Proceedings of the 2006 Winter Simulation Conference
L. F. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, eds.

DETECTION OF NUCLEAR MATERIAL AT BORDER CROSSINGS USING MOTION CORRELATION

David M. Nicol

Department of Electrical and Computer Engineering
University of Illinois, Urbana-Champaign
Urbana, IL 61801, U.S.A.

Rose Tsang
Heidi Ammerlahn
Michael Johnson

Sandia National Laboratories
Livermore, CA 94551, U.S.A.

ABSTRACT

This paper considers the problem that arises when a vehicle carrying nuclear material is detected approaching a border crossing. As quickly as possible, and with automation we wish to identify which vehicle among all those in the area is likely to be carrying the source. We show that if the border crossing area has technology for tracking the position of vehicles, we can correlate observed movements with observed changes in levels of detected radiation—for as the vehicle carrying the material gets closer to the detector, the stronger will be the detected radiation. We use a simulation model that captures the stop-and-go dynamics of a border crossing area to evaluate our ideas, and find a highly successful technique that tracks which vehicles move just when detected radiation changes, coupled with fitting radiation intensity/distance observations to an inverse-square law. This method almost always isolates the sought vehicle just as soon as the minimum number of data observations is obtained.

1 INTRODUCTION

The Department of Homeland Security has a heightened awareness of the threat that nuclear material might be illegally brought into this country. An industry has grown up around the problem of detecting radiation that is consistent with material that might be used by terrorists (Labov 2004).

We focus on the problem of detecting such material at border crossings. The state of the practice now is to employ a high quality detector, configured to raise an alarm when a threshold becomes high enough. Human responders are sent out with lower-grade hand-held detectors to isolate the vehicle carrying the material. Effective response then depends on the presence and timely reaction by border patrol officers on site. We seek to better automate the isolation procedure, and quickly identify which vehicle is carrying radioactive material. Our key insight is that by correlating observed changes in detection levels with observed changes in vehicle position, we can identify vehicles whose movements are consistent with changes in radiation detection. We say that approaches based on this observation use *motion correlation*.

In our approach we take advantage of the structure of large border crossings, where vehicles queue up in lanes that are more or less straight. As the occupants of each vehicle at the head of the lane are always stopped and questioned, a vehicle in queue engages in a number of start and stop maneuvers before reaching the head. We assume the use of technology capable of tracking vehicles as they move through the lanes. It is well within the capability of video surveillance systems today to track a vehicle under these circumstances and to detect and report times at which it moves (Jiao, Wu, Wu, Chang, and Wang 2006, Gardner and Lawton 1996). RFID is another technology that could do this. At the entrance to the crossing area, drivers would be given customs forms embedded with very cheap RFID tags; and detectors could be deployed along the lanes to detect the passing of vehicles. By correlating measured position data with measured radiation data, we can identity those vehicles that have *almost always* moved when the detection levels changed, and those that have *almost always* been stationary when the levels did not change.

Physics dictates that the detected intensity of a radioactive source decrease in inverse proportion to the square of the distance between source and detector. Even though we don't know what the source radiation level is, we can measure the distance a vehicle moves, correlate that with changes in detected radiation, and fit an inverse-squared model to the observed data. Intuition suggests that the model fitting error will distinguish the vehicle carrying the source from others; we use model-fit error to differentiate between vehicles whose movement are largely temporally correlated with changes in detection level.

We consider three ways of determining the source-carrying vehicle, and compare these on the basis of their ability to identify that vehicle, and the speed at which it is identified. We find that use of model fitting error is particularly effective. We use a movement monitoring system to correlate observed vehicle motion with observed changes in radiation level, filtering vehicles based on highly correlated changes in position and detected radiation. Among filtered vehicles we use model fitting to identity the one that best describes the inverse-squared law. In all the experiments we conducted, this technique *almost immediately* identifies the source-carrying vehicle.

2 MODEL

We model the border crossing approach area as a set of straight, parallel lanes. Entering the area a vehicle joins a queue in one of the lanes. The vehicle at the head of the queue experiences a service time, to account for the delay due to a border crossing official interviewing its occupants. Our model accounts for the fact that vehicles in real queues do not move forward simultaneously when the head vehicle leaves. The queue is modeled as a sequence of cells; vehicles move from cell to cell in a time-stepped fashion. Each time-step, for each lane, if a vehicle is at the head of the lane, we cause it to vacate its cell with probability p_h. Visiting each cell once per time-step, moving from front to back, if the cell in front of a vehicle is empty, the vehicle moves forward into the empty cell with probability p_f. If the very last cell is found to be empty, a new vehicle is introduced there with probability p_a. In this model it is possible for a number of queued vehicles to move forward "simultaneously" as if in lock-step, but also possible that some may not. In our experiments we run traffic through the border area for many time-steps, and then introduce a source-bearing vehicle (SV) into the entry cell of one of the lanes. As SV moves forward we consider how quickly each detection mechanism studied identifies SV's location.

