Using Container Inspection History to Improve Interdiction Logistics for Illicit Nuclear Materials

Chenhua Li,[1] Gary M. Gaukler,[2] Yu Ding[1]

[1] *Department of Industrial and Systems Engineering, Texas A&M University, College Station, Texas 77843*

[2] *Peter F. Drucker & Masatoshi Ito Graduate School of Management, Claremont Graduate University, Claremont, California 91711*

Received 5 June 2012; accepted 11 May 2013
DOI 10.1002/nav.21542
Published online 19 July 2013 in Wiley Online Library (wileyonlinelibrary.com).

Abstract: Strengthening the United States' ability to prevent adversaries from smuggling nuclear materials into the country is a vital and ongoing issue. The prospect of additional countries, such as Iran, obtaining the know-how and equipment to produce these special nuclear materials in the near future underscores the need for efficient and effective inspection policies at ports and border crossings. In addition, the reduction of defense and homeland security budgets in recent years has made it increasingly important to accomplish the interdiction mission with fewer funds. Addressing these complications, in this article, we present a novel two-port interdiction model. We propose using prior inspection data as a low-cost way of increasing overall interdiction performance. We provide insights into two primary questions: first, how should a decision maker at a domestic port use detection data from the foreign port to improve the overall detection capability? Second, what are potential limitations to the usefulness of prior inspection data—is it possible that using prior data actually harms decision making at the domestic port? We find that a boundary curve policy (BCP) that takes into account both foreign and domestic inspection data can provide a significant improvement in detection probability. This BCP also proves to be surprisingly robust, even if adversaries are able to infiltrate shipments during transit. © 2013 Wiley Periodicals, Inc. Naval Research Logistics 60: 433–448, 2013

Keywords: port security; nuclear materials; container shipping; inspection policies

1. INTRODUCTION

Over the past decade, as terrorist activities have spread throughout the world, the United States has become increasingly concerned about unsecured special nuclear materials (SNM), for example, plutonium and highly enriched uranium (HEU). Agencies such as the Department of Homeland Security fear that these materials may be used to build relatively simple but effective nuclear explosive devices [21]. To prevent adversaries from smuggling SNM into the U.S., the U.S. government has not only equipped its own ports with radiation detection equipment, but also helped install radiation detection equipment at foreign countries, through a number of global initiatives [1]. With the possibility of having radiation detection data at both a foreign port and a domestic port, this article intends to provide insights into two primary questions: First, how should a decision maker at a domestic port use detection data from the foreign port to improve the overall detection capability? Second, what are potential limitations to the usefulness of prior inspection data—is it possible that using prior data actually harms decision making at the domestic port?

To investigate these issues, we consider a two-port, seaborne container inspection problem. The port of origin (or embarkation) is a foreign port, the port of entry (or debarkation) is a domestic port, and in between, the containers travel on a container ship. This transportation and inspection setting is illustrated by the upper graph in Fig. 1. At each port, the container goes through a layered inspection system, different varieties of which have been described in the literature [9, 10, 23]. In this article, we consider an inspection system comprised of four detection components: the intelligence-based-automated target system (ATS), radiography equipment (such as X-ray machines), passive radiation detection equipment [radiation portal monitors (RPM)], and manual inspection (namely, opening up a container and doing item-by-item manual search). This part is illustrated in the lower graph in Fig. 1. The details of this layered inspection system are described in Gaukler et al. [10]. For this article to be self-contained, we will provide a brief review of how the single-port inspection system is modeled in Section 2.

Correspondence to: G.M. Gaukler (gary.gaukler@cgu.edu)

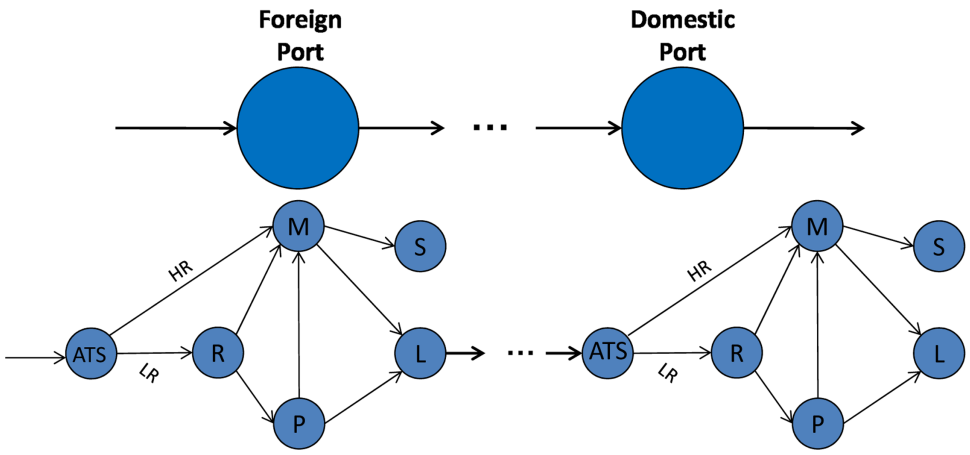

Figure 1. Two-port inspection network. ATS: automated target system; r: radiography; P: passive detection; M: manual inspection; S: successful detection; L: loading to a ship, or leaving the system; HR: high risk, LR: low risk. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

As the containers go through the layered inspection system at the foreign port, there are three pieces of detection-related data: the ATS score, the radiography output, and the gross radiation count from the passive detector (RPM). The ATS score is a risk assessment of the containers based on their manifest, documents, and other intelligence information. Here, we will consider the situation where the radiography output and the gross radiation count are shared between the two port inspection authorities. We note that it is not reasonable for ATS scores to be shared in this fashion because as additional intelligence information is entered into the system, ATS scores are continually updated throughout the container's or vessel's voyage. Hence, the ATS score at the time the container arrives at the domestic port contains the most up-to-date, and presumably most useful information. Therefore, the container inspection history, or prior detection data, which we will frequently refer to later, consists only of the container's radiography data and the gross radiation count obtained at the foreign port.

Given the obvious significance of detecting illicit nuclear materials, there has been a rich body of literature investigating related problems. The current literature and practice, however, primarily focus on proposing and analyzing inspection policies at a single port, using detection data obtained at the very same port. The current research can be largely grouped under two categories: the tactical problem, which is to find the optimal inspection policy under delay time and/or budget constraints [7, 9, 10, 14, 23, 24], and the strategic problem, which is to decide the locations of inspection stations and detectors along a long border, assuming that the inspection policy at each station is given [6, 16, 17]. We are not aware of any tactical problems that involve the use of container inspection history. The strategic problem naturally involves multiple ports/locations. But the current treatments [6, 16, 17] assume that each location conducts their inspection independently without considering the possibility of carrying over detection data already obtained from upstream locations.

Our research focuses on a tactical problem but involves the use of detection data at two ports. Our conjecture is that the container inspection history is valuable and should be used for making the final inspection decision at the domestic port. In this article, we explore several possible ways of using the prior detection data and evaluate their relative effectiveness. We recommend one particular scheme which we believe offers sensible improvement and is also robust in the presence of possible infiltration, that is, when an adversary breaches a container while it is in transit from the foreign port to the domestic port.

The article unfolds as follows. Section 2 provides a brief overview of the layered inspection system employed at each single port. Section 3 presents the modeling details, model formulation, and solution of the two-port inspection problem. Section 4 analyzes the proposed way of using the prior detection data and compares its performance with a possible alternative as well as the single-port inspection system, which ignores the container inspection history. Section 5 studies how the proposed method may perform in the presence of infiltration. Finally, we conclude the article in Section 6.

2. OVERVIEW OF THE LAYERED INSPECTION SYSTEM AT A SINGLE PORT

Figure 2 shows the container inspection practice commonly used at many US ports [3, 22]. A container goes through the intelligence-based ATS screening first, and if it is flagged as "high-risk," it will go through a more stringent, time-consuming manual search process. If the container is

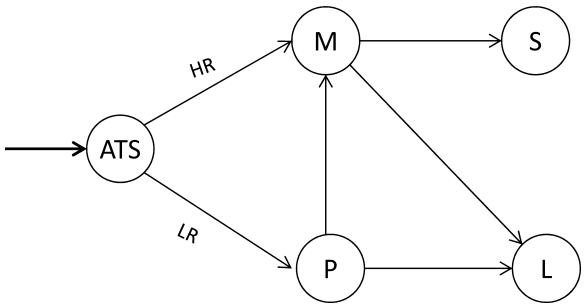

Figure 2. Standard version of the single port inspection with no radiography machines. Notations follow those in Fig. 1.

deemed "low-risk," then it will go through a passive radiation detection step, where gross radiation particle counts (usually gammas) are collected by a detector while the container is in the inspection station for about 30 to 60 s. These particle counts are used to decide whether the container needs to be escalated to a manual search. A threshold policy is used to make that decision, and the threshold is usually decided based on the distribution of the background radiation, the distribution of the elevated radiation when SNM of a certain quantity is present, and a tradeoff between the resulting probabilities of missed detections and false alarms.

Figure 3 provides a simple illustration of how the detection probability and false alarm probability are calculated. In the figure, Ψ denotes the gamma emissions stemming from the natural background, approximated as a Normal distribution, and Ω denotes the gamma emissions when SNM is present, which includes both the SNM emission and the background radiation. In practice, a high false alarm probability translates to a long delay time at ports, so in the latter sections of this paper, the criterion of false alarm is replaced by the expected delay time.

