# Interdicting Nuclear Material on Cargo Containers using Knapsack Problem Models

Laura A. McLay and Jamie D. Lloyd
Department of Statistical Sciences & Operations Research
Virginia Commonwealth University
Richmond, Virginia
lamclay@vcu.edu

Emily Niman
Industrial & Manufacturing Engineering Department
Pennsylvania State University
State College, Pennsylvania

December 22, 2008

**Abstract**

This paper introduces a framework for screening cargo containers for nuclear material at security stations throughout the United States using knapsack problem, reliability, and Bayesian probability models. The approach investigates how to define a *system alarm* given a set of screening devices, and hence, designs and analyzes next-generation security system architectures. Containers that yield a system alarm undergo secondary screening, where more effective and intrusive screening devices are used to further examine containers for nuclear and radiological material. It is assumed that there is a budget for performing secondary screening on containers that yield a system alarm. This paper explores the relationships and tradeoffs between prescreening, secondary screening costs, and the efficacy of radiation detectors. The key contribution of this analysis is that it provides a risk-based framework for determining how to define a system alarm for screening cargo containers given limited screening resources. The analysis suggests that highly accurate prescreening is the most important factor for effective screening, particularly when screening tests are highly dependent, and that moderately accurate prescreening may not be an improvement over treating all cargo containers the same. Moreover, it indicates that screening tests with high true alarm rates may mitigate some of the risk associated with low prescreening intelligence.

## 1 Introduction

Interdicting nuclear material being smuggled into the United States on cargo containers is an issue of vital national interest, since it is a critical aspect of protecting the United States from nuclear attacks. Our economic well-being is intrinsically linked with the success and security of the international trade system, and there are enormous economic consequences when this security system is compromised. International trade accounts for more than thirty percent of the United States economy (Rooney 2005). Ninety-five percent of international goods that enter the United States come through one of 361 ports, adding up to more than 11.4M containers every year (Fritelli 2005, Rooney 2005, US DOT 2007). Port security has emerged as a critically important yet vulnerable component in the homeland security system.

Despite the importance of port security to our nation's economy, a small proportion of cargo entering United States ports is inspected for nuclear and radiological material, since it is expensive to inspect cargo by physically unpacking the containers. Instead, nearly all cargo containers are screened by less expensive radiation portal monitors (RPMs) (Lava 2008, Bakir 2009). More advanced screening and inspection technologies, such as nonintrisive inspection and unpacking containers, are used more sparingly and are targeted at high-risk containers (US CBP 2007). The Automated Targeting System (ATS) is used to prescreen each cargo container and classify it as high-risk or low-risk (Strohm 2006). Although a risk-based approach to cargo screening is part of the U.S. Customs and Border Patrol (CBP) plan for security, few guidelines are given to implement and assess such a strategy.

It is difficult to screen many cargo containers as they enter the United States. Screening cargo containers that enter the United States at land border crossings (as opposed to ports) and those that are transported by trains or barges are particularly difficult to screen, and hence, are often not screened at all (Parrish 2008, Lava 2008). Cargo containers can be screened at security stations that are not limited to the points of entry to the United States or at foreign ports, where most screening is currently performed. This paper considers such a scenario where security stations are generally located, and it focuses on the screening operations within a single station. The methodology used in this paper can be used as part of a diverse security system to intercept nuclear material at security stations located at truck weigh stations along interstates, loading docks, train stations, or at ports.

This paper introduces the Container Reliability Knapsack Problem (CRKP), a linear programming model for using existing screening technologies (e.g., RPMs) to screen cargo containers at a security station using knapsack problem, reliability, and Bayesian probability models. Screening occurs at a security station at a specific location (e.g., the exit lanes at a single port). The approach determines how to define a *system alarm* and hence, designs and analyzes security system architectures. Containers that yield a system alarm undergo secondary screening, where more effective inspection methods are used to further examine containers for nuclear and radiological material. It is assumed that there is a budget for performing secondary screening on containers that yield a system alarm. This paper explores the relationships and tradeoffs between prescreening intelligence, secondary screening costs, and the efficacy of radiation detectors. The key contribution of this analysis is that it provides a risk-based framework for determining how to define a system alarm when screening cargo containers given limited secondary screening resources. A computational example indicates that there are conditions under which there are no screening differences between high-risk and low-risk containers. The analysis suggests that limitations in screening technologies can be mitigated in part by highly effective prescreening intelligence.

This paper is organized as follows. Section 2 provides a literature review for security screening problems and research models for interdicting nuclear material. Section 3 introduces parameters and notation used in the models. CRKP is introduced in Section 4, and its structural properties are illustrated in Section 5. A computational example is analyzed in Section 6. Concluding remarks and directions for future research are given in Section 7.

## 2 Background

There is a dearth of research that applies operations research methodologies to detecting nuclear material in cargo containers. Wein et al. (2007) analyze cargo containers on truck trailers passing by a series of screening devices (RPMs) at the port of Hong Kong. They apply queuing theory and optimization to determine the optimal placement and scanning time for RPMs such that a desired detection probability is achieved. Their analysis is based on a fixed cost for the total screening budget and variable passing times for each truck. Wein et al. (2007) focus on the spatial positioning of RPMs in a security station at a port to improve detection, and they define a system alarm by optimizing the RPM alarm thresholds. However, they do not consider the effects of prescreening to identify high-risk cargo containers.

Wein et al. (2006) analyze an 11-layer screening system for containers entering the

United States by considering a fixed budget and port congestion. They determine an alternative screening design that allows the weapon placement in the truck to vary and the detection capabilities of the system to be improved relative to the current design. Furthermore, they consider the effects of prescreening from ATS and whether a terrorist enrolls in the Customs-Trade Partnership Against Terrorism (C-TPAT) program. Morton et al. (2007) and Pan (2005) use stochastic network models to detect smugglers and nuclear material based on paths traversed as part of the Second Line of Defense program. Bakir (2009) presents a decision tree model to analyze the screening of cargo containers at commercial truck crossings on the U.S. border with Mexico. They do not recommend routine screening at such commercial truck crossings. Their results largely depend on the probability of an attack, and they suggest the need for new RPMs.

Several research papers examine inspection strategies for cargo containers that use several types of screening tests. Ramirez-Marquez (2008) proposes inspection strategies of cargo containers that minimize costs of inspection at ports using decision trees. The strategies involve selecting different sensors that have varying reliability and associated costs. Using decision trees, a minimum cost inspection strategy is presented that maintains the required detection rate. Additionally, an algorithm for efficiently determining an optimal solution is included. Note that Ramirez-Marquez (2008) assume that various sensors screen containers in a selected order, whereas order is not a factor in CRKP. Moffitt et al. (2005) develops a model using information gap decision making to determine how to inspect a number of targets to shed light on robust decisions. They find that robustness to protect against a minimum level of the expected utility is not always monotonic in the number of vessels to inspect, and they recommend developing robust inspection policies that depend on inspection costs and expected utilities. Boros et al. (2006) develop a large scale linear programming model to determine how to optimally inspect cargo containers. Goldberg et al. (2008) expand this work to develop optimal inspection policies using decision trees and knapsack problem models. They present a dynamic programming algorithm that enumerates the efficient frontier of inspection policies in cost-detection space. Kantor and Boros (2007) consider mixed inspection strategies for determining when to unpack and inspect cargo containers using principles of game theory. Note that none of these efforts explicitly consider the effects of prescreening to identify high-risk cargo containers.

Several papers investigate screening paradigms for inspecting aviation passenger baggage for explosives. Although the application is somewhat different than screening cargo containers for nuclear material, in both domains a risk assessment is performed (on passengers or containers) and security devices are used to find prohibited items. In the aviation domain, many of the prohibited items are not likely to pose a threat, such as lighters and bottles of shampoo, which makes the "threat" subgroup difficult to define. In the port security domain, a threat is defined more narrowly as weapons-grade nuclear material, which is highly unlikely to be accidentally placed in a cargo container. In addition, there are more security layers in the port security domain, with more information potentially shared between these layers.