In this model lanes act independently of each other, although all share the same parameters p_h, p_f, and p_a. We assume that at each time-step we can observe the positions of all vehicles, sample the detected radiation at all the detectors we employ, and time-stamp all these measures from a common clock or synchronized clocks.

We suppose there is a detector at the side of the line of front cells; one of our techniques assumes that there is a second detector at the far side of the line of front cells. We suppose that radiation detection readings are taken periodically, and are time-stamped. We suppose that a surveillance or tracking system can trace a vehicle's movement through a lane, providing a series of positions, each time-stamped with a clock that is synchronized with that of the detectors.

We model radiation detection in accordance with an inverse law, plus Poisson noise (Knoll 2000). If the radioactive source has intensity S, we suppose that the detected level at detector D_1 is

$$s_1 = c\frac{S}{d_1^{\alpha}} + N \tag{1}$$

where c is a constant of proportionality, d_1 is the Euclidean distance between source and detector D_1, α describes the rate of detection degradation (normally $\alpha = 2$), and N is random noise, positive, drawn from a Poisson distribution. The noise term is included to account for detected background radiation that is not due to the source of interest.

In this paper we assume that the noise component is small enough compared to detected signal strength so that if the source moves one or more cells between detector observations, the change will be detected with high probability. Specifically, suppose that

- we must detect sources of intensity S_0 or larger in the border crossing area;
- with high probability we must detect movement of a source of at least distance Δ_y within a lane.

In the absence of background noise, the smallest change in detected radiation occurs when the weakest source to be detected enters the border crossing area at a point furthest from the detector, then moves Δ_y in that lane. Choosing coordinates so that the detector is at location $(0,0)$ and the furthest entry point is at (x_L, y_L), then we require the detector to be sensitive enough to distinguish between $cS_0/(x_L^2 + y_L^2) + N_1$ and $cS_0/(x_L^2 + (y_L - \Delta_y)^2) + N_2$, with high probability. This implies that the detector is close enough to the source so that "signal" is large enough relative to "noise" (e.g. the border crossing area is small enough) so that the probability of observing a change (assuming $\alpha = 2$),

$$\begin{aligned}\Pr\left\{\frac{cS_0}{x_L^2 + (y_L - \Delta_y)^2} + N_2 > \frac{cS_0}{x_L^2 + y_L^2} + N_1\right\} \\ = \Pr\left\{N_1 - N_2 < \frac{cS_0}{x_L^2 + (y_L - \Delta_y)^2} - \frac{cS_0}{x_L^2 + y_L^2}\right\},\end{aligned}$$

is as high as one needs. Assuming that N_1 and N_2 are i.i.d. Poisson random variables, one can compute their common parameter λ to provide the desired bound. A back-of-the-envelope technique is to use the Chebyshev Inequality (Larson and Shubert 1979). Applied here, it states that

$$\Pr\{N_1 - N_2 > k\sigma\} \le \frac{1}{k^2},$$

where $\sigma = \sqrt{2/\lambda}$ is the standard deviation of $N_1 - N_2$. Setting $k = 10$, and $k\sigma$ equal to the minimal difference in detected signal strength after a move, we can solve for $1/\lambda$—the mean background noise—for that ensures that the probability of noise factors confusing detection of real movement is less than 1%. Using this method we computed a mean noise value that is approximately 250 times smaller than the detected radiation of a source with strength S_0 at the furthest entry point from the detector. Future work will consider how to deal with environments with comparatively more significant background noise.

We also explicitly incorporate error in movement measurement in our simulations. When a vehicle moves, on the next snapshot of the movement sensors that movement will not be detected, with probability p_{me}.

One of the questions we consider is the degree to which our detection mechanisms are sensitive to background noise, and to movement measurement error.

3 SOURCE DETECTION

We now outline the three key ideas used to determine the source location. One of these correlates movement with changes in detected intensity, another is to use dual detectors to constrain where the source might lie, and the third is to use model-fitting information. We describe each in turn.

3.1 Correlating Movement and Source Changes

Recall our assumption that each time-step the positions of all vehicles may be observed, and the detected radiation is also observed. We detect a change in detection value with high confidence, and so each time-step can note for each vehicle whether it behaved as did the (unknown) source—if the source changed position (as evidenced by a change in detection value) did the vehicle also move? If the source did not change, did the vehicle remain stationary?

Thus we can associate with each observed vehicle a *motion correlation* counter that records the number of time-steps in which the vehicle's movement agreed with changes in the source. The counter is initialized to zero when the vehicle enters the border crossing area, and is updated during each time-step in which the vehicle remains in the area.

