TUTORIALS IN
OPERATIONS RESEARCH
INFORMS 2014

© 2014 INFORMS | ISBN 978-0-9843378-5-9
http://dx.doi.org/10.1287/educ.2014.0131

Assessing and Improving Operational Resilience of Critical Infrastructures and Other Systems

David L. Alderson, Gerald G. Brown, W. Matthew Carlyle
Operations Research Department, Naval Postgraduate School, Monterey, California 93943
{dlalders@nps.edu, gbrown@nps.edu, mcarlyle@nps.edu}

Abstract In this tutorial, we quantify resilience for an infrastructure system to a set of disruptive events in terms of degradation of system function. We show how to build and solve a sequence of models to assess and improve the resilience of an infrastructure system to those disruptions. Through simple examples and real-world case studies, we provide motivation, details of the models, and solution algorithms.

Keywords infrastructure defense; resilience; bi-level optimization; tri-level optimization; Stakelberg game; optimization; operational model; attacker model; defender model

1. Introduction

In the last 15 years, a number of disasters, some deliberately caused, some not, have inflicted serious losses of human life, significant damage to property, and massive interruptions in service for a number of large infrastructure systems. In the wake of these events, the concept of *resilience* is now frequently used to characterize how well these infrastructure systems, their operators, and their users react to and recover from disruptive events. However, there are many different facets of resilience (e.g., Zolli and Healy [96]), and much of the literature on infrastructure resilience is qualitative in nature and does not suggest how the resilience of real systems can be improved (e.g., Haimes et al. [47], Madni and Jackson [61], Park et al. [68]). This tutorial provides a guide to recent work that applies constrained optimization (combined with models of system function and management) to assess and improve the resilience of critical infrastructure systems to disruptive events.

1.1. Infrastructure, Risk, and Resilience: An Abbreviated History

An early example of the importance of infrastructure to society is Ancient Rome, where the development of roads and aqueducts enabled unprecedented economic prosperity that lasted for hundreds of years. Over time, these infrastructure systems went from being conveniences to necessities upon which citizens and government depended, and their vulnerability to deliberate attack contributed to Rome's ultimate decline (Assante [11]).

The mid-1990s and the explosive growth of the Internet started a modern phase in the study of infrastructure systems because many formerly stand-alone physical systems—including transportation, energy, and water—quickly became interconnected via a common "central nervous system" that was not only a source of great efficiency but a new avenue of potential vulnerability. In particular, there was concern at the highest levels of the United States (U.S.) government (President's Commission on Critical Infrastructure Protection [69]) about the potential for large-scale disruptions that could cause serious harm to society (for a concise history, see Brown [26]).

The terrorist attacks of September 11, 2001 inflicted considerable impact on the infrastructure systems of the New York City metropolitan area, with ripple effects around the world. The primary response of the U.S. government to these attacks was to establish laws with a focus

on *security*, specifically apprehending those responsible and preventing future incidents. The term *critical infrastructure* was defined in Section 1016(e) of the USA PATRIOT Act of 2001 as "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, national economic security, national public health or safety, or any combination of those matters" (Title 42 U.S. Code, Section 5195c et seq. 2006 Supp. IV [87]). The Homeland Security Act of 2002 established the Department of Homeland Security (DHS) with its primary mission to "(A) prevent terrorist attacks within the United States; (B) reduce the vulnerability of the United States to terrorism; (C) minimize the damage, and assist in the recovery, from terrorist attacks that do occur within the United States" (Public Law 107-296, 117 Stat. 745 [70]).

However, the accidental electric power outage in the Northeastern United States in August 2003 demonstrated that the United States also needs to be concerned with nondeliberate events, such as technological failures and human accidents. In December 2003, the White House issued Homeland Security Presidential Directive (HSPD)-7: "Directive on Critical Infrastructure Identification, Prioritization, and Protection," which directed the use of *risk-based* strategies for assessing hazards and prioritizing investment (The White House [83]). DHS codified this guidance in the National Infrastructure Protection Plan (NIPP), first issued in 2006.

In the years following HSPD-7, there were a number of unprecedented natural disasters, including the Indonesian tsunami in December 2004, Hurricanes Irene and Katrina in Summer 2005, and the magnitude-7.6 earthquake in Pakistan in October 2005. The overwhelming need for emergency response to these events led to the following recognition in the 2007 National Strategy for Homeland Security (Homeland Security Council (HSC) [51]):

> We will not be able to deter all terrorist threats, and it is impossible to deter or prevent natural catastrophes. We can, however, mitigate the nation's vulnerability to acts of terrorism, other man-made threats, and natural disasters by ensuring the structural and operational resilience of our critical infrastructure and key resources. HSC [51, p. 27]

A number of unprecedented accidents recently, including the Deepwater Horizon oil spill in 2010 and the Fukushima Daiichi nuclear disaster in 2011, along with the devastation caused by Hurricane "Superstorm" Sandy in 2012, have reinforced the need for resilience of infrastructure systems and communities. And the April 2013 attack on the Pacific Gas and Electric Metcalf electricity substation in San Jose, California (Smith [79]) serves as a reminder that deliberate threats to infrastructure still persist.

Presidential Policy Directive (PPD)-21: "Critical Infrastructure Security and Resilience," signed in February 2013 (The White House [84]), defines *resilience* explicitly to mean "the ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions. Resilience includes the ability to withstand and recover from deliberate attacks, accidents, or naturally occurring threats or incidents." The most recent edition of the NIPP, released in December 2013 (Department of Homeland Security (DHS) [38]), now features the subtitle "Partnering for Critical Infrastructure Security and Resilience."

1.2. Goals of This Tutorial

Our goal is to show by example how to develop quantitative models to assess infrastructure resilience and, even more importantly, determine how to improve the operation of critical infrastructures and other systems in the presence of disruptive events.

An earlier tutorial (Brown et al. [22]) describes two classes of applications of bi-level programming models in the study of critical infrastructure systems. One class, called *attacker-defender* models, is used to analyze the vulnerability of infrastructure to worst-case attack. Another class of models, called *defender-attacker* models, is used to plan infrastructure defenses against a known adversary. These models can be combined to create tri-level problems,

called *defender-attacker-defender* models (Brown et al. [21]). Recent treatments of these and other *system interdiction models* can be found in Lim and Smith [59], Alderson et al. [7, 8], Wood [91], and Dimitrov and Morton [40].

This tutorial (1) synthesizes the most essential material in these many papers, (2) provides a step-by-step explanation of how and why we build these models as we do, (3) introduces a general solution technique for solving them, and (4) establishes connections to other related work. Throughout this tutorial, our focus is *operational resilience*—a term introduced in the 2007 National Strategy for Homeland Security (Homeland Security Council (HSC) [51]), but never formally defined—in which we restrict attention to the function, or operation, of infrastructure systems. We define operational resilience of an infrastructure system to a set of disruptive events in terms of the worst-case degradation of system function, and we show how to build and solve a sequence of models that allows us to assess and improve the operational resilience of an infrastructure system to those disruptions. We have used the concepts and techniques presented here for more than a decade and have applied them to a wide variety of critical infrastructures and other systems.

2. Terminology and Notation

We assume the reader is familiar with optimization models at the level of Rardin [71] and also network flow notation at the level of Ahuja et al. [1]. We often think of infrastructure systems as moving some commodity through time and space, and this makes them ideal subjects for network flow models. However, in our study of these systems, we find that there are always additional constraints on system operation, or a need to model multiple commodities, or nonlinearities in the system that require the use of something more sophisticated than a pure network flow model. Nonetheless, because underlying network structure is ubiquitous in these more complicated models, we retain the language of network models for its easily visualized function and indicate where we might need to extend the definitions or standard models as appropriate.

We define the following terms for use throughout. We treat an infrastructure as a system of interconnected *components* that work together to provide a particular *function*. Examples of function include traffic conveyance, electric power transmission, fuel delivery, manufacturing, supply chains, and communication. We often represent system function by using multicommodity network flow models of conveyance over space and/or time, where *nodes* and *arcs* represent parts of the system, and the flow of commodities represents the "function" provided by the infrastructure. For instance, if we are modeling the traffic flow function of a road network, then the commodities might be the traffic (i.e., vehicles) on the roads, partitioned based on their eventual destinations. The performance of the system might be a function of the number of vehicles, or the people in the vehicles, their economic value, or the carbon footprint of their travel, etc.

Our mathematical models frequently use a directed graph $G = (N, A)$ comprised of nodes and directed arcs to represent the connectivity (at the nodes) of the components (the arcs) in the system. Each component engages in one or more *activities*. Collectively, the activities of all components specify the *operation* of the system as a whole. We use arcs exclusively to model components with activities on them. If a component is more naturally associated with a node, such as a junction, but we need to model an activity, we use the standard node-splitting technique (see, for example, pp. 41–43 of Ahuja et al. [1]) to replace the original node with two surrogate nodes and a single arc representing the activity at the original node. We typically model activity at an arc by a decision variable. In a network flow model, this might be $y_{ij} \geq 0$ for each arc $(i, j) \in A$. One component could have multiple activities; for instance, there might be multiple commodities, indexed by $k \in K$, flowing through the same pipeline. Each would have its own decision variable, say, $y_{ij}^k \geq 0$. We use the term *operator* to refer to the decision-making entity who chooses activities in the system.

The activities in the system are limited by *operational constraints* that reflect, for example, the laws of physics, limited budgets, or dependencies between components. We typically express these as the constraints of a mathematical programming formulation. For example, a pipeline $(i,j) \in A$ might have a capacity, u_{ij}. One operational constraint could say that the total directional flow on the pipeline cannot exceed its capacity: $\sum_k y_{ij}^k \leq u_{ij}$.

A *design* refers to an existing or proposed nominal system configuration; this could include the selection of particular versions of arcs (changing their costs and capacities) or the addition of new arcs (i.e., adding new activities to the system). We model the choice of a design through a vector w of design decisions, one element for each component. A particular design is then specified as $\hat{w}$.

An *operational setting*, denoted $\hat{x}$, specifies the status of each component in the system, any environmental factors that can affect operations, and any exogenous supplies of or demands for the function provided by the system. An *event* is any change in the operational setting. If multiple elements of the operating conditions change at once, we still refer to this collective occurrence as "an event." However, when discussing the corresponding changes to individual components (e.g., simultaneous damage to two distinct pipes, each represented as its own component), we sometimes refer to "simultaneous events," depending on the context. We explicitly do *not* use the phrase "set of events," which has a different meaning, described below. Events can be deliberate or nondeliberate. Accidents and failures are examples of nondeliberate events, whereas attacks and repairs are examples of deliberate ones.

A *performance measure* evaluates how well a particular design of the system functions in a particular operational setting. We evaluate the performance of a system by solving an *operational model* that takes as input a specification of the operational setting and provides as output a specific set of activities and the corresponding performance for the resulting system operation. Operational models can be parameterized by the design of the system, to allow for the evaluation of alternative designs.

Optimizing system performance means using an operational model to determine a maximal-performance operation of the system. As a mathematical program, we often write the formulation as

$$z^* = \max_{y \in Y(\hat{w})} f(\hat{w}, \hat{x}, y),$$

where $f(\cdot)$ measures system performance and where $y \in Y(\hat{w})$ indicates that the feasible activities y depend on the design $\hat{w}$. Here, $y^* = \arg\max_{y \in Y(\hat{w})} f(\hat{w}, \hat{x}, y)$ is an optimal way to operate the system for design $\hat{w}$ under operational setting $\hat{x}$, and results in performance z^*.

The *consequence* of an event is the change in performance that results when the system operates in the modified setting. If the performance of the system becomes worse, we say the performance of the system has been degraded by the event. We use the consequence of an event to quantify the system's *operational resilience to an event*, because it quantifies the ability of the system to adapt its operations following the event. If the event results in a relatively large (respectively, small) consequence, we say that the system has relatively low (respectively, high) operational resilience to that event.

When studying infrastructure systems, our interests are generally twofold. First, we need to assess the operational resilience for the current (or proposed) design of a system. For executive decision makers, this typically amounts to answering the question "How bad can things get for this system?" The second thing we need to do is to identify budget-limited investments that can improve the operational resilience for this system. This essentially requires that we consider how best to (re-)design the system, answering the executive decision-maker's question "How should we spend our limited funds?"

3. Example: The Operational Resilience of a Rail Network

We illustrate the basic steps in assessing and improving operational resilience using a simple example, adapted from a historical case study on the Soviet rail network (Harris and Ross [49],

FIGURE 1. Network representation of the Soviet Rail system circa 1955 (from Alderson et al. [7]).

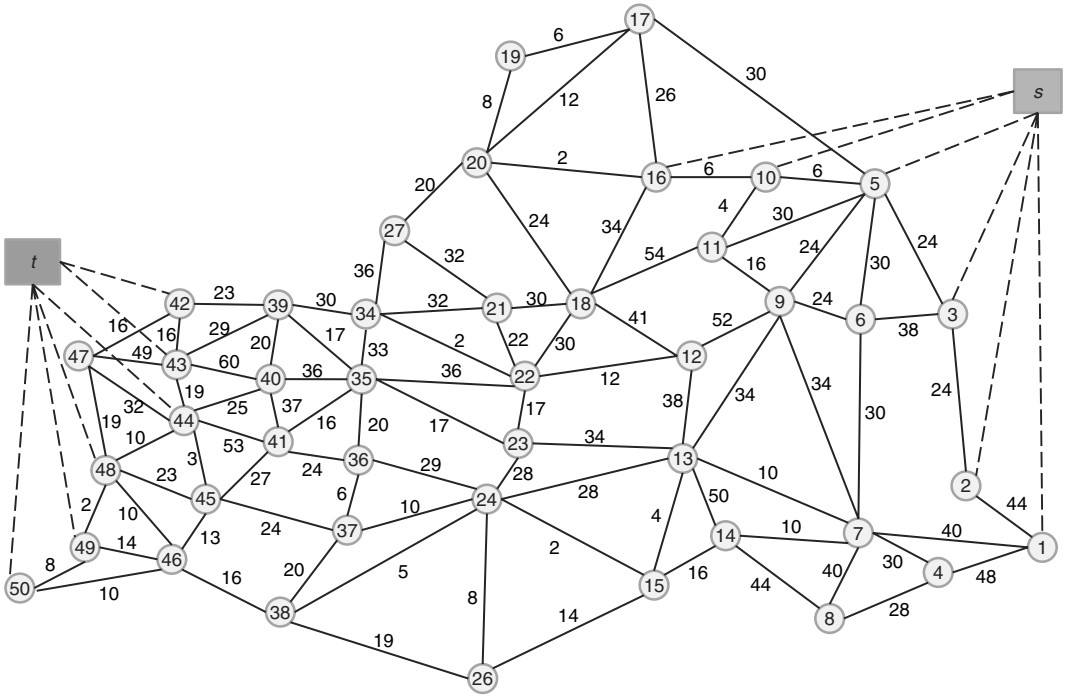

Notes. Each arc is annotated with its capacity (in 1,000s of tons). The maximum flow through the undisturbed network from node "s" to node "t" is 163,000 tons.

