

Solving Defender-Attacker-Defender Models for Infrastructure Defense

David L. Alderson, Gerald G. Brown, W. Matthew Carlyle, R. Kevin Wood

Center for Infrastructure Defense, Operations Research Department, Naval Postgraduate School, Monterey, California 93943 {dlalders@nps.edu, gbrown@nps.edu, mcarlyle@nps.edu, kwood@nps.edu}

Abstract This paper (a) describes a defender-attacker-defender sequential game model (**DAD**) to plan defenses for an infrastructure system that will enhance that system's resilience against attacks by an intelligent adversary, (b) describes a realistic formulation of **DAD** for defending a transportation network, (c) develops a decomposition algorithm for solving this instance of **DAD** and others, and (d) demonstrates the solution of a small transportation-network example. A **DAD** model generally evaluates system operation through the solution of an optimization model, and the decomposition algorithm developed here requires only that this system-operation model be continuous and convex. For example, our transportation-network example incorporates a congestion model with a (convex) nonlinear objective function and linear constraints.

Keywords infrastructure defense; infrastructure protection; homeland defense; intelligent adversary; game theory; optimization; defender-attacker-defender model; trilevel game; Stackelberg game; probabilistic risk analysis; traffic equilibrium

1. Introduction

Because of recent terrorist attacks that have destroyed public and private infrastructure (e.g., the World Trade Center attacks in New York in 2001, the train bombings in Madrid in 2004, the public-transport bombings in London in 2005), and because of continuing threats, the United States and other countries have directed substantial efforts toward (a) assessing threats to critical infrastructure from attacks by an intelligent adversary, (b) developing defenses that help prevent attacks, and (c) developing defenses that enhance system resilience, that is, defenses that mitigate the damage caused by successful attacks. This paper concerns itself with items (a) and (c).

As defined by the U.S. Government [71], *critical infrastructure* consists of “systems and assets, whether physical or virtual, so vital to the United States that the incapacity or destruction of such systems and assets would have a debilitating impact on security, the national economy, national public health or safety, or any combination of those matters.” The U.S. National Strategy for Homeland Security states the infrastructure mission unambiguously: “We must now focus on the resilience of the system as a whole—an approach that centers on investments that make the system better able to absorb the impact of an event without losing the capacity to function” (Homeland Security Council [43, p. 28]). Using limited investment resources to support this mission challenges infrastructure decision-makers at all levels of government, industry, and the military. This paper shows how to model and solve such investment problems.

One technique advocated for analyzing infrastructure defenses against a deliberate adversary builds on a long tradition of risk assessment for nondeliberate threats such as natural disasters, technological failures, and accidents: “probabilistic risk assessment” (“PRA,” also

called “probabilistic risk analysis”) is a conglomeration of techniques that many organizations, including the U.S. Department of Homeland Security (DHS), are using in an attempt to improve the resilience of infrastructure to attack. (See Garrick et al. [36], Parnell et al. [58], Ezell et al. [30] for general discussions of PRA; see DHS [29] regarding PRA’s application at DHS.) In the simplest case, risk assessment amounts to scoring the risk associated with individual attack scenarios by defining $Risk = Threat \times Vulnerability \times Consequence$; more complicated cases apply more complicated functions that are represented generically as $Risk = f(Threat, Vulnerability, Consequence)$. Roughly speaking, *Threat* is the probability of a particular attack, *Vulnerability* is the probability that such an attack would be successful, and *Consequence* measures the damage incurred by a successful attack, in terms of lives lost, economic damage, etc. Subject-matter experts must be involved in assessing all of these quantities (see Willis [76], ASME [5]). Once evaluated, risk scores become the basis for prioritized investment that aims to reduce those scores (Paté-Cornell and Guikema [59], Bier [8], Willis [76], Bier et al. [9]).

PRA models require that event probabilities be defined as static inputs. For a “terrorism risk analysis” of some infrastructure system, for instance, one input might be the probability that component X of a system will be attacked, and another might be the conditional probability that component X will be damaged to a specified degree if it is attacked. (Results of attacks are stated in terms of expected consequences, e.g., expected economic losses.)

Growing evidence indicates, however, what game theorists know intuitively: static probabilities are inappropriate for modeling the behavior of an intelligent adversary (Cox [25, 26], Golany et al. [39], Brown and Cox [11, 12]). Indeed, two National Research Council studies harshly criticize DHS’s use of PRA, especially in the context of terrorism (NRC [56, 57]). Further, even if PRA could measure *Risk* correctly through static inputs, PRA offers no general, computationally viable method for allocating limited resources to minimize risk. In particular, the standard method of “spending down the prioritized list” until a budget limit is reached is unlikely to be optimal. The only way to overcome the difficulty of minimizing risk within the PRA framework would be to develop an efficient method to compute, at least implicitly, the risk to each possible *set* of vulnerable components. But no such method currently exists. In several ways, then, PRA is the wrong tool for planning infrastructure defenses against an intelligent adversary.

Game theory, in contrast to PRA, models the actions of interacting “players” and therefore offers a more appropriate framework for modeling (a) a society that wants to protect its infrastructure from attack by building defenses, (b) an adversary who is likely to see those defenses and to attack in a maximally harmful way, and (c) a society that will observe the results of any attacks and operate to the best of its reduced ability. We propose such a model here, with the goal of maximizing resilience of infrastructure, i.e., minimizing disruption, against worst-case attacks. Disruption is evaluated quantitatively.

The rest of this paper is outlined as follows. Section 2 describes the paradigm of a sequential (Stackelberg) game for planning infrastructure defense, namely, a “defender-attacker-defender model” (**DAD**); we survey the literature here, also. Apparently, the literature reports computational results for only one instance of **DAD** for a realistically modeled infrastructure system, namely, an electric power transmission grid (Salmerón et al. [65]). That paper does not fully explain its solution methods, however. Therefore, §3 describes a realistic **DAD** model for planning defense of municipal road infrastructure, and §4 develops a simple, general algorithm for solving it. Section 5 presents computational results and analysis of a small example; that section, together with the appendix, specifies all problem data. To the best of our knowledge, this work describes the first use of a nonlinear system-operation model within the **DAD** framework. Section 6 presents conclusions and suggests directions for future work.

2. The **DAD** Model

A number of researchers have proposed the use of optimization-based models to represent a “defender’s” and an “attacker’s” sequential decisions for the purpose of defending infrastructure (Brown et al. [16]; Morton et al. [55]; Scaparra and Church [66]; Salmerón et al. [65, ?]). Brown et al. [16] formulate a model of defense, attack, and operation of an infrastructure system using a three-stage, sequential game, called a *defender-attacker-defender model* (**DAD**). This model, which is a type of Stackelberg game (see von Stackelberg [73]), commonly takes this form:

$$\mathbf{DAD}: \min_{\mathbf{w} \in W} \max_{\mathbf{x} \in X(\mathbf{w})} \min_{\mathbf{y} \in Y(\mathbf{w}, \mathbf{x})} f(\mathbf{y}).$$

In the first stage of this model, the “defender” chooses infrastructure investments $\mathbf{w} \in W$; in the second stage, the “attacker” sees those investments and attacks using attack plan $\mathbf{x} \in X(\mathbf{w})$; in the third stage, the defender, as “operator” of the system, sees attacks \mathbf{x} and infrastructure investments \mathbf{w} , and operates the system by choosing activities $\mathbf{y} \in Y(\mathbf{w}, \mathbf{x})$ that minimize operating cost measured through $f(\mathbf{y})$. More details follow.

2.1. The Operator **D**

The innermost minimization of $f(\mathbf{y})$ represents the actions of the *defender-as-operator*, or simply *operator*, who chooses a set of activities $\mathbf{y} \in Y(\mathbf{w}, \mathbf{x})$ to minimize the cost of operating the system. (The defender and operator may not be the same entity, but they share the same goals.) The notation $Y(\mathbf{w}, \mathbf{x})$ indicates that activities may be affected by both defensive investments \mathbf{w} and attacks \mathbf{x} .

Cost for the operator should be construed broadly, and can cover dollar cost, lives lost, delay of travelers, and so on. Of course, “negative output” or “negative throughput” can be used here if the operator’s goal is actually to maximize output, throughput or some similar measure. The model can also be generalized, typically for solution purposes, to include an objective function of the form $f(\mathbf{w}, \mathbf{x}, \mathbf{y})$. More important is the fact that the system need not have an actual operator. For instance, as demonstrated in this paper, **D** can represent the solution of an equilibrium model of a cost- or delay-minimizing population of travelers. The keys here are that (a) a validated model represents optimal system operation, and (b) the model can be manipulated easily to reflect parameters and/or constraints that change as a function of attacks that damage or destroy components, and as a function of defensive actions that protect existing system components, add capacity to such components, or even construct new ones.

Numerous authors propose the use of abstract, surrogate models for system operation (or for evaluating the effects of attacks), and never validate their models’ predictions using “prevalidated,” industry-standard models. For instance, Albert et al. [2], Chassin and Posse [18], Lewis [49, pp. 263–284], and Wang and Rong [74] make claims about the vulnerability of an electric power grid to attack using surrogate models that essentially ignore the physics of alternating current. Also, a number of authors make claims about the resilience of the Internet to attack or random disruptions (e.g., Albert et al. [3], Cohen et al. [21]), but none attempt to validate their work using an industry-standard network-simulation package (e.g., Lucio et al. [51]), or attempt to validate with experiments on real networks (e.g., Zaragoza and Belo [81]). These surrogate models might be useful, but we do not know.

For real infrastructure systems, operator models often exist that represent “best practices” within a particular engineering or industrial domain. When available, these models ought to be adapted and used. For example, when considering the value of components in electric power infrastructure, one ought to use an industry-standard model of power flow and supply (Salmerón et al. [64, 65]); when considering components of a water-distribution system, one ought to use a standard hydraulic model (Collins et al. [22]; Bhave and Gupta [7, pp. 115–151]); and when considering the value of components of road network, one

ought to use standard traffic-flow model (Beckmann [6], Gazis [37, pp. 185–236], Boyce and Bar-Gera [10]). If such an approach is used, then validation is essentially automatic, since the relevant industry has already performed the required validation. (Models involving multiple infrastructures or simulation certainly warrant investigation, but exceed the scope of this paper.)