A number of recent papers have aimed at improving this standard version of container inspection systems. A chief complaint about the current system, made in Gaukler et al. [9, 10] for instance, is that the ATS scores are not reliable enough to identify high-risk containers, and the passive detectors are inherently incapable of detecting the existence of small quantities of SNM. This is a particular concern when the SNM in question is HEU, which is not a very bright material, that is, it emits a relatively low level of radiation. Shielding the SNM by surrounding it with high-density materials like metals for example, further reduces the emissions from a particular container. To help increase the detection odds, it has been suggested to further classify the containers according to their contents and then subject those which have high-density materials (high-Z materials) to more stringent manual inspection, while letting containers with low-Z materials go through the passive detector [9, 10]. The rationale for this proposed change is that containers with high-Z materials provide a significant amount of shielding to SNM. Hence, should such materials be hidden in the container, the container will likely pass the passive detector with a very low probability of being stopped. On the other hand, containers with low-Z materials do not provide much shielding so that there is a reasonable chance for the SNM to be detected by passive RPM.

Such radiography technology that can obtain the Z-values of materials inside a container does exist. Doing this usually requires x-ray equipment, sometimes known as a Z portal monitor [2], to scan the container. The grayness of the objects in the returned image is proportional to the Z-value of an object. Such systems are deployed, for example, at the San Ysidro, CA border checkpoint [5].

Gaukler et al. [10] further propose a simple measure, called "hardness" and denoted by h_s, to characterize how hard it would be for a given passive detector to detect SNM inside a particular container type s (a container type here refers to a categorization of containers that are laden with different materials). For a given container type, the hardness measure h_s is chosen to be the misclassification error if the detection threshold was set at the intersection point of the two probability density functions, and it corresponds to the area of overlap between the two probability density functions; see Area A in Fig. 4. This misclassification error is the smallest possible among all choices of the detection threshold, had the two probability density functions been given. It is therefore chosen to characterize how hard (or easy) it is to detect the presence of SNM for a given container type. A container with a large h_s is referred to as a "hard" container, whereas that of a small h_s is referred to as a "soft" container; the larger this overlap area is, the harder the container, and the more difficult it is for the passive detector to detect the SNM, if there is any.

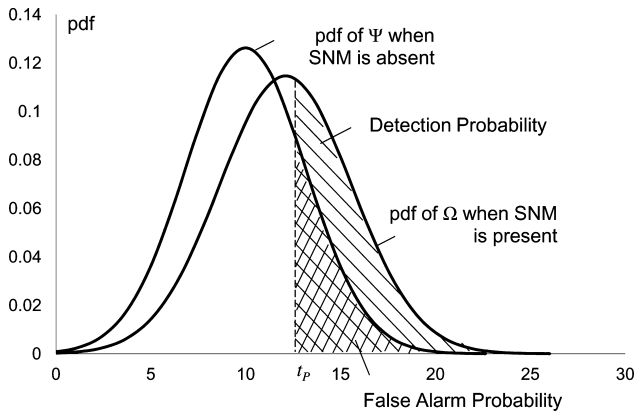

Figure 3. Detection probability, false alarm probability, and overlap of the probability density functions (pdf). t_p is the threshold used at the passive detection.

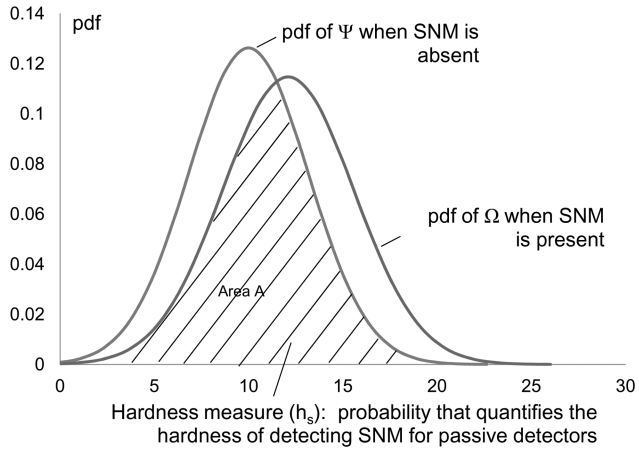

Figure 4. Determination of the hardness measure.

We can incorporate this new technology changes into the current system by inserting a radiography step (denoted by "R") after the ATS but before the passive node; please see Fig. 5 for the new system flow. This new system is labeled as the hybrid inspection system by Gaukler et al. [10], meaning that the decision is based on a hybrid of intelligence (via ATS) and the container content information. In this hybrid inspection system, the two distributions in Fig. 4 are computed using Monte Carlo N-Particle Code (MCNP) [15] simulations, a software package developed at Los Alamos National Laboratory and considered to be the *de-facto* industry standard for performing nuclear transport calculations for radiation particles. Thus, once the radiographic image of the container is obtained together with its geometry, MCNP simulates the performance of a given passive radiation portal monitor, with and without a particular quantity of SNM in the container. The output from MCNP are the two distributions in Fig. 4, and it is thus possible to compute the container hardness h_s for that particular container.

In the above description, we have used the concept of "container type," which was previously introduced in Gaukler et al. [9, 10]. The container type refers to a categorization of containers that are laden with different materials, and therefore, have different ability to shield any potentially existing SNM. In principle, every single container could form its own container type. In practice, however, containers of similar hardness levels will be grouped together to form one container type for simplicity, so that the total number of container types will be relatively small.

Our subsequent model development in this article will use this hybrid inspection system as a primary baseline. In the analysis section (Section 4), we will again touch on the implications of using prior detection data when the standard version of the inspection system (in Fig. 2) is used in both ports.

3. MODELING THE TWO-PORT PROBLEM

When considering a two-port inspection system, such as the one shown in Fig. 1, the inspection decision at the foreign port is no different from a single-port system. What we focus on in this article is how the new decision is made at the domestic port, taking into account both the local (domestic) and prior detection data.

Recall that the container inspection history includes the radiography output and passive detector's gross count data obtained at the foreign port. This container inspection history impacts the R node and P node at the domestic port. To reflect that, we use the notation R(F,D) and P(F,D), respectively, to make it explicit that the decisions associated with the two nodes are now a function of the data from both ports, where "F" refers to the foreign port information and "D" refers to the domestic port information. We can therefore simplify the two-port network by concentrating on the domestic port only, using the above notation to indicate the carrying-over of container inspection history (see Fig. 6); doing this differentiates the inspection process associated with this layout from a pure single-port circumstance (in Fig. 5).

3.1. Modeling Passive Detection

We first describe the passive detection stage, because the decision at the radiography stage is based on the perceived performance of passive detection.

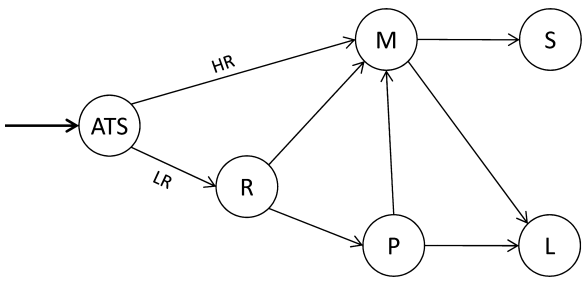

Figure 5. Hybrid inspection system in a single port with the radiography machines. Notations follow those in Fig. 1.

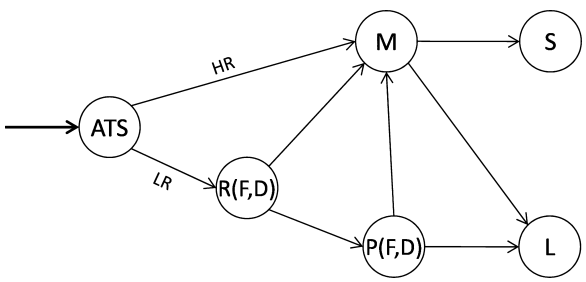

Figure 6. Inspection process at the domestic port with the historical detection data carried over.

In the nuclear engineering literature (for example, Fetter et al. [8]), a normal distribution (where the variance is set equal to the mean) is often used to approximate the Poisson distribution of particle incidence on a detector. We follow this convention here for simplicity as well. Consider a container type s, and let X_s^d and X_s^f denote the normal random variables that describe, respectively, the gross particle counts obtained at the domestic port detector and the foreign port detector. In the presence of SNM, the means of X_s^d and X_s^f are represented by ν_s^d and ν_s^f, respectively, whereas in the absence of SNM, the corresponding means are represented by μ_s^d and μ_s^f, respectively. Naturally, $\nu_s^d > \mu_s^d$ and $\nu_s^f > \mu_s^f$.

Given the two gross particle count random variables for the two ports, the random vector $\mathbf{X}_s = [X_s^d \quad X_s^f]'$ follows a bivariate normal distribution. When SNM is present, the mean vector is $\boldsymbol{\nu}_s = [\nu_s^d \quad \nu_s^f]'$ and the covariance matrix is

$$\boldsymbol{\Sigma}_{1s} = \begin{bmatrix} \nu_s^d & \mathrm{Cov}(X_s^d, X_s^f) \\ \mathrm{Cov}(X_s^d, X_s^f) & \nu_s^f \end{bmatrix},$$

and the corresponding pdf is:

$$f_{1s}(\mathbf{x}_s) = \frac{1}{2\pi\,|\boldsymbol{\Sigma}_{1s}|^{1/2}} e^{-\frac{1}{2}(\mathbf{x}_s-\boldsymbol{\nu}_s)'\boldsymbol{\Sigma}_{1s}^{-1}(\mathbf{x}_s-\boldsymbol{\nu}_s)}, \qquad (1)$$

where the subscript 1 is used to indicate the presence of SNM and $\mathbf{x}_s = [x_s^d \quad x_s^f]'$ denotes a realization of the random vector $\mathbf{X}_s$. When SNM is absent, the mean vector is $\boldsymbol{\mu}_s = [\mu_s^d \quad \mu_s^f]'$ and the covariance matrix is

$$\boldsymbol{\Sigma}_{2s} = \begin{bmatrix} \mu_s^d & \mathrm{Cov}(X_s^d, X_s^f) \\ \mathrm{Cov}(X_s^d, X_s^f) & \mu_s^f \end{bmatrix},$$

and the corresponding pdf is:

$$f_{2s}(\mathbf{x}_s) = \frac{1}{2\pi\,|\boldsymbol{\Sigma}_{2s}|^{1/2}} e^{-\frac{1}{2}(\mathbf{x}_s-\boldsymbol{\mu}_s)'\boldsymbol{\Sigma}_{2s}^{-1}(\mathbf{x}_s-\boldsymbol{\mu}_s)}, \qquad (2)$$

where the subscript 2 indicates the absence of SNM.

