Interdiction Models and Applications

Nedialko B. Dimitrov
Operations Research Department
Naval Postgraduate School

David P. Morton
Graduate Program in Operations Research
The University of Texas at Austin

February 25, 2013

Abstract

Through interdiction models, we infer the vulnerabilities inherent in an operational system. This chapter presents four applications of interdiction modeling: (i) to delay an adversary's development of a first nuclear weapon; (ii) to understand vulnerabilities in an electric power system; (iii) to locate sensors in a municipal water network; and (iv) to secure a border against a nuclear smuggler. In each case, we detail and interpret the mathematical model, and characterize insights gained from solving instances of the model. We point to special structures that sometimes arise in interdiction models and the associated implications for analyses. From these examples, themes emerge on how one should model, and defend against, an intelligent adversary.

This chapter describes how to assess the vulnerabilities of operational systems by using interdiction models. We do so in the context of four applications from the literature: delaying an adversary's development of a first nuclear weapon; understanding vulnerabilities in an electric power system; locating sensors to rapidly detect an illicit contaminant injected in a municipal water system; and locating radiation sensors to detect a nuclear smuggler. The key steps in this approach involve answering the following questions: (1) How is the system operated? and (2) What are the vulnerabilities of that system? Operations research has a rich history of developing mathematical models to answer question (1). Key to our approach is that we must be able to answer question (1) when any subset of the system's components has been interdicted.

We may be the operators of the system of interest, or an adversary may operate the system or be operating within a system we own. The former case arises when the system involves critical infrastructure, such as an electric power system or municipal water system. In such situations, a third question arises: (3) How can we invest to make the system more resilient? The latter case arises, for example, when an adversary is managing a project or attempting to transport illicit material across our transportation network.

In answering question (1), we assume operation of the system optimally adapts after interdiction of a subset of system components. Here, interpret the term *interdiction* liberally. It can mean an action that removes or degrades one more more system components, e.g., damaging a generator, substation, or transmission line in an electric power system, or it can mean an action that delays completion of a task in a project. Interdiction can also mean detecting illicit operations on, or threats to, a system that we own. Interdiction models identify a set of system components to interdict, subject to resource limits, so that system performance is optimally degraded. The system components identified indicate the vulnerabilities of the system, answering question (2). Again, central to the analysis is the recognition that system operation will optimally adapt, post-interdiction, to the residual system.

The sections that follow develop four applications of interdiction modeling. We motivate each application, describe a mathematical model, discuss important modeling choices, discuss computa-

tional tractability, and describe insights from analysis utilizing the model. In Section 1, we discuss delaying the development of a nuclear weapon; in Section 2, we discuss identifying vulnerabilities in the nation's power grid; in Section 3, we discuss detecting deliberate contamination of our drinking water supply; and in Section 4, we discuss securing our nation's borders against illicit smuggling of nuclear material.

1 Delaying an Adversary's Nuclear Weapons Project

Preventing a nation from covertly developing a first nuclear weapon is an international priority. More countries are pursuing civilian nuclear energy, and that growth may continue given concerns both with volatility in fossil fuel supplies and global warming. There is apprehension that, with more states having civilian nuclear power programs, some states might pursue a clandestine enrichment and reprocessing program for the purpose of developing a nuclear weapon [41]. Once the international community detects an illicit program, the tools available to stop or delay proliferation include diplomatic actions, economic embargoes, embargoes of key technologies, poaching of key personnel, sabotage, and military strikes. How should the potential effectiveness of such options be evaluated?

In a pair of papers, Harney et al. [30] first build a detailed operational model of how a "proliferator" would manage the complex project of building a first nuclear weapon, or rather, a small batch of such weapons. Then, Brown et al. [16] formulate a model on top of that operational model in which an "interdictor" selects a resource-constrained set of tasks to interdict so as to maximally delay completion of the nuclear weapons project. Reed [60] describes a lower-fidelity model but one with a similar notion of interdiction. We summarize this line of work in this section.

1.1 Project Management for a First Nuclear Weapon

The natural modeling framework for representing how the proliferator would manage the project of building a first nuclear weapon is that of the program evaluation review technique (PERT)/critical path method (CPM). O'Brien [52] discusses the origins of CPM and PERT, with the former beginning at Du Pont in 1956, and the U.S. Navy developing the latter in 1958 in conjunction with the Polaris missile program. These tools have since been employed pervasively in industry and government [66].

In its simplest form a PERT network models a collection of tasks represented by nodes that have specified durations and precedence relationships represented by directed arcs that indicate prerequisites. The length of the longest path in this network, called the critical path, indicates the minimum time required to complete the project. This time is achieved if all tasks on the critical path experience no delay, start as soon as possible, and tasks that are off the critical path are not sufficiently delayed.

Over the last several decades, PERT/CPM methods have evolved to capture the features needed to schedule and manage large, complex projects. Figure 1 depicts a simple PERT network with one such improvement: The "decision node" (node D in the figure) points to three alternative means to accomplishing a task. Brown et al. model three alternative technologies available to the proliferator to enrich uranium (gas centrifuges, gaseous diffusion, and aerodynamic enrichment) as well as these further enhancements:

1. In addition to the standard finish-to-start precedence relationship, start-to-start, finish-to-finish, and start-to-finish relationships are included. So in Figure 1, task C might have a

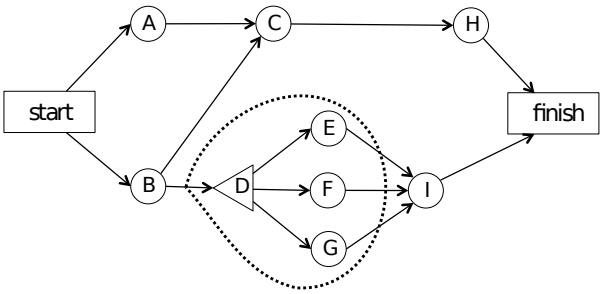

Figure 1: An example PERT network. The nodes, A through I, denote tasks that have a duration. Arcs denote precedence relationships. For example, to start task C both tasks A and B must finish. The triangular node, D, and its successors, E through G, denote a decision task. Only one of the three tasks E, F, or G must be completed to begin task I, and in order to finish the project.

start-to-start precedence relationship with task B with the addition of a 10-week lead time, meaning the earliest task C can start is 10 weeks after the start of task B.

2. In addition to consuming time, tasks consume resources including energy, raw materials, and three types of labor: scientific, skilled, and unskilled. Consumption of these resources, in turn, consumes a monetary budget. We assume that the proliferator's *capability*, i.e., his level of each of these resources, is known. To understand the sensitivity of the results to the assumed capability, we run several analyses, varying the capability assumptions.

3. The proliferator can expedite, or "crash," tasks, subject to resource limits. A task has a nominal duration, $\bar{d}$, which the proliferator can decrease to $\underline{d}$. Crashing is assumed to be linear so that c_r units of resource r are consumed per week, say, that the task is expedited.

To simplify presentation of the proliferator's operational model, we choose to neglect some of the key fidelity introduced in Brown et al. [16]. We do not include start-to-start, finish-to-finish, and start-to-finish precedence relationships. We further neglect a combinatorial aspect of the proliferator's problem in which he must make decisions as to alternative means of accomplishing tasks, as indicated in the decision node D of Figure 1. The simplified proliferator's operational model can be formulated as follows.

Sets

$i, j \in N$	nodes representing tasks
$(i, j) \in A$	precedence relationships; task i must finish before j can start
$r \in R^P$	proliferator resources

Data [units]

$\bar{d}_i$	nominal duration of task i [weeks]
$\underline{d}_i$	duration of task i if it is maximally expedited [weeks]
c_{ir}	unit consumption of resource r for expediting task i [\$/week]
b_r	budget of resource r for expediting tasks [\$]
lag_{ij}	lag time between the completion of i and the start of j [weeks]

Proliferator's decision variables [units]

S_i	earliest start time of task i [weeks]
E_i	time by which task i is expedited [weeks]

We distinguish two special nodes in N denoted "start" and "finish" as shown in Figure 1. These artificial tasks have zero duration and consume no resources.

$$\min_{S,E} \quad S_{\text{finish}} \tag{1a}$$

$$\text{s.t.} \quad S_j - S_i \geq (\bar{d}_i - E_i) + lag_{ij},\ (i,j) \in A \tag{1b}$$

$$\sum_{i \in N} c_{ir} E_i \leq b_r,\ r \in R^P \tag{1c}$$

$$0 \leq E_i \leq \bar{d}_i - \underline{d}_i,\ i \in N \tag{1d}$$

$$S_i \geq 0,\ i \in N \tag{1e}$$

$$S_{\text{start}} \equiv 0. \tag{1f}$$

Based on a finish-to-start precedence relationship, constraint (1b) indicates that the earliest start time for task j is the sum of: the earliest start time for task i, the duration of task i and any additional lag time between the completion of i and the start of j. Constraint (1c) limits the consumption of resources, which have been allocated for expediting tasks. These resources include energy, raw materials, and different types of labor. Carrying out the tasks at their nominal durations also consumes resources, but that consumption has already been accounted in the b_r values. Constraint (1d) limits the magnitude by which each task can be expedited, and the time at which the project finishes is minimized in (1a).

1.2 Formulation and Tractability in Delaying a Project

Given model (1), under which the proliferator is presumed to operate, the interdictor seeks to maximally delay completion of the proliferator's project. The interdictor is the nation, or group of nations, determined to delay the proliferator. The interdictor is limited by a monetary budget and diplomatic constraints as well as limits on economic and environmental consequences. Logical constraints on interdiction are also easily incorporated. For example, perhaps at most one of the tasks A, C, and E can be interdicted. The interdictor's model requires the following additional constructs.