McLay et al. (2008) examine risk-based issues in detecting explosives in aviation security baggage screening models. They examine the tradeoff between intelligence and screening technology capabilities for aviation baggage security screening systems using a cost-benefit analysis when there are two types of screening technologies, one for low-risk baggage and another for high-risk baggage. The more accurate and expensive baggage screening technology is targeted at passenger baggage classified as high-risk. It is concluded that more expensive screening technologies are warranted only if effective prescreening is available. There are several key differences between the model employed by McLay et al. (2008) and CRKP. First, McLay et al. (2008) evaluate the scenario where baggage is screened by one of two available types of technology, whereas CRKP evaluates the scenario when multiple screening devices are used to screen all cargo containers for nuclear material. Moreover, CRKP assumes that the same screening devices are used to screen both high-risk and low-risk containers, whereas McLay et al. (2008) assume that there are two different types of screening technologies available and that baggage is screened by one of the two available technologies. A third difference is that CRKP assumes that screening costs are limited whereas McLay et al. (2008) assess the screening costs but do not limit them.

Kobza and Jacobson (1997) consider the design of security system architectures using reliability models in the context of aviation security baggage screening systems. Different objects (aviation bags) can take different paths through the system, and hence, are screened by varying subsets of screening devices. Their model is analyzed based on Type I (a false alarm is given) and Type II (a threat is not detected) errors, and it is formulated for a series of dependent devices. Kobza and Jacobson (1997) define a system alarm in one of two possible ways: at least one device alarm signals a system alarm, or all device alarms signal a system alarm. Their results indicate that multi-device systems can be more effective than single-device systems, taking into account the probability of errors by each sub-system. CRKP generalizes this framework by considering system alarms to be defined more generally as a $k$-of-$n$ reliability model, where a system alarm is signaled by at least $k$ of $n$ alarms. Therefore, in CRKP, a system alarm is defined by the number of devices that yield an alarm response and by classification status (i.e., high-risk or low-risk).

## 3 Screening Framework

In this section, terminology and parameters are introduced for the Container Reliability Knapsack Problem (CRKP). In CRKP, all cargo containers first undergo *prescreening* to classify each cargo container as high-risk or low-risk. Cargo containers enter a security station to undergo *primary screening*. It is assumed that each container is on a truck trailer, although CRKP can be interpreted more generally to screen any types of objects using dependent screening devices. When a cargo container enters a security station, it is screened by several *sensors*. These sensors could be radiation detectors such as RPMs, which screen each cargo container for radiation that is emitted by nuclear material such as plutonium and highly enriched uranium (HEU). Sensors can be interpreted more generally to be any type of screening device or procedure that yields a binary response. Each sensor yields an alarm or clear response, based on how the sensor operates and the characteristics of the cargo container. Each truck trailer drives a cargo container through the security station by all the sensors sequentially, and after each cargo container is screened, the total number of sensor alarms is known, and based on this total number of sensor alarms, a system response is given, which allows the system response to be defined in one of several ways (Kobza and Jacobson 1996, 1997). The system response has one of two outcomes, either a system alarm is given or the container is cleared. If the cargo container is cleared, it exits the security station and continues along its path to its destination. The cargo containers that yield a system alarm undergo *secondary screening*. All cargo containers undergo primary screening and CRKP is used to determine the subset of these cargo containers that undergo secondary screening. Note that it has been observed that the costs associated with secondary screening dominate the costs associated with primary screening (Wein et al. 2007, Bakir 2009).

The parameters for CRKP are classified into two groups: (1) cost and screening parameters and (2) probability parameters.

(1) Cost and screening parameters

- $n$ = number of sensors in the security station,
- $N$ = number of cargo containers to be screened in the security station,
- $B$ = total secondary screening budget,
- $C_{SS}$ = cost to perform secondary screening on a container,
- $\beta$ = ratio of high-risk containers that are threats to low-risk containers that are threats (i.e., $\beta = P_{T|HR}/P_{T|LR}$).

The number of sensors is a deterministic value based on security station operations. The total number of containers is a deterministic value that represents the number of cargo containers that pass through a given station in a year, or another period of time. The budget for secondary screening is a deterministic value based on available resources, and it can be defined to implicitly capture the costs of delays and congestion. Note that this budget reflects only the direct cost of secondary screening, and hence, the additional costs that are incurred by false clears are not assessed against the budget. The cost to perform secondary screening is a deterministic value based on information collected and analyzed by Department of Homeland Security (DHS) and CBP. It is in part based on salaries paid to the employees hired to perform secondary screening. It is assumed that the cost to perform secondary screening is the same for threats and nonthreats.

(2) Probability parameters

- $P_{HR}$ ($P_{LR}$) = the probability that a cargo container is classified as high-risk (low-risk),
- $P_T$ ($P_{NT}$) = the probability that a cargo container is a threat, i.e., contains nuclear material (non-threat),

- $P_{kA}$ = the probability that a cargo container yields $k$ alarms (of the $n$ sensors), $k = 0, 1, ..., n$,
- $P^i_{A|T}$ =1 $-$ $P^i_{NA|T}$ = the probability that a threat container yields an alarm at sensor $i$, $i = 1, 2, ..., n$,
- $P^i_{A|NT}$ =1 $-$ $P^i_{NA|NT}$ = the probability that a non-threat container yields an alarm at sensor $i$, $i = 1, 2, ..., n$.

The screening process yields one of four possible outcomes: a true alarm, false clear, false alarm, or true clear. The probability of these outcomes occurring depends on how the sensors are operated, as well as the size, type, location, and shielding of the source for threat containers. If each sensor operates identically, then the single sensor alarm probabilities for all sensors are $P_{A|T}$ and $P_{A|NT}$ for threat and non-threat containers, respectively. Each cargo container is classified as high-risk or low-risk, based on a pre-screening system such as ATS. The characteristics that determine whether a container is high-risk or low-risk are classified. The probability that a container is classified as high-risk is based on the proportion of containers passing through a security station that are classified as high-risk, once a large number of cargo containers has been evaluated. The probability that a cargo container is a threat is assessed by personnel within the DHS based on the perceived threat level. This value is considered highly sensitive and may change based on changes in national or international situations, intelligence information, or the risk level of the Homeland Security Advisory System. The probability that a cargo container yields $k$ (of $n$) alarms depends on how the sensors operate, and it is assumed to only depend on whether a container is a threat or non-threat.

## 4 The Container Reliability Knapsack Problem

The objective of CRKP is to determine which high-risk and low-risk containers yield a system alarm in order to maximize the expected number of threats selected for secondary screening, given the number of alarms from the $n$ sensors and subject to a screening budget. Although selecting a threat container for secondary screening does not guarantee that it is detected, procedures for secondary screening, such as unpacking a cargo container and using non-intrusive inspection technologies, have a high probability of detecting nuclear material. Cargo containers that are not selected for secondary screening are cleared, and hence, there is no chance of interdicting the nuclear material if it is indeed present.

It is assumed that each container is screened independently of the other containers. Rather, sensor operation depends only on whether a container is a threat. After each container is screened by the sensors, the number of alarms is known, and a decision is made about whether secondary screening is used to screen each container. It is also assumed that the sensors work the same regardless of whether a container is classified as high-risk and low-risk.

Decision Variables:

- $x^k_{HR} = P_{SS|kA\cap HR}$ = fraction of high-risk containers that yield exactly $k$-of-$n$ alarms selected to undergo secondary screening,
- $x^k_{LR} = P_{SS|kA\cap LR}$ = fraction of low-risk containers that yield exactly $k$-of-$n$ alarms selected to undergo secondary screening.