At the passive detection stage, the decision is to classify a particular container type to be either regular or suspicious, where a suspicious container will be further investigated by more stringent manual inspection. For the single-port case, as illustrated in Fig. 3, this decision is made by implementing a single-value threshold (t_p in Fig. 3) achieving a sensible tradeoff between the false alarm and missed detection probabilities. The same philosophy applies to the two-port problem, except that we now have a two-dimensional (2D) decision space. Consequently, a decision boundary curve, rather than a single value, is sought to split the decision space into two regions R_1 and R_2, such that depending on which region a pair of observed gross particle counts falls into, a proper disposition of the corresponding container can be reached; please see Fig. 7 for illustration, where the thick-lined curve is a decision boundary splitting the decision space. Note that this particular curve is arbitrarily drawn for the sake of illustration.

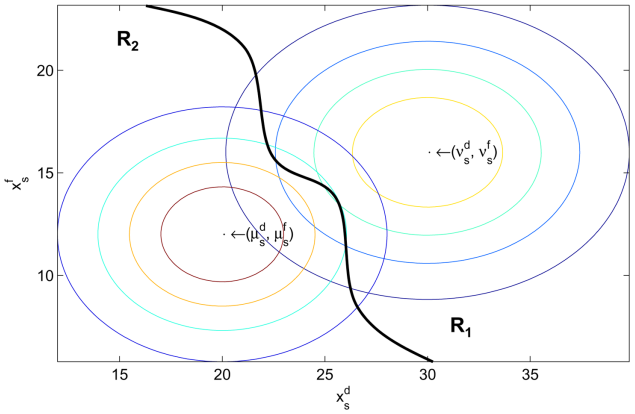

Figure 7. An arbitrary decision boundary in a 2D decision space. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

For simplicity, we will assume that $\mathrm{Cov}(X_s^d, X_s^f) = 0$ in the above covariance matrices. Note that the two detection processes happen independently: hence when SNM is absent, the background radiation is uncorrelated. When SNM is present and the SNM signal is small, the detection processes act like two random draws from a population of large variance, leading naturally to a small correlation between them. In our numerical example (described in Section 4.1), the correlation between measurements at foreign and domestic ports for a "hard" container type ($s = 3$), is 0.18, justifying the independence assumption. For "softer" container types with larger SNM signal, correlations are higher, as expected. However, due to the max-min formulation of the inspection policies (see Section 3.5), the hard container types determine the detection probability and therefore the inspection policy. Thus, we assume uncorrelatedness for simplicity.

3.2. Rectangular Boundary Curve Policy: Considering Two Detection Events Separately

A straightforward way of classifying a container type into "regular" versus "suspicious" is to set individual thresholds for each detector separately. Denote by q_s^d and q_s^f, respectively, individual thresholds for each detector, such that if a container type has either one of its gross particle counts greater than the corresponding threshold, namely either $x_s^d \geq q_s^d$ or $x_s^f \geq q_s^f$, then the container type is deemed "suspicious"; otherwise, it is treated as "regular." The resulting decision boundary is of a rectangular shape; see Fig. 8.

Same as in the single-port case, the optimal decision involves finding the best trade off between the missed detection and false alarm probabilities. For the case of a rectangular

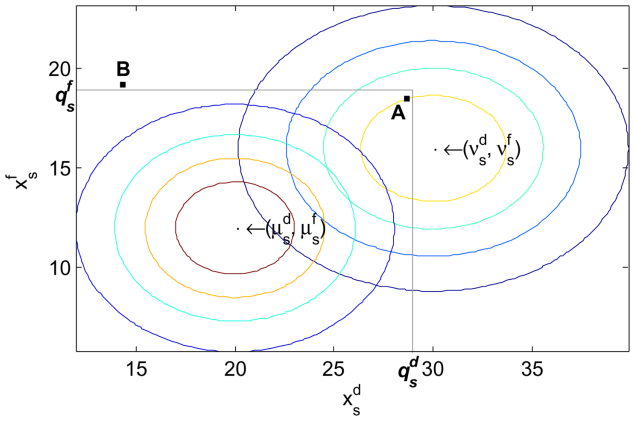

Figure 8. Rectangular decision boundary. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

decision boundary curve, the missed detection probability can be computed as:

$$\int_0^{q_s^d} \int_0^{q_s^f} f_{1s}(\mathbf{x}_s)\, \mathrm{d}x_s^d\, \mathrm{d}x_s^f, \tag{3}$$

when SNM is present.

The false alarm probability is calculated as:

$$1 - \int_0^{q_s^d} \int_0^{q_s^f} f_{2s}(\mathbf{x}_s)\, \mathrm{d}x_s^d\, \mathrm{d}x_s^f, \tag{4}$$

when SNM is absent.

In Section 3.5, we present an optimization formulation where these two measures are incorporated, and its solution yields the optimal rectangular decision boundary; we call the resulting inspection policy the rectangular boundary curve policy (RBCP).

3.3. Proposed BCP: Considering Two Detection Events Simultaneously

Using the rectangular decision, boundary curve is the same as combining the two single-port decisions in a sequential fashion. Doing this, however, may lead to counterintuitive and suboptimal decision-making. Consider Point A in Fig. 8, which represents a container having relatively high-gross particle counts at both detectors. This should have raised enough suspicion on this particular container. But because Point A is still within the rectangular decision boundary, the container will be deemed as "regular." Consider another Point B, corresponding to a container for which the detector at the foreign port returns a gross particle count slightly larger than that of A, while the domestic one returns a very small count. Yet, the container represented by B would have been deemed more suspicious than that of A. This simple example demonstrates a major drawback of the sequential decision making practice—it allows too much margin at the right-top area close to the decision boundary, where common sense indicates that more stringent inspection is warranted.

Given that we use the normal approximation to model the bivariate detection outcome $\mathbf{x}_s$, we propose that a sensible decision boundary is to minimize the expected cost of misclassification (ECM) [11], a criterion widely used in statistics and decision theory. The minimized ECM naturally leads to the optimal trade-off between missed detections and false alarms associated with the detection process.

The ECM is defined as

$$\mathrm{ECM} = c(2|1)P(2|1)p_1 + c(1|2)P(1|2)p_2, \tag{5}$$

where $c(2|1)$ and $P(2|1)$ are the cost and the probability of missed detections, p_1 is the prior probability that SNM is present in a container, $c(1|2)$ and $P(1|2)$ are the cost and the probability of false alarms, and p_2 is the prior probability of SNM is absent in the container; $p_1 + p_2 = 1$. Johnson and Wichern [11] provide the following conditions that minimize the ECM, leading to the optimal splitting between regions R_1 and R_2:

$$R_1 : \frac{f_{1s}(\mathbf{x}_s)}{f_{2s}(\mathbf{x}_s)} \geq \left(\frac{c(1|2)}{c(2|1)}\right) \cdot \left(\frac{p_2}{p_1}\right)$$
$$\text{i.e.,} \begin{pmatrix}\text{density}\\ \text{ratio}\end{pmatrix} \geq \begin{pmatrix}\text{cost}\\ \text{ratio}\end{pmatrix} \cdot \begin{pmatrix}\text{prior}\\ \text{probability}\\ \text{ratio}\end{pmatrix}, \tag{6}$$

and R_2 is the complement of region R_1.

Using different cost ratios and prior probabilities, one can generate a series of decision boundaries, parallel to each

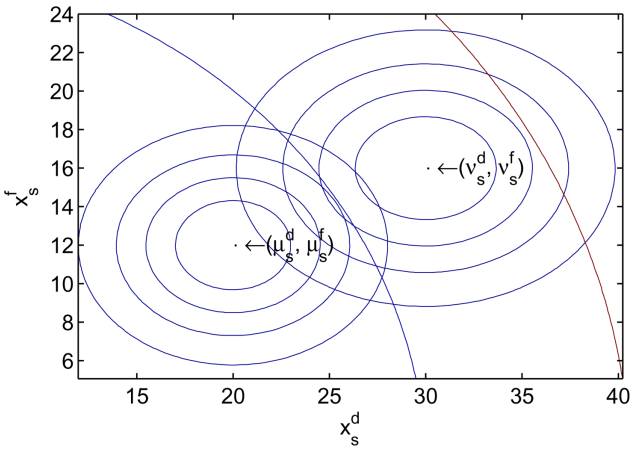

Figure 9. Proposed decision boundary curve. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

other, as shown in Fig. 9. The shape of each of these decision boundaries is elliptic, distinctively different from the rectangular ones. Like the RBCP, the ultimate inspection policy based on this new decision boundary will be obtained by solving an optimization problem, to be introduced in Section 3.5; we will refer to the resulting optimal decision boundary as the boundary curve policy (BCP).

One difficulty this proposed decision boundary brings is that the minimized ECM depends on the choices of misclassification costs and prior probabilities of events, which are in reality difficult to decide. But inspecting the above inequalities, one will notice that the right hand sides, namely the multiplication of the cost ratio and the prior probability ratio, do not depend on the actual reading of the detectors, manifested in $\mathbf{x}_s$. As such, we can treat the multiplication of the ratios as a constant, but unknown, quantity. Denote by τ_s the natural logarithm of this constant, that is, $\tau_s = \ln\left[\left(\frac{c(1|2)}{c(2|1)}\right)\cdot\left(\frac{p_2}{p_1}\right)\right]$. We then include this τ_s as one of the decision variables in the optimization problem to be presented in Section 3.5.