Sets

$r' \in R^I$ resources for interdiction

Data [units]

$v_{ir'}$ consumption of resource r' from interdicting task i [$]

$w_{r'}$ budget of resource r' for interdicting tasks [$]

$delay_i$ delay in completing task i from its interdiction [weeks]

Interdictor's decision variables

X_i binary variable that takes value 1 if task i is interdicted and 0 otherwise

The set of feasible interdiction plans is given by

$$\mathcal{X} = \left\{ X \,:\, \sum_{i \in N} v_{ir'} X_i \leq w_{r'}, r' \in R^I, X_i \in \{0,1\}, i \in N \right\}.$$

The set $\mathcal{X}$ can include further constraints such as the logical implications just mentioned. We let $f(X)$ denote the optimal value of model (1), except that the right-hand side of constraint (1b) is replaced by

$$(\bar{d}_i - E_i + delay_i X_i) + lag_{ij},$$

in which the duration to complete task i is modified to include $delay_i$ if task i is interdicted, i.e., if decision variable $X_i = 1$. The interdictor's optimization problem is then

$$\max_{X \in \mathcal{X}} f(X). \tag{2}$$

Subject to the constraints dictated by $\mathcal{X}$, the interdictor in model (2) seeks to maximally delay completion of the proliferator's project. The nested max-min structure of model (2), with $f(X)$ defined via the optimal value of model (1), modified to incorporate delays, means the interdictor:

1. First chooses one or more tasks to interdict; and,
2. After the interdictor's plans are revealed, the proliferator chooses a plan to best expedite his project and, in the full version of model (1) in [16], the proliferator further chooses an enrichment technology associated with a decision node in the PERT network.

There are a number of questions that one can pose, regarding the appropriateness of the model just put forward. They include:

- Will the proliferator actually use PERT/CPM in planning his project?
- What if the proliferator behaves differently?
- Can the interdictor adapt his interdiction plan over time?
- Has the proliferator already committed to some decisions?

While the tools of PERT/CPM have been well known since the 1950s, and there are widely available commercial products that ease their application, we do not know whether the proliferator will employ these off-the-shelf tools. That said, the proliferator *is* building a nuclear weapon, and so it is arguably reasonable to assume he will use such tools to manage that project. If the proliferator behaves suboptimally, he will finish the project later than what we predict, and in this sense, our prediction of the induced delay is appropriately conservative. In reality, the interdictor could adapt his interdiction plan over time, and the proliferator may commit to a partial course of action prior to when some decisions must be made by the interdictor. The two-stage model we have sketched is more computationally tractable than a richer multi-stage model. The two-stage model is also conservative in the same sense that we have just mentioned: If the model predicts a completion time induced by a set of interdiction activities then the actual completion time is at least that long.

1.3 Practical Implications, and Insights

Harney et al. [30] and Brown et al. [16] describe a PERT/CPM model instance with about 200 tasks and 600 precedence relationships for the proliferator's project. The model is similar to model (1) but with the enhancements discussed above. The proliferator has a budget of $380 million that can be allocated to the five types of resources: energy, raw materials, and scientific, skilled, and unskilled labor. Brown et al. build an interdiction model in the form of (2); they label the types of interdiction they consider as mild, non-military delays; and, they introduce a limit on the total number of tasks that can be interdicted.

If the interdictor does nothing, then the proliferator completes the project in about 260 weeks—about 5 years. The same result holds if the proliferator's budget grows from $380 million to $480 million.

Next, Brown et al. [16] assume that the proliferator plans his project, ignoring the possibility of interdiction. Knowing this, the interdictor selects two tasks to interdict. Then, post-interdiction, the proliferator does not adapt his plans. In this case, the proliferator completes his project in 356 weeks, a delay of 37%. This analysis runs counter to the notion of planning conservatively, which we sketch at the end of Section 1.2, but its value will become clear shortly.

Now suppose the interdictor selects two tasks to interdict, assuming the proliferator will adapt his plan optimally post-interdiction, and the proliferator does indeed respond optimally post-interdiction; i.e., we are in the setting of model (2). In this case, the proliferator completes his project in 348 weeks, a delay of 34% over the nominal 260 weeks. The proliferator saves only 8 weeks by reacting optimally to the interdiction of two tasks, relative to not reacting at all. This may counter our intuition that optimal adaptation by the proliferator should better insulate his project from delay. As Brown et al. indicate, this means that the interdictor has uncovered "unavoidable fragilities" in the proliferator's project.

Further analysis shows that as the number of tasks to interdict ranges from one task to four tasks, the interdictor can delay the project from 1.1 to 2.4 years when the proliferator has a budget of $380 million. This range changes to 1.1 to 2.25 years when the proliferator has an additional $100 million. When examining the tasks selected for interdiction, we see that these tasks need not be on the (original) critical path. Finally, the specific task of "cascade loading" is a task that is interdicted in all the variants that the authors consider, a fact clearly of interest to decision makers.

Since the publication of the pair of papers [16, 30] discussed above, there have been a number of related developments reported in the popular press. From January 2010 through January 2012, five Iranian nuclear scientists have been attacked, and four killed, most with a bomb magnetically attached to their cars by motorcyclists; see, e.g., [13]. The so-called Stuxnet worm targeted control software for centrifuges used to enrich uranium, and it has been reported as being the most sophisticated malware ever developed; see, e.g., [15, 35]. In September 2007, Israeli Air Force jets bombed what is reported to have been a partially developed nuclear reactor in Syria; see, e.g., [36]. According to the website WikiLeaks, in a diplomatic cable dated April 20, 2008, between the U.S. Embassy in Riyadh and Washington, the Saudi ambassador to the United States, Adel al-Jubeir, is reported as having "recalled the King's frequent exhortations to the U.S. to attack Iran and so put an end to its nuclear weapons program." al-Jubeir is reported to have said, "He told you to cut off the head of the snake." WikiLeaks documents further indicate that then Secretary of Defense Robert Gates indicated that any such strike on Iran would not eliminate their nuclear program. Rather, it would only delay their pursuit of a nuclear weapon "by one to three years."

2 Vulnerabilities in the Electric Power Grid

In August of 2003, a power surge blacked out parts of eight states in the northeast United States [8]. The nation's essential demand for electricity and our dependence on the electric power grid to deliver electricity has been recognized in congressional assessments [71], and presidential policy planning groups [28]. A congressional assessment dating back to 1990 states that "The bulk power system is vulnerable to terrorist attacks targeted on key facilities. Major metropolitan areas and even multi-state regions could lose virtually all power following simultaneous attacks on three to eight sites..." [71]. In addition, the growth in demand for electric power is outpacing the development of new generation [28], while new sources of power such as wind and solar have highly variable

generation and are difficult to integrate into the system [24]. All of these developments highlight the need to understand the vulnerabilities in our electric power grid.

At a high level, there are two approaches to analyzing the sustained operation of an electric power grid. The first approach studies the *reliability* of the power grid against random component failures [59]. Such analysis uses data on the failure of individual components, such as wind turbines [70], to analyze the system as a whole [69]. However, it has been observed that the power grid can be quite robust to random failures, and at the same time be quite susceptible to the failure of a small number of select components [3]. This leads to the second analysis approach, to study the *vulnerability* of the power grid against worst case component failure—rooted in the idea that an adversary could select components to attack to induce maximum disruption.

In the literature, there are two main methods of analyzing power-grid vulnerability. The first method defines some intuitive measures of grid component criticality, and then ranks the grid components based on those measures [3, 18, 23, 58]. The advantages of this method are that it does not require extensive computation, and the criticality measures can be adapted to, or borrowed from, other common networks. The disadvantage of this method is that it is not based on the physical properties of the electric power grid; that is, it is not based on the actual performance of the system. In particular, often these methods may not be validated with, or derived from, power-flow models for electricity distribution. Furthermore, as we discuss shortly with a specific example, the concept of ranking components based on criticality is flawed because criticality is a property of sets of power-grid components. A single component may not be critical on its own, but in a set of two components it may be highly critical.

The second approach is to use interdiction models for optimal or near-optimal interdiction of power-flow as in a pair of papers by Salmerón, Wood, and Baldick [63, 65]. The benefit of this approach is that it is indeed based on the physical properties of a given electric grid. An initial drawback of the approach is that it required algorithmic and computational tools that were not available when these researchers began studying the problem. Finally, to make the interdiction computations tractable, a steady-state optimal power-flow model is typically assumed. Specifically, the power-flow models used in interdiction do not include cascading failures. Steady-state power-flow models are good for modeling mid-term or long-term failures of the electric grid, on the range of weeks or months, as opposed to models of cascading failures [43], which cause outages on much shorter time scales.

For the remainder of this section, we outline the use of interdiction models to assess the vulnerabilities of the electric power grid [63, 65]. We focus on the intuitive interpretation of the models, basic mathematical formulation, and results from such analysis.

2.1 Interdicting the Electric Power Grid

To assess the vulnerability of an electric power grid, we seek to identify components that are critical to the continued operation of that grid. At a basic level, the electric power grid consists of transmission lines, buses, generators, and substations. The core generation and transmission components of the electric power grid connect to local power distribution networks, which operate at a lower voltage and deliver power to consumers. Some components of the grid, such as a transmission line, are relatively easy to repair. Others, such as generators or substations, could take weeks or months to repair, depending on the nature of the damage. Each component integrates with the grid in a unique fashion, based on the specific structure of the grid in question.

A natural way to measure the criticality of power-grid components is to understand how much disruption is caused by their loss. For example, the loss of a local power distribution line may cause a power outage to a dozen houses for a few hours, while the loss of a substation or generator may

cause a power outage to a small city for weeks. We capture this intuitive measure of criticality through the *load shedding*, or the total demand for power that is not met if the component is lost and requires repair. We can define the criticality of a set of grid components in similar fashion. The criticality of a set of components is the total load shedding if that set of components is lost and requires repair.

As is common in most interdiction models, the criticality of a set of grid components is not simply the sum of each component's criticality, but is instead determined by the structure of the grid in question and the corresponding system performance. For example, consider a town that is powered by two identical generators that each have enough capacity to satisfy the town's entire demand on their own. In addition, suppose that each generator is connected to the town's distribution network through its own transmission lines. The criticality of each generator on its own is rather low, because the other generator serves as a backup and there is no load shedding in the event of the loss of a single generator. However, the criticality of both generators as a set is very high, because if both generators are lost at once, all of the town's demand for power goes unserved.