2.2. The Attacker A

The maximization in **DAD** represents the actions of an *attacker* who observes defensive preparations and then chooses an *attack plan* \mathbf{x} , for example, $x_\ell = 1$ if component ℓ of the system is attacked, and $x_\ell = 0$, otherwise. Defenses will influence attacks and/or their effects; hence $\mathbf{x} \in X(\mathbf{w})$. The attacker seeks to maximize damage to the operator by maximizing the operator’s cost of operating the system.

We denote the model that results from fixing \mathbf{w} in **DAD** as **AD**(\mathbf{w}), or generically as **AD**: this is an *attacker-defender model*. Danskin [28] describes min-max models that resemble **AD** except that he uses only continuous variables; Moore and Bard [54] describe a more general framework that does allow for integer variables and which includes **AD** as a special case. Unlike **AD**, Moore and Bard’s model does not require that each player’s objective be diametrically opposed. This generalization does not seem useful in our context, however, as its use would force us to infer the attacker’s “true” objective, probably through the impossible-to-validate beliefs of subject-matter experts.

There is a long history in the development of interdiction models to assess the vulnerability of a system, typically a network, to attack. As documented by Schrijver [67], the famous max-flow/min-cut theorem has its origins in a 1955 study of how to interdict the Russian railroad network that, in the event of a war with the West, would have carried materiel from various staging points into eastern Europe (Harris and Ross [42]). That model may be viewed as a specialized instance of **AD** with a binary model of system operation: the system enables positive flow, or it does not. The models described below are more general, and each may be viewed as a full-fledged instance of **AD**.

Fulkerson and Harding [34], Golden [40], and Israeli and Wood [45] formulate and solve network-interdiction problems that maximize the shortest-path length between two designated nodes in the subject network. The first two papers model continuous reductions in capacity with “interdiction effort,” while the last models binary interdictions. Wood [78] minimizes the maximum flow in a capacitated network through interdiction (see also Wollmer [77], Ratliff et al. [62], Phillips [61]); Cormican et al. [23] model and solve a stochastic version of Wood’s problem that minimizes the expected maximum flow through a capacitated network given uncertain arc capacities and/or uncertain attack successes. Lim and Smith [50] present and solve a multicommodity-flow network-interdiction problem. Smith [68] and Wood [79] present overviews of interdiction models. Although Cormican et al. do not use the following terminology, they show how a model formulated with “capacity interdiction” can be reformulated usefully as a model with “cost interdiction,” that is, as a model in which interdiction increases the cost of an activity. This reformulation is often important for efficient solution of **AD**.

Early work on network-interdiction models of the form **AD** was not construed as identifying vulnerabilities in critical infrastructure. Much new work on **AD** models has that explicit purpose, however; for example, see Salmerón et al. [64] and Brown et al. [15, 14]. Such **AD** models have also served as the basis for over 150 “red-team exercises” performed by students at the Naval Postgraduate School. Brown et al. [16] document some insights on the vulnerability of infrastructure from those exercises.

2.3. The Defender D

The outermost minimization in **DAD**, i.e., “D,” represents the actions of a *defender* who takes the first step in this game model by choosing a defensive investment plan, or simply

defense plan, $\mathbf{w} \in W$. This plan may include mounting active defenses, hardening infrastructure against attack, building new infrastructure that is less vulnerable to attack, or adding redundancy. The constraints reflected by W normally include one or more important resource constraints, so by controlling $\mathbf{w} \in W$, the defender seeks to allocate limited resources that make his infrastructure system as resilient as possible to attack.

When defense plans correspond to resource-constrained component hardening, the solution to **DAD** identifies which system components should be protected to minimize the worst-case disruption to operation. In the context of facility location, Church et al. [20] introduce the r -interdiction median problem, a variation of the classical p -median location problem in which individual facilities are unprotected and subject to attack: such a model might help to identify the most important facilities in a supply system. Church and Scaparra [19] and Scaparra and Church [66] extend this work to allocate defensive (“fortification”) resources in order to minimize the impact of interdiction. The p -median problem is, however, only a surrogate for the operation of a real distribution system. One would hope that real investment in the protection of warehouses or other parts of a supply chain would follow from a realistic, validated model of supply, production and distribution (e.g., Geoffrion and Graves [38], Arntzen et al. [4], Brown et al. [17]).

Brown et al. [16] pose, but do not solve, instances of **DAD** in the context of several infrastructure-defense problems. Salmerón et al. [65] develop a “global Benders decomposition algorithm” to solve such models, and apply that algorithm to identifying optimal defensive investments in electric power systems. They solve some large, realistic problems, but their description lacks details and does not cover new construction or capacity expansion as our paper does.

When defense plans $\mathbf{w} \in W$ correspond to capacity expansion or the construction of new infrastructure, **DAD** represents a special type of system-design problem. An extensive literature exists on the design of “survivable networks,” where the objective is often posed as a generalization of the k -node or k -edge connected network problem; see Kerivin and Mahjoub [46] and Grötschel et al. [41] for surveys. Much of this literature uses abstract models as surrogates for real system operation, for example, requiring at least two node-disjoint (or edge-disjoint) paths between all node pairs in a telecommunications network (e.g., Fortz and Labbé [33]). Other network-design papers in telecommunications use simple, flow-based operator models (e.g., Mateus and Patrocínio [53]); these are close in spirit to the operator models we propose for use in **DAD**.

Smith et al. [69] formulate and solve a **DAD**-type model for designing a multicommodity flow network that is robust to optimal attacks. (They also consider models with heuristically planned attacks.) Their network-design constructs resemble ours, and could represent hardening of existing construction as well as new construction. And, similar to our work, they develop a decomposition algorithm for finding an optimal design, i.e., an optimal defense plan. Their algorithm has one key limitation, however: it depends heavily on (a) attacks being represented by bounded, continuous variables (which reduce flow capacity), and on (b) total attack effort being limited by a single knapsack constraint. Even with this limitation, generation of a single constraint (“cut”) for their algorithm’s master problem may require solution of $|A|$ mixed-integer subproblems, where A denotes the set of network arcs. Our methods do not inherently restrict the types of constraints that can be placed on the attacker, except that attacks are presumed to be binary. We also note that our methods work with convex, nonlinear operator models as well as more standard linear programs. The work by Smith et al. might be difficult to extend to the nonlinear case because of its explicit use of dual extreme points from the linear program.

2.4. Deliberate Actions vs. Random Events

In the form described in here, **DAD** models deliberate actions, not random events like natural disasters or accidents. Cormican et al. [23] and Morton et al. [55] extend deterministic

AD models to incorporate random events, however, and **DAD** will extend similarly. For instance, no conceptual barrier exists to modeling the random lifetime of an “emergency spare” that is used to replace a system component damaged by a deliberate attack.

3. **DAD** for Defending a Municipal Transportation Network

To illustrate the **DAD** approach, this section describes an application to protecting a specific infrastructure system. We consider the challenge of officials in a city government who must (a) assess the resilience of their city’s transportation network of roads and bridges to terrorist attacks, and (b) identify cost-effective means to improve that resilience by defending key links or adding redundant infrastructure. The key links in the network are bridges because of the need to connect several islands and, unlike road segments, bridges could take many years to replace. Thus, only bridges are vulnerable to attack in this example.

The operator’s model in this case is a convex, nonlinear program that evaluates total (or, equivalently, average) travel time for a population of travelers traversing a network. The nonlinear program implements the Wardrop traffic-equilibrium model (Wardrop [75], Beckmann [6]), which is employed commonly by traffic engineers (e.g., Gazis [37], Boyce and Bar-Gera [10]). Indeed, commercial traffic-analysis software provides traffic engineers with solutions of this equilibrium model (Correa and Stier-Moses [24]).

In the example model presented here, the cost of system operation is measured in terms of total user travel time for a single period like “the morning commute.” A more detailed model might integrate cost over time until the system’s damaged components are repaired or replaced and the system returns to normal. In effect then, our example assumes that (a) any component that is attacked will be repaired in the same amount of time, (b) any period of peak traffic is like any other, and (c) nonpeak traffic is of no interest. We present a complete model next, but warn the reader that some explanations are left to §4 where an algorithmic framework simplifies those explanations.

Indices and index sets

- $i, j, p \in N$ nodes in a transportation network (intersections, or city areas treated as a single locations in a transportation network); p denotes a population origin for trips;
- $(i, j) \in E$ undirected edges (“links”), i.e., bridges and road segments; $i < j$ is assumed;
- $E_B \subset E$ bridges;
- $(i, j) \in A$ directed arcs (edges with direction of travel included, which may be viewed as traffic lanes); $(i, j) \in E \Leftrightarrow (i, j) \in A \wedge (j, i) \in A$; and
- $d \in D$ defense options; $d \in D_{ij} \subseteq D$ denotes options available for edge $(i, j) \in E$; $d_0 \in D_{ij} \subseteq D$ is a “no-defense” option that leaves edge $(i, j) \in E$ unchanged (i.e., undefended).

Data [units, if applicable]

- b_{pi} for $p \neq i$, $-b_{pi}$ is the number of travelers at p who wish to travel to i [persons]; b_{pp} is the total supply of travelers originating at p ;
- c_{ij}^d length of arc $(i, j) \in A$ under defense option d [kilometers];
- q_{ij}^d “equivalent travel length” added to arc $(i, j) \in A$ under defense option d if the associated edge is attacked [kilometers] (used to penalize travel across attacked edges);
- α_{ij}^d linear term: empirically fit objective-function coefficient for $(i, j) \in A$ under defense option d [minutes/(persons \times kilometers)]; and
- β_{ij}^d quadratic term: empirically fit objective-function coefficient for $(i, j) \in A$ under defense option d [minutes/(persons² \times kilometers)].

Decision variables [units, if applicable]

- w_{ij}^d 1 if $(i, j) \in E$ is defended using defense option $d \in D_{ij}$, and 0 otherwise;
- x_{ij} 1 if $(i, j) \in E$ is attacked, and 0 otherwise; and

y_{pij}^d : traffic volume (over a fixed time window) originating from node p that traverses arc $(i, j) \in A$ under defense option d [persons].

