Ann Oper Res (2011) 187:207–228
DOI 10.1007/s10479-009-0677-2

Network deployment of radiation detectors with physics-based detection probability calculations

Nedialko B. Dimitrov · Dennis P. Michalopoulos · David P. Morton · Michael V. Nehme · Feng Pan · Elmira Popova · Erich A. Schneider · Gregory G. Thoreson

Published online: 10 December 2009

Abstract We describe a model for deploying radiation detectors on a transportation network consisting of two adversaries: a nuclear-material smuggler and an interdictor. The interdictor first installs the detectors. These installations are transparent to the smuggler, and are made under an uncertain threat scenario, which specifies the smuggler's origin and destination, the nature of the material being smuggled, the manner in which it is shielded, and the mechanism by which the smuggler selects a route. The interdictor's goal is to minimize the probability the smuggler evades detection. The performance of the detection equipment depends on the material being sensed, geometric attenuation, shielding, cargo and container type, background, time allotted for sensing and a number of other factors. Using a stochastic radiation transport code (MCNPX), we estimate detection probabilities for a specific set of such parameters, and inform the interdiction model with these estimates.

1 Introduction

The Department of Homeland Security (DHS) has been installing portal detectors in the US and these installations will continue.[1] The Second Line of Defense (SLD) program of the US Department of Energy (DOE) seeks to reduce the risk of illicit trafficking of nuclear material through international airports, seaports, and border crossings. The program's initial efforts were in Russia but have grown to include other key transit states in Eurasia.[2]

[1]General Accounting Office, Combatting Nuclear Smuggling: DHS's Decision to Procure and Deploy the Next Generation of Radiation Detection Equipment is not Supported by its Cost-Benefit Analysis, GAO-07-581T, March 2007.

[2]Department of Energy and National Nuclear Security Administration, NNSA's Second Line of Defense Program, Fact Sheet, September 2009, http://www.nnsa.energy.gov/news/2299.htm.

N.B. Dimitrov · D.P. Michalopoulos · D.P. Morton (✉) · M.V. Nehme · E. Popova · E.A. Schneider · G.G. Thoreson
The University of Texas at Austin, Austin, TX 78712, USA
e-mail: morton@mail.utexas.edu

F. Pan
Los Alamos National Laboratory, Los Alamos, NM 87545, USA

The DHS and DOE are addressing a real threat. In the early 1990s, Russia inherited roughly 600–850 metric tons of highly-enriched uranium (HEU) and plutonium, enough material to make over 50,000 explosive devices,[3] and the nuclear ambitions of rogue nations make daily news. An International Atomic Energy Agency (IAEA) database includes over 1000 incidents of trafficking of nuclear and radioactive material from 1993–2006 that have been confirmed by a member country's government.[4] Fifty-five percent of these involved nuclear material and 18 involved weapons-grade uranium or plutonium. Sometimes a smuggler's intent is difficult to discern, but according to the IAEA report, many of the thefts of material were motivated by profit and a perceived demand on the illegal market. Other smuggling attempts were apparently motivated by malicious intent. US efforts to assist the Former Soviet Union in securing nuclear material are ongoing,[5] but by themselves, insufficient. An accurate inventory of the nuclear material that existed at the beginning of the 1990s seems impossible.

SLD's first detector installation was at Moscow's Sheremetyevo International Airport in September 1998. Then Secretary of Energy Bill Richardson and his Russian counterpart Chairman Valeriy Draganov of the Russian Federation State Customs Committee dedicated the equipment's installation with a ribbon-cutting ceremony.[6] According to the DOE, such detector installations have two purposes: (i) to deter potential theft and smuggling of nuclear material and (ii) to detect and therefore prevent actual smuggling attempts.

Importantly, considerable effort is being devoted to developing more sophisticated radiation detectors. Less attention is devoted to how to best deploy these detectors over a system-wide network to deter and interdict the smuggling of nuclear material. Well-designed deployment can significantly improve system performance, and in this paper, we describe a stochastic network interdiction model for locating radiation detectors. There is considerable transparency in our strategies and would-be smugglers can access information that aids their ability to circumvent our systems, and our model can capture this issue. A key input to our interdiction model is the ability of radiation detectors to sense nuclear material. Using plastic scintillator radiation portal monitors (RPMs) as a baseline detector, we describe an efficient method for parametrically computing detection probabilities using a limited number of computations with the much slower, standard radiation transport code MCNPX (Monte Carlo N-Particle Extended) (MCNPX user's manual 2008). Our approach to efficiently computing detection probabilities builds a framework for hedging against a large number of threat scenarios, without the computational bottleneck of standard radiation transport simulations.

[3]General Accounting Office, Nuclear Nonproliferation: US Efforts to Help Other Countries Combat Nuclear Smuggling Need Strengthened Coordination and Planning, Report to the Ranking Minority Member, Subcommittee on Emerging Threats, and Capabilities, Committee on Armed Services, US Senate, GAO-02-426, May 2002.

[4]International Atomic Energy Agency, Illicit Trafficking Database, Fact Sheet: January 1993–December 2006.

[5]Department of Energy and National Nuclear Security Administration, NNSA's Second Line of Defense Program, Fact Sheet, October 2009, http://www.nnsa.energy.gov/news/2676.htm.

[6]Department of Energy Press Release: Bill Richardson, Russian Federation Dedicate "Second Line of Defense", US Nuclear Detection Technology to Help Secure Russian Borders, September 2, 1998.

2 Stochastic network interdiction problem

First, we describe our stochastic network interdiction problem (SNIP) on a general transportation network, which could involve multiple countries. Then, we specialize to the computationally simpler case of placing detectors on the border of one country. We call the latter model BiSNIP, as it is SNIP specialized to a bipartite network, as we describe in detail below. In Sect. 5 we consider two sets of single-country model instances, one involving Russia and the other the United States.

We model two adversaries, an interdictor and a smuggler, and an underlying transportation network $G(N, A)$. The smuggler starts at origin node $o \in N$ and wishes to reach destination node $d \in N$. The probability that the smuggler evades detection while traversing arc $(i, j) \in A$ is q_{ij} if (i, j) has a detector and $p_{ij} > q_{ij}$ if not. A smuggler can be caught by indigenous law enforcement without detection equipment, and so $p_{ij} < 1$. Detection events on distinct arcs are assumed to be mutually independent. The smuggler chooses an o–d path to maximize his evasion probability. With limited resources, the interdictor must select arcs on which to install detectors in order to minimize this probability.

The *threat scenario*, indexed by $\omega \in \Omega$, is unknown when detectors are installed, but is governed by a probability mass function, $p^\omega, \omega \in \Omega$, which is assumed to be known. A threat scenario specifies the origin-destination pair, (o^ω, d^ω), as well as other details about the nuclear material being smuggled and the manner in which it is shielded. (We discuss this issue in more detail in Sect. 3.) So, the probability a smuggler evades detection if a detector is installed on arc (i, j) is scenario dependent, i.e., $q_{ij} = q_{ij}^\omega$. In general, the indigenous evasion probabilities, p_{ij}, could also depend on the threat scenario. The bulk of what we present is valid when $p_{ij} = p_{ij}^\omega$, but in Sect. 4 we discuss a computationally valuable scenario-aggregation scheme that arises naturally when p_{ij} does not depend on ω. In what follows, "threat scenario ω" is often referred to as simply "smuggler ω."

The probability the smuggler traverses the network undetected is a sum of conditional evasion probabilities, each weighted by p^ω, over all threat scenarios. The timing of the interdictor's and smuggler's decisions and the realization of the threat scenario is as follows: First, the interdictor installs detectors on a subset of the network's arcs subject to a budget constraint. Then, a threat scenario is revealed and the smuggler selects an o^ω–d^ω path.

Naturally, the smuggler's path choice is crucial in determining the interdictor's detector placement choice. Models with limited information and gaming are important to consider, since they may lead to different placement solutions, but in this paper we conservatively assume the smuggler selects a path with full knowledge of the detector locations and evasion probabilities. This approach is conservative in the following sense. If solving our model yields a detector-placement plan and an associated system-wide evasion probability of, say, p^* then under that placement plan, the smuggler cannot achieve an evasion probability higher than p^* regardless of his path choice. This is true even if the smuggler decides to take an arbitrary route, takes a second-best route, or chooses a route by placing a probability distribution on the possible paths. Of course, for a particular smuggler movement model, there may exist alternate placement plans that achieve a significantly better objective function value than p^*.