Finding the set of k most critical components requires us to solve a complex combinatorial optimization problem. First, we require a model of how load shedding is affected by the loss of subsets of grid components. Second, we have to compute the set of k components that maximizes that load shedding. Salmerón, Wood, and Baldick [63, 65] develop such a model and the associated algorithms for computing the most critical components.

2.2 Formulation and Tractability

The sequence of papers leading to the ability to analyze vulnerabilities of realistic electric power grids provides an excellent example of the development of interdiction models over time. The sequence begins with a simplified physics-based power-flow formulation [63], and builds to higher-fidelity power-flow models and the ability to analyze grids on the order of 5000 buses, 5000 transmission lines, and 1000 transformers—the size of a large regional grid [65].

At the most basic level, a DC power-flow model is used to measure grid performance. Let $i \in I$ denote buses, $g \in G$ generators, $\ell \in L$ transmission lines, $c \in C$ consumer demand sectors, and $s \in S$ substations. Additionally, let $i \in I(s)$ denote buses at substation s, $g \in G(i)$ generation units connected to bus i, $\ell \in L_i^{\text{bus}}$ lines connected to bus i, $\ell \in L_s^{\text{sub}}$ lines connected to substation s, and $\ell \in L_{\ell'}^{\text{par}}$ lines running parallel to line ℓ'. In the model, transformers are represented by lines.

To fully specify the grid's structure, we also require a number of data parameters. Let $o(\ell)$ and $d(\ell)$ denote the origin and destination buses of line ℓ. Let $i(g)$ denote the bus for generator g, $\bar{P}_\ell^{\text{line}}$ be the transmission capacity for line ℓ, and $\bar{P}_g^{\text{gen}}$ be the maximum output of generator g. Let r_ℓ and x_ℓ be the resistance and reactance of line ℓ, giving a susceptance of $B_\ell = x_\ell/(r_\ell^2 + x_\ell^2)$. Finally, let d_{ic} be the demand for power (load) of consumer sector c at bus i and f_{ic} be the load-shedding cost for consumer sector c at bus i. We can also think of f_{ic} as specifying the relative importance of unmet demand in different consumer sectors, where shedding load at a hospital may be more costly than in another sector.

The DC power-flow model solves for the required generation from each generator (P_g^{gen}), the power-flow on each line (P_ℓ^{line}), the phase angle at bus i (θ_i), and the load shedding in consumer sector c at bus i (S_{ic}). Generating power in some generators is cheaper than others. It is also possible to introduce costs per generator and include those in the model, but we leave out that detail for simplicity. For brevity, let $\boldsymbol{P}$ denote the vector of variables $P_g^{\text{gen}}, P_\ell^{\text{line}}$, $\boldsymbol{\theta}$ denote the vector with components θ_i, and $\boldsymbol{S}$ denote that of S_{ic}. The DC power-flow can be computed using

the following linear program:

$$\min_{\boldsymbol{P},\boldsymbol{\theta},\boldsymbol{S}} \quad \sum_{i\in I}\sum_{c\in C} f_{ic}S_{ic} \tag{3a}$$

$$\text{s.t.} \quad P_\ell^{\text{line}} = B_\ell(\theta_{o(\ell)} - \theta_{d(\ell)}), \quad \ell \in L \tag{3b}$$

$$\sum_{g\in G(i)} P_g^{\text{gen}} - \sum_{\ell|o(\ell)=i} P_\ell^{\text{line}} + \sum_{\ell|d(\ell)=i} P_\ell^{\text{line}} = \sum_{c\in C}(d_{ic} - S_{ic}), \quad i \in I \tag{3c}$$

$$-\bar{P}_\ell^{\text{line}} \le P_\ell^{\text{line}} \le \bar{P}_\ell^{\text{line}}, \quad \ell \in L \tag{3d}$$

$$0 \le P_g^{\text{gen}} \le \bar{P}_g^{\text{gen}}, \quad g \in G \tag{3e}$$

$$0 \le S_{ic} \le d_{ic}, \quad i \in I, c \in C. \tag{3f}$$

The objective function of the linear program, (3a), minimizes load shedding. Constraint (3b) approximates active power flow on each line through a linear approximation involving the phase angles on the right-hand side; constraint (3c) maintains power balance at each bus; constraint (3d) maintains the transmission capacity of each line; constraint (3e) maintains the generating capacity of each generator; and, constraint (3f) enforces that load shedding cannot exceed demand.

Once the power-flow model is formulated, we place an interdiction model on top of that model to compute the k most critical components. For the interdiction model, we introduce binary variables that indicate whether each component is functional (binary variable is 0) or not (binary variable is 1). For the power-flow model (3), let the binary interdiction variables be δ_g^{gen}, $\delta_\ell^{\text{line}}$, δ_i^{bus}, and δ_s^{sub}, each indicating whether the corresponding system component is functional. Using the interdiction variables, we can compute if a transmission line ℓ is down and store the result in a binary variable d_ℓ (value of 0 for "no power," 1 for "may have power") as follows

$$d_\ell = (1-\delta_\ell^{\text{line}})(1-\delta_{o(\ell)}^{\text{bus}})(1-\delta_{d(\ell)}^{\text{bus}}) \prod_{s|\ell\in L_s^{\text{sub}}} (1-\delta_s^{\text{sub}}) \prod_{\ell'|\ell'\in L_\ell^{\text{par}}} (1-\delta_{\ell'}^{\text{line}}). \tag{4}$$

Equation (4) states that transmission line ℓ cannot have power if line ℓ itself is nonfunctional; its origin or destination buses are nonfunctional; any substation that the line connects to is nonfunctional; or any parallel line is nonfunctional. The variable d_ℓ is only for notational convenience and is not an interdiction variable itself. For brevity, let the vector of interdiction variables for all components be $\boldsymbol{\delta}$. Using the interdiction variables and the notational convenience of d_ℓ, we can compute the k most critical components by solving the following optimization problem:

$$\max_{\boldsymbol{\delta},\boldsymbol{d}} \quad \min_{\boldsymbol{P},\boldsymbol{\theta},\boldsymbol{S}} \sum_{i\in I}\sum_{c\in C} f_{ic}S_{ic} \tag{5a}$$

$$\text{s.t.} \quad \sum_{g\in G}\delta_g^{\text{gen}} + \sum_{\ell\in L}\delta_\ell^{\text{line}} + \sum_{i\in I}\delta_i^{\text{bus}} + \sum_{s\in S}\delta_s^{\text{sub}} = k \tag{5b}$$

$$d_\ell = (1-\delta_\ell^{\text{line}})(1-\delta_{o(\ell)}^{\text{bus}})(1-\delta_{d(\ell)}^{\text{bus}}) \prod_{s|\ell\in L_s^{\text{sub}}}(1-\delta_s^{\text{sub}}) \prod_{\ell'|\ell'\in L_\ell^{\text{par}}}(1-\delta_{\ell'}^{\text{line}}), \quad \ell\in L \tag{5c}$$

$$\delta_g^{\text{gen}}, \delta_\ell^{\text{line}}, \delta_i^{\text{bus}}, \delta_s^{\text{sub}}, d_\ell \in \{0,1\}, \quad g\in G, \ell\in L, i\in I, s\in S \tag{5d}$$

$$P_\ell^{\text{line}} = B_\ell(\theta_{o(\ell)} - \theta_{d(\ell)})d_\ell, \quad \ell\in L \tag{5e}$$

$$\sum_{g\in G(i)} P_g^{\text{gen}} - \sum_{\ell|o(\ell)=i} P_\ell^{\text{line}} + \sum_{\ell|d(\ell)=i} P_\ell^{\text{line}} = \sum_{c\in C}(d_{ic} - S_{ic}), \quad i\in I \tag{5f}$$

$$-\bar{P}_\ell^{\text{line}} d_\ell \le P_\ell^{\text{line}} \le \bar{P}_\ell^{\text{line}} d_\ell, \quad \ell\in L \tag{5g}$$

$$0 \le P_g^{\text{gen}} \le \bar{P}_g^{\text{gen}}(1-\delta_{i(g)}^{\text{bus}})(1-\delta_g^{\text{gen}}), \quad g\in G \tag{5h}$$

$$0 \le S_{ic} \le d_{ic}, \quad i\in I, c\in C. \tag{5i}$$

Like model (2) in Section 1, the interdiction model (5) is a bi-level program known as a Stackelberg game, with a nested "min-max." First, components are removed from the power grid. This is accomplished by binary decision variables $\boldsymbol{\delta}$ subject to the cardinality constraint (5b) and binary restrictions (5d) along with the notational convenience, $\boldsymbol{d}$, defined via constraint (5c). Second, using the remaining components, i.e., the residual power grid, variables $\boldsymbol{P}$, $\boldsymbol{\theta}$, and $\boldsymbol{S}$ compute an optimal power flow to minimize load shedding. These variables are subject to constraints (5e)-(5i), which are similar to those of model (3), with additional parameterization in $\boldsymbol{d}$. Specifically, the d_ℓ in constraints (5e) and (5g) ensure that no down transmission line can have power. Constraint (5h) similarly ensures that a disconnected or nonfunctional generator cannot generate power.

We can alter constraint (5b) to make interdicting some components more costly than others; for example, interdicting a substation may be more costly than interdicting a single transmission line. Under such an alteration, k would be replaced with a total interdiction budget, and the optimization model would find the best interdiction plan for the specified budget.

Model (5) is amenable to interpretation; however, it is not immediately tractable as written. First, it is not possible to solve the problem with standard optimization software because some of the variables in the model are attempting to maximize the objective function, while others are attempting to minimize it. Second, the constraints of the model are nonlinear. Developing effective algorithms to solve the interdiction model is central to both the practicality of interdiction modeling and the majority of research in the area.