Generic constraints

$A^{\text{DEF}} \mathbf{w} \leq \mathbf{b}^{\text{DEF}}$ generic linear constraints on defense plans; and
 $A^{\text{ATK}} \mathbf{x} \leq \mathbf{b}^{\text{ATK}}$ generic linear constraints on attack plans.

Formulation “DAD-Transport”:

$$z^* = \min_w \max_{x \in X} \min_{y \in Y(w)} f(\mathbf{x}, \mathbf{y}), \quad \text{where} \quad (1)$$

$$f(\mathbf{x}, \mathbf{y}) = \sum_{\substack{(i,j) \in E \\ d \in D_{ij}}} \left[(c_{ij}^d + q_{ij}^d x_{ij}) \left(\alpha_{ij}^d \sum_{p \in N} y_{pij}^d + \beta_{ij}^d \left(\sum_{p \in N} y_{pij}^d \right)^2 \right) \right. \\ \left. + (c_{ji}^d + q_{ji}^d x_{ij}) \left(\alpha_{ji}^d \sum_{p \in N} y_{pji}^d + \beta_{ji}^d \left(\sum_{p \in N} y_{pji}^d \right)^2 \right) \right], \quad (2)$$

$$W = \left\{ \mathbf{w} \in \{0, 1\}^{|E|} \mid A^{\text{DEF}} \mathbf{w} \leq \mathbf{b}^{\text{DEF}}, \sum_{d \in D_{ij}} w_{ij}^d = 1 \quad \forall (i, j) \in E \right\}, \quad (3)$$

$$X = \{ \mathbf{x} \in \{0, 1\}^{|E|} \mid A^{\text{ATK}} \mathbf{x} \leq \mathbf{b}^{\text{ATK}} \}, \quad \text{and} \quad (4)$$

$$Y(\mathbf{w}) = \left\{ \mathbf{y} \in R_+^{|N||A||D|} \mid \sum_{\substack{j \mid (i,j) \in A \\ d \in D_{ij}}} y_{pij}^d - \sum_{\substack{j \mid (j,i) \in A \\ d \in D_{ij}}} y_{pji}^d = b_{pi} \quad \forall p, i \in N, \right. \\ \left. y_{pij}^d + y_{pji}^d \leq b_{pp} w_{ij}^d \quad \forall p \in N, (i, j) \in E, d \in D_{ij} \right\}. \quad (5)$$

With one caveat discussed below, the objective function (1) in **DAD-Transport** measures total travel time for all travelers, given a defense plan, an attack plan, and a set of “traveler flows” in the network. Total time on each arc increases quadratically with the volume of traffic on that arc; §4 provides more details.

The constraint set W (see Equation (3)) will limit total defense expenditures to a posited budget, represented as a simple cardinality constraint.

The constraint set X (see Equation (4)) will limit the total number of edge attacks to a reasonable, worst-case, upper bound.

The first set of constraints in $Y(\mathbf{w})$ (see Equation (5)) define standard, multicommodity flow-balance constraints that ensure that all b_{pp} travelers originating at each $p \in N$ arrive at appropriate destinations. The second set of constraints in $Y(\mathbf{w})$ requires that all travelers traversing an arc use the “version” of that arc, d , that has been prepared by the selected defense option. That is, if $w_{ij}^d = 1$ for edge (i, j) , then all travelers traversing arcs (i, j) and (j, i) are governed by parameters determined by defense option d for edge (i, j) and by whether or not the edge has been attacked. A caveat pertains, however. If a vulnerable edge (i, j) is attacked in our examples, it is destroyed. In this case, the corresponding arc parameters are set so that all flow on (i, j) and (j, i) is 0, unless positive flows are required for feasibility. Positive flow on “destroyed arcs” indicates that the network is disconnected and that total travel time is effectively infinite.

Constraints on defense expenditures will be known, and traffic engineers should have a good model of traffic flow in the region. Thus, parts **D** and **D** of this model, i.e., W and $Y(\mathbf{w})$ should be well understood. “**A**,” i.e., constraint set X , will be modeled in generic terms, and potential attackers and their capabilities will be studied using “capabilities analysis” (e.g., Cragin and Daly [27, pp. 39–57], Steinhäusler et al. [70]). This analysis should provide a

reasonable range for the maximum number of bridges that might be attacked simultaneously. In practice, results within that range would be produced using **DAD** and presented to decision-makers for final action; the examples in §5 illustrate.

Notes. (a) A more detailed model might measure total travel time per person as above, but would adjust for vehicles and the number of persons per vehicle. In effect, **DAD-Transport** assumes pedestrian traffic or one person per vehicle.

(b) “Supply of travelers” assumes a period of time over which all travel will take place, such as during a peak morning commute period of two hours. Parameters c_{ij}^d , α_{ij}^d , and β_{ij}^d are set accordingly.

(c) The constraints $A^{\text{DEF}}\mathbf{w} \leq \mathbf{b}^{\text{DEF}}$ and $A^{\text{ATK}}\mathbf{x} \leq \mathbf{b}^{\text{ATK}}$ represent arbitrary linear constraints on defense plans and attack plans, respectively. The only such constraints used in our examples are (a) a cardinality constraint on the number of bridges defended, n^{DEF} , (b) a cardinality constraint on the number of bridges attacked, n^{ATK} , and (c) constraints to reflect that fact that “nonbridges” are invulnerable to attack and need not be defended. Constraints $A^{\text{DEF}}\mathbf{w} \leq \mathbf{b}^{\text{DEF}}$ and $A^{\text{ATK}}\mathbf{x} \leq \mathbf{b}^{\text{ATK}}$ could reflect limited budgets covering several categories (e.g., money, labor resources, energy resources), logical conditions between attacks or between defenses, and so on.

4. A Decomposition Algorithm to Solve **DAD**

This section develops a decomposition algorithm to solve **DAD-Transport** and more general instances of **DAD**. We first present additional detail on the operator’s model for this problem, and describe several subsidiary formulations used in the algorithm.

4.1. The Operator’s Problem

Given a fixed infrastructure-defense plan $\widehat{\mathbf{w}} \in W$, and a fixed attack plan $\widehat{\mathbf{x}} \in X$, the following model defines the operator’s problem:

$$\mathbf{DAD}(\widehat{\mathbf{w}}, \widehat{\mathbf{x}}, \cdot): z_{\mathbf{D}}^*(\widehat{\mathbf{w}}, \widehat{\mathbf{x}}) = \min_{\mathbf{y} \in Y(\widehat{\mathbf{w}})} f(\widehat{\mathbf{x}}, \mathbf{y}).$$

The notation required to describe the full **DAD** model in Equations (1)–(5) makes the operator’s problem appear more complicated than it is. Ignoring the caveat on penalties used to discourage use of destroyed bridges, $\mathbf{DAD}(\widehat{\mathbf{w}}, \widehat{\mathbf{x}}, \cdot)$ is a simple multicommodity network-flow model with a quadratic objective function. Each commodity is defined in terms of the origin of a group of travelers, but could be based on destination, or a commodity could be defined for each origin-destination (O-D) pair. The objective function measures total travel time over all travelers, over some normalizing interval of time, by summing travel time on each arc $(i, j) \in A$ (through a summation over $(i, j) \in E$). Because of congestion effects, total travel time for users of arc (i, j) depends quadratically on the total number of travelers that traverse that arc, $\bar{y}_{ij}^d \equiv \sum_{p \in N} y_{p ij}^d$.

Total travel time on arc (i, j) may be expressed as $\bar{y}_{ij}^d g(\bar{y}_{ij}^d)$, where $g(\cdot)$ is called a “delay function.” Numerous delay functions have been used in the literature, but simple polynomial functions are standard and have been validated experimentally (e.g., LeBlanc et al. [48]). We use a linear delay function for computational simplicity here, yielding a quadratic objective function; future work will investigate the use of using higher-degree polynomials and, perhaps, other functional forms.

The nonlinear program $\mathbf{DAD}(\widehat{\mathbf{w}}, \widehat{\mathbf{x}}, \cdot)$ (Beckmann [6]) derives from the basic *traffic-equilibrium model* (or *traffic assignment model*) described by Wardrop [75]. Florian and Nguyen [32] provide one the first validations of the model. Many refinements of the basic model have appeared since the 1960s (e.g., Boyce and Bar-Gera [10]), but the basic model is still in use (Correa and Stier-Moses [24]). We note that early traffic-equilibrium models defined commodities through O-D pairs, and formulations are still often described in that

manner. Defining commodities by origin or by destination, however, is clearly more efficient (Petersen [60], Leblanc et al. [48]).

O-D demands are estimated by sampling actual traffic and statistical estimation. The complete estimation step is often referred to as “trip generation”; for example, see Van Zuylen and Willumsen [72], and Mannering et al. [52, pp. 293–298]. As described by Wardrop [75], each traveler is assumed to follow an O-D path such that total travel time for all travelers is minimized: this maximizes “societal good.” Wardrop’s equilibrium conditions imply that total travel time on a link is a convex increasing function of traffic density; empirical work verifies the validity of this functional form. A more refined model might replace travel time with a “generalized cost of travel,” which could include travel time, tolls, out-of-vehicle time, and other factors; see Abrahamsson and Lundqvist [1] and Boyce and Bar-Gera [10].

4.2. The Attacker-Defender Subproblem

The overall decomposition algorithm for **DAD** will solve a sequence of attacker-defender subproblems that result by fixing $\widehat{\mathbf{w}} \in W$:

$$\mathbf{DAD}(\widehat{\mathbf{w}}, \cdot, \cdot): z_{\mathbf{AD}}^*(\widehat{\mathbf{w}}) = \max_{\mathbf{x} \in X} \min_{\mathbf{y} \in Y(\widehat{\mathbf{w}})} f(\mathbf{x}, \mathbf{y}).$$

An optimal or near-optimal solution to this problem is denoted $\mathbf{x}^*(\widehat{\mathbf{w}})$.

