Sensor Networks for the Detection and Tracking of Radiation and Other Threats in Cities

Annie H. Liu
Computer Science
Caltech
Pasadena, CA, USA
aliu@cs.caltech.edu

Julian J. Bunn
CACR
Caltech
Pasadena, CA, USA
Julian.Bunn@caltech.edu

K. Mani Chandy
Computer Science
Caltech
Pasadena, CA, USA
mani@cs.caltech.edu

ABSTRACT

This paper presents results from experiments, mathematical analysis, and simulations of a network of static and mobile sensors for detecting threats on city streets and in open areas such as parks. The paper focuses on the detection of nuclear radiation threats and shows how the analysis can be extended to other classes of threat. The paper evaluates algorithms that integrate methods of parametric and Bayesian statistics. A pure Bayesian approach is difficult because obtaining prior distributions on the large number of parameters is challenging. The results of analyses and simulations are compared against measurements made on a reduced scale testbed. A survey of background radiation in the city of Sacramento is used to quantify the efficacy of police patrols to detect threats. The paper also presents algorithms that optimize network parameters such as sensor placement.

Categories and Subject Descriptors

G.3 [**Probability and Statistics**]: Nonparametric Statistics; C.2 [**Computer-Communication Networks**]: Network Architecture and Design; I.6 [**Simulation and Modeling**]: Model Validation and Analysis; G.1.6 [**Optimization**]

1. INTRODUCTION

1.1 Overview

Countries around the world are concerned with detecting radiation, chemical and biological threats. This paper describes laboratory experiments, mathematical analyses, and simulations of systems that use static and mobile detectors coupled with algorithms for detecting and tracking radiation threats. The paper also suggests ways of extending the results to deal with other types of threats. The problem is described in the context of the following threat scenarios.

(Scenario 1) Static Threat in an Open Space: Security officials are given a tip that a threat is stored in an open space such as a park. Officers place a network of static sensors in the park to detect the presence of a threat (if one is present). In addition, officers may carry sensors with them as they search the park.

(Scenario 2) Static Threat on a Sidewalk: Security officers use static sensors and patrol cars to detect a possible threat on a sidewalk (for instance, in a garbage can).

(Scenario 3) Mobile Threat: A person carrying a threat in a park or on city streets must be interdicted.

1.2 Organization of the Paper

We address **Scenario 1** in Section 2 and 3. In Section 2 we present an integrated algorithm that first uses parametric statistics to get *a priori* estimates, which are then used in a Bayesian algorithm. The integrated algorithm gives better results than algorithms based exclusively on parametric statistics or purely on Bayesian approaches when prior distributions cannot be estimated accurately. Section 3 describes testbed experiments in a laboratory setup for the case of a static threat and static sensors. Laboratory measurements are compared against model-based simulations to validate our model and analysis.

Scenario 3 is addressed in Section 4 and 5. Section 4 presents algorithms for tracking a moving threat with a collection of stationary sensors in an open space. Section 5 describes testbed experiments in which a moving radiation source is introduced, and the section presents comparison of measurements with analytic results for detecting and tracking mobile threats.

In section 6, we address **Scenario 2** by presenting algorithms that help patrol cars detect threats in city streets. We use the City of Sacramento as an example, and analyze results for the given configuration of city streets and real background radiation measurements. We also presents strategies to be used by patrol cars to detect threats quickly.

Section 7 goes back to **Scenario 1** and presents a general detection confidence function for sensor network and provides an algorithms for preferred placement of static sensors. Section 8 discusses this work in the context of earlier work.

2. BAYESIAN, PARAMETRIC, AND INTEGRATED ALGORITHMS

2.1 The Problem

We are given a threat and a network of identical sensors. The network consists of a base station and M sensors indexed j where $j \in \{1, .., M\}$. Let S_j be the location of sensor j in a given 3D space and let S be the M-element vector of locations. The base station knows the location of

IPSN'11, April 12–14, 2011, Chicago, Illinois.

each sensor. Let $n_j(t)$ be the measurements made by sensor j in the time interval $[0, t]$. The sensors and the base station have synchronized clocks. At arbitrary points in time t sensor j sends a message containing (a) its measurement $n_j(t)$ and (b) the current time t to the base station.

The region in which sensors are placed generates noise signals. For the problem of detecting radiation threats, the noise consists largely of radiation generated by objects such as bricks, rocks and concrete; in addition, sensors themselves generate noise. When the search region is small and homogeneous we may assume that background noise is uniform throughout the region. When the region is heterogeneous, for example containing rocks with radioactive material, we may be unable to make reasonable *a priori* estimates without knowing something about the material in the area.

Based on the information received from the sensors, the base station (a) determines whether a threat is present in the region, and (b) if a threat is estimated to be present then the location and magnitude of the threat. We postpone, for the time being, the problem of identifying the type of threat.

The quality of the detection system is measured by the metrics employed for such problems:

1. The ROC (Receiver Operating Characteristic) curve [12] which shows the relationship between false positive rates (FPR) and true positive rates (TPR),
2. the DOCA (Distance Of the Closest Approach) curve that plots the probability of localizing within a certain error measured by the distance between the true threat location and the estimated location, and
3. the absolute difference between the estimated magnitude and true magnitude of the threat, if one is present.

Consider a single threat with intensity μ (we use intensity and magnitude interchangeably) at a location D in the space. The intensity is the rate at which the threat generates signals; for the case of radiation the intensity is the photons/second generated by the threat.

The measurement $n_j(t)$ made by sensor j in an interval from time 0 to time t is a random variable with a probability density function $h[\mu, D, S_j, t]$. Generally, if the distance between D and S_j is small then sensor j is likely to get stronger measurements; however, other factors such as obstacles between the source and sensor, or environmental factors such as wind may play a role.

A threat may be isotropic, transmitting signals equally in all directions, or it may be anisotropic. For an isotropic source with uniform background noise, the probability density function h depends only on the *distance* $|D - S_j|$, and not on their absolute values. By contrast, the density function h for an anisotropic source varies with the actual values of D and S_j. For the radiation detection problem the h function may vary with time when objects, such as cars, come between the source and the sensor, and thus impact the rate at which sensor readings change. For chemical and biological threats, the density function may change with time due to changes in the wind field.

The region in which sensors are placed generates noise which is unknown, and the sensors themselves generate noise which can be calibrated to some extent though noise characteristics may change. Given the locations S of sensors, the measurements $n(t)$, and the probability density function h, the problem is to determine whether sources are present and to estimate the sources' locations D and intensities μ. We restrict attention here to the presence of at most one threat in the search region.

2.2 Algorithms

Bayesian algorithms generally give good results when the prior estimates of unknown parameters are accurate. They give bad results when the prior estimates are poor and when the amount of signal data is too small to overcome the poor prior estimates. Since the goal of the system is to detect and interdict threats quickly, in most situations the system cannot wait for sensors to gather sufficient information.

A Bayesian algorithm requires prior distributions on the likelihood of a source being present, the intensity of the background, as well as the source's intensity and location. In addition, the algorithm requires prior distributions on the degree of anisotropy of the source and the objects preventing free flow of radiation from source to sensor. The algorithm must then evaluate the probability of a given set of measurements for each possible combination of source location, source intensity, and background intensity, location of occlusions, and type of anisotropy. The computation is intractable because the number of degrees of freedom is too large.

We present a heuristic that combines classical parametric statistical algorithms with Bayesian algorithms to deal with this problem. We first use parametric algorithms to estimate the values of some of the parameters; then we use distributions based on these estimates as the prior distributions in Bayesian algorithms. This heuristic could suffer from incorrect estimations from the parametric algorithm becoming accentuated by the Bayes algorithm; however, extensive comparisons between a Bayesian approach, a parametric approach, and the integrated heuristic show that the heuristic works well in different scenarios.