Schrijver [78], Alderson et al. [7]), in which the operator's goal is to move as much cargo as possible from one particular station to another (see Figure 1). In this figure, each line represents a section of track having a finite tonnage limit, and each circle represents a station where adjacent sections of track meet. The performance of the system is the maximum tons of cargo that can be moved (i.e., its *capacity*), and the concern is maintaining as much capacity as possible even in the face of an event that eliminates one or more sections of track from the system.

Following the seminal work of Harris and Ross [49], we model aggregate flows through the network in a single time period instead of using a detailed model of train scheduling, track siding usage, etc. In the original study, the key operational goal was to assess the potential for a rapid movement of materiel, and their model (and ours) has sufficient fidelity to represent this capacity.

Our modeling and analysis follows a script that covers each of several steps:

1. Formulate an operational model, the *operator model*, to determine the system activities and the corresponding performance of system operation.

2. Define the set of events that can disrupt the system, and identify how each event modifies the operational setting.

3. Modify the operator model to incorporate events and their impact on system operation.

4. Formulate a bi-level *attacker model* to identify worst-case events, minimizing best-case performance following a worst-case event.

5. Define design decisions that can change the system.

6. Modify the operator and attacker models to include the effect of any design on the operations.

7. Formulate a tri-level *defender model* to choose the best design in anticipation of a worst-case event.

The following subsections describe how to assess and improve the resilience of this rail network, and are numbered to correspond directly to the steps in our script above.

3.1. Formulate the Operator Model

The function of the system is source-to-destination cargo delivery, and the components of the system are the sections of track, modeled as the edges in a network. The activities on the components are the flows of cargo on the tracks, and the two sets of operational constraints are the single-edge capacities given by a tonnage limit on each section of track and materiel balance constraints, one for each station (represented by a node) in the network. The performance of the system is measured by the total amount of cargo that flows from the start node to the destination node (all capacities are measured *per day*). We formulate the problem as follows.

Indices and Sets

$n, i, j \in N$ stations (ordered set of nodes);
$s, t \in N$ distinguished start and end stations;
$[i, j] \in E$ undirected edge between nodes i and j; where $i < j, \forall\, [i, j] \in E$;
$(i, j) \in A$ directed arc from i to node j; $[i, j] \in E \Leftrightarrow i < j \wedge ((i, j) \in A \wedge (j, i) \in A)$ (for every undirected edge there is a pair of antiparallel directed arcs between the same pair of nodes).

Data [*units*]

u_{ij} upper bound on total (undirected) flow on edge $[i, j] \in E$ [tons].

Decision variables [*units*]

y_{ij} directional flow of cargo on arc $(i, j) \in A$ [tons];
y_{ts} total flow through network from s to t [tons].

RAIL-NET-CAPACITY

$$\max_{y} \; y_{ts} \tag{1}$$

$$\text{s.t.} \sum_{j:\,(n,j)\in A} y_{nj} - \sum_{i:\,(i,n)\in A} y_{in} = \begin{cases} y_{ts} & n = s, \\ 0 & n \neq s, t, \\ -y_{ts} & n = t, \end{cases} \qquad \forall\, n \in N; \tag{2}$$

$$y_{ij} + y_{ji} \leq u_{ij} \quad \forall\, [i, j] \in E; \tag{3}$$

$$y_{ij} \geq 0 \quad \forall\, (i, j) \in A; \tag{4}$$

$$y_{ts} \geq 0. \tag{5}$$

Discussion. We model this system as a maximum flow problem on a directed graph $G = (N, A)$, where each (undirected) edge $[i, j] \in E$ represents a section of track with an overall flow capacity, u_{ij}, the maximum tonnage of cargo that can move over that track. We model directional flow along each track section using a pair of antiparallel arcs. Each node $n \in N$ represents a station where track sections meet, and where cargo can be redirected but cannot accumulate. We have a distinguished node, s, representing the source station, and a distinguished node, t, representing the destination station.

The activity, y_{ij} on each track section is simply the tons of cargo shipped on that track from i to j. The operation of this system is the selection of these activities. The flow variable y_{ts} is not associated with any actual track section; it represents an artificial arc from t back to s, and the constraints in the model ensure that its flow value represents the total tonnage of cargo that flows from s to t through the network. We determine the capacity of this rail network by operating the system to maximize y_{ts} without exceeding the capacity of any track in the network and while maintaining balance of flow at every station.

There are two sets of operational constraints. Constraints (2) say that any cargo that flows into a station has to flow out, possibly using the artificial flow variable y_{ts}, which is therefore always set equal to the net flow of cargo out of s (and the net flow into t). Constraints (3)–(5) restrict the directional cargo flows on each section of track to be nonnegative, and also restrict the total flow to be no larger than that section's tonnage limit.

We refer to the operational model *RAIL-NET-CAPACITY* as an operator model, because it represents the system from the point of view of the operator of that system, who has a vested interest in its performance. This formulation is valid for any static configuration of the rail network, but if there are events that can modify the operational setting, we need to update the model to take these into account.

3.2. Define the Events

For this rail system, the events of concern to us involve the simultaneous damage of one or more edges. Damage to a node can be captured in the same way that we would model activity at a node, through the use of "node splitting." To represent an event, we use a binary vector $\hat{x}$ with an element $\hat{x}_{ij}$ for each edge indicating whether or not that section of track has been damaged. Subsequently, there are many ways to represent the possible events of concern. We can enumerate a set of, say, p, possible events, $S = \{\hat{x}^1, \hat{x}^2, \ldots, \hat{x}^p\}$, or we could define the set of possible events S using constraints on the elements of $\hat{x}$:

$$S = \left\{ \hat{x} \colon \hat{x} \in \{0,1\}^{|E|}, \sum_{(i,j)\in A} \hat{x}_{ij} \le atk\_budget \right\},$$

where in this example we have defined S to be the set of all events consisting of combinations of no more than *atk_budget* damaged edges. We could use any number of constraints to define S in this way, including constraints that render pairs of edges to be mutually exclusive in an attack, or constraints that bind edges together because of their proximity.

3.3. Incorporate Events Into the Operator Model

If a section of track is damaged, i.e., $\hat{x}_{ij} = 1$, we assume that it is unavailable to carry cargo, and that the operator will not get any benefit from using that track segment. In other applications, it might make sense to consider partial damage to an edge in the system, but we restrict attention to the simpler case here.

One way to incorporate a change in the availability of a track segment is to modify its capacity, specifically, replace (3) with the following:

$$y_{ij} + y_{ji} \le (1 - \hat{x}_{ij}) u_{ij} \quad \forall\, [i,j] \in E.$$

Thus, when $\hat{x}_{ij} = 1$, the capacity on edge $[i,j]$ is zero. Although this type of modification is natural, it can lead to computational complications if the incumbent solution to system operation suddenly becomes infeasible.

In our experience, there is a better way to account for this change in the operational setting. Instead of having the $\hat{x}$ directly affect the capacities of the track sections, we leave (3) unchanged and modify the objective function to penalize any flow across an attacked track section so that it will have a negative impact on the performance measure. Specifically, we rewrite the objective function (1) as follows:

$$\max_{y} \; y_{ts} - \sum_{[i,j]\in E} 2(y_{ij} + y_{ji})\hat{x}_{ij}.$$

Any flow (y_{ij} or y_{ji}) across a damaged edge (i.e., having $\hat{x}_{ij} = 1$) can potentially contribute to the total flow value, y_{ts}, but *twice* that flow value will be subtracted from the performance

measure, resulting in a net loss of cargo getting through the network. (Any scalar coefficient greater than or equal to 1 will work here; using a penalty of 1 admits alternative optimal solutions with flows across attacked arcs. We use a coefficient of 2 because it is an integer and keeps the objective coefficients simple.) With this modified objective function it is clearly always better to operate the system by not moving cargo over a track segment $[i, j]$ with $\hat{x}_{ij} = 1$,. Thus, the operational setting (which reflects the availability of the track sections through $\hat{x}$) is directly accounted for in the operator model by parameterizing the performance measure of the system by the event that occurs. This use of *cost-based interdiction* is perhaps less natural but computationally important.

3.4. Formulate the Attacker Model

To assess the extent to which system operation is resilient to the events in the set S, we must be able to identify the event(s) in S that reduce(s) the capacity of the system to the lowest possible point.

If the number of elements in S is small, then finding the worst one can be achieved by a simple enumeration, where for each event $\hat{x} \in S$ we must solve the operator model explicitly. The amount of time required to find the worst event is then directly proportional to the size of S. In practice, finding the worst event by exhaustive enumeration is impractical when the number of elements in S is large, and it is impossible if the set S is infinite in size.

An alternate, and intuitively appealing, means of identifying the worst-case event in S is to consider the *decision* of a hypothetical, intelligent adversary who deliberately selects the event in S that disrupts system function the most. We refer to this adversary as the *attacker*, and this leads us to the attacker model, an optimization model that takes the modified operator model and uses decision variables to identify the "worst-case" event $x^* \in S$. In the case of the Soviet rail system, this means identifying the track sections to damage in order to minimize the residual capacity of the system, operated as best as possible after attack. For simplicity, we hereafter use the term *attack* synonymously with *event* even if it was not deliberate, and we refer to an individual damaged arc as a *target* of the attack.

We formulate the attacker model as follows.

New data [*units*]
atk_budget maximum number of track sections targeted in an attack [cardinality].

New decision variables [*units*]
x_{ij} = 1 if track section $[i, j] \in E$ is attacked, =0 otherwise [binary].

ATTACK-RAIL-NET

$$\min_{x} \max_{y} \; y_{ts} - \sum_{[i,j]\in E} 2(y_{ij} + y_{ji})x_{ij} \tag{6}$$

$$\text{s.t.} \;\; (2), (3), (4), (5),$$

$$\sum_{[i,j]\in E} x_{ij} \le atk\_budget, \tag{7}$$

$$x_{ij} \in \{0, 1\} \quad \forall [i, j] \in E. \tag{8}$$

Discussion. Each variable x_{ij} represents an attacker's decision to target edge $[i, j]$ in an attack. The simple cardinality constraint (7) is a surrogate for any set of restrictions we might impose on the attacker's choice of targets. For example, we could include budget constraints faced by an attacker, logical constraints that prohibit infeasible combinations of targets, or constraints that require collocated targets to always be in the same attack. Here, different choices for *atk_budget* can be interpreted as different attacker *capabilities* (with larger values of *atk_budget* being associated with greater capability), and we can evaluate the consequences of a worst-case attack for a *range* of attacker capabilities by solving this model several times.

3.5. Define the Design Decisions

A natural and mathematically straightforward way to improve the resilience of our train network is to add new sections of track that provide alternative routes (with new capacities) to the operator.

If we use a vector $\hat{w}$ to represent the track sections that are built (i.e., $\hat{w}_{ij} = 1$), then, in much the same way that we defined the set S, the set of all possible designs we consider can be defined either through explicit enumeration, $\Delta = \{\hat{w}^1, \hat{w}^2, \ldots, \hat{w}^q\}$, or implicitly, perhaps through a set of constraints:

$$\Delta = \left\{ \hat{w}\colon \hat{w} \in \{0,1\}^{|E|}, \sum_{[i,j]\in E} \mathit{def\_cost}_{ij}\hat{w}_{ij} \le \mathit{def\_budget} \right\},$$

where $\mathit{def\_cost}_{ij}$ represents the cost to build track section $[i,j] \in E$, where the sets E and A have been extended to include all existing *and potential* track sections, and where *def_budget* represents an overall defense budget available to the defender of the rail system. Existing track sections can have $\mathit{def\_cost}_{ij} = 0$, and then all original sections will be available, and only new sections will count against the construction budget.

We could also model design decisions that defend existing sections to make attacks targeting them less effective, or completely ineffective. Because our overall design motivation is to improve the performance of the system in the presence of disruptions, we refer to a particular $\hat{w}$ as a *defense*.

3.6. Incorporating Design Decisions Into the Models

We can modify the capacity constraints (3) in the operator model (and in the attacker model) to only allow cargo to be transported across track that has been built:

$$y_{ij} + y_{ji} \le u_{ij}\hat{w}_{ij} \quad \forall [i,j] \in E.$$

Here we can use a capacity control because the design and operational variables are both controlled by the same "side," namely, the defender and operator, who both wish to maximize flow. This is in contrast to the attacker, who wishes to *minimize* flow. Now we can evaluate any particular design $\hat{w} \in \Delta$ by solving the modified attacker model for that particular design, and then determining the resulting capacity after the attacker chooses the worst-case attack from S *for that design*.

3.7. Formulating the Defender Model

We want to model a defender (for example, the owner of the rail network) who wishes to spend a limited construction budget to build new track so that the rail network that results from these design decisions (in Δ) will have the maximum residual capacity after a worst-case attack (from S) occurs. We formulate this defender model as follows.

New data [*units*]
def_budget defense construction budget [$];
$\mathit{def\_cost}_{ij}$ defense construction cost of track section $[i,j] \in E$ [$].

New decision variables [*units*]
$w_{ij} = 1$ if we decide to build track section $[i,j] \in E$, =0 otherwise [binary].

DEFEND-RAIL-NET

$$\max_{w} \min_{x} \max_{y} \; y_{ts} - \sum_{[i,j]\in E} 2(y_{ij} + y_{ji})x_{ij} \tag{9}$$

$$\text{s.t. } (2), (4), (5), (7), (8),$$

$$y_{ij} + y_{ji} \le u_{ij}\hat{w}_{ij} \quad \forall [i,j] \in E, \tag{10}$$

$$\sum_{[i,j]\in E} def\_cost_{ij}\hat{w}_{ij} \le def\_budget, \tag{11}$$

$$w_{ij} \in \{0,1\} \quad \forall [i,j] \in E. \tag{12}$$

Discussion. Although the design decisions do not directly affect the objective function (9) in the defender model, they enable flows through new components via constraints (10). Constraints (11) ensure that the defender only builds track sections that he can afford. If the track sections in the original network have construction costs $def\_cost_{ij} = 0$, then the defender can build them all at no cost, and only spends his budget *def_budget* on new construction.

With the formulation of this defender model, we have completed all of the basic steps to assess and improve the resilience (at least in an instant of time) of an infrastructure system to a set of attacks. For any defender capability represented through the set Δ, whether through a construction budget *def_budget* or any more complicated set of restrictions on feasible designs, and for any attacker capability represented through the set S, we can find the system design that maximizes performance of the system after any attack in the set S.

Extension of DEFEND-RAIL-NET. If we are able to reinforce an existing track section so that an attack would be (essentially) ineffective, the model only requires a slight modification in which we represent the new (defended) version of the track section as a parallel edge with similar (but possibly modified) capacity, and with zero penalty on cargo flow on that edge's arc in the objective. To represent these parallel edges we introduce a new index, d, representing a set of *defense options* for each edge in the network. We define new data representing whether an edge is susceptible to attack, and reformulate as follows.