The objective function value of CRKP is stated as follows.

$$\begin{aligned}\max \quad & E[\text{Number of threats selected for secondary screening}] \\ & = NP\{\text{A threat container is selected for secondary screening}\} \\ & = N\sum_{k=0}^{n}\left(P_{SS\cap T\cap kA\cap HR} + P_{SS\cap T\cap kA\cap LR}\right) \\ & = N\sum_{k=0}^{n}\left(P_{kA|T\cap HR}P_{HR|T}P_{T}P_{SS|T\cap kA\cap HR} + P_{kA|T\cap HR}P_{LR|T}P_{T}P_{SS|T\cap kA\cap LR}\right).\end{aligned}$$

By assumption, $P_{SS|T\cap kA\cap HR} = x^k_{HR}$ and $P_{SS|T\cap kA\cap LR} = x^k_{LR}$, $k = 0, 1, ..., n$, and $P_{kA|T\cap HR} = P_{kA|T\cap LR} = P_{kA|T}$, and therefore, the objective function is rewritten as

$$\max \qquad NP_T\sum_{k=0}^{n}P_{kA|T}\left(P_{HR|T}x^k_{HR} + P_{LR|T}x^k_{LR}\right). \tag{1}$$

Bayes rule can be used to compute $P_{HR|T}$ and $P_{\mathrm{HR|NT}}$,

$$P_{\mathrm{HR|T}} = \frac{\beta P_{\mathrm{HR}}}{1 - P_{\mathrm{HR}} + \beta P_{\mathrm{HR}}}, \tag{2}$$

$$P_{\mathrm{HR|NT}} = \frac{P_{\mathrm{HR}} - P_{\mathrm{HR|T}}P_{\mathrm{T}}}{1 - P_{\mathrm{T}}}. \tag{3}$$

There is a single budget constraint in CRKP that ensures that the expected number of containers that undergo secondary screening is less than $B/C_{SS}$. Using the same reasoning as to construct the objective function value (1), the budget constraint is

$$N \sum_{k=0}^{n} \Big( \big( P_{kA|T} P_T P_{HR|T} + P_{kA|NT} P_{NT} P_{HR|NT} \big) x_{HR}^k + \big( P_{kA|T} P_T P_{LR|T} + P_{kA|NT} P_{NT} P_{LR|NT} \big) x_{LR}^k \Big) \leq B/C_{SS}. \tag{4}$$

CRKP is formulated as a linear programming model, using the objective function (1), budget constraint (4), and two sets of constraints that set the decision variable upper and lower bounds, (5) and (6).

$$0 \leq x_{HR}^k \leq 1, k = 0, 1, ..., n \tag{5}$$

$$0 \leq x_{LR}^k \leq 1, k = 0, 1, ..., n. \tag{6}$$

In the case when each sensor operates independently and identically with the probability of a single sensor true alarm $P_{A|T}$ and the probability of a single sensor false alarm $P_{A|NT}$, then $P_{kA|T} = C_k^n P_{A|T}^k (1 - P_{A|T})^{n-k}$ and $P_{kA|NT} = C_k^n P_{A|NT}^k (1 - P_{A|NT})^{n-k}$ using the Binomial distribution with parameters $n$ and $P_{A|T}$ for threat containers, parameters $n$ and $P_{A|NT}$ for non-threat containers, and $C_k^n$ is the Binomial coefficient.

Note that computing the conditional probability that there are $k$ alarms given that a container is a threat or non-threat is not trivial if there is dependence between the sensors, which is likely to hold in practice. Finding these conditional probabilities can be accomplished by computing the reliability of a $k$-out-of-$n$ reliability system, in which the system yields an alarm response if at least $k$ sensors yield an alarm (Koucky 2003). Define the following parameters. The state of the system is defined by the vector $\mathbf{Y} = (Y_1, Y_2, ..., Y_n)$, where $Y_i = 1$ if sensor $i$ yields an alarm and 0 otherwise.

- $R_{(k,n)}$ = reliability of the $k$-out-of-$n$ system,
- $U(s) = P\{\prod_{i \in s} Y_i = 0\}$ = joint probability of the components of a subset $s$ of the $n$ sensors ($s \subseteq \{1, 2, ..., n\}$),
- $S(r)$ = family of cardinality $r$-subsets of $\{1, 2, ..., n\}$,
- $C_k^n$ = Binomial coefficient.

The reliability of the system is given by

$$R_{(k,n)} = 1 - \sum_{i=0}^{k-1} (-1)^i C_i^{n-k+i} U_{(n-k+i+1)},$$

where $U_{(r)} = \sum_{s \in S(r)} U(s)$ (Koucky 2003). Then, the probability that there are exactly $k$ alarms is given by $P_{kA} = R_{(k,n)} - R_{(k+1,n)}, k = 0, 1, ..., n-1$, with $P_{nA} = 1 - \sum_{k=0}^{n-1} P_{kA}$. Note that the number of alarms can be computed separately for threat and nonthreat containers, yielding $P_{kA|T}$ and $P_{kA|NT}$.

# 5 Structural Properties

This section summarizes the structural properties of CRKP. CRKP is identical to the linear programming relaxation to the 0-1 Knapsack Problem (KP). In KP, there are $m$ items with a reward $r_i$ and weight $w_i$, $i = 1, 2, ..., m$, and a knapsack capacity $c$. The linear programming relaxation to KP can be solved in O($m$) time. The items are sorted in decreasing order of the ratio of the item reward to weight (i.e., $r_1/w_1 \geq r_2/w_2 \geq ... \geq r_m/w_m$), which is defined as the *optimal knapsack sequence*. Starting with the first item, items are greedily inserted into the knapsack in order until there is no remaining capacity in the knapsack. Therefore, the variables are all one or zero for all items except the critical item $s$ (where $s = \arg\min_j \{\sum_{i=1}^{j} w_i > c\}$). KP corresponds to CRKP with $m = 2(n+1)$, capacity $c = B/C_{SS}$, and rewards equal to the expected number of high-risk and low-risk threat containers that yield $k$ alarms, and weight equal to the expected number of high-risk and low-risk containers that yield $k$ alarms, $k = 0, 1, ..., n$.

In CRKP, define the rewards for high-risk and low-risk containers as

$$r_{HR}^k = NP_T P_{kA|T} P_{HR|T}, k = 0, 1, ..., n,$$

$$r_{LR}^k = NP_T P_{kA|T} P_{LR|T}, k = 0, 1, ..., n,$$

respectively. Likewise, define the weights for high-risk and low-risk containers as

$$w_{HR}^k = N\left(P_{kA|T} P_T P_{HR|T} + P_{kA|NT} P_{NT} P_{HR|NT}\right), k = 0, 1, ..., n,$$

$$w_{LR}^k = N\left(P_{kA|T} P_T P_{LR|T} + P_{kA|NT} P_{NT} P_{LR|NT}\right), k = 0, 1, ..., n,$$

respectively. Therefore, CRKP can be rewritten as

$$\begin{aligned} \max \quad & \sum_{k=0}^{n} (r_{HR}^k x_{HR}^k + r_{LR}^k x_{LR}^k) \\ \text{subject to} \quad & \sum_{k=0}^{n} (w_{HR}^k x_{HR}^k + w_{LR}^k x_{LR}^k) \leq B/C_{SS} \\ & 0 \leq x_{HR}^k \leq 1, k = 0, 1, ..., n \\ & 0 \leq x_{LR}^k \leq 1, k = 0, 1, ..., n. \end{aligned}$$

The screening scenario captured by CRKP can be viewed in terms of Bayesian probabilities, with the prescreening risk classification defining the prior probabilities, and the number of alarms defining the posterior probabilities. The prior probabilities that high-risk and low-risk cargo containers are a threat are computed using (2) and (3) and are

$$P_{T|HR} = \frac{\beta P_T}{1 - P_{HR} + \beta P_{HR}},$$
$$P_{T|LR} = P_{T|HR}/\beta = \frac{P_T}{1 - P_{HR} + \beta P_{HR}},$$

respectively.

The posterior probabilities are the conditional probabilities that a cargo container is a threat given that it is classified as high-risk (low-risk) and yields $k$ alarms. Theorem 1 defines the posterior probabilities.

**Theorem 1** *The posterior probabilities $P_{T|kA\cap HR}$ and $P_{T|kA\cap LR}$ are defined as the ratio of the CRKP reward to the weight, $r_{HR}^k/w_{HR}^k$ and $r_{LR}^k/w_{LR}^k$, $k = 0, 1, ..., n$, respectively.*

Proof. First consider high-risk cargo containers. The posterior probability that a high-risk cargo container yielding $k$ alarms is a threat is

$$P_{T|kA\cap HR} = \frac{P_{T\cap kA\cap HR}}{P_{kA\cap HR}} = \frac{P_{T\cap kA\cap HR}}{P_{T\cap kA\cap HR} + P_{NT\cap kA\cap HR}} = \frac{r_{HR}^k/N}{w_{HR}^k/N}$$

The posterior probabilities for low-risk cargo containers are computed in a similar manner. □

For practical reasons, it is desirable for CRKP to identify containers for secondary screening that yield more alarms rather than fewer alarms, resulting in a threshold policy. Theorem 2 indicates the conditions under which a high-risk (low-risk) container yielding more alarms makes it more likely to be selected for secondary screening. Note that among only high-risk (low-risk) containers, the order that items are put into the knapsack (i.e., the order in which containers are selected for secondary screening) depends only on how the sensors work together and not on prescreening intelligence, the underlying probability of a threat, or the proportion of containers classified as high-risk.