Assuming transparency of detector locations is arguably reasonable. In Russia, the initial installation was accompanied by a ribbon-cutting ceremony, and subsequent installations were reported in the news. Completely sealing Russia's 12,500 miles of borders is impractical, and so, in addition to catching nuclear smugglers, the SLD program seeks to deter

would-be smugglers, e.g., who seek financial gain. Issues surrounding detector installation in the US have also been openly reported.[7]

The SNIP formulation follows:

Network and sets

$G(N, A)$	directed network with nodes N and arcs A
$FS(i)$	set of arcs leaving node i
$RS(i)$	set of arcs entering node i
$AD \subset A$	arcs on which detectors may be placed

Data

f	total budget for installing detectors
c_{ij}	cost of installing a detector on arc $(i, j) \in AD$

Random elements

$\omega \in \Omega$	sample point and sample space for threat scenarios
(o^ω, d^ω)	realization of random origin–destination pair
p^ω	probability mass function
p^ω_{ij}	evasion probability for smuggler ω on (i, j); no detector installed
$q^\omega_{ij} < p^\omega_{ij}$	evasion probability for smuggler ω on (i, j); detector installed

Interdictor's decision variables

x_{ij}	1 if a detector is installed on arc (i, j) and 0 otherwise

Smuggler's decision variables

y_{ij}	positive only if smuggler traverses (i, j) and no detector is installed
z_{ij}	positive only if smuggler traverses (i, j) and a detector is installed
y_{d^ω}	probability that smuggler ω, starting at o^ω, can reach d^ω undetected

Boundary conditions

$x_{ij}, z_{ij} \equiv 0\ (i, j) \notin AD$

Formulation

$$\min_{x \in X} \sum_{\omega \in \Omega} p^\omega h(x, \omega), \tag{1}$$

where $X = \{x : \sum_{(i,j)\in AD} c_{ij} x_{ij} \le f, x_{ij} \in \{0, 1\}, (i, j) \in AD\}$, and where

$$h(x, \omega) = \max_{y,z \ge 0} y_{d^\omega} \tag{2a}$$

$$\text{s.t.} \sum_{(o^\omega, j)\in FS(o^\omega)} (y_{o^\omega j} + z_{o^\omega j}) = 1 \tag{2b}$$

$$\sum_{(i,j)\in FS(i)} (y_{ij} + z_{ij}) = \sum_{(j,i)\in RS(i)} \left(p^\omega_{ji} y_{ji} + q^\omega_{ji} z_{ji}\right), \quad i \in N \setminus \{o^\omega, d^\omega\} \tag{2c}$$

[7] National Journal Group. Global security newswire: Daily news on nuclear, biological and chemical weapons, terrorism and related issues. http://gsn.nti.org/gsn/.

$$y_{d^\omega} = \sum_{(j,d^\omega)\in RS(d^\omega)} \left(p^\omega_{jd^\omega} y_{jd^\omega} + q^\omega_{jd^\omega} z_{jd^\omega}\right) \tag{2d}$$

$$y_{ij} \le 1 - x_{ij}, \quad (i,j) \in AD \tag{2e}$$

$$z_{ij} \le x_{ij}, \quad (i,j) \in AD. \tag{2f}$$

The value, $h(x,\omega)$, is the conditional evasion probability, given ω. The goal is to install detectors, under the budget constraint $x \in X$, to minimize the evasion probability over all threat scenarios, i.e., to minimize the objective function in (1). Each transportation link that can receive a detector is modeled as two parallel arcs. If $x_{ij} = 1$ then flow may occur only on the "detector" arc (z_{ij}), and otherwise only on the "no detector" arc (y_{ij}). Flow on arc (i,j) is multiplied by that arc's gain, either p^ω_{ij} or q^ω_{ij}. So, if P_{o^ω,d^ω} is a path from o^ω to d^ω then

$$y_{d^\omega} = \prod_{(i,j)\in P_{o^\omega,d^\omega}} \left[p^\omega_{ij}(1-x_{ij}) + q^\omega_{ij} x_{ij}\right]$$

is the probability that smuggler ω can traverse P_{o^ω,d^ω} without being detected. The smuggler's subproblem finds a path P_{o^ω,d^ω} maximizing y_{d^ω} by forcing one unit of flow out of o^ω in (2b), enforcing flow conservation at all intermediate nodes in (2c), defining the flow that reaches d^ω as y_{d^ω} in (2d) and maximizing that value in (2a). Flow is forced on the appropriate arc, and incurs the associated gain (actually, loss), by the interdictor's decision variable x_{ij} in constraints (2e) and (2f). We note that constraint (2f) can be dropped since $p^\omega_{ij} > q^\omega_{ij}$ means z_{ij} is positive only if $x_{ij} = 1$ in an optimal solution to the subproblem (2).

Installing a detector or a system of detectors, e.g., at an international customs border crossing may be more naturally thought of as installing a detector on a node, not an arc. However, we split such a border-crossing node into two nodes with an arc representing travel through the checkpoint. A more general approach would allow the model to choose among detectors of different types, e.g., based on different technologies, to be installed at each location. While model (1) does not insist all detectors be identical, it does assume that for each candidate location a specific type of detector has been specified.

When locating detectors, the interdictor knows: (i) the network topology, (ii) the indigenous detection probability on each arc, (iii) the detection probability given the presence of a detector, (iv) the resource constraint, (v) the probability distribution governing the random (o,d) pair, and (vi) the method by which the smuggler selects a path. After the threat scenario ω is revealed, the smuggler selects an o^ω–d^ω path that maximizes the evasion probability, knowing (i), (ii) and (iii) as well as the detector locations. SNIP with h defined in (2) is a bilevel stochastic mixed-integer program (MIP). In bilevel programs (e.g., Bard 1998; Ben-Ayed 1993; Ishizuka et al. 1997) each player has an objective function, and these can differ because the players' motives differ. In our case, the objective function is the same for both players, but the interdictor is minimizing that function and the smuggler is maximizing it. In Pan and Morton (2008) we describe a decomposition algorithm that allows us to effectively solve problem instances in which we install between 30 and 100 detectors on a subset of 300 candidate locations with 3,000 transportation arcs and 500 threat scenarios.

In what follows we restrict attention to a simplified version of SNIP that arises when we can only place detectors at border crossings of a single country. The most pressing problem in our initial SLD work had potential detector locations restricted to customs checkpoints leaving Russia. Nuclear material originally in Russia may no longer be there, and Russia is not the only source for such material. So, another single-country model of interest is to install detectors to minimize the probability a smuggler could enter the US with nuclear material.

The key to simplifying the formulation for the single-country case is that each o^ω–d^ω path has exactly one arc on which the smuggler could encounter a detector. Let K be the set of *checkpoint* arcs, i.e., arcs that a smuggler could traverse depending on the selected path, that could contain a detector. For each ω, and checkpoint arc k, we compute the value of the maximum-reliability path from o^ω to the tail of k, say $\gamma^\omega_{k,1}$, and the value of the maximum-reliability path from the head of k to d^ω, say $\gamma^\omega_{k,2}$. Call the product of these two probabilities $\gamma^\omega_k = \gamma^\omega_{k,1}\gamma^\omega_{k,2}$, for all pairs of ω and $k = (i, j) \in K$. Then,

$$h(x, \omega) = \max_{k \in K} \left\{ \gamma^\omega_k p^\omega_k (1 - x_k), \gamma^\omega_k q^\omega_k x_k \right\} \tag{3}$$

is the probability smuggler ω avoids detection. Linearizing (3), we obtain the MIP:

$$\min_{x,\theta} \quad \sum_{\omega \in \Omega} p^\omega \theta^\omega \tag{4a}$$

$$\text{s.t.} \quad x \in X \tag{4b}$$

$$\theta^\omega \geq \gamma^\omega_k p^\omega_k (1 - x_k), \quad k \in K, \omega \in \Omega \tag{4c}$$

$$\theta^\omega \geq \gamma^\omega_k q^\omega_k x_k, \quad k \in K, \omega \in \Omega. \tag{4d}$$

BiSNIP (4) may be viewed on a bipartite network with arcs (ω, k) linking each threat scenario with its checkpoints. Excluding the checkpoint, γ^ω_k is the smuggler's probability of traveling from o^ω to d^ω, via k, undetected. This is multiplied by q^ω_k or p^ω_k depending on whether a detector is installed at k. Variable θ^ω is the conditional probability the smuggler avoids detection, given ω, and model (4) minimizes the (unconditional) probability the smuggler avoids detection.

The set K could be indexed by ω because for a particular origin-destination pair, some checkpoints are either impossible or unreasonable, and this is how we implement BiSNIP in practice. That said, we can drop the dependence on ω without loss of generality because for an impossible ω–k combination, we can set $\gamma^\omega_k = 0$. This notational simplicity will be convenient in Sect. 4. We also note that we can replace constraint (4d) by the constraint $\theta^\omega \geq q^\omega_{\max}$, $\omega \in \Omega$, where $q^\omega_{\max} \equiv \max_{k \in K} \gamma^\omega_k q^\omega_k$, $\omega \in \Omega$. We return to this and more involved model improvements in Sect. 4.

There is a modest but growing literature on network interdiction. We do not attempt to review this here, but point only to recent stochastic models of network interdiction (Bailey et al. 2006; Cormican et al. 1998; Hemmecke et al. 2003) and references contained therein. The base model we describe in this section was first proposed in Pan et al. (2003) and more fully developed in Pan and Morton (2008) as was its bipartite special case in Morton et al. (2007). See Atkinson and Wein (2008), Behrens et al. (2007), Brown et al. (2006, 2009) for further related work.