New indices and sets
$d \in D$ defense option.

New data [*units*]
v_{ij}^d vulnerability of defense option d for track section $[i,j] \in E$ [tons lost/tons shipped];
$v_{ij}^d = 2$ if an attack disables the track using defense d, and is zero otherwise;
u_{ij}^d capacity of track section $[i,j] \in E$ for defense option d [tons];
$def\_cost_{ij}^d$ construction cost of defense option d for track section $[i,j] \in E$ [$].

New decision variables [*units*]
y_{ij}^d cargo flowing across directed arc $(i,j) \in A$ with defense option d [tons];
w_{ij}^d =1 if we select defense option d for track section $[i,j] \in E$, =0 otherwise [binary].

DEFEND-RAIL-NET

$$\max_{w}\min_{x}\max_{y} \; y_{ts} - \sum_{[i,j]\in E}\sum_{d\in D}(v_{ij}^d y_{ij}^d + v_{ij}^d y_{ji}^d)x_{ij} \tag{13}$$

$$\text{s.t.} \sum_{d\in D}\Bigg[\sum_{j:\,(n,j)\in A} y_{nj}^d - \sum_{i:\,(i,n)\in A} y_{in}^d\Bigg] = \begin{cases} y_{ts} & n = s, \\ 0 & n \ne s,t, \\ -y_{ts} & n = t, \end{cases} \quad \forall n \in N; \tag{14}$$

$$(5), (7), (8);$$

$$y_{ij}^d + y_{ji}^d \le u_{ij}^d w_{ij}^d \quad \forall [i,j] \in E, d \in D; \tag{15}$$

$$y_{ij}^d \ge 0 \quad \forall (i,j) \in A, d \in D; \tag{16}$$

$$\sum_{d\in D}\sum_{[i,j]\in E} def\_cost_{ij}^d w_{ij}^d \le def\_budget; \tag{17}$$

$$\sum_{d\in D} w_{ij}^d = 1 \quad \forall [i,j] \in E; \tag{18}$$

$$w_{ij}^d \in \{0,1\} \quad \forall [i,j] \in E, d \in D. \tag{19}$$

FIGURE 2. *Resilience curves* showing (instantaneous) throughput as a function of the number of attacks (damaged track sections) for varying numbers of defended rail sections.

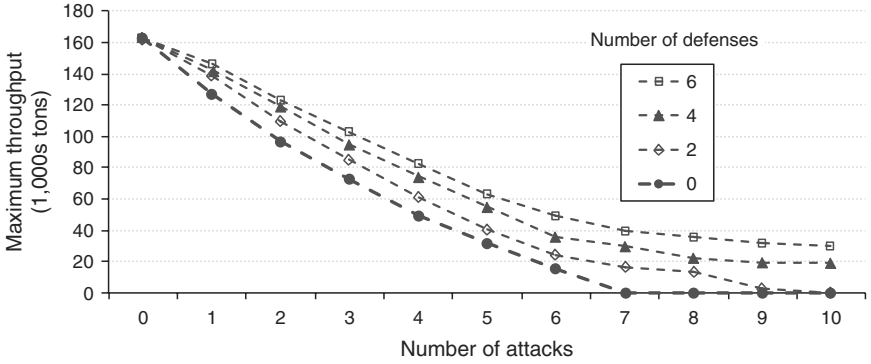

Notes. For example, four attacks with no defense reduces maximum throughput from 163,000 tons to about 50,000. Addition of six defenses raises throughput from 50,000 tons to just over 80,000. With no defense, seven attacks reduce the capacity of the system to zero. As the number of defended track sections increases, the rate of degradation inflicted by additional attacks is diminished.

Discussion. The objective function (13) is a generalization of (9) with attack vulnerability parameters v_{ij}^d. A track section design with $v_{ij}^d = 2$ penalizes system operation and will not be used if attacked (as before), but a track section with $v_{ij}^d = 0$ is effectively *invulnerable* because attacking it does not affect system operation. Constraint (18) forces exactly one design to be selected for each arc, and, as a consequence, constraints (14) simply maintain balance of flow at each node. In this model, defense decisions are made independently of each other (i.e., the defense decision for one track section does not affect the defense decision for another track segment). Constraints (15) enforce the upper limit on aggregate flow on each arc. Constraint (17) requires defensive investments to adhere to a general budget constraint.

3.8. Example Analysis

Alderson et al. [7] analyze the resilience of the Soviet rail example from Harris and Ross [49] by adding bidirectional flows on track sections and modifying the defensive options and attacks to influence both antiparallel arcs representing flow on a track section, as described above. The assumption is that following an attack, each targeted track section has zero capacity. Defensively, this analysis restricts attention to "hardening" of existing track sections so that they are invulnerable to attack.

Selected results from this analysis are repeated in Figure 2, which shows the capacity of the rail system for a varying number of attacks, and the improvements that result from a varying number of defenses. We refer to each individual line in Figure 2 as a *resilience curve* because it characterizes the performance of the system in response to disruptive events of different magnitude. In particular, it shows how well the system can respond to disruptive events of increasing magnitude.

The resilience curves in Figure 2 are similar in concept to earlier notions of resilience in physical science that characterize the amount of elastic deformation in a solid material (Park et al. [68]) or refer to resilience as the feature "that allows a system to return to its original form, position, or configuration after being bent, compressed or stretched" (Madni and Jackson [61], p. 185). In essence, a single curve in Figure 2 shows how well a particular design of the system performs in response to "being hit with different amounts of force."

From the perspective of our hypothetical attacker, the shape of each curve can also be interpreted as the return on investment (measured in terms of degraded system performance) for increasing the budget *atk_budget*. Specifically, we observe that when attacking an undefended system, the attacker reduces the capacity of the network approximately linearly

with each additional attack, and with seven attacks can reduce this capacity to zero (i.e., complete interdiction).

Conversely, the difference between individual curves in Figure 2 shows the improvement in resilience that comes with each additional defense, and thus it can be interpreted as a crude return on defensive investment. For example, with two defenses, 10 attacks are required to achieve complete interdiction. With six defenses, the system retains approximately 20% of its capacity in the presence of 10 attacks, no matter where they occur. This is an effective way to communicate with decision makers about the operational resilience of the existing system and how it can be improved.

We discuss additional insights available from these tri-level models, along with how to present them to decision makers, in §6.

4. Generalizations

The modeling technique presented here is very general and can be applied to almost any system with a well-defined performance measure and an appropriate operator model. Most broadly, we consider general *activities* that are constrained by limited *resources*, and we measure system performance in terms of an *operating cost* that is to be minimized. In such cases, we need to specify a cost for each possible activity in the system, and we include an incremental penalty cost for an activity that uses a damaged component. The form of the objective function will depend on the system itself.

4.1. General Cost-Based Formulation

The following formulation is a general representation of our tri-level model for a minimum cost problem.

Index use

$i \in I$	resource;
$j \in J$	activity;
$g \in G$	target group;
$j \in J_g \subseteq J$	set of activities in target group g;
$s \in S$	defense strategy;
$d \in D$	defense option;
$(j,d) \in \Theta_s$	set of activity j and defense option d pairs enabled by defense strategy s.

Data [*units*]

str_cost_s	fixed cost to adopt defense strategy s [\$];
$def\_cost_j^d$	fixed cost to adopt defense option d for activity j [\$];
def_budget	budget available for defense strategies and options [\$];
pen_attack_j	increased cost of activity j if it is attacked [\$/j-unit];
grp_cost_g	fixed cost to enable attacking target group g [\$];
atk_cost_j	cost of attack on activity j [\$];
atk_budget	maximum cost of attack groups and attacks [\$];
$act\_cost_j^d$	cost of activity j under defense option d [\$/j-unit];
a_{ij}	amount of resource i consumed by (operator) activity j [i-unit/j-unit];
b_i	available units of resource i [i-unit];
$\bar{x}_j^d$	binary indicator that activity j is vulnerable under defense option d;
$y_j^d, \bar{y}_j^d$	lower and upper bounds on activity j under defense option d [activity j].

Decision variables [*units*]

$q_s = 1$	if defender selects defense strategy $s \in S$, $= 0$ otherwise [binary];
$w_j^d = 1$	if defender selects defense option d for activity j, $= 0$ otherwise [binary];
$m_g = 1$	if attacker targets group g, $= 0$ otherwise [binary];

$x_j = 1$ if attacker attacks activity j, $= 0$ otherwise [binary];
$y_j^d \in Y_j^d$ level of activity j under defense option d [j-unit].

Formulation DEFEND-CRIT-INF

$$\min_{q,\,w}\ \max_{m,\,x}\ \min_{y}\ z = \sum_{j\in J,\, d\in D} (act\_cost_j^d + [pen\_attack_j x_j]_{\bar{x}_j^d=1}) y_j^d \tag{D0}$$

$$\text{s.t.} \sum_{s\in S} str\_cost_s q_s + \sum_{j\in J,\, d\in D} def\_cost_j^d w_j^d \le def\_budget; \tag{D1}$$

$$w_j^d \le \sum_{s\,|\,(j,\,d)\in\Theta_s} q_s \quad \forall\, j\in J,\, d\in D; \tag{D2}$$

$$\sum_{d\in D} w_j^d = 1 \quad \forall\, j\in J; \tag{D3}$$

$$x_j \le \sum_{g\,|\,j\in J_g} m_g \quad \forall\, j\in J; \tag{D4}$$

$$\sum_{g\in G} grp\_cost_g m_g + \sum_{j\in J} atk\_cost_j x_j \le atk\_budget; \tag{D5}$$

$$\sum_{j\in J,\, d\in D} a_{ij} y_j^d = b_i \quad \forall\, i\in I; \tag{D6}$$

$$\underline{y}_j^d w_j^d \le y_j^d \le \bar{y}_j^d w_j^d \quad \forall\, j\in J,\, d\in D; \tag{D7}$$

$$q_s \in \{0,1\} \quad \forall\, s\in S; \tag{D8}$$

$$w_j^d \in \{0,1\} \quad \forall\, j\in J,\, d\in D;$$

$$m_g \in \{0,1\} \quad \forall\, g\in G;$$

$$x_j \in \{0,1\} \quad \forall\, j\in J;$$

$$y_j^d \in Y_j^d \quad \forall\, j\in J,\, d\in D.$$

Discussion. The objective (D0) assesses the total cost of choosing defense strategies, and for each activity choosing a defense option and level of activity under that option. Some activities cost more because they have been attacked. Constraint (D1) limits the total cost of defense strategies and defense options chosen. Each constraint (D2) permits a defense option to be chosen for an activity only if a defense strategy has been chosen that enables such a selection. Each constraint (D3) requires that some single defense option be chosen for an activity. Each constraint (D4) permits an activity to be attacked only if a target group containing that activity has been attacked. Constraint (D5) limits the total cost of enabling target groups and attacking targets. Each constraint (D6) limits the sources and uses of a resource by activities. Each pair of constraints (D7) determines the domain limits for an activity under a defense option. Stipulations (D8) give domain limits for the decision variables, with the domain of operator activity variables y possibly being integral.

This formulation could also have been used for the Soviet rail example, albeit obscuring some of the special structure of that maximum-flow interdiction. The objective could also be nonlinear in some applications. For example, we have used a quadratic model of traffic congestion to determine optimal attacks and defenses in a municipal transportation network Alderson et al. [8], and we have used a piecewise linear approximation to higher-degree polynomial functions for a larger example Alderson et al. [5].

A defense strategy enables each of a set of activities to adopt its own defense option. An example of a defense strategy might be "approve new construction of security checkpoints and offset fencing," and the enabled set of activities and defense options might include "build a

checkpoint and offset fence for this activity." A defense strategy also provides a means to represent dependence between defense options.

A defense option might be "do nothing," but it might also be any of a range of measures to harden, defend, or otherwise render the activity less (or more) vulnerable to attack. Activity under each defense option may exact a distinct activity cost. For instance, the activity may be more expensive because of a security checkpoint and offset fencing.

A defense option may also be used to build, establish, enable, or otherwise initiate a completely new activity. For example, using (D7) for a new candidate activity, the do nothing defense option would be associated with activity variables that are restricted to zero, and any other defense option would enable some range of activity at a fixed defense cost, and with an option-dependent variable cost of operation.

Constraints (D1) can include limitations on many defender resources, and interactions among defense options, expressing realistic constraints on defense courses of action.

A target group is a set of activities that, for instance, might be susceptible to a particular kind of attack, or might be a collocated set of activities all vulnerable to a single kinetic attack. An example of a target group might be a "refinery" consisting of pipes, pumps, tanks, and manifolds, and the targets in that group would be associated with particular activities within that refinery. We have also used the notion of a target complex to model the common dependence of several mechanical components on a shared electricity source (e.g., Montgomery [62]); in this case, an "attack" on the target complex represents the loss of power that, in turn, disables all components in the target complex. Target complexes can thus be used to represent a variety of kinetic and nonkinetic vulnerabilities, including energy, computing, communication, and controls (i.e., a *cyber* layer). Constraints (D5) can include limitations on many attacker resources, and interactions among attack options, expressing realistic constraints on attacker courses of action. If target attacks need not be governed by target groups, simply drop constraints (D4) and the first term of constraint (D5).

Defense options can influence efficiency and the amounts of resources available. We have the ability to alter the efficiency of an activity and/or the amount of resources available for an activity as a function of its defense option (e.g., by replacing a_{ij} with a_{ij}^d and/or replacing b_i with b_i^d in (D6) and adding defense option as a summand).

This optimization model needs to be instrumented so that *no matter the data, a feasible solution is always rendered.* This requires some obvious preprocessing of, e.g., costs and budgets (to ensure nonnegativity), and in particular making constraints (D6) elastic (e.g., Brown and Dell [25]). In the event any elastic penalty is incurred, our solution algorithms (to come) will still function, but this is a signal of outright "system failure," a catastrophic collapse that indicates operating the system for basic functionality is no longer feasible.

4.2. Building the Tri-Level Model

Despite the relatively compact nature of formulation *DEFEND-CRIT-INF*, we do not recommend trying to write the formulation all at once.

Instead, our experience is that it is better to follow the script in §3, and build these models from the inside out, starting with the operator model. In the case of a minimum cost problem, this is going to assume the form

$$\min_{y \in Y(\hat{w})} f(\hat{w}, \hat{x}, y). \tag{20}$$

Our experience is that most of the effort in building these models lies in the development of an operator model that is instrumented to allow for all the attack and defense options we want to consider, and to always yield a feasible solution, albeit perhaps one with telltale penalties that indicate the system has broken down (for example, all paths from source to sink have been blocked by at least one edge targeted in an attack). See, for example, Brown and Dell [25].