We begin by considering a single radiological isotropic point source placed in a region without occluding objects.

Bayesian Algorithm.

Bayesian algorithms compute an *a posteriori* probability distribution based on a given prior distribution and likelihood values calculated from measurements. Let θ be the vector of parameters we want to estimate. Assuming a prior PDF π_0 for θ, the posterior PDF of π is:

$$\pi(\theta) = L(\theta; n)\pi_0(\theta) \tag{1}$$

where L is the likelihood of observing $n = [n_1, n_2, ..., n_m]$ photons at detector $j = 1, 2, ..., m$ in time interval $[0, t]$, and θ is the vector of parameters. (Herafter we assume that the time interval t during which measurements are made is the same for all sensors and therefore omit t writing n_j rather than $n_j(t)$.) L is found by the detector measurement equation as:

$$L(\theta; n) = \prod_{j=1}^{m} f(n_j; \Lambda_j(\theta)) \tag{2}$$

In the case of radiation detection, sensors measure the number of photons striking the sensors. Photon emission is a Poisson process, therefore $f(n_j; \Lambda_j(\theta)) = \Lambda^{n_j} e^{-\Lambda_j}/n!$ is the Poisson probability mass function where $\Lambda(\theta)$ is the expected number of photons measured at sensor j in time t. θ includes parameters such as source strength μ, source location D, and expected number of photons Γ from the background in time

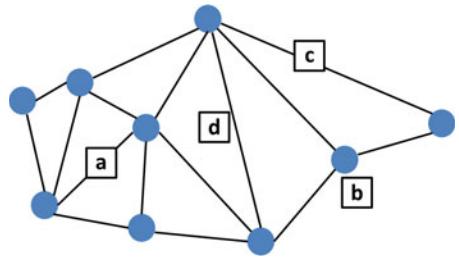

Figure 1: Snapshot of Delaunay triangulation as part of the kSigma routine. The partitioning is used to compute the {single, edge, triangle, cell} groups for data fusion. Four possible source positions relative to the sensor network are shown.

t. The likelihood function can be modified to incorporate additional information, such as photon energy [2].

In the scenario where the background noise is profiled beforehand, given *a priori* probability $\pi_0(\theta)$ and the measured data n, the algorithm computes the *a posteriori* probability that a source of intensity υ is located at position x. The posteriori probability that a source is present is the summed probability over all x and υ. If the algorithm must make a binary decision — a threat is present or is not present — then it decides that a threat is present if the posterior probability that a source is present exceeds a threshold where the threshold is determined by the tolerance for false positives.

K-Sigma Algorithm.

Next we present a parametric algorithm called the kSigma algorithm for estimating background noise level, source location and intensity, and the probability that a source is present; these estimates will be used in the integrated heuristic as the *a priori* estimates for the Bayesian algorithm.

Dynamic sensor grouping The signal to noise ratio (SNR) drops as the distance between sensors and source increases. Therefore simply combining measurements from all sensors may decrease the overall SNR [15]. Instead, we group measurements from proximate sensors by using a Delaunay triangulation, an efficient technique from computational geometry that partitions a space into triangles, while maximizing the minimum angle of the triangles. Figure 1 is an example of a triangulations.

Estimating background radiation We estimate the background noise based on the assumption that the background radiation is uniform within the region. The sensor flux drops off faster than $1/r^2$ due to absorption of photons in the air, where r is the distance from the radiation source. Therefore, sensors far from a source receive negligible flux from a threat. We estimate the background flux throughout the region by computing the average rate of photons received by these distant sensors. We do this by breaking the region up into quadrilateral cells (pairs of triangular cells that share a common edge), where each cell is identified by the sensors at its four vertices. The total photon count received in each cell is called its *cell count.* We use the average of the four sensors within the cell having the lowest cell count, after corrected by order statistics, as the background rate estimate Γ for all sensors.

Estimating whether a source is present We can coarsely characterize the relative source position in the sensor network by one of the following configurations: the source is close to (a) the center of a quadrilateral cell of sensors, (b) one of the sensors, (c) an edge between two sensors, and (d) the center of a triangle, as illustrated in Figure 1. If a source is near the center of a cell then the average counts measured by the sensors at the four corners of the cell is likely to be higher than the counts from the background. Likewise, if a source is very near a sensor then the photon count measured by that sensor is likely to be higher than the count from the background. So, we estimate the average counts from each *single* sensor, the pair of sensors along each *edge*, and the sensors at the corners of each *triangle* and each *cell.* We compute the number of standard deviations, called *kSigma values*, of the measured counts from the estimated counts if only the background were present; and we compute kSigma for each and every group (singleton, edge, triangle, quad) in the field.

$$kSigma = (N - \Gamma)/\sqrt{\Gamma} \tag{3}$$

where N is the group's aggregate radiation count in time t, and Γ is the estimated aggregate count in the group from only the background in time t. The aggregate radiation count in a group is a Poisson process: $\sqrt{\Gamma}$ is the standard deviation of the process. $kSigma$ is thus the number of standard deviatiosns that the measured group count has from that expected if only background radiation is present.

The kSigma values are then corrected using order statistics to account for the bias inherent in ordering the groups by aggregate counts. The corrected K for kSigma are calculated from the uncorrected values using the following factors for singletons, edges, triangles, and quads $K_1, \ldots K_4$ (for $N = 9$ sensors) as follows:

$$K_1 = 1.01K, K_2 = 1.1K, K_3 = 1.3K, K_4 = 1.5K$$

Estimating source intensity and location Let x be a d-dimensional vector representing a point in space. Consider an experiment conducted over an interval of duration T. Let Γ be the expected number of photons measured at a sensor generated by the background in this interval, and n_j the number of photons measured at sensor j in this interval. The flux at a point decreases roughly as the square of the distance ignoring absorption in the air. If there is a source with intensity μ present at location x, then the number of photons from the source measured at sensor j in the interval T is approximately $C \cdot \mu \cdot T/|S_j - x|^2$ where C is a constant of proportionality that depends on the sensitivity of the sensor. For convenience we introduce the sensor intensity variable $\upsilon_j(x)$ where: $\upsilon_j(x) = C \cdot \mu_j(x) \cdot T$.

If $n_j > \Gamma$ then the difference $n_j - \Gamma$ is attributed to a source at any location x and $\upsilon_j(x)$ specified by

$$n_j - \Gamma = \upsilon_j(x)/|S_j - x|^2$$

To estimate the source position, we make use of the sensor quad that exhibits the largest value of kSigma. We compute $\upsilon_j(x)$ for all sensors j in this quad, for all locations x. The location that is most consistent with the observed sensor readings is most likely to be the source location. To quantify this, we compute a variance estimate

$$L(x) = \sum_{j=1}^{m} (\upsilon_j(x) - \overline{\upsilon(x)})^2$$

where $\overline{\upsilon(x)}$ is the average $\upsilon_j(x)$ for j ranging over the vertices of the quad. We postulate that the likelihood of a

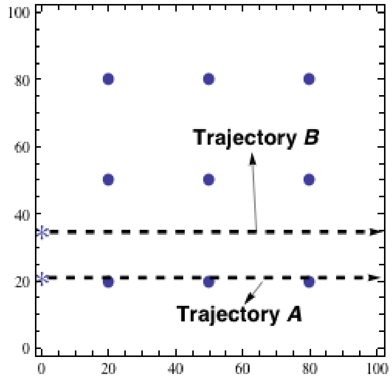

Figure 2: Simulation experiment setup. The circles mark the positions of the sensors. The trajectories are used in experiments with moving sources.

source at a point x is inversely proportional to this variance function. sThe source location is thus the x that maximizes

$$1/(1 + L(x))$$

Once the source location has been determined, the estimation of its intensity is straightforward: it is derived knowing the distance of each sensor from the source and the elapsed time T.