This paper's main contributions are as follows:

1. We use physics-based detection probability (DP) calculations to populate our interdiction model, BiSNIP. An improved alarm algorithm accounts for baseline suppression (Sect. 3.1), and we investigate the system-wide benefit of this algorithm on BiSNIP instances in which detectors are placed at customs checkpoints leaving Russia (Sect. 5.1).
2. We demonstrate the tractability of using DP calculations coming from a computationally expensive radiation transport code. Using relatively few radiation transport calculations, we parameterize DPs, allowing us to obtain DPs for a large number of scenarios involving the thickness of lead shielding placed around the nuclear material (Sect. 3.2).

3. In tightening BiSNIP's formulation, we provide qualitative insight on the role DPs play under an intelligent, informed smuggler (Sect. 4.1). We provide conditions under which we can aggregate threat scenarios in BiSNIP (Sect. 4.2) and hence handle many shielding scenarios. This aggregation significantly reduces computational effort (Sect. 5.1).
4. We apply BiSNIP to a model instance under a range of budgets in which detectors are placed at motor crossings entering the United States from Mexico and Canada. We obtain important qualitative insights on the nature of optimal detector-deployment plans with respect to natural geographic groupings of checkpoints (see Sect. 5.2).

The rest of the paper is organized as follows. In Sect. 3, we consider the computation of DPs for a single detector. In Sect. 4, we tighten BiSNIP's formulation and introduce our technique for aggregating threat scenarios. In Sect. 5, we provide computational results and analysis on the model instances involving the US and Russia. Finally, we provide concluding remarks in Sect. 6.

3 Detection probabilities

We now turn to estimating a detector's detection probability, under various threat scenarios, ω. The probability that smuggler ω traveling through checkpoint $k \in K$ is detected by an installed radiation detector is $\bar{q} \equiv 1 - q_k^\omega$, where we temporarily suppress dependence of $\bar{q}$ on ω and k for notational simplicity. We focus on two issues. First, we discuss how the transporting vehicle suppresses background radiation and its effect on DP. We describe an adjustment to the alarm algorithm to account for baseline suppression. Using this as a surrogate for a higher-fidelity detector, we demonstrate the system-wide implications in Sect. 5.1. Second, we discuss lead shielding placed directly around the special nuclear material (SNM). We describe how to parametrically characterize DP using a limited number of computationally-expensive runs of a standard radiation transport code. Our parameterization can be used to hedge against a continuum of threat scenarios, without the computational bottleneck of standard radiation simulations. Our method has the advantages that it is fast, physics-based, and is a good approximation to the more detailed simulation. We show how this parameterization can be put to use in BiSNIP in Sect. 4.2. Our description here of these two issues in computing DPs is brief, and we refer to Thoreson (2009) and Thoreson and Schneider (2009) for a more thorough development.

The standard way of deriving DPs is by simulating detector response to a radioactive source via a stochastic radiation transport code such as MCNPX (MCNPX user's manual 2008). Alarm algorithms are then applied to the simulation results to translate count rate probability density functions (pdfs) into detection probabilities (Geelhood 2003). Alarm thresholds are typically selected to control the false alarm probability (FAP). The DP and FAP also depend on interrogation time, which is a primary concern for portal detectors, where many short measurements generate a spatial spectrum profile. Checkpoints with portal detectors therefore employ a two-tier screening process where the primary detector's alarm produces a secondary inspection. Our focus is on the primary detector as opposed to detailed modeling of multi-tier systems (see, e.g., Boros et al. 2009; McLay et al. 2008; Wein et al. 2006).

3.1 Baseline suppression

Vehicle profiles must be considered when estimating DP because they shield background radiation. This baseline suppression can be approximately compensated for in real-time detection. An alarm algorithm can adjust alarm thresholds by vehicle type as predicted using a

particle transport model or empirical data. Lo Presti et al. (2006) have developed an extensive library of suppressions by vehicle type and detector location.

Given the threshold, count rates, spectra and alarm algorithm, we compute DP using the method of Geelhood (2003), which applies to both active and passive systems and involves two steps. First, the pdf of the detection metric for a vehicle containing no source is established. The metric may be gross count rate with or without baseline suppression correction, photopeak count rate, or energy bin count ratio, depending on the alarm algorithm. Second, the detection metric pdf for the smuggled cargo is calculated. With these pdfs in hand, the alarm threshold is chosen so the FAP is acceptable in view of traffic flow and secondary screening capacity.

We denote DP as $\bar{q}$, FAP as $\underline{q}$, observation time as Δ [sec.], and alarm threshold count rate as t [#/sec.]. The random number of counts of gamma ray photons that reach the detector is modeled as a homogenous Poisson process. The rate of that process depends on whether a vehicle is present and whether a special nuclear material (SNM) source is in that vehicle. The background radiation process with no vehicle has expected count rate b [#/sec.]. When a vehicle is being interrogated, baseline suppression reduces the background count rate from b to $\underline{b}$ [#/sec.]. The expected count rate for the SNM source is denoted s [#/sec.]. So, for a vehicle with an SNM source, the random number of counts reaching the detector in interval Δ has expectation and variance equal to $(s + \underline{b})\Delta$. Most contemporary alarm algorithms establish thresholds based on the FAP using unsuppressed background because real-time adjustment of the threshold requires knowledge of container geometry and cargo content. Nonetheless, the increase in sensitivity from taking baseline suppression into account is considerable, and algorithms to adjust the alarm threshold in real-time are being matured. The Radiation Portal Monitor Project at Pacific Northwest National Laboratory (Robinson et al. 2008) is developing so-called "injection-study" procedures that allow count rate profiles for commercial vehicles to be drawn from a set of pre-computed profiles.

Given a specific scenario, we use MCNPX to compute, in three separate radiation transport calculations, s, b and $\underline{b}$. The detector measures a superposition of the independent suppressed background and SNM signals. The background calculation assumes typical radionuclide concentrations in soil and concrete. For our test case (described further below), the baseline-suppressed background is found to be 13% lower than the unsuppressed background, consistent with experimentally observed values (Lo Presti et al. 2006). The count rates are sufficiently large that we make the normal-distribution approximation to the Poisson. We denote by $F(\cdot, \sigma^2, \mu)$ the cumulative distribution function of a normal random variable with mean μ and variance σ^2. With this notation, we can represent the DP and FAP using the standard gross-count algorithm as

$$\bar{q} = 1 - F(t\Delta, (s + \underline{b})\Delta, (s + \underline{b})\Delta) \tag{5a}$$

$$\underline{q} = 1 - F(t\Delta, b\Delta, b\Delta). \tag{5b}$$

We use (5) as follows: We first select a value for $\underline{q}$ (e.g., 0.01) and then use (5b) to determine the alarm threshold, t. Using this value of t in (5a) we determine the DP, $\bar{q}$. Here, t is computed via (5b) using the standard background, b. To evaluate the detection probability in (5a), we use the (actual) suppressed background rate, $\underline{b}$, superpositioned with that of the SNM source.

We can refine the alarm algorithm for a given false alarm probability if we instead use the suppressed background process in both equations. That is, we use equations:

$$\bar{q} = 1 - F(t\Delta, (s + \underline{b})\Delta, (s + \underline{b})\Delta) \tag{6a}$$

$$\underline{q} = 1 - F(t\Delta, \underline{b}\Delta, \underline{b}\Delta). \tag{6b}$$

Following the same steps described above with the same initial value of $\underline{q}$, we find a smaller value of t from (6b) and hence a larger DP, $\bar{q}$, from (6a). This improved algorithm takes baseline suppression into account and would be feasible if the real-time procedures described in Robinson et al. (2008), combined with the ability to obtain data regarding container specifications and declared cargo, were put in place. In the analysis of Sect. 5.1 we discuss the implications of this improved alarm algorithm on the system-wide evasion probability.

3.2 Parameterized detection probabilities: shielding thickness

The results from a small number of Monte Carlo transport calculations can be parameterized to take into account variations in sampling time, alarm threshold, alarm algorithm, position of the source relative to the detector, and shielding thickness. Here, we restrict attention to shielding thickness. The case of a gross-count alarm algorithm is in fact the most difficult to parameterize, since in contrast to an energy-windowed algorithm, Compton scattered, Bremsstrahlung and other secondary radiation contribute appreciably to the signal. The source count rate, s, discussed above is now a function of the shielding thickness around the source. The standard method for finding s is to take a given source configuration, mass and shape and perform a radiation transport calculation for a specific shielding material and thickness.