There are a number of ways of reformulating and solving a model like (5) in a tractable fashion. The major methods to gain tractability are:

1. Use a heuristic search to find the critical components [63].
2. Linearize the products of binary variables and take the dual of the inner problem to obtain a resulting MIP with a single maximization operator [64].
3. Apply Benders' decomposition [4].
4. Rewrite the inner minimization model by forming its optimality (KKT) conditions as constraints [46]. This method also allows the inner minimization problem to have a different

objective function than the outer maximization problem [7], which is sometimes desirable.

5. Develop custom, problem-specific algorithms, which subsequently may generalize to handle other problems [65].

It is often the custom, problem-specific algorithms that lead to truly large-scale tractability of the interdiction problem—as is the case for interdicting an electric power grid.

Salmerón, Wood, and Baldick [65] develop what they call a global Benders' decomposition algorithm to solve for the most critical components of the electric power grid. The need for such an algorithm arises because the optimal value of the inner minimization is not a concave function in the interdiction variables (or, rather, on the convex hull of their domain). This issue arises frequently in interdiction models. For example, the time to complete the adversary's project, $f(\cdot)$, in the model of Section 1, is convex on the convex hull of $\mathcal{X}$, yet the interdictor seeks to maximize that function. Such a setup does not naturally lend itself to Benders' decomposition, which, by design, forms an outer linearization of the objective function of a convex program. In some cases, because of the binary nature of the interdiction variables, it is possible to reformulate the inner problem with the interdiction variables instead in the objective function, leading to maximization of a concave function [20, 45]. However, this is not easily done in the case of the inner model in (5), largely because of constraint (5e). The ability to apply the global Benders' decomposition algorithm of Salmerón et al. hinges on being able to: (a) evaluate the optimal load shedding of the inner problem given $\boldsymbol{\delta}$ and (b) form an affine majorizing function of optimal value of the inner minimization, even though it is not concave in $\boldsymbol{\delta}$. Salmerón et al. [65] show how to do so using properties of the optimal power-flow model under two empirically verified assumptions. This allows us to solve for the most critical components of large-scale instances, and yields explicit optimality gaps if the algorithm is terminated prematurely.

2.3 Practical Implications, and Insights

Practically, a number of aspects of model (5) can be altered to yield higher fidelity results. As indicated above, it may be more difficult to disable an entire substation than a single transmission line. This can be reflected by altering constraint (5b) to take into account the relative ease of disabling each component. With such a modification, one can derive realistic efficiency curves, measuring the vulnerability of the grid, in load shedding, as a function of an adversary's capability, captured by their budget for interdiction. If the vulnerability curve is relatively flat, the grid can maintain functionality as we increase the number of failed components, or rather the interdiction budget. If the vulnerability curve increases sharply, it indicates a fragile grid in which the load shedding increases sharply with a few failed components.

It is also possible to restrict the set of components that can be interdicted. As described, every component has an associated interdiction variable. However, if it is of interest to consider the vulnerability of the entire grid to the failure of components in a particular geographic region, we can restrict the interdiction model to select only components in a particular region. Another useful restriction involves restricting interdiction to a particular type of component. In this way, we can identify the most critical power generators, for example.

The ability of interdiction algorithms to scale to grids with thousands of components allows us to analyze realistic scenarios on large-scale problems. We can use the interdiction model to answer questions such as: What are the three most critical substations in California? Is there a small set of five to ten components whose failure can cause a large amount of load shedding for a long period of time? And finally, if several candidate plans for electric power grid expansion, or hardening, are proposed, which ones decrease the vulnerability of the grid most effectively?

3 Monitoring Our Drinking Water Supply

Following the attacks of September 11, 2001, the U.S. government distributed responsibility of the nation's critical infrastructure to both newly founded and existing federal agencies. The Environmental Protection Agency (EPA) is charged with leading protection of the nation's water supply [17]. At about the same time, the U.S. Government Accountability Office identified "distribution systems as among the most vulnerable physical components of a drinking water utility," placing highest priority on the need to develop new technologies to monitor and "quickly detect contaminants in treated drinking water on its way to consumers [25]." These government priorities have directed a decade-long research effort to develop and deploy early warning systems for rapid detection of contaminants in our drinking water.

A central problem in designing a warning system to detect contamination is selecting the best locations for a limited number of water-quality sensors. Historically, optimization models for selecting sensor locations explicitly contained constraints that model water flow [39]. However, modeling water flow using simulators such as EPANET can produce higher fidelity, physics-based predictions of contaminant flow [62]. Because of the availability of high-fidelity water flow simulations, optimization models switched to exploiting the results of the simulators in selecting sensor locations, instead of modeling water flow through constraints [54].

A long-running collaboration between the EPA and operations researchers at Sandia National Laboratories has led to the practical application and deployment of well-designed contamination warning systems. Key to this success was representing the sensor placement problem as a well-known operations research problem—the p-median facility location problem [10, 11]. Subsequent and significant improvements to this initial modeling step have led to the development and distribution of the TEVA-SPOT software toolkit, a set of tools to help municipal water utilities locate sensors in their water networks [12, 9, 47, 48, 72].

For the remainder of this section, we describe the key steps to locating sensors that monitor our drinking water supply, focusing on some of the optimization models available in the TEVA-SPOT software toolkit. We provide an intuitive interpretation of the models, the basic mathematical formulation, and the results from such analyses. See Hart and Murray [32] for a review of a number of different optimization models for placing sensors in water distribution systems.

3.1 Locating Sensors to Monitor Drinking Water Networks

A drinking water network can be represented as a set of nodes connected by pipes. The level of resolution of the network can vary from application to application. For example, a single node could represent a single house in some applications, while it may represent an entire neighborhood in others. The water flow throughout the network can be quite complex, and is determined by time-dependent demand patterns and operation of pumps and tanks.

Consider a contaminant injected at a single node in the network. The injected contaminant would then move through the network, following a water system's complex flow patterns. If the contaminant flows past an installed sensor, the contaminant may be detected and actions mitigating the contamination can be taken. Formulating an optimization model to locate sensors requires clarifying what makes one placement of sensors preferable to another, and this is complicated by a number of factors including a system's complex flow patterns, questions regarding where the contaminant may be injected, and the inherent stochasticity of sensor equipment and detection events.

The initial models of the sensor placement problem assume perfect sensors and simply seek to maximize *contamination detection coverage* [39]. With this objective, a node v in the network is

considered covered, if a contaminant injection at v is detected at any point in the future by some sensor. While such objectives provide a good starting point for investigation, they can produce somewhat unrealistic results. For example, consider a simple network with two nodes in a line—an upstream node and a downstream node. Suppose that we have to place a single sensor at one of the two nodes. For a contaminant injection at the upstream node, a detection coverage objective would evaluate placing the sensor in either node as equally good because both locations detect the injection. This objective function misses the fact that, in reality, a detection after contaminated water has reached many households is not as valuable as a detection before contaminated water has reached a large segment of the population.

That is why modern sensor placement formulations consider an objective function that minimizes the impact of a contamination event [11]. The *impact* of a contamination can be defined in terms of the number of people exposed to the contaminant, the key facilities exposed to the contaminant, the length of time of the exposure, or even a combination of such factors [31]. The sensor placement analysis takes as input a set of contamination scenarios, with each scenario providing a contaminant injection point, injection rate, and length of time for the injection. The analysis can seek to position sensors throughout the network to minimize the *expected impact of contamination*, taken over the provided contamination scenarios [11]. It is possible that placing sensors to minimize expected impact does not adequately detect a few contamination scenarios with severe impacts. For this reason, we may seek to minimize the *impact of the worst contamination* scenario or to minimize other risk measures, such as the value at risk or the tail-conditional expectation [72].

Computing good sensor locations in large water networks, for many contamination scenarios, using meaningful objective functions, as informed by water flow simulations, leads to a complex combinatorial optimization problem. A key insight by Berry et al. [11] shows how a variant of this problem can be reduced to the well known p-median problem. This initial step assumes perfect sensors and focuses on minimizing expected impact. Subsequent work expands on the insight to incorporate imperfect sensors and objective functions that incorporate other risk measures [9, 72].

3.2 Formulation and Tractability

A basic p-median model for computing good sensor locations can be constructed as follows. Let $\mathcal{A}$ be a set of contamination scenarios. Each scenario $a \in \mathcal{A}$ completely describes a contamination event, including details such as injection point(s), injection rate, start and stop times, type of contaminant, etc. For each scenario, a high-fidelity water flow simulator such as EPANET [62] is used to compute a time series of the impact of the scenario. Let $d_a(t)$ denote the impact of scenario a at time t after the start of the simulation.

Let V denote the set of potential sensor locations. For this initial model, we assume perfect sensors that detect the contaminant when concentrations exceed a given threshold. For a contamination scenario a and a sensor location j, let γ_{aj} denote the earliest time at which concentrations of the contaminant at location j exceed the detection threshold. Under the perfect sensor assumption, γ_{aj} is the time at which a sensor installed at j sounds an alarm for scenario a. Let $d_{aj} = d(\gamma_{aj})$ be the impact of scenario a if it is first detected by a sensor at location j. Some sensor locations may never detect the contamination scenario. For such locations, we set d_{aj} to be the total impact of the undetected contamination scenario. In reality, if a scenario is not detected by any sensor, it might be detected by another means, such as reported illnesses.

