Models and Algorithms for Radiation Detection

K. Mani Chandy
Computer Science
California Institute of Technology
Email: mani@cs.caltech.edu

Julian J. Bunn
Center For Advanced Computing Research
California Institute of Technology
Email: Julian.Bunn@caltech.edu

Annie H. Liu
Computer Science
California Institute of Technology
Email: aliu@cs.caltech.edu

***Abstract*—The objective of this work is to develop systems, models and algorithms that help distinguish signatures of dangerous radiation material from background and other (e.g. medical) sources of radiation. The challenge is to do so rapidly and with an extremely low probability of false alarms. Radiological detection architectures will be deployed in a variety of settings such as monitoring political rallies and Coast Guard maritime boarding parties. This paper presents models and algorithms by which sensor networks and mobile agents collaborate to detect dangerous radiation sources. Extensive simulations using the algorithms have been carried out and some results are presented here. The paper explores ways in which the models can be extended from radiation detection to detecting other types of threats such as chemical threats.**

I. Introduction

Goals Homeland security deals with a variety of scenarios associated with with detecting and neutralizing threats [1] [3]. Scenarios include searching areas for hidden threats and maintaining continuous surveillance [2] [4] [7]. The problem in all scenarios is to distinguish signatures of dangerous material from background and other (e.g., medical) material [12]. Dangerous threats must be detected, located and identified rapidly with a low probability of false alarms. Detection systems are expected to minimize inconvenience to the public; for instance, security personnel can better prevent dangerous material from being carried into an area by requiring everybody to enter the area through narrow gates containing accurate sensors — but such restrictions are often unacceptable. Goals of threat detection systems are to reduce: (a) the times to detect, locate, identify, and neutralize dangerous material; (b) the rate of false alarms; (c) cost; (d) degree of public inconvenience and (e) problems in deployment and maintenance. This paper presents models and analytic techniques that help to tradeoff these conflicting goals in a systematic manner. The only threats this paper considers are radiation sources; however, the basic concepts apply to other types of threats.

The Basic Problem Radiation sensors register photons striking the sensor. More sensitive (and more expensive) sensors also register the energy of photons. Figure 1 shows the energy spectrums of radiation from a background and from Cesium, a dangerous source. We study the problem of using sensors that only record the number of photons striking the sensor and we later study how to exploit sensors that also register energy spectra.

Consider a single point source of radiation that is not shielded in an open field without obstacles. A sensor that is r meters away from the source detects photons at a rate that decreases as $1/r^2$ if there is no absorption of photons in the air. Absorption in the air results in a multiplicative $e^{-\theta r}$ term where θ is the rate of absorption per unit distance. For instance, the count rate for a detector 33 meters away from a source is more than a thousand times lower than that of a detector a meter away. Therefore, a key challenge for security personnel is to get sensors as close to dangerous sources as quickly as possible even though they don't know where the sources are.

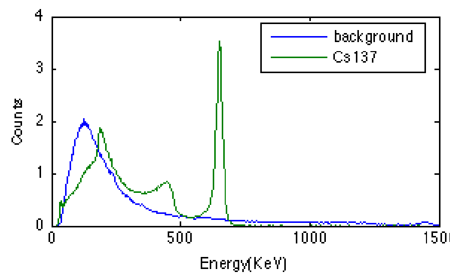

Fig. 1. Energy spectrum of some typical background and Cs137

Exacerbating Difficulties The problem is exacerbated by the following difficulties:

1) The source may be shielded by lead or other material.
2) Material such as cars and buildings absorb photons.
3) Background radiation is ubiquitous and may not be uniform across a region.
4) The source may be moving, and there may be more than one source.

Characteristics of Sensors We assume that sensors, or people carrying sensors, know their locations and can communicate with each other. We have carried out studies of the substantial advantage obtained by having sensors that are mobile [11]. Some sensors have limited directionality — we are given a probability distribution of the direction from which a photon came — and directionality also has substantial benefits. We first present results for stationary, non-directional sensors and consider mobility and directionality later.

