Ann Oper Res (2011) 187:65–87
DOI 10.1007/s10479-010-0717-y

Detecting nuclear materials smuggling: using radiography to improve container inspection policies

Gary M. Gaukler · Chenhua Li · Rory Cannaday · Sunil S. Chirayath · Yu Ding

Published online: 2 March 2010

Abstract This paper proposes a layered container inspection system for detecting illicit nuclear materials using radiography information. We argue that the current inspection system, relying heavily on the Automated Targeting System (ATS) and passive radiation detectors, is inherently incapable of reliably detecting shielded radioactive materials, especially highly enriched uranium (HEU). This motivates the development of a new inspection system, which is designed to address a fundamental flaw of the ATS-based system, allowing for improved defense against sophisticated adversaries. In the proposed inspection system, all cargo containers go through x-ray imaging equipment first. From the x-ray image, a hardness measure of the container is computed. This hardness measure characterizes how likely it is that shielded HEU, if it does exist in the container, will not be detected in a subsequent passive detection step. Depending on the value of the hardness, the lower-hardness containers are sent to passive detection and the high-hardness containers are sent directly to active detection. This paper explores the trade-off between the detection probability of the new inspection system and the expected sojourn time a container spends in the system. The solution details and decision-making tools for using such a system are provided. Comparisons are made between the proposed system and the current ATS-based nuclear inspection system.

Keywords Inspection system · Nuclear materials · Radiography information · Queueing network

G.M. Gaukler (✉) · C. Li · R. Cannaday · Y. Ding
Department of Industrial and Systems Engineering, TAMU 3131, Texas A&M University, College Station, TX 77843, USA
e-mail: gaukler@tamu.edu

S.S. Chirayath
Nuclear Security Science & Policy Institute, TAMU 3473, Texas A&M University, College Station, TX 77843, USA

1 Introduction

Following the events of September 11, the US government has taken an increasingly proactive role in attempting to prevent any further terrorist acts on American soil (O'Hanlon et al. 2003). One of the most fearsome scenarios is a possible attack by a terrorist group using a nuclear weapon made of highly enriched uranium (HEU) or plutonium. There are several ways in which a nuclear weapon could be delivered to a target inside the United States. One possibility is that the attackers attempt to smuggle an already assembled device into the United States; the other possibility is to smuggle only the nuclear materials and assemble the weapon on US soil. The second scenario poses the greater challenge to the defender, because smaller (sub-critical) quantities of nuclear materials emit significantly less radiation that can be detected by sensors, compared to a fully assembled weapon. Therefore, the US government has become increasingly concerned about the threat posed by unsecured weapons-grade nuclear materials.

In response to this concern, the US has undertaken a number of initiatives in partnership with Russia and other former Soviet states to prevent the theft of fissile material (Ball 1998). The Material Protection, Control and Accounting Program (MPC&A), which was begun in 1993 to prevent the theft of nuclear materials from Russian civilian complexes, is considered the "First Line of Defense". The Second Line of Defense (SLD) program at the Department of Energy's National Nuclear Security Administration (NNSA) seeks to interdict illicit trafficking of nuclear and radiological materials through airports, seaports, and border crossings in Russia and other key transit states. It strives to do this by helping states install and use radiation detection equipment at these sites, and providing associated training and support.

Despite these efforts, according to the International Atomic Energy Agency (IAEA) database (IAEA 2006), between 1993 and 2005, there were 827 confirmed incidents of illicit trafficking and other unauthorized activities involving nuclear and radioactive materials. Of these 250 involved nuclear materials and 16 involved HEU and plutonium. Even small amounts of nuclear materials are worrisome because as little as 15 kilograms of HEU or 6 kilograms of plutonium could be used to build a simple implosion weapon (UCS 2004).

In line with the perceived threat of nuclear smuggling, and the potential vulnerability of the existing defense system, this research focuses on developing an advanced inspection system for detecting nuclear materials smuggling. We apply this inspection system in a numerical study to the particular case of detecting smuggled HEU (70% enriched). This material was chosen because it arguably presents the greatest current threat (Cochran and McKinzie 2008): HEU is much simpler to assemble into a weapon than plutonium, and while plutonium can fairly easily be passively detected by its bright gamma and neutron emissions, the gamma and neutron emissions from HEU are several orders of magnitude lower. Thus, HEU can be easily shielded and is harder to detect than Plutonium. Even though we focus on detecting HEU, our inspection policy applies to plutonium detection or the detection of complete/assembled nuclear devices as well.

One of the major ways in which such nuclear materials may be transported is via container shipping, in particular through seaports. Seaports are critical gateways for international commerce, and maritime shipping containers play a vital role in the movement of cargo between global trading partners. More than 11.4 million shipping containers arrive at US ports every year carrying more than 95 percent of US imports by weight from outside North America (Fritelli 2005; US DOT 2007).

The current container inspection system at US and foreign seaports is a multi-layered system, as illustrated in Fig 1. It consists of prescreening through the Automated Targeting System (ATS); primary screening via passive radiation portal monitors; secondary screening (possibly active radiation detection); and manual inspection.

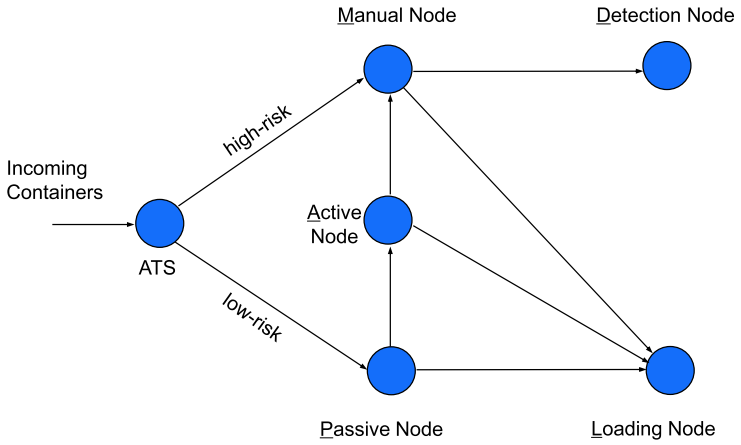

Fig. 1 Flowchart of current container inspection system

US Customs and Border Protection (CBP) uses the ATS to review documentation, including electronic manifest information submitted by the ocean carriers on all arriving shipments, to produce a risk score (Stana 2006; Caldwell 2008). CBP requires the carriers to submit manifest information 24 hours prior to a US-bound sea container being loaded onto a vessel in a foreign port. Based on an automated assessment of all these manifest data and other intelligence information, the ATS categorizes all US-bound containers into "high-risk" and "low-risk" containers.

All high-risk containers are sent directly to manual inspection, involving a physical search, during which a container is opened and its contents are removed for review. Low-risk containers, on the other hand, go through primary screening, using passive radiation detectors, such as the Radiation Portal Monitors (RPMs) (Aloise 2007). RPMs can pick up radiation emitted from nuclear materials (or from natural background radiation), while the container is passing through the monitors. Containers that fail in primary screening are then sent to secondary screening for further scrutiny or examination. The secondary screening may use handheld isotope scanning devices or active detectors, prototypes of which are available now but not broadly deployed yet. Active detectors bombard a container with low-energy neutrons, creating a telltale gamma-radiation signal that may be used to distinguish HEU from other naturally occurring radioactive materials (NORM). If a container fails the secondary screening, it will be manually examined by CBP officials.

Passive detection is the most time-efficient inspection method available. Active detection is more time-intensive and requires significantly more expensive equipment. Manual inspection is arguably the most capable detection method, but takes significantly greater efforts and time than all other methods. The logic behind the above inspection policy is to expend the best and highest inspection effort on containers that are "suspicious", and to keep overall throughput manageable by using less stringent checks on containers that are deemed not suspicious.

A major drawback of the current ATS-based inspection system is the possibility of mis-categorizing a dangerous container which contains HEU. Under a perfect ATS system, using perfect intelligence, the probability that a container with HEU inside is categorized as high-risk is 1. In that case, the ATS-based system will achieve the highest possible detection probability, because it sends every such container to manual inspection. In a real-world

scenario with imperfect intelligence however, this probability will be less than 1, opening the door for sending a dangerous container to less stringent inspection. At the same time, an intelligent adversary with knowledge of the existence of the ATS system will likely attempt to game this system, e.g. by placing the HEU in a container originating from an unsuspicious country, in a shipment from a respected company using an established carrier etc.