Next, we simplify the left-hand side in the above inequalities, which is the likelihood ratio between two pdf's. Let r_s be the natural logarithm of this likelihood ratio, namely:

$$r_s(\mathbf{x}_s) = \ln\left[\frac{f_{1s}(\mathbf{x}_s)}{f_{2s}(\mathbf{x}_s)}\right] = -\frac{1}{2}\mathbf{x}_s'\left(\boldsymbol{\Sigma}_{1s}^{-1} - \boldsymbol{\Sigma}_{2s}^{-1}\right)\mathbf{x}_s + \left(\boldsymbol{\nu}_s'\boldsymbol{\Sigma}_{1s}^{-1} - \boldsymbol{\mu}_s'\boldsymbol{\Sigma}_2^{-1}\right)\mathbf{x}_s - k, \tag{7}$$

where

$$k = \frac{1}{2}\ln\left(\frac{|\boldsymbol{\Sigma}_{1s}|}{|\boldsymbol{\Sigma}_{2s}|}\right) + \frac{1}{2}\left(\boldsymbol{\nu}_s'\boldsymbol{\Sigma}_{1s}^{-1}\boldsymbol{\nu}_s - \boldsymbol{\mu}_s{}'\boldsymbol{\Sigma}_{2s}^{-1}\boldsymbol{\mu}_s\right) = \frac{1}{2}(\ln(\nu_s^d\nu_s^f) - \ln(\mu_s^d\mu_s^f)) + \frac{1}{2}(\nu_s^d + \nu_s^f - \mu_s^d - \mu_s^f).$$

The resulting BCP is as follows. Given a new pair of gross particle counts $\mathbf{x}_s$ for container type s, if the corresponding $r_s \geq \tau_s$, then this container type is classified as "suspicious"; otherwise, it is classified as "regular."

The missed detection probability associated with the BCP can be computed as:

$$\int_{r_s<\tau_s} f_{1s}(\mathbf{x}_s)\,d\mathbf{x}_s = 1 - \int_{r_s\geq\tau_s} f_{1s}(\mathbf{x}_s)\,d\mathbf{x}_s = 1 - P(r_s(\mathbf{X}_s) \geq \tau_s|\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\nu}_s, \boldsymbol{\Sigma}_{1s})). \tag{8}$$

The false alarm rate of the BCP is:

$$\int_{r_s\geq\tau_s} f_{2s}(\mathbf{x}_s)\,d\mathbf{x}_s = P(r_s(\mathbf{X}_s) \geq \tau_s|\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_{2s})). \tag{9}$$

To obtain the above two probabilities, the conditional tail probability $P(r_s(\mathbf{X}_s) \geq \tau_s|\mathbf{X}_s)$ will need to be derived. In the following, we show how to derive the conditional tail probability when $\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\nu}_s, \boldsymbol{\Sigma}_{1s})$, that is, when SNM is present. The tail probability when $\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_{2s})$ can be obtained likewise.

We introduce a new random variable Y_s, defined as: $Y_s := -\frac{1}{2}\mathbf{X}_s'(\boldsymbol{\Sigma}_{1s}^{-1} - \boldsymbol{\Sigma}_{2s}^{-1})\mathbf{X}_s + (\boldsymbol{\nu}_s'\boldsymbol{\Sigma}_{1s}^{-1} - \boldsymbol{\mu}_s'\boldsymbol{\Sigma}_{2s}^{-1})\mathbf{X}_s$. Then, the conditional tail probability is:

$$\begin{aligned} P(r_s(\mathbf{X}_s) \geq \tau_s|\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\nu}_s, \boldsymbol{\Sigma}_{1s})) &= P(Y_s - k \geq \tau_s|\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\nu}_s, \boldsymbol{\Sigma}_{1s})) \\ &= P(Y_s \geq \tau_s + k|\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\nu}_s, \boldsymbol{\Sigma}_{1s})). \end{aligned} \tag{10}$$

Noticing $\boldsymbol{\nu}_s'\boldsymbol{\Sigma}_{1s}^{-1} - \boldsymbol{\mu}_s'\boldsymbol{\Sigma}_{2s}^{-1} = [\nu_s^d \quad \nu_s^f]\cdot\begin{bmatrix}\frac{1}{\nu_s^d} & 0\\ 0 & \frac{1}{\nu_s^f}\end{bmatrix} - [\mu_s^d \quad \mu_s^f]\cdot\begin{bmatrix}\frac{1}{\mu_s^d} & 0\\ 0 & \frac{1}{\mu_s^f}\end{bmatrix} = [0 \quad 0]$, we can simplify Y_s to:

$$\begin{aligned} Y_s &= -\frac{1}{2}\mathbf{X}_s'(\boldsymbol{\Sigma}_{1s}^{-1} - \boldsymbol{\Sigma}_{2s}^{-1})\mathbf{X}_s \\ &= \frac{1}{2}\left(\frac{1}{\mu_s^d} - \frac{1}{\nu_s^d}\right)(X_s^d)^2 + \frac{1}{2}\left(\frac{1}{\mu_s^f} - \frac{1}{\nu_s^f}\right)(X_s^f)^2. \end{aligned} \tag{11}$$

The above equation shows that Y_s is a quadratic function of the bivariate normal random variables in $\mathbf{X}_s$, suggesting that Y_s follows a generalized χ^2 distribution. This is expected because given $\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\nu}_s, \boldsymbol{\Sigma}_{1s})$, we know that $\frac{(X_s^i)^2}{\nu_s^i}$, for $i = d$ or f, follows a noncentral chi-square distribution $\chi^2_{\omega_s^i}(\delta_s^i)$, with degree of freedom $\omega_s^i = 1$ and noncentral parameter $\delta_s^i = \frac{(\nu_s^i)^2}{\nu_s^i} = \nu_s^i$ [12]. So Y_s is basically the summation of two noncentral χ^2 variables with their respective coefficients. As such, we can obtain the following:

$$\begin{aligned} Y_s &= \frac{1}{2}\left(\frac{1}{\mu_s^d} - \frac{1}{\nu_s^d}\right)(X_s^d)^2 + \frac{1}{2}\left(\frac{1}{\mu_s^f} - \frac{1}{\nu_s^f}\right)(X_s^f)^2 \\ &\sim \lambda_s^d\,\chi^2_{\omega_s^d}(\delta_s^d) + \lambda_s^f\,\chi^2_{\omega_s^f}(\delta_s^f) \end{aligned} \tag{12}$$

where $\lambda_s^i = \frac{1}{2}\cdot(\frac{1}{\mu_s^i} - \frac{1}{\nu_s^i})\cdot\nu_s^i$, for $i = d$ or f.

Likewise, we can compute the tail probability $P(r_s(\mathbf{X}_s) \geq \tau_s|\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_{2s}))$ by using two similar noncentral χ^2 distributions. The difference is that when $\mathbf{X}_s \sim \mathcal{N}(\boldsymbol{\mu}_s, \boldsymbol{\Sigma}_{2s})$, $Y_s \sim \lambda_s^d\,\chi^2_{\omega_s^d}(\delta_s^d) + \lambda_s^f\,\chi^2_{\omega_s^f}(\delta_s^f)$, where $\delta_s^i = \mu_s^i$ and $\lambda_s^i = \frac{1}{2}\cdot(\frac{1}{\mu_s^i} - \frac{1}{\nu_s^i})\cdot\mu_s^i$, for $i = d$ or f.

For computing the above noncentral χ^2 distributions, we use the approximation developed by Liu et al. [12], which provides superior accuracy and performance compared to Pearson's approximation method [18].

3.4. Hardness Measure using Prior Radiography Information

In the two-port problem, the hardness of a container type at the domestic port is determined based on the radiography information obtained at both foreign and domestic ports.

The hardness of a container type is still defined as the smallest misclassification error, but now under two bivariate normal distributions. According to the results from Johnson and Wichern [11] regarding the ECM, the misclassification error can be computed when the detection threshold is set to $\tau_s = 0$ [or equivalently, the right-hand side of Inequality (6) equals 1]. The decision boundary curve corresponding to $\tau_s = 0$ is the curve on which the two bivariate pdfs intersect with each other.

We again use the MCNP software package to simulate a given passive detection process. The difference is that we now use the radiography information of the same container type obtained at both ports. This is equivalent to simulating the two bivariate normal density functions, as illustrated in the previous subsection (Fig. 9), with and without SNM in the container. Let the resulting simulated pdfs be denoted by $\tilde{f}_{1s}$ (for SNM present) and $\tilde{f}_{2s}$ (for SNM absent). Then the new hardness $h_s(F, D)$ at the domestic port is given by:

$$\begin{aligned} h_s(F, D) &= \text{missed detection rate} + \text{false alarm rate} \\ &= \int_{r_s<0} \tilde{f}_{1s}(\mathbf{x}_s)\, d\mathbf{x}_s + \int_{r_s\geq 0} \tilde{f}_{2s}(\mathbf{x}_s)\, d\mathbf{x}_s \end{aligned} \quad (13)$$

3.5. Optimization Problem

The setup of the optimization problem that solves for the decision variables is similar to the formulations used in Gaukler et al. [9, 10]. We briefly summarize the major considerations.

We use a simulation (MCNP) to find the appropriate distributions of gross particle counts for containers with and without SNM. We then evaluate the performance of the combined inspection system by solving a nonlinear optimization problem using a genetic algorithm. Details of the optimization problem setup are provided in Section 4.1.