II. A Simple Scenario

A simple set of experiments which may seem artificial, but provides insight, is as follows. The set consists of two groups of experiments: the first group has a source present and the second group has no source. In the first group of experiments a heavily-shielded dangerous radiation source is uncovered (the shield is removed) at time $t = 0$, and the experiment records whether the collection of sensors correctly determines: (a) the presence of the source in the field by some time $t = T$, (b) the location of the source at a later time $t = T'$ and (c) the identity (isotope) of the source by time $t = T''$. The fraction of time that the system correctly detects the presence of the source is the estimate of the probability of a true positive. The experiments also measure the accuracy of location and isotope identification. In the second group of experiments the field has no source; the fraction of times that the system claims that a source is present is the estimate of the probability of a false positive.

The experimental setup is simple. A flat rectangular field has a collection of non-directional stationary sensors. The sensors can communicate with each other. The experiments determine the probabilities of false positives and true positives for different values of T and the location and identification accuracy as a function of T' and T''.

A. What we Learn from a Single Sensor: Key Tradeoffs

First, let's determine what a single sensor can detect and later consider what can be detected by collections of cooperating sensors. We begin with simple sensors that only record the total number of photons detected and later consider sensors that also record energy spectra.

The system issues an alert exactly when the number of photons detected in time T by the (single) sensor is greater than or equal to some threshold value B. The greater the value of B the lower the probability of false positives but also the lower the probability of true positives. Let's vary B to look at the tradeoff between the probabilities of false positives and true positives as shown by ROC (Receiver Operating Characteristic) curve for this simple model.

Notation Let $poisson(r, t, n)$ be the probability of n events in time t when events are generated in a Poisson manner at rate r:

$$poisson(r, t, n) = \frac{(rt)^n e^{-rt}}{n!}$$

We begin by comparing probabilities of only two alternatives: (i) no source is present or (ii) there are sources present that produce a *given* total detection rate of Λ photons from the source. A comparison between just these two alternatives is academic because we don't know Λ. The academic exercise will, however, give us insight into the problem.

Let *pFP* be the probability of a false positive given that there is no source present; this is the probability of B or more photons from the background of rate Γ in time T:

$$pFP = \sum_{n \geq B} poisson(\Gamma, T, n)$$

Let *pTP* be the probability of a true positive given that there is a source present; this is the probability of B or more photons from the background or the source in time T:

$$pTP = \sum_{n \geq B} poisson(\Gamma + \Lambda, T, n)$$

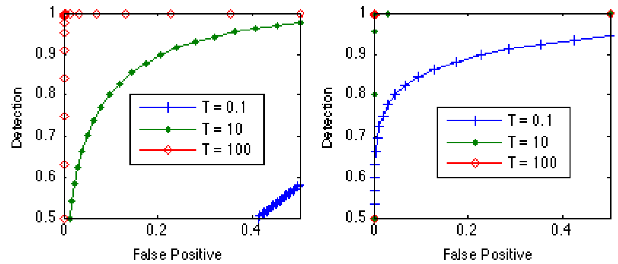

(a) $\Lambda = 2, \Gamma = 8$ (b) $\Lambda = 32, \Gamma = 8$

Fig. 2. ROC curves of the probability of detection.

Figure 2 shows *pTP* as a function of *pFP* by varying the threshold B with the time T, Γ and Λ fixed.

Tradeoff True Positives, False Positives, Time and Cost The figures in Figure 2 show that ROC curves improve (lower false positives and higher true positives) with increasing T — the time at which the system determines whether a source is present or absent. The figures also show that moving the source further away from the sensor or increasing the background intensity reduces the quality of the ROC curve. Figure 3 plots the time to detection versus the distance between the source and the sensor. In these graphs, the background rate is 8 photons per second and the source generates 200 photons per second when it is one meter away from the source. The probability of a false positive is less than 0.01 and the probability of a true positive is at least 0.99.

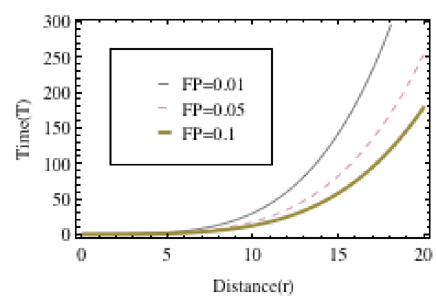

Fig. 3. Time to detect as a function of distance.

