An Analysis of Data Fusion For Radiation Detection and Localization

Annie H. Liu
Computer Science
California Institute of Technology
Pasadena, CA, U.S.A.
aliu@cs.caltech.edu

Julian J. Bunn
CACR
California Institute of Technology
Pasadena, CA, U.S.A.
Julian.Bunn@caltech.edu

K. Mani Chandy
Computer Science
California Institute of Technology
Pasadena, CA, U.S.A.
mani@cs.caltech.edu

Abstract – *This paper explores fundamental relations between critical parameters of distributed sensor networks (DSN) that detect and locate radiation sources. The paper presents mathematical analyses and Monte Carlo methods that help understand fundamental tradeoffs between the time to detect radiation sources, the probabilities of false positives and negatives, system cost, numbers of sensors, locations of sensors, mixes of sensors of different capabilities, benefits of data fusion, and communication load.*

Keywords: Radiation detection; distributed sensor network; data fusion; sensor density; threat localization; threat identification; surveillance

1 Introduction

Goals: In this paper we study distributed sensor networks (DSN) that detect and locate radiation sources. We develop formulae, asymptotic analyses, and Monte Carlo simulations to help understand basic tradeoffs between critical parameters such as the time to detect radiation sources, the probabilities of false positives and negatives, system cost, numbers of sensors, mixes of sensors of different capabilities, and communication load. Governments monitor borders, ports, and public gatherings to ensure that the public is not exposed to dangerous nuclear radiation; however, this paper does not deal with concepts of operations because our goal is to gain insight into key issues, and we do so by analyzing simple problems.

How much does data fusion help in detection? We study basic questions: How much does data fusion help in detecting dangerous radiation? And how should we quantify the benefit of fusion? If data fusion doesn't help much then a simple algorithm can be used for detection: each sensor independently claims that a threat is present when the radiation (number of photons) measured by a sensor exceeds a threshold. We suggest that an effective way of quantifying the benefit of fusion is the time gained by using fusion: How much quicker can a system using fusion detect a threat than a system in which sensors operate independently? We explore these issues using both mathematical analyses and Monte Carlo simulations.

How dense a sensor network is required? Another important design criterion is the density of sensors, i.e., the number of sensors per unit area. Is there a minimum density required for the system to detect threats in reasonable time? Costs of high density networks can be contained only if each sensor is inexpensive. Isotope identification requires sensors that monitor and analyze energy spectra, and such sensors are not cheap. This raises the question: should combinations of sensors with different capabilities be used? And what combinations are effective? We use mathematics and simulations to study these issues.

Where should static sensors be placed and how are threats localized? An obvious placement of sensors is in the form of a regular grid. We explore the question: are there better placements, and what algorithms help in producing better placements? We present algorithms and experimental results for determining the locations of threats, and we suggest an approach for identifying mixtures of isotopes using fewer expensive sensors.

2 Related Work

DSNs are ideal in situations where the area to cover is too large for traditional portal style detectors [21][1]. DSN have the advantage that they are portable, flexible in configuration, easier to deploy, and are more fault-tolerant [14]. Their main disadvantage comes as the need to deal with the spatial and temporal noise introduced by combining heterogeneous sensor data from various locations: overcoming this issue an active area of research. Indeed, data fusion is not always beneficial, as we will show in later sections.

The major challenge of detecting radiation in an urban environment is to distinguish signal (e.g. photons from radiation sources of interest) from noise (e.g.

background radiation from natural gamma ray emitting materials). The detection process is very often coupled with localization. Robust localization algorithms can enhance detection accuracy [12][17][7]. A collection of methods have been presented to solve the detection and localization problem, including deterministic solutions such as inverse-law inference [4], Maximum Likelihood Estimator [6], or probabilistic solutions such as 2-dimensional least squares fitting (LS) [8][7], sequential probability testing (SPRT) [9][16][15], Bayesian posterior estimation [3][20][11], Extended Kalman Filter and its variants [7]. Data fusion from multiple sensors is either incorporated into these algorithms or added as an additional step.