Using the error function, defined in terms of the standard normal distribution function as $\operatorname{erf}(x) = 2F(\sqrt{2}x, 1, 0) - 1$, we can re-express (6) as

$$\bar{q} = \frac{1}{2} - \frac{1}{2}\operatorname{erf}\left(\frac{t - \underline{b} - s}{\sqrt{2(\underline{b} + s)}}\right) \tag{7a}$$

$$t = \underline{b} + \sqrt{2\underline{b}}\,\operatorname{erf}^{-1}(1 - 2\underline{q}), \tag{7b}$$

where for simplicity we set $\Delta = 1$. For the moment, the expected source count rate, s, corresponds to a specific shielding thickness. Combining (7a) and (7b) yields the DP as a function of the source, given a FAP and suppressed background:

$$\bar{q} = \frac{1}{2} - \frac{1}{2}\operatorname{erf}\left(\frac{\sqrt{2\underline{b}}\,\operatorname{erf}^{-1}(1 - 2\underline{q}) - s}{\sqrt{2\underline{b} + 2s}}\right). \tag{8}$$

Repeating this derivation beginning with (5), i.e., assuming the alarm threshold is computed ignoring baseline suppression, leads to the decreased detection probability of

$$\bar{q} = \frac{1}{2} - \frac{1}{2}\operatorname{erf}\left(\frac{b + \sqrt{2b}\,\operatorname{erf}^{-1}(1 - 2\underline{q}) - \underline{b} - s}{\sqrt{2\underline{b} + 2s}}\right). \tag{9}$$

Instead of performing a computationally-expensive transport calculation for each shielding thickness, we parametrize source count rate s as a function of shielding thickness τ [cm] via

$$s(\tau) = \alpha_1 e^{-\beta_1 \tau} + \alpha_2 e^{-\beta_2 \tau}, \tag{10}$$

where $\alpha_1 + \alpha_2$ is the count rate in the absence of shielding, and β_1 [1/cm] and β_2 [1/cm] are attenuation coefficients. For a narrowly energy-windowed detection algorithm, we could

have $\alpha_2 = 0$ and β_1 would be the true attenuation coefficient of the material at that energy, since any scattering and absorption events would remove the radiation from the window being considered. Otherwise, the additional exponential term in (10) helps capture different attenuation coefficients at different photon energy levels. Under this parameterization, we can run relatively few transport calculations to obtain data $(s(\tau), \tau)$, and perform a nonlinear regression to estimate the four parameters $\alpha_1, \alpha_2, \beta_1$ and β_2. With these parameters in hand, we now parameterize DP in terms of shielding thickness τ by substituting $s = s(\tau)$ from (10) into (8) and (9) for the two alarm algorithms described in Sect. 3.1.

We now turn to an example calculation, involving 8 kg of HEU smuggled in a standard 53-foot truck-trailer. The detector consists of two polyvinyl toluene (PVT) panels with dimensions 3.8 cm × 36 cm × 173 cm. The shielding is a lead shell whose thickness we vary. The source spectrum is generated using the RadSrc code package (Gronberg 2007). Average US terrestrial background and an interrogation time of $\Delta = 1$ second are assumed (consisting of a sum of counts from ten 0.1 second interrogation intervals) with a vehicle speed of 2.2 mph. We compute detection probabilities using a false alarm probability of $\underline{q} = 0.01$.

We run a series of MCNPX simulations at various shielding thickness of a concentric lead sphere. The data collected are $(s(\tau), \tau)$ for shielding thicknesses 0.1–10.0 cm, and are plotted in Fig. 1a. We estimate α_1, β_1, α_2, and β_2 as 2.75×10^6 [#], 15.2 [1/cm], 6.08×10^4 [#], and 1.10 [1/cm] respectively. Figure 1a also plots the corresponding equation (10). Visually, the fit appears to be reasonable and the adjusted R^2 from fitting is 0.99. Figure 1b then plots the detection probability $\bar{q}$ versus shielding thickness as calculated with baseline suppression correction (8), and in its absence (9).

Of course, the shielding thickness a smuggler would use is unknown to us when we must plan our detector deployment. The optimal system-wide deployment of detectors depends on the distribution of threat scenarios we choose to hedge against. The parameterization method we have described allows us to consider many more, even a continuum, of scenarios by providing physics-based calculations for rapidly computing detection probabilities.

4 Improving the BiSNIP formulation

Our initial attempts to solve the BiSNIP model (4) using a branch-and-bound solution method indicated that BiSNIP's linear programming relaxation can produce very weak lower bounds. In this section, we begin by describing a straightforward procedure to tighten the formulation. Then, we describe how certain threat scenarios can be aggregated and still obtain an equivalent optimization model.

4.1 Tightening BiSNIP

The following proposition tightens constraints (4c) and effectively eliminates constraints (4d) in the BiSNIP model.

Proposition 1 *Consider the BiSNIP model* (4), *let* $q^\omega_{\max} \equiv \max_{k \in K} \gamma^\omega_k q^\omega_k$, *and assume* $0 \le q^\omega_k \le p^\omega_k \le 1$ *and* $0 \le \gamma^\omega_k \le 1$ *for all* $k \in K$, $\omega \in \Omega$. *Then the inequalities*

$$\theta^\omega \ge \gamma^\omega_k p^\omega_k - (\gamma^\omega_k p^\omega_k - q^\omega_{\max}) x_k, \quad k \in K,\ \omega \in \Omega \tag{11a}$$

$$\theta^\omega \ge q^\omega_{\max}, \quad \omega \in \Omega \tag{11b}$$

are valid for BiSNIP.

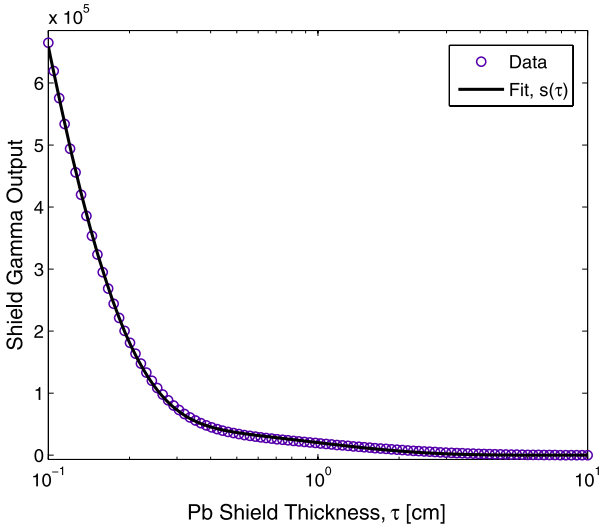

(a)

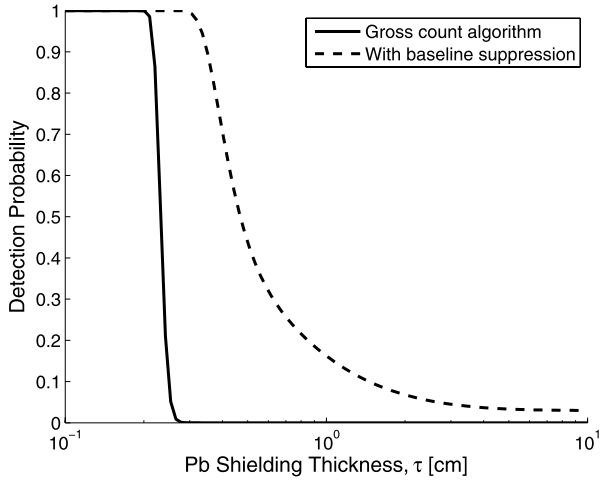

(b)

Fig. 1 Part (**a**) of the figure shows the number of counts received at the detector for an 8 kg sphere of HEU as a function of thickness of lead shielding. The *solid line* represents the fit function, $s(\tau)$, while the *circles* correspond to the data values obtained in MCNPX experiments. Part (**b**) uses $s(\tau)$ in (8) and (9) to plot detection probability versus shielding thickness under the algorithms that ignore, and account, for baseline suppression

Proof Let $k^* \in \operatorname{argmax}_{k\in K} \gamma_k^\omega q_k^\omega$ for some $\omega \in \Omega$. If $x_{k^*} = 1$, then constraint (4d) dominates (4c) and yields $\theta^\omega \geq \gamma_{k^*}^\omega q_{k^*}^\omega = q_{\max}^\omega$. And, if $x_{k^*} = 0$ then constraint (4c) dominates (4d) and yields $\theta^\omega \geq \gamma_{k^*}^\omega p_{k^*}^\omega \geq \gamma_{k^*}^\omega q_{k^*}^\omega = q_{\max}^\omega$. This proves the validity of (11b). Now for any $k \in K$, if $x_k = 1$ then (11a) becomes (11b), and if $x_k = 0$ then (11a) is equivalent to (4c). Thus (11a) is valid as well. □

We can view the right-hand side of (11b) as providing an optimistic bound, from the interdictor's perspective, on the evasion probability of smuggler ω. Then (11a) is simply a strengthened version of (4c) in which the right-hand side drops down to the lower bound $q_{\max}^\omega$ instead of zero when $x_k = 1$.