Next we determine asymptotic relationships between ROC curves, T, distance from the source and background intensities; these asymptotic relationships will give insight into how to design detection systems.

A Poisson distribution can be approximated, with error low enough for our purposes, by a normal distribution when the count N is 10 or more. Let $\sigma_{source}[T]$ be the standard deviation of the number of photons detected by a sensor in time T when a source is present. In this case photons arrive from both background and source. Since the combined background + source process is Poisson with intensity $\Lambda + \Gamma$:

$$\sigma_{source}[T] = \sqrt{(\Lambda + \Gamma)T}$$

Let $\sigma_{null}[T]$ be the standard deviation of the number of photons detected by a sensor in time T when no source is present. In this case photons arrive from the background:

$$\sigma_{null}[T] = \sqrt{\Gamma T}$$

Let *bFP* be the upper bound on the probability of false positives; for example, this values may be 0.01. Let z_{FP} be a positive real value such that the probability of values in a standard normal distribution exceeding z_{FP} is *bFP*. Then

$$B - \Gamma T \geq z_{FP}\sigma_{null}[T]$$

For example, when *bFP* = 0.01 then $z_{FP} = 2.33$. Similarly, let *bFN* be the upper bound on the probability of false negatives, and let the probability of values in a standard normal distribution being less than $-$*bFN* be z_{FN}. Then

$$(\Lambda + \Gamma)T - B \geq z_{FN}\sigma_{source}[T]$$

Consider the point on the ROC curve where $z_{FN} = z_{FP}$, and let's call this value z. Thus, the higher the value of z the better the ROC curve. From the above equations we get

$$z = \frac{\Lambda\sqrt{T}}{\sqrt{\Gamma} + \sqrt{\Lambda + \Gamma}}$$

For values of Λ much lower than the background intensity Γ:

$$\Lambda \ll \Gamma : \quad z \approx \frac{\Lambda\sqrt{T}}{2\sqrt{\Gamma}} \tag{1}$$

and for values of Λ much greater than the background intensity:

$$\Lambda \gg \Gamma : \quad z \approx \sqrt{\Lambda T} \tag{2}$$

We shall see that these equations provide insight into fundamental design tradeoffs.

Impact of Distance between the Sensor and the Source Consider a source that generates μ photons per unit time. When the source is r meters away from the sensor, the rate of photons from the source detected by the sensor is Λ where

$$\Lambda = \frac{A\mu e^{-\alpha r}}{r^2} \tag{3}$$

α is the photon absorption rate, per meter, in air and A is a proportionality constant which depends on factors such as the size of the sensor's detection crystal. Let's ignore absorption in the air. Then, from the previous equations

$$\frac{A\mu}{\Gamma} \ll r^2 : \quad z \approx \frac{A\mu\sqrt{T}}{2r^2\sqrt{\Gamma}} \tag{4}$$

This tells us that if we fix the bounds on the probabilities of false positives and true positives (i.e., fix the ROC curve) and thus fix z then the relationship between key design parameters when the source is far from the sensor is:

$$\frac{A\mu\sqrt{T}}{r^2\sqrt{\Gamma}} = K$$

where K is a constant. Therefore:

$$A^2\mu^2 T = K^2 r^4 \Gamma \tag{5}$$

If the distance r from the source to the sensor is doubled then the time T to make a decision must go up by a factor of 16, if everything else remains unchanged. If we want to keep T unchanged, then the sensitivity of the sensor, parameterized by the sensor area A, must go up by a factor of 4. If the background intensity is doubled then the time to make a decision is doubled as well if other parameters remain unchanged; if we want to make decisions just as quickly then A must increase by $\sqrt{2}$ or 1.4.

