008.

[12] Y. Zou and K. Chakrabarty, "Sensor deployment and target localization based on virtual forces," in *Proc. IEEE INFOCOM*, 2003.