Unlike the source detection mechanisms we previously studied (Nicol, Tsang, Ammerlahn, and Johnson 2006), this one makes no assumptions about the form or parameters of how detected radiation declines as a function of distance. It does however assume that the background noise is small enough relative to signal strength that movement by the source is with high probability detected as changes occur in the detected radiation.

3.2 Using Dual Detectors

Reconsider Equation (1). If we have two detectors and N is small relative to cS/d_1^α, without knowledge of c, S, or d_1 it is possible to determine a curve through the border crossing area on which the source must lie. Ignoring N we write $s_1 d_1^\alpha = cS$; and note that also $s_2 d_2^\alpha = cS$ if s_2 is the detected radiation level at detector D_2, and d_2 is the distance between the source and D_2. From the equation $s_1 d_1^\alpha = s_2 d_2^\alpha$ we determine $(s_1/s_2)^{1/\alpha} = (d_2/d_1)$. The left-hand-side we can compute. As the positions of D_1 and D_2 are known, the set of points (x, y) whose distances from D_1 and D_2 satisfy $d_2/d_1 = (s_1/s_2)^{1/\alpha}$ form a *source curve* through the border crossing area, on which the source must lie. Figure 1 illustrates the concept. In fact, in our previous work we showed that with *three* detectors we can construct three such curves (there being three distinct pairings of three detectors) and precisely identify the source location as the unique intersection of those three curves. We now are interested in isolating the source using fewer detectors, and to determine how quickly a source can be identified once it enters the border crossing area.

Just as in the case of one detector, when a vehicle enters the border crossing area we initialize the correlated motion counter to zero. If in a time-step the source is detected to move, then we examine only vehicles in cells through which the source curve passes, and cells that are "close" to the source curve (e.g., cells that are within some distance r of any point on the source curve.) We increment the motion correlation counter only of vehicles that moved and are close to the source curve.

3.3 Model Fitting

We can use Equation (1) to determine how well a vehicle's movements, temporally correlated with changes in (single) detector measurements, obeys the assumed law. Suppose we have n data points $\{(x_i, m_i)\}$, $i = 1, 2, \ldots, n$, where x_i is the measured distance of a vehicle from the detector, and m_i is the detector reading when the vehicle is at that position. We take care to ensure that $x_i \neq x_j$ and $m_i \neq m_j$ for all i, j. If the data were to fit Equation (1) exactly and there were no noise, then

$$\log(m_i) = \log(cS) - \alpha \log(x_i)$$

for each i. However, this cannot be explicitly checked, as we don't know the value cS. We can however *estimate* cS with a value b. Given b and α, the residual error associated with (x_i, m_i) is the difference between the value predicted using b and α, and the value observed:

$$e_i(b, \alpha) = b - \alpha \log(x_i) - \log(m_i).$$

A commonly used measure of the overall fit is the sum (over all observations) of squared errors :

$$ss_e(b, \alpha, n) = \sum_{i=1}^{n} e_i(b, \alpha)^2.$$

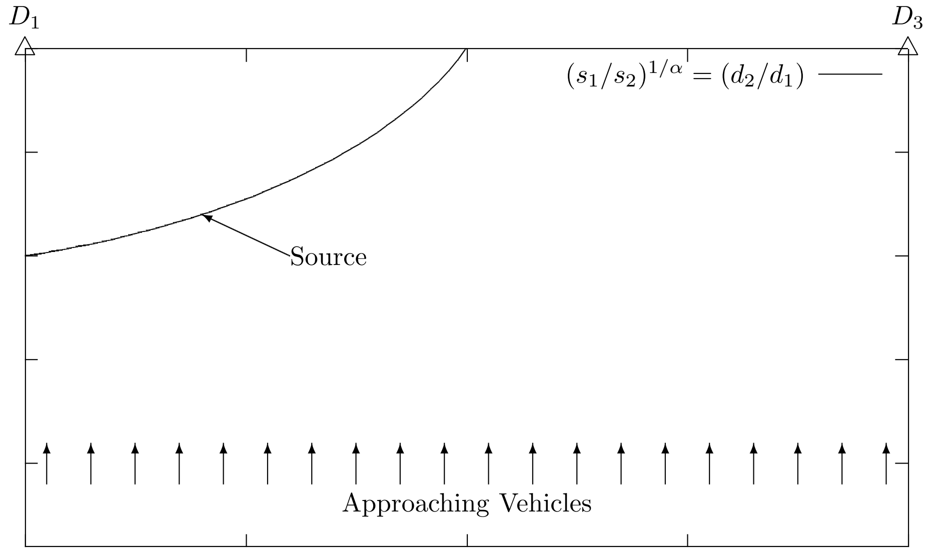

Figure 1: Using Two Detectors We Can Compute a Line upon Which the Source Must Lie