B. What we can Learn from Multiple Sensors

Let's look at how we detect a source faster and locate a source by using data from multiple sensors. Let's start with the simplest case of a square region with four sensors, one at each corner of the square. We discriminate between the hypotheses based on the counts of photons received by each of the sensors at the corners of the region. Since the processes by which photons are detected at each sensor are independent Poisson processes, the total across all sensors is also a Poisson process; therefore we declare that the null hypothesis holds if the total count across all sensors in time T is less than a threshold value B, and choose the alternative hypothesis if the total count exceeds or is equal to B. Thus, the analysis is exactly the same as for the single sensor.

Figure 4 shows the time T to detect a source as a function of the location $[x, y]$ of the source. In this diagram the sensors are at locations $[0, 0], [0, 30], [30, 30], [30, 0]$. The value of T for a point $[x, y]$ is the time required to discriminate between the hypotheses given that the sensor is either at point $[x, y]$ or does not exist. The time required for discrimination increases with points further away from sensors, with the worst case being at the midpoint $[15, 15]$ of the bounding box.

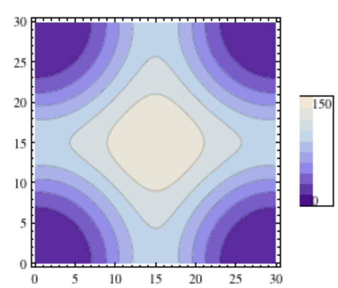

Fig. 4. Contour plot of time to detect on the 2D planes. Four detectors are placed on the corners of the bounding box.

C. Dealing with Unknown Backgrounds

Assume that sensors are laid in some sort of grid. We assume that a source, if one is present, is most likely to be surrounded by the four sensors with the highest total photon count. Likewise, we assume that the region surrounded by the four sensors with the lowest photon count is most likely to contain no source, and therefore the counts are due primarily to the background. Therefore, we estimate the background rate as the average of these four sensors. We then carry out the calculation to estimate whether a source is present or not present in the box with the highest photon count.

III. K-Sigma For Detection and Localization With Multiple Sensors

In this detection method the algorithm inspects the aggregate count of radiation measured by each detector for the elapsed experiment time. It then sums these counts in detector groups, which can contain one or more detectors. In our algorithm we use 1, 2 or 4 detectors. Given the measured (known) or deduced background rate Γ, the difference between the aggregate group count and Γ can be expressed as number, k, standard deviations (sigma) from the background. Detection of the source is announced if the value of k exceeds a threshold.

The choice of k determines, for a given physical arrangement of detectors, source strength, and prevailing background, the True Positive and False Positive rates, and these can be examined by using a ROC curve. An example is shown in Figure 5 where the detector type and arrangement is the same as described previously in this paper, a 1mCi source is positioned randomly in the field, the background rate was set to approximately 8 counts per second at each detector, and 10,000 experiments were simulated with a source randomly placed in the field, and 10,000 with no source.

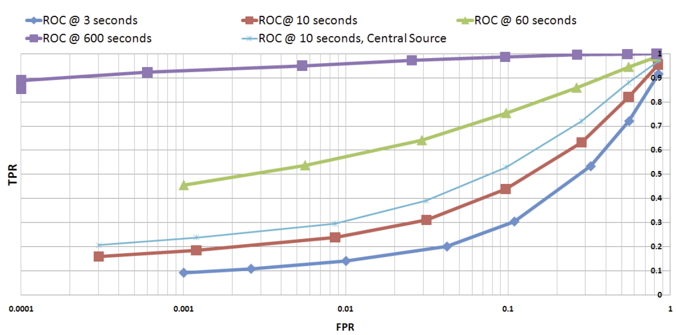

Fig. 5. ROC curves for the k.Sigma detection method, at elapsed times of 3, 10, 60 and 600 seconds. To achieve a False Positive Rate of 1 in 10,000 and a True Positive Rate greater than 90% the detection algorithm requires 600 seconds of elapsed time. Note the considerable improvement for the ROC curve when the source is confined to the bounding box formed by the outer detectors.