Integrated Algorithm.

Computation of the integrated algorithm is carried out in two steps — parametric estimation and Bayesian update. We first use the kSigma algorithm to estimate: the background rate Γ, the probability that a source is present at a location x, for each x, and the source intensity μ. Next, we use these estimates as *a priori* estimates in a Bayes algorithm.

Results of parametric estimation reduce the range of values that need to be considered by the Bayesian algorithm, and thus makes the Bayesian calculations tractable. We give results from several experiments evaluating the integrated heuristic in Section 2.3 and 3.2.

2.3 Simulation Results

All simulation parameters used in this paper are based on measurements taken by IPRL (Intelligent Portable Radiation Locater) sensors developed by *Smiths Detection*. IPRL is a high efficiency CdZnTe radiation sensor that is about the size of a camera (see Figure 5). The IPRL can be equipped with one (IPRL-1) to six crystals (IPRL-6).

Setup.

In the simulation, we simulated a background of 48 counts per second (cps) and a source of 1200 cps at one meter away. These values are based on measurements made with IPRL-6 sensors detecting a 1mCi Cesium-137 source. We simulated detection and localization in a 100 x 100 m field with nine of the sensors placed in a grid formation as in Figure 2. 20,000 simulation runs were carried out for each ROC curve, of which 10,000 were with a source randomly placed in the field and the other 10,000 without. We recorded the detection and localization results at T=9 and T=60 seconds. The data were fused and processed using four algorithms — (1) Bayes with correct priors ("bayes"), (2) the integrated heuristic, (3) parametric ("kSigma"), and (4) Bayes with incorrect priors that assumed a source five times stronger than was actually simulated ("bayes 5x"). For the Bayesian algorithms, we assumend an a prior probability that a source is present of 0.1.

Results.

We evaluated the integrated algorithm's performance using ROC and DOCA curves. Figure 3 and 4 shows the detection and localization results. We use the bayesian algorithm with correct priors as the upper bound and compare outputs from the other algorithms to it. At T=9, the integrated approach has similar detection performance as the parametric approach. There is a considerable improvement in localization performance: where the integrated algorithm localizes to within 10m $\approx$40% of the time as opposed to 20% with kSigma only. At T=60, the improvement is even clearer. The integrated algorithm performs nearly as well as the upper bound both in detection and localization. This jump in improvement over time is the characteristic of Bayesian type of algorithms — they can tolerate small errors in priors but when the errors get larger (as the ones generated by kSigma at T=9), the improvement by combining is small. In fact, all three algorithms perform considerably better than the Bayesian algorithm with large errors in priors.

3. TESTBED MODEL VALIDATION FOR STATIONARY SOURCE

3.1 Setup

A testbed, setup in the laboratory at IOS Pasadena, consists of six IPRL-6 sensors arranged in a 3x2 grid, as shown in Figure 5. An exempt 9.5 μCi Cesium-137 source can be placed anywhere within, or nearby, the grid. The spacings between sensors are chosen so that the setup will emulate the detection of a 1 mCi source in a 30x20 m field. The schematic is shown in Figure 6.

This particular setup was chosen based on on the number of sensors available, the type and strengths of sources, and the amount of time we were able to access the equipment and space. With these constraints in mind, only a small number of representative experiments were performed.

Experimental data were obtained for the three source positions in Figure 6. Position 1 has the source placed at 1 meter away from a sensor, at (x,y) coordinates (0,19). Position 2 is midway between two sensors at (15,10). Position 3 is in the middle of four sensors at (7.5,10). For each position, we ran 100 experiments of 120 seconds each. Between T=0 and T=60, the source was completely shielded (and therefore not detectable by the sensors). At T=60, the shielding was removed and the system observed till T=120. The data were analyzed with the Bayesian algorithm with correct priors on the source and background strength. The prior probability that a source is present was set at 0.1, and a detection threshold is at 0.5. Localization was computed after a detection was made.

3.2 Testbed Results

Limited by the number of experiments that can be repeated within a reasonable amount of time, the results are evaluated in terms of averaged detection confidence and localization error over the time span, shown in Figure 7, in-

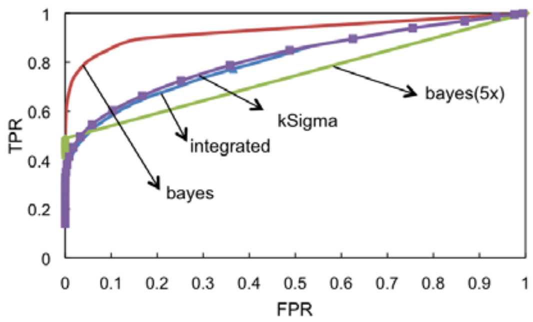

(a) T = 9

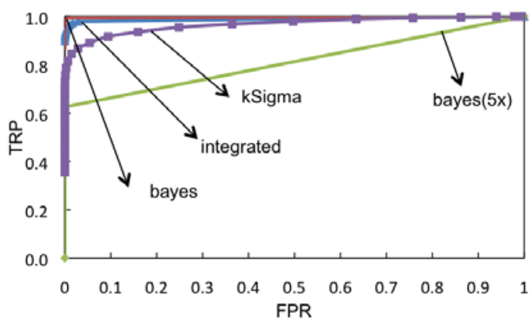

(b) T = 60

Figure 3: Detection performance compared using ROC snapshots for 20,000 simulation runs at two time intervals, T=9 and T=60 sec. The Bayesian algorithm was repeated with an incorrect source prior distribution (bayes 5x), as a comparison.

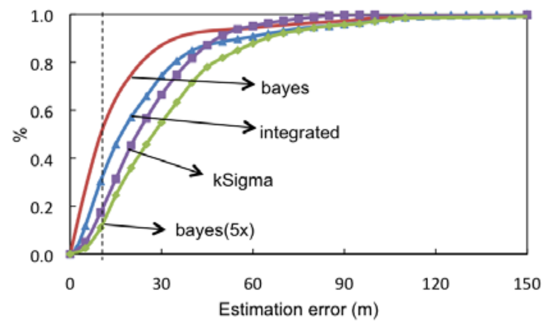

(a) T = 9

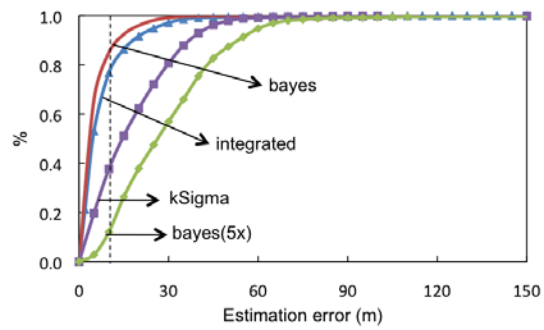

(b) T = 60

Figure 4: Localization performance compared using DOCA snapshots for 10,000 simulation runs at two time intervals, T=9 and T=60 secs. The Bayesian algorithm was repeated with an incorrect source prior distribution (bayes 5x), as a comparison.

stead of ROC and DOCA as used in Section 2.3. The averaged time to detect varied for the three source positions, therefore the localization sequences start at different times. In the first 60 seconds for all 100 runs, no false positives were generated. After the source was unshielded at T=60, the source was detected the fastest at Position 1 ($\approx$ 5 sec), followed by Position 2 (10 sec), and 3 (15 sec). This order is inverted in localization: the source was best localized when it was at the center of sensors. This result is not surprising as measurements from multiple sensors improve target localization. In these experiments, the localization errors converge at about 1 meter. Also observe the dip in localization error for Position 1 at $T = 65$. This is the result from high variance in error among the 100 runs when localization is attempted with too few data from a single sensor. These results not only validate the practicality of detection of a static weak source with networked sensors as many past work has shown, but also highlight the difference in algorithm performance for the three representative source-sensors layouts that have not been studied before.