When both cS and α are unknown, techniques of linear regression(Weisberg 2005) identify estimates $\hat{b}$ and $\hat{\alpha}$ that minimize this error measure. In our previous work (Nicol, Tsang, Ammerlahn, and Johnson 2006) we observed that this estimator, taken over *all* of a vehicle's positions up to the border crossing, was highly effective in identifying the source—the error associated with the source vehicle was several orders of magnitude smaller than any other. In the present paper we do not fit different values of α to different vehicles—it is after all the same law of physics applying to each—and we are interested in using model fit error much earlier in the source vehicle's trajectory through the border crossing. Also, the traffic model in the earlier paper assumed that all cells have vehicles, and that vehicles move forward one cell in lock-step, each time-step.

For a fixed (and known) value of α, we recognize the relationship of

$$ss_e(b, \alpha, n) = \sum_{i=1}^{n} \Big(b - (\alpha \log(x_i) + \log(m_i)) \Big)^2$$

to the second central moment of the data values $\{(\alpha \log(x_i) + \log(m_i)\}$. We know then that the value of b which minimizes $ss_e(b, \alpha)$ is the sample mean (Larson and Shubert 1979)

$$\hat{b} = (1/n) \sum_{i=1}^{n} (\alpha \log(x_i) + \log(m_i)),$$

and corresponding assess the fitness of a vehicle's trajectory with a normalized minimized model fit error $(1/n) ss_e(\hat{b}, \alpha, n)$. The normalization enables us to compare the fits of vehicles with different numbers of observations. In the absence of noise, with a correct value of α, and with perfect position information, the error associated with the source will be nearly 0.

3.4 Detection Algorithms

We bring the detection ideas together to form three algorithms. The *single-detector, model-free* algorithm (SD) assumes one detector, and does not use model-fitting information. For each time-step we construct a set of vehicles whose movements deviate from changes in detected intensity since arriving by no more than a threshold T_{move}. We call these vehicles "suspects", and are interested in how the number of suspects changes in time as the source moves through the border crossing. Ideally the number of suspects will have shrunk to 1 before the source reaches the crossing.

The *dual-detector, model-free* algorithm (DD) is identical to SD, except that the motion correlation counters are updated only for vehicles on the source curve. Like SD, we maintain a set of suspect vehicles.

The third algorithm is constructed by bringing model-fitting information to SD. (It could also be brought to DD, but that turns out to be unnecessary.) We further analyze the suspect groups defined by motion correlation, creating for each vehicle a tuple comprised of (i) the number of deviations of movement relative to the source, and (ii) the normalized model-fit error based on observations from detector D_1. One tuple A is considered to *dominate* tuple B if each of A's components is no larger than the corresponding component in B. Tuples that are not dominated by any other are candidates for the source-carrier. In the *single-detector,* *model-used* (SDM) method we count the number of vehicles with non-dominated tuples. When that figure goes to 1, we suppose that the unique vehicle with non-dominated tuple is the source-carrier.

4 EVALUATION

We are interested in how well different detection techniques work, and scenario parameters that may significant impact on that performance. Our evaluation study is built around a simulator, whose workings have already been described. Experimentation has shown us that the detection mechanisms are sensitive to how many vehicles are in the border crossing area, and how many "snapshots" are used as vehicles move through the area. We control the number of vehicles in the area through parameters p_h and p_a, respectively the probability that at a time-step a vehicle at the head of the line departs, and that a vacancy at the entrance to a lane is filled by a new vehicle. We set the probability of a vehicle moving ahead into an empty cell to be 0.8, assume $p_h = p_a$, and vary the former parameters. By equivalencing p_h and p_a we encourage a certain amount of queueing, but do not have the arrival rate so high that a full system is ensured. Our experiments include "slow" service with $p_h = 0.05$, and "fast" service with $p_h = 0.4$. In any experiment all servers use the same parameter.

We consider three sizes of domain—ten lanes each twenty vehicles deep, ten lanes each ten lanes deep, and five lanes each ten vehicles deep.

Each experiment initially "warms up" the state of the border crossing by running 1000 time-steps, introducing vehicles and having them move though the system. On the 1000^{th} time-step the source is introduced. Thereafter, until the source departs, we apply the detection algorithms. Each time-step we identify the current set of suspect vehicles. When (and if) that set dwindles in size to 1 vehicle we consider to have found the source. The simulator checks to ensure that the vehicle so identified is indeed the source, and in no experiment run was it ever the case that the wrong vehicle was selected.

The SD strategy employs a threshold parameter n_s of the number of time-steps in which a vehicle's detected movement can differ from detected changes in source intensity, and still be in the set of suspected vehicles. In all the experiments reported here, $n_s = 3$.