$\mathbf{DAD}(\widehat{\mathbf{w}}, \cdot, \cdot)$ is a “simple” attacker-defender model, which we solve through Benders decomposition. This solution method is standard for such problems, so we omit a description (see Cormican et al. [23], Israeli and Wood [45] for examples). We find it computationally advantageous to specify a nonzero optimality gap in the solution of $\mathbf{DAD}(\widehat{\mathbf{w}}, \cdot, \cdot)$, however, and this complicates the decomposition algorithm for **DAD**. Zakeri et al. [80] describe and overcome a similar complication, which arises when solving a linear program through Benders decomposition, and not solving subproblems to optimality. Section 4.4 will explain how to handle the **DAD** version of this issue, and we make several definitions in advance for that explanation:

$\varepsilon_{\mathbf{AD}}$ user-specified, nonnegative, relative optimality gap for the decomposition algorithm that solves $\mathbf{DAD}(\widehat{\mathbf{w}}, \cdot, \cdot)$,
 $z_{\mathbf{AD}}^{\text{LO}}, z_{\mathbf{AD}}^{\text{UP}}$ lower and upper bounds provided at termination of the decomposition algorithm that solves $\mathbf{DAD}(\widehat{\mathbf{w}}, \cdot, \cdot)$; these values must satisfy $z_{\mathbf{AD}}^{\text{UP}} - z_{\mathbf{AD}}^{\text{LO}} \leq \varepsilon_{\mathbf{AD}} z_{\mathbf{AD}}^{\text{LO}}$.

4.3. A Detailed Decomposition Algorithm

We develop a decomposition algorithm here to solve **DAD-Transport** and similar problems. Some additional definitions follow:

\widehat{X} the set of all feasible attack plans viewed as an enumerated set; and
 $y_{p|ij}^{dk}$ traffic volume originating from node p that traverses arc $(i, j) \in A$ under defense option d , in response to attack plan $\widehat{\mathbf{x}}^k \in \widehat{X}$ [persons].

Letting \mathbf{y}^k denote the vector form of $y_{p|ij}^{dk}$, we may now reformulate **DAD-Transport** as

$$z^* = \min_{\mathbf{w} \in W} \max_{\widehat{\mathbf{x}}^k \in \widehat{X}} \min_{\mathbf{y}^k \in Y(\mathbf{w})} f(\widehat{\mathbf{x}}^k, \mathbf{y}^k). \quad (6)$$

Note that $\mathbf{y}^k \in Y(\mathbf{w})$ implies a separate set of (identical) constraints for each flow vector \mathbf{y}^k .

Because the defender in formulation (6) can now choose each set of flows \mathbf{y}^k independently in anticipation of each feasible attack plan, we can exchange the innermost “min” and “max” to obtain a conceptually simpler min–max problem:

$$z^* = \min_{\mathbf{w} \in W, \mathbf{y}^k \in Y(\mathbf{w})} \max_{\widehat{\mathbf{x}}^k \in \widehat{X}} f(\widehat{\mathbf{x}}^k, \mathbf{y}^k). \quad (7)$$

Naturally, we cannot hope to solve realistic instances of (7) by enumerating all attack plans, and creating a separate traffic-flow problem for each. But, this formulation leads to a decomposition algorithm that generates attack plans on an as-needed basis, and we hope to identify an ε -optimal solution long before enumerating all attack plans. Given a set of K feasible attack plans $\widehat{X}^K = \{\widehat{\mathbf{x}}^1, \dots, \widehat{\mathbf{x}}^K\}$, we formulate a “relaxed **DAD** master problem” for this decomposition algorithm as follows:

DAD-MP(\widehat{X}^K):

$$z^*(\widehat{X}^K) = \min_{w \in W, \mathbf{y}^1, \dots, \mathbf{y}^K} z \quad (8)$$

$$\text{s.t. } z \geq f(\widehat{\mathbf{x}}^k, \mathbf{y}^k) \quad \forall \widehat{\mathbf{x}}^k \in \widehat{X}^K, \quad (9)$$

$$\mathbf{y}^k \in Y(\widehat{\mathbf{w}}) \quad \text{for } k = 1, \dots, K. \quad (10)$$

This model is a quadratically constrained integer program that may not be solved exactly, and the following parameter and output values must therefore be defined:

ε_{MP} user-specified, nonnegative, relative optimality gap for the algorithm that solves **DAD-MP**(\widehat{X}^K); and
 $z_{\text{MP}}^{\text{LO}}, z_{\text{MP}}^{\text{UP}}$ lower and upper bounds provide at termination of the algorithm that solves **DAD-MP**(\widehat{X}^K); these values must satisfy $z_{\text{MP}}^{\text{UP}} - z_{\text{MP}}^{\text{LO}} \leq \varepsilon_{\text{MP}} z_{\text{MP}}^{\text{LO}}$.

We can now state a full decomposition algorithm for solving **DAD**.

Algorithm **DAD-Decomp**

Input: Full **DAD** problem data and optimality tolerances $\varepsilon, \varepsilon_{\text{MP}}, \varepsilon_{\text{AD}} \geq 0$ for the overall decomposition, the **DAD** master problem, and the **AD** subproblem, respectively.

*/** $\varepsilon \geq \varepsilon_{\text{MP}}$ is assumed.**/*

Output: ε -optimal defense plan \mathbf{w}^* and corresponding attack plan \mathbf{x}^* ;

1. $\text{LB} \leftarrow -\infty; \text{UB} \leftarrow \infty; K \leftarrow 1;$
2. for (all $(i, j) \in E$) $\{\widehat{w}_{ij}^{d_0 K} \leftarrow 1; \widehat{w}_{ij}^{d_0 K} \leftarrow 0, d \neq d_0; \}; \mathbf{w}^* \leftarrow \widehat{\mathbf{w}}^K;$

*/**That is, choose “no defense” as the initial defense plan and as the incumbent solution.**/*

3. **Subproblem:** Solve **DAD**($\widehat{\mathbf{w}}^K, \cdot, \cdot$) to determine attack plan $\widehat{\mathbf{x}}^K$ given defense plan $\widehat{\mathbf{w}}^K$ such that $z_{\text{AD}}^{\text{UP}} - z_{\text{AD}}^{\text{LO}} \leq \varepsilon_{\text{AD}} z_{\text{AD}}^{\text{LO}};$

*/**We assume $z_{\text{AD}}^{\text{UP}}, z_{\text{AD}}^{\text{LO}} \geq 0$ **/*

4. if $(z_{\text{AD}}^{\text{UP}} < \text{UB}) \{ \text{UB} \leftarrow z_{\text{AD}}^{\text{UP}}; \mathbf{w}^* \leftarrow \widehat{\mathbf{w}}^K; \mathbf{x}^* \leftarrow \widehat{\mathbf{x}}^K; \}$
5. if $(\text{UB} - \text{LB} \leq \varepsilon \text{LB})$ go to **End**;

6. if $\widehat{\mathbf{x}}^K$ repeats any prior attack, i.e., $\widehat{\mathbf{x}}^K \in \widehat{X}^K$, temporarily add one “solution-elimination constraint” to **DAD**($\widehat{\mathbf{w}}^K, \cdot, \cdot$) for each $\widehat{\mathbf{x}}^k \in \widehat{X}^K$, and re-solve for a new $\widehat{\mathbf{x}}^K$;

*/**Solution-elimination constraints are described below. For simplicity, the algorithm ignores the possibility that problem in Step 7 could be infeasible.**/*

7. $\widehat{X}^K \leftarrow \widehat{X}^{K-1} \cup \{\widehat{\mathbf{x}}^K\};$

8. **Master Problem:** Solve **DAD-MP** (\widehat{X}^K) to determine defense plan $\widehat{\mathbf{w}}^{K+1}$ such that $z_{\text{MP}}^{\text{UP}} - z_{\text{MP}}^{\text{LO}} \leq \varepsilon_{\text{MP}} z_{\text{MP}}^{\text{LO}};$

*/**We assume $z_{\text{MP}}^{\text{UP}}, z_{\text{MP}}^{\text{LO}} \geq 0$ **/*

9. if $(z_{\text{MP}}^{\text{LO}} > \text{LB}) \text{LB} \leftarrow z_{\text{MP}}^{\text{LO}};$
10. if $(\text{UB} - \text{LB} \leq \varepsilon \text{LB})$ go to **End**;

11. $K \leftarrow K + 1;$ go to **Subproblem**;

12. **End:** print (“ ε -optimal defense plan and corresponding attack plan are,” $\mathbf{w}^*, \mathbf{x}^*$).

If the **AD** subproblems are not solved to optimality in each step, the algorithm can repeat an attack plan. This can lead to cycling, because the bounds will not change, the master

problem’s feasible region will not change, no new defense-plan solutions need be generated, and no new attack plan need be generated in response. There are a few ways to handle this difficulty, the simplest of which is to record every subproblem solution (i.e., attack plan) observed, and, if one repeats, reduce the tolerances on the respective problem(s) until a new solution is found. This may cause run times to increase dramatically, however.

Another approach, the one we take, forces the generation of a new attack plan (the algorithm assumes one exists) by adding a set of K “solution-elimination constraints” (SECs) whenever an attack plan $\hat{\mathbf{x}}^K$ repeats:

$$\sum_{(i,j) \in E: \hat{x}_{ij}^k = 0} x_{ij} \geq 1 \quad \forall \hat{\mathbf{x}}^k \in \hat{X}^K. \tag{11}$$

The SEC in (11) based on a specific attack plan $\hat{\mathbf{x}}^k$ makes that plan infeasible in the master problem, along with any “dominated” plans $\hat{\mathbf{x}} \leq \hat{\mathbf{x}}^k$. Thus, no solution $\hat{\mathbf{x}} \in \hat{X}^K$ can be regenerated at Step 7 of the algorithm. (Note that no bounds are updated in the algorithm when using SECs because the validity of those bounds cannot be guaranteed.) The SEC used here, which is a special case of a “super-valid inequality” (Israeli and Wood [45]), requires that any attack plan that targets multiple components dominate all plans that target a strict subset of those components (i.e., the attacker will always prefer to target more components of the system than fewer). If no such dominance relationship exists, constraints that enforce a lower bound of 1 on the Hamming distance between a new solution and each $\hat{\mathbf{x}}^k \in \hat{X}^K$ could replace (11); see Brown and Dell [13]. We note that adding SECs for $\hat{\mathbf{w}}$ in the master problem provides a third approach to handling nonzero optimality gaps in the decomposition algorithm, but we have not yet explored that possibility.