In short, we argue that the adversary will always try to infiltrate a container that is highly unlikely to be categorized as high-risk. Of course, the adversary may fail to infiltrate such a container successfully. However, if he does succeed, the current inspection system (which will send such a container to passive inspection) may not be adequate. The reason for this is that the radiation (gamma and neutron) of HEU in particular is easily shielded with lead or other high-z materials (the z-value, also called atomic number, describes the number of protons in an atomic nucleus). It is therefore extremely difficult for passive radiation detectors to distinguish a small, properly shielded quantity of HEU from the regular background radiation present everywhere (Srichrishna et al. 2005). This natural background radiation can vary significantly depending on a container's contents (e.g., through the presence of NORM, including cat litter, medical supplies, ceramics).

The inability of passive detectors to differentiate the shielded HEU from background radiation or NORM leads inevitably to either a high false positive rate (i.e., false alarms), or a high false negative rate (i.e., a missed detection). A high rate of false positives would send too many containers to the active or manual inspection stages, resulting in congestion and deteriorating port operations. On the other hand, a high rate of false negatives (or even just one false negative!) would mean that a dangerous container may escape from the system, potentially leading to a successful nuclear attack. However, it is also not realistic to indiscriminately send many containers to active or manual inspection, because active detectors and manual inspections are much more expensive and there typically are not enough active detectors installed (and manual inspection teams available) to handle all incoming containers.

In this paper, we introduce a new inspection policy system that is designed to avoid the fundamental flaw of the ATS-based system, allowing for improved defense against sophisticated adversaries. This new system is also formulated to effectively incorporate detection of shielded nuclear materials. At the core of this system is the ability to obtain a radiography image of every container. Based on this image, material densities and z-values of the cargo inside the container are obtained. From this information, a measure is computed that corresponds to how "hard" (or difficult) it is for a passive detector to differentiate between this particular container having a given quantity of HEU inside, and having no HEU inside. This measure takes into account the materials that make up the contents of the container and considers both the potential absorption of HEU radiation as well as the natural background radiation. We call this measure the "hardness" of that particular container. By adding this new hardness measure into the current practice, the incoming container stream is split into "hard" containers and "easy" containers. Easy containers (containers with small hardness value), are sent to passive detectors directly, while hard containers (containers with high hardness value), are sent directly for active detection. Containers that fail in the passive detection undergo active detection for further examination (see Fig. 2).

Ideally, the hardness measure is obtained based on radiography input. Since x-ray radiography technology is already being used in field practice at several US and foreign ports (Caldwell 2008; US Customs House Guide 2008), it is technologically plausible to add this hardness measure in the very near future. Even if the radiography input is not available, the bill of lading (BOL), which itemizes the contents of the container, may present a workable substitute. For this paper, we will assume that the hardness measure is computed based on radiography data.

This paper analyzes the proposed hardness-based system and explores the trade-off between the detection probability and the expected sojourn time a container spends in the system. The solution details and decision-making tools for using this system are provided. Comparisons are made between the proposed system and the current practice system.

The remainder of this paper is organized as follows. Section 2 provides a literature review of related research that focuses on preventing nuclear materials smuggling. Section 3 develops the model of the detection system, explains the details of each node in the detection system and derives the detection probability of the system. The queuing network model is described in Sect. 4. Section 5 presents numerical examples and explores the trade-off between the detection probability and the sojourn time. We conclude the paper in Sect 6.

2 Related work

Preventing illicit nuclear materials from entering the US has become a prominent issue not only in the security community (Cochran and McKinzie 2008; Fritelli 2005; Moffitt et al. 2005), but also the academic community. Several researchers have applied operations research methodologies to the problem of preventing nuclear materials smuggling.

The inspection system in this paper utilizes different types of detection technology. The work by Wein et al. (2006) is similar to our model in that it also takes into consideration both passive technology and active detection technology in the security system. Wein et al. describe an 11-layer security system consisting of shipper certification, container seals, and a targeting software system, followed by passive, active, and manual testing at overseas and domestic ports. A Stackelberg game is used to formulate the attacker-defender strategies in light of this layered inspection system, and the optimal inspection strategy (for the defender) is obtained while considering constraints on budget and port congestion. The key difference between our work and Wein et al.'s work is that our proposed system utilizes radiography information of the containers and includes a hardness criterion for dispatching the containers to the subsequent detection steps. This radiography node is not part of the 11-layer system considered by Wein et al. (2006).

As part of the SLD program, Pan (2005) and Morton et al. (2007) develop a class of stochastic network interdiction models that attempt to answer the question of where to build inspection stations or place detectors along a long border. This line of work is to address the strategic problem in stopping nuclear materials smuggling, while our work is a tactical problem of deciding the specific inspection policy at a given location (in this paper it is at a sea port, but in principle it can be at a border crossing as well). In their work, the inspection policy is assumed to be given.

Other research on container inspection systems at port-of-entry, but not necessarily with a focus on nuclear materials, is due to Boros et al. (2006), Stroud and Saeger (2003), Madigan et al. (2007), and Boros et al. (2008). Boros et al. (2006) develop a large scale linear programming model for sequential container inspection to determine an optimal inspection strategy under various constraints (budget, sensor capacity, time limits, etc.). Stroud and Saeger (2003) and Madigan et al. (2007) formulate the inspection sequencing task as a problem of finding an optimal binary decision tree for an appropriate Boolean decision function. Boros et al. (2008) review and present the research work for improving container inspection procedures using optimization techniques.

There are other security and inspection systems considering different layers of detection, such as in aviation security systems. Kobza and Jacobson (1997) consider bags entering the aviation security system, and taking several sub-paths with different detection technologies. However, their work does not consider the queuing results presented as in this paper

and in Wein et al. (2006). Other examples of research on inspecting aviation passenger baggage for explosives are Kobza and Jacobson (1996), and McLay et al. (2008). Kobza and Jacobson (1996) present a method to quantify the effect of dependence in security system architectures. They examine the Type I and Type II errors of a multi-device system, and present probability models for access control security system architectures. McLay et al. (2008) examine risk-based issues in detecting explosives in aviation security baggage screening models. The trade-off between intelligence and screening technology capabilities is explored using a cost-benefit analysis. All these papers focus on detecting explosives, other than nuclear materials.

3 Nuclear detection modeling

In this and the following sections we provide a detailed analysis of the proposed nuclear detection system with a radiography node and using a hardness measure. Figure 2 shows the layout of the proposed system, which is composed of the following: the radiography node (R-node), a passive node (P-node), an active node (A-node), a manual node (M-node), the loading node (L-node), and the detection node (D-node). The passive, active, and manual nodes function in the same way as their counterparts in Fig. 1. The R-node is the place where imaging x-ray equipment scans the contents of a container, leading to the computation of the container's hardness value. Containers leave the system from either the L-node or the D-node.

The inputs to this model are: the container scenario setting q_s (to be described later) and their prior distribution $P(q_s)$, the arrival rate of incoming containers λ, the service rate of each node μ_i, and the number of servers at each node m_i, where $i = \{R, P, A, M\}$. We assume that a binary decision will be made at each of the three detection nodes (i.e., R, P, and A nodes), and that the decision is characterized by a threshold value. More specifically, after each detection node, there are two possible paths (or actions to be taken), and which path to take depends on the comparison between the reading or calculation from a detector and the corresponding threshold; thus, our inspection policy is a threshold policy. The decision variables are the thresholds at each detection node t_i, where $i = \{R, P, A\}$.