We are interested in whether the source is isolated, how quickly the source is isolated, and the impact that modest degrees of noise and movement measurement error have on detection performance. Figures 2 and 3 describe the result. For each parameter setting we ran 100 independent replications. Each replication we measured and stored the size of suspect sets as a function of time since the source entered the area. From these we computed mean and standard deviations, and plot these against time, indexed against the number of movements forward the source has taken (i.e., the size of the suspect set is measured at a time-step in which the source vehicle moves, and the statistics are taken from these snapshots.) We require three observations for a model fit (in order to minimize spurious "good fits"), and so begin the evaluation upon the third detected movement of the source.

Figure 2 describes performance when there is no background noise, and no movement measurement error. The high variance in suspect set size in early time-steps is a distinctive characteristic of all the data. Even so, there are clearly discernable trends.

- The suspect set size for the SD technique is in the early time-steps strongly influence by both the domain size and the vehicle arrival rate. This is understood as reflecting how many vehicles there are in the domain area at a time. The more there are, the more suspects exist early on.
- The SD technique does not reliably completely isolate the source on domains that are only 10 vehicles deep. To converge to source isolation (as seen in the 10×20 domains) it needs more observation points. This would be accomplished in a given domain simply by taking more movement and radiation change samples.
- While the DD technique immediately defines smaller suspect sets and improves the chances of isolating the source, cases exist where it has poor performance—consider in particular the 5×10 case with fast service. Both SD and DD isolate the source in fewer than 1/2 the cases.
- SDM is the star of the show. In the absence of noise and movement error it *immediately* converges on the source. Indeed, considering the cost of sensitive detectors, there really is no point in using two detectors if one must deploy a movement tracking system for it anyway.

Figure 3 considers the impact of limited noise, and some movement measurement error. The mean background noise level was selected as described earlier, using Chebyshev's inequality to ensure that at least 99% of the time simple comparison of detection levels before and after the source moves will reveal the move. The probability that a vehicle's actual movement is missed on one time-step is set to 0.1. For reference, the noise/error free data for SD is also plotted. The take-home message of the data shown in Figure 3 has two parts. The first is that this limited amount of background noise and measurement error doesn't affect SD or SDM very much. The noise/error can keep SDM from perfectly isolating the source on occasion; in those replications the (random) movement error occurred to the source, causing the "movement error" component of its tuple to be bested