4. TRACKING MOBILE SOURCES IN OPEN FIELDS

4.1 Modified Algorithm

The three algorithms described in Section 2 can be modified to accommodate mobility. In this study, we kept the sensors stationary for ease of comparison, but the results can be applied to both mobile source and mobile sensors.

Bayesian Algorithm.

The integration of the Bayesian algorithm with motion estimation models such as a *Kalman filter* is widely studied in object tracking. Here we implement a much simplified version of the Kalman filter. Instead of a sophisticated motion model, the filter assumes that the source can move in all directions with equal probability in the next second in a reasonable speed (e.g. 2 m/s if on foot) and allow for smooth weight transfer by redistributing the weight of the current source position estimate $\bar{x}$ to its neighbor y within a radius R using the following rules:

$$\pi(D = y) \leftarrow \pi(D = y) + \frac{Z}{R^2} \times \pi(D = \bar{x})$$

$$\pi(D = \bar{x}) \leftarrow (1 - Z) \times \pi(D = \bar{x})$$

where Z is the redistribution factor. R and Z can be dynamically adjusted according to the estimated speed and intensity of the source.

Parametric Algorithm.

When the threat source and sensors are stationary, measurements in the distant past and more recent past are equally useful; however, when the source or sensors move then measurements in the distant past are less relevant. The modified

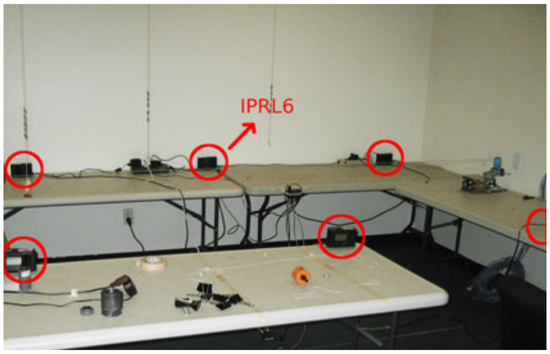

Figure 5: Testbed setup at IOS Pasadena. Six IPRL-6 are placed in two arrays of three, each circled out in red.

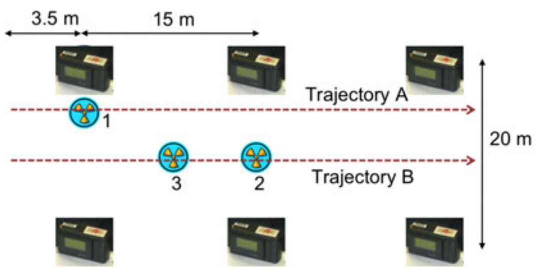

Figure 6: Laboratory testbed schematic, showing the six IPRL-6 sensors, the three source positions, and the two motion trajectories.

kSigma algorithm gives higher weight to more recent measurements by aging the relative weights of measurements exponentially with time. At a given elapsed time T, we compute an adjusted weighted sum W of all measurements, as follows:

$$W = \sum_{t=0}^{T} n(t)e^{\frac{t-T}{T_0}} \tag{4}$$

where T_0 is a decay constant. The value of T_0 is adjusted based on the estimated source and sensor speeds: the smaller the value, the more sensitive the algorithm is to quicker motion.

Integrated Algorithm.

The integrated algorithm remains the same as in the stationary scenario. However, the source and background estimates and the likelihood map are generated by kSigma algorithm with count aging.

4.2 Simulation Results

Using the same sensor grid setup as in Section 2.3, we simulated a 10 mCi Cesium-137 source traveling in a straight line along two trajectories, at a speed of 0.5 m/s (a slow human walking speed), as illustrated in Figure 2. Results compiled from 100 runs for each trajectory are shown in Figure 8. Results using the unmodified Bayesian algorithm for the no-motion assumption are plotted for reference. For both types of trajectory, all three algorithms are able to closely track a moderately strong source traveling at this slow speed. Moderate improvement is also observed from the integrated algorithm compared to the pure parametric algorithm.

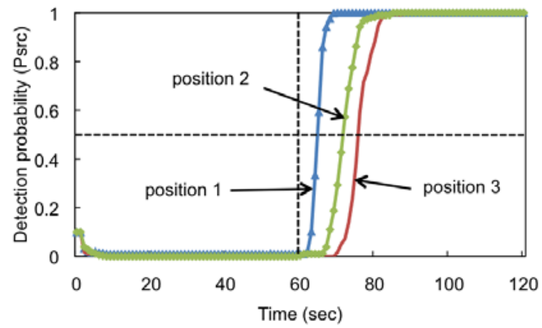

(a) Detection results.

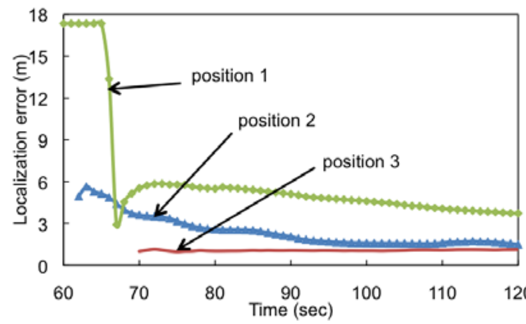

(b) Localization results.

Figure 7: Laboratory testbed results computed using a Bayesian filter with correct priors. The results are averaged over 100 test runs.

5. TESTBED MODEL VALIDATION FOR MOBILE SOURCES

5.1 Setup

Using the same setup as in Figure 6, we ran a set of experiments in which the source was moved at a constant velocity equivalent to 0.5 m/s along Trajectory A or B. At $T = 0$, the source starts moving from 3.5 meters left to the grid. At approximately $T = 70$ second, the source reaches the end of the grid and we stop the experiment. For each trajectory, 5 runs of 70 seconds each were recorded.

5.2 Testbed Results

Motion tracking results for the two trajectories are shown in Figure 9. The horizontal axis marks the actual location of the source when it is moving in the runs. The vertical axis shows the localization errors for the sequence of locations. For Trajectory A, three local minima at 0, 15, 30, occur when the source passes through those corresponding sensor locations at 1 m away. For Trajectory B, the fluctuations are smaller because the source is further away from all sensors. In both setups, all three algorithms demonstrated good capability of tracking a considerably weak 1mCi source. These results correspond closely to those generated in simulation and validates the feasibility of the algorithms in target tracking. Limited by low statistics (5 runs for each trajectory) no

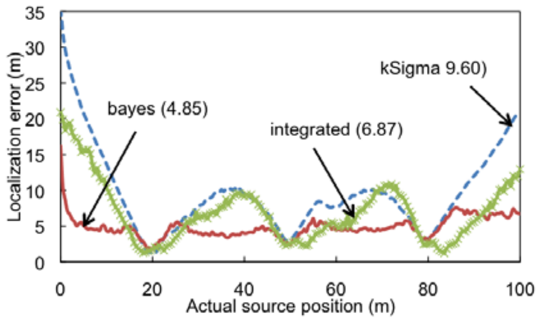

(a) Trajectory A.

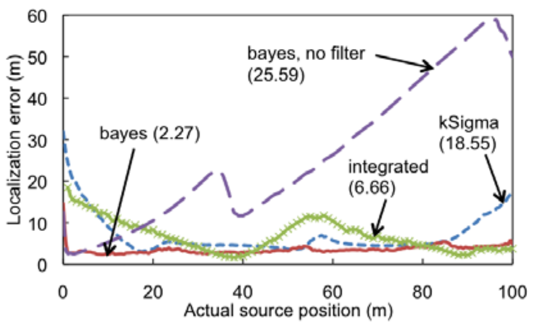

(b) Trajectory B.