We can replace constraints (4c) and (4d) in BiSNIP with (11a) and (11b) since every constraint in the former set is dominated by some constraint in the latter. In doing so we obtain a model with half as many structural constraints and at least as strong a linear programming relaxation. Furthermore, defining $\bar{\theta}^\omega = \theta^\omega - q^\omega_{\max}$, we can transform BiSNIP into a model in which $\bar{\theta}^\omega$ has simple lower bounds of zero:

$$\min_{x,\bar{\theta}} \quad \sum_{\omega\in\Omega} p^\omega \bar{\theta}^\omega \tag{12a}$$

$$\text{s.t.} \quad x \in X \tag{12b}$$

$$\bar{\theta}^\omega \geq r^\omega_k (1 - x_k), \quad k \in K,\ \omega \in \Omega, \tag{12c}$$

where $r^\omega_k = (\gamma^\omega_k p^\omega_k - q^\omega_{\max})^+$ and where $(\cdot)^+ = \max(\cdot, 0)$. Here $\gamma^\omega_k p^\omega_k \leq q^\omega_{\max}$ implies that $r^\omega_k = 0$ and in this case, the corresponding constraint (12c) reduces to a nonnegativity constraint. This occurs when smuggler ω prefers the checkpoint with evasion probability $q^\omega_{\max}$ to that of checkpoint k. Model (12) implicitly ignores such checkpoint-smuggler pairs.

Model (12) is equivalent to BiSNIP in that both models have the same set of optimal solutions for locating the detectors, but their objective functions differ by the constant $\sum_{\omega\in\Omega} p^\omega q^\omega_{\max}$. We can view this as a transformation to a model in which the radiation detectors are perfectly reliable, i.e., model (12) has the form of model (4) with $q^\omega_k = 0$.

In the next section we discuss how smuggler ω rank orders the checkpoints in set K. Looking at model (12) smuggler ω would rank order the checkpoints k by sorting the associated values r^ω_k, $k \in K$. Note that only the largest of the $\gamma^\omega_k q^\omega_k$ values for each $\omega \in \Omega$ actually contributes to the transformed model (through $q^\omega_{\max}$) since each smuggler is guaranteed an evasion probability which is at least that high. This implies, perhaps counterintuitively, that a smuggler only considers the probability of being caught by indigenous law enforcement and not the effectiveness of detection equipment at each checkpoint when ranking checkpoints. Detection probabilities do, however, influence the number of positive r^ω_k values and consequently the number of checkpoints a smuggler could traverse with a positive probability of evading detection in the transformed model (or with a probability that exceeds that of $q^\omega_{\max}$ in the original model).

4.2 Scenario aggregation

We now focus our discussion on the transformed model (12) but suppress the "bar" notation on θ^ω for simplicity. Suppose that for some pair of smugglers $\omega, \omega' \in \Omega$ we can index the checkpoints in K, $k_1, k_2, \ldots, k_{|K|}$, such that $r^\omega_{k_1} \geq r^\omega_{k_2} \geq \cdots \geq r^\omega_{k_{|K|}}$ and $r^{\omega'}_{k_1} \geq r^{\omega'}_{k_2} \geq \cdots \geq r^{\omega'}_{k_{|K|}}$. That is to say both smugglers, while they may have different evasion probabilities at some or all checkpoints, can rank-order the checkpoints in an identical manner.

The motivation for considering the above situation arises as follows. Suppose the indigenous evasion probabilities do not depend on the threat scenario. Consider two smugglers, ω and ω', that are identical in every way, including their origin-destination pair, the mass and type of material they smuggle, etc., except that smuggler ω shields his material better than does smuggler ω'. Then, for each checkpoint the indigenous evasion probabilities associated with traveling from origin to destination via that checkpoint are identical for both smugglers, $p^\omega_k \gamma^\omega_k = p^{\omega'}_k \gamma^{\omega'}_k$ for all $k \in K$. And, the evasion probability at each checkpoint is larger for the smuggler with better shielding, $q^\omega_k > q^{\omega'}_k$ for all $k \in K$. This then results in smugglers ω and ω' ordering their checkpoints in an identical manner. As suggested above, there may be fewer positive values of r^ω_k, $k \in K$, than of $r^{\omega'}_k$, $k \in K$, but they satisfy the requisite

(inclusive) ordering condition. The same result can arise, e.g., when the two smugglers are carrying different masses of nuclear material, and it can arise for distinct origin-destination pairs, typically in close geographic proximity. It can also arise when the indigenous evasion probabilities depend on the threat scenario, as long as their ordering is identical. Specifically, since $r_k^\omega = (\gamma_k^\omega p_k^\omega - q_{\max}^\omega)^+$ two smugglers with different but identically ordered $\gamma_k^\omega p_k^\omega$ values can be aggregated.

For ω and ω' satisfying the identical-ordering assumption we have $\theta^\omega = r_{k^*}^\omega$ and $\theta^{\omega'} = r_{k^*}^{\omega'}$, where $k^* \in \operatorname{argmax}_{k\in K} r_k^\omega(1-x_k)$ and $k^* \in \operatorname{argmax}_{k\in K} r_k^{\omega'}(1-x_k)$ can be taken to be the same checkpoint. The contribution of θ^ω and $\theta^{\omega'}$ to the objective function (12a) is $p^\omega\theta^\omega + p^{\omega'}\theta^{\omega'}$. So, we can replace ω and ω' with a single scenario, say $\bar\omega$. The objective function coefficient of $\theta^{\bar\omega}$ is equal to $p^\omega + p^{\omega'}$, and the evasion probability at each checkpoint $k \in K$ is

$$\frac{p^\omega r_k^\omega + p^{\omega'} r_k^{\omega'}}{p^\omega + p^{\omega'}}$$

for scenario $\bar\omega$. Extending this idea to more scenarios yields the following proposition.

Proposition 2 *Consider model* (12) *and let* $x \in X$. *Suppose there exists a partition,* Ω^n, $n \in \mathcal{N}$, *of* Ω *such that every smuggler in a particular subset* Ω^n *orders his evasion probabilities in an identical fashion. That is, for each* $n \in \mathcal{N}$ *there exists* $k_1^n, k_2^n, \ldots, k_{|K|}^n$ *such that* $r_{k_1^n}^\omega \geq r_{k_2^n}^\omega \geq \cdots \geq r_{k_{|K|}^n}^\omega \geq 0$ *for all* $\omega \in \Omega^n$. *Let* $\theta^{\omega_n} = \max_{k\in K} r_k^{\omega_n}(1-x_k)$ *where* $r_k^{\omega_n} = \sum_{\omega\in\Omega^n} p^\omega r_k^\omega / p^{\omega_n}$ *and where* $p^{\omega_n} = \sum_{\omega\in\Omega^n} p^\omega$. *Then* $p^{\omega_n}\theta^{\omega_n} = \sum_{\omega\in\Omega^n} p^\omega\theta^\omega$.

Proof Under the ordering assumption for r_k^ω, $\omega \in \Omega$, for each $x \in X$ and $n \in \mathcal{N}$, there exists a k^* such that $r_{k^*}^\omega = \max_{k\in K} r_k^\omega(1-x_k)$, $\forall\omega \in \Omega^n$. Since the p^ω are nonnegative, the same k^* also maximizes $\sum_{\omega\in\Omega^n} p^\omega r_k^\omega(1-x_k)$. Thus,

$$\begin{aligned} p^{\omega_n}\theta^{\omega_n} &= \max_{k\in K} p^{\omega_n} r_k^{\omega_n}(1-x_k) \\ &= \max_{k\in K} \sum_{\omega\in\Omega^n} p^\omega r_k^\omega(1-x_k) \\ &= \sum_{\omega\in\Omega^n} p^\omega \max_{k\in K} r_k^\omega(1-x_k) \\ &= \sum_{\omega\in\Omega^n} p^\omega\theta^\omega. \end{aligned}$$

□

Corollary 3 *Under the hypotheses of Proposition* 2, *the following model is equivalent to model* (12):

$$\begin{aligned} \min_{x,\theta} \quad & \sum_{n\in\mathcal{N}} p^{\omega_n}\theta^{\omega_n} \\ \text{s.t.} \quad & x \in X \\ & \theta^{\omega_n} \geq r_k^{\omega_n}(1-x_k), \quad k \in K,\ n \in \mathcal{N}. \end{aligned} \tag{13}$$

In the equivalent aggregated model (13), $r_k^{\omega_n}$ and θ^{ω_n} are still conditional evasion probabilities but are now conditioned on the event $\omega \in \Omega^n$ whereas their counterparts in (12) were

conditioned on the realization of a single threat scenario. Similarly $p^{\omega_n} = P(\Omega^n)$ is the probability that a threat scenario in Ω^n is realized.

Finally, this aggregation can be performed even if Ω^n has a continuum of scenarios. This could arise if we have a continuous distribution on the thickness of lead shielding as considered at the close of Sect. 3. To make this concrete, let Ω^n correspond to all shielding scenarios for a specific origin-destination pair for a specific type and mass of SNM, and assume: the shielding thickness, τ, has pdf $\phi(\cdot)$, the indigenous detection probabilities γ_k and p_k do not depend on the threat scenario, and $q_k^\tau = q(\tau) = 1 - \bar{q}(\tau)$ is identical for all k (e.g., because identical detectors can be deployed in an identical manner in geographically distinct but otherwise identical settings). Then the continuous analog of $r_k^{\omega_n}$ from Proposition 2 is given by

$$r_k^{\omega_n} = \int \left(\gamma_k p_k - \left[\max_{k \in K} \gamma_k \right] q(u) \right)^+ \phi(u) du / P(\Omega^n).$$

When we do not account for baseline suppression, $q(\tau) = 1 - \bar{q}(\tau)$ is defined by (9) and (10), and when we do account for baseline suppression, $q(\tau)$ is defined by (8) and (10). Given this parameterized q, and τ's density ϕ, we can numerically compute this one-dimensional integral.