Detection and localization of radiation source(s) is essentially limited by the $1/r^2$ factor. The signal to noise ratio (SNR) can be dramatically improved if the sensors are close to the source. This observation suggests the usefulness of a mobile sensor network. The recent work on mobile sensor networks focuses mainly on directed information driven redeployment algorithms [18][19]. The same technique can be applied to fast sample the background radiation in a field [5]. It can also be extended to efficiently place sensors for monitoring purpose as we will show in the later section.

Our work builds on top of the existing research and goes towards a fundamental understanding on the characteristics of radiation sensor networks. These analyses will help provide insights into the factors affecting the design and deployment of successful DSN in real world situation.

3 Does Fusion Help For Detection?

Data fusion is clearly beneficial for locating and identifying a source, but it is not obvious how much it helps for detection, if at all. The argument lies in the fundamental measure for detection - the *signal to noise ratio* (SNR). Naively combining data (e.g. summing all data together without any filtering) improves signal strength but also adds in noise, which may decrease SNR. We thus ask the questions: How much does data fusion help to detect the presence of a source of radiation? Or, is independent analysis by each sensor in the network adequate? If independent analysis is almost as good as data fusion then communication between sensors can be simplified until detection takes place. For example, if background radiation has been characterized in the area under surveillance and if independent analysis is sufficient then the algorithm for each sensor is simple: Raise an alarm when its count of photons exceeds a threshold. (Or, if spectral analysis is used, raise an alarm based on the counts of photons in different energy windows.) If, however, data fusion is required then each sensor's measurements must be communicated to nodes of the sensor network where measurements are aggregated. If data fusion does help then a follow-up question is: How much does it help? And this question raises another one: How do we quantify the degree of help? We explore these issues next.

3.1 Classical Statistics Analysis

We consider the case where the background has been characterized and remains unchanged for the duration of the experiment. The experiment consists of a number of tests for detecting a radiation source in a given region which we assume is homogeneous without structures that absorb photons. On a given test either no radiation source is introduced into the region, or a shielded radiation source within the region is unshielded when the test starts. The problem is to detect the radiation source, if one is present, with probability at least $bPTP$ (lower bound on the probability of a true positive), and to claim erroneously that a radiation source is present, when one is not present, with probability at most $bPFP$ (upper bound on the probability of a false positive).

We first consider systems that use simple detectors that count the total number of photons but do not measure their energies. The null hypothesis is that no source is present. The alternate hypothesis is that there is at least one source present. The null hypothesis is rejected if and only if some function of the detectors' photon counts exceeds threshold(s); different functions and thresholds are used according to whether decisions are being made with or without data fusion.

Assume that the estimation of whether a source is present or not is made at a given fixed time T. Let $c[j]$ be the random variable which is the number of photons from the background detected at sensor j in time T. $c[j]$ is identically distributed for all sensors j. Let Γ and σ be the mean and standard deviation of this random variable.

We begin with an analysis of a situation where the field is a square with a sensor at each corner. We compare the case where each of the four sensors makes independent decisions with the case where the data from the four sensors is fused.

No fusion: The null hypothesis is rejected if the photon count for any sensor exceeds $\Gamma + K\sigma$ where K is a constant determined by the bounds, $bPTP$ and $bPFP$, on the probabilities of true positives and false positives.

With data fusion: The photon count data from each of the sensors can be fused in several ways. We explore a simple way which builds upon the idea used for the no-fusion case. The null hypothesis is rejected if and only if any of the following conditions described hold:

1. The photon count for any sensor exceeds $\Gamma + K_1\sigma$, where K_1 is a constant. This corresponds to the no-fusion case except that the threshold K_1 is different from (and is larger than) K. If the source is

near a corner of the square then the photon count for the sensor at that corner is likely to exceed this threshold. If, however, the source is far away from every corner then this threshold is unlikely to be exceeded.

2. We compute the average of the counts of each pair of neighboring sensors - there are four such pairs. If the average of any of the pairs exceeds $\Gamma + K_2\sigma$ then reject the null hypothesis. The reason for fusing data from adjacent pairs of sensors is to deal with the possibility of a source being along an edge of the square, near the middle of the edge. For example, if the length of the square is 20 meters, and the source is midway along an edge of the square then r^2 is 100 for the two nearest sensors and is 500 for the two farthest sensors, where r is the distance between the source and sensor.