5. A Computational Example: The Seven Bridges of Königsberg

In 1758, Leonhard Euler published a paper using as a motivating example the propensity of city residents to traverse the seven bridges of Königsberg (Euler [31]); see Figure 1(a). Using the graph shown in Figure 1(b), Euler proved that no walking path existed through the city that crossed each bridge exactly once. We adapt this well-known example from seminal graph theory to more modern concerns.

City officials are concerned about the disruptions to city traffic, and thereby to the local economy, that would result from the destruction of one or more bridges by terrorists. Officials




FIGURE 1. The seven bridges of Königsberg.

Notes. (a) A drawing of the seven bridges (Kraitchik [47, pp. 209–211]). (b) In Euler’s graph representation, each vertex is a land mass and each undirected edge is a bridge. (c) For illustrative purposes, we adopt a network representation that reflects the bridges (heavy lines), normal road segments (horizontal lines at top and bottom) and artificial, “intra-island edges” represented by the graph cliques on islands A and D. Bridges are subject to attack and congestion; road segments are subject to congestion but not attack; and the intra-island edges are subject to neither attack nor congestion, but do require a fixed amount of time to traverse. For an indication of the scale here, note that the small central island is about one half kilometer in its longest dimension.

TABLE 1. Node data for **DAD** modeling of the Königsberg transportation network.

Nodes $p \in N$	Supply b_{pp} (persons)	Demand $-b_{pi}$ (persons)
Aa, Ab, Ac, Ae, Af, De, Df, Dg	200	Proportional to supply: $-b_{pi} \equiv b_{pp} b_{ii} / \sum_{j \neq p} b_{jj} \quad \forall p, i \in N$
Ba, Bf, Bg Cc, Cd, Cg	800 1,200	

Notes. The data here apply to the **DAD-Transport** model of the network shown in Figure 1(c). (In 1700, Königsberg had a population of about 40,000, so these numbers are plausible.)

also want to know if worst-case disruptions could be reduced, i.e., resilience enhanced, by defending bridges from attack or making other infrastructure improvements. We measure functionality of the transportation network in terms of the average travel time that a citizen would experience in moving about the city on a busy morning.

The data requirements for this problem are modest. Figure 1(c) shows an abstract representation of the main routes in the city. Table 1 provides data on the nodes for this problem; Table 2 provides basic edge and arc data in the absence of attack; the appendix presents detailed arc data for all examples. The **DAD** examples presented here are small enough that they could be solved by total enumeration, that is, by solving the nonlinear program **DAD**($\widehat{\mathbf{w}}, \widehat{\mathbf{x}}, \cdot$) for each feasible combination of $\widehat{\mathbf{w}}$ and $\widehat{\mathbf{x}}$. The decomposition applies broadly, however, and we trust that the examples serve well to illustrate its use.

We must specify X and W for this problem, also. Capabilities analysis indicates that bridges are key targets, and that at most three could be attacked at one time. Thus,

$$X \equiv \left\{ \mathbf{x} \in \{0, 1\}^{|E|} \mid \sum_{(i,j) \in E_B} x_{ij} \leq n^{\text{ATK}}, x_{ij} \equiv 0 \quad \forall (i,j) \in E \setminus E_B \right\} \quad \text{for } n^{\text{ATK}} = 0, 1, 2 \text{ or } 3.$$

Planners believe that the city budget will allow for the defense of up to four bridges, and thus


$$W \equiv \left\{ \mathbf{w} \in \{0, 1\}^{|E|} \mid \sum_{(i,j) \in E_B} w_{ij}^{d_1} = n^{\text{DEF}}, \sum_{d \in D_{ij}} w_{ij}^d = 1 \quad \forall (i,j) \in E, \right. \\ \left. w_{ij}^{d_0} \equiv 1 \quad \forall (i,j) \in E \setminus E_B \right\} \quad \text{for } n^{\text{DEF}} = 0, 1, 2, 3 \text{ or } 4.$$

TABLE 2. Edge and arc data for **DAD** modeling of the Königsberg transportation network: nominal system parameters (i.e., assuming $\widehat{\mathbf{w}} = \mathbf{0}$, $\widehat{\mathbf{x}} = \mathbf{0}$ and $d = d_0$).

Edge type	Edges (i, j)	$l_{ij}^{d_0}, l_{ji}^{d_0}$	$\alpha_{ij}^{d_0}, \alpha_{ji}^{d_0}$	$\beta_{ij}^{d_0}, \beta_{ji}^{d_0}$
Bridge	(Aa, Ba), (Ab, Bb), (Ac, Cc), (Ad, Cd), (Ae, De), (Bf, Df), (Cg, Dg)	1	5	0.020
Road segment	(Ba, Bb), (Bb, Bf), (Cc, Cd), (Cd, Cg)	1	15	0.005
Intra-island	All having form (Ax, Ay) or (Dx, Dy)	1	5	0.000

Notes. The data in this table covers all edges in the network shown in Figure 1(c), and all implied antiparallel arcs. All edge lengths are 1 kilometer for simplicity. The intra-island edges (i, j) represent a complex network which is assumed free of congestion delays, i.e., $\beta_{ij}^{d_0} = 0$. However, traversing any of these edges does incur five minutes of travel time, i.e., $\alpha_{ij}^{d_0} = 5$. (Note that bridge edges can also be designated by single letters, i.e., **a, b, . . . , g**.)

FIGURE 2. Nominal, optimal operation of the Königsberg transportation network bridges.




Notes. (a) Optimal number of travelers using each traffic lane, on each road and bridge, under nominal conditions. (The double lines here represent the two arcs for each edge.) The average travel time is 37.6 minutes. (b) Optimal routes followed by the 1,200 travelers originating at node Cc. For clarity, we omit arcs without flow.

In the absence of attack, travelers may use any of the bridges to convey traffic. The minimum-time solution incurs an average travel time of 37.6 minutes; Figure 2 depicts the optimal solution. We cannot easily illustrate the individual routes followed by each of the 7,600 travelers, but Figure 2(a) provides a sense of congestion. In addition, for travelers originating at one selected node, Cc, Figure 2(b) shows the total number of travelers on each road and bridge.

5.1. Königsberg’s Bridges Attacked with No Defenses


It is worthwhile investigating, by solving AD for various values of n^{ATK} , how attacks might affect Königsberg’s traffic flow if no bridges are defended. It turns out that the most disruptive single-bridge attack ($n^{ATK} = 1$) is on bridge c, but this results in an increase in average travel time of only 9.2 minutes, about 24%; see Figure 3(a). If capabilities analysis shows that terrorists could destroy only a single bridge, we might conclude that the city is already “well defended.” The optimal two-bridge attack targets bridges c and d (Figure 3(b)), and average travel time increases by 44.5 minutes. Perhaps officials should become worried if a two-bridge attack appears possible.

FIGURE 3. Traffic flow in Königsberg resulting from the most disruptive one- and two-bridge attacks on an undefended system.



Notes. (a) The worst-case one-bridge attack destroys c, resulting in an average travel time of 46.8 minutes. (b) The worst-case two-bridge attack destroys bridges c and d, resulting in an average travel time of 82.1 minutes. In each case, the figures indicate the optimal rerouted flows in response to the attacks.

FIGURE 4. Traffic flow in Königsberg resulting from all possible one-, two-, and three-bridge attacks on an undefended system.



Notes. (a) The plot depicts average travel time resulting from each attack plan having $n^{\text{ATK}} = 1, 2,$ or 3 ; results are ranked, from most to least disruptive for each value of n^{ATK} . There are three three-bridge attacks that disconnect the network, so the resulting average travel time, denoted “140*,” could be arbitrarily high. The inset expands results for single-bridge attacks and shows the actual bridge involved in each. (b) Considerable variation appears in the disruption caused by attacks involving one, two, or three bridges. In particular, the most disruptive two-bridge and three-bridge attacks result in substantially more travel delay than a random attack plan having the same number of bridges.

Figure 4 shows, given no defenses and for different values of n^{ATK} , rank-ordered lists of optimal attack plans and their outcomes. We observe several features. First, for any value of n^{ATK} , a considerable difference can appear in the disruption caused by an optimal attack plan versus a plan chosen randomly. For $n^{\text{ATK}} = 1$, the loss of bridge **c** increases the average travel time by 9.2 minutes, while the expected increase in average travel time is 7.1 minutes for a “random attack,” that is, if a “dumb” attacker were to choose to attack each bridge with probability $1/7$. For $n^{\text{ATK}} = 2$, the loss of bridges **c** and **d** increases the average travel time by 44.5 minutes, while a random attack increases expected average travel time by only 18.9 minutes. When $n^{\text{ATK}} = 3$, three attack plans, namely **[a, b, f]**, **[c, d, g]**, and **[e, f, g]**, can disconnect the network, and average travel time can become arbitrarily long. In contrast, a random three-bridge attack among the other combinations results in an expected increase to average travel time of (only) 36.0 minutes.

In this example, the optimal one-bridge attack **[c]**, optimal two-bridge attack **[c, d]**, and optimal three-bridge attack **[c, d, f]** define monotonic attack plans, i.e., $\mathbf{x}_1^* \leq \mathbf{x}_2^* \leq \mathbf{x}_3^*$, where the subscript corresponds to n^{ATK} . This is good news for an attacker of Königsberg, who can follow a simple prioritized list of attacks: **c** then **d** then **f**. If he is stopped (for example, caught or killed) after attacking $n < 3$ bridges, he has been maximally disruptive given that $n^{\text{ATK}} = n$. For larger infrastructure systems, however, we typically find that simple prioritized lists yield a sequence of suboptimal attack plans.