Once we have the parameterized operational model, it is a short trip to build a model that reveals how an attacker could maximize system operating costs by choice of a worst-case feasible attack:

$$\max_{x\in X} \min_{y\in Y(\hat{w})} f(\hat{w}, x, y). \tag{21}$$

Finding the best defense against the worst-case disruption from an attacker yields our tri-level model:

$$\min_{w\in W} \max_{x\in X} \min_{y\in Y(w)} f(w, x, y). \tag{22}$$

The key issues to address in this development are the identification of components, system configuration, the decisions available to the operator, the performance metric for the system, the potential attacks, and the design options. Table 1 illustrates the key features of an infrastructure "system" and the corresponding defender-attacker-defender model for each of

Table 1. Examples of various modeling elements, respectively, required for an electric power transmission grid, a highway network, and undersea fiber-optic communication systems.

	Electric power transmission grid	Highway network	Undersea communications cables
System components	Generators; buses; transmission lines; transformers; substations	Road segments; tunnels; bridges; interchanges	Landing stations; branching units; repeaters; fiber-optic cables ("links")
System configuration	Inter-component connections; line thermal capacities; generating capacities	Inter-component connections; component lengths, capacities, and speed limits	Inter-component connections; router capacities; link capacities
Relevant operating environment	During one or more weekday time periods: generation costs; customer classes; load-shedding costs; demands at each bus	During one or more peak travel periods: demands for vehicular travel between origin-destination pairs	During one or more periods of high demand: user requirements for end-to-end communications
Operator	Independent system operator makes centralized, near-real-time generating decisions to balance supply with demand	Drivers select routes in a decentralized but "smart" fashion (implicitly following the tenets of game-theoretic, equilibrium model)	Undersea cable operator establishes end-to-end "lightpath" connections, and "grooms" network traffic (e.g., Zhu and Mukherjee [94])
Operator model	A "DC optimal power-flow model" (a linear program) that system operators use to optimize generation to meet demands (e.g., Wood and Wollenberg [90], pp. 108–111)	A traffic-equilibrium model (solved as a nonlinear program) for origin-destination routing decisions and travel times (e.g., Beckmann et al. [12])	A multicommodity transportation model to route customer traffic (e.g., Mukherjee et al. [63])
System performance metric	Minimize: generation costs plus the economic cost of unserved demand over the course of a typical work day (e.g., Salmerón et al. [74])	Minimize: average travel time during for network users during a peak commute period (e.g., Alderson et al. [8])	Minimize: traffic delays and shortage penalties for unmet end-to-end traffic demands (e.g., Crain [36])
Attacks on components	Generators, buses, etc., damaged or destroyed by explosives, gunfire, etc.	Road segments, tunnels, etc., damaged or destroyed by explosives, burning liquids, etc.	Cables severed by accident, natural disaster, or deliberate attack; landing stations attacked
Design (defenses)	Offset fencing at substations; physical or electromagnetic shielding; surplus component capacity (e.g., new generators, upgraded transmission lines)	Vehicle inspections at bridge entrances; structural reinforcement; increased police patrols; surplus component capacity (e.g., new bridges, widened roads)	Construction of additional redundant pathways; Enhanced physical security at landing stations

three applications: an electric power transmission grid, a highway network, and an undersea telecommunication system.

4.3. Stochastic Models and Nondeliberate Hazards

The models presented thus far are deterministic mathematical programs; however, this is not a requirement. If the infrastructure system's operation has significant randomness in it, the innermost operator model could be a stochastic program or even a complex simulation model, provided that the simulation used is susceptible to optimization (see Fu [42] for a comprehensive review of simulation optimization, and Subramanian et al. [81] for an example of simulation optimization applied to a pipeline infrastructure planning problem).

If the events of concern are all nondeliberate (e.g., natural disasters, accidents, or random failures), we might try to define a probability distribution over the set of disruptive events X, and the appropriate expression of the expected disruption could then take the form of a stochastic optimization problem:

$$\mathbb{E}_{\tilde{x}}\Big[\min_{y\in Y(\hat{w})} f(\hat{w},\tilde{x},y)\Big]. \tag{23}$$

Here, $\tilde{x}$ is a random variable taking on values in the set X, and $\mathbb{E}_{\tilde{x}}$ denotes the expectation with regard to $\tilde{x}$. This expectation can be evaluated using Monte Carlo methods: for a fixed design $\hat{w}$, iteratively generate realizations of $\tilde{x}$ and solve the (deterministic) operator model for each, then use these trials to calculate the expectation, construct confidence intervals, etc. See Chen and Miller-Hooks [33] for an example of this type of formulation and how to solve it, as applied to intermodal freight transport.

The corresponding model for defending against random disruptions is

$$\min_{w\in W} \mathbb{E}_{\tilde{x}}\Big[\min_{y\in Y(w)} f(w,\tilde{x},y)\Big]. \tag{24}$$

In this problem, the defender makes an up-front investment w before the random event $\tilde{x}$ is realized, and then the operator chooses y to operate the system at minimum cost. This is just two-stage stochastic programming with recourse (e.g., see Kall and Wallace [54] or Birge and Louveaux [15]).

We have two concerns for modeling disruptive events using random variables. The first is validation of the probability distribution for $\tilde{x}$. For many nondeliberate events (e.g., weather), there is sufficient historical data to build validated models of random events. However, the National Research Council has criticized the use of random variables to represent the deliberate action of an intelligent adversary (e.g., terrorists), because such models cannot be validated National Research Council (NRC) [65, 66]; see also the discussion in Brown and Cox [23, 24].

The second concern with an expected performance measure, such as (23), is that it is more consistent with measures of *risk* and does not really characterize resilience. Our view is that in order to be resilient to a set of events, the system must perform well for all events in the set, not just the "most likely" ones.

A complicating issue for the defender of an infrastructure system is that the parts of the system identified as "most critical" by stochastic models often differ from those identified using worst-case analysis. In practice, a defender must consider both nondeliberate (e.g., Mother Nature) and deliberate (e.g., terrorist) disruptions. Determining how to combine these different objectives in a sensible manner is an open topic for research.

4.4. Choosing an Operator Model

The general formulation *DEFEND-CRIT-INF* is flexible enough to represent a diversity of critical infrastructures and other large systems. As we have advised, developing an operator

model that reflects anticipated responses to induced failures of system components, perhaps failures to numbers of these, takes a lot of time and effort. But, with a good operator model in place, developing the additional layers for the attacker model and defender model is a relatively straightforward exercise (although the task of solving these models and communicating insight about their results can be nontrivial; we defer these issues until §§5 and 6).

Rather than trying to prescribe how to build a good representation of every possible infrastructure behavior (something beyond the scope of this tutorial), this section provides an informal taxonomy for some of the variations of models we have built or studied, and the applications in which we have found them to be effective. The hope is that this discussion of a variety of applications, along with references, provides a good starting point for any particular modeling effort.

4.4.1. Shortest-Path Problems. Shortest-path problems are perhaps the simplest to understand, formulate, and interdict. They are most useful when the operator is concerned with the movement of a single entity from one location to another. For example, in a follow-up study to the seminal analysis of the Soviet rail system, Wollmer [89, p. v] discusses strategies to "attack the link whose disruption would force trains to take the longest route." The inherent relationship between distance and time for physical systems makes shortest-path problems also relevant to emergency management applications, for instance, where the operator might want to route an emergency response vehicle from a fire station to one (or more) potential disaster or accident locations with the goal of finding the route(s) with minimum travel time, or where the operator's goal is to minimize evacuation clearing times. In both cases, the worst-case interdiction is the one that delays operations the most.

Time and distance are not the only choices for performance measures that lead to shortest-path formulations, though. If the "length" of an arc in the network is taken to be the probability of a successful transit, then the product of the lengths of the arcs in any path is the independent probability of successfully making it from one end of the path to another. A logarithmic transformation of the data converts the product of probabilities into a sum of (negative) values. Multiplying those transformed values by -1 and converting the original "max" to a "min" yields a shortest-path problem whose optimal solution is a path of minimum risk; it has the maximum probability of success in the original network. Applications of minimum-risk paths (and extensions of that model) include routing aircraft over air defense threats (Carlyle et al. [30], Royset et al. [73]), placing patrol boats to protect a valuable ship in a port from a small boat attack (Brown et al. [19]), or infiltration models including border security applications (Brown et al. [22], Dimitrov and Morton [40]).

4.4.2. Maximum-Flow Problems. In any situation where the operating costs are significantly less important than the value the infrastructure function provided to the consumer of that function (one recurring example we have seen is the supply of fuel through an existing system to support military operations), the operator will not be concerned about the costs (or time, or probability of success) of providing function. Instead, the focus will be on delivering *as much of* that function to a particular location (or from a particular location, or both) as the system will bear, and a maximum-flow formulation is the standard way to estimate this *capacity* of an infrastructure system. The Soviet rail system (Harris and Ross [49], Alderson et al. [7]) is a prime example of this; the operator wanted to know how much military equipment and personnel could be moved into Europe (a single "demand node") from various starting locations that can be supplied by a single "source node," limited only by the tonnage restrictions on tracks connecting adjacent locations. Another maximum-flow application is figuring out how much fuel the United States can extract in a fixed time period from the strategic petroleum reserve (Brown et al. [22]) during a fuel crisis.

4.4.3. Minimum-Cost Network Flow Problems. The shortest-path and maximum-flow problems are both special cases of the more general minimum-cost network flow problem, (or simply the *network flow* problem). For us, it is by far the most common starting point for modeling an infrastructure system, whether it be a set of storage tanks and pumps that can deliver fuel to various consumers (e.g., pp. 112–114 of Brown et al. [22], Ileto [52], or Chankij [32]), simple supply chain interdiction (e.g., see Snyder et al. [80], for a separate TutORial on this topic), or interdiction of military logistics networks based on time-phased force deployment data (TPFDD) (e.g., Brown [28], Koprowski [56]). Often, the transshipment networks of interest describe movement of goods in both space and time, and application of the attacker-defender techniques to this type of expanded network is straightforward (e.g., Derbes [39]).

Network flow problems can also be used to represent military problems such as ballistic missile defense, in which the operator wants to advantageously move launchers and then launch a simultaneous volley of missiles at a set of targets, in order to do as much damage to that list of targets as possible (e.g., Brown et al. [16], Repass [72]).

4.4.4. Multicommodity Flow Problems. Although a network flow formulation is a standard starting point for building an operator model, we frequently need more than one basic commodity to represent infrastructure function. This happens either because there are actually multiple commodities moving in the system (for example, the movement of both refrigerated and nonrefrigerated cargo through a port terminal in Delacruz [37]), or because the individual entities flowing through the system, although they look the same, are actually specified by both an origin and a destination, and each entity has to make a trip through the system from that origin to that destination, without substitution by other entities. Multicommodity flows have been used in this way to represent the operation (and interdiction) of public mass transit systems (e.g., pp. 114–115 of Brown et al. [22]), regional highway commuting systems (e.g., Alderson et al. [8]), ocean shipping lanes (e.g., Garcia Olalla [43]), and undersea fiberoptic communication cables (e.g., Crain [36]). Murray et al. [64] formulate and solve a simplified path-based interdiction problem for telecommunication flows. See Lim and Smith [59] for an introduction to interdiction models for multicommodity flow problems.

4.4.5. Project Scheduling Problems. If the operator is managing a large industrial project, the operator model can take the form of a project scheduling problem, which can represent the operator's decisions and resource allocations to keep the project moving (e.g., Dimitrov and Morton [40]). This type of model has application in many settings that are a step beyond a single infrastructure system, such as nuclear proliferation (e.g., Brown et al. [17, 20]). In this case the operator tries to build their first nuclear weapon as quickly as possible, possibly hiding his actual development plan by starting multiple similar projects.

4.4.6. Linear Programs. The preceding models appeal because they are easy to illustrate with nodes and arcs, Gantt charts, etc. but additional "side constraints" and complicating activities almost always arise that generalize beyond these network specializations. Nonetheless, a network characterization or a Gantt chart is a powerful device to help visualize the operator problem. Linear programs are not always as easy to illustrate, but they are quite general and powerful in representing the resource- and protocol-limited actions of a system operator. For example, although our graphical depiction of a petroleum system might be well illustrated by nodes and arcs, when we have to model operations of a refinery we need general technological constraints (e.g., Montgomery [62]). An electrical generation and distribution grid looks great on a map of the components and their connections, but the physics of its operation requires considerable modeling beyond networks (e.g., Salmerón et al. [75]). The multimodal transport (e.g., over barge, rail, truck) of multiple commodities (e.g., different types of coal (Alderson et al. [2]), or different types of petroleum-based fuels (Burton [29], Long [60])) is intuitively appealing as a system of interconnected networks, but the additional requirements on transfer

and storage of goods between modes requires more general constraints. We also find that linear programs are required in situations that incorporate long-term capacity expansion and protection (Brown [27]), use target groups (as described in §4.1) to represent interdependencies between components (Chankij [32], Montgomery [62], Burton [29], Long [60]), or model general interdependencies between infrastructure systems (Dixon [41]).

4.4.7. Integer-Linear Programs. If the operator must commit discrete decisions, we need to model these. An integer-linear program can express an extensive diversity of operator decision problems. Fortunately, we have at hand some very powerful off-the-shelf commercial optimization packages that can be used artfully to solve very large-scale integer-linear programs. However, we cannot always guarantee exact solutions to such operator models. We must usually admit some integrality tolerance between solutions and bounds on the quality of solutions not yet discovered, and we lose the ability to advantageously replace a primal linear program by its dual. We will discuss solution tactics in the next section to accommodate these complications.

Facility location problems are one example where discrete decisions are essential to the formulation. Scaparra and Church [76] formulate and solve tri-level optimization for facility location. Snyder et al. [80] consider facility location problems, network design models, and fortification models for supply chains.

Undersea warfare is another example requiring discrete decisions, because submarines can operate in either a "passive" (listening) mode or an "active" (pinging) search mode. BASTION (Thomas [86], Scherer [77], Brown et al. [18]) is a tri-level decision support system to defend an ocean area against submarine attackers. BASTION allocates defending marine patrol aircraft and their passive sonobuoys, and stealthy attack submarines, which cannot be sensed by an attacking submarine, as well as noisy surface ship searchers, their sonobuoy-dipping helicopters, and attack submarines that decide to switch to much more effective "active" search, all of which an attacking submarine can sense.

4.4.8. Nonlinear and Nonlinear-Integer Programs. Some operator models necessarily involve nonlinear phenomena such as congestion. We can accommodate these, with the caveat that they be solvable. For example, the operator model in Alderson et al. [8] is a convex, nonlinear program that evaluates total travel time for a population of travelers traversing a network. Alderson et al. [5] use a piecewise-linear approximation to an industry-standard sixth-degree polynomial function to represent the congestion of vehicles in a regional highway network.