3. Compute the average of the counts from all four sensors and reject the null hypothesis if the count exceeds $\Gamma + K_3\sigma$ where K_3 is a constant. This threshold deals with the possibility of a source near the center of the square.

The values of K_1, K_2 and K_3 are determined to satisfy the given bounds on the true positive and false positive probabilities. Note that by setting K_1 to K and K_2 and K_3 to infinity, the data fusion algorithm becomes the same as the independent algorithm. Therefore, we can ensure that the fusion algorithm is at least as good as the independent algorithm; we will, however, set the values of K_i so that the fusion algorithm performs better than the independent algorithm.

Analytical Results We derive formulae for the case of two detectors and use Monte Carlo methods for the case of four detectors. First consider two detectors a distance R apart. Let G_μ and F_μ be the cumulative distribution and probability mass functions for a Poisson random variable Y with parameter μ; thus $G_\mu(n)$ is the probability of $Y \leq n$, and $F_\mu(n)$ is the probability of $Y = n$. Let Γ be the expected count of photons from the background at a sensor in time T. The decision strategy is to declare that a source is present if and only if:

1. When data is not fused: the count of any sensor exceeds a threshold Q. The threshold Q is $\Gamma + K\sigma$ (rounding off to integer values).

2. When data is fused: the count of any sensor exceeds a threshold Q_1 or the sum of the counts of the two sensors exceeds a threshold Q_2 (or equivalently, the average of the two values exceeds $Q_2/2$).

Probability of declaring that a source is present: The number of photons received by a sensor in time T is a Poisson random variable with parameter μ where $\mu = \Gamma T$ when there is no source present because the sensor only records background photons which arrive at rate Γ, and $\mu = (\Gamma + \lambda)T$ when photons from radiation sources arrive at rate λ in addition to the background rate of Γ. The value of λ is determined by the location of the source and its strength.

Let $ps[\mu]$ be the probability that an algorithm declares that a source is present when the count at each sensor is a Poisson random variable with parameter μ. Since a false positive is a declaration that a source is present when there is only background radiation, and a true positive is a declaration that a source is present when there is both background and source radiation, we have the following relationship:

$$(PFP = ps[\Gamma]) \wedge (PTP = ps[\Gamma + \lambda])$$

No fusion case: In this case a source is not declared when $c[1]$ and $c[2]$ are both less than Q. Therefore:

$$ps[\mu] = 1 - G_\mu(Q)^2$$

A ROC (Receiver Operating Characteristic) curve plots $ps[\Gamma]$ on the x-axis and $ps[\Gamma + \lambda]$ on the y-axis for different values of Q and a given fixed λ.

Fusion case: In this case the algorithm claims a source is present when $c[1]$ or $c[2]$ exceeds Q_1 or their sum exceeds $Q2$. Therefore, the algorithm declares the presence of a source when (a) $c[1]$ is greater than or equal to Q_1 , or (b) $c[1]$ is less than Q_1 and $c[2]$ is greater than either $Q[2] - c[1]$ or Q_1. Since the two disjuncts are mutually exclusive their probabilities are summed:

$$ps[\mu] = (1 - G_\mu(Q_1)) + \int_0^{Q_1} F_\mu(c[1]) \cdot (1 - G_\mu(Min\{Q_2 - c[1], Q_1\}))\, \mathrm{d}c[1]$$

Figure 1 shows the ROC curves for the independent and fusion algorithms for T = 3, 15, and 50 seconds. The values of the thresholds, Q_1 and Q_2 could be optimized for each value of PFP and PTP; however, in these plots, the same values of Q_1 and Q_2 were used throughout a plot.The value of λ used in these plots corresponds to a 1 millicurie point source of Cesium placed along the line joining two sensors 20 meters apart, where the location of the source is equally likely to be any point on the line between 1 and 19 meters. Absorption in the air is ignored in this calculation. The graphs show that the fusion algorithm provides better results. For $T = 3$, the fusion algorithm behaves slightly worse than the independent algorithm for extremely high false probabilities; this is because the values of Q_i were not optimized for that case.