Results in Table 3 also indicate the dubiousness of basing infrastructure-defense analysis on a single, heuristically chosen attack plan. For instance, a reasonable greedy heuristic would first attack the link with the largest nominal traveler flow, but such a choice is only the third best for the attacker. See Smith et al. [69] for a related discussion.

5.2. Optimal Defenses for Königsberg

We conclude from the previous section that (a) a small number of attacks can cause substantial disruption to travel in Königsberg, (b) a defensive model that assumes random attacks could leave the city open for a highly disruptive optimal attack plan, and (c) a defensive model that plans against a heuristically derived attack plan is also open to making a serious error. So, solving **DAD** near-optimally could give some important information to city officials, and any solution needs to be based on, or at least imply, near-optimal solutions of **AD(w)**.

TABLE 3. Heuristically chosen and optimal one-bridge attacks.

Bridge	Total baseline traffic (persons)	Rank in baseline traffic	Increase in average travel time if destroyed (minutes)	Rank in disruption
a	1,190	5	6.9	4
b	1,444	2	6.4	6
c	1,407	3	9.2	1
d	1,687	1	8.3	3
e	661	7	3.1	7
f	1,070	6	6.9	5
g	1,205	4	8.9	2

Notes. The bridge that carries the most traffic is not necessarily the bridge that, if lost, results in the greatest disruption (increase in average travel time).

Assume now that analysts believe that the worst possible attack on the Königsberg bridges would destroy two or three bridges. City officials are unsure of their budget for bridge defenses, and would like to know the optimal set of bridges to defend for each “budget level” $n^{\text{DEF}} \in \{1, \dots, 4\}$. Solution of **DAD-Transport** will provide the answers.

We use GAMS (GAMS [35]) to formulate all models in the decomposition for **DAD-Transport** and solve them using CPLEX 12.02 (IBM [44]) on a Lenovo T510 laptop computer. Master problems and nonlinear subproblems are solved by specifying the quadratically constrained programming option in GAMS (“QCP = CPLEX”), which also handles quadratic objective functions. We run the full decomposition algorithm on the Königsberg data using tolerances $\varepsilon = 0.01$, $\varepsilon_{\text{MP}} = 0.01$, and $\varepsilon_{\text{AD}} = 0.001$. No individual problem requires more than 10 minutes to solve and, in total, results reported in Table 4 require less than 20 minutes to produce.

Table 4 presents initial results. The table shows the variety of optimal defense plans that arise when multiple bridges can be defended and when multiple bridges may be attacked. Note that the optimal defense-plan vector is not necessarily monotonic in the number of bridges defended. Specifically, for neither $n^{\text{ATK}} = 2$ nor $n^{\text{ATK}} = 3$ does an optimal one-bridge

TABLE 4. Optimal bridges to defend in Königsberg.

Number of bridges attacked n^{ATK}	Number of bridges defended n^{DEF}	Bridges optimally defended	Bridges attacked after defense	Minimized average travel time (minutes)	Num. of AD problems solved in Alg. DAD-Decomp	Num. of AD problems solved if using enumeration
2	1	c	a, b	75.9	3	7
	2	b, d	c, g	65.3	5	21
	3	b, c, d	a, f	58.9	7	35
	4	b, c, f, g	a, d	55.0	12	35
3	1	d*	a, b, f	∞	3	7
	2	c, f	a, b, g	103.4	6	21
	3	b, d, g	c, e, f	70.5	9	35
	4	b, d, f, g	a, c, e	59.2	12	35

Notes. For each number of attacks and defenses this table presents an optimal defense, and a resulting optimal attack, determined using the decomposition algorithm. For three attacks and one defense, the optimal solution (defend bridge **d**) is arbitrary (denoted by an asterisk): if any single bridge is defended, three bridges can always be attacked so that some travelers cannot reach their destinations. In such a case, the resulting objective-function value is arbitrarily large.

TABLE 5. Optimal bridges to defend and road segments to upgrade in Königsberg.

Number of bridges attacked n^{ATK}	Number of bridges defended n^{DEF}	Optimal bridges to defend	Optimal road segments to upgrade	Bridges attacked after defense	Minimized average travel time (minutes)	Avg. travel time decrease beyond bridge defense alone (minutes)
2	1	c	(Bb, Bf), (Cd, Cg)	a, b	68.5	7.3
	2	b, d	(Cc, Cd), (Cd, Cg)	c, g	59.0	6.3
	3	b, c, d	(Bb, Bf), (Cd, Cg)	a, f	54.4	4.6
	4	a, d, f, g	(Cc, Cd), (Cd, Cg)	b, c	49.3	5.7
3	1	d*	(Cc, Cd), (Cd, Cg)	a, b, f	∞	—
	2	c, f	(Bb, Bf), (Cd, Cg)	a, b, g	96.1	7.4
	3	b, d, g	(Cc, Cd), (Cd, Cg)	a, c, f	64.2	6.3
	4	a, d, f, g	(Cc, Cd), (Cd, Cg)	b, c, e	52.7	6.5

Notes. Again, the asterisk in the first row for three attacks indicates that the network becomes disconnected. With the exception of the shaded rows involving 4 defenses, the optimal bridge defenses and the optimal attacks remain the same with or without road-segment upgrades. Average delays are reduced by between 7% and 11% compared to Table 4.

defense define a subset of an optimal two-bridge defense. Thus, no optimal prioritized list of defenses can be created. (Alternate optimal solutions might make this possible, but do not.)

As one would expect, solutions reflect “diminishing returns” as the number of defended bridges grows. For example, for both values of n^{ATK} , the difference between defending two bridges and defending three exceeds the difference between defending three and four.

5.3. Optimal Defenses for Königsberg: Extensions to New Construction

In addition to considering defenses on some of the city’s bridges, a separate line item exists in the Königsberg city budget to upgrade any two of the road segments (**Ba, Bb**), (**Bb, Bf**), (**Cc, Cd**), (**Cd, Cg**) for less congestion and faster travel. The question is then: which combination of n^{DEF} defended bridges and two upgraded road segments creates the most resilient transportation system, given that two or three bridges might be attacked? For our purposes, an upgrade on a road segment reduces $\alpha_{ij} = 15$ to $\alpha_{ij} = 10$ and reduces $\beta_{ij} = 0.005$ to $\beta_{ij} = 0.001$. To evaluate these alternatives, we add an edge for each candidate road improvement to the base model, and add an ad-hoc constraint limiting the number of these new defensive improvements to two.

Algorithm DAD-Decomp extends to this new situation easily and produces the results shown in Table 5 in about 16 minutes of computation time. Unfortunately, those results also show that that the city cannot substantially improve resilience to attack of its transportation network by upgrading road segments.

One city planner therefore asks “What if we shift those road-upgrade funds into building a new, invulnerable bridge (**Ba, Cc**)?” (Actually, current-day Kaliningrad possesses such a bridge.) We assume that the budget remains unspecified for protecting the other bridges, and compute results analogous to Tables 4 and 5 with an invulnerable bridge $(i, j) = (\mathbf{Ba}, \mathbf{Cc})$ in place, having parameters $c_{ij} = 3$, $\alpha_{ij} = 5$, and $\beta_{ij} = 0.01$. (The new bridge will be three times longer than the other bridges, but subject to less congestion; compare to values in Table 3.) This requires adding just one edge to the base case to represent the new bridge. The results in Table 6, computed in less than four minutes, show that the new bridge would enhance resilience of the Königsberg road-and-bridge network substantially, and that option is much better than upgrading any two road segments.

TABLE 6. Optimal bridges to defend in Königsberg given that an invulnerable bridge (**Ba, Cc**) is built, and no other new construction is possible.

Number of bridges attacked n_{ATK}	Number of bridges defended n^{DEF}	Optimal bridges to defend	Bridges attacked after defense	Minimized average travel time (minutes)	Avg. travel time decrease beyond bridge defense alone (minutes)
2	1	d	e, g	53.5	22.3
	2	d, g	b, f	52.2	13.1
	3	d, f, g	a, b	48.8	10.2
	4	b, d, f, g	a, c	43.8	11.2
3	1	g	a, b, f	75.1	∞
	2	f, g	c, d, e	60.5	43.0
	3	b, d, e	c, f, g	53.3	17.3
	4	b, d, f, g	a, c, e	46.1	13.1

Notes. Note the much-reduced travel times compared to Tables 4 and 5. With the new bridge, a one-bridge defense also suffices to prevent a three-bridge attack from disconnecting the network. The shaded rows identify optimal defense and attack plans that differ from solutions obtained without the new-bridge option.

6. Conclusions and Areas for Future Research

This paper has demonstrated how a three-stage, sequential game provides an appropriate paradigm for planning budget-limited defenses and/or new construction that will maximize the resilience of a critical infrastructure system subject to attack by an intelligent adversary. A defender-attacker-defender model (**DAD**) represents the following: (1) a *defender* makes budget-limited investments to improve an infrastructure system; (2) an *attacker* sees those investments and attacks the system so as to maximize damage; and (3) damage is measured in terms of the cost (or increased cost, decreased value, etc.) that the *defender-as-operator*, or simply *operator*, incurs when operating the system optimally. Cost is evaluated by solving an *operator's model* which, rather than using untested surrogate measures of operational effectiveness, reflects real measures such as travel delay, unserved demand, throughput, etc. Operations of an electric power grid, for instance, should be modeled using an industry-standard power-flow model, and road congestion should be measured using an industry-standard traffic-flow model (at least until the usefulness of simpler surrogate models is established). Tests using a standard traffic-equilibrium model show how no actual operator of the system may be necessary, as this model represents the actions of delay-minimizing travelers.

The paper has also developed a general decomposition algorithm for solving **DAD** models. We solve model instances with a relative optimality tolerance of 1% (i.e., $\varepsilon = 0.01$) to enable interested researchers to reproduce our results. Even though our algorithm's master problem is an integer nonlinear programming problem, this tight tolerance leads to scenario solution times of only a few seconds to a few minutes.

The attacks envisaged in this paper are primarily physical, but communications networks like the Internet are subject to "cyber-attacks." Defense against cyber-attacks may be amenable to study via **DAD**, and this area needs investigation. We have investigated single infrastructure systems, yet attacks on one system may affect another; for example, an electric power line carried by a bridge may be lost if the bridge is attacked, and the resulting power outage may increase traffic delays through a loss of traffic signals. This topic certainly warrants study, also.