4.4.9. Simulation, Simulation Optimization, Heuristics, Standard Protocols, and Advice from Subject Matter Experts. These can each be an essential source of advice for how an operator would respond to any contingency. However, these sources of suggested operator actions lack any quantitative bound on what better action may be available that we have not yet discovered. Also, one can anticipate interest in scenarios that range far outside of the normal operating domain, and rigidly abiding by standard operating procedures in such scenarios is not always good policy, so these resilience modeling exercises ought to consider how the system *could* respond to extreme situations, as opposed to how the system normally behaves.

5. Solution Tactics

As we can see in *DEFEND-CRIT-INF*, the typical defender, or defender-attacker-defender (DAD), model is a sequence of nested decisions that share the same objective function, albeit with opposing intent in the successive stages. There is, as yet, no standard software to solve tri-level decision problems in min-max-min (or max-min-max) form, but these problems can be solved. Though most published examples to date deal with decisions advised or forecast by

models that are linear programs, perhaps with integer and/or nonlinear embellishments, one has to be prepared to admit any decision source.

For infrastructure systems, developing a useful operator model (represented by the rightmost "D" in DAD) is crucial: it must represent how the infrastructure will operate in the presence of any foreseeable disruption with reasonable fidelity, and provide decision advice that is credible to the infrastructure operators, compliant with operational protocols, and consistent with the best-known characterizations of its operation.

The operator model is frequently a conventional optimization model seeking to minimize operating costs subject to constraints on system operation, one that may be solved by an algorithm or heuristic search. But it may also be a simulation, a prepackaged set of operator responses, a systems dynamics model, or even a human subject matter expert decision maker (preferably the system operator). We need to accommodate any and all operator models from the panoply of those accepted and in use. Fortunately, there is often already an industry-standard model or approved protocol available.

It is also possible that the attacker model (the "A" in DAD) is quite complex, as is the case in military planning. Similarly, the defender model (the left-most D) may be governed by complicated constraints on affordable and politically acceptable options to defend, harden, or otherwise improve resilience of the system under consideration.

5.1. Solving the Defender-Attacker-Defender Model with Decomposition

We present a general algorithm for solving DAD models like *DEFEND-CRIT-INF*. This algorithm was first presented by Alderson et al. [3]. Our algorithm requires a few new formulations, which we explain. We then offer some specializations of this algorithm for specific circumstances.

We refer to a generic DAD model as $\mathbf{DAD}(w, x, y)$

$$z^* = \min_{w \in W} \max_{x \in X} \min_{y \in Y(w)} f(w, x, y)$$

assuming (a) finite defense plans $w \in W$ (for notational simplicity, we have folded variables q from *DEFEND-CRIT-INF* into w), (b) finite attack plans $x \in X$ (folding variables m into x), and (c) operational decisions $y \in Y(w)$. With this notation we can then refer to the attacker model for a fixed defense, $\hat{w}$, as $\mathbf{DAD}(\hat{w}, x, y)$, and, similarly, the operator model for a fixed defense, $\hat{w}$, and a fixed attack, $\hat{x}$, as $\mathbf{DAD}(\hat{w}, \hat{x}, y)$. If we have a fixed attack, $\hat{x}$, we refer to the *design model* $\mathbf{DAD}(w, \hat{x}, y)$ that represents the (idealistic) situation in which the targets of the attack are known in advance, and the defender can design and operate the system optimally for that attack. Although this does not model a real-world situation, it provides a valid lower bound on system cost.

Defender-Attacker-Defender Decomposition. Because X is a finite set of attacks, we can define a set of vectors of variables $\{y^k\}$, where each element y^k denotes the operator's response to a specific attack $\hat{x}^k \in X$. In theory, $\mathbf{DAD}(w, x, y)$ may then be reformulated as

$$z^* = \min_{w \in W} \max_{\hat{x}^k \in X} \min_{y^k \in Y(w)} f(w, \hat{x}^k, y^k),$$

where the max operator now simply ranges over the discrete set of attacks and where the objective function $f(\,)$ is calculated separately for each of these attacks. The operator can now choose each y^k *in anticipation* of attack $\hat{x}^k$, so we can exchange the innermost min and max to obtain a formulation equivalent to **DAD**:

$$z^* = \min_{\substack{w \in W, \\ y^k \in Y(w)}} \max_{\hat{x}^k \in X} f(w, \hat{x}^k, y^k).$$

In practice this equivalent formulation is far too large to solve (although in some cases the enumeration of X is not prohibitive: see Tarvin et al. [82]). However, given an enumeration of a subset of K attacks, $\hat{x}^1, \hat{x}^2, \ldots, \hat{x}^K$, we can create a *relaxed master* problem *DAD-Master*:

$$\begin{aligned} z* = \min_{\substack{z,\, w \in W \\ y^k \in Y(w)}} \quad & z \\ \text{s.t.} \quad & z \geq f(w, \hat{x}^k, y^k) \quad \forall\, k = 1, \ldots, K. \end{aligned} \tag{DADC1}$$

A solution to *DAD-Master* provides a feasible defense $\hat{w}$ and, for each attack $\hat{x}^k$, an optimal operational response $\hat{y}^k$. For any fixed defense $\hat{w}$, a subroutine solves the *attacker subproblem* $\mathbf{DAD}(\hat{w}, x, y)$ for the resulting optimal attack, and provides a new cut (DADC1) for each such attack. (For *DEFEND-CRIT-INF*, this master problem turns out to be an integer-linear program.)

Our decomposition algorithm solves *DAD-master* to within a prescribed optimality tolerance ϵ_D, or terminates after a prescribed number of iterations, $K_{\mathrm{DAD}}^{\max}$, and we denote the resulting solution as $w^*(\hat{x})$, its objective value as z^{UP}, and the tightest upper bound obtained during its solution as z^{LO}. If $z^{\mathrm{LO}} = z^{\mathrm{UP}}$, then $w^*(\hat{x})$ is an optimal defense for the given set of attacks. After obtaining the solution, whether or not it is optimal, we have either $z^{\mathrm{UP}} - z^{\mathrm{LO}} \leq \epsilon_D$ or an iteration-limited solution. A description of the algorithm follows.

Algorithm DAD-Decomp

Input: Full **DAD** problem data, iteration limit $K_{\mathrm{DAD}}^{\max} > 0$ and optimality tolerance ϵ for the decomposition, tolerance ϵ_D for the master problem, and iteration limit $L_{AD}^{\max}$ and tolerance $\epsilon_{AD} \geq 0$ for the subproblem, with $\epsilon \geq \epsilon_D + \epsilon_{AD}$.

Output: ϵ-optimal iteration-limited defense plan w^* and corresponding attack plan x^*;

(1) Select a feasible initial attack, $\hat{x}^0 \in X$;

(2) Solve design model $\mathbf{DAD}(w, \hat{x}^0, y)$ to determine the optimal defense, $\hat{w}^1$, and operational cost, z^*, given this known attack;

(3) Initialize $\mathrm{LB} \leftarrow z^*$; $\mathrm{UB} \leftarrow \infty$; iteration counter $K \leftarrow 1$;

(4) *Subproblem*: Solve attacker subproblem $\mathbf{DAD}(\hat{w}^K, x, y)$ with solution tolerance ϵ_{AD} and iteration limit $L_{AD}^{\max}$ for attack $\hat{x}^K$;

(5) If $(\mathrm{UB} > z_{AD}^{\mathrm{UP}})\{\mathrm{UB} \leftarrow Z_{AD}^{\mathrm{UP}}; w^* \leftarrow \hat{w}^K; x^* \leftarrow \hat{x}^K;\}$

(6) If $(\mathrm{UB} - \mathrm{LB} \leq \epsilon|\mathrm{LB}|$ or $K \geq K_{\mathrm{DAD}}^{\max})$ goto **End**;

(7) *Master problem*: Given attack plans $\hat{x}^k, l = 1, \ldots, K$, solve *DAD-master* to determine defense plan $\hat{w}^{K+1}$ and $z^{\mathrm{UP}}, z^{\mathrm{LO}}$ such that $z^{\mathrm{UP}} - z^{\mathrm{LO}} \leq \epsilon_D|\mathrm{LB}|$;

(8) If $(\mathrm{LB} < z^{\mathrm{LO}})\mathrm{LB} \leftarrow z^{\mathrm{LO}}$;

(9) If $(\mathrm{UB} - \mathrm{LB} \leq \epsilon|\mathrm{LB}|)$ or iteration limit exceeded goto **End**;

(10) $K \leftarrow K + 1$; goto (4);

(11) **End**: Solve $\mathbf{DAD}(\hat{w}^*, \hat{x}^*, y)$ to determine optimal flows $\hat{y}^*$. Print("Best found defense and corresponding attack and flows are," $\hat{w}^*$, $\hat{x}^*$, $\hat{y}^*$);

It is possible that some attack $\hat{x}^K$ will repeat a prior attack; this could happen because the master or subproblem is not solved to optimality, but when the subproblem is also an integer linear program it can happen regardless of the quality of the solutions. Because of this, every version of this algorithm we have implemented for real infrastructure systems contains a mechanism to detect repeated solutions and prohibit them from reappearing in subsequent solves. When attack decisions are binary vectors, we can simply check each successive attack $\hat{x}^K$ against the attacks seen so far, and if it is a repeated attack, we can ignore it and add one solution elimination constraint (SEC) to the attacker subproblem:

$$\sum_{(i,j):\, \hat{x}_{ij}^k = 0} x_{ij} + \sum_{(i,j):\, \hat{x}_{ij}^k = 1} (1 - x_{ij}) \geq 1,$$

for each prior attack. Re-solving the restricted version of $\mathbf{DAD}(\hat{w}^K, x, y)$ yields a new, distinct $\hat{x}^K$ that we can use to build a new cut in the master. These SECs do not yield valid bounds on the overall decomposition, and could lead to premature termination. Therefore, when these SECs are enforced, steps 5 and 6 are skipped. See Alderson et al. [8] and its references. We must be careful to enforce these constraints only when we need them; they are used immediately after a repeated attack, to force the generation of a new attack, and then are removed in successive solves.

Similarly, step 7 may yield a repeated defense plan $\hat{w}^{K+1}$. Again, we enforce SECs for w until a new defense plan is discovered, and while these temporary SECs are in place, steps 8 and 9 are skipped.

5.2. Attacker Subproblem

Our algorithm assumes that there is a subroutine to solve the attacker model, $\mathbf{DAD}(\hat{w}, x, y)$, with two parameters chosen by the user: a solution tolerance ϵ_{AD}, and an iteration limit $L_{AD}^{\max}$. Regardless of the method used to solve the attacker subproblem, we must have termination criteria that guarantee a finite algorithm; in the absence of solution quality guarantees, the iteration limit $L_{AD}^{\max}$ can be used as a bound on the number of steps in a heuristic algorithm, or on the number of replications in a regenerative simulation, etc.

Attacker-Defender Subproblem and Solution via Decomposition. Given a fixed defense plan $\hat{w} \in W$, $\mathbf{DAD}(\hat{w}, x, y)$ represents the resulting attacker model as a subproblem of the defender model:

$$z_{AD}^*(\hat{w}) = \max_{x \in X} \min_{y \in Y(\hat{w})} f(\hat{w}, x, y).$$

We solve $\mathbf{DAD}(\hat{w}, x, y)$ using Benders decomposition (Geoffrion [44]). Given an enumeration of L feasible operational plans $\hat{y}^1, \ldots, \hat{y}^L$, the *AD-master* problem at iteration L is

$$\begin{aligned} z_{AD}^L(\hat{w}, \hat{y}) = \max_{z,\, x \in X} \quad & z \\ \text{s.t.} \quad & z \le f(\hat{w}, x, \hat{y}^l) \quad \forall\, l = 1, \ldots, L. \end{aligned} \tag{ADC1}$$

For any fixed attack $\hat{x}$, we solve the *operator subproblem* $\mathbf{DAD}(\hat{w}, \hat{x}, y)$ to obtain the resulting optimal operational plan, and add a new constraint, or "cut," (ADC1) for each such plan. We maintain an upper bound z_{AD}^{UP} from the successive master problem objective values and a lower bound z_{AD}^{LO} from the sequence of incumbent attacks, x^K, and their associated values when evaluated in the operator subproblem. The attacker subproblem Benders decomposition algorithm would then terminate when $z_{AD}^{\mathrm{UP}} - z_{AD}^{\mathrm{LO}} < \epsilon_{AD} |z_{AD}^{\mathrm{LO}}|$ or when the number of attacker master iterations L exceeds a prescribed bound $L_{AD}^{\max}$. We denote the resulting feasible attack as $x^*(\hat{w})$.

See Cormican et al. [35] as well as Israeli and Wood [53] for more details about solutions of *AD* models of this form by Benders decomposition. If we solve the attacker subproblem using a heuristic method that does not provide guaranteed, improving bounds on solution quality, we might need to ignore the optimality tolerance, but we still need some sort of iteration limit to guarantee a reasonable completion time for the algorithm (e.g., Scherer [77]).

More complicated subproblems (e.g., nonlinear and convex) can be handled as well, but we do not discuss those generalizations in this tutorial. See Geoffrion [44] for a discussion of models that could be handled by decomposition algorithms similar to ours.

5.3. Specializations

If a decision model is not solved exactly, and if we can develop no bound on how much better a solution might be, the termination conditions resort to primitive iteration limits. For example, if we solve the *AD* subproblem in step 4 with a method that provides no bound, we have $Z_{AD}^{\mathrm{UP}} = \infty$ and thus need to terminate the solution effort with the iteration limit $L_{AD}^{\max}$.

FIGURE 3. Random attacks on the Soviet railway compared with optimal ones.

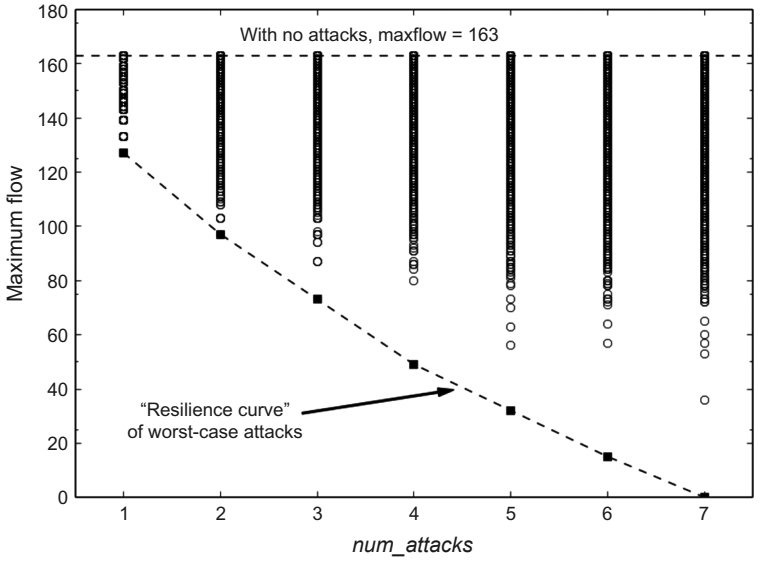

Notes. The maximum-flow (system performance) degrades with an increasing number of damaged activities. For an attack budget of $num\_attacks = 1, 2, \ldots, 7$, we present the worst-case disruption, along with 10,000 randomly generated attacks. As the number of possible attack combinations increases, it becomes harder and harder to find the worst-case attack by random sampling. The dashed line connects the worst-case disruptions for this system is our resilience curve, where the term "curve" refers here to a discrete frontier of points. (Figure from Alderson et al. [7], Figure 5.)