Figure 8: Simulated source tracking of a 10mCi Cesium-137 source traveling at 0.5 m/s along two horizontal trajectories. Parameters $T_0 = 10$, $R = 10$, and $Z = 0.3$ are used. The averaged localization error over the total 200 seconds for all 100 runs is included in the parentheses next to each algorithm label.

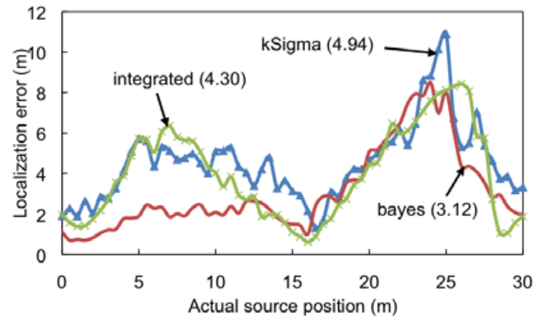

(a) Trajectory A.

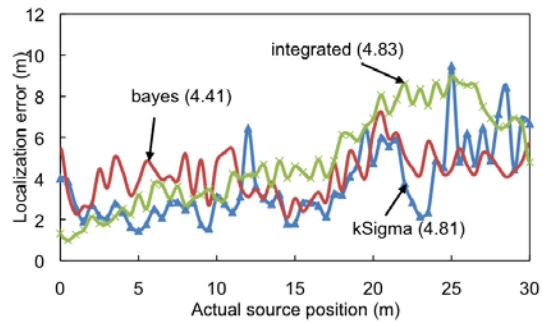

(b) Trajectory B.

Figure 9: Testbed source tracking of a 1mCi Cesium-137 source traveling at 0.5 m/s along two horizontal trajectories. Parameters $T_0 = 2$, $R = 10$, and $S = 0.3$ are used. The averaged position error over the total 70 seconds run for all 5 runs is included in the parentheses next to each algorithm label.

clear comparison can be made among the three algorithms.

6. THREATS DETECTION IN CITY STREETS

Threat detection in cities is a very challenging problem. In the past it has been studied in simulation as the problem of passively detecting a radiation source inside a city with a set of taxis equipped with sensors [4, 9]. In this study, we present an active detection approach using sensors mounted in police patrol cars, based on an extensive data set.

6.1 Sacramento City Radiation Map

Background radiation data in the downtown Sacramento city were collected by scientists from Lawrence Livermore National Laboratory. This data set consists of measurements made over a 5-day period of ≈ 6 hours on each day, measured by two large NaI scintillators placed in the rear of a van. Each sensor had dimensions 2x4x16 cm and an energy resolution of 6%. The van's GPS position, together with spectrum data, was recorded at one second intervals. The data covers about a 12x12 blocks area with multiple revisits for each point (at least 5 times) at different times. The data from the two sensors were combined, cleaned and corrected for calibration errors. For our study, an area of 5x5 block within the data set is selected. Inside this region, there are obvious fluctuations of background noise, ranging from a minimum of 500 cps to a maximum of 1200 cps. The data are visualized in Figure 10.

6.2 Simulated Roadside Detection

We simulated a set of detectors that travel the city streets at an arbitrary constant speed, respecting the one-way system where present. For each simulation, a source of strength 1mCi was placed at random somewhere along one of the streets, and at 4 meters from the center of the street. The intent was to simulate a dirty bomb or similar source placed, for example, in a backpack at the side of the street. For ease of analysis, we assumed an unshielded, isotropic radiation emission pattern. At $T = 0$, the sensors (i.e. patrol cars) initially placed at random positions in the grid started moving at a constant speed d.

The simulated counts were based on real noise profile in the Sacramento streets at a one-meter resolution. A simple threshold on the kSigma value was used to determine if the source was detected by the detector in question and if so, the simulation run was terminated, and a record made of the elapsed time to detection t. We assume that if a patrol car detects the probable presence of a source then the officers in the car search the nearby area and detect the source, with probability 1, if a source is present.

When a detector reached a junction, a decision was made as to which street the detector would start moving along.

(a) Full set of 30 hour data. Lighter blue indicates higher noise. The sub-region used in this study is circled in red.

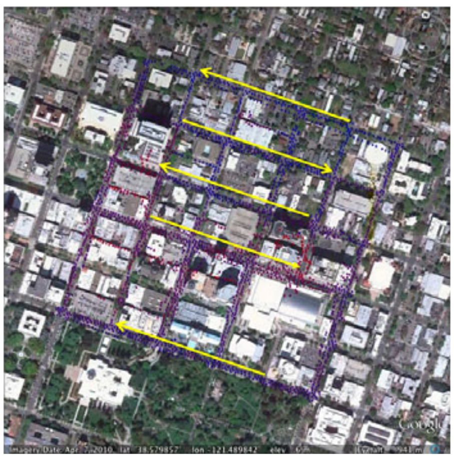

(b) Zoomed-in street layout. All E-W streets are one-way as marked by the arrows. All N-S streets are two-way.

Figure 10: Sacramento background radiation data visualized using GoogleEarth.

The decision took into account the one-way system, and we prevented detectors from doubling back, unless that was the only option. Otherwise, the detector's new street was assigned at random from the available (up to three) possibilities.

6.2.1 *Optimizing Sensor Number and Speed*

Figure 11 shows the average elapsed time to detection for sets of between one and 32 detectors in search of a 1mCi source. Two curves are displayed: one where the detectors move at 10 m/s ($\approx$ 22mph) and the other where they move at 25 m/s ($\approx$ 56mph). The improvement in detection time is marginal above half a dozen or so detectors, and detection times are shortest when the detectors travel at the higher speed.

There is, however, a tradeoff between speed and detection rate: the faster the detector moves, the smaller the effective sensor integration window is [19], and the less likely that the source will be detected. This is illustrated in Figure 12, which shows, for a set of eight detectors, the average detec-

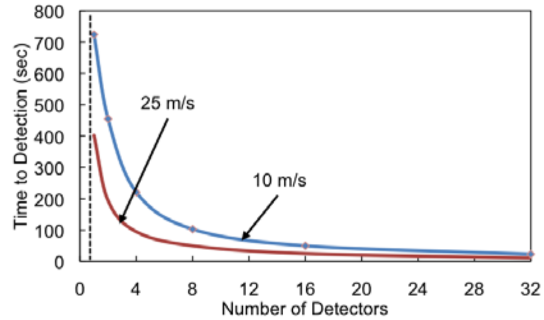

Figure 11: Time to detect as a function of the number of sensors at two speeds.

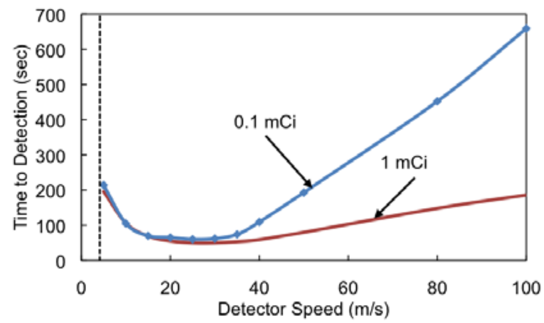

Figure 12: Time to detect as a function of sensor speed using two levels of source strength, for N=8 sensors.

tion time as a function of speed for sources of two strengths: 1 mCi and 0.1 mCi. We observe that if the detector speed is increased from 40 to 50 m/s the time to detect a source increases. Likewise, if the detector speed is decreased from 20 to 10 m/s the time to detect a source increases again. In fact, the speed that minimizes detection time in both cases is around 25 m/s (about 56mph). This interesting observation suggests that, in a situation where eight police patrol cars carrying detectors are dispatched in response to a tip-off that a dirty bomb is located somewhere in the city region, they should drive as fast as possible consistent with safety, up to a limit of 56mph, in order to detect the bomb in the shortest time possible time. They should not, however, go faster than 56 mph.