**Theorem 2** *High-risk (low-risk) containers that yield $k$ alarms occur before high-risk (low-risk) containers that yield $k-1$ alarms in the optimal knapsack sequence,*

$$\frac{r_{HR}^k}{w_{HR}^k} \geq \frac{r_{HR}^{k-1}}{w_{HR}^{k-1}} \left( \frac{r_{LR}^k}{w_{LR}^k} \geq \frac{r_{LR}^{k-1}}{w_{LR}^{k-1}} \right)$$

*only if*

$$\frac{P_{kA|T}}{P_{(k-1)A|T}} \geq \frac{P_{kA|NT}}{P_{(k-1)A|NT}}.$$

Proof. First, consider the high-risk containers. By definition,

$$\frac{P_{kA|T}P_T P_{HR|T}}{P_{kA|T}P_T P_{HR|T} + P_{kA|NT}P_{NT}P_{HR|NT}} \geq \frac{P_{(k-1)A|T}P_T P_{HR|T}}{P_{(k-1)A|T}P_T P_{HR|T} + P_{(k-1)A|NT}P_{NT}P_{HR|NT}}.$$

Rearranging yields the desired result. The same approach can be taken for the low-risk containers. □

Corollary 1 illustrates when the conditions in Theorem 2 hold for the particular case when each sensor operates independently and identically. It indicates that the single sensor true alarm probability must be higher than the single sensor false alarm probability in order for a cargo container yielding $k$ alarms to occur earlier in the optimal knapsack sequence before a cargo container yielding $k-1$ alarms, $k = 1, 2, ..., n$.

**Corollary 1** *When sensors alarms are independently and identically distributed with the probability of a true alarm $P_{A|T}$ and the probability of a false alarm $P_{A|NT}$, then*

$$\frac{r_{HR}^k}{w_{HR}^k} \geq \frac{r_{HR}^{k-1}}{w_{HR}^{k-1}} \text{ and } \frac{r_{LR}^k}{w_{LR}^k} \geq \frac{r_{LR}^{k-1}}{w_{LR}^{k-1}}$$

*only if $P_{A|T} \geq P_{A|NT}$.*

Proof. First, consider the high-risk containers. Using Theorem 2, then

$$\frac{P_{kA|T}}{P_{(k-1)A|T}} = \frac{C_k^n P_{A|T}^k (1 - P_{A|T}^{n-k})}{C_{k-1}^n P_{A|T}^{k-1} (1 - P_{A|T}^{n-k+1})} \geq \frac{C_k^n P_{A|NT}^k (1 - P_{A|NT}^{n-k})}{C_{k-1}^n P_{A|NT}^{k-1} (1 - P_{A|NT}^{n-k+1})} = \frac{P_{kA|NT}}{P_{(k-1)A|NT}}.$$

Rearranging yields

$$\frac{P_{A|T}}{1 - P_{A|T}} \geq \frac{P_{A|NT}}{1 - P_{A|NT}},$$

and simplifying yields $P_{A|T} \geq P_{A|NT}$. The same approach can be taken for the low-risk containers. □

Lemma 1 and Corollaries 2 and 3 quantify the relationships between the CRKP rewards.

**Lemma 1** *The objective coefficients $r_{HR}^k > r_{HR}^{k-1}$ and $r_{LR}^k > r_{LR}^{k-1}$ only if $P_{kA|T} > P_{(k-1)A|T}$.*

Proof. Follows from the objective function. □

**Corollary 2** *When sensors alarms operate independently and identically with the probability of a true alarm* $P_{A|T}$*, then* $P_{kA|T} > P_{(k-1)A|T}$ *only if* $P_{A|T} > \frac{k}{n+1}$*,* $k = 1, 2, ..., n$*.*

Proof. The number of alarms can be modeled as a Binomial random variable with parameters $n$ and $P_{A|T}$. Then

$$P_{kA|T} = \frac{n!}{k!(n-k)!} P_{A|T}^{k}(1-P_{A|T})^{n-k} > \frac{n!}{(k-1)!(n-k+1)!} P_{A|T}^{k-1}(1-P_{A|T})^{n-k+1} = P_{(k-1)A|T}$$

and rearranging yields

$$P_{A|T} > \frac{k}{n+1}. \ \square$$

**Corollary 3** *When sensors alarms operate independently and identically with the probability of a true alarm* $P_{A|T}$*, then* $P_{kA|T} > P_{(k-1)A|T}$ *for all* $k = 1, 2, ..., n$ *only if* $P_{A|T} > \frac{n}{n+1}$*.*

Proof. Follows from Corollary 2. □

Lemma 2 indicates that the ratio of the rewards for high-risk to low-risk containers is a constant factor for each $k$, $k = 0, 1, ..., n$, that depends only on prescreening intelligence and the proportion of containers that are classified as high-risk.

**Lemma 2** *The ratio of rewards for high-risk to low-risk containers is*

$$\frac{r_{HR}^{k}}{r_{LR}^{k}} = \frac{P_{HR|T}}{P_{LR|T}} = \frac{\beta P_{HR}}{1 - P_{HR}},$$

*for* $k = 0, 1, ..., n$*.*

Proof. Follows from the definition of the rewards, since $\frac{r_{HR}^{k}}{r_{LR}^{k}} = \frac{P_{HR|T}}{P_{LR|T}}$. □

## 6 Computational Example and Results

This section reports results for a computational example to assess the theoretical properties of CRKP and to understand the tradeoffs between prescreening intelligence (i.e., $\beta$), secondary screening costs, and the false alarm and false clear rates associated with each sensor. The analysis considers two cases. The first case (Case 1) considers cargo containers on truck trailers that drive by a series of $n$ sensors that are independent and operate identically. Therefore, the number of alarms for threat containers are modeled using a Binomial distribution with parameters $n$ and $P_{A|T}$, and the number of alarms for threat containers are modeled using a Binomial distribution with parameters $n$ and $P_{A|NT}$. The second case (Case 2) considers a series of $n$ sensors that have a degree of dependence between the sensors.

CRKP is analyzed for a single security station over a time horizon of one year. Table 1 contains the base case input parameters for CRKP, which remain constant unless otherwise specified. It is assumed that $N = 100{,}000$ containers enter the security station during the time horizon. The probability that a container is a threat is $1/N$, which was selected such that one threat is expected to pass through the security station. Five percent of all containers are assumed to be high-risk, which is consistent with what is reported in the public domain (Robinson et al. 2005, Strohm 2006, Lava 2008, The Royal Society 2008). The cost of secondary screening is set to $C_{SS} = \$50$ per container, although this value appears to be highly variable (Wein et al. 2007, Bakir 2009).

In the analysis, the objective function represents the expected number of true alarms. The expected number of threats in the system is $P_T N = 1$, and hence, 1.0 is an upper bound on the objective function value for the base case. Define the *detection probability* as the conditional probability that a threat is selected for secondary screening. The detection probability is computed as the objective function value divided by $P_T N$. For all scenarios considered, CRKP is solved to determine the minimum cost (i.e., level of the budget) needed to ensure a detection probability of 0.95 (Wein et al. 2007, Levi 2007).

The minimum cost to achieve a detection probability of 0.95 depends on the costs associated with secondary screening as well as the total number of cargo containers passing through the security station. The cost of secondary screening depends on many factors, such as labor, delay, and offsite testing costs. To avoid these factors, the proportion of containers selected for secondary screening is reported as a proxy for cost. Therefore, the total cost to achieve a detection probability of 0.95 is rescaled by $C_{SS}N$ to reflect the proportion of containers that are selected for secondary screening. Let $Q_{SS}[\mathrm{DP}]$ denote the proportion of cargo containers selected for secondary screening in order to achieve a specified detection probability (DP), with DP$= 0.95$ for the cases considered.

Note that CRKP is an instance of the linear programming relaxation to KP, and hence, there is at most one fractional variable in an optimal solution. A fractional variable is interpreted to represent the fraction of containers yielding the particular number of alarms that is randomly selected for secondary screening. For example, $x^4_{LR} = 0.2$ is interpreted to mean that a low-risk container yielding four alarms has a probability of 0.2 of being selected for secondary screening and a probability of 0.8 of being cleared. All other variables are zero (meaning that no containers are selected for secondary screening) or one (meaning that all containers are selected for secondary screening).