The objective is to maximize the worst-case detection probability (DP) among all possible container types (a max-min problem), subject to two constraints. The first constraint is the expected delay time (DT), which is the real-life consequence of having false alarms. The delay time can be computed through modeling the queueing network associated with the inspection processes, and is a function of the arrival rates and false alarm probabilities at each node. This queueing network model has been developed and described in detail in earlier research; for details, please refer to Gaukler et al. [9, 10]. The second constraint is on the utilization rate at the manual node, denoted by ρ_M. The second constraint is motivated by the observation that any inspection policy attempts to use the manual inspection as much as possible for obvious reasons. Yet the manual inspection is the most time-consuming step and takes significantly longer to complete than any other inspection steps. As a consequence, the queueing system is very sensitive to even the smallest unanticipated changes to the arrival stream and could easily become unstable. Putting the constraint on ρ_M helps obtain a more robust and practicable system.

In this article, we optimize inspection policies under three settings: the single-port policy (SP) that ignores the container inspection history, the two-port policy that uses the rectangular decision boundary (RBCP), and the two-port policy that is derived using the expected misclassification cost (BCP). SP and RBCP serve as the alternatives to the BCP. In Section 4, we will demonstrate the merit of BCP by comparing its performance with that of SP and RBCP.

The three inspection policies use a similar set of decision variables, which include:

- the threshold(s) at the passive detection node. For BCP, τ_s is the threshold; for RBCP, the passive detection has two thresholds q_s^d and q_s^f; for SP, the passive detection threshold is represented by γ_s.
- the threshold t_R at the radiography node; and
- a randomization factor a, $0 < a < 1$.

The randomization factor works as follows: if the hardness measure for container type s is higher than the prescribed threshold t_R, then a randomly selected proportion a of container type s is elevated to manual inspection. If a container is not selected for manual inspection, it will be sent to passive inspection. This randomization factor a was introduced in Gaukler et al. [10] to produce a smooth operation of inspections when the number of container types is relatively small. Consider the situation when there are only four container types. As a result, t_R can only take on five discrete values. Without the randomization factor a, t_R would be set at 1 (the maximum possible value) whenever the volume of the hardest container types is too large for the manual inspection stage to handle. Introducing the randomization factor a gives the system the flexibility to escalate a subset of containers whose $h_s \geq t_R$, the volume of which can be handled within the capacity restriction of the manual inspection stage.

With the above thoughts and notations, we present the optimization formulation below. For the two-port problem with the proposed elliptic decision boundary (BCP), the optimization formulation is:

$$\begin{aligned} &\max_{t_R, a, \tau_s} \min_s \mathrm{DP}_s^{\mathrm{BCP}} \\ &s.t. \quad \mathrm{DT}^{\mathrm{BCP}} \leq T_0, \\ &\rho_M^{\mathrm{BCP}} \leq \rho_0, \end{aligned} \quad (14)$$

where T_0 is the limit set for the system delay time, and ρ_0 is the utilization limit for the manual inspection.

For the single-port problem (SP), the optimization formulation is:

$$\max_{t_R,a,\gamma_s} \min_{s} \mathrm{DP}_s^{\mathrm{SP}}$$
$$s.t. \quad \mathrm{DT}^{\mathrm{SP}} \le T_0,$$
$$\rho_M^{\mathrm{SP}} \le \rho_0. \tag{15}$$

For the two-port problem with the rectangular decision boundary curve (RBCP), the formulation is:

$$\max_{t_R,a,q_s^d,q_s^f} \min_{s} \mathrm{DP}_s^{\mathrm{RBCP}}$$
$$s.t. \quad \mathrm{DT}^{\mathrm{RBCP}} \le T_0, \qquad \rho_M^{\mathrm{RBCP}} \le \rho_0. \tag{16}$$

In the above optimization problems, $0 < \rho_0 < 1$ is a fixed constant chosen *a priori*. Each optimization problem is solved sequentially for different values of delay time T_0, using a genetic algorithm [13]. Ultimately, this sequence of optimizations produces for each inspection process an efficient frontier between the detection probability and the delay time. We use the efficient frontier to characterize the performance of individual inspection policies. Which decision variable setting to choose out of an efficient frontier in an actual implementation is left for a policy maker to decide.

4. ANALYSIS OF INSPECTION SYSTEMS

In this section, we conduct a numerical analysis of the afore-mentioned three inspection policies and compare their performance. First, we discuss the setup of the numerical study. Then, we present the performance comparison among all three policies. Last, we discuss additional implications if all three policies are implemented in the standard version of the inspection systems, which does not use radiography equipment (i.e., the inspection system shown in Fig. 2).

4.1. Setup of the Numerical Study

In this study, without loss of generality, we assume that the two ports involved use the same detection equipment and have the same level of background radiation, absent any SNM in a container. We set the arrival rate of containers at the domestic port as 90 per h, the same as in Wein et al. [23]. On arrival, the containers first go through the ATS step, which does not consume any time in the queueing system because the ATS scores can be assigned before a container ship arrives at the port. We assume that the percentage of "high-risk" containers that is directly escalated to manual inspection is 2%.

The next step is the radiography stage. We assume there are three radiography machines at each port, each of which takes an average scan time of 90 s. In literature, the imaging and analysis time of stationary X-ray inspection ranges from 75 s to 90 s [20]. We model two passive detectors at each port, each of which takes an average scan time of 45 s. Those containers that are flagged by either the ATS, the radiography stage, or a passive detector, will go through manual inspection, which is assumed to take an average of 1 h. We assume there are six manual inspection teams that can work in parallel. The inspection time of either the radiography machines, the passive detectors, or the manual inspection stage is modeled as following an exponential distribution. We set the maximum allowable capacity utilization for the manual inspection stage to be $\rho_0 = 0.95$; the same value is used in all three inspection models.

Table 1. The port operation parameters.

Parameter	Description	Value
λ	Arrival rate of container	90/h
δ	Percentage of containers elevated to manual inspection by ATS	2%
μ_R	Service rate at radiography node (R-node)	40/h
m_R	Number of servers at R-node	3
μ_P	Service rate at passive detection node (P-node)	80/h
m_P	Number of servers at P-node	2
μ_M	Service rate at manual detection node (M-node)	1/h
m_M	Number of servers at M-node	6
ρ_0	Maximum utilization at M-node	0.95

We further assume that when a container with SNM is manually inspected, the probability that the SNM is discovered is one. In reality, it is certainly possible that a manual team could miss the SNM (see the story in Scientific American [4]). When the detection probability at the manual inspection is less than one, then all the detection probabilities and efficient frontiers that we report need to be scaled by a factor equaling the manual detection probability, but the general insights arrived at in this study will remain.

The choices of the above parameters are consistent with those used in previous studies, for example in Wein et al. [23] and Gaukler et al. [9, 10]. To the best of our knowledge, these parameter choices reflect a reasonable setup and operation of the inspection steps at a port. The set of port operation parameters is summarized in Table 1.

Under the advice of our nuclear engineering collaborators, we model four container types in this analysis. The first container type consists of only low Z-value materials, whose Z-value is less than 10 (e.g., textiles, plastic, wood). The contents of the second container type is a mixture of medium Z-value materials, whose Z-value is between 10 and 20 (e.g., aluminum), and low Z-value materials. The third container type has both high Z-value materials, whose Z-value is greater

than 20 (e.g., steel), and low Z-value materials. The fourth and last container type is laden with naturally occurring radiative materials (NORM), such as fertilizer. To account for the effect of NORM, the background emission level of the fourth container type is set to be five times that of the first three container types.

The four container types do not arrive at a port with equal proportion. We estimate the proportion of each container type, denoted by p_s, based on a 2007 listing of the top 100 U.S. container importers and the industry (and thus cargo) segments represented by these top 100 importers [19]. It turns out that there are relatively fewer hard containers than there are soft containers. For all four container types, MCNP is used to obtain the background radiation particle counts as well as the particle counts when SNM is present. Here, our goal is to detect the presence of a solid ball of HEU (40% U-235 and 60% U-238), with one centimeter of lead shielding around it. This shielded SNM is strategically located in the center of the highest-Z-value area of the container. Once MCNP obtains the two pdfs, the hardness associated with each container type can be computed. Table 2 lists the container types, their proportion in the arrival stream, the gross particle count means obtained from the MCNP simulations, and the hardness measures. In this table, we report the hardness based on information obtained only at the foreign or domestic ports (h_s^f and h_s^d, respectively), as well as the hardness based on information obtained at both ports ($h_s(F, D)$).