In a further study, we examined the detection sensitivity as a function of the source position. In other words, whereabouts in this 5x5 block city region are dirty bombs most likely to be detected quickly. This study involved 1,000,000 separate simulation runs, each using 32 detectors traveling at 25 m/s (an optimally realistic scenario). Detection times at each of the one-meter spaced grid positions in the region were averaged to obtain a sensitivity map, shown in Figure 13. In this map, the green band alongside each street (colored blue through red showing background intensity) indicates the sensitivity: the more intense the green, the more sensitive the detection capability if a source were placed at that location.

The result in Figure 13 shows that the *sensitivity* of a particular point depends predominantly on how often it is passed by, and not so much on the background level (given

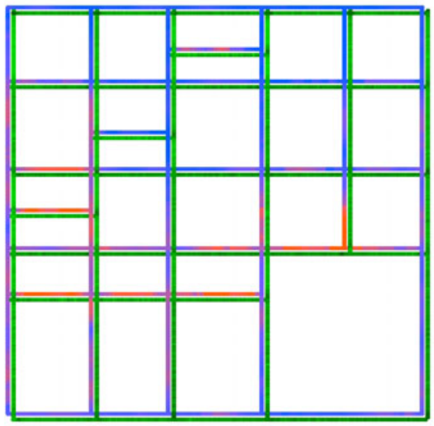

Figure 13: Detection sensitivity. The level of green indicates how long it will take for the network to detect a 1mCi source placed at that point. This is shown alongside the background radiation level (blue lowest, red highest) along each street.

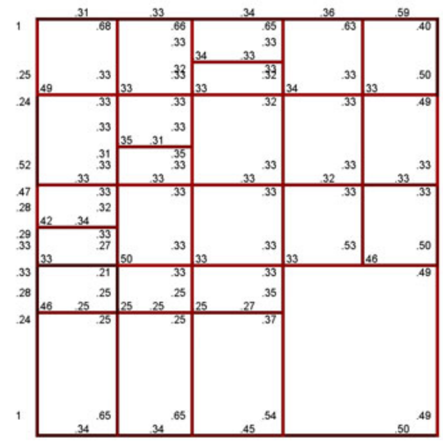

Figure 14: Street selection probabilities that minimize the range of traffic flows, as derived by a Genetic Algorithm. The numbers adjacent to each street at each junction show the probability that should be used when choosing that street when leaving the junction.

that the sensors are as large as the ones in this data set). However, this may change when the source to detect is heavily shielded: we intend to explore this possibility in future work.

6.2.2 *Foiling a Terrorist*

The detection problem essentially becomes a search problem when the source is bright enough. Then what is the best strategy for patrol cars to use in order to foil a terrorist? A random traveling strategy in which all exits from a junction are taken with equal probability is sub-optimal since it leaves some streets such as those in the South East corner of the city relatively poorly covered; detectors pass there infrequently simply because there are limited routes to get to those streets. We outline below how to calculate an optimal strategy. We make the standard game theoretic assumption of rational adversaries and compute a strategy that obtains the best outcome in the worst case.

Consider a graph in which each stretch of street, without turns, is modeled as an edge. Edges are directed. A junction of streets is represented by a vertex. We want each meter of street to be traveled with the same frequency, i.e., the time interval between repeated visits for each meter segment of street should be the same. This is because if a street segment were not traveled for a long time then a terrorist would game the system by placing a threat on that segment. Likewise, if streets were patrolled in a deterministic manner, such as first patrol East-West streets from North to South, then the terrorist could estimate the instant in the patrolling schedule with the greatest duration for the arrival of the next patrol car. Therefore we use a probabilistic strategy in which patrol cars make turns at junctions randomly, where the probability of a given turn is specified. A practical result of this strategy analysis will be instructions to patrol cars, for example "When you come to the junction of Avenue X and Street Y traveling East, go North 60% of the time, South 10% of the time, make a U-turn 5% of the time, and keep going East the rest of the time".

The analysis explores the flow of patrol cars on the city streets. Let $\mathit{flow}[j,k]$ be the number of patrol cars per hour that travel along street (edge) $[j,k]$, i.e., the street from junction j to junction k. For a two-way street between junctions j and k the total flow of patrol cars along the street is the sum $\mathit{flow}[j,k] + \mathit{flow}[k,j]$. For each junction j, conservation of cars gives the equation: $\sum_i \mathit{flow}[i,j] \;=\; \sum_k \mathit{flow}[j,k]$ Our goal is to maximize the minimum flow amongst all the streets; this can be formulated as a linear programming problem. We used a fast heuristic to solve this problem and we plan to compare run times of the heuristic against run times of a linear program solver. Figure 14 shows that the heuristic increased the minimum flow rate from 0.45 cars per unit time to 0.71 cars per unit time.

7. OPTIMIZING SENSOR PLACEMENT

In this section, we go back to **Scenario 1** and study the problem of placing sensors in a fixed size field to optimize system performance. This analysis can be extended to the detection of a wide range of anomalies.

7.1 Detection Function Definition

In many anomaly detection scenarios, the ability of an individual sensor to make a correct detection decision drops as the sensor moves away from the anomaly. This behavior can be described with a sensor detection characteristic function Φ that is unique to the sensor type. Φ denotes how the sensor detection probability TPR_{sensor} changes with varying distance r to the anomaly, and desired sensor false positive rate FPR_{sensor}. In the example of radiation detection, Φ is derived below:

$$TPR_{sensor} = \Phi(\lambda, \Gamma) = \frac{1}{2} - \frac{1}{2} erf\Big[\frac{Z_{TPR}}{\sqrt{2}}\Big]$$
$$= \frac{1}{2}\left(1 + erf\left[\frac{\sqrt{2}\lambda - 2\sqrt{\Gamma} \times erf^{-1}[FPR_{sensor}]}{2\sqrt{\Gamma + \lambda}}\right]\right)$$

λ and Γ are the expected signal and noise strength, with respectively, at the sensor. erf is the error function. λ is a function of the absolute signal strength μ, the position

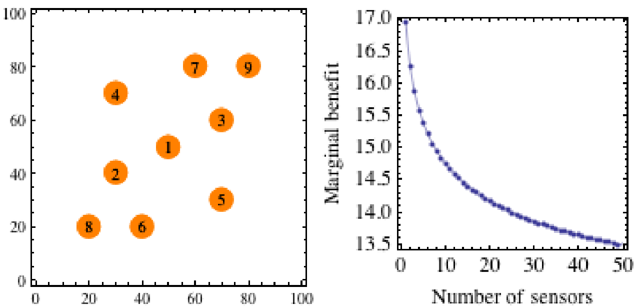

(a) Greedy placement of 9 IPRL-1 sensors. (b) Marginal benefit for each additional sensor.

Figure 15: Sensor placement for detection of a 1 mCi source with 9 sensors in 60 seconds. The result was computed using a greedy approach. Each sensor is labeled by the order of which it was added into the field.

of the source x, and the position of the sensor s. Γ is a function of the sensor position s. The system false positive rate FPR_{map} can be defined in two ways, 1) per sensor, and 2) per map. Using definition 1), Φ does not change with the number of sensors. However, using definition 2), FPR_{sensor} needs to decrease as the number of sensors in the map increases, because FPR_{map} has to stay constant. FPR_{sensor} can be adjusted as a function of N, the total number of sensors.

$$FPR_{sensor} = 1 - Exp\Big[\frac{\ln(1 - FPR_{map})}{N}\Big]$$

.