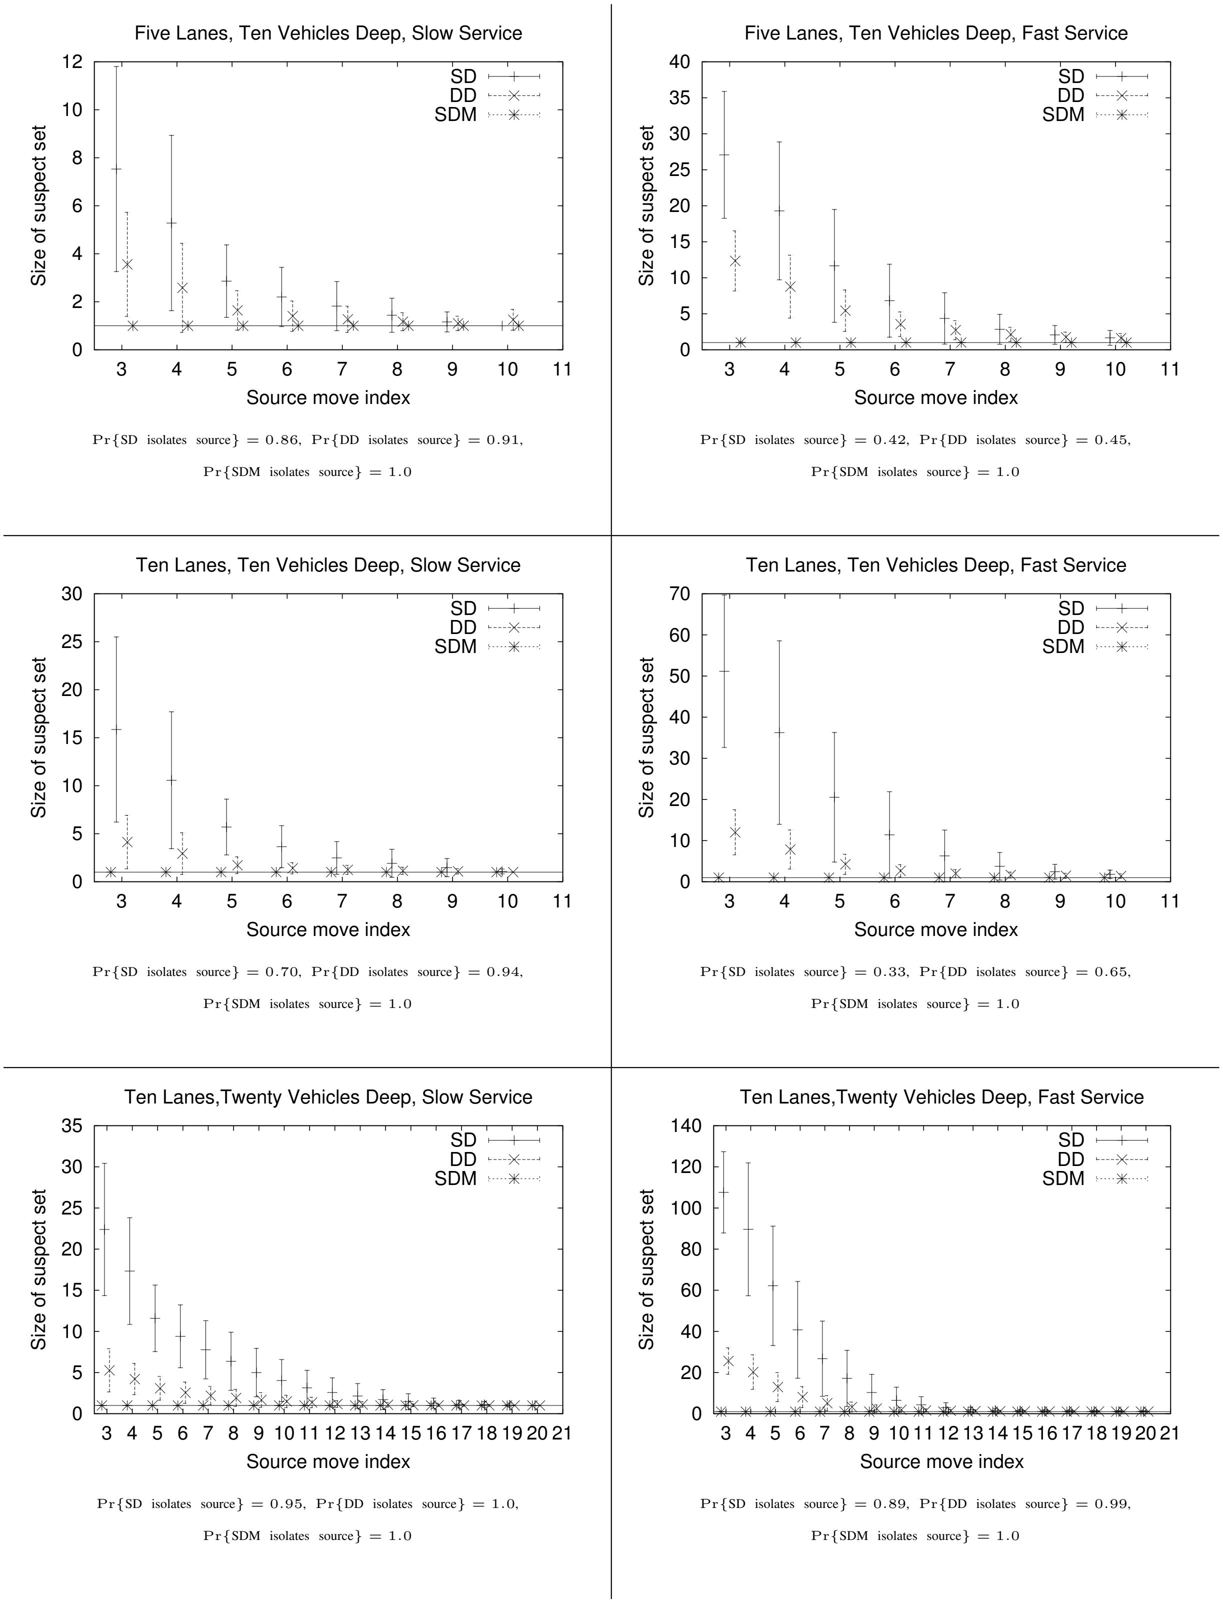

Figure 2: Source Isolation Behavior without Error in Movement Measurement and without Background Noise

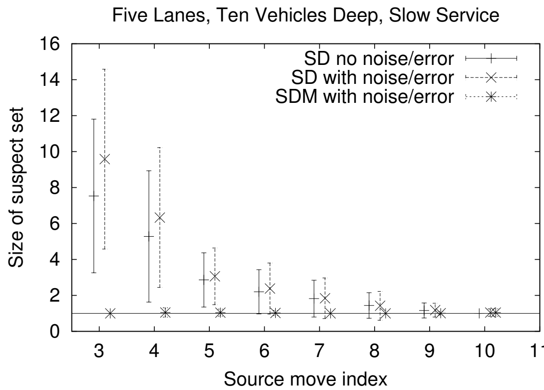

$\Pr\{\text{SD isolates source}\} = 0.84$, $\Pr\{\text{SD loses source}\} = 0.0$,