Table 2. Container information, simulated gross counts, and hardness.

s	p_s	μ_s^d	μ_s^f	ν_s^d	ν_s^f	h_s^d	h_s^f	$h_s(F, D)$
1	0.60	10	10	29.52	29.52	0.022	0.022	0.001
2	0.30	10	10	15.75	15.75	0.416	0.416	0.247
3	0.08	10	10	12.10	12.10	0.749	0.749	0.649
4	0.02	50	50	52.06	52.06	0.888	0.888	0.842

Using the data in the above table, we generate the rectangular decision boundaries and the proposed elliptic decision boundaries for each container type; please see Fig. 10. Multiple decision boundaries are plotted, corresponding

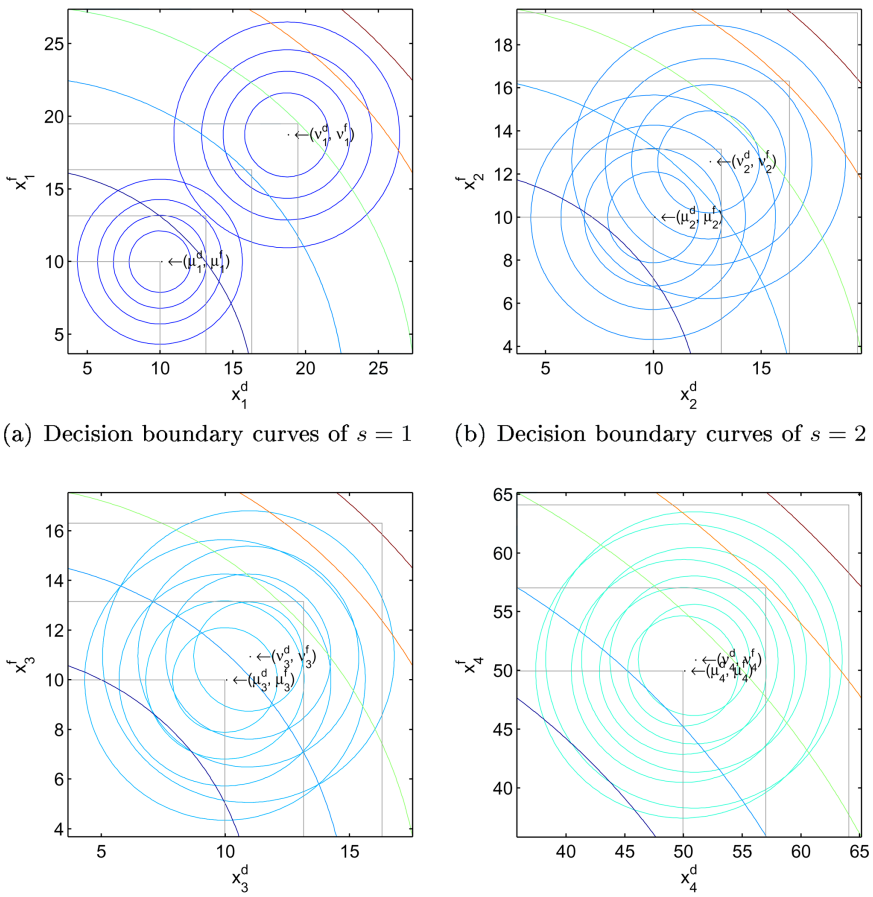

(a) Decision boundary curves of $s = 1$ (b) Decision boundary curves of $s = 2$

(c) Decision boundary curves of $s = 3$ (d) Decision boundary curves of $s = 4$

Figure 10. Decision boundary curves for each container type. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

to different tradeoffs between missed detections and false alarms. Obviously, container type 4 is the hardest container as the overlap between the two pdf is enormous. According to Table 2, this overlap translates quantitatively to a hardness of 0.842.

4.2. System Comparison

We use the optimization formulations in Section 3.5 to solve for the efficient frontiers associated with the three inspection policies (i.e., SP, RBCP, and BCP). The resulting efficient frontiers are presented in Fig. 11.

The efficient frontier of BCP dominates those of RBCP and SP. The difference in terms of detection probability between BCP and SP for a large portion of delay time is about 10 percentage points. Considering that the SP has a detection probability less than 0.6, the 10 percentage points translates to an improvement of more than 17% . This outcome supports our earlier conjecture that using the container inspection history enhances the overall detection capability. RBCP also improves over SP but not as much as BCP. This echoes our argument that using the detection data sequentially for decision-making (as in RBCP) is not as good as using them simultaneously.

The above analysis is conducted based on the assumption that radiography equipment is available at both ports. We acknowledge that in current practice, the standard version of the inspection system without radiography capability, as outlined in Fig. 2, is still much more common than the hybrid inspection system with radiography equipment. Hence we

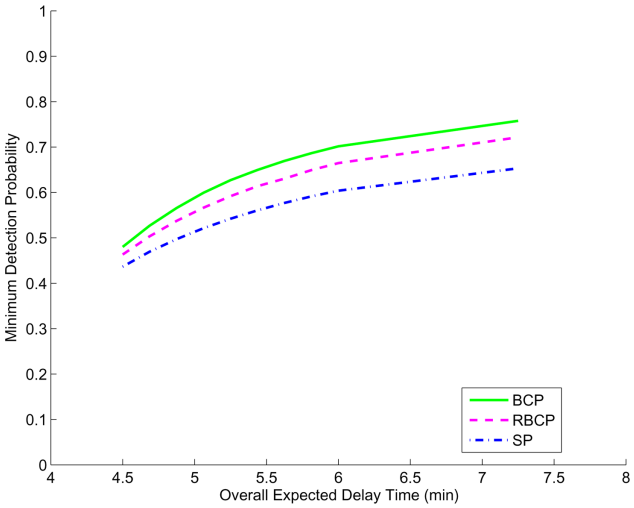

Figure 11. Efficient frontiers of the three inspection policies when radiography machines are available at both ports. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

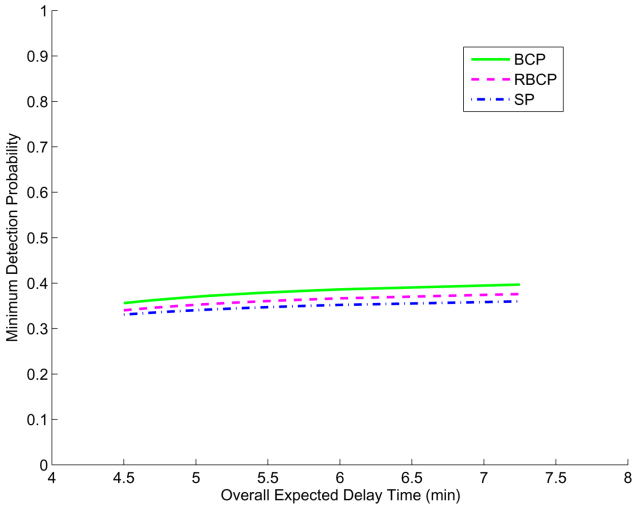

Figure 12. Efficient frontiers of the three inspection policies when radiography machines are not available at either port. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

are interested in seeing what difference, if any, prior inspection data may make on the standard version of the inspection system.

The absence of radiography information eliminates the possibility of classifying containers into different hardness types. We will have to treat all the containers as coming from one pooled population, a mixture of the distributions of distinct container types, weighted by their respective proportions p_s. The resulting efficient frontiers of the same three inspection policies are shown in Fig. 12. Although BCP still outperforms RBCP and SP, the benefit diminishes dramatically without the radiography information. On top of that, all policies attain significantly reduced detection probabilities: for instance, at a delay time of 7 min, the detection probability of SP decreases from around 0.6 with radiography information to less than 0.4 without radiography information.

This result is not really surprising, because without radiography information or the classification of container types, the inspection system loses much if its power of discrimination and is is less effective in allocating its most precious resources. With all containers pooled together, the variability of the gross particle counts increases significantly. The impact of the extra detection data from the previous port is inevitably dwarfed by the presence of a large variability. Thus, this analysis also supports the inclusion of the radiography machines in the detection process and the use of the container type as an informative proxy to focus the limited resources on those containers that need most of the attention, an argument initially made by Gaukler et al. [9] for the single-port case.

5. TWO-PORT PROBLEM CONSIDERING THE POSSIBILITY OF INFILTRATION

The analysis presented in Section 4 assumes that the contents in the containers are not altered en route from one port to the other. In this case, it intuitively makes sense to use the prior detection data because for a container in which SNM is present, its two gross particle counts are likely both large, while for a container without SNM, its two counts are likely both small. This pattern explains the intuition why BCP and RBCP perform better than SP, because another reading of a large gross particle count helps reinforce the belief that SNM is present.

But what if the contents in a container may be altered during transit? More specifically, we want to consider the case where a container is infiltrated, that is, a certain amount of SNM is placed and hidden in a non-SNM container some time after the container is loaded onto the ship at the foreign port, but before it starts the inspection process at the domestic port. Would considering the prior detection data mislead us to the degree that we would be better off completely ignoring the inspection history?

We believe that the answer depends on the likelihood of infiltration. We presented the question of how likely it might be that an adversary could infiltrate a container during transit to a group of colleagues who work on port security and container inspection problems. Although a majority of the respondents believed that infiltrating a container during transit is unlikely for the reason that doing so is not as easy as generally imagined, some others did believe that this could possibly be done if a smart adversary was determined to do so. The existing port security literature is silent on this aspect, possibly because there have never been any recorded incidents of this type of infiltration in the past. This lack of records to some degree supports the claim that the likelihood of infiltration is small.

There are two possible ways that SNM may get into a container: either before the container is loaded onto the ship at the foreign port or in between the two ports ("during transit"); see Fig. 13.

To represent a broad spectrum of possibilities, we choose to model infiltration using a simple probability parameter. Denote by P_I the likelihood of infiltration during transit, and $1 - P_I$ the likelihood of SNM getting into a container before embarkation at the foreign port. We here only consider the possibility that an adversary puts SNM into a non-SNM container rather than someone taking SNM out of a container. We also assume that the content change is small enough so that the container type remains the same.

We need to slightly modify the optimization formulation presented earlier to incorporate the likelihood of infiltration. Because infiltration only happens to a certain container, it will not change the false alarm probability. Consequently, the expected delay time is the same with or without infiltration.

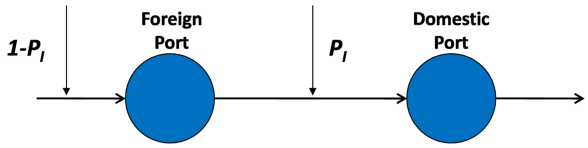

Figure 13. Possibility of infiltration. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

Existence of infiltration does impact the detection probability of an inspection system, however. We use $\mathrm{DP}_s^{\mathrm{BCP},I}$ and $\mathrm{DP}_s^{\mathrm{BCP},nI}$, respectively, to present the detection probability of BCP for container type s with and without infiltration. The final detection probability of the inspection system is weighted by the infiltration probability, P_I.