Similarly, if we solve the master problem in step 7 with no bound, we have $Z_{LO} = -\infty$, step 7 is ineffective, and we need to terminate this solution effort with the iteration limit $K_{AD}^{\max}$. In practice, we might adopt slightly more sophisticated stopping rules than a mere iteration limit, but such limits at least serve to bound the algorithms.

For cases in which we cannot get a bound on the optimal solution, the tests for repeated decisions are crucial, as well as a termination condition sensing complete enumeration. (We are reminded that when we cannot solve a problem exactly, a bound on the achievable value of any undiscovered solution is as important as the best solution found.) If our algorithm is likely to explore a large number of attacks, then detecting a repeat attack can become burdensome; over the course of the algorithm, after K iterations of the master problem we will have spent $O(|E|K^2)$ work to detect repeats if we simply search a linear list of dense vectors $\{\hat{x}^k\}$, each of length $|E|$. We can reduce this effort to about $O(K)$ with sparse storage schemes for the attack vectors and a hash function to rapidly look up each new attack.

5.4. Sampling-Based Solution Techniques

One may be tempted to just try many random attack or defense decisions in a Monte Carlo simulation, returning the best solution found. Figure 3 shows an example comparing the optimal solution to $DAD(\hat{w}, x, y)$ for various numbers of targets with the results of randomly sampling 10,000 attacks. Even for a small, simple maximum-flow problem, the most damaging random result is far from optimal, and there is no way of identifying this without knowledge of the true optimal attack decision. (This caution also applies to local search heuristics.)

Exhaustive decision enumeration can be quite useful. Consider a hypothetical problem to decide which three bridges to protect in the San Francisco Bay area against a worst-case pair of simultaneous attacks. There are only seven bridges in the area, so if you have an operator model (say, a traffic congestion model), you need only solve it $\binom{7}{3}\binom{7}{2} = 35 \times 21 = 735$ times (e.g., Tarvin et al. [82]) to find the exact best solution (and with the sole termination test being exhaustive enumeration).

5.5. Linear Programs Perhaps with Integer Features

Many of our decision models are linear programs, and linear programs can be replaced and solved via their duals. A frequent instance of this is solving $\mathbf{DAD}(\hat{w}, x, y)$ when the attacker model is a linear or linear-integer program with decisions x, and the operator model is a linear program with continuous decisions y. Replacing the operator decision model linear program by its dual yields a standard (i.e., single-level) integer-linear program that is equivalent to the attacker subproblem, but does not contain the operator variables explicitly. If we use this to solve the attacker subproblem to determine a (near-) optimal attack, $\hat{x}$, then a single solve of *DAD*$(\hat{w}, \hat{x}, y)$ recovers optimal operator decisions, y^* (e.g., Brown et al. [22], [16], Scaparra and Church [76], Zhao and Zeng [93]).

If the attacker decision model is a linear program, we can replace it by its dual, yielding an integer-linear program that is equivalent to the design model but without attack variables represented explicitly. After solving this monolithic optimization for w^* and y^* (with attacker influence expressed via its dual variables and constraints), we can recover the continuous attacker actions by solving *DAD*(w^*, x, y^*) (e.g., Brown et al. [18]).

5.6. Importance of Cost-Based Interdiction

We are tempted to present our tri-level formulation as

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

However, although attacks x do influence the domain of operations, Y, in practice, we do not model this directly through constraints, but rather by penalty costs—what we call "cost interdiction." Penalty terms in the objective function are associated with each attack variable, and high penalty costs on an activity component targeted by an attack will "encourage" the operator to avoid using that component. A large enough penalty renders such a component essentially unusable by the operator. Similarly, defenses w influence the domain of attacks, X, but again use "cost interdiction." Specifically, the design decisions determine which version of each activity is available to the operator; these different versions can have different penalty costs, and therefore the same attack will have a different effect in the objective for different design choices. These cost interdictions work well in practice, though they are not explicit constraints, and the above-mentioned influences of defenses on attacks, and of attacks on operation, are implicit in our notation.

6. Key Insights and How to Present Them

Our mathematical modeling results are not of much use unless we find a way to convey the "what" of our predictions so clients can draw informed conclusions about the "why," and decide what to do to increase resilience either by defensive measures or changes to operating procedures.

6.1. Performance and System Reconstitution Over Time

Operators of infrastructure systems are often concerned not just with the instantaneous loss of function from an attack, but also the way in which and the rate at which we can best recover over time. That is, we may care about the total loss time, and so we need to follow the best operational recovery over time.

If the system is able to recover from an event, we model this *reconstitution* in discrete stages (time *epochs*) over each of which the system performance is constant; at the end of each successive epoch, the performance improves as repairs are completed until reaching a fully reconstituted state. Each epoch can have a different duration depending on the difficulty of the repairs being modeled, and the overall consequence of the event can then be an integral over time (really just a summation) of the consequences over the finite set of epochs.

FIGURE 4. Reconstitution of a notional system following two different events.

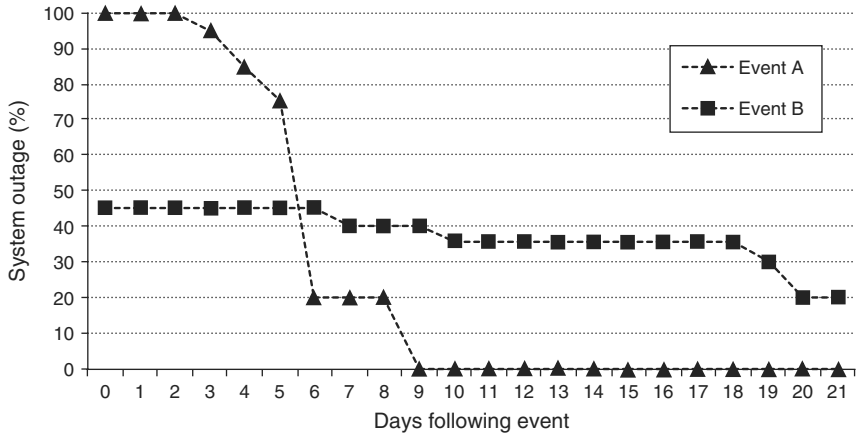

Notes. Event A causes a complete (100%) outage of system function for the first couple days, after which the system recovers quickly and achieves full reconstitution (0% outage) after nine days. Event B causes a smaller (45%) outage, but the system recovers more slowly; it still experiences a 20% outage three weeks following the event.

Figure 4 compares the reconstitution of a notional system following two distinct events. Event A causes a complete shutdown (100% outage), but the system is able to recover relatively quickly. Event B causes a smaller outage, but the system recovers more slowly. For cases in which the operator's objective is a function of system outage over the entire time horizon of interest (i.e., calculated as a summation over discrete time epochs), we face trade-offs between severity and duration in infrastructure disruption; these trade-offs arise frequently in attacker and defender models.

We typically represent the time required for reconstitution of each component as deterministic input data to the operator model. For cases in which reconstitution efforts compete for shared resources (e.g., when there is a limited number of repair crews), prioritization of reconstitution and recovery efforts is itself an important topic (e.g., Ang [10], Lee et al. [57], Nurre et al. [67], Cavdaroglu et al. [31]) and can additionally be included as part of the operator model (e.g., Gong et al. [45], Coffrin et al. [34], Thiébaux et al. [85]).

6.2. Moving Beyond Single Points of Failure

In our experience, it is common for decision makers to focus on the loss of single components and/or single infrastructure sites. In a recent, private conversation, we asked an independent system operator of an electric grid about not just $N-1$ reliability (i.e., the ability to provide service after the loss of any single component), but worst-case $N-K$ reliability (after losing K components). His initial response was "Oh, that just can't happen," because in his experience single equipment failures had occurred so infrequently, and could be repaired so quickly, that the idea of planning for two or three happening simultaneously was not worth the effort. Consistent with this mindset, electrical systems in the United States are regulated to be $N-1$ reliable, but there is no requirement for continuity of function after losing more than one major component.

However, there are cases where infrastructure systems suffer multiple losses, often due to unexpected circumstances, and operators can suffer serious consequences if they are unable to restore system function quickly. For example, from June 1997 to December 1998, the Union Pacific Railroad suffered a severe disruption in its ability to transport cargo because of extreme congestion caused by a confluence of unrelated events that included (a) a derailment in a key rail yard outside of Houston, Texas; (b) downtime of an important regional rail corridor because of unscheduled maintenance; (c) labor troubles in the neighboring Mexican

railroad because of political instability; and (d) unusual operating conditions at competing railroads (see Alderson [6], p. 4). This "perfect storm" of events was so improbable that it was beyond the imagination of the system operator before it occurred; yet it nearly caused the company to go bankrupt. A conservative worst-case analysis assessing the impact of losing a small number of track segments might have revealed this potential vulnerability in advance.

Of course, natural disasters and extreme weather events can cause damage to large numbers of system components, but in such cases the concern for continuity of function is typically replaced by a major mobilization to recover function.

6.3. Interpreting Resilience Curves

Solving the attacker model for different levels of *atk_budget* allows us to draw a resilience curve that provides a useful, quantitative characterization of operational resilience. Figure 5 shows the notional resilience curves for three separate systems, labeled A, B, and C. The relative shape of each curve reveals that these systems have very different resilience. Specifically, the performance of system B degrades approximately linearly with each component lost. By contrast, system A is "more resilient" because it suffers relatively little performance degradation with up to four component losses, and system C is "less resilient" because it suffers considerable performance degradation with only a single lost component.

6.4. Identifying Attack Sets and Defending Against Them

The contribution of a component to overall system function often depends on the availability of other components, and therefore we have argued that the focus when studying critical infrastructure systems should be on sets of components. However, presenting the results of such analysis can be challenging.

Tables 2 and 3—taken from a real-world, but anonymous, electric power system—follow a form that we have found to be useful in both written documents and presentations. Consistent with Salmerón et al. [74, 75], the assumption in the analysis underlying these figures is that the attacker's budget is measured in terms of number of human attackers available, and that each component has a minimum number of attackers needed (i.e., a cost) to interdict it.

Table 2 shows that the worst-case attack depends on the available attack budget and that the target list associated with the worst-case attack is not monotonic in that budget. We see that many components appear on some, but not all, target lists as the attacker's budget increases. This is evidence that it is the combined function of sets of components that

FIGURE 5. Resilience curves for three notional systems, and for disruptions that include the loss of up to five components.

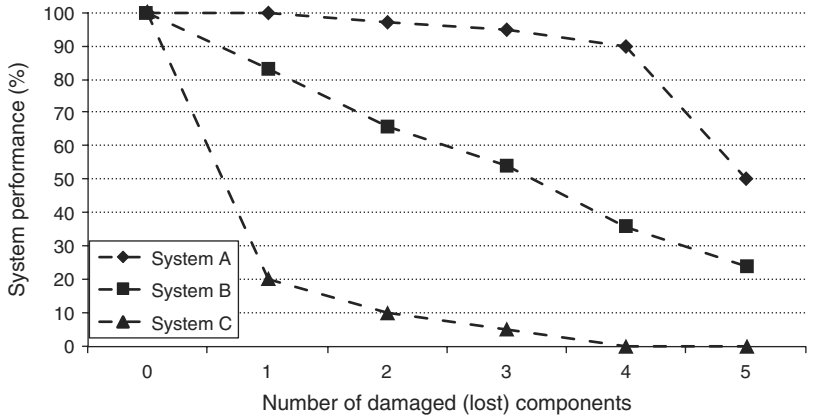

Notes. The performance of system B degrades approximately linearly with the size of the attack budget. System A is more resilient, whereas system C is less resilient for this range of disruption.

Table 2. Most-disruptive interdictions by attack budget.

		atk_budget											
Component name	*atk_cost*	1	2	3	4	5	6	7	8	9	10	11	12
Line1	1	X		X									
Line2	1										X		
Substation 1	2		X	X	X	X			X			X	
Substation 2	2				X								
Substation 3	3					X	X		X				
Substation 4	3						X	X	X	X	X	X	X
Substation 5	4							X		X	X	X	X
Substation 6	2									X	X	X	X
Substation 7	3												X

Note. Each component has a different cost to interdict; specifically, an overhead line can be interdicted by a single attacker, whereas the number of attackers required to interdict a substation is two or more. An "X" indicates the component is targeted in the worst-case attack. Note that as the number of attackers increases (this is a simple surrogate for an *atk_budget*), the components attacked do not simply accumulate. Rather, sets of components arise whose simultaneous loss cause most systemic damage. The lesson here is that the contribution of individual components to system function is only a first-order effect; much more important is the synergetic operation of component sets. To discover these synergistic effects, one must evaluate the entire system.

matters, rather than the contribution of any single component. In practice, we do not know in advance what the attacker's budget will be (or more generally, how large an incident we can reasonably expect). In many cases, however, the infrastructure operator will have a sense of how many components could reasonably be lost simultaneously, and this can provide some insight into the sets of components that are of most concern. Even if this is not the case, we observe an informal measure of component "importance" by the frequency with which individual components appear in the target lists for each worst-case attack. In Table 2, we observe substation 1 appears in many worst-case attacks, particularly when the attacker budget is small, whereas substation 4 appears in all worst-case attacks when the attacker budget is at least six.

Table 3. Optimal defensive "hardening" of links can mitigate the worst-case attack.

		def_budget					
Component name	*atk_cost*	0	1	2	3	4	5
Substation 1	4	X					
Substation 2	3	X	O	O	O	O	O
Substation 3	2	X					
Substation 4	3		X	X	X	X	X
Substation 5	2		X	O	O	O	O
Substation 6	3		X	X	X	X	O
Substation 7	2		X	X	X	O	O
Substation 8	2			X	O	O	O
Substation 9	2				X	X	X
Substation 10	2					X	X
Substation 11	3						X

Note. Here, an "O" represents the protection of a link visible to the attacker, and an "X" represents an attack. We consider the case from Table 2 when $atk\_budget = 10$. For a given number of attacks, an optimal defense "breaks up" the worst-case set of attacks, and the attacker finds the next-worst set of attacks. The activities in this system cannot be ranked in a simple priority list of importance, however, the frequency with which an activity appears in attack or defense solutions provides an indication of relative importance.