$\Pr\{\text{SDM isolates source}\} = 0.99$

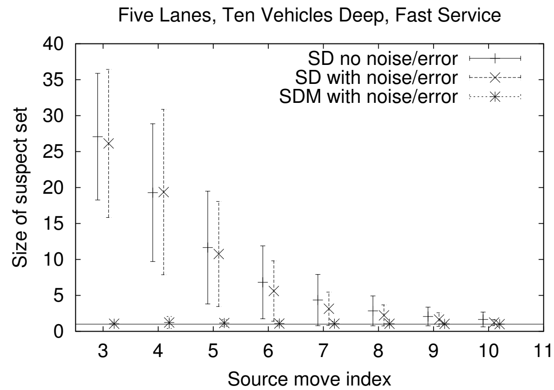

$\Pr\{\text{SD isolates source}\} = 0.67$, $\Pr\{\text{SD loses source}\} = 0.0$,

$\Pr\{\text{SDM isolates source}\} = 0.97$

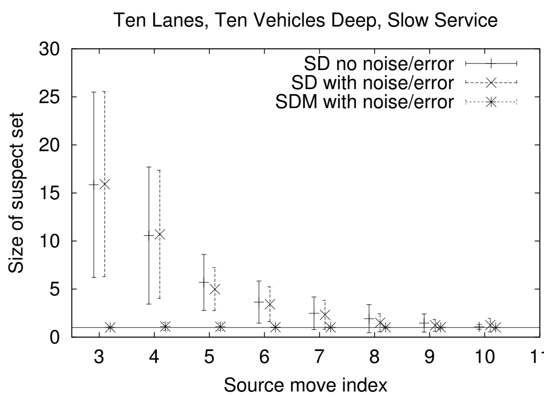

$\Pr\{\text{SD isolates source}\} = 0.87$, $\Pr\{\text{SD loses source}\} = 0.0$,

$\Pr\{\text{SDM isolates source}\} = 0.98$

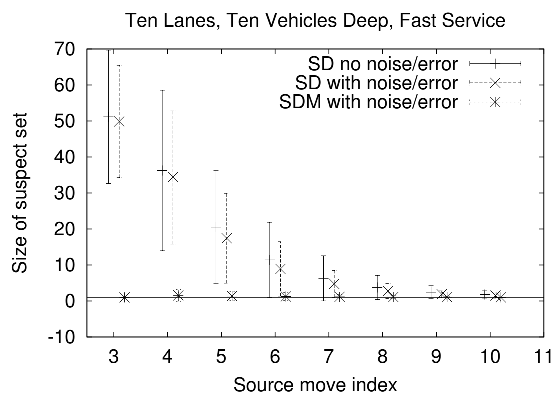

$\Pr\{\text{SD isolates source}\} = 0.50$, $\Pr\{\text{SD loses source}\} = 0.01$,

$\Pr\{\text{SDM isolates source}\} = 0.97$

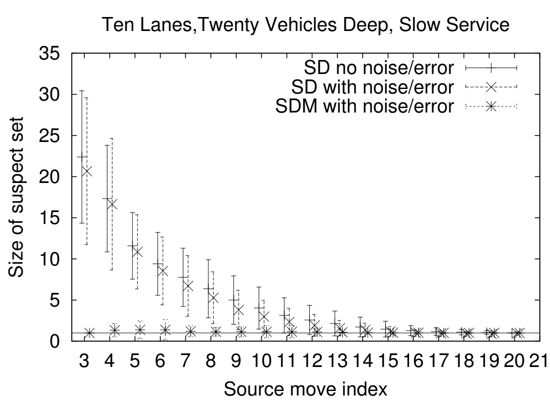

$\Pr\{\text{SD isolates source}\} = 0.86$, $\Pr\{\text{SD loses source}\} = 0.13$,

$\Pr\{\text{SDM isolates source}\} = 0.99$

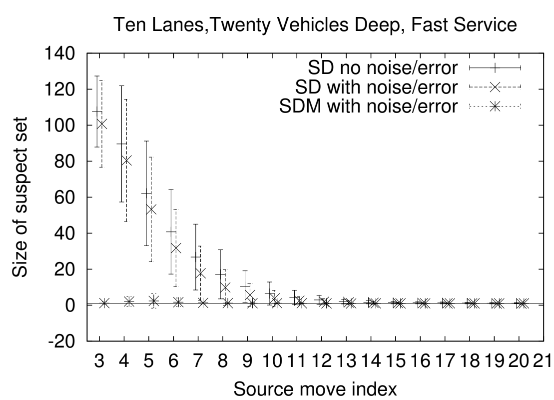

$\Pr\{\text{SD isolates source}\} = 0.9$, $\Pr\{\text{SD loses source}\} = 0.08$,

$\Pr\{\text{SDM isolates source}\} = 0.98$

Figure 3: Source Isolation Behavior with Probability of Missed Moved = 0.1 and with Background Noise

by another tuple. However, the model-fitting error of the source still isolates the source.