The analytic results suggest that fusion does help, but not by much. Why? Fusion helps because the values of Q_1 and Q_2 can be adjusted so that the fusion

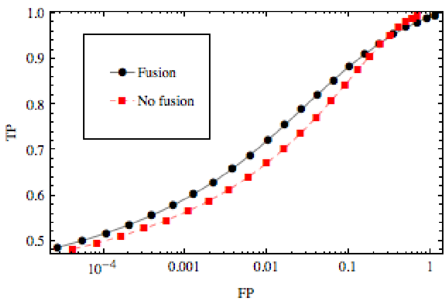

(a) $T = 3$

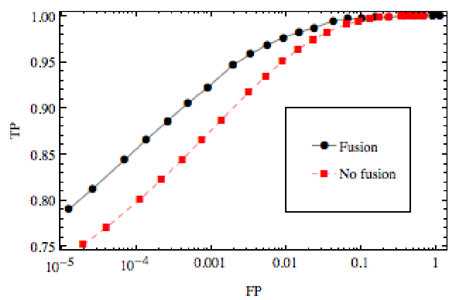

(b) $T = 15$

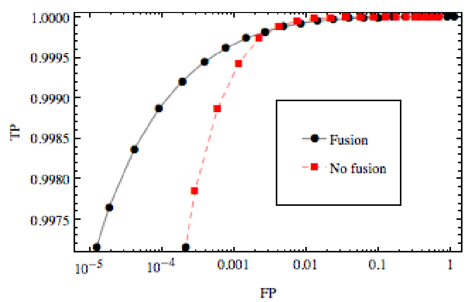

(c) $T = 50$

Figure 1: ROC curves of detection with two sensors with and without data fusion. $\Gamma = 8T$, $\Lambda = 200T$, $R = 20$. Note the x-axis is on log scale.

algorithm is the same as the independent algorithm for some values of these parameters and better for other values — fusion provides more degrees of freedom. The following example illustrates why fusion doesn't help much. When a source is within 6 meters of a sensor (and therefore more than 14 meters from the other sensor), the ratio of the counts of the two sensors is greater than $(14/6)^2 = 5.44$; so the distant sensor adds relatively little value. Thus for at least 12 of the 20 meters, data fusion helps but not by much. Moreover, the single sensor threshold, Q_1 has to be larger in the fusion algorithm than in the independent algorithm to ensure that the false alarm rates don't increase; thus, for some identical counts of the nearer sensor, the independent algorithm will detect the source whereas the fusion algorithm will not.

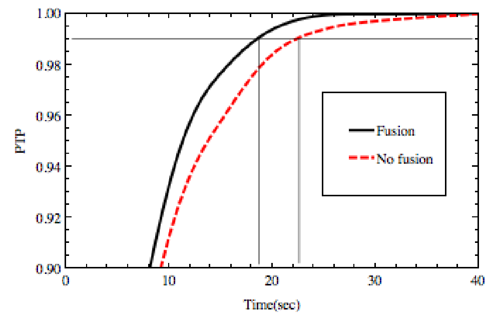

Figure 2: PTP as a function of time to detect while fixing $PFP = 0.01$ using optimized Q_1 and Q_2

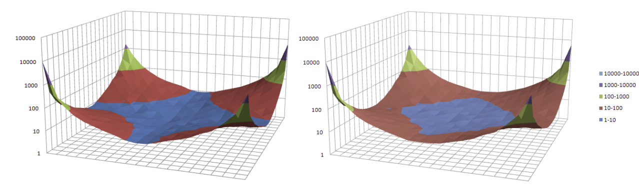

Figure 3: Source detection capability for four sensors arranged in a square, without (left) and with fusion (right).

We can measure the performance gain in terms of the amount of time. How much time do we save to reach a desired PTP and PFP by using fusion? Figure 2 shows that it takes approximately 19 and 23 seconds to reach $PTP = 0.99$ with and without fusion, respectively. In other words, the gain from fusion is 4 seconds.