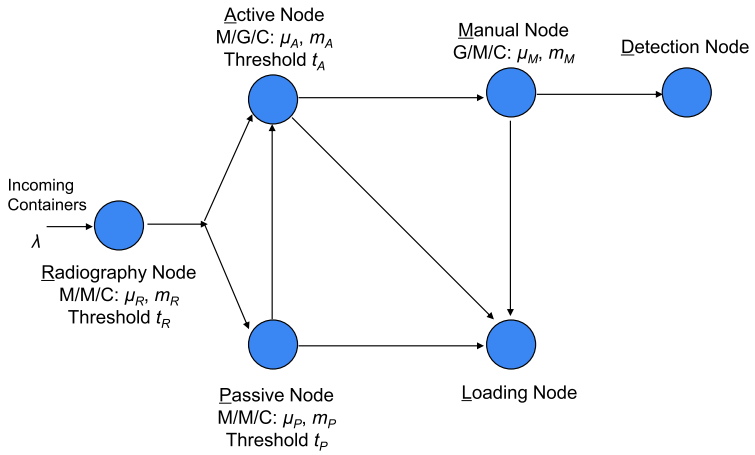

Fig. 2 Flowchart of the proposed inspection system

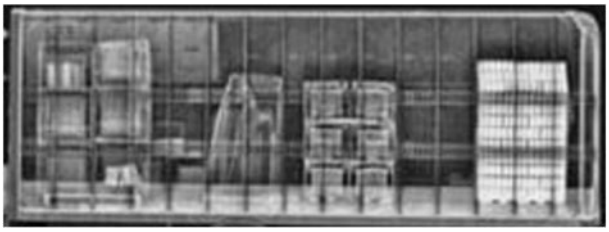

Fig. 3 A sample Z portal image

In our system, as well as in other inspection systems, two performance measures are of primary concern: one is the detection probability, i.e., the probability that a container with HEU inside finally arrives at the D-node, and the other is the system throughput, characterized by the expected sojourn time a container spends in passing through the inspection system. The second measure of system throughput is an implicit measure of false positives: If there are too many false positives, many containers will be sent to time-consuming detection nodes, which may subsequently significantly congest the system and delay the inspection process. Thus the two performance measures usually conflict with each other; when all other conditions are kept the same, one can only be improved at the expense of degrading the other.

Sections 3 and 4 are devoted to studying these two measures, respectively; Sect. 3 provides the detection probability models and Sect. 4 provides the models to compute the expected sojourn time. The way we determine the detection probability for the whole system is through analyzing the probability that a container will take a certain path at any given detection node. For this reason, Sects. 3.1 to 3.3 present this type of analysis at three detection nodes, respectively; and Sect. 3.4 calculates the probability that a container arrives at the D-node.

3.1 R-node

All incoming containers are sent to the R-node first. At the R-node, radiography equipment, such as Z portal and Z Backscatter systems (for details, see, e.g., AS&E 2008) are used to obtain x-ray images of each container. United States CBP have begun using Z portal systems at the San Ysidro, CA, border check point to inspect incoming vehicles for threats (US Customs House Guide 2008).

The Z portal or other types of radiography equipment provides a multi-view, photo-like image of the container's contents; one such example image is shown in Fig. 3 (taken from AS&E 2008). Based on the content images, containers are classified into a finite number of scenarios, denoted by $s \in S$, where S is the scenario set. The subsequent analysis at the R-node is scenario-based: for a given container scenario, we compute its hardness score, which indicates how difficult it is for a passive detector to differentiate between this particular container having a given quantity of HEU inside, and having no HEU inside. To complete the analysis, the following three issues need to be resolved: (a) the precise definition of a container scenario; (b) how to obtain the probability distribution of radiation emissions of a particular container scenario; and (c) how to compute the hardness of a particular container scenario.

(a) *The definition of a container scenario*:

The radiation emitting from the x-ray source is scattered and absorbed when passing through the container contents. The nuclear charge of the content materials determines the

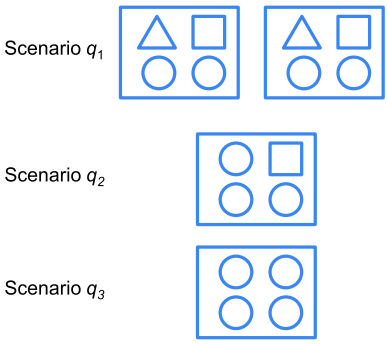

Fig. 4 Sample abstract container scenarios

scattering and absorption of the x-ray radiation. A high z-value material (having high density) absorbs more x-ray radiation and leaves a dark area in the resulting image; while a low z-value material (having low density) leaves a bright area for the same reason. From the darkness of the items (or areas) in the x-ray image, it is possible to infer the z-value of the corresponding contents in the container. When this is done, the container contents can be converted into a three-dimensional z-value matrix, which can be considered as the mathematical expression of the container scenario q_s. Thus, in principle, there could be a one-to-one relationship between container scenarios, and the x-ray images of the actual containers.

In reality, the container scenarios will likely be an abstraction of the actual x-ray images. The abstraction gets rid of some insignificant details in the x-ray image and shows regularized item shapes or boundaries according to the z-values in the surrounding areas. This abstraction can be achieved, for example, if an automated object recognition software is used to extract the shapes and boundaries of individual items. Such an abstraction step may be desirable to simplify the modeling of the container contents (for example, to facilitate the subsequent computation of the probability distribution of radiation readings). Figure 4 illustrates this type of abstract definition of container scenarios, where the big rectangle represents a container, and each shape within it represents one type of contents. Admittedly, the examples used in Fig. 4 vastly simplify what a container could contain; the actual container scenarios, even after some details are eliminated, are going to be much more complicated.

(b) *How to obtain the probability distribution of radiation emissions*:

Given a scenario q_s, the radiation reading in terms of gross particle counts that would be returned by a passive detector (before the actual detection takes place) can be described by a probability density function (pdf).

In our research, we use a nuclear physics simulation software, the Monte Carlo N-Particle (MCNP) code (X-5 Monte Carlo Team 2005), to obtain the pdf of gross particle counts of gamma radiation that is resulting from a container scenario.

MCNP is a software package developed at Los Alamos National Lab. The software is used for performing nuclear transport calculations for particles such as photons, neutrons and electrons. The software simulates the creation of particle emissions, particle trajectories, and particle interactions with surrounding materials, including absorption, reflection, scattering, and attenuation. For our purpose of evaluating nuclear radiation being emitted from a shipping container, we provide MCNP with the geometric and structural properties of a shipping container, as well as the z-value matrix representing the contents of the container. Radiation fluxes are computed by MCNP for defined regions around the container, and the flux values are then incorporated into a detector model to arrive at gross particle counts.

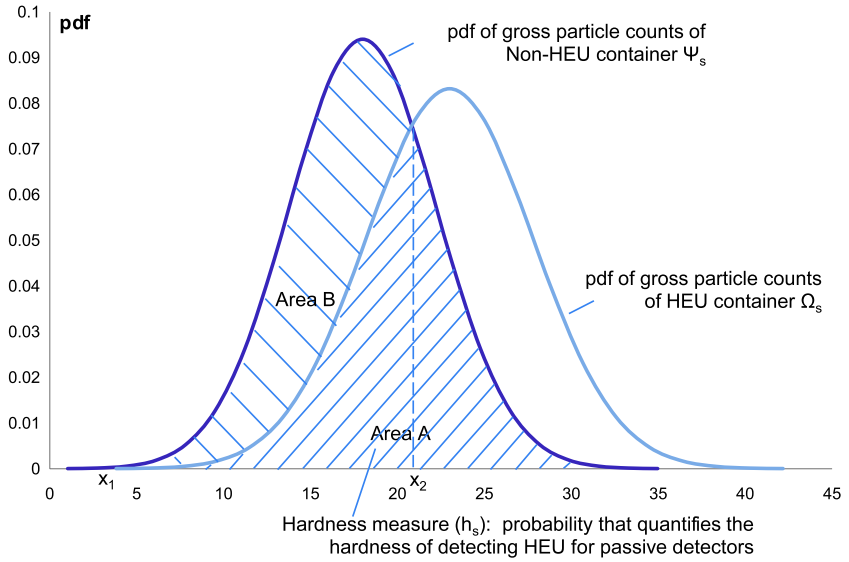

Fig. 5 Derivation of the hardness measure

Thus, the MCNP code takes the z-value matrix associated with a container scenario as its input, and mimics the working of a given type of passive (or active) detector. The MCNP code counts how many particle hits (i.e., influx) the simulated detector will receive for a specified exposure time. The MNCP code outputs the average influx per unit time per unit detector area.

For each given scenario q_s, the MCNP code will be used twice, one time when HEU is absent, and the other time when a given quantity of HEU is present (scenario q_s^{HEU}), corresponding to obtaining the two pdf's in Fig. 5. The hardness measure then corresponds to the overlapping area between the two pdf's, as illustrated by Area A in Fig. 5.

When HEU is absent, we assume that the natural background emission for scenario q_s is a normally distributed random variable Ψ_s. The normal distribution is the result of the normal approximation to the Poisson distributed particle counts. According to the nuclear engineering literature (e.g., Fetter et al. 1990), Ψ_s have the mean of $\mu_{\Psi_s} = A\epsilon\psi_s\tau_P$ and the standard deviation of $\sigma_{\Psi_s} = \sqrt{A\epsilon\psi_s\tau_P}$, where A is the area of the passive detector, ϵ is the efficiency of the detector, τ_P is the passive detection time, and $\psi_{\mathbf{s}}$ is the average photon flux per unit time per unit area produced by the MCNP code.