5 Numerical results and analysis

This section presents results using two sets of BiSNIP instances, the first involving Russia and the second the United States. Our goal is to illustrate the types of insights that BiSNIP can provide and to understand the computational benefit of the scenario aggregation procedure described in Sect. 4.2.

The Russian model involves a transportation network over 79 oblasts with 33 storage sites of SNM, which we model as being vulnerable to theft, along with eight destinations outside of Russia to where a smuggler may wish to travel. We consider five equally-likely lead shielding scenarios of the SNM, and we view these scenarios as surrogates for the smuggler's level of sophistication. This leads to $|\Omega| = 33 \cdot 8 \cdot 5 = 1320$ threat scenarios, prior to scenario aggregation. In the problem instances we consider we restrict attention to motor vehicle crossings, and consider 231 such customs checkpoints departing Russia. A subset of these checkpoints receive SAIC Exploranium AT-900 PVT detectors, i.e., the type of detector for which we presented DP calculations in Sect. 3.

In Sect. 5.1, we consider the Russian model under the two detection algorithms described in Sect. 3 to illustrate the system-wide effect of having enhanced detection capability. The first detection algorithm is the standard gross-count algorithm, which does not account for baseline suppression; see (5) and the associated discussion. The second detection algorithm accounts for baseline suppression of the transporting cargo truck-trailer; see (6). The results presented here are based on HEU, and we again use an 8 kg sphere shielded in a concentric sphere of lead. Section 5.1 also demonstrates the computational value of scenario aggregation on these same Russian BiSNIP instances.

Section 5.2 discusses results from US model instances, restricting attention to land border crossings entering the continental US from Mexico and Canada. Using a North American road network, we model 7 origins in Mexico, 7 origins in Canada and 10 destinations in the US, giving a total of $|\Omega| = 140$ threat scenarios. In both models, since all detectors are identical, we use a cardinality-constrained special case of the BiSNIP model, i.e., $c_k = 1$, for all $k \in K$, in constraint set X, and we solve the model for various budget values, f,

representing the number of border crossings equipped with detectors. Each checkpoint has an indigenous evasion probability based on its perceived vulnerability, p_k, and this varies by checkpoint, k. However, facing the same threat specified by ω, we assume detectors in distinct locations behave identically and the probability a smuggler evades detection, by the detector, is q^ω, which does not depend on k. If a detector is installed at k we assume both the indigenous detection capability and the detector technology are independently employed so that q_k^ω used in (4d) is given by $q_k^\omega = q^\omega p_k$. If ω only specifies the origin-destination pair (e.g., because distinct shielding scenarios have already been aggregated) then we can drop q's dependence on ω so that $q_k^\omega = q p_k$ for some constant q. The indigenous evasion probabilities in both models are based on a multi-attribute factor model described in detail in Witt (2003). All of the MIPs associated with both sets of BiSNIP instances were solved via the commercially-available CPLEX software (ILOG CPLEX 2008).

5.1 Russian model

Figure 2 shows the evasion probability as a function of the number of detectors we install in the Russian model, and it does so for both alarm algorithms, i.e., with, and without, accounting for baseline suppression. The former alarm algorithm yields smaller evasion probabilities and the gap between the evasion probabilities for the two systems grows with the number of detectors installed. Generally speaking, both sets of results exhibit diminishing returns with respect to decreasing the evasion probability as the number of detectors grows. That said, there are some interesting features, such as the larger drop in evasion probability as we go from 20 to 25 detectors under the baseline suppression alarm algorithm. The objective function value on the y-axis in Fig. 2 has been scaled to one if no detectors are installed, and hence instead of an evasion probability, the y-axis is the ratio of the evasion probability when installing a number of detectors to that when no detectors are installed.

We further investigate the computational value of the scenario aggregation scheme described in Sect. 4.2. We use the same two sets of Russian BiSNIP instances, i.e., one with detectors that use the baseline suppression alarm algorithm and one that does not. These models have 1320 scenarios based on five equally-likely shielding scenarios and 264 origin-destination pairs. So, aggregating the shielding scenarios alone yields a problem with 264 scenarios. However, it turns out that further reduction is possible. Threat scenarios with different origin-destination pairs still sometimes have identical rank-orderings for checkpoints due to geographic proximity and sparse road networks in parts of Russia. The result is that the 264 scenarios can be further reduced to an equivalent model with a total of 21 scenarios.

Table 1 reports computation times for problems with $f = 10, 20, \ldots, 120$ detector installations for the problem instances with, and without, the baseline suppression alarm algorithm and with, and without, the scenario aggregation procedure. The runtimes for the model instances with baseline suppression exceed those for the model instances without baseline suppression. This is not surprising, as the smaller evasion probabilities, q_k^ω, for the former model tend to lead to weaker linear programming relaxations as the binary x_k variables fractionate over a larger number of checkpoints. Entries in the table labeled as "×" indicate that the runtime exceeded 2 hours. The computation times reported in this section were on a 3.73 GHz Dell Xeon dual-processor machine with 8 GB of memory, running CPLEX version 9.1 with an absolute tolerance of 0.0001.

The results of Table 1 suggest there is substantial value in the scenario aggregation procedure. The bulk of the unaggregated models result in runtimes exceeding two hours. All the runtimes for the aggregated model instances are under 15 minutes and more than half are under one minute. We note that the original and aggregated models have the same number

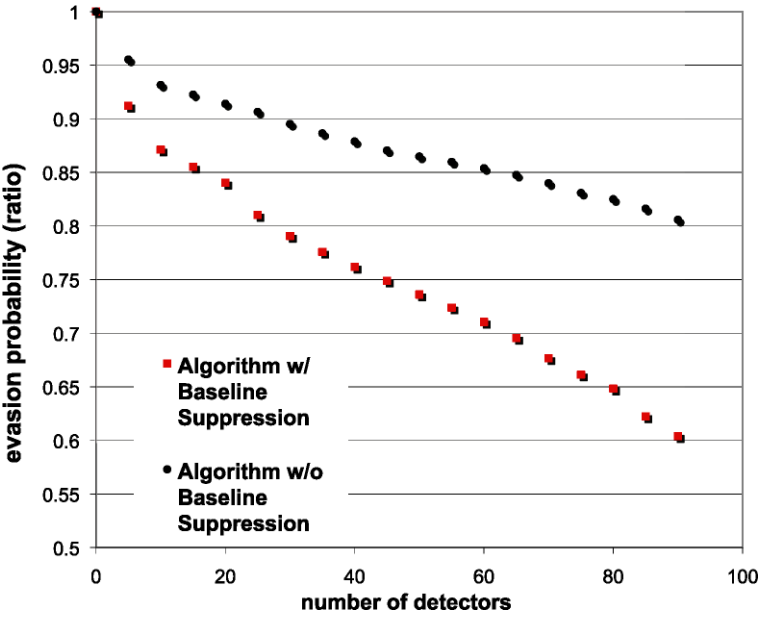

Fig. 2 The figure shows the improvement factor as a function of the number of detectors installed for two alarm algorithms, one with and the other without accounting for baseline suppression. The results come from the BiSNIP model restricted to motor crossings in Russia

of binary variables but the latter has 21 instead of 1320 continuous variables. As shown in Sect. 4.2 the original and aggregated models are mathematically equivalent, i.e., there is no loss of fidelity as a result of aggregation.

5.2 US model

Figure 3 shows the 136 motor-crossing checkpoints we consider in the US model instances. The figure also indicates four clusters of checkpoints important in results we describe below. We again solve the associated BiSNIP instances for a range of values of f, the number of detectors we can install. These hedge against 140 origin-destination threat scenarios, with half originating in Canada and the other half in Mexico. In addition to ranging f we assume the effectiveness of the detection equipment is independent of the scenario and checkpoint, that is $q_k^\omega = qp_k$ for some constant q. We create multiple model instances by ranging the value of q.

Figure 4 shows the optimal evasion probability over all threat scenarios versus the budget for four values of the detector effectiveness, q. As in Sect. 5.1, the evasion probability is reported as a fraction of that when no detectors are installed. Significant jumps in the graph occur when we are given just enough detectors to interdict an entire cluster of checkpoints. For example, we notice a large decrease in the evasion probability as the budget increases to $f = 34$ as such a budget allows us to interdict every checkpoint along the Mexican border. Smaller but still significant jumps occur when the budget increases to 11, allowing us to interdict all checkpoints in Mexico east of Big Bend (see Fig. 5), and when the budget increases to $f = 97$, allowing us to interdict all checkpoints in Mexico and all checkpoints in Canada west of Lake Huron (see Fig. 6a and 6b).