Fusion with Four Sensors Next we present the case where four sensors are arranged in a square and a radiation source is placed within that square. Figure 3 shows the detection capability of the four sensor ensemble, with and without fusion, as a function of the source position. The vertical axis is the logarithm of the detection variable K; the higher the value the better. The points at which it is most difficult to detect a source are those along the edges and the middle of the square, and the plots show that fusion does help, but like the 2-sensor case, it doesn't help a great deal. Quantitatively, in this experiment, fusion decreases the time to detect by approximately 10%.

4 How Dense a Sensor Network is Required?

The results for data fusion algorithms are better than, but similar to, results for independent sensor algorithms. Therefore, we carry out an asymptotic analysis of sensor density requirements based on the independent sensor algorithm [2]. Let Γ be the rate at which photons from the background are measured by the sensor, and let λ be the rate from a source if one is present. The value of λ depends on the location and strength of the source, but for the time being let's assume that λ is given. Let μ_{null} and σ_{null} be the mean and standard deviation of the number of photons measured by the sensor in time T when there is no source present; and let μ_{source} and σ_{source} be the corresponding values when the source is present. Using Gaussian approximation of Poisson distribution with high mean, we have:

$$(\mu_{null} = \Gamma T) \ \wedge \ (\mu_{source} = (\Gamma + \lambda)T)$$

$$(\sigma_{null} = \sqrt{\Gamma T}) \ \wedge \ (\sigma_{source} = \sqrt{(\Gamma + \lambda)T})$$

The no-fusion algorithm rejects the null hypothesis (and claims that a source is present) when the count in time T exceeds a threshold Q. Recall that $bPFP$ is the given upper bound on the probability of false positives and $bPTP$ is the given lower bound on the probability of true positives. For this analysis assume that false positives and false negatives have equal weight, i.e., $bPTP = 1 - bPFP$. Let K be a positive real value such that the probability of values in a standard normal distribution exceeding z is $bPFP$. Then

$$Q - \Gamma T = K\sigma_{null}[T]$$

Since false positives and negatives have equal weight:

$$(\lambda + \Gamma)T - Q = K\sigma_{source}[T]$$

Summing these two equations and substituting for σ:

$$K = \frac{\lambda\sqrt{T}}{\sqrt{\Gamma} + \sqrt{\lambda + \Gamma}}$$

When the source intensity λ is much lower than the background intensity Γ:

$$\lambda \ll \Gamma \quad \Rightarrow \quad K \approx \frac{\lambda\sqrt{T}}{2\sqrt{\Gamma}} \qquad (1)$$

When a source is r meters away from the sensor, for large r, the rate of photons from the source detected by the sensor is λ, where

$$\lambda = \frac{\Lambda e^{-\alpha r}}{r^2} \qquad (2)$$

Λ is a constant determined by the strength of the source, and α is the photon absorption rate, per meter, in air. To a first approximation, Λ is the rate of photons per second measured by a sensor when the source is one meter away. From the previous equations

$$\frac{\Lambda}{\Gamma} \ll e^{\alpha r} r^2 : \quad \Rightarrow \quad T \approx \frac{4K^2 \Gamma r^4 e^{2\alpha r}}{\Lambda^2} \qquad (3)$$

We use this equation to analyze the following situations.

1. **Increasing distance between sensors**: When r is large, *the time T to make a decision increases more rapidly than the fourth power of r*. For example, with $\Lambda = 200$, $K = 3.2$, and $r = 20$ meters, then a 25% increase in r requires a 150% increase in T. Thus timely detection forces dense deployments of sensors.

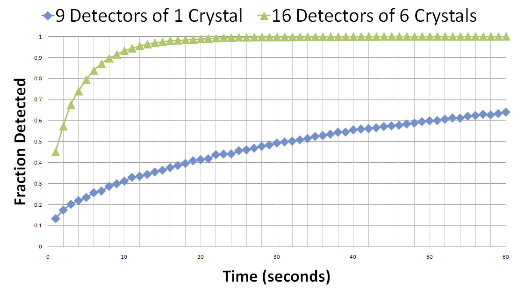

Figure 4: Increasing the density of detectors improves detection. The graph compares detection efficiency using 9 single-crystal detectors with that obtained using 16 six-crystal detectors, under otherwise identical conditions.