The value of the prescreening multiplier $\beta$ determines the probability that a threat container is classified as high-risk for a given proportion of containers classified as high-risk $P_{HR}$. It is difficult to estimate $\beta$, since there are no known attempts of terrorists smuggling weapons-grade nuclear material into the United States (IAEA 2007). However, it has been reported that ATS may not be effective in identifying threats (GAO 2006). In the aviation domain, McLay et al. (2008) report that for $P_{HR} = 0.05$, $\beta = 10$ (i.e., $P_{HR|T} = 0.34$) is realistic and that $\beta = 100$ (i.e., $P_{HR|T} = 0.84$) is an upper bound for an improved prescreening system. Since $\beta$ is a function of $P_{HR}$, it is difficult to compare scenarios with a given $\beta$ across different values of $P_{HR}$. As $P_{HR}$ increases for a fixed value of $\beta > 1$, the ratio of the number of threat containers classified as high-risk to the number of threat containers classified as low-risk is constant. However, $P_{HR|T}$ increases as a result of more containers being classified as high-risk, not as a result of an improvement in prescreening intelligence. To avoid this problem, scenarios with a fixed value of $P_{HR} = 0.05$ are compared across different values of $\beta$. In this case, $P_{HR|T}$ increases with $\beta$ as a result of improvements in prescreening intelligence. Note that $\beta = 1$ corresponds to the *random prescreening* case. When $\beta = 1$, $P_{T|HR} = P_{T|LR}$, so screening is random and independent of the risk classification, where a proportion $P_{HR}$ of cargo containers are randomly classified as high-risk.

The parameters $P_{A|T}$ and $P_{A|NT}$ represent the probability of a single sensor true alarm and false alarm, respectively. The base case true alarm and false alarm values used in the analysis are set to 0.7 and 0.05, respectively. The false alarm probability is set to 0.05 to be consistent with high false alarm rates experienced at ports (Slaughter et al. 2003, Rooney 2005, The Royal Society 2008). Publicly reported estimates for the true alarm probability have widely varied, and hence, the true alarm probability is set to 0.7 (Levi 2007, Cochran and McKinzie 2008). Note that Corollary 3 indicates that the

Table 1: Base case parameter values

| Parameter | Value(s) |
|---|---|
| $N$ | 100,000 |
| $P_T$ | $1/N = 0.00001$ |
| $n$ | 1, 3, 5 |
| $P_{HR}$ | 0.05 |
| $C_{SS}$ | \$50 |
| $\beta$ | 1, 10, 100 |
| $P_{A\|T}$ | 0.7 |
| $P_{A\|NT}$ | 0.05 |

true alarm probability should be greater than the false alarm probability when sensors operate identically and independently, and this is consistent with the parameters used in this analysis.

### 6.1 Case 1: Identical and Independent Sensors

In order to assess CRKP, the values of $P_{A|T}$, $P_{A|NT}$, and $\beta$ are varied for the case when the sensors operate independently and identically. Although there is dependency between sensors currently used to screen containers for nuclear material, Case 1 assumes independence to shed light on how to optimally screen cargo containers using multiple sensors under ideal conditions using next-generation screening technologies with fewer dependencies. In this case, assume that each sensor is a unique type of screening technology.

Each cargo container is screened by $n = 1, 3, 5$ sensors. Each sensor operates independently and identically, with the single sensor true and false alarm probabilities being 0.7 and 0.05, respectively. Note that under these conditions, the conditions under Theorem 2 and Corollary 1 are satisfied, and hence, containers that yield more alarms are selected for secondary screening before containers that yield fewer alarms, which indicates that the optimal screening policy is a threshold policy. Table 2 shows the proportion of cargo containers selected for secondary screening in order to achieve a detection probability of 0.95 as well as the corresponding cost per cargo container.

Figure 1 shows $Q_{SS}[0.95]$ as $\beta$ varies from 1 to 100 for $P_{A|T} = 0.7$ (the base case) and $P_{A|T} = 0.1$ in order to illustrate the effect of having inaccurate sensors, since it has been reported that RPMs do not consistently identify nuclear material (Levi 2007, Cochran and McKinzie 2008). Figure 1 suggests that having accurate sensors is more important for keeping secondary screening costs to a minimum than highly accurate prescreening

Table 2: Base case costs and minimum proportion of cargo containers selected for secondary screening for a detection probability of 0.95

| $n$ | $\beta$ | $Q_{SS}[0.95]$ | Cost per Container ($) |
|---|---|---|---|
| 1 | 1 | 0.842 | 42.08 |
| | 10 | 0.770 | 38.52 |
| | 100 | 0.097 | 4.83 |
| 2 | 1 | 0.499 | 24.93 |
| | 10 | 0.273 | 13.65 |
| | 100 | 0.095 | 4.73 |
| 3 | 1 | 0.126 | 6.31 |
| | 10 | 0.119 | 5.94 |
| | 100 | 0.028 | 1.38 |
| 4 | 1 | 0.090 | 4.52 |
| | 10 | 0.048 | 2.38 |
| | 100 | 0.014 | 0.68 |
| 5 | 1 | 0.019 | 0.97 |
| | 10 | 0.018 | 0.90 |
| | 100 | 0.003 | 0.15 |

intelligence or having many sensors. A single sensor with $P_{A|T} = 0.7$ using random prescreening (i.e., $\beta = 1$) has lower secondary screening costs compared to three sets of scenarios with $P_{A|T} = 0.1$ ($n = 1$ and $\beta \leq 41$, $n = 3$ and $\beta \leq 35$, $n = 5$ and $\beta \leq 29$), which suggests that sensor inaccuracies can be offset by highly accurate prescreening intelligence.

To better understand secondary screening costs, sensitivity analysis was performed for $P_{A|T}$ and $P_{A|NT}$. Figure 2 shows $Q_{SS}[0.95]$ as a function of the probability of a single sensor true alarm. Figure 2(a) illustrates the case with $n = 1$, Figure 2(b) illustrates the case with $n = 3$, and Figure 2(c) illustrates the case with $n = 5$. As the probability of a single sensor true alarm approaches 1.0, the proportion of containers that require secondary screening to maintain a detection probability of 0.95 decreases drastically, which suggests that highly effective sensors can counteract less accurate prescreening intelligence. However, for more moderate values of $P_{A|T}$, prescreening intelligence is necessary to reduce secondary screening costs, particularly when there is a single sensor.

Figure 3 shows the system alarm threshold as a function of $P_{A|T}$ for the case with $n = 5$ sensors, with Figure 3(a) showing the sensor alarms for the $\beta = 10$ case and Figure 3(b) showing the sensor alarms for the $\beta = 100$ case. In Figure 3, if the number of observed alarms is greater than the system alarm threshold, then the cargo container is selected for secondary screening. If the number of observed alarms is less than or equal to than the system alarm threshold, then the cargo container is cleared. Note

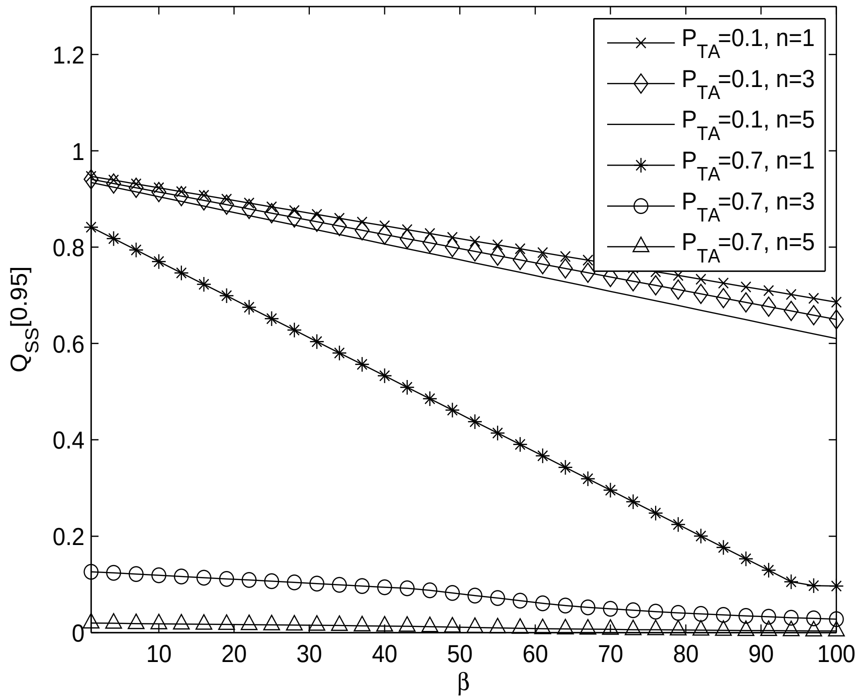


Figure 1: Minimum proportion of cargo containers selected for secondary screening for a detection probability of 0.95 as a function of $\beta$