Table 1 Solve times (in seconds) for the Russian BiSNIP instances for a range of the number of checkpoints, f, which receive detectors. The second and third columns correspond to the problem in which we account for baseline suppression in the alarm algorithm, and the fourth and fifth columns correspond to the problem in which we do not account for baseline suppression. All "Original" problem instances have 1320 scenarios while all "Aggregated" instances have 21 scenarios

f	Baseline Suppression		No Baseline Suppression	
	Original	Aggregated	Original	Aggregated
10	1038.00	12.18	679.30	4.75
20	×	44.50	2665.54	13.69
30	×	28.17	1598.41	12.87
40	×	38.68	4282.93	34.03
50	×	38.68	5437.07	60.39
60	×	575.93	×	174.01
70	×	209.17	×	180.37
80	×	302.47	×	115.34
90	×	51.74	×	56.99
100	×	101.40	×	53.02
110	×	62.11	×	52.44
120	×	868.06	×	319.51

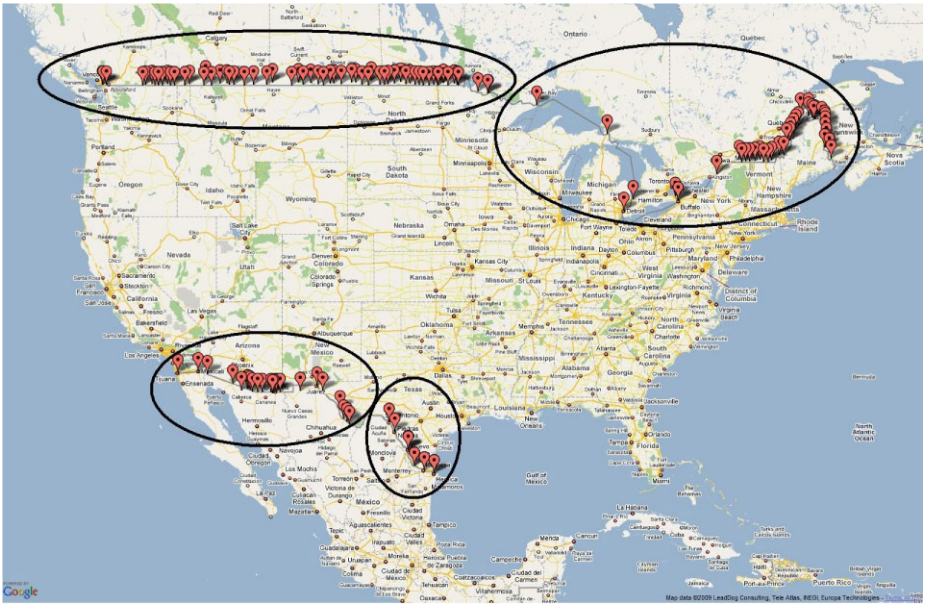

Fig. 3 The figure shows 136 motor-crossing checkpoints from Canada and Mexico into the continental United States, and groups the checkpoints into four clusters

Also noteworthy is the fact that for small values of the budget ($f < 11$), the optimal solution interdicts checkpoints along the Great Lakes (see Fig. 5a). Intuitively this is because there are more gaps between those checkpoints than there are anywhere else. Finally, we

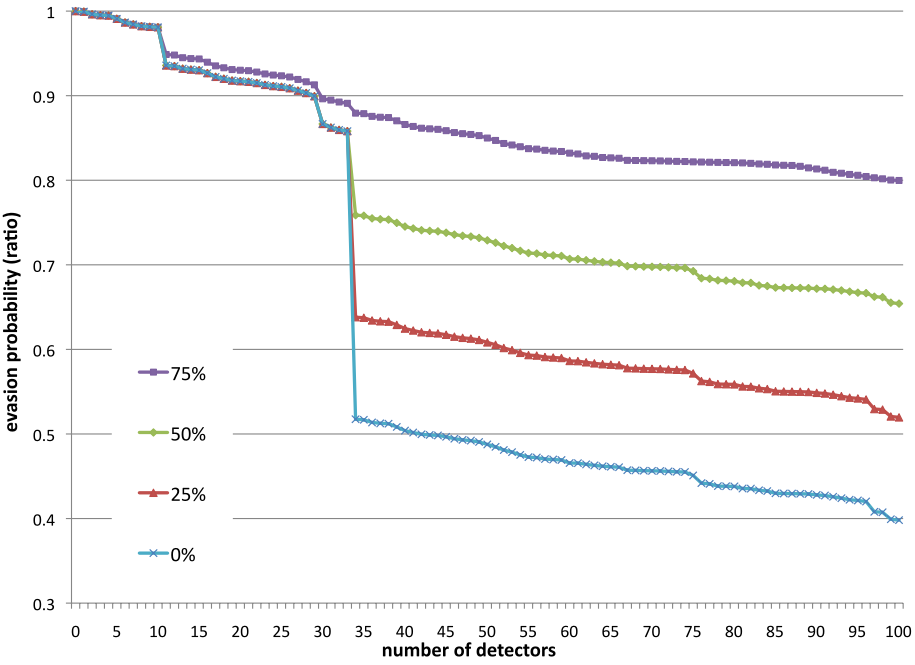

Fig. 4 The figure shows the improvement factor as a function of the number of detectors installed for the US model. The four plots correspond to different levels of effectiveness of the detectors, specifically, with $q = 0.75, 0.50, 0.25$ and 0 in $q_k^o = qp_k$

note that the solutions did vary as detectors become less effective. A notable example of this is that with more effective detectors ($q = 0, 0.25, 0.5$), there is an incentive to shift all detectors from eastern Canada to western Canada when the budget increases from 96 to 97. This was not the case with the most ineffective detectors ($q = 0.75$) as such detectors could not convince smugglers with origins in western Canada to travel around the Great Lakes to traverse a detector-free checkpoint (see Fig. 6c).

6 Summary

We have presented a stochastic network interdiction model, which can provide decision support for locating radiation detectors to thwart nuclear smuggling. The model takes as input detection probabilities, which we estimate via physics-based simulations using the industry-standard MCNPX radiation transport code.

We described an alarm algorithm that accounts for shielding of terrestrial radiation by the transporting vehicle's undercarriage and cargo, an effect known as baseline suppression. We compared the detection probability for a PVT detector using this algorithm versus one that ignores baseline suppression. We view this as two alternative systems and compare their relative performance in decreasing evasion probability on a set of model instances involving motor-crossing checkpoints leaving Russia. Our aim is not to focus on the benefits of using the baseline suppression algorithm, per se, but rather to show how our interdiction model can compare the merits of alternative systems. For example, we could compare the potential benefits of a novel detector with that of an existing system.

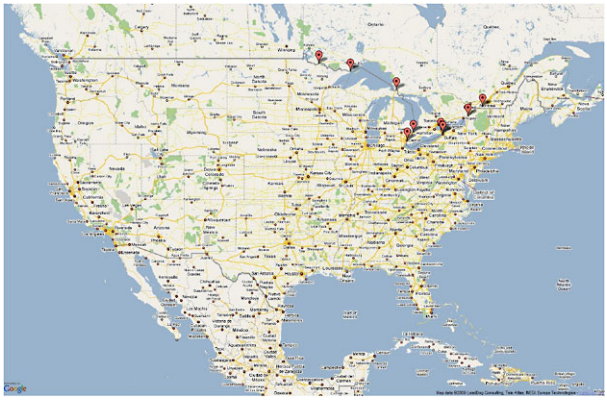

(a)

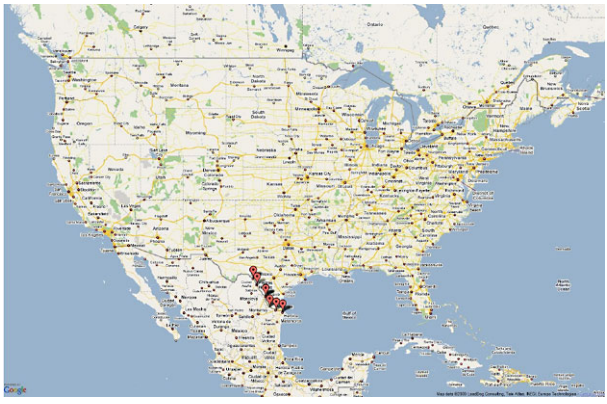

(b)

Fig. 5 Part (**a**) of the figure shows the optimal solution to the US model instance with perfectly reliable detectors, $q = 0$, and with a budget to install detectors at $f = 10$ locations. Part (**b**) of the figure is identical but for $f = 11$. Note that the full number of checkpoints is not visible in the map due to their close proximity

We proposed a method to parameterize a detector's performance, in terms of detection probability, with respect to the thickness of shielding. Using regression and MCNPX experiments we found an exponential decrease in count-rate to be a plausible model for thickness of lead shielding, even under a gross-count algorithm, so that more complicated treatment of radiation build-up is not necessary. We showed analytically how our uncertainty on the shielding thickness that a smuggler may use affects uncertainty on the detection probability. With this as motivation, we described a scenario aggregation scheme for our single-country stochastic network interdiction model. This scheme allows us to aggregate multiple smugglers provided that they rank-order checkpoints identically. Such aggregation is possible even when we consider a continuum of shielding scenarios. Our computational results suggest the benefits of aggregation can be substantial.