When a given quantity of HEU is present, it is also assumed that the radiation emission for scenario q_s^{HEU} is a normally distributed random variable Ω_s with mean $\mu_{\Omega_s} = A\epsilon\omega_s\tau_P$ and standard deviation $\sigma_{\Omega_s} = \sqrt{A\epsilon\omega_s\tau_P}$. Here, ω_s, also produced by the MCNP code, is the average photon flux when HEU is present, including both the emission from the HEU and the natural background radiation.

Thus, the two pdf's in Fig. 5 are described by the following two functions:

$$f_{\Psi_s}(x) = \frac{1}{\sqrt{2\pi}\sigma_{\Psi_s}} e^{-(x-\mu_{\Psi_s})^2/2\sigma_{\Psi_s}^2}$$

$$f_{\Omega_s}(x) = \frac{1}{\sqrt{2\pi}\sigma_{\Omega_s}} e^{-(x-\mu_{\Omega_s})^2/2\sigma_{\Omega_s}^2}$$

(c) *How to compute the hardness of a particular container scenario*:

The hardness h_s of a scenario q_s is defined as the misclassification error that is determined by the two distribution functions in Fig. 5, namely

$$h_s := P(\text{conclude that HEU does not exist} \mid \text{HEU is present}) + P(\text{conclude that HEU exists} \mid \text{HEU is absent})$$

It follows that h_s is the intersection area of the two distributions, shown as Area A in Fig. 5. Since $\mu_{\Psi_s} < \mu_{\Omega_s}$, and $\sigma_{\Psi_s} < \sigma_{\Omega_s}$, the intersection points can be obtained as the solution to the following equation:

$$\text{Find } x \text{ such that } f_{\Psi_s}(x) = f_{\Omega_s}(x) \tag{1}$$

The intersection points are given by:

$$x_1 = \frac{-R_1 - \sqrt{R_1^2 - R_2}}{2(\sigma_{\Omega_s}^2 - \sigma_{\Psi_s}^2)}$$

$$x_2 = \frac{-R_1 + \sqrt{R_1^2 - R_2}}{2(\sigma_{\Omega_s}^2 - \sigma_{\Psi_s}^2)}$$

where $R_1 = 2\mu_{\Omega_s}\sigma_{\Psi_s}^2 - 2\mu_{\Psi_s}\sigma_{\Omega_s}^2$, and $R_2 = 4(\sigma_{\Omega_s}^2 - \sigma_{\Psi_s}^2)(\mu_{\Psi_s}^2\sigma_{\Omega_s}^2 - \mu_{\Omega_s}^2\sigma_{\Psi_s}^2 + 2\sigma_{\Psi_s}^2\sigma_{\Omega_s}^2(\ln(\sigma_{\Psi_s}) - \ln(\sigma_{\Omega_s})))$.

The hardness of scenario q_s is then calculated as:

$$h_s = 1 - \int_{x_1}^{x_2} (f_{\Psi_s} - f_{\Omega_s})dx$$

where the integral in the above equation corresponds to Area B in Fig. 5.

The hardness h_s takes on a value in $(0, 1]$. When $h_s = 1$, meaning that the two pdf's are completely overlapping, it is impossible to make a confident distinction between the alternative hypotheses of HEU being present, or not. This corresponds to the hardest scenario for detection. On the other end of the spectrum, $h_s \to 0$, means that the two pdf's have almost no overlap. This corresponds to the situation that one can detect the presence of HEU with near certainty and represents the easiest case for detection.

At the R-node, we set t_R, a threshold such that any scenario with $h_s < t_R$ is sent to the P-node. Any scenario with $h_s \geq t_R$ is defined as "hard" , and is sent directly to the A-node. Thus, there are two sets of scenarios:

$$S_P = \{s : h_s < t_R,\ s \in S\}, \quad S_P \subset S$$

$$S_A = \{s : h_s \geq t_R,\ s \in S\}, \quad S_A \subset S$$

For a given q_s, the probability of going from the R-node directly to the P-node, denoted by $P(RP|q_s)$, is

$$P(RP|q_s) = \begin{cases} 1 & h_s < t_R \\ 0 & h_s \geq t_R \end{cases}$$

and the probability of going from the R-node directly to the A-node, denoted by $P(RA|q_s)$, is $P(RA|q_s) = 1 - P(RP|q_s)$.

Let α and β be the proportion of containers that go from the R-node to the A-node and the P-node, respectively. Then

$$\alpha = P(RA) = \sum_{s\in S} P(RA|q_s)P(q_s)$$

$$\beta = P(RP) = \sum_{s\in S} P(RP|q_s)P(q_s)$$

3.2 P-node

The container scenarios in the subset S_P, namely $q_s, s \in S_P$, go through passive radiation monitoring. A threshold value t_P is used to split the container stream to A-node or L-node. If the gross particle count obtained from the passive detector is greater than t_P, the container is sent to the A-node for active detection. Otherwise, it is sent to the L-node for loading.

Based on the normality assumption made in the previous section, the probability that scenario $q_s, s \in S_P$ goes to the A-node is:

$$P(PA|q_s) = P(\Psi_s > t_P) = 1 - \Phi_{\mu_{\Psi_s},\sigma_{\Psi_s}}(t_P) \quad \forall s \in S_P$$

where $\Phi_{\mu,\sigma}(\cdot)$ is the cumulative distribution function (CDF) of a normal distribution with mean μ and standard deviation σ. Thus, the proportion of containers going from the P-node to the A-node is:

$$P(PA) = \sum_{s\in S} P(PA|q_s)P(q_s|RP).$$

Denote by q_s^{HEU} a container scenario q_s with the presence of a given quantity of HEU. The probability that a container with HEU is correctly flagged at the P-node (i.e., the probability of going to the A-node) is:

$$P(PA|q_s^{\mathrm{HEU}}) = P(\Omega_s > t_P) = 1 - \Phi_{\mu_{\Omega_s},\sigma_{\Omega_s}}(t_P) \quad \forall s \in S_P$$

3.3 A-node

The A-node receives two streams of containers, one directly from the R-node and the other from the P-node. Denote by $P(A|q_s)$ the probability that a given scenario q_s reaches the A-node, which is computed as:

$$P(A|q_s) = \begin{cases} P(RA|q_s) & s \in S_A \\ P(RP|q_s) \cdot P(PA|q_s) & s \in S_P \end{cases}$$

Using Bayes' theorem, the conditional probability that a container reaching the A-node is of scenario q_s, is

$$P(q_s|A) = \frac{P(A|q_s)P(q_s)}{P(A)} = \frac{P(A|q_s)P(q_s)}{\sum_{s\in S} P(A|q_s)P(q_s)}$$

We still assume a normal distribution for the gross particle counts of gamma radiation to be detected at an active detector as we did for the passive detector. Hence, gross particle

counts under absence of HEU for scenario q_s follow a normal random variable Υ_s with mean μ_{Υ_s} and standard deviation $\sigma_{\Upsilon_s} = \sqrt{\mu_{\Upsilon_s}}$, while the gamma particle counts under presence of HEU follow a normal random variable Θ_s with mean μ_{Θ_s} and standard deviation $\sigma_{\Theta_s} = \sqrt{\mu_{\Theta_s}}$. These distribution parameters are also obtained using the MCNP code.

Suppose the threshold used at the A-node is t_A. The probability for a container scenario q_s to go from the A-node to the M-node is

$$P(AM|q_s) = P(\Upsilon_s > t_A) = 1 - \Phi_{\mu_{\Upsilon_s}, \sigma_{\Upsilon_s}}(t_A) \quad \forall s \in S$$

Thus, the proportion of containers going from the A-node to the M-node is

$$P(AM) = \sum_{s \in S} P(AM|q_s) \cdot P(q_s|A)$$

Further, the probability that a container with HEU is correctly flagged at the A-node (i.e., the probability of going to the M-node) is:

$$P(AM|q_s^{\text{HEU}}) = P(\Theta_s > t_A) = 1 - \Phi_{\mu_{\Theta_s}, \sigma_{\Theta_s}}(t_A) \quad \forall s \in S$$

3.4 Detection probability

We let d_M denote the probability of detecting the presence of HEU during the manual inspection at the M-node when HEU does exist. It is possible that the manual inspection team does not find the HEU even if that particular container arrives at the M-node. In that case, $d_M < 1$ describes the residual risk. In the numerical example presented in Sect. 5, we assume that $d_M = 1$ for the sake of simplification.