that in all instances, high-risk containers require fewer sensor alarms to be selected for secondary screening, and as prescreening intelligence improves, the difference between low-risk and high-risk containers is accentuated. As the probability of a single sensor true alarm increases, the screening process more accurately identifies threat containers, and hence, containers with fewer sensor alarms are less likely to be selected for secondary screening. The system alarm threshold is identical between high-risk and low-risk containers across all random prescreening scenarios with $\beta = 1$, meaning that there is no practical screening differences for randomly classified high-risk and low-risk cargo containers.

Note that in several cases, there is no difference in $Q_{SS}[0.95]$ between the $\beta = 1$ and $\beta = 10$ case (such as $n = 3$ and $P_{A|T} = 0.63$ in Figure 2(b), $n = 5$ and $P_{A|T} = 0.45$ in Figure 2(c)), which indicates that efforts made to moderately improve prescreening intelligence over random may not have any impact on security. This observation is counter-intuitive. In the aviation security domain, the opposite conclusion has been

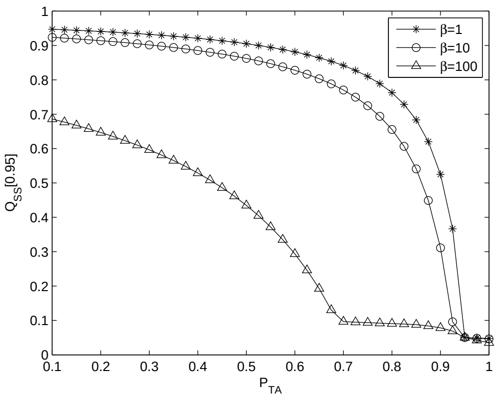


(a) $n = 1$

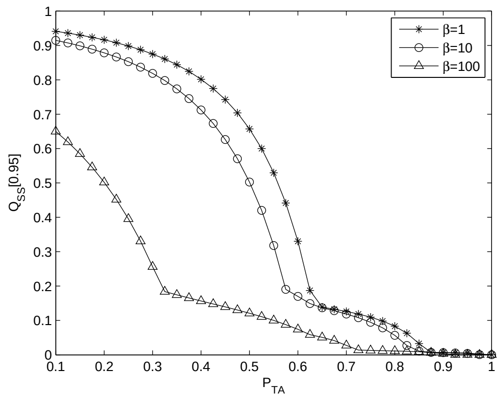


(b) $n = 3$

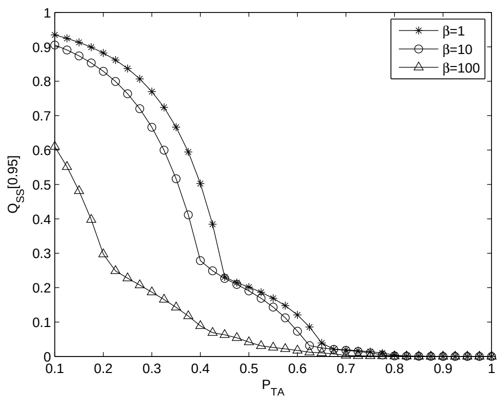


(c) $n = 5$

Figure 2: Proportion of cargo containers selected for secondary screening for a detection probability of 0.95 as a function of the probability of a single sensor true alarm

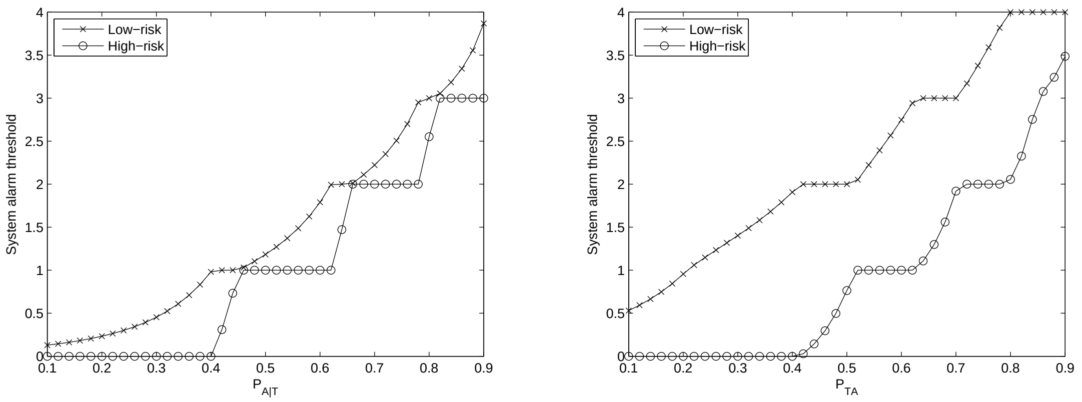


(a) $\beta = 10$

(b) $\beta = 100$

Figure 3: System alarms as a function of the probability of a single sensor true alarm for $n = 5$ scenarios

drawn, namely that moderate increases in prescreening intelligence have large effects on security (McLay et al. 2008). In CRKP, this occurs when CRKP defines an identical system alarm threshold for both high-risk and low-risk cargo containers (see Figure 3(a)), which suggests that how sensors operate should also be considered when designing screening systems.

Figure 4 shows $Q_{SS}[0.95]$ as a function of $P_{A|NT}$. As the probability of a single sensor false alarm increases, so does the proportion of containers that are selected for secondary screening. When $\beta = 1$, and $n = 5$, $Q_{SS}[0.95]$ is lower than the corresponding scenarios with $\beta = 10$ and $n = 1, 3, 5$ for all values of $P_{A|NT}$, $\beta = 100$ and $n = 1$ for all values of $P_{A|NT}$, and $\beta = 100$ and $n = 3$ for $P_{A|NT} \leq 0.06$. This suggests that multiple, independent screening technologies can counteract poor prescreening intelligence. Improving $\beta$ from 10 to 100 significantly reduces the proportion of containers that are selected for secondary screening, which suggests that highly effective prescreening counteracts sensors with high false alarm rates. Note that the proportion of containers requiring secondary screening when $P_{A|NT} = 0.1$ and $n = 3, 5$ is smaller than the proportion of containers requiring secondary screening when $P_{A|NT} = 0.01$ and $n = 1$ for corresponding values of $\beta = 1, 10$, which suggests that using multiple sensors with high false alarm rates may be more effective than using a single sensor with a low false alarm rate.