A second single-country model considered how to locate radiation detectors on the US-Canada and US-Mexico borders to minimize the probability a nuclear-material smuggler can travel through a motor crossing into the US. Results from these model instances suggest significant improvements in reducing the evasion probability when enough resources are allocated to equip an entire geographic cluster of border crossings, and also illustrate qualitative differences in solutions that arise depending on the effectiveness of the detection equipment.

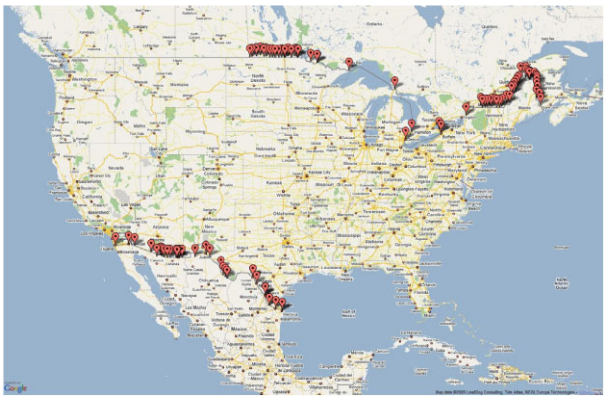

(a)

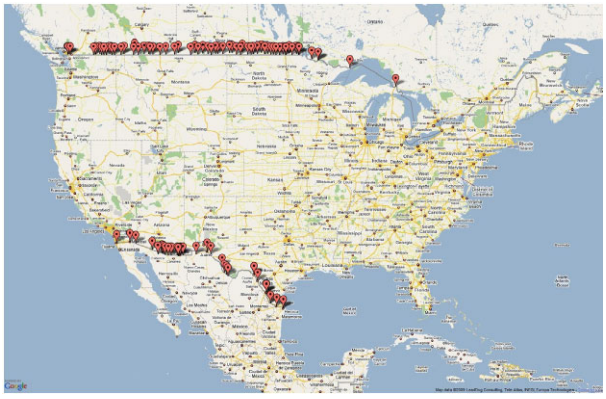

(b)

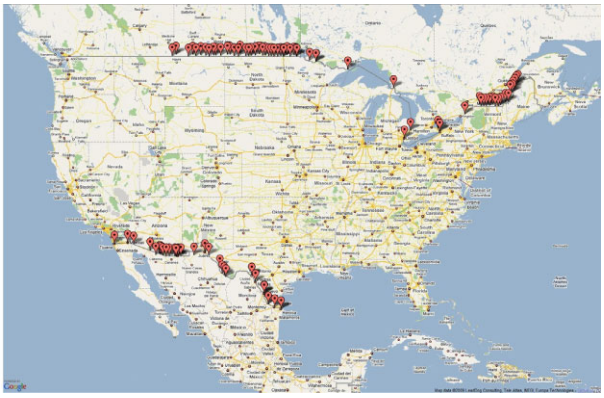

(c)

Fig. 6 Part (**a**) of the figure shows the optimal solution to the US model instance with perfectly reliable detectors, $q = 0$, and with a budget to install detectors at $f = 96$ locations. Part (**b**) of the figure is identical but for $f = 97$. Part (**c**) of the figure is for $f = 97$ and $q = 0.75$

In this paper's problem instances: We restrict attention to motor crossings; we assume a single type of SNM is smuggled in a 53-foot truck trailer; we model a single type of PVT detector; and, we restrict attention to Russia or the United States, as opposed to a more global transportation network. In future work, we will improve the model's scope and fidelity in

multiple ways. Our goal in this paper is to demonstrate the type of analyses that can be performed with our interdiction model, and indicate how high-fidelity detection probability calculations can be employed in such a model.

Acknowledgements The authors thank two anonymous referees for helpful comments that improved the paper. This work has been supported by the National Science Foundation through grants CMMI-0653916 and CMMI-0855577, the Defense Threat Reduction Agency through grant HDTRA1-08-1-0029, and the US Department of Homeland Security under Grant Award Number 2008-DN-077-ARI001-02. 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 US Department of Homeland Security.

References

Atkinson, M. P., & Wein, L. M. (2008). Spatial queueing analysis of an interdiction system to protect cities from a nuclear terrorist attack. *Operations Research*, *56*, 247–254.

Bailey, M. D., Shechter, S. M., & Schaefer, A. J. (2006). SPAR: stochastic programming with adversarial recourse. *Operations Research Letters*, *34*, 307–315.

Bard, J. F. (1998). *Practical bilevel optimization: algorithms and applications*. Boston: Kluwer Academic.

Behrens, D. A., Caulkins, J. P., Feichtinger, G., & Tragler, G. (2007). Incentive Stackelberg strategies for a dynamic game on terrorism. In S. Jørgensen, M. Quincampoix, & T. L. Vincent (Eds.), *Advances in dynamic game theory* (pp. 459–486). Boston: Birkhäuser.

Ben-Ayed, O. (1993). Bi-level linear programming. *Computers and Operations Research*, *20*, 485–501.

Boros, E., Fedzhora, L., Kantor, P. B., Saeger, K. J., & Stroud, P. (2009). Large scale linear programming model for finding optimal container inspection strategies. *Naval Research Logistics*, *56*, 404–420.

Brown, G. G., Carlyle, W. M., Harney, R., Skroch, E., & Wood, R. K. (2006). Anatomy of a project to produce a first nuclear weapon. *Science and Global Security*, *14*, 163–182.

Brown, G. G., Carlyle, W. M., Harney, R., Skroch, E., & Wood, R. K. (2009). Interdicting a nuclear-weapons project. *Operations Research*, *57*, 866–877.

Cormican, K., Morton, D. P., & Wood, R. K. (1998). Stochastic network interdiction. *Operations Research*, *46*, 184–197.

Geelhood, B. *Evaluation of the use of energy thresholds to enhance detection sensitivity using data from Detroit's Fort Street cargo facility* (Technical Report PNNL-14282 TM-054). Pacific Northwest National Laboratory, Richland, Washington, 2003.

Gronberg, J., Hiller, L., Gosnell, T., & Wright, D. (2007). Calculating gamma-ray signatures from aged mixtures of heavy nuclides. In *IEEE nuclear science symposium conference record, NSS '07* (Vol. 2, pp. 1138–1142), Honolulu, HI, 2007.

Hemmecke, R., Schultz, R., & Woodruff, D. L. (2003). Interdicting stochastic networks with binary interdiction effort. In D. L. Woodruff (Ed.), *Network interdiction and stochastic integer programming* Boston: Kluwer Academic.

ILOG CPLEX 9.0 user's manual, 2008.

Ishizuka, Y., Shimizu, K., & Bard, J. F. (1997). *Nondifferentiable and two-level programming*. Boston: Kluwer Academic.

Lo Presti, C., Weier, D., Kouzes, R., & Schweppe, J. (2006). Baseline suppression of vehicle portal monitor gamma count profiles: a characterization study. *Nuclear Instruments and Methods in Physics Research A*, *562*, 281–297.

McLay, L. A., Lloyd, J. D., & Niman, E. (2008). *Interdicting nuclear material on cargo containers using knapsack problem and Bayesian probability models* (Technical report). Department of Statistical Sciences & Operations Research, Virginia Commonwealth University.

MCNPX user's manual v2.6.0 (Los Alamos National Laboratory Report LA-CP-07-1473), April 2008.

Morton, D. P., Pan, F., & Saeger, K. J. (2007). Models for nuclear smuggling interdiction. *IIE Transactions on Operations Engineering*, *38*, 3–14.

Pan, F., & Morton, D. P. (2008). Minimizing a stochastic maximum-reliability path. *Networks*, *52*, 111–119.

Pan, F., Charlton, W., & Morton, D. P. (2003). Interdicting smuggled nuclear material. In D. L. Woodruff (Ed.), *Network interdiction and stochastic integer programming* (pp. 1–20). Boston: Kluwer Academic.

Robinson, S. M., Siciliano, E. R., & Schweppe, J. E. (2008). Source injection distribution functions for alarm algorithm testing. *Journal of Radioanalytical and Nuclear Chemistry*, *276*, 447–453.

Thoreson, G. G. (2009). *A framework for efficient detection probability computation in smuggled nuclear material interdiction*. Nuclear and Radiation Engineering, The University of Texas at Austin, M.S. Thesis.

Thoreson, G. G., & Schneider, E. A. (2009). Efficient calculation of detection probabilities. *Nuclear Instruments and Methods in Physics Research A*, under revision.

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*, 1377–1393.

Witt, K. M. (2003). *Development of a probabilistic network model to simulate the smuggling of nuclear materials*. Nuclear and Radiation Engineering, The University of Texas at Austin, M.S. Thesis.