Appendix

This appendix presents the arc data which, along with the node data in §3, suffices to reproduce all results in this paper. See §2.1 for definitions.

TABLE A.1. Base-case arc data for undefended roads and bridges of Königsberg.

Arc tail i	Arc head j	$c_{ij}^{d_0}$	$q_{ij}^{d_0}$	$\alpha_{ij}^{d_0}$	$\beta_{ij}^{d_0}$	Arc tail i	Arc head j	$c_{ij}^{d_0}$	$q_{ij}^{d_0}$	$\alpha_{ij}^{d_0}$	$\beta_{ij}^{d_0}$
Aa	Ab	1	0	5	0	Ab	Aa	1	0	5	0
Aa	Ac	1	0	5	0	Ac	Aa	1	0	5	0
Aa	Ad	1	0	5	0	Ad	Aa	1	0	5	0
Aa	Ae	1	0	5	0	Ae	Aa	1	0	5	0
Ab	Ac	1	0	5	0	Ac	Ab	1	0	5	0
Ab	Ad	1	0	5	0	Ad	Ab	1	0	5	0
Ab	Ae	1	0	5	0	Ae	Ab	1	0	5	0
Ac	Ad	1	0	5	0	Ad	Ac	1	0	5	0
Ac	Ae	1	0	5	0	Ae	Ac	1	0	5	0
Ad	Ae	1	0	5	0	Ae	Ad	1	0	5	0
Aa	Ba	1	1,000	5	0.020	Ba	Aa	1	1,000	5	0.020
Ab	Bb	1	1,000	5	0.020	Bb	Ab	1	1,000	5	0.020
Ac	Cc	1	1,000	5	0.020	Cc	Ac	1	1,000	5	0.020
Ad	Cd	1	1,000	5	0.020	Cd	Ad	1	1,000	5	0.020
Ae	De	1	1,000	5	0.020	De	Ae	1	1,000	5	0.020
Ba	Bb	1	0	15	0.005	Bb	Ba	1	0	15	0.005
Bb	Bf	1	0	15	0.005	Bf	Bb	1	0	15	0.005
Bf	Df	1	1,000	5	0.020	Df	Bf	1	1,000	5	0.020
Cc	Cd	1	0	15	0.005	Cd	Cc	1	0	15	0.005
Cd	Cg	1	0	15	0.005	Cg	Cd	1	0	15	0.005
Cg	Dg	1	1,000	5	0.020	Dg	Cg	1	1,000	5	0.020
De	Df	1	0	15	0.005	Df	De	1	0	15	0.005
De	Dg	1	0	15	0.005	Dg	De	1	0	15	0.005
Df	Dg	1	0	15	0.005	Dg	Df	1	0	15	0.005

Note. These data all correspond to the “do-nothing” defense option d_0 .

TABLE A.2. Arc data for “hardening” defenses on Königsberg bridges.

Arc tail i	Arc head j	$c_{ij}^{d_1}$	$q_{ij}^{d_1}$	$\alpha_{ij}^{d_1}$	$\beta_{ij}^{d_1}$	Arc tail i	Arc head j	$c_{ij}^{d_1}$	$q_{ij}^{d_1}$	$\alpha_{ij}^{d_1}$	$\beta_{ij}^{d_1}$
Aa	Ba	1	0	5	0.02	Ba	Aa	1	0	5	0.02
Ab	Bb	1	0	5	0.02	Bb	Ab	1	0	5	0.02
Ac	Cc	1	0	5	0.02	Cc	Ac	1	0	5	0.02
Ad	Cd	1	0	5	0.02	Cd	Ad	1	0	5	0.02
Ae	De	1	0	5	0.02	De	Ae	1	0	5	0.02
Bf	Df	1	0	5	0.02	Df	Bf	1	0	5	0.02
Cg	Dg	1	0	5	0.02	Dg	Cg	1	0	5	0.02

Notes. (See results in Table 4.) These data represent defense option d_1 , which applies only to bridges, initially.

TABLE A.3. Arc data added to base case for upgrading condition of Königsberg roads.

Arc tail i	Arc head j	$c_{ij}^{d_1}$	$q_{ij}^{d_1}$	$\alpha_{ij}^{d_1}$	$\beta_{ij}^{d_1}$	Arc tail i	Arc head j	$c_{ij}^{d_1}$	$q_{ij}^{d_1}$	$\alpha_{ij}^{d_1}$	$\beta_{ij}^{d_1}$
Ba	Bb	1	0	10	0.001	Bb	Ba	1	0	10	0.001
Bb	Bf	1	0	10	0.001	Bf	Bb	1	0	10	0.001
Cc	Cd	1	0	10	0.001	Cd	Cc	1	0	10	0.001
Cd	Cg	1	0	10	0.001	Cg	Cd	1	0	10	0.001

Notes. (See results in Table 5.) These data correspond to defense option d_1 .

TABLE A.4. Arc data for adding a new, invulnerable bridge to Königsberg.

Arc tail i	Arc head j	$c_{ij}^{d_0}$	$q_{ij}^{d_0}$	$\alpha_{ij}^{d_0}$	$\beta_{ij}^{d_0}$	Arc tail i	Arc head j	$c_{ij}^{d_0}$	$q_{ij}^{d_0}$	$\alpha_{ij}^{d_0}$	$\beta_{ij}^{d_0}$
Ba	Cc	3	0	5	0.01	Cc	Ba	3	0	5	0.01

Notes. (See results in Table 6.) This bridge is invulnerable to attack, so only defense option d_0 applies.

References

- [1] T. Abrahamsson, and L. Lundqvist. Formulation and estimation of combined network equilibrium models with applications to Stockholm. *Transportation Science* 33:80–100, 1999.
- [2] R. Albert, I. Albert, and G. Nakarado. Structural vulnerability of the North American power grid. *Physics Review E* 69(025103(R)):1–4, 2004.
- [3] R. Albert, H. Jeong, and A. Barabási. Attack and error tolerance of complex networks. *Nature* 406:378–382, 2000.
- [4] B. C. Arntzen, G. G. Brown, T. P. Harrison, and L. L. Trafton. Global supply chain management at Digital Equipment Corporation. *Interfaces* 25:69–93, 1995.
- [5] ASME Innovative Technologies Institute. RAMCAP, risk analysis and management for critical asset protection. <http://www.asme-iti.org/RAMCAP/>. Accessed August 8, 2008.
- [6] M. J. Beckmann. On the theory of traffic flows in networks. *Traffic Quarterly* 21:109–116, 1967.
- [7] R. Bhave and R. Gupta. *Analysis of Water Distribution Networks*. Alpha Science International, Ltd., Oxford, UK, 2006.
- [8] V. M. Bier. Choosing what to protect. *Risk Analysis* 27:607–620, 2007.
- [9] V. M. Bier, N. Haphuriwat, J. Menoyo, R. Zimmerman, and A. M. Culpen. Optimal resource allocation for defense of targets based on differing measures of attractiveness. *Risk Analysis* 8:763–770, 2008.
- [10] D. Boyce and H. Bar-Gera. Validation of multiclass urban travel forecasting models combining origin-destination, mode, and route choices. *Journal of Regional Science* 43:517–540, 2003.
- [11] G. Brown and A. Cox. How probabilistic risk assessment can mislead terrorism risk analysis. *Risk Analysis* 31:196–204, 2011.
- [12] G. Brown and A. Cox. Making terrorism risk analysis less harmful and more useful: Another try. *Risk Analysis*. 31(2):193–195, 2011.
- [13] G. Brown and R. Dell. Formulating integer and linear programs: A Rogues’ gallery. *INFORMS Transactions on Education* 7(2):153–159, 2007.
- [14] G. Brown, M. Carlyle, and K. Wood. Optimizing Department of Homeland Security Defense investments: Applying defender-attacker (-defender) optimization to terror risk assessment and mitigation. Appendix E of “Department of Homeland Security Bioterrorism Risk Assessment—A Call for Change,” National Research Council, Washington, DC, 90–102, 2008.
- [15] G. Brown, M. Carlyle, J. Salmerón, and K. Wood. Analyzing the vulnerability of critical infrastructure to attack, and planning defenses. H. Greenberg, and J. Smith, eds. *Tutorials in Operations Research: Emerging Theory, Methods, and Applications*. Institute for Operations Research and Management Science, Hanover, MD, 102–123, 2005.
- [16] G. Brown, M. Carlyle, J. Salmerón, and K. Wood. Defending critical infrastructure. *Interfaces* 36:530–544, 2006.
- [17] G. Brown, J. Keegan, B. Vigus, and K. Wood. The Kellogg Company optimizes production, inventory, and distribution. *Interfaces* 31:1–15, 2001.
- [18] D. P. Chassin and C. Posse. Evaluating North American electric grid reliability using the Barabasi-Albert network model. *Physica A* 355(April):667–677, 2005.
- [19] R. L. Church and M. P. Scaparra. Protecting critical assets: The r -interdiction median problem with fortification. *Geographical Analysis* 39:129–146, 2007.
- [20] R. L. Church, M. P. Scaparra, and R. Middleton. Identifying critical infrastructure: The median and covering facility interdiction problems. *Annals of the AAG* 9:491–502, 2004.
- [21] R. Cohen, K. Erez, D. ben-Avraham, and S. Havlin. Resilience of the Internet to random breakdowns. *Physical Review Letters* 81:4626–4628, 2000.
- [22] M. Collins, L. Cooper, R. Helgason, J. Kennington, and L. LeBlanc. Solving the pipe network analysis problem using optimization techniques. *Management Science* 24:747–760, 1978.
- [23] K. J. Cormican, D. P. Morton, and R. K. Wood. Stochastic network interdiction. *Operations Research* 46:184–197, 1998.
- [24] J. Correa and N. Stier-Moses. Wardrop equilibria. J. Cochran, ed. *Encyclopedia of Operations Research and Management Science*. Wiley, Hoboken, NJ, Forthcoming, 2011.
- [25] A. Cox. Some limitations of “Risk = Threat \times Vulnerability \times Consequence” for risk analysis of terrorist attacks. *Risk Analysis* 28:1749–1761, 2008.