Figure 5 shows the system alarm threshold defined for the case with $n = 5$ sensors,

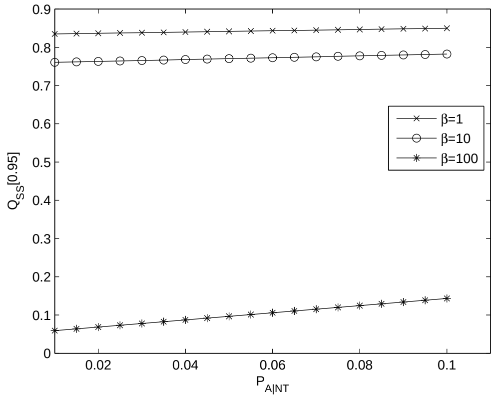


(a) $n = 1$

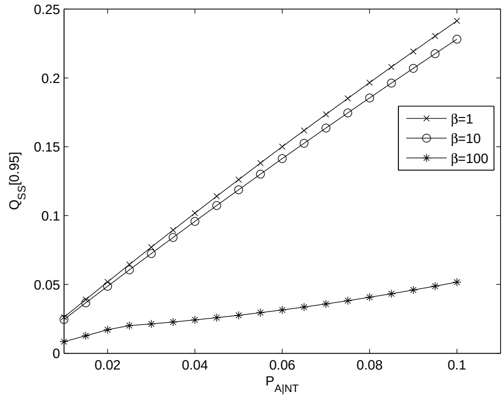


(b) $n = 3$

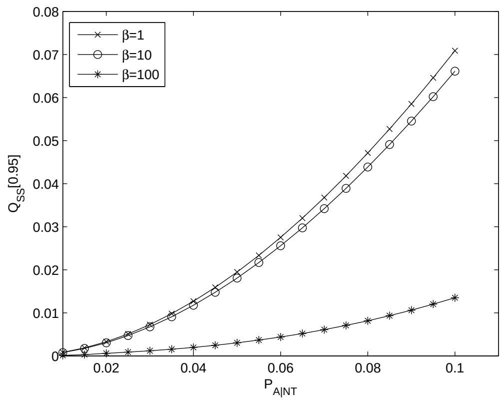


(c) $n = 5$

Figure 4: Proportion of cargo containers selected for secondary screening for a detection probability of 0.95 as a function of the probability of a single sensor false alarm

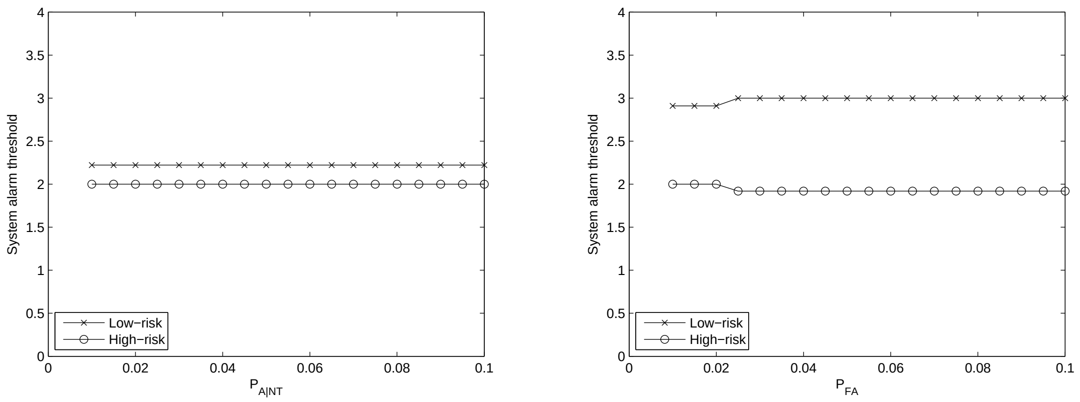


(a) $\beta = 10$ (b) $\beta = 100$

Figure 5: System alarms as a function of the probability of a single sensor false alarm for $n = 5$ scenarios

with Figure 5(a) showing the system alarm threshold for the $\beta = 10$ case and Figure 5(b) showing the sensor alarm threshold for the $\beta = 100$ case. Note that in all instances, high-risk containers require fewer sensor alarms to be selected for secondary screening, and as prescreening accuracy improves, the difference between low-risk and high-risk containers is accentuated, largely by raising the system alarm threshold for low-risk containers.

### 6.2 Case 2: Dependent Sensors

In practice, there is likely to be a high level of dependence between sensors for detecting nuclear material (Fetter et al. 1990, Levi 2007). With highly dependent sensors, a sensor is extremely likely to yield an alarm (clear) response if other sensors yield alarm (clear) responses. In order to determine system performance when there are multiple, dependent sensors, the following criteria are used to specify the number of alarms. All sensors are assumed to work identically but not independently. The probability of observing an alarm at the second and subsequent sensors is assumed to be conditional on the response of the first sensor for threat and non-threat containers. Given the response of the first sensor, the remaining $n-1$ sensors are assumed to operate independently and identically, given the level of dependence. Let $D$ define the level of dependence between the first sensor and the remaining $n-1$ sensors, $0 \leq D \leq 1$. Define $Aj$ $(NAj)$ as the event that the $j$th sensor yields an alarm (clear) response. The true alarm and false alarm probabilities for the first sensor ($j = 1$) are $P_{A|T}$ and $P_{A|NT}$, respectively. If the first sensor yields an alarm response, then the true alarm and false alarm probabilities for the remaining $n - 1$ sensors are defined as

$$\begin{aligned} P_{Aj|A1\cap T} &= P_{A|T} + D(1 - P_{A|T}),\ j = 2, 3, ..., n, \\ P_{Aj|A1\cap NT} &= P_{A|NT} + D(1 - P_{A|NT}),\ j = 2, 3, ..., n, \end{aligned}$$

respectively. In other words, given an alarm by the first sensor for a threat (non-threat) container, the probability that subsequent sensors yield an alarm response is linearly scaled between $P_{A|T}$ ($P_{A|NT}$) and one by $D$. If the first sensor yields a clear response, then the true alarm and false alarm probabilities for the remaining $n - 1$ sensors are defined as

$$\begin{aligned} P_{Aj|NA1\cap T} &= (1 - D)P_{A|T},\ j = 2, 3, ..., n, \\ P_{Aj|NA1\cap NT} &= (1 - D)P_{A|NT},\ j = 2, 3, ..., n, \end{aligned}$$

respectively. In other words, given a clear response by the first sensor for a threat (non-threat) container, the probability that subsequent sensors yield an alarm response is linearly scaled between zero and $P_{A|T}$ ($P_{A|NT}$) by $D$.

Figure 6 shows the proportion of cargo containers selected for secondary screening to achieve a detection probability of 0.95 with $n = 5$ sensors. Note that when $D = 0$, the proportion of cargo containers selected for secondary screening is identical to the case when sensors operate identically and independently. When $D = 1$, the $n = 5$ sensors yield identical outcomes, and hence, the proportion of cargo containers selected for secondary screening is identical to the case with one sensor. Therefore, when there is a high level of dependence between sensors, using additional sensors for screening cargo containers offers few benefits as compared to using a single sensor. Figure 6 suggests that the problems with highly dependent sensors can be mitigated in part when $\beta$ is large. Note that when $D = 0.48$ and $D \leq 0.12$, there is no practical difference between the $\beta = 1, 10$ scenarios. This suggests that low and moderate levels of $\beta$ may not be adequate to improve the detection probability.

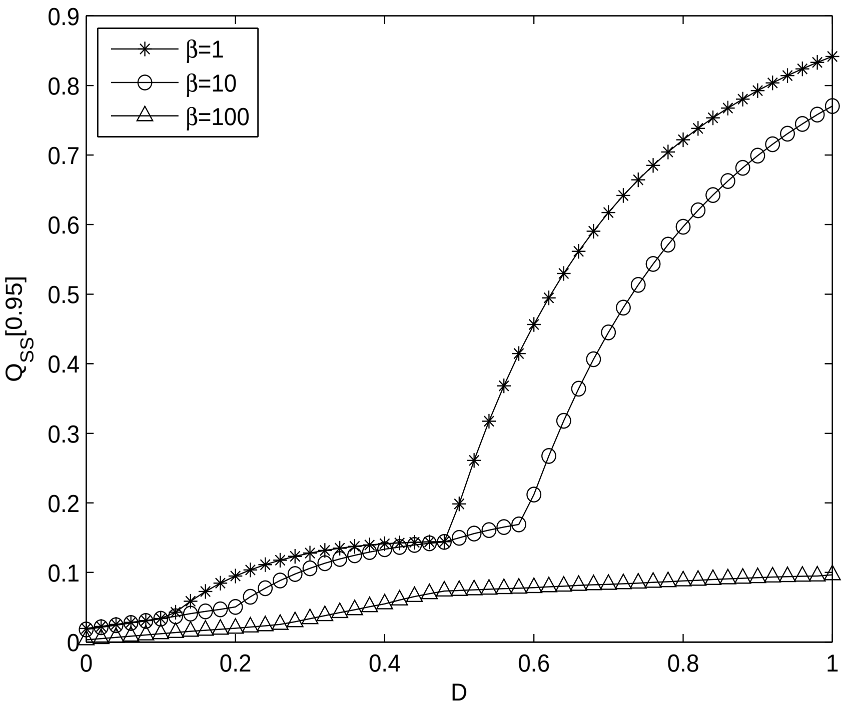


Figure 6: Minimum proportion of cargo containers selected for secondary screening for a detection probability of 0.95 as a function of the level of dependence $D$ for $n = 5$

## 7 Conclusions

This paper introduces CRKP, a linear progrmaming model for screening cargo containers for nuclear material at security stations using knapsack problem, reliability, and Bayesian probability models. The approach designs and analyzes security system architectures, and it provides a risk-based framework for determining how to define a system alarm when screening cargo containers given limited screening resources. Analysis of the models suggests that accurate prescreening intelligence is the most important factor for effective screening, particularly when sensors are highly dependent, and that sensors with high true alarm rates can mitigate some of the risk associated with low prescreening intelligence and sensor dependencies.