2. **Shielding radiation sources**: If the source is shielded so that the rate at which photons are detected is halved (and therefore Λ is halved), then the time T to detection is quadrupled.

3. **Increasing sensor density**: Let D be the sensor density: it is the surveillance area in meters-squared divided by the number of sensors. To a first approximation sensor density is proportional to $1/r^2$ where r is the average spacing between sensors laid in a grid. For large r we see that:
T increases more rapidly than $\Gamma/D^2\Lambda^2$.

Figure 4 shows the benefits of dense networks. A sparse network (9 sensors with a single crystal per sensor) won't detect threats rapidly, whereas dense networks (16 sensors with 6 crystals per sensor, increasing photon measurement rates by a factor of 6) can. Dense networks are expensive if each sensor is carried by a security officer; however, if sensors can be placed on street lights or in traffic cones, and the communication demands are low then dense networks are feasible.

Dense coverage and limited budgets suggest the use of inexpensive sensors; however, isotope identification requires more expensive sensors that can monitor energy spectra accurately. We discuss source localization and identification, and the impact of cost constraints on network design, later in the paper.

5 Where Should Sensors Be Placed?

Sensor placement is a difficult optimization problem because the objective functions, such as minimizing time to detection, are not convex and have multiple minima. Next, we compare placement of identical sensors on a uniform grid with a greedy algorithm and a genetic algorithm. We study a greedy algorithm for two reasons: (1) If sensors are added to the network incrementally — as for example, if security personnel equipped with sensors join teams that have already been deployed at a site — then placing additional sensors where they can do the most good may be more feasible than rearranging the locations of all the existing sensors. (2) The greedy algorithm is simple and we can prove a bound on the goodness of the algorithm compared with the optimum solution.

5.1 Greedy Algorithm

Assume that the surveillance area is gridded, and associate a random variable with each cell in the grid. The random variable takes on the value 1 if there is a source in that cell and 0 otherwise. The *a priori* probability of a source in each cell is given. In this calculation we assume that each cell has the same probability p of containing a source. Consider a point on the grid and assume that there is a sensor at distance d from it. Measurements by the sensor over some time interval T will allow an algorithm to compute the *posteriori* probability that the cell contains a source. The smaller the distance d the more certain we are about the presence or absence of a source at that point; equivalently the posteriori probability density will be heavier near 0 or 1, and thus the entropy associated with that random variable is reduced.

Assume that M sensors are deployed over the surveillance region. Let D be the vector of sensor locations with $D[j]$ being the cell of the j-th sensor. Let $f_i(D, p)$ be the *posteriori* probability of there being a source at cell i given sensor placements D and *a priori* probability p. We seek to compute D that minimizes the total entropy $F(D) =$

$$\sum_i f_i(D,p) \log f_i(D,p) + (1-f_i(D,p)) \log(1-f_i(D,p))$$

The greedy algorithm adds one sensor at a time, increasing the length of D by one. We determine the location of the first sensor to minimize $f(D[1])$, i.e., assuming that there is only one sensor in the field at location $D[1]$. The first security officer coming to a field, assuming that he or she is the only one, will move to location $D[1]$. Given the locations $D[j]$ of M sensors we determine the location of the $(M+1)$-th sensor as follows: We assume that $D[1], \ldots, D[M]$ are fixed and we determine the location l of the $(M+1)$-th sensor that minimizes $F(D+l)$. This corresponds to the $(M+1)$-th security officer moving to the location that best helps the existing deployment of the previous M officers.

Function F has two important properties that ensure the greedy algorithm is within $(1-1/e) \approx 63\%$ of the optimum [13]: (1) F is monotone, i.e., adding sensors doesn't make the objective function worse and (2) F is *submodular* [10], i.e., there is a decreasing marginal benefit from adding more sensors.