Let α_a denote the probability of encountering contamination scenario a, and suppose we are limited to installing at most p sensors. We can formulate the problem of finding sensor locations

that minimize the expected impact of contamination over all scenarios as:

$$\min_{\boldsymbol{x},\boldsymbol{s}} \quad \sum_{a\in\mathcal{A}} \alpha_a \sum_{j\in V} d_{aj} x_{aj} \tag{6a}$$

$$\text{s.t.} \quad \sum_{j\in V} x_{aj} = 1, \quad a \in \mathcal{A} \tag{6b}$$

$$x_{aj} \le s_j, \quad a \in \mathcal{A}, j \in V \tag{6c}$$

$$\sum_{j\in V} s_j \le p \tag{6d}$$

$$s_j \in \{0,1\}, \quad j \in V \tag{6e}$$

$$x_{aj} \in \{0,1\}, \quad a \in \mathcal{A}, j \in V. \tag{6f}$$

The decision variables s_j denote whether location j is chosen for sensor installation, with 1 meaning a sensor is installed and 0 meaning a sensor is not installed. The variable x_{aj} is an auxiliary variable that has value 1 if contamination scenario a is first detected by a sensor at location j, and is 0 otherwise. The objective function in (6a) computes the expected impact of contamination over all scenarios, with the inner sum computing the impact of scenario a. Constraint (6b) ensures that each scenario is first detected by exactly one sensor; constraint (6c) ensures that scenario a can only be detected by a sensor at location j if a sensor is installed at location j; constraint (6d) ensures that no more than p sensors are installed; and, constraints (6e) and (6f) ensure binary decision variables.

To gain an intuitive understanding of model (6), imagine having a set of five potential sensor locations, $V = \{1, \ldots, 5\}$, and facing a single contamination scenario, a. Suppose we are given an installation plan, for example $s_1 = s_2 = s_3 = 1$ installing sensors in the first three locations, and $s_4 = s_5 = 0$. The x_{aj} variables are simply accounting variables that help us compute the impact of scenario a under the given sensor installation plan. Let j^* be the location with an installed sensor—either 1, 2, or 3 in our example—with minimum impact d_{aj}. When the model computes values for the variables x_{aj}, because of (6a), (6b) and (6c), all x_{aj} are set to zero except x_{aj^*}. Thus, model (6) calculates the impact of the scenario under the given sensor installation plan as being equal to $\min_{j|s_j=1} d_{aj}$. This makes intuitive sense since the sensor that sounds the contamination alarm first is the installed sensor giving minimum impact to the contamination scenario. The reasoning in this example also shows that we could relax binary constraint (6f) in favor of continuous bounds between 0 and 1.

Model (6) has the form of the classic p-median facility location problem; see, e.g., [21]. Sensor locations correspond to facility locations, scenarios correspond to demand points, and impacts d_{aj} correspond to distances between demands and facilities. Recognizing this gives us access to a rich set of tractability improvements based on a large literature devoted to the p-median problem. It is not our purpose here to review such results for the p-median problem. But, we do note that reformulating model (6) to aggregate similarly performing sensor locations, dual-based methods employing Lagrangian relaxation, integer-programming based model reductions, and special-purpose heuristics have been widely studied.

Model (6) overlooks two important factors in sensor placement. First, sensors are not perfect. They can fail to sense a contaminant (a false negative), and they can alarm when there is no contaminant (a false positive). Berry et al. [9] show how to incorporate such imperfect sensors into model (6). The basic idea of the reformulation is to change the meaning of the variables x_{aj} to the probability that scenario a is first detected by a sensor at location j. These probabilities can be computed by ordering sensor locations in a temporal manner, with locations that have

an opportunity to detect the contamination first coming first in the ordering. The probability calculations lead to a nonlinear program, which can be linearized at the cost of a significant increase in the number of decision variables. Nevertheless, the linearized program is able to compute good sensor locations when the sensors are imperfect.

The second issue with model (6) is that an objective function minimizing the expected impact over all scenarios can leave some contamination scenarios with extraordinarily high impacts undetected. This is especially a problem if there is reason to believe an adversary could observe the design of our system and exploit it. Even if there is no adversary observing our designs, simply determining scenario probabilities can be extremely difficult. In either case, we may still have interest in the vulnerability of our system to a collection of posited attack scenarios. In this setting, we should minimize the impact of the worst contamination scenario. In other words, we would like to replace objective (6a) to change the formulation to

$$\begin{array}{lll} \min\limits_{\boldsymbol{x},\boldsymbol{s}} \ \max\limits_{a\in\mathcal{A}} & \displaystyle\sum_{j\in V} d_{aj}x_{aj} & \\ \text{s.t.} & \text{constraints (6b)-(6f).} & \end{array} \tag{7}$$

Model (7) has a natural interpretation: First we select sensor locations $\boldsymbol{s}$; second, an adversary, knowing our sensor locations, selects the worst contamination scenario. Such a model more naturally applies to a terrorist action than does model (6). Solving model (7) is possible through a standard linearization

$$\begin{array}{lll} \min\limits_{\boldsymbol{x},\boldsymbol{s},y} & y & \\ \text{s.t.} & \displaystyle\sum_{j\in V} d_{aj}x_{aj} \le y, \quad a\in\mathcal{A} & \\ & \text{constraints (6b)-(6f),} & \end{array}$$

that is equivalent to the p-center problem. Watson et al. [72] indicate that it is possible to have both a low expected impact over all scenarios, and a low impact for the worst scenario. While we do not detail it here, additional risk measures including value at risk and tail-conditional expectation are explored in Watson et al. [72].

3.3 Practical Implications, and Insights

Practical instances of the sensor location problem can be so large that they do not fit in the memory of a typical 32-bit workstation. One example in the literature involves a network with about 12,000 nodes, with sensor locations hedging against 39,000 contamination scenarios. A naive formulation of models (6) or (7) would require about a half billion variables [11, 47]. Even heuristics for sensor placement require on the order of 8 gigabytes of memory to solve such an instance. Through careful reformulations, and special purpose algorithms, researchers have been able to find near-optimal solutions to such instances in seconds on a standard laptop computer.

The EPA and Sandia National Laboratories have packaged these methods for computing water sensor locations for municipal water systems in a software package called TEVA-SPOT [12], which is available for free download on the Internet [38]. The package has been used to analyze the networks of at least 18 water utilities, and the results have been used to operationally deploy sensors in at least eight utilities. The mean savings from sensor deployment, in terms of reduction in the expected economic impact of a contamination incident if one were to occur, ranged from $1 billion to $33.4 billion with a median of $5.8 billion. The expected economic impact of the worst

contaminations, those in the 95th percentile, dropped by a median of $19 billion. In more than half of the utilities studied, the expected number of fatalities expected from a contamination dropped by at least 50% [47].

The long-term research efforts in water sensor placement have significantly increased the use of operations research in the water resource planning community. Interactions between government, academia, and industry, have prompted a realistic mathematical model design and resulted in theoretical, computational, and operational advances. The resulting software continues to be improved and employed by the EPA, with the goal of securing the more than 50,000 water utilities across the United States.

4 Securing a Border Against a Nuclear Smuggler

The International Atomic Energy Agency (IAEA) maintains an Illicit Trafficking Database [2] to which over 100 nation states contribute by reporting events involving illicit trafficking, or other unauthorized possession, of nuclear material and other radioactive material. From 1993 to 2011 over 2000 such incidents were reported, and about 400 incidents involved criminal activity. During this same time period, 16 cases involved weapons grade material, i.e., highly-enriched uranium (HEU) or plutonium. Some of these seizures involved kilograms of material, and some represented small samples from a larger unsecured stockpile. The IAEA reports that, when such information is available, the majority of the cases concerned traffickers seeking financial gain by attempting to sell illicit material. That said, the motives of transporters of illicit nuclear material may change as the material changes hands following its theft and moves along the "supply chain" required to form a weapon. Many of these cases involve perpetrators characterized as being amateurs, but the IAEA reports that some incidents involve organized, professional groups with a history of illicit trafficking in nuclear material. The cases involving HEU and plutonium appear to have originated in Russia or neighboring states, where material was not adequately secured after the fall of the Soviet Union.

The U.S. Department of Homeland Security's Domestic Nuclear Detection Office (DNDO) is charged with developing the Global Nuclear Detection Architecture (GNDA). This involves coordination with multiple federal agencies, including the Department of Energy (DOE), the Department of Defense, the Department of State, the Nuclear Regulatory Commission, and coordination with foreign partners. For example, the DOE's National Nuclear Security Administration (NNSA) works with foreign governments to

> deter, detect, and interdict illicit trafficking in nuclear and other radioactive materials across international borders and through the global maritime shipping system. The goal is to reduce the probability of these materials being fashioned into a weapon of mass destruction or a radiological dispersal device ("dirty bomb") to be used against the United States or its key allies and international partners [1].

There is a strong need for developing better radiation detectors that can sense material like HEU, which can be difficult to detect. At the same time, these detectors should be able to differentiate threats from naturally occurring radioactive material. There is much research in developing effective detectors of radioactive material. That said, there is also important research in how to best deploy and operate these detectors on a large-scale transportation network. Much of the GNDA deployment effort to date, both domestically and abroad, has involved NNSA and DNDO installing radiation portal monitors (RPMs) at seaports, airports, and rail and road border crossings. DNDO also equips Customs and Border Protection officers with mobile detectors and has

proposed development of additional mobile detection units, which could be deployed in a surge operation, informed by shorter time-scale intelligence. There are initiatives that seek to secure cities and to deal with difficult challenges such as detecting nuclear smuggling between authorized ports of entry, with small maritime craft, and via general aviation [26].

4.1 Locating Radiation Detectors

DNDO has indicated an effort to incorporate increased analytical rigor in its development and analysis of the GNDA [53]. There is a small but growing literature in operations research concerning rigorous analytical models for detecting nuclear material. Wein et al. [73] propose improvements to an existing spatial deployment of RPMs at a foreign port to increase effectiveness of the system without increasing congestion. Gaulker et al. [27] and Wein et al. [74] both employ queueing network models to characterize congestion in a multi-layered security system at a seaport, and they seek to optimize the inspection strategy, understanding the tradeoff between detection probability and congestion. For further work on inspection strategies at a single port, see [14, 40, 42, 67]. Atkinson et al. [5] develop a model of a radiation detection system in and around a city, wherein an adversary attempts to get as close as possible to a target in a city center before detonating a nuclear weapon. Cheng et al. [19] and Hochbaum and Fishbain [34] analyze mobile distributed detection systems, in which nuclear detectors are mounted on a fleet of many cars—e.g., taxi cabs and/or police cars.