Table 3 lists the optimal defenses against an attack with $atk\_budget = 10$. Here, we have a case where the defenses are monotonic: as more defensive resources become available, new components are added to the defended set, and remain in that set. This is because defenses in this simple example all have the same cost, and this system does not have alternate low-order sets of vulnerable components that are particularly critical. Such is not always the case, but here this suggests a simple defensive strategy.

The optimal defenses here can be understood intuitively by observing the role that each defense plays in "breaking up" the target list for the worst-case attack. Specifically, we observe that in the absence of any defenses, the worst-case attack targets substations 1–3. The best single defense must protect one of these substations (or it will not mitigate the worst-case attack), and in the example here it protects substation 2. With substation 2 defended, the worst-case attack now targets substations 4–7 (note that it does not target substation 1 or substation 3). The best defensive strategy for two substations is to protect substation 2 and substation 5, thereby breaking up each of the target lists associated with these observed worst-case attacks. Reading the columns of Table 3 from left to right, we observe what can be interpreted as a type of "iterative and incremental fictitious play" between attacker and defender, where the defender breaks up the current worst target list, and the attacker then finds a new worst target list. This is not what is happening in our model or solution algorithm; each of the entries in Table 3 comes from a separate run of our tri-level model. However, the form of this table reveals key features regarding the relative importance of sets of components.

If necessary, we can force attack and/or defense monotonicity by sequentially fixing variables as we increase *atk_budget* and/or *def_budget*. Although this may render solutions that are easier to brief, you never know how much better you can do without solving the unrestricted case. See Koç et al. [55] for a discussion of how to build priority lists that are robust to uncertainties in available budget, as well as when such priority lists can be optimal.

6.5. Uniqueness and Relative Quality of Solutions

The solution to *ATTACK-RAIL-NET* identifies a single worst-case attack, for a given level of *atk_budget* (attacker capability). However, this solution on its own does not reveal whether there are many such attacks that can achieve the same consequence, or how different the "worst" is from other attacker options. We can solve for this information by the use of solution elimination constraints (§5.1). Investigating these attack alternatives informs us whether we can mount defenses visible to the attacker that are robust. Figure 6 shows a notional example.

We can do the same when solving for best defenses. If there are many defensive solutions that are about the same in quality, this allows the decision maker to select the one that best fits other criteria perhaps not represented explicitly in the formulation (e.g., the most politically palatable).

6.6. Accepting Exogenous Advice

It is important that the decision domains admit exogenous guidance. We have found it essential to be able to evaluate a defender decision that has been suggested as policy. We have also had to use defender decisions dictated by standard defense planning guidance (doctrine). We have been pleased to reassure decision makers when a suggested defense policy is a good one, and we have attempted to change what we have discovered to be poor doctrine. For example, if we are analyzing the resilience of the San Francisco Bay Area transportation network to attacks against bridges, our optimal solution for a particular scenario might be to defend the Bay Bridge and the San Mateo Bridge. If a city planner demands that we defend the Golden Gate Bridge because of its historic value and status as a national icon, we can fix that single component in w in the defender model and determine how to best expend any remaining resources. We can quantify exactly how much resilience we lose if we commit to this suboptimal defense, and a decision maker can then weigh this cost against the political value of defending an "obvious" target.

FIGURE 6. Top five rank-ordered attacks for target lists containing one to three components.

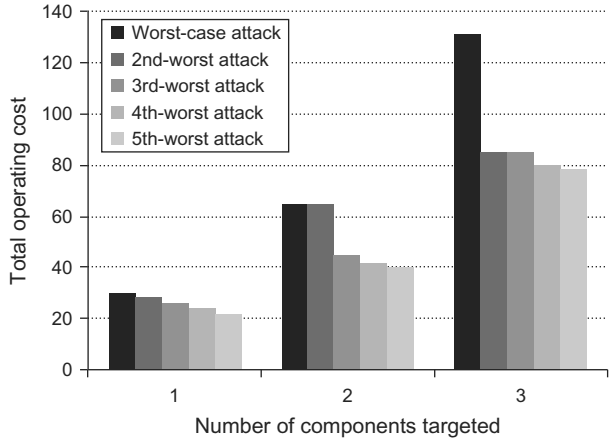

Notes. For a given number of components in a target list, what are the best worse-case alternatives for the attacker? Here, with just one component targeted, the rank ordering of the attacker's best five alternatives differs only slightly from the best to the fifth-best attack. With two components targeted, the best two attacks are much more damaging than the lesser alternatives. With three components targeted, the best attack is much better that the lessor alternatives (here, a target list with three components has discovered a particular set of components that, together, are acutely vulnerable).

6.7. Assessing Attacker and Defender Capabilities

Many critical physical infrastructures are quite vulnerable to primitive physical attacks (e.g., Smith [79]). In our unclassified work, we use reasonable first-order estimates of the resources and level of effort required to attack some component or target complex of components, and we usually assume an attacked vulnerable component is disabled. These estimates turn out to be reliable enough for the sort of resilience drills we run. In particular, these quickly reveal synergistic effects among sets of components, worthy of attention whether or not our estimates of attacker capability are precise. There is a huge literature available on this topic: a simple Google search of "weapons effects" yields many thousands of relevant hits.

In the military, we have engineers whose profession is testing and cataloging for our uniformed forces the kinetic means required to achieve any end in terms of physical damage. We also have special operations forces available to help us plan small-unit attacks. That our cyber infrastructure is also vulnerable is not lost on us, and we are assisting our Cyber Corps to assess threats there.

Similarly, our unclassified assessments of defensive measures are based on our judgment of things such as the ability to defend components, to harden them, to identify and localize component damage, the availability of spare components and repair parts, the preparation of personnel and equipment to restore damage, and the affordability of the measures themselves. Here, we strive to advise wise investments to mitigate the worst vulnerabilities.

7. Through the Looking Glass: Resilience from the "Top-Down"

Given the ultimate objective to improve the operational resilience of an infrastructure system, it is tempting to start with an optimization problem of the form

$$\max_{w \in W} h(w), \tag{25}$$

where the function $h(w)$ represents the resilience of the system after choosing design w.

The simplicity of this optimization problem is appealing, but two challenges arise in the definition of $h(w)$. First, determining the appropriate units of "resilience" is not straightforward, and second, even if the units of h have been decided, how does the function $h(w)$ map a design decision to the resilience to the as-yet unrepresented disruption(s)?

One of the major insights we have gained while studying critical infrastructure systems is that the term resilience itself is meaningless without context: in order to begin studying the resilience of any specific system, we always find ourselves asking the same question: "Resilience to what?" That is, without specifying the source of the disruption, and, more specifically, some limitations on the magnitude of the disruption, one can only talk about resilience in the most generic, relative terms. This follows the basic argument in Alderson et al. [7] (a similar argument for discussing "robustness" was made in Alderson and Doyle [9]).

If we define the set of potential disruptive events X, then we might introduce the function $g(w, x)$ to represent the system resilience to the disruptive event $x \in X$ when the system is designed according to w. Taking a conservative approach, our overall goal might be to come up with a design that maximizes the resilience in the presence of the worst possible disruptive event:

$$\max_{w \in W} \min_{x \in X} g(w, x). \tag{26}$$

The corresponding case for nondeliberate events modeled stochastically follows from (23):

$$\min_{w \in W} \mathbb{E}_{\tilde{x}}[g(w, \tilde{x})]. \tag{27}$$

The common requirement for (26) and (27) is specification of a function $g(w, x)$ that directly measures the consequence to the system resulting from disruption x applied to design w. For systems in which damage can be directly assessed in a deterministic manner, this might be straightforward. For instance, if the only "consequence" is that the operator immediately pays a replacement cost to completely and instantaneously recover the full operation of his system, $g(w, x)$ simply measures the total replacement cost of the components damaged or destroyed by the disruption x. This is the basic perspective in Bier [13], Zhuang and Bier [95], Bier et al. [14], Hauksen et al. [50], and Levitin and Hauksen [58].

For many systems, the consequence associated with a disruption depends not only on the initiating event but also on how the system responds to it. Indeed, system response is a common theme across many recent discussions of resilience in engineering systems (e.g., Hale and Heier [48], Woods [92], Haimes [46], Madni and Jackson [61], Vugrin et al. [88], Park et al. [68]). However, when writing a descriptive function $g(w, x)$, this creates the additional challenge of specifying in advance the response to every possible disruption. Such a task can be daunting, if not impossible, because the set of feasible disruptions, X, can be enormous.

Our second and possibly more important insight in our study of critical infrastructure systems is that we can drastically simplify the modeling of the response of the system to each disruption if we explicitly model the operation of the system by an operator who is actively working post-event to minimize consequences that result from a known disruption. This requires an operational model of the system accounting for design decisions and disruptions. Given such a model as a starting point, we arrive at the third, and final, level of the formulation:

$$\min_{w \in W} \max_{x \in X} \min_{y \in Y(w)} f(w, x, y). \tag{28}$$

Of course, this tri-level formulation is the starting point for this tutorial. The corresponding model for random disruptions is

$$\min_{w \in W} E_{\tilde{x}} \min_{y \in Y(w)} f(w, \tilde{x}, y). \tag{29}$$

We could walk through this line of reasoning for every infrastructure system whose resilience we wish to improve, but we always end up in the same place—we need an operational model that explicitly accounts for any design decisions that have been made, allows for the system to respond to disruptive events, and also explicitly accounts for any consequence that could result from any combination of these. Building such models would be arduous, were it not

FIGURE 7. In our experience, the most productive path to assessing and improving operational resilience is to model infrastructure systems "from the bottom up" starting with an operator model that is parameterized to account for any system design and any operational setting.

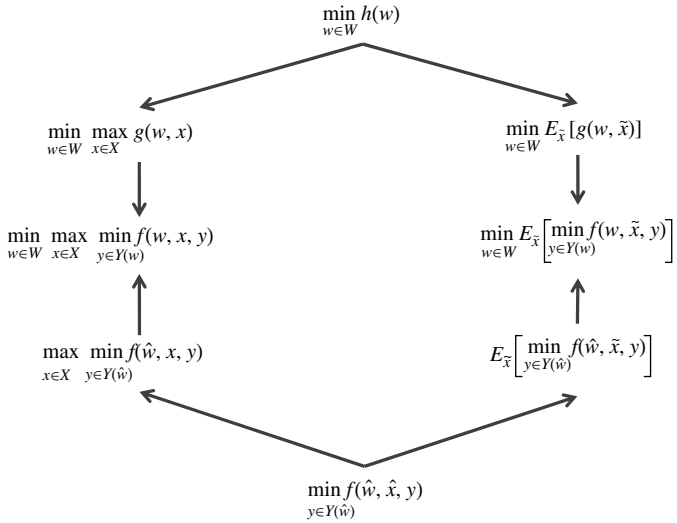

Notes. The function f measures the performance of the system as viewed and experienced by the system operator. Our attempts to represent system resilience directly (via the functions h or g) are incomplete or are not useful when it comes to informing defensive or operational decisions. [Figure adapted from Alderson et al. [4]].

for our third, and final, insight: the operator of the system is constantly making decisions about system behavior not only in response to disruptive events, but under normal operating conditions as well. Assessing and improving operational resilience is greatly facilitated by modeling the decisions that govern this "normal" case and then parameterizing the model to account for any changes in design and for any disruptions that may occur.

Figure 7 summarizes the relationship between these "top-down" and "bottom-up" views of assessing and improving the operational resilience of critical infrastructure systems.

8. Summary

Modern society depends on a multitude of critical infrastructure systems at the national, regional, and local levels. Owners, operators, and managers of these systems face economic pressure to increase the efficiency of these systems, often at the expense of operational resilience. Often, we discover this only when things go badly wrong, and we are left asking, "How could this have happened?"

The U.S. government has stated that we must consider resilience when allocating investment in infrastructure systems. There is a need for analysts to assess the resilience of these infrastructure systems and to identify investments that improve them.

The technique presented here requires knowledge of the objectives and the constraints for the system in question. It also requires some understanding of how disruptions to system components will impact the operational costs or constraints. However, we do not presume to know in advance the impact of an attack on the system as a whole; rather, our models solve for this.

We have applied models and methods like these with real-world, highly detailed, empirical data, and have had to scale up to larger scope and finer fidelity. We have confirmed with theory and a lot of empirical experience that our declared principles hold in all these circumstances.

Acknowledgments

This research was supported by the Air Force Office of Scientific Research, the Defense Threat Reduction Agency, and the Office of Naval Research. The authors thank Kevin Wood for stimulating conversations about resilience and for pointing out some key references. The authors thank Javier Salmerón and Ned Dimitrov for ongoing discussions about the use of system interdiction problems.

References

[1] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, Upper Saddle River, NJ, 1993.

[2] D. Alderson, G. Brown, J. DiRenzo, R. Engle, J. Jackson, B. Maule, and J. Onuska. Improving the resilience of coal transport in the Port of Pittsburgh—An example of defender-attacker-defender optimization-based decision support. Technical Report NSP-OR-12-004, Naval Postgraduate School, Monterey, CA, 2012.

[3] D. L. Alderson, G. G. Brown, W. M. Carlyle, and R. K. Wood. Optimal attack and defense of (network) infrastructures. Presentation, Military Operations Research Society Symposium, Fort Leavenworth, KS, June 15, 2009.

[4] D. L. Alderson, G. G. Brown, W. M. Carlyle, and R. K. Wood. Assessing and improving operational resilience of infrastructure systems. Presentation, INFORMS Computing Society Meeting, Santa Fe, NM, 6 January 2013.

[5] D. L. Alderson, G. G. Brown, W. M. Carlyle, and R. K. Wood. Assessing and optimizing the resilience of critical infrastructure to attack. Presentation, INFORMS Computing Society Meeting, Santa Fe, NM, 6 January 2013.

[6] D. L. Alderson. *Congestion-Induced Collapse in Networks: Managing Failure Cascades in Complex Systems and Infrastructure Protection*. Ph.D. thesis, Stanford University, Stanford, CA, 2003.

[7] D. L. Alderson, G. G. Brown, W. M. Carlyle, and L. A. Cox. Sometimes there is no "most vital" arc: Assessing and improving the operational resilience of systems. *Military Operations Research* 18(1):21–37, 2013.

[8] D. L. Alderson, G. G. Brown, W. M. Carlyle, and R. K. Wood. Solving defender-attacker-defender models for infrastructure defense. K. Wood and R. Dell, eds. *Operations Research, Computing and Homeland Defense*, INFORMS, Hanover, MD, 28–49, 2011.

[9] D. L. Alderson and J. C. Doyle. Contrasting views of complexity and their implications for network-centric infrastructures. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans* 40(4):839–852, 2010.

[10] C. C. Ang. Optimized recovery of damaged electrical power grids. Master's thesis, Naval Postgraduate School, Monterey, CA, 2006.

[11] M. J. Assante. Infrastructure Protection in the Ancient World. *Proceedings of the 42nd Hawaii International Conference on System Sciences—2009*. IEEE Computer Society, Washington, DC, 1–10, 2009.