In the calculation, we use the following heuristic.

$$f_i(D,p) = Min\left\{(p + \sum_{m=1}^{M} K(d_{im})), 1\right\}$$

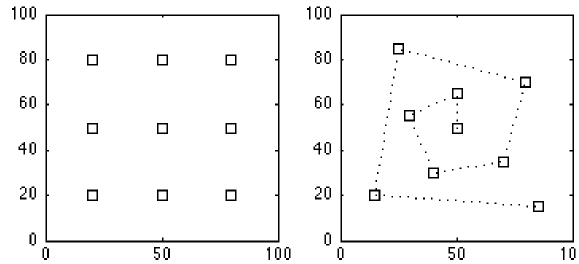

(a) Grid layout (b) Greedy entropy minimization layout. $K(d) = 2/d$

Figure 5: Sensor placement layouts. The dotted line connects positions in the order they are chosen by the greedy algorithm at each step. The order starts from the center.

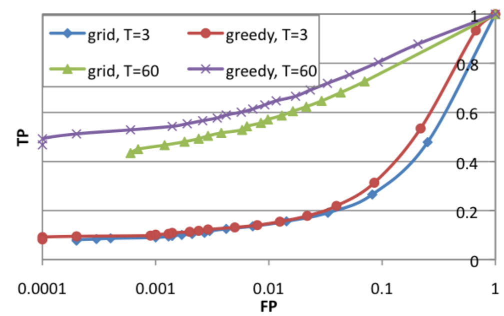

Figure 6: Comparison of greedy and grid sensor layout using ROC curves. The source strength is 200 cps at 1m. The background strength is 8 cps.

where K decreases as distance d.

The performance of this layout with 9 sensors in a 100 x 100 m field (Figure 5(b)) is compared against a grid layout (Figure 5(a)) in simulation. Each curve in Figure 6 is compiled from 10,000 runs of T seconds with a randomly placed source and 10,000 runs with no source. At $T = 60$ seconds, the new layout outperforms the grid layout by $\approx 5\%$.

5.2 Genetic Algorithm

Heuristics, such as genetic algorithms, are employed to get good solutions to difficult optimization problems. We compare the results of genetic algorithms with uniform grid layouts. In the experiment, the genetic algorithm (GA) uses a population of 400 chromosomes. Each chromosome encodes the coordinates of M sensors. Initially, the coordinates of the sensors are chosen at random. At each epoch, each chromosome determines the value of the objective function for its layout of sensors. The best chromosomes (those with largest values of the objective function) are mated, mutated and proceed to the next epoch. In these experiments, a heuristic was used to compute the objective function.

Figure 7 illustrates a typical result from the GA for

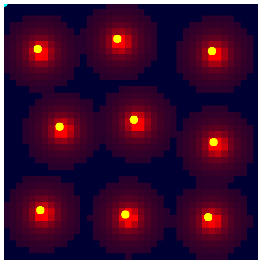

Figure 7: A candidate optimal layout for 9 sensors (shown as yellow discs), as suggested by the GA. The red shading indicates the sensitivity.detection sensitivity.

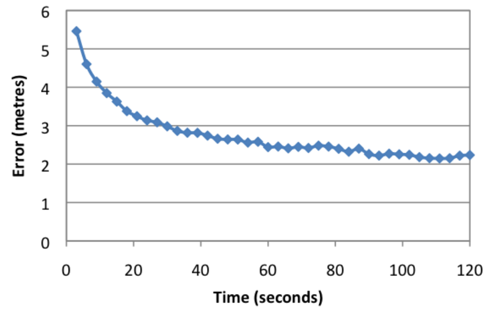

Figure 8: Localization error as a function of time using fusion with unknown background.

a set of 9 sensors in a field of 100 x 100m. The layout is equivalent to tesselation of the field by discs of radius proportional to the sensor sensitivity.

Summary Uniform grid layouts of sensors are not optimal. Simple greedy algorithms that minimize entropy do reasonably well, and these algorithms can be adopted when the number of security personnel available is not known for certain, and where security officers arrive at a scene incrementally, in small groups. Genetic algorithms and other heuristics do better than uniform grids, and we are exploring these heuristics further.