The detection probability for any given scenario q_s is the product of d_M and the probability that container scenario ends up at the M-node. Recall that for any scenario q_s, the container has two possible pathways to travel to the M-node: either R-node → A-node → M-node, or R-node → P-node → A-node → M-node, and that container scenarios are divided into two streams with S_A going to the A-node and S_P going to the P-node. Then, the probability that a container scenario q_s^{HEU} arrives at the M-node can be expressed as:

$$P(M|q_s^{\text{HEU}}) = \begin{cases} P(AM|q_s^{\text{HEU}}) & s \in S_A \\ P(PA|q_s^{\text{HEU}})P(AM|q_s^{\text{HEU}}) & s \in S_P \end{cases}$$

We define P_s to be the detection probability, conditioned on that q_s is given. Then:

$$P_s = d_M \cdot P(M|q_s^{\text{HEU}})$$

When the container scenario is not specified, we can compute the detection probability of the inspection system averaged over all possible scenarios:

$$P_{\text{overall}} = \sum_{s \in S} P_s \cdot P(q_s^{\text{HEU}})$$

where $P(q_s^{\text{HEU}})$ is the prior probability that a scenario q_s contains HEU. This prior probability may come from intelligence information. For the example in Sect. 5, we assume this prior probability to be uniform, meaning that all scenarios are equally likely to contain HEU. Incorporating $P(q_s^{\text{HEU}})$ also provides a way to model possible adversary behavior. For example, the adversary may choose to select a "hard" q_s (with larger h_s value) to transport

HEU in order to reduce detectable emissions, or to select a low-hardness q_s to reduce the likelihood to be singled out for more stringent interrogation by the radiography node.

We point out that our model not only provides the overall detection probability of the inspection system, but also the detection probability for each container scenario. To our knowledge, this is the first available model that allows the decision maker to obtain these detailed detection probabilities. This detailed information can be used by the decision maker to identify any weak points of the inspection system, or to manipulate the threshold values in order to obtain the desired detection probability for a certain container type.

4 Queuing network model

The second performance measure considered in our model is the expected sojourn time of a container in the system. More aggressive testing strategies result in longer sojourn time, thus higher delay cost. In order to find the trade-off between the detection probability and the sojourn time, a queuing network model is developed.

In this section, we explore the queuing results for all the nodes in the network. We assume that containers arrive at the port according to a Poisson process with rate λ. Since both the arrival process and the service process of the R-node and the P-node are Markov processes, the R-node and the P-node are modeled as $M/M/C$ queues, where the first M denotes the Markov arrival process, the second M denotes the Markov service process, and C denotes multiple servers. The service rate and arrival rate at the R-node are μ_R and $\lambda_R = \lambda$, respectively. We let m_R represent the number of servers at the R-node. The service rate at the P-node is μ_P, the number of servers at the P-node is m_P and the arrival rate at the P-node is $\lambda_P = \beta\lambda$. Thus, general $M/M/C$ queue results (Bolch et al. 2006) can be applied to obtain the expected waiting time at the R-node and the P-node, denoted by $E(W_R)$ and $E(W_P)$.

Although the arrival stream for the A-node is a Markov process, the service time is assumed to be a general distribution. This same assumption is made in Wein et al. (2006). Thus, queueing at the A-node is modeled as $M/G/C$, where G denotes the general distribution. Since the A-node is the only node upstream of the M-node, the arrival process, determined by the service process at the A-node, is assumed to be a general distribution. The M-node is modeled as a $G/G/C$ queue.

For the A-node and the M-node, the parametric-decomposition procedure is used to obtain the expected waiting times. In the parametric-decomposition procedure, each queue in a network is characterized by five parameters: the arrival rate λ_i, the coefficient of variation of inter-arrival time c_{ai}, the service rate μ_i for each server, the coefficient of variation of service time c_{si}, and the number of servers m_i, where $i = \{A,\ M\}$. Then the expected waiting time at a generic queue with those five parameters is approximated (Whitt 1983, 1993).

4.1 A-node queuing model ($M/G/C$)

At the A-node two streams merge, one directly from the R-node and another from the P-node. Let λ_A and μ_A denote the arrival rate and the service rate at the A-node and denote by m_A the number of servers at the A-node. The coefficient of variation of inter-arrival time and service time at the A-node are c_{aA} and c_{sA}, respectively. Since the A-node has two incoming streams, the arrival rate at the A-node is $\lambda_A = \alpha\lambda + P(PA) * \lambda_P$. Since a split or merging of a Poisson process is also a Poisson process, we assume $c_{aA} = 1$. Also, we assume that the active detector service time is an Erlang (order two) random variable with $c_{sA}^2 = 0.5$, which

is the same choice as in Wein et al. (2006). Then the coefficient of departure is obtained from those five parameters (Whitt 1983):

$$c_{dA}^2 \approx 1 + (1 - \rho_A^2)(c_{aA}^2 - 1) + \frac{\rho_A^2}{\sqrt{m_A}}(c_{sA}^2 - 1)$$

where $\rho_A = \frac{\lambda_A}{m_A \mu_A}$.

The expected waiting time is obtained as (Whitt 1983):

$$E(W_A) \approx \frac{c_{aA}^2 + c_{dA}^2}{2m_A\mu_A(1-\rho_A)(1+\sqrt{2\pi}\nu_A\Phi(\nu_A)e^{\nu_A^2/2})}$$

where $\Phi(\cdot)$ is the CDF of the standard normal distribution, and $\nu_A = \sqrt{m_A}(1 - \rho_A)$.

4.2 M-node queuing model $(G/G/C)$

The manual inspection station is downstream only from the active detection station. Therefore, the coefficient of variation of inter-arrival times is related to the proportion of containers that fail active interrogation, which is denoted by $P(AM)$. We can then find $\lambda_M = P(AM) \cdot \lambda_A$. Therefore the coefficient of variation of inter-arrival times is obtained as (Whitt 1983):

$$c_{aM}^2 = \frac{\lambda_M}{\lambda_A} c_{dA}^2 + 1 - \frac{\lambda_M}{\lambda_A}$$

The coefficient of variation of inter-departure times is obtained as (Whitt 1983):

$$c_{dM}^2 \approx 1 + (1 - \rho_M^2)(c_{aM}^2 - 1) + \frac{\rho_M^2}{\sqrt{m_M}}(c_{sM}^2 - 1)$$

where m_M is the number of servers at the M-node, and $\rho_M = \frac{\lambda_M}{m_M \mu_M}$, μ_M denote the service rate and arrival rate at M-node. The service time at the M-node is assumed to be general distribution. In the numerical example, it is assumed to be exponential, and therefore the coefficient of variation of service time at the M-node $c_{sM}^2 = 1$. Thus, the expected waiting time at the M-node is derived:

$$E(W_M) \approx \frac{c_{aM}^2 + c_{dM}^2}{2m_M\mu_M(1-\rho_M)(1+\sqrt{2\pi}\nu_M\Phi(\nu_M)e^{\nu_M^2/2})}$$

where $\nu_M = \sqrt{m_M}(1 - \rho_M)$, and it is necessary that $\rho_M < 1$ in order to make the queueing system stable.

4.3 Expected sojourn time

Let the expected time spent at node i be T_i, $i \in \{R, P, A, M\}$. The quantity T_i is the summation of the expected service time $1/\mu_i$ and the expected waiting time $E(W_i)$. For the proposed queuing network, there are five pathways that a container can take through the system: $R \to P \to L$, $R \to P \to A \to L$, $R \to P \to A \to M \to L$, $R \to A \to L$, and

$R \to A \to M \to L$. After calculating the probability that a given container scenario q_s takes each individual path, the expected sojourn time for scenario q_s is as follows:

$$T_s = \begin{cases} T_R + T_P & \text{with prob. } P(RP|q_s)(1 - P(PA|q_s)) \\ T_R + T_P + T_A & \text{with prob. } P(RP|q_s)P(PA|q_s)(1 - P(AM|q_s)) \\ T_R + T_P + T_A + T_M & \text{with prob. } P(RP|q_s)P(PA|q_s)P(AM|q_s) \\ T_R + T_A & \text{with prob. } P(RA|q_s)(1 - P(AM|q_s)) \\ T_R + T_A + T_M & \text{with prob. } P(RA|q_s)P(AM|q_s) \end{cases}$$

The expected sojourn time unconditional of the scenario is then obtained by averaging over all container scenarios:

$$T_{\text{overall}} = \sum_{s \in S} T_s \cdot P(q_s)$$

Similar to the detection probability in Sect. 3, both the overall expected sojourn time and the expected sojourn time for each container type are calculated in this model. The expected sojourn time for each container type is useful information for the shipper. This information may, for example, help in determining more accurate and granular predictions of shipping delays based on the type of container that is being shipped.