The optimization problem for BCP considering the possibility of infiltration is then described as:

$$\max_{t_R,a,\tau_s} \min_{s} \left[\mathrm{DP}_s^{\mathrm{BCP},I} \cdot P_I + \mathrm{DP}_s^{\mathrm{BCP,nI}} \cdot (1 - P_I)\right]$$

$$s.t. \quad \mathrm{DT}^{\mathrm{BCP}} \le T_0 \tag{17}$$

$$\rho_M^{\mathrm{BCP}} \le \rho_0 \tag{18}$$

where ρ_0 is still chosen as 0.95.

The same change can be done to the optimization for RBCP; for brevity, we omit the detailed expression. Because SP does not use the container inspection history in the first place, its detection outcome will not change in the presence of infiltration, and we need not modify the previous SP formulation.

The efficient frontiers from the solution of the revised optimization formulations are presented in Fig. 14. Fig. 14a shows the efficient frontiers for the same parameter settings as in Fig. 11. We can observe that when the probability of infiltration increases, the benefit of using prior inspection information diminishes, as one might expect. When infiltration becomes a certainty (Fig. 14f), it is clearly better not to use prior detection data. Interestingly, and somewhat unexpectedly though, the results also show that the benefit of using prior detection data in BCP does not go away until the probability of infiltration is $P_I > 0.5$. We consider this a pleasant surprise because it suggests a good degree of robustness of the proposed BCP, especially considering that many experts we asked believe that P_I is very small, if not next to nothing.

In accordance with our discussion in Section 4.2, we also investigate what happens if the hybrid inspection system with radiography is replaced with the standard version without radiography. It turns out that the insight garnered from Section 4.2 remains: without radiography, the difference among the inspection policies with and without prior detection information becomes smaller, and the detection probabilities of all the policies decrease considerably; please see Fig. 15 for the efficient frontiers of the three policies. This result reflects the nonideal state of the current practice of port security.

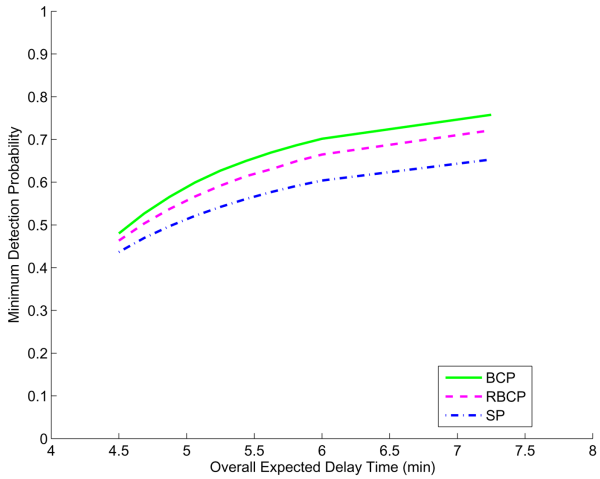

(a) No Infiltration: $P_I = 0$

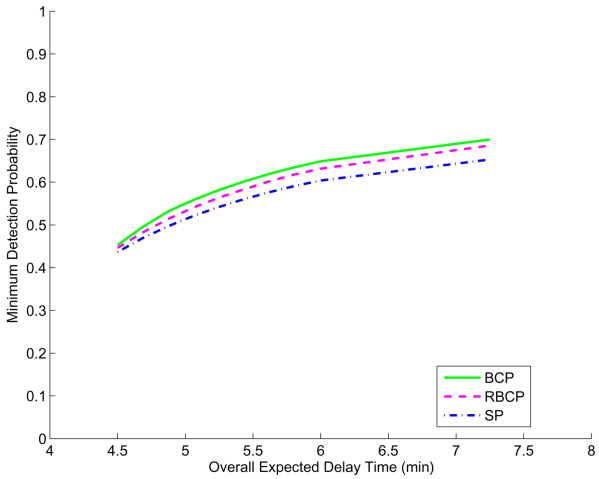

(b) $P_I = 0.36$

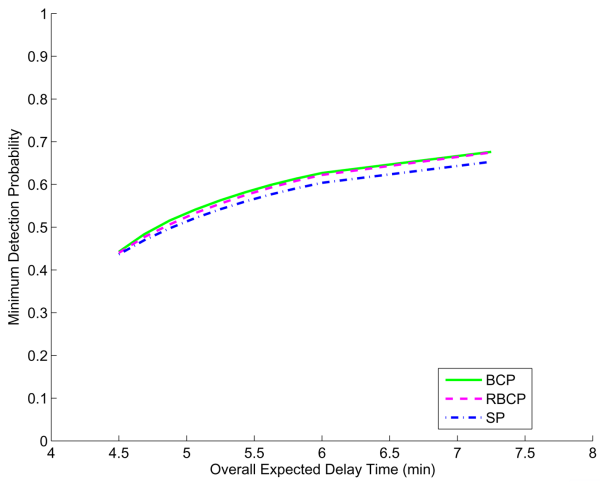

(c) $P_I = 0.5$

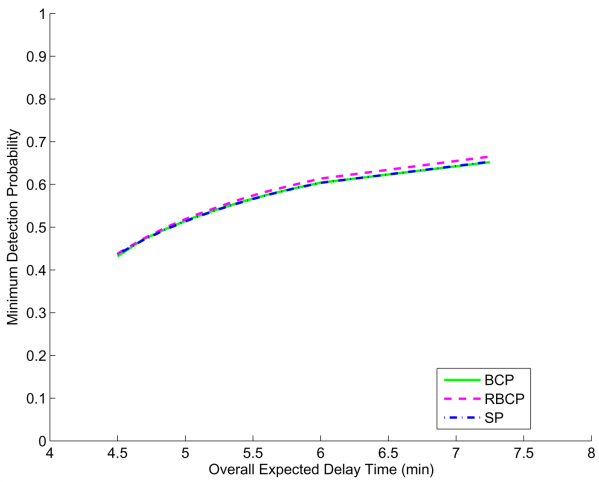

(d) $P_I = 0.64$

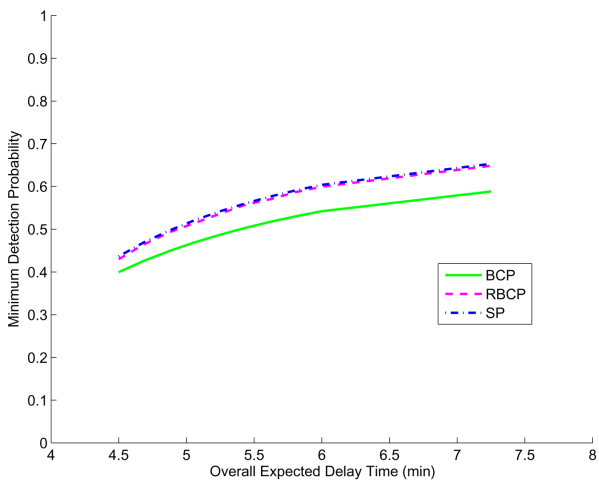

(e) 100% Infiltration: $P_I = 1$

Figure 14. Efficient frontiers of the three inspection policies under different infiltration probability P_I when radiography machines are available at both ports. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