[12] M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT, 1956.

[13] V. M. Bier. Choosing what to protect. *Risk Analysis* 27(3):607–620, 2007.

[14] 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* 28(3):763–770, 2008.

[15] J. R. Birge and F. Louveaux. *Introduction to Stochastic Programming*, 2nd edition, Springer, New York, 2011.

[16] G. Brown, M. Carlyle, D. Diehl, J. Kline, and K. Wood. A two-sided optimization for theater ballistic missile defense. *Operations Research* 53(5):263–275, 2005.

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

[18] G. Brown, J. Kline, A. Washburn, and K. Wood. A game-theoretic model for defense of an oceanic bastion against submarines. *Military Operations Research* 16(4):25–40, 2011.

[19] G. G. Brown, M. Carlyle, A. Abdul-Ghaffar, and J. Kline. A defender-attacker optimization of port radar surveillance. *Naval Research Logistics* 58(3):223–235, 2011.

[20] G. G. Brown, W. M. Carlyle, R. C. Harney, E. Skroch, and R. K. Wood. Interdicting a nuclear-weapons project. *Operations Research* 57(4):866–877, 2009.

[21] G. G. Brown, W. M. Carlyle, J. Salmerón, and K. Wood. Defending critical infrastructure. *Interfaces* 36(6):530–544, 2006.

[22] G. G. Brown, W. M. Carlyle, J. Salmerón, and R. K. Wood. Analyzing the vulnerability of critical infrastructure to attack, and planning defenses. J. C. Smith, ed. *Tutorials in Operations Research: Emerging Theory, Methods, and Applications*, INFORMS, Hanover, MD, 102–123, 2005.

[23] G. G. Brown and L. A. Cox. How probabilistic risk assessment can mislead terrorism risk analysis. *Risk Analysis* 31(2):196–204, 2011.

[24] G. G. Brown and L. A. Cox. Making terrorism risk analysis less harmful and more useful: Another try. *Risk Analysis* 31(2):193–195, 2011.

[25] G. G. Brown and R. F. Dell. Formulating linear and integer linear programs: A rogues' gallery. *INFORMS Transactions on Education* 7(2):153–159, 2007.

[26] K. A. Brown. *Critical Path: A Brief History of Critical Infrastructure Protection in the United States*. George Mason University, Fairfax, VA, 2006.

[27] P. S. Brown. Optimizing the long-term capacity expansion and protection of Iraqi oil infrastructure. Master's thesis, Naval Postgraduate School, Monterey, CA, 2005.

[28] W. A. Brown. Re-optimization of time-phased force deployment plans in response to emergent changes during deployment. Master's thesis, Naval Postgraduate School, Monterey, CA, 1999.

[29] C. W. Burton. Analyzing the U.S. Military fuel distribution network on Okinawa. Master's thesis, Naval Postgraduate School, Monterey, CA, 2013.

[30] W. M. Carlyle, J. O. Royset, and R. K. Wood. Lagrangian relaxation and enumeration for solving constrained shortest-path problems. *Networks* 52(4):256–270, 2008.

[31] B. Cavdaroglu, E. Hammel, T. C. Sharkey, and W. Wallace. Integrating restoration and scheduling decisions for disrupted interdependent infrastructure systems. *Annals of Operations Research* 203:279–294, 2013.

[32] M. Chankij. Assessing resilience of the Jet Propellant-8 (JP-8) distribution system on Guam. Master's thesis, Naval Postgraduate School, Monterey, CA, 2012.

[33] L. Chen and E. Miller-Hooks. Resilience: An indicator of recovery capability in intermodal freight transport. *Transportation Science* 46(1):109–123, 2012.

[34] C. Coffrin, P. van Hentenryck, and R. Bent. Last mile restoration for multiple interdependent infrastructures. *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*. AAAI Press, Palo Alto, CA, 2012.

[35] K. J. Cormican, D. P. Morton, and R. K. Wood. Stochastic network interdiction. *Operations Research* 46(2):184–197, 1999.

[36] J. K. Crain. Assessing resilience in the global undersea cable infrastructure. Master's thesis, Naval Postgraduate School, Monterey, CA, 2012.

[37] C. F. Delacruz. Defending the maritime transport of cargo for the Hawaiian Islands. Master's thesis, Naval Postgraduate School, Monterey, CA, 2011.

[38] Department of Homeland Security (DHS), National Infrastructure Protection Plan, Washington, DC, 2013.

[39] H. D. Derbes. Efficiently interdicting a time-expanded transshipment network. Master's thesis, Naval Postgraduate School, Monterey, CA, 1997.

[40] N. B. Dimitrov and D. P. Morton. Interdiction models and applications. J. W. Hermmann, ed. *Handbook of Operations Research for Homeland Security*, Springer, New York, 73–103, 2013.

[41] C. A. Dixon. Assessing vulnerabilities in interdependent infrastructures using attacker-defender models. Master's thesis, Naval Postgraduate School, Monterey, CA, 2011.

[42] M. C. Fu. Optimization for simulation: Theory vs. practice. *INFORMS Journal on Computing* 14(3):192–215, 2002.

[43] O. R. Garcia Olalla. Assessing the resilience of global sea routes. Master's thesis, Naval Postgraduate School, Monterey, CA, 2012.

[44] A. M. Geoffrion. Generalized Benders decomposition. *Journal of Optimization Theory and Applications*, 10(4):237–260, 1972.

[45] J. Gong, E. E. Lee, J. E. Mitchell, and W. A. Wallace. Logic-based multiobjective optimization for restoration planning. *Optimization and Logistics Challenges in the Enterprise*, Springer, New York, 305–324, 2009.

[46] Y. Y. Haimes. On the definition of resilience in systems. *Risk Analysis* 29(4):498–501, 2009.

[47] Y. Y. Haimes, K. G. Crowther, and B. M. Horowitz. Homeland security preparedness: Balancing protection with resilience in emergent systems. *Systems Engineering* 11(4):287–308, 2006.

[48] A. Hale and T. Heijer. Defining resilience. E. Hollnagel, D. D. Woods, and N. Leveson, eds. *Resilience Engineering: Concepts and Precepts*, Ashgate Press, Aldershot, UK, 95–123, 2006.

[49] T. E. Harris and F. S. Ross. Fundamentals of a method for evaluating rail net capacities. The RAND Corporation, Santa Monica, CA, Research Memorandum RM-1573, 1955.

[50] K. Hausken, V. Bier, and J. Zhuang. Defending against terrorism, natural disaster, and all hazards. V. M. Bier and M. N. Azaiez, eds. *Game Theoretic Risk Analysis of Security Threats*, Springer, New York, 65–97, 2009.

[51] Homeland Security Council (HSC). *National Strategy for Homeland Security*, The White House, Washington, DC, 2007.

[52] J. V. Ileto. Improving resiliency of the petroleum supply chain for the Hawaiian Islands. Master's thesis, Naval Postgraduate School, Monterey, CA, 2011.

[53] E. Israeli and R. K. Wood. Shortest-path network interdiction. *Networks* 40(3):97–111, 2002.

[54] P. Kall and W. Wallace. *Stochastic Programming*. John Wiley and Sons, Chichester, UK, 1994.

[55] A. Koç, D. P. Morton, E. Popova, S. M. Hess, E. Kee, and D. Richards. Prioritizing project selection. *The Engineering Economist* 54:267–297, 2009.

[56] P. M. Koprowski. Interdicting a force deployment two-sided optimization of asset selection, lift scheduling, and multi-commodity load planning. Master's thesis, Naval Postgraduate School, Monterey, CA, 2005.

[57] E. E. Lee, J. E. Mitchell, and W. A. Wallace. Restoration of services in interdependent infrastructure systems: A network flows approach. *Systems, Man, and Cybernetics, Part C: Applications and Reviews, IEEE Transactions* 37(6):1303–1317, 2007.

[58] G. Levitin and K. Hausken. Resource distribution in multiple attacks against a single target. *Risk Analysis* 30(8):1231–1239, 2010.

[59] C. Lim and J. C. Smith. Algorithms for discrete and continuous multicommodity flow network interdiction problems. *IIE Transactions* 39(1):15–26, 2007.

[60] C. S. Long. Analyzing the resilience of the U.S. Military fuel distribution system for mainland Japan. Master's thesis, Naval Postgraduate School, Monterey, CA, 2013.

[61] A. M. Madni and S. Jackson. Towards a conceptual framework for resilience engineering. *IEEE Systems Journal* 3(2):181–191, 2009.

[62] J. S. Montgomery. Oahu Petroleum infrastructure resilience. Master's thesis, Naval Postgraduate School, Monterey, CA, 2013.

[63] B. Mukherjee, B. Banerjee, S. Ramamurthy, and A. Mukherjee. Some principles for designing a wide-area WDM optical network. *IEEE/ACM Transactions on Networking* 4(5):684–706, 1996.

[64] A. T. Murray, T. C. Matisziw, and T. H. Grubesic. Critical network infrastructure analysis: Interdiction and system flow. *Journal of Geographical Systems* 9(2):103–117, 2007.

[65] National Research Council (NRC). *Department of Homeland Security Bioterrorist Risk Assessment: A Call for Change*, National Academies Press, Washington, DC, 2008.

[66] National Research Council (NRC). *Review of the Department of Homeland Securitys Approach to Risk Analysis*, National Academies Press, Washington, DC, 2010.

[67] S. G. Nurre, B. Cavdaroglu, J. E. Mitchell, T. C. Sharkey, and W. A. Wallace. Restoring infrastructure systems: An integrated network design and scheduling (INDS) problem. *European Journal of Operational Research* 223(3):794–806, 2012.

[68] J. Park, T.P. Seager, P.S.C. Rao, M. Convertino, and I. Linkov. Integrating risk and resilience approaches to catastrophe management in engineering systems. *Risk Analysis* 23(3):356–367, 2013.

[69] President's Commission on Critical Infrastructure Protection. Critical Foundations. Technical report, The White House, Washington, DC, 1997.

[70] Public Law 107-296, 117 Stat. 745. Homeland Security Act of 2002. http://www.dhs.gov/homeland-security-act-2002; Accessed: 2/1/2014.

[71] R. R. Rardin. *Optimization in Operations Research*. Prentice Hall, Upper Saddle River, NJ, 1997.

[72] L. M. Repass. Optimal stationing of radar pickets and anti-ballistic missile defenders for long range surveillance and tracking (LRSandT) and ballistic missile defense (BMD). Master's thesis, Naval Postgraduate School, Monterey, CA, 2006.

[73] J. O. Royset, W. M. Carlyle, and R. K. Wood. Routing military aircraft with a constrained shortest-path algorithm. *Military Operations Research* 14(3):31–52, 2009.

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

[75] 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.

[76] M. P. Scaparra and R. L. Church. A bilevel mixed-integer program for critical infrastructure protection planning. *Computers and Operations Research* 35:1905–1923, 2008.

[77] S. D. Scherer. Game-theoretic anti-submarine warfare mission planner (heuristic-based, fully excel capable). Master's thesis, Naval Postgraduate School, Monterey, CA, 2009.

[78] A. Schrijver. On the history of the transportation and maximum flow problems. *Mathematical Programming, Series B* 91:437–445, 2002.

[79] R. Smith. Assault on California Power Station Raises Alarm on Potential for Terrorism. *The Wall Street Journal*, February 4, 2014. Accessed February 11, 2014, http://online.wsj.com/news/articles/SB10001424052702304851104579359141941621778.

[80] L. V. Snyder, M. P. Scaparra, M. S. Daskin, and R. L. Church. Planning for disruptions in supply chain networks. M. P. Johnson, B. Norman, and N. Secomandi, eds. *Tutorials in Operations Research: Models, Methods, and Applications for Innovative Decision Making*, INFORMS, Hanover, MD, 234–257, 2006.

[81] D. Subramanian, J. F. Pekny, and G. V. Reklaitis. A simulation-optimization framework for addressing combinatorial and stochastic aspects of an R&D pipeline management problem. *Computers and Chemical Engineering* 24(2–7):1005–1011, 2000.

[82] D. Tarvin, K. Wood, and A. Newman. Using enumeration to solve binary master problems in Benders decomposition. Working paper, Colorado School of Mines, Golden, CO, 2014.

[83] The White House. Homeland Security Presidential Directive 7: Critical Infrastructure Identification, Prioritization, and Protection, Washington, DC, 2003.

[84] The White House. Presidential Policy Directive 21—Critical Infrastructure Security and Resilience, Washington, DC, 2013.

[85] S. Thiébaux, C. Coffrin, H. Hijazi, and J. Slaney. Planning with MIP for supply restoration in power distribution systems. *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence, IJCAI'13*, AAAI Press, Palo Alto, CA, 2900–2907, 2013.

[86] A. J. Thomas. Tri-level optimization for anti-submarine warfare mission planning. Master's thesis, Naval Postgraduate School, Monterey, CA, 2008.

[87] Title 42 U.S. Code, Sec. 5195c et seq. 2006 Supp. IV. Critical infrastructures protection, 2011. The Stafford Act. Accessed September 8, 2013, http://www.gpo.gov/.

[88] E. D. Vugrin, D. E. Warren, M. A. Ehlen, and R. C. Camphouse. A framework for assessing the resilience of infrastructure and economic systems. K. Gopalakrishnan and S. Peeta, eds. *Sustainable and Resilient Critical Infrastructure Systems: Simulation, Modeling, and Intelligent Engineering*, Springer-Verlag, Berlin, 77–116, 2010.

[89] R. D. Wollmer. Some methods for determining the most vital link in a railway network. The RAND Corporation, Santa Monica, CA, Research Memorandum RM-3321-ISA, 1963.

[90] A. J. Wood and B. F. Wollenberg. *Power Generation, Operation and Control*, 2nd edition. John Wiley & Sons, New York, 1996.

[91] R. K. Wood. Bilevel network interdiction models: Formulations and solutions. J. J. Cochran, ed. *Wiley Encyclopedia of Operations Research and Management Science*, John Wiley & Sons, Hoboken, NJ, 1–11, 2011.

[92] D. D. Woods. Essential characteristics of resilience. E. Hollnagel, D. D. Woods, and N. Leveson, eds. *Resilience Engineering: Concepts and Precepts*, Ashgate Press, Aldershot, UK, 49–60, 2006.

[93] L. Zhao and B. Zeng. Robust unit commitment problem with demand response and wind energy. Technical report, University of South Florida, Tampa, FL, October 2010.

[94] K. Zhu and B. Mukherjee. Traffic grooming in an optical WDM mesh network. *IEEE Journal on Selected Areas in Communication* 20(1):122–133, 2002.

[95] J. Zhuang and V. M. Bier. Balancing terrorism and natural disasters-defensive strategy with endogenous attacker effort. *Operations Research* 55(5):976–991, 2007.

[96] A. Zolli and A. M. Healy. *Resilience: Why Things Bounce Back*. Free Press, New York, 2012.