Once a source has been detected it is typically required to estimate its location. For a single detector, i, we can express the detected counts C_i at a time T in terms of the distance of the detector from the source, and the background rate:

$$C_i(T) = \frac{\Lambda_s c T}{(x_i - x_s)^2 + (y_i - y_s)^2} + \Gamma T \tag{6}$$

where Λ_s is the source radiation rate, c is a constant, Γ is the background rate at the detector, and (x_i, y_i), (x_s, y_s) are the locations of the detector and source, respectively. (For the physical parameters used in experiments of interest, and for simplicity, we ignore the attenuation length of the radiation, since at a value of 80 metres it introduces a negligible correction to the detected source rate for the distances we are interested in.) The three unknowns are the source location (two variables) and the constant term $Lambda_s c$. If the detected count rate is measured at three detectors whose locations are known, we obtain three simultaneous equations, which can be solved to obtain the source location (x_s, y_s) (for convenience, we have used Mathematica). In general, for N detectors, where $N \geq 4$, the problem is over-constrained, and we use the following heuristic approach to obtain the coordinates of the source.

First, we select the group of four detectors which show the greatest aggregate value of k.Sigma: at least one of these detectors must be the closest to the source. Taking each detector in this group in turn, we evaluate the following function, Υ_i for each grid position (x, y) in the field:

$$\Upsilon_i = (C_i/T - \Gamma)((x_i - x)^2 + (y_i - y)^2) \tag{7}$$

The value of Υ_i at this grid position should be the same for each detector i, if the source is located at that position. Thus a measure of the likelihood of the source being located at (x, y) is:

$$L(x, y) = \sum_{i=1}^{i=4} (\Upsilon_i(x, y) - \overline{\Upsilon})^2 \tag{8}$$

After calculating $L(x, y)$ for all grid positions, the most likely position of the source is the grid position with the maximum likelihood.

IV. Bayesian Estimate Method And Priors

The detection and localization of radiation source(s) can also be formulated as a parameter estimation problem [5] [6] [8]. 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:

$$\pi(\theta) = L(\theta; y)\pi_0(\theta) \tag{9}$$

where L is the likelihood of observing $y_1, y_2, ..., y_m$ photons at detector $j = 1, 2, ..., m$ given the vector of parameters θ. L is found by the detector measurement equation as:

$$L(\theta; y) = \prod_{j=1}^{m} f(y_j; \Lambda_j(\theta)) \tag{10}$$

$f(y; \Lambda(\theta)) = \Lambda^y e^{-\Lambda}/y!$ is the Poisson statistics function. Note that Λ is dependent of θ as illustrated in (3). L can be further modified to account for spectrum data.

By observing the posterior PDF $\pi(\theta)$, we can infer the probability that a source is present, the probability distribution of source location as well as source and background intensities.

Since the posterior PDF of the Bayes method has no closed-form solution, iterative approximation procedure is necessary. In each iteration, $\pi(\theta)$ is updated using (9).

In the rest of the section, we investigate how prior PDF affect time to detect in simulation. We restrict our focus to the scenario illustrated in Figure 6 with zero or one 1mCi Cs-137 source. Without loss of generality, we assume isotropic point stationary source, spherical ideal sensor and location/time-independent background model.

Set $\pi_0(p_{source} = true) = 0.1$, where p_{source} is the random variable that takes the probability of a source is present. We want to examine how π for p_{source} evolves over time with different priors π_0 for source and background intensity.

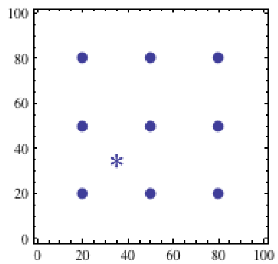

Fig. 6. Detector and source arrangement in a 100 x 100 m field. "*" marks the source location, if present.

Fix $\pi_0(\Gamma = 8) = 1$. Let $\pi_0(\Lambda = g) = 1/3$, where g can take the following 4 sets of 3 values. a) g = {0.1 1 10}. b) g = {0.2 2 20}. c) g = {0.05 0.5 5}, d) g = {exp(k), k ∈ [log(0.5) log(5)]}. The simulation is run and recorded at every 0.2 second intervals for 30 seconds. All 4 settings use the same set of records, therefore the comparison is exact (Figure 7).