Although the framework introduced in this paper is idealized, it addresses important screening issues faced at cargo container security stations in ports. For example, it has been reported that cargo containers at a port security station are screened by an RPM, and those that yield an alarm response are screened by a second RPM. Cargo containers that yield a second RPM alarm response are selected for secondary screening (Lava 2008). In this context, cargo containers are screened by $n = 2$ sensors, with the system alarm threshold defined as two sensor alarms. CRKP can be used as a general framework to determine how to design next-generation security screening system as well as define a system alarm for any type of problem that relies on a series of screening devices or methods, risk assessments, and a limited secondary screening budget.

There are several possible extensions to CRKP. One extension is to consider CRKP as one component in a larger access security system, with secondary screening as additional components in the system (Kobza and Jacobson 1997, Christer 1994), and to investigate the dependencies between the components.

A second extension to CRKP is to consider a second level of classification for each of the containers. CRKP assumes that each container is classified as high-risk or low-risk, which quantifies the likelihood of the container containing nuclear material. However, the vast majority of system alarms encountered by our nation's ports are due to naturally occurring radioactive material (NORM) alarms, not nuclear materials (Huizenga 2005, The Royal Society 2008). Prescreening can be used to identify which containers have high levels of naturally occurring radiation, and hence, each cargo container can be classified as NORM or non-NORM as well as high-risk or low-risk. Analyzing how the two levels of classification as well as their interaction may shed light on the tradeoffs between prescreening intelligence and the physical contents and characteristics of the containers.

A third extension to CRKP is to differentiate the type of threat, affecting the probability of a true alarm at a given sensor. The probability that a threat container yields an alarm response at a sensor depends on the type of the source, the size of the source, the amount of shielding, and the location of the nuclear material within the container (Fetter et al. 1990, Levi 2007). This can be addressed by identifying a spectrum of threat scenarios as well as the likelihood of each scenario occurring. Work is in progress to address all of these extensions.

## Acknowledgements

The authors would like to acknowledge Colonel William H. Parrish and Dr. Jason R.W. Merrick at Virginia Commonwealth University for the discussions that motivated this research. This material is based upon work supported by the U.S. Department of Homeland Security under Grant Award Number 2008-DN-077-ARI001-02 and by the National Science Foundation (CBET-0735735). The computational work was done at Virginia Commonwealth University. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the U.S. Department of Homeland Security or the National Science Foundation.

## References

[1] N. O. Bakir. A decision tree model for evaluating countermeasures to secure cargo at United States Southwestern ports of entry. *Decision Analysis*, (to appear), 2009.

[2] E. Boros, L. Fedzhora, P. B. Kantor, K. Saeger, and P. Stroud. Large scale LP model for finding optimal container inspection strategies. Rutcor Research Report RRR 26-2006, Rutgers University, Piscataway, New Jersey, 2006.

[3] T. B. Cochran and M. G. McKinzie. Detecting nuclear smuggling: Radiation monitors at U.S. ports cannot reliably detect highly enriched uranium, which onshore terroists could assemble into a nuclear bomb. *Scientific American*, March 2008.

[4] S. Fetter, V. A. Frolov, M. Miller, R. Mozley, O. F. Prilutsky, S. N. Rodionov, and R. Z. Sagdeev. Detecting nuclear warheads. *Science & Global Security*, 1:225 – 302, 1990.

[5] J. F. Fritelli. Port and maritime security: Background issues for congress. CRS Report for Congress, Congressional Research Service, The Library of Congress, RL31733, 2005.

[6] N. Goldberg, J. Word, E. Boros, and P. Kantor. Optimal sequential inspection policies. Rutcor Research Report RRR 14, October 1, 2008, Rutgers University, Piscataway, New Jersey, 2008.

[7] D. Huizenga. Detecting nuclear weapons and radiological material: How effective is available technology? Statement before the Subcommittee on Prevention of Nuclear and Biological Attacks and Subcommittee on Emergency Preparedness, Science and Technology, The House Committee on Homeland Security, June 21 2005.

[8] International Atomic Energy Agency. IAEA illicit trafficking database. Fact Sheet, 2007.

[9] P. Kantor and E. Boros. Deceptive detection methods for optimal security with inadequate budgets: The screening power index. Rutcor Research Report RRR 26-2007, Rutgers University, Piscataway, New Jersey, 2007.

[10] J. E. Kobza and S. H. Jacobson. Addressing the dependency problem in access security system architecture design. *Risk Analysis*, 16(6):801–812, 1996.

[11] J. E. Kobza and S. H. Jacobson. Probability models for access security system architectures. *Journal of the Operational Research Society*, 48(3):255–263, 1997.

[12] M. Koucky. Exact reliability formula and bounds for general $k$-out-of-$n$ systems. *Reliability Engineering and System Safety*, 82:229 – 231, 2003.

[13] J. Lava. U.S. Customs and Border Protection. DIMANCS/DyDAn/LPS Workshop on Port Security/Safety, Inspection, Risk Analysis, and Modeling, Piscataway, NJ, November 17-18, 2008, 2008.

[14] M. Levi. *On Nuclear Terrorism.* Harvard University Press, Cambridge, Mass., 2007.

[15] L. A. McLay, S. H. Jacobson, and John E. Kobza. The tradeoff between technology and prescreening intelligence in checked baggage screening for aviation security. *Journal of Transportation Security*, 1(2):107 – 126, 2008.

[16] L. J. Moffitt, J. K. Stranlund, and B. C. Field. Inspections to avert terrorism: Robustness under severe uncertainty. *Journal of Homeland Security and Emergency Management*, 2(3), 2005. Available at www.bepress.com/jhsem/vol2/iss3/3 (accessed 11/22/2006).

[17] D. P. Morton, F. Pan, and K. J. Saeger. Models for nuclear smuggling interdiction. *IIE Transactions*, 39:3–14, 2007.

[18] F. Pan. *Stoachastic Network Interdiction: Models and Methods.* PhD thesis, University of Texas, Austin, TX, 2005.

[19] W. H. Parrish. Personal Interview, 2008, May 5.

[20] Royal Society, The. Detecting nuclear and radiological materials. RS policy document 07/08, The Royal Society, London, UK, 2008.

[21] D. Slaughter, M. Accatino, A. Bernstein, J. Candy, A. Dougan, J. Hall, A. Loshak, D. Manatt, A. Meyer, B. Pohl, S. Prussin, R. Walling, and D. Weirup. Detection of special nuclear material in cargo containers using neutron interrogation. Report UCRL-ID-155315, Lawrence Livermore National Laboratory, Livermore, CA, 2003.

[22] C. Strohm. Investigators call cargo security program unreliable. *GovExec.com*, April 5 2006. available at www.govexec.com/dailyfed/0406/040506cdam3.htm, accessed on November 22, 2006.

[23] United States Customs and Border Protection. Secure Freight with CSI, Megaports. Fact sheet, Washington, D.C., 2007.

[24] United States Department of Transportation. America's container ports: Delivering the goods. Technical report, Research and Innovative Technology Administration, Bureau of Transportation Statistics, Washington, D.C., March 2007.

[25] United States Government Accountability Office. Preliminary observations on the status of efforts to improve the Automated Transport System. Gao-06-591t, Government Accountability Office, Washington, D.C., March 30 2006.

[26] L. M. Wein, Y. Liu, Z. Cao, and S. E. Flynn. The optimal spatiotemporal deployment of radiation portal monitors can improve nuclear detection at overseas ports. *Science and Global Security*, 15:211–233, 2007.

[27] L. M. Wein, A. H. Wilkins, M. Baveja, and S. E. Flynn. Preventing the importation of illicit nuclear materials in shipping containers. *Risk Analysis*, 26(5):1377–1393, 2006.