Therefore, the sensor's true positive function Φ drops as N increases. We denote Φ^N as the per sensor true positive function when there are a total of N sensors. We define our detection function F as the sum of true positives over the whole map while keeping the false positive rate constant.

$$F = TPR_{map} = \int_x \left(1 - \prod_{n=1}^{N}(1 - \Phi^N(\lambda(x), \Gamma(s)))\right) dx \quad (5)$$

7.2 Greedy Approximate Algorithm

The optimal sensor placements for detection will maximize Equation 5. We approximate the optimal solution using a greedy approach: placing the sensors one at a time with the goal of maximizing Equation 5 at each step. The greedy approach is guaranteed to perform at least a fraction $(1 - 1/e)$ of the optimal solution if the target function is monotone and *submodular*, i.e. the function has diminishing returns [18]. These conditions are formally stated below:

1. Submodularity (diminishing returns): $A \subseteq A' \subseteq V$ and $y \in V \setminus A'$, $F(A \cup y) - F(A) \geq F(A' \cup y) - F(A')$
2. Monotonicity: $A \subseteq A' \subseteq V$, $F(A') \geq F(A)$

Here, we prove that the detection function F is monotone and argue why it is also submodular. Observe that if $Log[F(N+1)] - Log[F(N)] \geq 0$, then $F(N+1) - F(N) \geq 0$

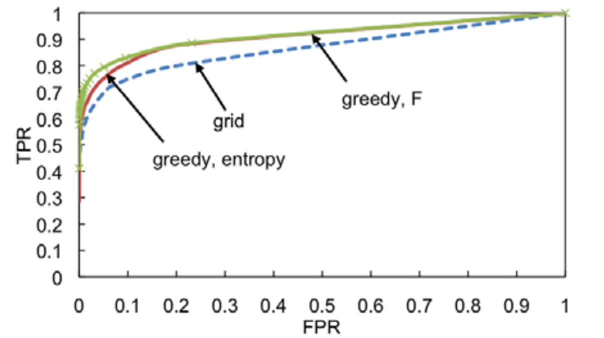

Figure 16: ROC comparison of three configurations of 9 IPRL-1 sensors for detection of a 1mCi source in T=60 seconds.

implies that F is monotone. We simplify $Log[F(N+1)] - Log[F(N)]$ by combining terms that involve the same sensor d_n.

$$\int_x \left(\sum_{n=1}^{N} Log\left[\frac{1 - \Phi^N(d_n)}{1 - \Phi^{N+1}(d_n)}\right] - Log\left[1 - \Phi^{N+1}(d_{N+1})\right]\right) dx \quad (6)$$

The first Log is the sum of a series of small negative values because the fraction is always less than, but close to, unity. The second Log term evaluates to a negative value because $1 - \Phi^{N+1}(X_y) \leq 0$. Depending on where the (N+1)th sensor d_{N+1} is placed, the two terms inside the integral do not always evaluate to a nonnegative value for all possible source positions x. However, observe that if the discretization of the distance is fine enough, that is, there exists an x such that $|x - d_{N+1}| \leq \epsilon$, then $Log[1 - \Phi^{N+1}(d_{N+1})] \to -\infty$. By adjusting ϵ, Equation 6 is never negative, thus F is monotone under the constraint that the per map false positive rate stays constant.

We argue that F is submodular from the following observations. If the field is fully covered with N sensors in even square meter, the difference between $F(N+1)$ and $F(N)$ is greater than the difference between $F(1)$ and $F(0)$ because in the former case the accumulated TPR is saturated for each position in the map. This argument is further supported by Figure 15(b). With the proof of monotonicity and the argument for submodularity, we claim that the greedy approach is guaranteed to produce a solution at least 63% as good as the optimal solution.

7.3 Results

Using the detection function F as the objective function, we computed the placement of nine sensors in a field of 100x100 meter with a desired $FPR_{map} = 0.01$. The source used in the calculation was a 1mCi Cesium-137 with an equivalent source strength $\mu = 200$ counts per second measured by an IPRL-1 at 1 meter distance. The expected background was uniform at 8 cps. The time allowed before a decision is made was 60 seconds. We restricted the minimum distance between two sensors to 10 meters. There are $\binom{11^2=121}{N}$ possible ways to place N sensors. The placement results for placing 9 sensors are shown in Figure 15(a). The marginal return for each additional sensor for up to 50 sensors is also plotted in Figure 15(b).

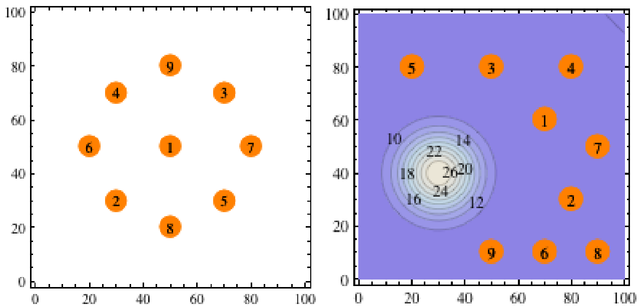

(a) Nine IPRL-1 sensor placement computed at T=600

(b) Nine IPRL-1 sensor placement computed with a non-uniform background, modeled as a bivariate Gaussian distribution.

Figure 17: Different layouts are computed as being optimal when the prior information changes.

The performance of this layout (Figure 15(a)) was evaluated in simulations using Bayesian statistics with correct priors. The resulting ROC curve is compared to those acquired from a simple grid layout (Figure 2), and a layout computed using entropy as the objective function [15]. As shown in Figure 16, the new layout in Figure 15(a) performs considerably better than the grid layout. There is also a slight improvement compared to previous results that used entropy as objective function. Note that the sensor model used here is IPRL-1. At T=60, the results would've been much better with IPRL-6 as shown in previous experiments in this paper.

7.4 Long Observation Time and Varying Background

The sensor's capability to make a correct decision (Φ) improves as T, the amount of time allowed before the decision is made. The objective function 5 therefore depends on T as well. Figure 17 shows two results when we compute the placement using different sets of prior information. In Figure 17(a), the amount of time allowed for detection has been increased from 60 to 600 seconds. The placement computed approaches a uniform grid layout. In Figure 17(b), we introduced a non-uniform background, modeled as a bivariate Gaussian distribution centered at coordinates (30,40), where the highest expected count was three times (25 cps) of the lowest count (8 cps). The placement of sensors, taking into account this prior information, first covers the area with lower noise variation before placing sensors in the noisier area. This result is not surprising since the marginal benefit of placing a sensor in an area of low SNR area is smaller than placing it in an area with high SNR.

7.5 Discussion

Given that the prior information can be computed in real time using a parametric method such as kSigma, it is possible to calculate a near-optimal sensor placement dynamically in order to optimize the detection capability of a sensor network. The sensor characteristics function Φ is general enough to be modified and applied to detection of anomalies other than radiation sources. By examining the output of the detection function F, one can compute the number of sensors required in an area to achieve a certain detection confidence within a certain amount of time, without the need for extensive simulation. We intend to investigate whether this placement method also optimizes localization performance.

8. RELATED WORK

Different detection scenarios require different sensor deployment strategies. Large portal style of sensors are ideal at places with specific entry and exit points such as airports and hospitals [28]; distributed sensor network (DSN) on the other hand, is a lot more flexible in configurations and easier to deploy in urban setting [1, 19]. The challenge in DSN is then to fuse data from multiple sensors either at a centralized server or distributed nodes [26].