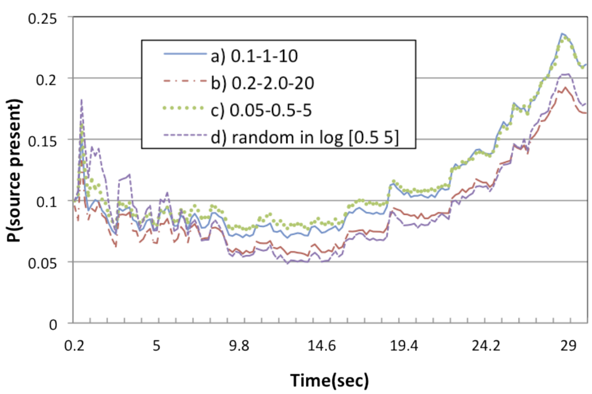

Fig. 7. $p(truePositive)$ as a function of time under four different prior settings over 30 seconds.

We can infer Λ by taking the weighted average of posterior PDF $\Sigma_i \pi(\Lambda = g_i) g_i$. The weighted average for Λ in setting d) is 1.1825 mCi, which is close to the true value. This result shows that the posterior distribution of parameters we are estimating converges to correct values over time. In general, the more knowledge we have about the threats, the faster the system can make decisions. Nonetheless, the algorithm converges eventually given enough time.

V. Mobile Sensors

Sensitivity of stationary detectors are limited by the $1/r^2$ drop-off in Λ. Since placing arbitrarily many detectors in a large field is not feasible given a budget constraint, the detectors simply have to move in order to achieve the desired speed and accuracy to detect, localize, and identify [10]. Figure 8 shows the time to detect as a function of detector speed in a 1000 x 1000 m field with 16 detectors in 4 x 4 grid formation. When the detectors are stationary, it takes on average of 28 minutes to detect a source in the center of 4 detectors. But when the detectors are allowed to move, even at only 1 m/s, the time to detect dramatically decreases to 4 minutes.

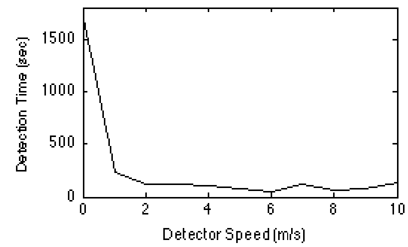

Fig. 8. Detection time as a function of detector speed.

This result uses Bayesian approximation described in Section IV with a simple heuristic that moves detectors towards the location that has the highest posterior probability over 2D space at each iteration at a constant speed. This algorithm uses no coordination between detectors. Note that the performance can be further improved with intelligent sensor placement algorithms such as information gain driven detector redeployment algorithm [9].

VI. Finite Horizon Algorithms: Optimizing Cost/Benefit Payoffs

In previous sections we saw how the success of source detection and localization algorithms depend strongly on the length of time available to accumulate data. Typically, the greater the elapsed time, the greater is the confidence in detection and localization. In real world situations time is at a premium. Thus time to detection is an important metric: the shorter, the better. Indeed, there may be little or no value in the detection of a source if this process takes too long. We can define a "payoff" value which quantifies the value of making a true positive detection of a source at a given time after the detection process begins. Similarly, we can define a "penalty" function which quantifies the negative effect of making a false positive detection. We combine these functions in a recursive Finite Horizon algorithm that derives an optimum set of detection thresholds - Claim Boundaries - as a function of time. These can be shown to maximize the performance and payoff of the chosen detection algorithm.

VII. Summary

We have described several strategies and algorithms that can be used to help in the detection and localisation of radiation sources in the presence of background. The fundamental variable is time: given sufficient time, sources can be detected and localised with confidence and accuracy. The problem is challenging because there is typically little time available, the sources may be weak and/or distant from the deployed detectors, and the background unknown. Detector mobility and the choice of a suitable prior probability of source presence ameliorate the detection/localization task somewhat, and an intelligent choice of detection thresholds as a function of time helps to maximize the performance of the system.

REFERENCES

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

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

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

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

[5] M Morelande, A Parkville, A Skvortsov, and A Fishermen's .... Radiation field estimation using a gaussian mixture. *International Conference on Information Fusion*.

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

[7] 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.

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

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

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

[11] 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.

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