- [26] A. Cox. Improving risk-based decision making for terrorism applications. *Risk Analysis* 29:336–341, 2009.
- [27] K. Cragin and S. Daly. *The Dynamic Terrorist Threat: An Assessment of Group Motivations and Capabilities in a Changing World*. RAND Corporation, Santa Monica, CA, 2004.
- [28] J. Danskin. The theory of max–min, with applications. *SIAM Journal on Applied Mathematics* 14:641–664, 1966.
- [29] Department of Homeland Security (DHS). National infrastructure protection plan. Department of Homeland Security, Washington, DC, 2009.
- [30] B. C. Ezell, S. Bennett, D. von Winterfeldt, J. Sokolowski, and A. Collins. Probabilistic risk analysis and terrorism risk. *Risk Analysis* 30:575–589, 2010.
- [31] L. Euler. Solutio problematis ad geometriam situs pertinentis. *Novi Commentarii Academiae Scientiarum Imperialis Petropolitanae* 7:9-28, 1758. (English translation available from N. Biggs, E. Lloyd, and R. Wilson. *Graph Theory: 1736–1936*, Clarendon Press, Oxford, UK, 3–8, 1976.)
- [32] M. Florian and S. Nguyen. An application and validation of equilibrium trip assignment methods. *Transportation Science* 10:374–389, 1976.
- [33] B. Fortz and M. Labbé. Polyhedral approaches to the design of survivable networks. M. G. C. Resende and P. M. Pardalos, eds. *Handbook of Optimization In Telecommunications*. Springer, New York, 367–389, 2006.
- [34] D. R. Fulkerson and G. Harding. Maximizing the minimum source-sink path subject to a budget constraint. *Mathematical Programming* 13:116–118, 1977.
- [35] GAMS. <http://www.gams.com>. Accessed December 21, 2010.
- [36] B. Garrick, J. Hall, M. Kilger, J. McDonald, T. O’Toole, P. Probst, E. Parker, R. Rosenthal, A. Trivelpiece, L. Van Arsdale, and E. Zebroski. Confronting the risks of terrorism: Making the right decisions. *Reliability Engineering and System Safety* 86:129–176, 2004.
- [37] D. C. Gazis. *Traffic Theory*. Kluwer Academic Publishers, Boston, 2002.
- [38] A. M. Geoffrion and G. W. Graves. Multicommodity distribution system design by Benders decomposition. *Management Science* 20:822–844, 1974.
- [39] B. Golany, E. Kaplan, A. Marmur, and U. Rothblum. Nature plays with dice—Terrorists do not: Allocating resources to counter strategic versus probabilistic risks. *European Journal of Operations Research* 192:198–208, 2009.
- [40] B. Golden. A problem in network interdiction. *Naval Research Logistics Quarterly* 25:711–713, 1978.
- [41] M. Grötschel, C. L. Monma, and M. Stoer. Design of survivable networks. M. O. Ball, T. L. Magnanti, C. L. Monma, and G. L. Nemhauser, eds. *Network Models. Handbooks in Operations Research and Management Science*, Vol. 7. Elsevier, Amsterdam, 617–669, 1995.
- [42] T. Harris and F. Ross. Fundamentals of a method for evaluating rail net capacities. Research Memorandum RM-1573, RAND Corporation, Santa Monica, CA, 1955.
- [43] Homeland Security Council. National strategy for homeland security. Homeland Security Council, The White House, Washington, DC, 2007.
- [44] IBM. CPLEX 12.02. <http://www.ibm.com>. Accessed February 1, 2011.
- [45] E. Israeli and R. K. Wood. Shortest-path network interdiction. *Networks* 40:97–111, 2002.
- [46] H. Kerivin and A. R. Mahjoub. Design of survivable networks: A survey. *Networks* 46:1–21, 2005.
- [47] M. Kraitchik. *Mathematical Recreations*. Norton, New York, 1942.
- [48] L. LeBlanc, E. Morlok, and W. Pierskalla. An efficient approach to solving the road network equilibrium traffic assignment problem. *Transportation Research* 9:309–318, 1975.
- [49] T. G. Lewis. *Critical Infrastructure Protection in Homeland Security*. Wiley-Interscience, Hoboken, NJ, 2006.
- [50] C. Lim and J. C. Smith. Algorithms for discrete and continuous multicommodity flow network interdiction problems. *IIE Transactions* 39:15–26, 2007.
- [51] G. F. Lucio, M. Paredes-Farrera, E. Jammeh, M. Fleury, and M. J. Reed. OPNET modeler and ns-2: Comparing the accuracy of network simulators for packet-level analysis using a network testbed. *WSEAS Transactions on Computers* 2:700–707, 2003.

- [52] F. L. Mannering, S. S. Washburn, W. P. Kilareski. *Principles of Highway Engineering and Traffic Analysis*. John Wiley & Sons, Hoboken, NJ, 2009.
- [53] G. R. Mateus and Z. K. G. Patrocínio. Optimization issues in distribution network design. M. G. C. Resende and P. M. Pardalos, eds. *Handbook of Optimization in Telecommunications*. Springer, New York, 341–366, 2006.
- [54] J. Moore and J. Bard. The mixed integer linear bilevel programming problem. *Operations Research* 38:911–921, 1990.
- [55] D. P. Morton, F. Pan, and K. J. Saeger. Models for nuclear smuggling interdiction. *IIE Transactions* 39:3–14, 2007.
- [56] National Research Council (NRC). *Department of Homeland Security Bioterrorist Risk Assessment: A Call for Change*. NRC, Washington, DC, 2008.
- [57] National Research Council (NRC). *Review of the Department of Homeland Security’s Approach to Risk Analysis*. NRC, Washington, DC, 2010.
- [58] G. Parnell, R. Liebe, R. Dillon-Merrill, D. Buede, J. Scouras, B. Colletti, M. Cummings et al. *Homeland Security Risk Assessment: Volume I—An Illustrative Framework and Volume II—Appendices of Methods*. Homeland Security Institute, Washington, DC, 2005.
- [59] M. Paté-Cornell and S. Guikema. Probabilistic modeling of terrorist threats: A systems analysis approach to setting priorities among countermeasures. *Military Operations Research* 7:5–23, 2002.
- [60] E. R. Petersen. A primal-dual traffic assignment algorithm. *Management Science* 22:87–95, 1975.
- [61] C. Phillips. The network inhibition problem. *Proceedings of the 25th Annual ACM Symposium on the Theory of Computing*, 776–785, 1993.
- [62] H. D. Ratliff, G. T. Sicilia, S. H. Lubore. Finding the n most vital links in flow networks. *Management Science* 21:531–539, 1975.
- [63] J. Salmerón, R. Baldick, and K. Wood. Worst-case interdiction, and Defense, of large-scale electrical power grids. *20th International Symposium on Mathematical Programming, Chicago*, Mathematical Programming Society, 2009.
- [64] J. Salmerón, K. Wood, and R. Baldick. Analysis of electric grid security under terrorist threat. *IEEE Transactions on Power Systems* 19:905–912, 2004.
- [65] J. Salmerón, K. Wood, and R. Baldick. Worst-case interdiction analysis of large-scale electric power grids. *IEEE Transactions on Power Systems* 24:96–104, 2009.
- [66] M. P. Scaparra and R. L. Church. A bilevel mixed-integer program for critical infrastructure protection planning. *Computers & Operations Research* 35:1905–1923, 2008.
- [67] A. Schrijver. On the history of the transportation and maximum flow problems. *Mathematical Programming, Series B* 91:437–445, 2002.
- [68] J. C. Smith. Basic interdiction models. J. Cochran, ed. *Encyclopedia of Operations Research and Management Science*. Wiley, Hoboken, NJ. Forthcoming, 2011.
- [69] J. C. Smith, C. Lim, and F. Sudargho. Survivable network design under optimal and heuristic interdiction scenarios. *Journal of Global Optimization* 38:181–199, 2007.
- [70] F. Steinhäusler, P. Furthner, W. Heidegger, S. Rydell, and L. Zaitseva. Security risks to the oil and gas industry: Terrorist capabilities. *Strategic Insights* 7(1), 2008. Available at <http://www.nps.edu/Academics/centers/ccp/publications/OnlineJournal/2008/Feb/steinhauslerFeb08.html>.
- [71] U.S. Government. Critical Infrastructures Protection Act of 2001. Title 42, Chapter 68, Subchapter IV-B§5195c. Also known as Section 1016(e) of the USA PATRIOT Act of 2001. U.S. Government, Washington, DC, 2001.
- [72] H. J. Van Zuylen and L. G. Willumsen. The most likely trip matrix estimated from traffic counts. *Transportation Research Part B: Methodological* 14:281–293, 1980.
- [73] H. von Stackelberg. *The Theory of the Market Economy*. William Hodge, London, 1952.
- [74] J. Wang and L. Rong. Cascade-based attack vulnerability on the US power grid. *Safety Science* 47:1332–1336, 2009.
- [75] J. G. Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers Part II* 1:325–378, 1952.
- [76] H. Willis. Guiding resource allocations based on terrorism risk. *Risk Analysis* 27:597–606, 2007.
- [77] R. Wollmer. Removing arcs from a network. *Operations Research* 12:934–940, 1964.

- [78] R. K. Wood. Deterministic network interdiction. *Mathematical and Computer Modelling* 17:1–18, 1993.
- [79] R. K. Wood. Bilevel network-interdiction models: Formulations and solutions. J. Cochran, ed. *Encyclopedia of Operations Research and Management Science*. Wiley, Hoboken, NJ, Forthcoming, 2011.
- [80] G. Zakeri, A. B. Philpott, and D. M. Ryan. Inexact cuts in benders decomposition. *SIAM Journal on Optimization* 10:643–657, 1999.
- [81] D. Zaragoza and C. Belo. Experimental validation of the ON-OFF packet-level model for IP traffic. *Computer Communications* 30:975–989, 2007.