5 Numerical example

In this section, we present a numerical example. We compare the proposed hardness-based inspection system with the current ATS-based inspection system and investigate the trade-off between the detection probability and the expected sojourn time.

5.1 Description of container scenarios

Container capacity is often expressed in terms of twenty-foot equivalent units (TEUs). A TEU is a measure of containerized cargo capacity equal to one standard 20 foot (length) × 8 foot (width) × 8 foot (height) container. Commonly used cargo containers are 20-foot containers (one TEU) and 40-foot containers (two TEUs).

For our numerical example, we define four container scenarios, each one corresponding to a 20-foot container. The first container scenario 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 scenario is a mixture of medium z-value materials, whose z-value is between 10 to 20 (e.g., concrete), and low z-value materials. The third container scenario has high z-value materials, whose z-value is greater than 20 (e.g., iron), mixed with other medium and low z-value materials. It is assumed that the detector is well shielded from below to suppress the natural background gamma radiation coming from the ground. The average natural background photon flux rate for the first three scenarios is assumed to be $b_s = 5 \times 10^{-5}$ gamma rays per cm^2 per second. This reflects the background level when the ground radiation is shielded. The fourth and last scenario considers the case where other particle emissions (from NORM) exist. To account for this, we double the natural background emission level for the fourth scenario.

For the container scenarios where HEU is present (i.e., q_s^{HEU}), a one kg HEU sphere (30% of U-238, and 70% of U-235) with one cm lead shielding is placed in the center of the container. In each scenario, the HEU is placed such that the highest z-value materials are

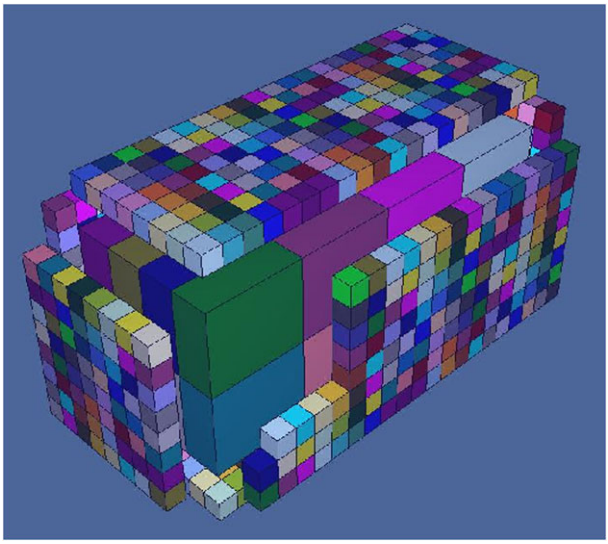

Fig. 6 3-D image of container model in MCNP

Table 1 Distribution parameters of gross counts of passive and active detection

	$P(q_s)$	Passive					Active			
		μ_{Ψ_s}	σ_{Ψ_s}	μ_{Ω_s}	σ_{Ω_s}	h_s	μ_{Υ_s}	σ_{Υ_s}	μ_{Θ_s}	σ_{Θ_s}
q_1	0.40	3.72	1.93	24.68	4.97	0.002	1.24	1.11	19.75	4.44
q_2	0.35	3.72	1.93	15.58	3.95	0.040	1.24	1.11	12.43	3.53
q_3	0.15	3.72	1.93	11.84	3.44	0.123	1.24	1.11	9.60	3.10
q_4	0.10	7.43	2.73	15.55	3.94	0.218	2.48	1.57	11.10	3.33

adjacent to it, thus modeling the adversary's attempt to take advantage of the innate shielding power provided by the container contents. Each container scenario is modeled in the MCNP code by using a z-value matrix (please recall the description in Sect. 3.1). Figure 6 is a 3-D picture illustrating a container scenario model used in the MCNP code. The large boxes that make up the container are the areas of the different z-value materials inside the container. The small boxes, along the outside of the container, represent the volumes in which the radiation flux is captured in MCNP.

We model the radiation detectors as follows. Detectors are placed on both sides of the container, with 1 foot distance from the wall of the container. The portal radiation monitor we use here has a height of 8 feet and width of 1 foot. Thus, the detector area is $A = 16$ ft^2. The efficiency of the detector is set at $\epsilon = 0.2$ (Glaser 2007), and the passive detection exposure time is assumed to be $\tau_P = \frac{1}{\mu_P} = 25$ seconds.

Based on the radiation model in Sect. 3.1, the distribution of gross counts at the passive detector and the associated hardness can be calculated for each scenario. The results are summarized in Table 1, where the distributions are parameterized by their means and standard deviations. Both the cases with and without the presence of HEU are computed. Also shown in Table 1 is the prior distribution of the container scenarios ($P(q_s)$).

Table 1 also includes the distribution parameters for the radiation attained by the active detector. Ideally, this information would be obtained from MCNP code. Since an active detector model in MCNP was not available to us, the active detector parameters for our

Table 2 Values for the model parameters

Parameter	Description	Value
A	Area of passive radiation detector	16 ft^2
ϵ	Efficiency of passive detector	0.2
b_s	Natural background radiation	5×10^{-5} gamma rays/cm^2 sec
λ	Arrival rate of container	90/hour
μ_R	Service rate at R-node	80/hour
m_R	Number of servers at R-node	2
μ_P	Service rate at P-node	144/hour
m_P	Number of servers at P-node	1
μ_A	Service rate at A-node	20/hour
m_A	Number of servers at A-node	2
μ_M	Service rate at M-node	1/hour
m_M	Number of servers at M-node	8

numerical example are assigned based on educated assumptions. We assume that the background level of an active detector is $1/3$ of that at the passive detector, and the hardness for detecting the presence of HEU at an active detector is $1/3$ of that of the passive detector.

In this study, we use four different container scenarios for the purpose of testing the concept. At any actual port, the number of scenarios is clearly going to be much higher. Nevertheless, the procedure of using that type of information and the analysis leading to the final decision should be the same as we present.

5.2 Choice of model parameters

In our example, we assume that the arrival rate of containers at the port is 90/hour, the same arrival value as used in Wein et al. (2006). The service time for the radiography machine is exponentially distributed with a mean of 45 seconds. We also assume that two radiography machines are available. For the passive detection, the service time is assumed to be exponentially distributed with a mean of 25 seconds; the number of passive detectors is assumed to be one. The service time for active detection is assumed to follow an Erlang distribution of order 2 with a mean time of 3 minutes, the same assumption as made in Wein et al. (2006). The service time at the manual detection stage is an exponential distribution with a mean time of one hour. Further, there are 2 active detectors and 8 manual service teams. These server numbers at the A-node and M-node are chosen to be the smallest possible numbers that make the queuing system stable. The values of the model parameters used in this numerical example are summarized in Table 2.

5.3 Performance measures and generation of efficient frontier

Recall that we have chosen two measures, the detection probability and the expected container sojourn time, to characterize the performance of an inspection system. In this paper, we choose to present the efficient frontier rather than provide a set of optimized decision variables. The reason is that we feel that using an efficient frontier illustrates better the trade-off between the two conflicting performance measures. We here mean to provide an operations research tool that enables the final decision maker to make an informed decision based on a complete description of the trade-off involved.

The efficient frontier is generated by solving a sequence of the following optimization problems for discrete choices of sojourn time limit t:

$$\max_{t_R,t_P,t_A} \text{Detection probability}$$
$$\text{s.t. Expected sojourn time} \leq t$$

One can also add other constraints to the above optimization problem in order to confine the length of queueing at each node if the waiting space is limited. In our example, the waiting area is assumed to be unconstrained. Execution times are rapid and not a limiting factor.

5.4 System comparison and discussion

In this section, we refer to our proposed system as the hardness computation system (HCS). We compare the proposed HCS with the current ATS-based system to obtain information on the potential merit of using the radiography information and hardness criterion.

Before proceeding with the comparison, we need to assign a measure of "trust" to the ATS scores. As we argued before, if the ATS is perfectly accurate, then it is (from a detection probability standpoint) the optimal system, since it will always assign any container with HEU to the set of high-risk containers. In reality, the intelligence used in the ATS is likely less ideal. To make a fair comparison, we need to assess and decide how reliable the ATS classification can be.