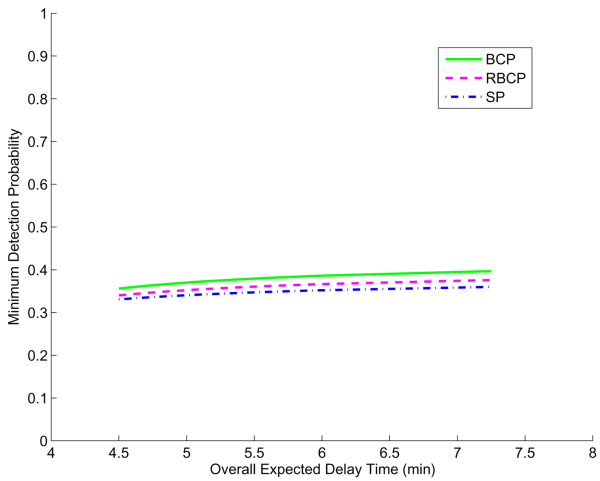

(a) No Infiltration: $P_I = 0$

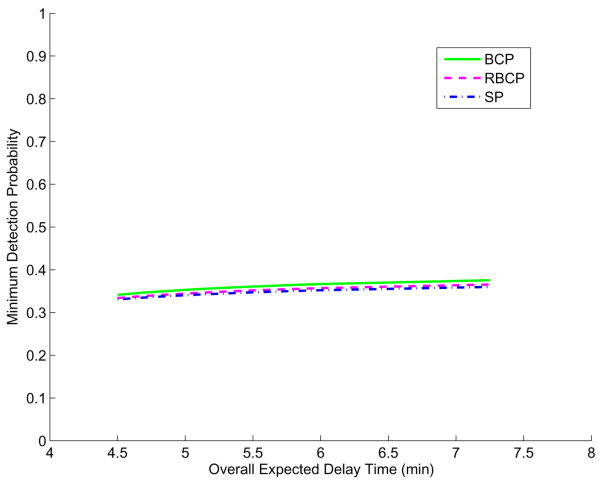

(b) $P_I = 0.36$

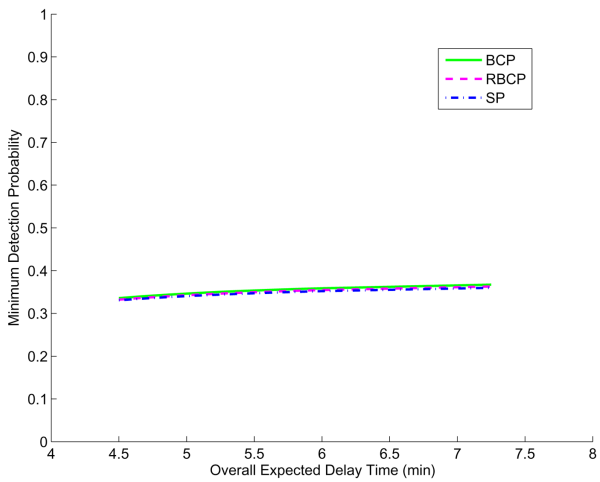

(c) $P_I = 0.5$

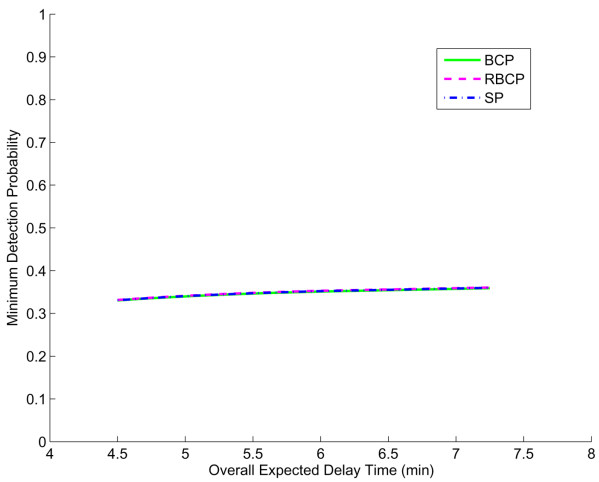

(d) $P_I = 0.64$

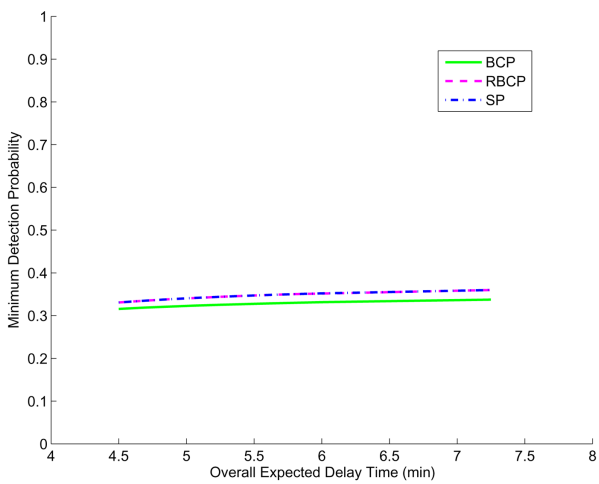

(e) 100% Infiltration: $P_I = 1$

Figure 15. Efficient frontiers of the three inspection policies under different infiltration probability P_I when radiography machines are not available at either port. [Color figure can be viewed in the online issue, which is available at wileyonlinelibrary.com.]

6. CONCLUDING REMARKS

The United States are faced with the grave threat of adversaries smuggling SNM into the country for nefarious purposes. Indeed, once an adversary obtains SNM such as weapons-grade enriched uranium, the step to assembling a nuclear device is a comparatively simple one. The recent development of enrichment facilities in Iran is particularly worrisome, because basic enrichment capability for power generation purposes can easily be extended to weapons grade enrichment: for centrifuge-type equipment as used in Iran, the enrichment percentage is largely a function of centrifuge runtime.

To thwart SNM smuggling efforts, efficient and effective inspection policies at ports and border crossings are necessary. In this article, we propose using prior inspection data as a low-cost way of increasing overall interdiction performance. First, under the assumption that prior inspection data are reliable, we investigate how a decision maker at a domestic port should use detection data from a foreign port to improve the overall detection capability. Then, we discuss the impact of prior inspection data that is potentially unreliable. Prior inspection data, for example, is unreliable if an adversary infiltrates a container downstream of the first inspection. Thus, there is the danger that prior data may harm decision making at the domestic port.

To investigate and quantify these issues, we consider a two-port, seaborne container inspection problem. Our research analyzes activities at the domestic port only, but involves the use of detection data from both foreign and domestic ports. This unifying approach is, to the best of our knowledge, novel and unique in the port security literature. Current work either focuses on a single port without taking into account prior information, or it focuses on multiple interdiction locations, but assumes that there is no sharing of data or information. Thus, our work breaks new ground here.

We derive several inspection policies that make use of prior detection data. In particular, we derive a BCP that is based on the concept of minimizing the ECM. We show that this container inspection history is valuable when used in conjunction with our control policy, and we argue that such information should be used for making the final inspection decision at the domestic port. Our study shows that using the BCP policy offers significant improvement in detection probablity. In particular, we demonstrate that the BCP policy performs best when radiography equipment is available at both ports. However, even in the absence of radiography, BCP offers good improvement over a SP, which does not use container inspection history information. We also find that the BCP policy is robust in the presence of possible infiltration, that is, when an adversary breaches a container while it is in transit from the foreign port to the domestic port.

Thus, our work provides two important suggestions to policy makers: first, radiography equipment should be an integral part of an inspection system; second, with or without radiography equipment, information on the container inspection history should always be used for inspection decisions. With BCP, we provide a robust inspection policy that shows how this prior detection data should be used, with and without radiography equipment, under different expectations on container infiltration.

REFERENCES

[1] G. Aloise, Nuclear nonproliferation: Comprehensive U.S. planning and better foreign cooperation needed to secure vulnerable nuclear materials worldwide, U.S. Government Accountability Office GAO-11-227, 2010.

[2] AS&E, Cargo & vehicle inspection, American Science & Engineering, Inc., Billerica, MA, USA. Available at: http://www.ase.com/products_solutions/cargo_vehicle_inspection.asp, accessed 26 April 2012.

[3] S.L. Caldwell, Supply chain security: Challenges to scanning 100 percent of U.S.-bound cargo containers, U.S. Government Accountability Office GAO-08-533T, January 2008.

[4] T.B. Cochran and M.G. McKinzie, Detecting nuclear smuggling: Radiation monitors at U.S. ports cannot reliably detect highly enriched uranium, which onshore terrorists could assemble into a nuclear bomb, Sci Am 298(4) (2008), 99–104.

[5] DefenseFile, AS&E's Z portal vehicle screening system at USA southwest border, October 2008, Carousel Publishing Group, London, UK. available at: http://www.defensefile.com, accessed 26 April 2012.

[6] N.B. Dimitrov, D.P. Michalopoulos, D.P. Morton, M.V. Nehme, F. Pan, E. Popova, E.A. Schneider, and G.G. Thoreson, Network deployment of radiation detectors with physics-based detection probability calculations, Ann Oper Res 187(1) (2011), 207–228.

[7] E.A. Elsayed, C.M. Young, M. Xie, H. Zhang, and Y. Zhu, Port-of-entry inspection: Sensor deployment policy optimization, IEEE Trans Auto Sci Eng 6(2) (2009), 265–276.

[8] S. Fetter, V.A. Frolov, M. Miller, R. Mozley, O.F. Prilutsky, and S.N. Rodionov, Detecting nuclear warheads, Sci Global Secu 1 (1990), 225–302.

[9] G.M. Gaukler, C. Li, R. Cannaday, S.S. Chirayath, and Y. Ding, Detecting nuclear materials smuggling: Using radiography to improve container inspection policies, Ann Oper Res 187(1) (2011), 65–87.

[10] G.M. Gaukler, C. Li, Y. Ding, and S.S. Chirayath, Detecting nuclear materials smuggling: Performance evaluation of container inspection policies, Risk Anal 32(3) (2012), 531–554.

[11] R.A. Johnson and D.W. Wichern, Applied multivariate statistical analysis, Pearson Prentice Hall, Upper Saddle River, NJ, 2007.

[12] H. Liu, Y. Tang, and H.H. Zhang, A new chi-square approximation to the distribution of non-negative definite quadratic forms in non-central normal variables, Comput Stat Data Anal 53 (2009), 853–856.

[13] MATLAB, R2010b documentation, global optimization toolbox: Using the Genetic Algorithm, 2010, The Mathworks, Inc., Natick, MA, USA, available at: http://www.mathworks.com/help/toolbox/gads/f6010dfi3.html, accessed 26 April 2012.

[14] L.A. McLay, J.D. Lloyd, and E. Niman, Interdicting nuclear material on cargo containers using knapsack problem models, Ann Oper Res 187(1) (2011), 185–205.
[15] X-5 MonteCarloTeam, MCNP—a general Monte Carlo N-particle transport code, version 5, Los Alamos National Labs, X-5 (2003).
[16] D.P. Morton, F. Pan, and K.J. Saeger, Models for nuclear smuggling interdiction, IIE Trans 39 (2007), 3–14.
[17] F. Pan, Stochastic network interdiction: Models and methods. PhD Thesis, University of Texas, Austin, TX, May 2005.
[18] E.S. Pearson, Note on an approximation to the distribution of non-central χ^2 distributions, Biometrika 46 (3/4) (1959), 364.
[19] PIERS, Top 100 U.S. importers via ocean container transport-2007, Manuf Technol News 15(14) (2008), 6.
[20] W.A. Reed and E. Haines, Throughput performance factors in X-ray cargo screening systems, Port Technology International 39 (2006), available at: http://www.porttechnology.org, accessed 20 January 2012.
[21] D.A. Shea, The global nuclear detection architecture: Issues for congress, March 2009. CRS Report for Congress, Congressional Research Services, RL34574.
[22] R.M. Stana, Cargo container inspections: Preliminary observations on the status of efforts to improve the automated targeting system, United States Government Accountability Office Report GAO-06-591T, March 2006.
[23] L.M. Wein, A.H. Wilkins, M. Baveja, and S.E. Flynn, Preventing the importation of illicit nuclear materials in shipping containers, Risk Analy 26(5) (2006), 1377–1393.
[24] C.M. Young, M. Li, Y. Zhu, M. Xie, E.A. Elsayed, and T. Asamov, Multiobjective optimization of a port-of-entry inspection policy, IEEE Trans Auto Sci Eng 7(2) (2010), 392–400.