In the remainder of this section, we review a strategic-level model that places RPM detectors at seaports, airports, and rail and road border crossings. The development of this model began as part of the DOE's Second Line of Defense Program [45, 55] and later was coupled with physics-based estimates of detection probabilities and adapted for U.S. ports of entry [22]. The model we review here is the simplest of a family of models that has been developed. The simple model addresses securing the border of a single country, for example that of Russia or the U.S.; deals with only stationary detectors; and assumes that both the interdictor and the nuclear smuggler have the same perception of the detection probabilities. Models that relax all of these assumptions have also been developed [45, 49, 56, 68].

A key aspect of our model is the transportation network used by the smuggler to move the nuclear material. A smuggler starts at some origin in the network and would like to move to some destination. The transportation network may involve multiple modes of transport; however, our assumption of securing a single country ensures that the smuggler crosses at most one border crossing on his way from the origin to the destination. We restrict attention to installing radiation detectors on the country's legitimate border crossings.

A smuggler may be detected both by indigenous law enforcement, without radiation detectors, and by detectors at border crossings. An intelligent smuggler chooses an origin-destination path to maximize the probability he evades detection, and we assume that he does so knowing the location of radiation detectors. The interdictor does not know the type of smuggler he might face, or the smuggler's origin or destination. We model this lack of complete information as a probability distribution over a range of possible threat scenarios, each specifying a possible smuggler adversary. A threat scenario specifies the smuggler's origin-destination pair; the type of material he smuggles, including its mass, isotopic composition, and geometry; the manner in which that material is shielded, for example by lead of a specified thickness; and, the fashion in which the material is transported. Each of these has further detail. For example, the manner in which it is transported can include its position in a rail car or a tractor-trailer container, whether it is a single mass or distributed in the container, and the nature of the accompanying material in the container. All these factors—and more, concerning the type of detector, the algorithm by which it alarms, background

radiation from pavement, and whether it has recently rained—contribute to the probability an RPM will detect smuggled material. Subject to resource limits, the interdictor selects sites to install detectors to minimize the system-wide evasion probability.

Following the structure of the models in the three previous sections, the timing of the interdictor's and smuggler's decisions, along with the realization of the threat scenario, is as follows: First, the interdictor installs detectors at a subset of border crossings, subject to a budget constraint. Then, a threat scenario unfolds and the smuggler selects an origin-destination path. The manner in which the smuggler chooses a path is important in determining the best placement of detectors. The model we describe is conservative in that it assumes the smuggler has full knowledge of detector locations and detection probabilities. It is possible to develop models with limited information or different strategies governing the smuggler's behavior. Solutions derived from the conservative model we present have a guaranteed level of performance against more limited adversaries; however, solutions from models from limited adversaries typically do not guarantee performance against an intelligent and informed adversary. While the model we describe specifies a smuggler origin and destination, mathematically, this includes as an important special case a smuggler who optimizes over origin, destination, or both.

4.2 Formulation and Tractability

We formulate the model just sketched using the following notation:

Set

$k \in K$	border checkpoints

Data

b	budget for installing detectors
c_k	cost of installing detector at k

Random Elements

$\omega \in \Omega$	threat scenarios
ϕ^ω	probability mass function on threat scenarios
p_k^ω	evasion probability at k, under ω, when no detector is installed
q_k^ω	evasion probability at k, under ω, when a detector is installed
γ_k^ω	evasion probability on origin-destination path through k, excluding checkpoint k

Decision Variables

x_k	binary variable indicating whether (1) or not (0) a detector is installed at k
θ^ω	evasion probability under threat scenario ω

The formulation of the one-country smuggler interdiction model is then:

$$\min_{x,\theta} \quad \sum_{\omega\in\Omega} \phi^\omega \theta^\omega \tag{9a}$$

$$\text{s.t.} \quad \sum_{k\in K} c_k x_k \le b \tag{9b}$$

$$\theta^\omega \ge \gamma_k^\omega p_k^\omega (1 - x_k), \quad k \in K, \omega \in \Omega \tag{9c}$$

$$\theta^\omega \ge \gamma_k^\omega q_k^\omega x_k, \quad k \in K, \omega \in \Omega \tag{9d}$$

$$x_k \in \{0, 1\}, \quad k \in K. \tag{9e}$$

Constraints (9b) and (9e) ensure yes-no detector installation decisions, which satisfy the budget constraint. Constraints (9c) and (9d), coupled with minimization of the objective function, define

the evasion probability for a smuggler, conditional on threat scenario ω, as

$$\theta^\omega = \max_{k \in K} \left\{ \gamma_k^\omega p_k^\omega (1 - x_k) \,,\, \gamma_k^\omega q_k^\omega x_k \right\},$$

which encodes the assumption that the smuggler chooses a border crossing to maximize his evasion probability. If the smuggler chooses checkpoint k then his evasion probability is the product of: (a) his evasion probability from his origin to the checkpoint, (b) his evasion probability from just past the checkpoint to his destination, and (c) his evasion probability through the checkpoint itself. The evasion probability (c) depends on whether a detector is installed at k, and hence is either p_k^ω if $x_k = 0$ or q_k^ω if $x_k = 1$. The product of the probabilities in (a) and (b) is γ_k^ω. The value of γ_k^ω can be precomputed by finding maximum evasion probabilities from the origin to each checkpoint, k, and from each checkpoint to the destination, using maximum-reliability-path calculations.

While there are an enormous number of factors that affect the detection probability of an RPM, under mild assumptions, we can aggregate many of these and achieve an equivalent model [22]. This significantly reduces model complexity. We can further simplify model (9) by replacing constraints (9c) and (9d) with

$$\theta^\omega \geq r_k^\omega (1 - x_k), \quad k \in K, \omega \in \Omega, \tag{10}$$

where $r_k^\omega = \max(\gamma_k^\omega p_k^\omega - q_{\max}^\omega, 0)$ and $q_{\max}^\omega = \max_{k \in K} \gamma_k^\omega q_k^\omega$. The resulting model is equivalent, with its optimal value differing from that of model (9) by the constant $\sum_{\omega \in \Omega} \phi^\omega q_{\max}^\omega$.

Still, the linear-programming relaxation of model (9) is so weak that the required computational effort to solve realistically-sized instances is prohibitive. We can gain computational traction in model (9) by observing that constraints (10) have the form of the so-called *mixing inequalities*; see Miller and Wolsey [44] and references therein. This opens two computationally-promising avenues. One is rooted in using an exponentially-sized class of valid inequalities [29, 57], which can be separated in polynomial time by solving an appropriately-defined shortest-path problem. This avenue has been pursued for model (9) [45] and for variants of model (9) where the smuggler and interdictor have differing perceptions of evasion probabilities, respectively [56, 68]. The second avenue is to use a so-called *extended formulation* for the mixing inequalities [44, 57] to tighten the formulation. This has been developed for model (9), and we describe this extended formulation next [50].

Thinking from smuggler ω's perspective, we sort the transformed evasion probabilities: $r_{k(1,\omega)}^\omega \geq r_{k(2,\omega)}^\omega \geq \cdots \geq r_{k(|K|,\omega)}^\omega$. Here, $k(i,\omega)$ denotes smuggler ω's i-th best checkpoint. We define $\Delta_{k(i,\omega)}^\omega = r_{k(i,\omega)}^\omega - r_{k(i+1,\omega)}^\omega$, $i = 1, \ldots, |K|$, as the reward the interdictor collects by forcing smuggler ω from his i-th to his $(i+1)$-st best checkpoint. And, we introduce decision variable u_k^ω, which takes value 1 if smuggler ω is forced to a checkpoint lower than k on this sorted list. With boundary conditions $r_{k(|K|+1,\omega)}^\omega = 0$ and $u_{k(0,\omega)}^\omega = 1$ we have:

$$\theta^\omega \;=\; \sum_{i=1}^{|K|} r_{k(i,\omega)}^\omega (u_{k(i-1,\omega)}^\omega - u_{k(i,\omega)}^\omega) \;=\; r_{k(1,\omega)}^\omega + \sum_{i=1}^{|K|} \underbrace{(r_{k(i+1,\omega)}^\omega - r_{k(i,\omega)}^\omega)}_{-\Delta_{k(i,\omega)}^\omega} u_{k(i,\omega)}^\omega. \tag{11}$$

Upon substituting (11) we have the following reformulation of model (9):

$$\max_{x,u} \quad \sum_{\omega\in\Omega}\sum_{k\in K} \phi^{\omega}\Delta_k^{\omega}u_k^{\omega} \tag{12a}$$

$$\text{s.t.} \quad \sum_{k\in K} c_k x_k \le b \tag{12b}$$

$$u_k^{\omega} \le x_k, \quad k\in K, \omega\in\Omega \tag{12c}$$

$$u_{k(i,\omega)}^{\omega} \le u_{k(i-1,\omega)}^{\omega}, \quad i=2,\ldots,|K|, \omega\in\Omega \tag{12d}$$

$$0 \le u_k^{\omega} \le 1, \quad k\in K, \omega\in\Omega \tag{12e}$$

$$x_k \in \{0,1\}, \quad k\in K. \tag{12f}$$

The constraints of model (12) capture those of model (9) with the addition of constraints (12c), (12d), and (12e) to define u_k^{ω}; i.e., they allow reward Δ_k^{ω} to be collected only if a detector is installed at checkpoint k and at all the checkpoints that smuggler ω ranks above k. The variable u_k^{ω} is naturally binary, given that we require x_k to be binary. Model (12) has a much tighter linear-programming relaxation than that of model (9), allowing us to solve large instances.