We first introduce the following notation:

- $P(\text{HR})$ is the probability that a container is classified as "high-risk",
- $P(\text{LR})$ is the probability that a container is classified as "low-risk",
- $P(\text{HEU})$ is the probability that a container contains HEU.

To quantify the reliability of the ATS system, we adopt the "trust" measure δ_{ATS} introduced in McLay et al. (2008), namely,

$$\delta_{\text{ATS}} = \frac{P(\text{HEU}|\text{HR})}{P(\text{HEU}|\text{LR})}.$$

Simply put, this trust measure is the ratio of how likely a high-risk container indeed contains HEU versus that a low-risk container contains HEU. If this ratio is one, meaning that high-risk containers and low-risk containers have the same likelihood to contain HEU, this indicates that the intelligence is no better than random sampling. On the other hand, if δ_{ATS} is large, it means that the high-risk containers are significantly more likely to contain HEU than the low-risk containers. In this circumstance, the value of intelligence is great. In our comparison, we consider three δ_{ATS} values: 1, 10, and 50.

Given the definition of the trust measure δ_{ATS} and $P(\text{HR}|\text{HEU}) + P(\text{LR}|\text{HEU}) = 1$, using Bayes' theorem, we obtain:

$$P(\text{HR}|\text{HEU}) = \frac{P(\text{HEU}|\text{HR})P(\text{HR})}{P(\text{HEU}|\text{LR})P(\text{LR}) + P(\text{HEU}|\text{HR})P(\text{HR})}$$
$$= \frac{\delta_{\text{ATS}} P(\text{HR})}{(1 - P(\text{HR})) + \delta_{\text{ATS}} P(\text{HR})}$$

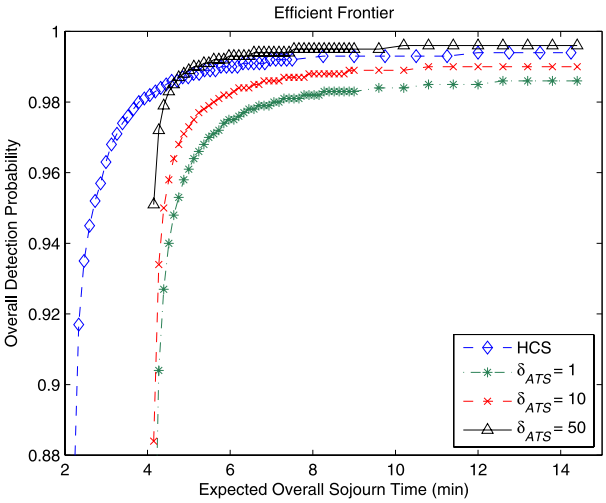

Fig. 7 Comparison of overall efficient frontier of HCS and ATS

Therefore, the detection probability of the current ATS-based system, denoted by P_s^{ATS}, is:

$$P_s^{\text{ATS}} = P(\text{HR}|\text{HEU}) \cdot d_M + P(\text{LR}|\text{HEU}) \cdot P(\text{PA}|q_s^{\text{HEU}}) \cdot P(\text{AM}|q_s^{\text{HEU}}) \cdot d_M$$

In our numerical example, we assume that the proportion of containers that are classified as "high-risk" is $P(\text{HR}) = 0.05$. The same value is used in McLay et al. (2008). In addition, we assume that the ATS classification is independent of the container content.

We present the system comparison in Figs. 7 and 8. Figure 7 shows the efficient frontiers of both systems by aggregating over the four scenarios, and Fig. 8 presents four plots, corresponding to the four individual scenarios.

Figure 7 shows four curves. There is only one efficient frontier for the HCS system, because the performance of HCS does not depend on the choices of δ_{ATS}. The other three curves are the efficient frontiers of the ATS system under $\delta_{\text{ATS}} = 1$, 10, and 50, respectively. The efficient frontiers of ATS have different length from that of the HCS because we choose the same spans of the threshold values t_R, t_P, and t_A for the HCS and ATS, leading to the different ranges of detection probability and expected sojourn time for the two systems.

From Fig. 7, we observe that for the mix of container scenarios in our numerical example, HCS outperforms ATS under almost all the circumstances because the efficient frontier of HCS dominates almost entirely the efficient frontier of ATS. One exception is for the case of $\delta_{\text{ATS}} = 50$: in the upper part of the efficient frontier, the detection probability of ATS is higher than that of HCS. Note that $\delta_{\text{ATS}} = 50$ represents a very high degree of trust in the ATS. Nevertheless, even under $\delta_{\text{ATS}} = 50$, the advantage of HCS in terms of having a short expected sojourn time is still obvious at the lower part of the curve. This is a pleasant surprise because decision makers tend to be reluctant to add any new detection node out of the worry that it may slow down the inspection process. What we show here is that if we add the right node and do it intelligently, it can help reduce inspection time, improve throughput, and maintain or even increase detection probability.

Figure 8 shows the four efficient frontiers corresponding to the four container scenarios under $\delta_{\text{ATS}} = 10$. We choose $\delta_{\text{ATS}} = 10$ because we believe $\delta_{\text{ATS}} = 10$ reflects more realistically the current quality of intelligence used in ATS. This same argument and choice of parameter value is also made by McLay et al. (2008).

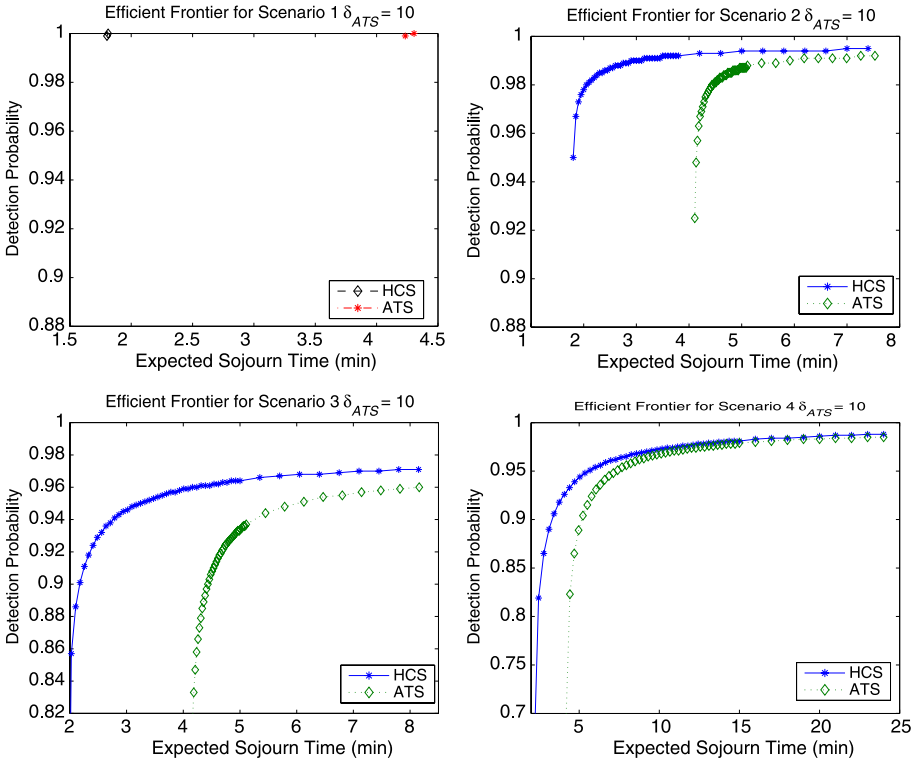

Fig. 8 Efficient frontier for different scenarios with $\delta_{\mathrm{ATS}} = 10$

Scenario 1 (top-left figure) represents a container full of textiles, plastic and wood, and thus $h_s = 0.002$ (refer to Table 1). In this case, it is very easy to detect the presence of HEU, even for a passive detector. The efficient frontier is a short line spanning a very narrow range of detection probabilities close to 1. Presumably, this is a type of container scenario that a sophisticated adversary would never select to hide HEU inside. The containers of scenario 1 without HEU most likely will take the path from $R \rightarrow P \rightarrow L$. Since there are only very slight changes in the detection probability and expected sojourn time, the efficient frontier degenerates to a short line on the top-left figure. The expected sojourn time of the HCS model is shorter than for the ATS model.

Scenario 2, with mixed contents of medium and low z-value materials, has $h_s = 0.04$, also an easy scenario.The efficient frontier is a smooth curve. In the lower part of the curve, for the same detection probability, the expected sojourn time of the HCS model is shorter than the expected sojourn time of the ATS model. Furthermore, in the upper part of the curve, the HCS model is superior to the ATS model because of its higher detection probability.