6 Localization and Identification

The location of the source can be estimated with or without data fusion. An algorithm in which each sensor operates independently estimates the source location to be at or near the sensor with the maximum counts (or the maximum standard deviation departure from background). Data fusion algorithms can aggregate data from multiple sensors to triangulate or otherwise estimate source locations more accurately.

The count $c_i(T)$ of photons at a single detector i in an interval of duration T is $(\Gamma+\lambda)T$ where Γ is the rate from the background and λ is the rate from the source. The rate λ at the detector is inversely proportional to the square of the distance, ignoring the attenuation due to absorption in the air (which plays only a small role at the distances considered here). Therefore

$$c_i(T) = (\frac{\Lambda}{(x_i - x_s)^2 + (y_i - y_s)^2} + \Gamma) \cdot T \qquad (4)$$

where Λ is a constant determined by the source strength. The equation has three unknowns, x_s, y_s, Λ. 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). In general, for N detectors, where $N \geq 4$, the problem is over-constrained, and we use the following heuristic to determine the unknowns.

First, we select the group of four detectors in the network with the greatest aggregate count. Taking each detector i in turn from this group we evaluate Υ_i, an estimate for Λ, for each grid position (x, y) in the field:

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

If the source is located at $[x, y]$ then the value of Υ_i should be the same for each detector i. Thus a measure of the likelihood of the source being located at (x, y) is:

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

where $\overline{\Upsilon}$ is the average over all (x, y). We estimate the position of the source to be the grid position with the maximum likelihood. Figure 8 plots localization error in a four detector square of 30 meters apart. This simple method, though lacking the sophistication and accuracy of other more complex methods, is shown to do reasonably well in a noisy environment.

Source identification Source identification is based on the energy spectrum of photons at a sensor. Identifying mixtures of isotopes can be carried out by analyzing the intensity recorded by sensors of spectra in energy intervals near the peaks for candidate isotopes. Sensors that record energy spectra accurately are more expensive than sensors that merely count the total number of photons over wide energy intervals. One way of dealing with cost constraints is to use dense inexpensive sensors to detect and localize sources, and then use a few, more expensive, sensors to identify isotopes of a source at a specified location.

An approach we are exploring utilizes directionality data provided by higher-end sensors. Sensors with directionality capability can reconstruct the direction of incoming photon stream either through intelligent placement of crystals or *Compton Scattering* effect. It is thus possible to orientate the sensors to get better SNR, which essentially determines the effective range of a sensor. For example, a directional sensor of resolution $\theta = 2\pi/36$ can detect a source at 30 meters away just as effectively as a source at 5 meters away using a nondirectional sensor ($\theta = 2\pi$) given that SNR increases as $r^2 e^{-\alpha r}/\theta\Gamma$. Directionality also allows for localization with as few as two sensors instead of three as current methods require.

The directional sensor is used only after a source is detected and localized to within a few meters by relatively inexpensive, dense, nondirectional sensors. The directional sensor, capable of measuring energy spectra accurately, gets measurements from the source with little interference from the background.

7 Conclusion

This paper explored basic questions about systems for identifying dangerous sources of radiation in the presence of natural background sources. This paper studied fundamental questions such as: How much does data fusion help and how dense a sensor network is necessary? The paper used mathematical analysis and Monte Carlo methods to understand fundamental tradeoffs. The insight that the analysis provides adds to the substantial amount of earlier work. A great deal of work remains to be done particularly with regard to different concepts of operations.

References

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

[2] M Chandy, J Bunn, and A Liu. Models and algorithms for radiation detection. *Modeling and Simulation Workshop for Homeland Security*, 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 Chin, D Yau, N Rao, Y Yang, and C Ma. Accurate localization of low-level radioactive source under noise and measurement errors. *Proceeding of the 6th ACM conference on Embedded network sensor*, Jan 2008.

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

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

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

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

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

[10] 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.

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

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

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

[14] 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.

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

[16] 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.

[17] 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.

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

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

[20] 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*, May 2008.

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