4.3 Practical Implications, and Insights

Rather than viewing constraint (12b) as a hard budget constraint, it typically makes sense to study the trade-off between system performance—in this case, the probability we detect a smuggler—and the cost of the associated system design. To do so, we can solve model (12) parametrically in the budget b to obtain the set of Pareto efficient solutions. If we modify model (12) by removing constraint (12b) and instead maximizing the objective function

$$\sum_{\omega\in\Omega}\sum_{k\in K} \phi^{\omega}\Delta_k^{\omega}u_k^{\omega} - \lambda\sum_{k\in K} c_k x_k,$$

where λ parametrically ranges over positive values, we can obtain a subset of Pareto efficient solutions. Specifically, we obtain those that are extreme points of the concave envelope of the efficient frontier [37, 51]. Note that when relaxing the model in this way, so that we have a soft budget constraint, model (12)'s constraint set has a dual-network structure. The constraint matrix is totally unimodular, as each structural constraint has one $+1$ and -1.

The relaxed model, with the soft budget constraint, has the form of the *selection problem* of Balinski [6] and Rhys [61]. This leads to a very special structure of the extreme point solutions of the concave envelope of the efficient frontier. In particular, these solutions are *nested* [33, 51, 75]; i.e., if $x^*(b)$ denotes checkpoints which receive detectors under budget b for one such extreme point and $x^*(b')$ for an extreme point under a larger budget, then $x^*(b) \le x^*(b')$ in the vector sense. In words, this notion of nestedness means that the optimal set of checkpoints to receive detectors at budget b is a subset of those at a larger budget b'. This has important practical implications because usually the border, or another system we seek to protect, is incrementally hardened over time as additional funds become available. It is typically impossible, or too expensive, to completely redesign the system as the budget grows. This result yields budget increments at which optimal solutions are naturally nested.

Additional, geographic structure of optimal solutions to model (12) exists as we parametrically range the budget, b [22]. In particular, as the budget grows, the checkpoints that receive detectors fall in geographic clusters. In model instances for installing detectors on the land border crossings of the contiguous U.S., four geographic clusters emerge: crossings east of Big Bend in Texas,

the remaining crossings on the U.S.-Mexico border, crossings in the Great Lakes region and the rest of the northeast, and crossings west of the Great Lakes. The reason for this structure in optimal solutions is as follows: If we are dealing with an intelligent and well-informed smuggler, then installing detectors at only a subset of nearly identical border crossings does not improve our ability to detect the smuggler. Instead, we must equip all checkpoints in a geographic cluster in order to force the smugger to select an alternate path with lower evasion probability.

5 Discussion and Conclusions

The four applications discussed above—delaying a nuclear weapons project, assessing vulnerabilities in the electric power grid, detecting contaminants in drinking water, and securing our border against nuclear smugglers—exemplify the utility and the development of interdiction modeling. Analyses using interdiction models have made important contributions at multiple levels of government. They can be used to analyze and harden our critical infrastructure systems as well as to look for vulnerabilities in an adversary's system. Standard pathways for developing interdiction models are a useful first step in analysis, often delivering key insights. In addition, thoughtful, problem-specific interdiction models and optimization methods can elevate the interdiction approach to directly applicable, large-scale settings.

Sometimes, as is the situation for some of the case studies we present, it may take a decade of research and a sequence of insights into a problem to develop the special-purpose methods required to deliver specific and timely guidance on large-scale interdiction problems. Building on initial, stylized models, such an effort can have a marked operational impact. Success can also depend on persistence in delivering the insights from analyzing interdiction models.

Having an impact in practice can further hinge on putting forward compelling arguments, perhaps even based on detailed analysis, showing that less principled approaches to interdiction can yield inferior results, to potentially devastating effect. The two foremost approaches that we categorize as being less principled involve: (a) ignoring the distinction between an intentional attack and a random disruption and (b) ignoring the underlying system. As we discuss above, there is a rich literature on assessing the reliability of a system to random component failure. However, in making a modeling error of type (a), we presume our adversaries will behave similarly. There is ample evidence that this is simply not the case, particularly when our adversaries have the will and means to become well informed as to our system's design, and defenses, and when they seek to inflict maximum damage.

Modeling errors of type (b) are all too pervasive in practice, and even in our literature. In a typical such setting, an analyst develops a measure of an individual component's value. The analyst then "scores" each of the components in the system and sorts to obtain a priority list for components that should be interdicted, or hardened against interdiction. In interdicting a maximum-flow network or in interdicting a shortest-path network, this amounts to forming a sorted list of arcs based on their capacities or lengths. This ignores the fact that a system's performance can depend in subtle, and sometimes surprising, ways on the manner in which the components interact and on key subsets of components, as opposed to individual components. That such subtleties and surprises emerge from our models with regularity is well recognized in operations research. We should not forget this *raison d'être* when seeking to understand the vulnerability of our systems to intentional attack.

Because of their utility, interdiction models have already become a standard element of educational curricula in many operations research programs. Sometimes, interdiction modeling is a part of advanced courses on optimization, and sometimes it is simply included in basic, required courses

on network modeling. Giving future operations researchers a good understanding of the principles of interdiction modeling, contrasting interdiction with less suitable approaches, and teaching the basic modeling techniques and computational tools for interdiction, ensures our ability to effectively detect vulnerabilities in the systems we build and uncover such vulnerabilities in our adversaries' systems.

Acknowledgements

The authors thank Regan Murray, Javier Samerón, Jean-Paul Watson, and Kevin Wood whose thoughtful comments improved this chapter. This work has been supported by the National Science Foundation through grants CMMI-0653916 and CMMI-0800676, 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-ARI021-05. 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

[1] National Nuclear Security Administration. Fact sheet: NNSA's Second Line of Defense Program. https://nnsa.energy.gov/mediaroom/factsheets/nnsassecondlineofdefenseprogram. Accessed on 2012-02-11.

[2] International Atomic Energy Agency. Fact sheet: IAEA Illicit Trafficking Database (ITDB). http://www-ns.iaea.org/security/itdb.asp. Accessed on 2012-02-11.

[3] R. Albert, I. Albert, and G. L. Nakarado. Structural vulnerability of the North American power grid. *Physical Review E*, 69(2):025103+, 2004.

[4] R. Alvarez. Interdicting electrical power grids. Master's thesis, Naval Postgraduate School, Monterey, CA, 2004.

[5] M. P. Atkinson, Z. Cao, and L. M. Wein. Optimal stopping analysis of a radiation detection system to protect cities from a nuclear terrorist attack. *Risk Analysis*, 28:353–371, 2008.

[6] M. L. Balinski. On a selection problem. *Management Science*, 17:230–231, 1970.

[7] J.F. Bard. *Practical Bilevel Optimization: Algorithms and Applications*. Kluwer Academic Publishers, Boston, MA, 1998.

[8] J. Barron. Power surge blacks out northeast. *The New York Times*, August 15, 2003.

[9] J. Berry, R. D. Carr, W. E. Hart, V. J. Leung, C. A. Phillips, and J. Watson. Designing contamination warning systems for municipal water networks using imperfect sensors. *Journal of Water Resources Planning and Management*, 135(4):253+, 2009.

[10] J. Berry, W. E. Hart, C. A. Phillips, and J. Uber. A general integer-programming-based framework for sensor placement in municipal water networks. In *Proceedings of ASCE/EWRI World Water and Environmental Resources Congress*, 2004.

[11] J. Berry, W. E. Hart, C. A. Phillips, J. G. Uber, and J. Watson. Sensor placement in municipal water networks with temporal integer programming models. *Journal of Water Resources Planning and Management*, 132(4):218+, 2006.

[12] J. W. Berry, E. Boman, L. A. Riesen, W. E. Hart, C. A. Phillips, J.-P. Watson, and R. Murray. User's manual: TEVA-SPOT toolkit version 2.4. Technical Report EPA/600/R-08/041, National Homeland Security Research Center, Office of Research and Development, U.S. Environmental Protection Agency, 2010.

[13] J. Borger. Who is responsible for the Iran nuclear scientists attacks? *The Guardian*, January 12, 2012.

[14] Endre Boros, L. Fedzhora, P. B. Kantor, K. J. Saeger, and P. Stroud. Large scale LP model for finding optimal container inspection strategies. *Naval Research Logistics*, 56:404–420, 2009.

[15] W. J. Broad, J. Markoff, and D. E. Sanger. Israeli test on worm called crucial in Iran nuclear delay. *The New York Times*, January 15, 2011.

[16] G. G. Brown, W. M. Carlyle, R. Harney, E. Skroch, and R. K. Wood. Interdicting a nuclear-weapons project. *Operations Research*, 57:866–877, 2009.

[17] G. W. Bush. Subject: Critical infrastructure identification, prioritization, and protection. http://www.dhs.gov/xabout/laws/gc_1214597989952.shtm, December 2003. Homeland Security Presidential Directive HSPD-7, Accessed on 2011-07-09.

[18] D. Chassin and C. Posse. Evaluating North American electric grid reliability using the Barabási–Albert network model. *Physica A: Statistical Mechanics and its Applications*, 355(2-4):667–677, 2005.

[19] J. Cheng, M. Xie, and F. Roberts. Design and deployment of a mobile sensor network for the surveillance of nuclear materials in metropolitan areas. In *Proceedings of 15th International Conference on Reliability and Quality of Design (ISSAT09)*, 2009.

[20] K. Cormican, D. P. Morton, and R. K. Wood. Stochastic network interdiction. *Operations Research*, 46:184–197, 1998.

[21] M. S. Daskin. *Network and Discrete Location: Models, Algorithms, and Applications*. John Wiley & Sons, Inc., New York, NY, 1995.

[22] N. B. Dimitrov, D. Michalopoulos, D. P. Morton, M. V. Nehme, F. Pan, E. Popova, E. A. Schneider, and G. G. Thoreson. Network deployment of radiation detectors with physics-based detection probability calculations. *Annals of Operations Research*, 187:207–228, 2011.

[23] J. Espiritu, D. Coit, and U. Prakash. Component criticality importance measures for the power industry. *Electric Power Systems Research*, 77(5-6):407–420, 2007.