Scenario 3 contains some high z-value contents (like iron), which makes it hard for passive detectors to detect the presence of HEU. The lower part of the HCS curve represents the situation where the containers are entirely sent to the P-node, which results in smaller sojourn time and lower detection probability, and the upper part of the curve represents the situation where the containers are sent directly to the A-node, elevating the already saturated efficient frontier from the lower part to a higher level. Under this scenario, the HCS model significantly outperforms the ATS model.

Scenario 4 is the most difficult case with the highest h_s among the four scenarios. The container holds an amount of NORM (such as fertilizer) that emits radiation, which increases the level of the natural background radiation. This makes it more difficult for the passive detector to decide whether HEU is present. For scenario 4, the HCS model outperforms the ATS model with a shorter sojourn time in the lower part and higher detection probability in the upper part.

6 Conclusion

In this paper, we propose, model, and analyze a container inspection policy for detecting nuclear materials smuggling. This new inspection system adds a radiography node before conducting any other radiation detection. The addition of this radiography node is to address the inherent inability of passive radiation detectors to detect shielded HEU in containers with a large quantity of high z-value materials and/or NORM. Based on the radiography information obtained, a "hardness" measure is computed. Containers with natural or artificial shielding materials (usually high z-value materials) are deemed "hard" containers and have a higher probability of being sent through more stringent inspection steps (such as through active detectors).

Through numerical examples and comparisons, we demonstrate that the proposed system can perform significantly more effectively than the ATS-based system. More surprisingly, adding the radiography node does not slow down the inspection process; in many circumstances, it actually improves the system throughput.

Despite the superior performance demonstrated by the proposed inspection system in the four chosen container scenarios, we are not advocating to completely replace the current ATS-based system. Field intelligence is valuable, and containers deemed to be "suspicious" deserve extra attention. However, decision makers ought to be aware that a pure ATS-based system has significant vulnerabilities that a sophisticated adversary can, and likely will, exploit. In our opinion, a useful step towards improving the nation's safety from nuclear terrorism would be to use a hybrid system that continues to use the ATS, but complements it with intelligent use of radiography information in a similar manner as presented here. A possible hybrid system approach could be the following: continue to use the ATS to categorize containers prior to departure and during the voyage. Then, at the domestic port, route high-risk containers directly to manual inspection, and route low-risk containers to the radiography node of the HCS system. It is easy to see that with this particular hybrid system structure, the detection probability is higher than either under ATS or HCS alone; however, this higher detection probability may come at the price of longer expected sojourn times than under the pure HCS system.

Acknowledgements The authors gratefully acknowledge financial support for this research, provided by the Domestic Nuclear Detection Office (DNDO) under grant ARI-LA 2007. We thank our colleagues Wolfgang Bangerth, David Boyle, Bill Charlton, Pete Miller, Paul Nelson, and Arnold Vedlitz for many fruitful discussions that have contributed to enhancing this paper.

References

Aloise, G. (2007). *Combating nuclear smuggling: DHS's decision to procure and deploy to the next generation of radiation detection equipment is not supported by its cost-benefit analysis* (US Government Accountability Office Report GAO-07-581T).

AS&E (2008). Cargo & vehicle inspection. http://www.as-e.com/products_solutions/cargo_vehicle_inspection.asp. Accessed 26 December 2008.

Ball, D. Y. (1998). The US second line of defense: preventing nuclear smuggling across Russia's borders. PONARS Policy Memo 50. Lawrence Livermore National Laboratory. http://www.csis.org/media/csis/pubs/pm_0050.pdf. Accessed 26 December 2008.

Bolch, G., Greiner, S., Meer, H. D., & Trivedi, K. S. (2006). *Queueing networks and Markov chains: modeling and performance evaluation with computer science application*. New Jersey: Wiley.

Boros, E., Fedzhora, L., Kantor, P. B., Saeger, K., & Stround, P. (2006). *Large scale LP model for finding optimal container inspection strategies* (Rutcor Research Report).

Boros, E., Elsayed, E., Kantor, P., Robert, F., & Xie, M. (2008). Optimization problems for port-of-entry detection systems. In Chen, H., & Yang, C. C. (Eds.), *Intelligence and security informatics: techniques and applications* (pp. 319–335). Berlin: Springer.

Caldwell, S. L. (2008). *Supply chain security: examinations of high-risk cargo at foreign seaports have increased, but improved data collection and performance measures are needed* (US Government Accountability Office Report GAO-08-187).

Cochran, T. B., & McKinzie, M. G. (2008). Detecting nuclear smuggling: radiation monitors at U.S. ports cannot reliably detect highly enriched uranium, which onshore terrorists could assemble into a nuclear bomb. *Scientific American*, *298*(4), 98–104.

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

Fritelli, J. F. (2005). *Port and maritime security: background issues for congress* (CRS Report for Congress, Congressional Research Service, the Library of Congress RL31733).

Glaser, A. (2007). Detection of special nuclear materials. Princeton University. http://www.princeton.edu/~aglaser/lecture2007_detection.pdf. Accessed 26 December 2008.

International Atomic Energy Agency (2006). Illicit trafficking and other unauthorized activities involving nuclear and radioactive materials fact sheet. http://www.iaea.org/NewsCenter/Features/RadSources/PDF/fact_figures2005.pdf. Accessed 26 December 2008.

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

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

Madigan, D., Mittal, S., & Roberts, F. (2007). Sequential decision making algorithms for port of entry inspection: overcoming computational challenges. In Muresan, G., Altiok, T., Melamed, B., & Zeng, D. (Eds.), *Proceedings of IEEE international conference on intelligence and security informatics (ISI-2007)* (pp. 1–7). New Jersey: IEEE Press.

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

Moffitt, L. J., Stranulund, J. K., & Field, B. C. (2005). Inspections to avert terrorism: Robustness under severe uncertainty. *Journal of Homeland Security and Emergency Management 2*(3).

Morton, D. P., Pan, F., & Sager, K. J. (2007). Models for nuclear smuggling interdiction. *IIE Transactions*, *39*, 3–14.

O'Hanlon, M. E., Orszag, P. R., Daalder, I. H., Destler, I. M., Gunter, D. L., Lindsay, J. M., Litan, R. E., & Steinberg, J. B. (2003). *Protecting the American homeland*. Washington: Brookings Institution Press.

Pan, F. (2005). *Stochastic network interdiction: models and methods*. Ph.D. thesis, University of Texas, Austin, TX.

Srichrishna, D., Chari, A. N., & Tisch, T. (2005). Nuclear detection: fixed detectors, portals, and NEST teams won't work for shielded HEU on a national scale; a distributed network of in-vehicle detectors is also needed to deter nuclear terrorism. http://www.devabhaktuni.us/research/disarm.pdf. Accessed 22 December 2008.

Stana, R. M. (2006). *Cargo container inspections: preliminary observations on the status of efforts to improve the automated targeting system* (US Government Accountability Office Report GAO-06-591T).

Stroud, P. D., & Saeger, K. J. (2003). Enumeration of increasing boolean expressions and alternative digraph implementations for diagnostic applications. In Chu, H., Ferrer, J., Nguyen, T., & Yu, Y. (Eds.), *Proceedings volume IV, computer, communication and control technologies: I, international institute of informatics and systematics*, Orlando, FL (pp. 328–333).

Union of Concerned Scientists (2004). Weapon materials basics. http://www.ucsusa.org/nuclear_weapons_and_global_security/nuclear_terrorism/technical_issues/fissile-materials-basics.html. Accessed 26 December 2008.

US Customs House Guide (2008). U.S. customs and border protection initiates use of AS&E's Z portal vehicle screening system at Southwest border. http://blogs.customhouseguide.com/news/?p=1786. Accessed 26 December 2008.

United States Department of Transportation (2007). *America's container ports: delivering the goods* (Technical report, Research and Innovative Technology Administration, Bureau of Transportation Statistics, Washington, D.C.).

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

Whitt, W. (1983). The queueing network analyzer. The. *Bell System Technical Journal*, *62*(9), 2779–2815.

Whitt, W. (1993). Approximations for the GI/G/m queue. *Production and Operations Management*, *2*(2), 114–161.

X-5 Monte Carlo Team. (2005). MCNP—a general Monte Carlo N-particle transport code, Version 5. http://mcnp-green.lanl.gov/pdf/MCNP5_Manual_Volume_I_LA-UR-03-1987.pdf. Accessed 26 December 2008.