Many detection algorithms were presented in the past. This includes deterministic solutions such as inverse-law inference [5], Maximum Likelihood Estimator [7], or probabilistic solutions such as 2-dimensional least squares fitting (LS) [10, 8], sequential probability testing (SPRT) [11, 22, 21], Bayesian posterior estimation [3, 27, 17], Extended Kalman Filter and its variants [8]. Of those algorithms, detection and localization are often coupled. In fact, robust localization leads to improvement in detection accuracy [16, 23, 20]. On comparison of the two groups of algorithms, deterministic solutions can be computed rapidly but they do not have the same level of flexibility or accuracy as probabilistic methods. Probabilistic solutions allow for complex sensor and environment models and often produce better estimations. However, because the algorithm needs to sweep a large probability space constructed from prior information that is at time missing, the actual feasibility in deployment is questionable. Our work showed that by combining the two methods carefully, one can achieve noticeable improvement in performance while significantly reducing the computation.

The principle behind all algorithms for anomaly detection is to distinguish signal from noise [13]. By combining data from sensors at multiple locations, we hope to increase the overall SNR. However, this is not always the case as shown in [15, 20]. The only reliable way to increase SNR is to decrease the distance between sensors and the source(s), either by deploying a large number of sensors, or by physically moving the sensors towards the source. Ristic et. al. looked at applying information theory to dynamically redeploy mobile sensors to detect strong sources in a low background environment [24, 25]. Similar techniques were used in [6, 14] to efficiently monitor spatial phenomena such as background radiation and temperature. The sensor placement problem under limited resources was first introduced in [15] in the context of radiation detection. In this paper, we extend from that and present a system level generalized detection probability function for other threat scenarios and provided a proof for theoretical bound for the greedy solution.

Detection in a city is extremely challenging. Cheng et. al. looked at clustering local sensor data for detection with a network of taxi cabs equipped with radiation sensors [4]. Hochbaum el. al. used a weighted concentrated alert approach in the same scenario in order to optimize the detection rate while minimizing the false positive rate [9]. Apart from stationary sources, Nemzek et. al. studied the problem of detecting a mobile source traveling on a street with an array of stationary sensors along the street [19]. However, these work were purely derived from computer simulations and lack real world implications. The study presented in this paper based all computations on real measurements collected over an extensive period of time and location. It further looked into how speed and the number of sensors affect the time to detect and devised strategies for fast detection.

9. ACKNOWLEDGMENTS

We would like to thank Bryan Cutler, Steve Foote, Andy Gooden, and Bozhil Makaveev of IOS, Pasadena for their assistance with performing the experiments described in the paper, and to thank IOS for the loan of their equipment. We also wish to thank David Campbell, David Trombino from the Lawrence Livermore National Laboratory for providing radiation measurements and analysis tools. Special thanks to Karl Nelson and Simon Labov from the Lawrence Livermore National Laboratory for their close guidance on analyzing radiation measurements. The Sacramento measurement work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Our research was supported in part by the National Science Foundation Cyber Physical Systems program.

10. REFERENCES

[1] S Brennan, A Mielke, and D Torney. Radiation detection with distributed sensor networks. *IEEE Computer*, Jan 2004.

[2] K M Chandy, J J Bunn, and A H Liu. Models and algorithms for radiation detection. *Modeling and Simulation Workshop for Homeland Security*, pages 1–6, Mar 2010.

[3] M Chandy, C Pilotto, and R McLean. Networked sensing systems for detecting people carrying radioactive material. *... International Conference on Networked Sensing Systems*, Jan 2008.

[4] J Cheng, M Xie, R Chen, and F Roberts. A mobile sensor network for the surveillance of nuclear materials in metropolitan areas. *DIMACS Technical Report*, 2009.

[5] J Chin, D Yau, N Rao, Y Yang, and C Ma. Accurate localization of low-level radioactive source under noise and measurement errors. *Proceedings of the 6th ACM conference on Embedded network sensor*, Jan 2008.

[6] R Cortez, H Tanner, R Lumia, and C Abdallah. Information surfing for radiation map building. *research.me.udel.edu*, 2010.

[7] A Gunatilaka, B Ristic, and R Gailis. On localisation of a radiological point source. *Information, Decision and Control*, Jan 2007.

[8] A Gunatilaka, B Ristic, and R Gailis. Radiological source localisation. *dspace.dsto.defence.gov.au*, Jan 2007.

[9] D Hochbaum and B Fishbain. Nuclear threat detection with mobile distributed sensor networks. *Annals of Operations Research*, Jan 2009.

[10] J Howse, L Ticknor, and K Muske. Least squares estimation techniques for position tracking of radioactive sources. *Automatica*, Jan 2001.

[11] K Jarman, L Smith, and D Carlson. Sequential probability ratio test for long-term radiation monitoring. *2003 IEEE Nuclear Science Symposium Conference*, Jan 2003.

[12] S Kay. Fundamentals of statistical signal processing, volume 2: Detection theory. *Lavoisier*, 2, Jan 1998.

[13] Glenn F. Knoll. Radiation detection and measurement. page 754, Jan 1989.

[14] A Krause, A Singh, and C Guestrin. Near-optimal sensor placements in gaussian processes: Theory, efficient algorithms and empirical studies. *The Journal of Machine Learning Research*, Jan 2008.

[15] A Liu, J Bunn, and K Chandy. An analysis of data fusion for radiation detection and localization. *International Conference on Information Fusion*, 2010.

[16] M Morelande, B Ristic, and A Gunatilaka. Detection and parameter estimation of multiple radioactive sources. *Information Fusion*, Jan 2007.

[17] M Morelande and A Skvortsov. Radiation field estimation using a gaussian mixture. *International Conference on Information Fusion*, 2009.

[18] G Nemhauser, L Wolsey, and M Fisher. An analysis of approximations for maximizing submodular set functions—i. *Mathematical Programming*.

[19] J S Dreicer D C Torney T T Warnock R J Nemzek. Distributed sensor networks for detection of mobile radioactive sources. *IEEE Transactions on Nuclear Science*, 5(4):1693–1699, Sep 2004.

[20] N Rao, J Chin, D Yau, and C Ma. Localization leads to improved distributed detection under non-smooth distributions. *International Conference on Information Fusion*, 2010.

[21] N Rao, C Glover, and M Shankar. Improved sprt detection using localization with application to radiation sources. *International Conference on Information Fusion*, Jan 2009.

[22] N Rao, M Shankar, J Chin, D Yau, S Srivathsan, SS Iyengar, Y Yang, and JC Hou. Identification of low-level point radiation sources using a sensor network. *Information Processing in Sensor Networks*, Jan 2008.

[23] N Rao, M Shankar, J Chin, D Yau, and Y Yang. Localization under random measurements with application to radiation sources. *International Conference on Information Fusion*, Jan 2008.

[24] B Ristic, A Gunatilaka, M Rutten, and M DSTO. An information gain driven search for a radioactive point source. *Information Fusion*, Jan 2007.

[25] B Ristic, M Morelande, and A Gunatilaka. A controlled search for radioactive point sources. *Information Fusion*, Jan 2008.

[26] R Waguespack, S Pellegrin, Millet B., and C G Wilson. Integrated system for wireless radiation detection and tracking. *IEEE Region 5 Technical Conference*, pages 1–3, Oct 2007.

[27] M Wu, A Liu, and M Chandy. Virtual environments for developing strategies for interdicting terrorists carrying dirty bombs. *International Society of Crisis Response and Management Conference*, pages 1–5, May 2008.

[28] K Ziock and W Goldstein. The lost source, varying backgrounds and why bigger may not be better. *AIP Conference Proceedings*, Jan 2002.