[24] GE Energy for The National Renewable Energy Laboratory. *Western Wind and Solar Integration Study*. The National Renewable Energy Laboratory, 2010. Report NREL/SR-550-47434.

[25] United States General Accounting Office (GAO). Drinking water: Experts' views on how future federal funding can best be spent to improve security. Technical Report GAO-04-29, 2003. Report to the Committee on Environment and Public Works, U.S. Senate.

[26] United States General Accounting Office (GAO). Combating nuclear smuggling: DHS has developed a strategic plan for its global nuclear detection architecture, but gaps remain. Technical Report GAO-11-869T, 2011. Testimony Before the Subcommittee on Cybersecurity, Infrastructure Protection, and Security Technologies, U.S. House of Representatives.

[27] G. M. Gaukler, C. Li, R. Cannaday, S. S. Chirayath, and Y. Ding. Detecting nuclear materials smuggling: using radiography to improve container inspection policies. *Annals of Operations Research*, 187:65–87, 2011.

[28] National Energy Policy Development Group. *National Energy Policy: Report of the National Energy Policy Development Group*. U.S. Government Printing Office, 2001.

[29] O. Günlük and Y. Pochet. Mixing MIR inequalities for mixed integer programs. *Mathematical Programming*, 90:429–458, 2001.

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

[31] W. E. Hart, J. W. Berry, R. Murray, C. A. Phillips, L. A. Riesen, and J. Watson. SPOT: A sensor placement optimization toolkit for drinking water contaminant warning system design. http://cfpub.epa.gov/si/si_public_record_report.cfm?dirEntryId=166528. Accessed on 2012-03-12.

[32] W. E. Hart and R. Murray. A review of sensor placement strategies for contamination warning systems. *Journal of Water Resources Planning and Management*, 136(6):611–619, 2010.

[33] D. S. Hochbaum. Dynamic evolution of economically preferred facilities. *European Journal of Operational Research*, 193:649–659, 2009.

[34] D. S. Hochbaum and B. Fishbain. Nuclear threat detection with mobile distributed sensor networks. *Annals of Operations Research*, 187:45–638, 2011.

[35] G. Keizer. Is Stuxnet the 'best' malware ever? *Computerworld*, September 16, 2010.

[36] G. Kessler and R. Wright. Israel, U.S. shared data on suspected nuclear site. *Washington Post*, September 21, 2007.

[37] H. W. Kuhn and A. W. Tucker. Nonlinear programming. In *Proceedings of the 2nd Berkeley Symposium on Mathematical Statistics and Probability*, pages 481–492, 1951.

[38] Sandia National Laboratories. TEVA-SPOT Toolkit: A sensor placement optimization tool for water security.
https://software.sandia.gov/trac/spot. Accessed on 2012-03-17.

[39] B. H. Lee and R. A. Deininger. Optimal locations of monitoring stations in water distribution system. *Journal of Environmental Engineering*, 118(1):4+, 1992.

[40] D. Madigan, S. Mittal, and F. Roberts. Sequential decision making algorithms for port of entry inspection: Overcoming computational challenges. In *Proceedings of IEEE International Conference on Intelligence and Security Informatics (ISI-2007)*, pages 1–7, 2007.

[41] F. McGoldrick. *Limiting Transfers of Enrichment and Reprocessing Technology: Issues, Constraints, Options*. Report for Belfer Center for Science and International Affairs, Harvard Kennedy School, Cambridge, MA, May, 2011.

[42] L. A. McLay, J. D. Lloyd, and E. Niman. Interdicting nuclear material on cargo containers using knapsack problem models. *Annals of Operations Research*, 187:185–205, 2011.

[43] L. Mili, Q. Qiu, and A. G. Phadke. Risk assessment of catastrophic failures in electric power systems. *International Journal of Critical Infrastructures*, 1(1):38–63, 2004.

[44] A. J. Miller and L. A.Wolsey. Tight formulations for some simple mixed integer programs and convex objective integer programs. *Mathematical Programming*, 98:73–88, 2003.

[45] D. P. Morton, F. Pan, and K. J. Saeger. Models for nuclear smuggling interdiction. *IIE Transactions on Operations Engineering*, 38:3–14, 2007.

[46] A. L. Motto, J. M. Arroyo, and F. D. Galiana. A mixed-integer LP procedure for the analysis of electric grid security under disruptive threat. *IEEE Transactions on Power Systems*, 20(3):1357–1365, 2005.

[47] R. Murray, W. E. Hart, C. A. Phillips, J. Berry, E. G. Boman, R. D. Carr, L. A. Riesen, J. Watson, T. Haxton, J. G. Herrmann, R. Janke, G. Gray, T. Taxon, J. G. Uber, and K. M. Morley. US Environmental Protection Agency uses operations research to reduce contamination risks in drinking water. *Interfaces*, 39(1):57–68, 2009.

[48] R. Murray, T. Haxton, R. Janke, W. E. Hart, J. Berry, and C. Phillips. Sensor network design for drinking water contamination warning systems: A compendium of research results and case studies using the TEVA-SPOT software. Technical Report EPA/600/R-09/141, National Homeland Security Research Center, Office of Research and Development, U.S. Environmental Protection Agency, 2010.

[49] M. V. Nehme. *Two-Person Games for Stochastic Network Interdiction: Models, Methods, and Complexities*. PhD thesis, The University of Texas at Austin, 2009.

[50] M. V. Nehme and D. P. Morton. Tightening a network interdiction model. In *Proceedings of the IIE Research Conference*, 2009.

[51] M. V. Nehme and D. P. Morton. Efficient nested solutions of the bipartite network interdiction problem. In *Proceedings of the IIE Research Conference*, 2010.

[52] J. J. O'Brien. *Scheduling Handbook*. McGraw-Hill Book Company, 1969.

[53] Domestic Nuclear Detection Office. The last line of defense: federal, state, and local efforts to prevent nuclear and radiological terrorism within the United States. http://www.dhs.gov/ynews/testimony/20110726-stern-last-line-of-defense.shtm. DNDO Director Warren Stern, before the House Committee on Homeland Security, Subcommittee on Cybersecurity, Infrastructure Protection, and Security Technologies, Accessed on 2012-02-11.

[54] A. Ostfeld and E. Salomons. Optimal layout of early warning detection stations for water distribution systems security. *Journal of Water Resources Planning and Management*, 130(5):377+, 2004.

[55] F. Pan, W. Charlton, and D. P. Morton. Interdicting smuggled nuclear material. In D.L. Woodruff, editor, *Network Interdiction and Stochastic Integer Programming*, pages 1–20. Kluwer Academic Publishers, Boston, 2003.

[56] F. Pan and D. P. Morton. Minimizing a stochastic maximum-reliability path. *Networks*, 52:111–119, 2008.

[57] Y. Pochet and L. A. Wolsey. Polyhedra for lotsizing with Wagner-Whitin costs. *Mathematical Programming*, 67:297–323, 1994.

[58] Q. Qiang and A. Nagurney. A unified network performance measure with importance identification and the ranking of network components. *Optimization Letters*, 2(1):127–142, 2008.

[59] M. Rausand and A. Høyland. *System Reliability Theory: Models, Statistical Methods, and Applications*. John Wiley & Sons, Inc., Hoboken, NJ, second edition, 2003.

[60] B. K. Reed. Models for proliferation interdiction response analysis. Master's thesis, Naval Postgraduate School, Monterey, CA, 1994.

[61] J. M. W. Rhys. A selection problem of shared fixed costs and network flows. *Management Science*, 17:200–207, 1970.

[62] L. A. Rossman. EPANET 2: Users manual. Technical Report EPA/600/R-00/057, United States Environmental Protection Agency, 2000.

[63] J. Salmerón, K. Wood, and R. Baldick. Analysis of electric grid security under terrorist threat. *IEEE Transactions on Power Systems*, 19(2):905–912, 2004.

[64] J. Salmerón, K. Wood, and R. Baldick. Optimizing electric grid design under asymmetric threat (II). Technical Report NPS-OR-04-001, Naval Postgraduate School, 2004. Prepared for U.S. Department of Justice, Office of Justice Programs and Office of Domestic Preparedness.

[65] J. Salmerón, K. Wood, and R. Baldick. Worst-case interdiction analysis of large-scale electric power grids. *IEEE Transactions on Power Systems*, 24(1):96–104, 2009.

[66] A. Shtub, J. F. Bard, and S. Globerson. *Project Management: Processes, Methodologies, and Economics*. Prentice Hall, Upper Saddle River, NJ, 2005.

[67] P. D. Stroud and K. J. Saeger. Enumeration of increasing Boolean expressions and alternative digraph implementations for diagnostic applications. In *Proceedings Volume IV, Computer, Communication and Control Technologies*, pages 328–333, 2003.

[68] K. M. Sullivan, D. P. Morton, F. Pan, and J. C. Smith. Interdicting stochastic evasion paths with asymmetric information on bipartite networks. *Naval Research Logistics*, 2012. under revision.

[69] S. Talukdar, J. Apt, M. Ilic, L. Lave, and M. Morgan. Cascading failures: Survival versus prevention. *The Electricity Journal*, 16(9):25–31, November 2003.

[70] P. J. Tavner, J. Xiang, and F. Spinato. Reliability analysis for wind turbines. *Wind Energy*, 10(1):1–18, 2007.

[71] Office of Technology Assessment U.S. Congress. *Physical Vulnerability of Electric Systems to Natural Disasters and Sabotage*. U.S. Government Printing Office, 1990. OTA-E-453.

[72] J. Watson, R. Murray, and W. E. Hart. Formulation and optimization of robust sensor placement problems for drinking water contamination warning systems. *Journal of Infrastructure Systems*, 15(4):330+, 2009.

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

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

[75] C. J. Witzgall and P. B. Saunders. Electronic mail and the "locator's" dilemma. In R. D. Ringeisen and F. S. Roberts, editors, *Applications of Discrete Mathematics*, pages 65–84. SIAM, Philadelphia, 1988.