It is interesting to note that on the larger domains the mean size of the SD suspect size decreases slightly when we include noise and measurement error. This is due the fact that random movement error only decreases the size of the suspect set. Our error model only keeps a moving vehicle from being noticed, it does not report that a stationary vehicle moved.

Figure 3 also records the observed frequency with which the source vehicle was lost from the suspect set.

The overall conclusion of the our experiments is that use of model-fitting data is a very good thing. However, the model-free approach of SD *can* often isolate the source, if enough concurrent movement and radiation change measurements can be made before the source reaches the border crossing.

5 CONCLUSIONS

This paper considers the problem of detecting and isolating a vehicle carrying nuclear material at a border crossing. We consider three ideas that contribute to isolation algorithms. One is to correlate in time vehicle movements (as observed through a visual or RFID tracking system) with observed changes of radiation level at one (or two) detectors; another is to use two detectors to determine a region of the domain in which a source must lie; a third is to use simultaneous position and detected radiation levels to fit an inverse-squared law to observed vehicle movement, with the expectation that the vehicle carrying the material will distinguish itself with very low model fitting error. Using a simple simulation that attempts to capture relevant dynamics of vehicles in a border crossing area, we find that is it usually possible to very quickly distinguish the vehicle carrying nuclear material from others in the area, with a combination of the first and third ideas.

ACKNOWLEDGMENTS

This research was supported in part by DARPA Contract N66001-96-C-8530, and NSF Grant CCR-0209144. Accordingly, the U.S. Government retains a non-exclusive, royalty-free license to publish or reproduce the published form of this contribution, or allow others to do so, for U.S. Government purposes. In addition, this research program is a part of the Institute for Security Technology Studies, supported under Award number 2000-DT-CX-K001 from the U.S. Department of Homeland Security, Science and Technology Directorate. Points of view in this document are those of the author(s) and do not necessarily represent the official position of the U.S. Department of Homeland Security or the Science and Technology Directorate.

REFERENCES

Gardner, W., and D. Lawton. 1996, November. Interactive model-based vehicle tracking. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 18 (11): 1115–1121.

Jiao, L., G. Wu, Y. Wu, E. Chang, and Y.-F. Wang. 2006. The anatomy of a multi-camera video surveillance system. *ACM Multimedia System Journal*. to appear.

Knoll, G. 2000. *Radiation detection and measurement, 3^{rd} edition*. Hoboken, NJ: Wiley.

Labov, S. 2004, Sept.. Radiation detection on the front lines. *Lawrence Livermore Science and Technology Review*. see http://www.llnl.gov/str/September04/Labov.html.

Larson, H., and B. Shubert. 1979. *Probabilistic models in engineering sciences, volume i*. Hoboken, NJ: Wiley.

Nicol, D., R. Tsang, H. Ammerlahn, and M. Johnson. 2006, April. Sensor fusion algorithms for the detection of nuclear material at border crossings. In *Proceedings of the 2006 SPIE Conference on Sensors, and Command, Control, Communications, and Intelligence (C3I) Technologies for Homeland Security and Homeland Defense V*. Orlando, FL.

Weisberg, S. 2005. *Applied linear regression, 3^{rd} edition*. Hoboken, NJ: Wiley.

AUTHOR BIOGRAPHIES

DAVID M. NICOL is Professor of Electrical and Computer Engineering at the University of Illinois, Urbana-Champaign. He received a B.A. in mathematics from Carleton College in 1979,and M.S. and Ph.D. degrees in computer science from the University of Virginia in 1983 and 1985 respectively. He has served as program and general chair of several simulation-oriented conferences, and from 1997-2003 was Editor-in-Chief of ACM Transactions on Modeling and Computer Simulation. He is a Fellow of the IEEE, and a Fellow of the ACM.

HEIDI AMMERLAHN the manager of the Computational Sciences and Mathematics Research (CSMR) Department at Sandia National Laboratories, California. She holds an M.S. in Computer Science from Stanford University and a B.S. in Computer Science and Mathematics from the University of Washington. Heidi has led the development of computer simulations supporting evaluation of technologies and procedures for detecting and responding to terrorist threats with weapons of mass destruction in metropolitan and U.S. Border regions.

MICHAEL M. JOHNSON is a program manager in the Homeland Security Systems and Development Center at Sandia National Laboratories, California. He received his M.Sc. degree in Computer Engineering from the University

of California, San Diego, California in 1991. His research interests include high performance computing, modeling and simulation, and embedded systems design.