



Calhoun: The NPS Institutional Archive
DSpace Repository

NPS Scholarship

Publications

2017-07-19

Assessing and Improving the Operational Resilience of a Large Highway Infrastructure System to Worst-Case Losses

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

INFORMS

Alderson, David L., et al. "Assessing and Improving the Operational Resilience of a Large Highway Infrastructure System to Worst-Case Losses." *Transportation Science* (2017).
<https://hdl.handle.net/10945/60908>

This publication is a work of the U.S. Government as defined in Title 17, United States Code, Section 101. Copyright protection is not available for this work in the United States.

Downloaded from NPS Archive: Calhoun



Calhoun is the Naval Postgraduate School's public access digital repository for research materials and institutional publications created by the NPS community. Calhoun is named for Professor of Mathematics Guy K. Calhoun, NPS's first appointed -- and published -- scholarly author.

Dudley Knox Library / Naval Postgraduate School
411 Dyer Road / 1 University Circle
Monterey, California USA 93943

<http://www.nps.edu/library>

Assessing and Improving the Operational Resilience of a Large Highway Infrastructure System to Worst-Case Losses

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

^a Operations Research Department, Naval Postgraduate School, Monterey, California 93943

Contact: dlalders@nps.edu (DLA); gbrown@nps.edu (GGB); mc Carlyle@nps.edu (WMC); kwood@nps.edu (RKW)

Received: November 18, 2015

Revised: September 8, 2016

Accepted: November 8, 2016

Published Online in Articles in Advance:
July 19, 2017

<https://doi.org/10.1287/trsc.2017.0749>

Copyright: © 2017 INFORMS

Abstract. This paper studies the resilience of the regional highway transportation system of the San Francisco Bay Area. Focusing on peak periods for commuter traffic, traffic patterns are computed from a model that includes nonlinear increases in travel times due to congestion and reflects actual travel demands as captured by U.S. Census demographic data. We consider the consequences associated with loss of one or more road, bridge, and/or tunnel segments, where travelers are allowed to reroute to avoid congestion or potentially not travel at all if traffic is bad. We use a sequential game to identify sets of road, bridge, or tunnel segments whose loss results in worst-case travel times. We also demonstrate how the model can be used to quantify the operational resilience of the system, as well as to characterize trade-offs in resilience for different defensive investments, thus providing concise information to guide planners and decision makers.

Funding: This research was supported by the Defense Threat Reduction Agency [Grant HDTRA1-10-1-0087], and also by the Office of Naval Research and the Air Force Office of Scientific Research.

Keywords: infrastructure • operational model • resilience • traffic congestion • defender-attacker-defender • game theory • optimization

1. Introduction

In the last decade, the term “resilience” has come to the forefront of the discussion in the United States regarding investments in major infrastructure systems such as highway networks. As the 2007 National Strategy for Homeland Security recognizes,

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 . . . We must now focus on the resilience of the system as a whole—an approach that centers on investments that make the system better able to absorb the impact of an event without losing the capacity to function (Homeland Security Council 2007, pp. 27–28).

“Critical infrastructure” in this quote refers to a complete infrastructure system such as a highway or pipeline network, and we assume that “structural resilience” corresponds to the ability of individual system components to withstand and recover from the effects of natural disasters, accidents, attacks, and other forces (e.g., Williamson and Winget 2005, Cimellaro, Reinhorn, and Bruneau 2010). We leave the study of structural resilience to structural, civil, mechanical, and other engineers and focus our efforts on operational resilience of critical infrastructure.

We find no definition of operational resilience in the National Strategy or in more recent policy documents

such as Presidential Decision Directive 21 (The White House 2013). Yet, the National Strategy does connect the concept with “capacity to function,” and measuring and optimizing a system’s capacity to function lies in the domain of operations research.

This paper takes up this challenge of operational resilience for regional transportation systems. Specifically, it develops a mathematical model that identifies resource-limited investments that maximize the performance of regional highway infrastructure in the presence of worst-case losses of its components. By extending an industry-standard model for traffic congestion to measure the effect of worst-case losses, and by demonstrating results for a realistic data set from the San Francisco Bay Area, we show the practicability of the model and related solution methods.

Section 2 reviews the relevant literature addressing resilience of transportation systems and describes the basic ideas and literature supporting our technique. Section 3 provides a detailed formulation for this “defender-attacker-defender (DAD) model” as applied to designing defenses for a regional highway network. Section 4 demonstrates the model in a detailed case study of the highway (and bridge and tunnel) network of the San Francisco Bay Area. That section also illustrates the types of results that can be shown to decision makers who, ultimately, must determine a budget or approve a system design. Section 5 adds some discussion and concludes the paper. Appendices A and B provide additional detail on models and data.

2. Literature Review

The modern study of roadway traffic dates back to Wardrop (1952) and Beckmann, McGuire, and Winsten (1956); see also Boyce, Mahmassani, and Nagurney (2005) for a retrospective. Dafermos and Sparrow (1969) formalized the distinction between user equilibrium (UE) and system optimization (SO) solutions, respectively, to the traffic assignment problem. See also Beckmann (1967); Florian and Nguyen (1976); Dowling, Singh, and Cheng (1998); Gazis (2002, pp. 185–236); Boyce and Bar-Gera (2003); and Correa and Stier-Moses (2010) for background on modeling and assessment of roadway systems.

More recently, there is a growing literature focused on the performance of transportation systems in the presence of disruptive events. Faturechi and Miller-Hooks (2014b) review over 200 articles focused on the performance of transportation infrastructure systems in the presence of disasters, categorizing them in terms of the applied performance measure, their modeling technique, their assessment methodology, and the strategy used for managing these infrastructures during crisis. In this section, we review selected works most relevant to our study to place our contribution in context.

2.1. Recent Studies on Transportation Resilience

Liu, Fan, and Ordóñez (2009, p. 1582) consider the challenge of “allocating limited retrofit resources over multiple highway bridges to improve the resilience and robustness of the entire transportation system in question.” Theirs is a multicommodity flow model defined by origin-destination (OD) pairs, where the travel time on a link depends on its flow, as described by the Bureau of Public Roads (BPR) function (Bureau of Public Roads 1964). Seismic damage to highway bridges is uncertain and modeled using a finite number of damage scenarios. Decisions and scenario damages are binary. They present a two-stage stochastic program with the first stage predisaster preparations for bridge retrofit and the second postdisaster operations, i.e., user network flows resulting from a SO model. The best retrofit solution is the one that minimizes a combination of expected reconstruction costs and travel delays. They solve their problem with an L-shaped decomposition (Van Slyke and Wets 1969) and generalized Benders decomposition (Geoffrion 1972). They apply their analysis to two examples: (1) a small Sioux Falls road network with 24 nodes, 76 arcs, and only six bridges vulnerable; and (2) an example from Alameda County with 510 nodes and 1,424 arcs (resulting in 2,401 OD traffic pairs), and 13 vulnerable bridges.

Fan and Liu (2010) extend this work for essentially the same problem by solving it with a progressive hedging technique (Rockafellar and Wets 1991). They use the BPR function for traffic congestion and also

include a penalty cost for unsatisfied travel demand. Here, traffic flows follow a UE model. They use the same Sioux Falls example with six vulnerable bridges, and they completely enumerate 2^6 scenarios. They admit the method requires careful tuning—computation times can be long and solutions can oscillate.

Peeta et al. (2010) study how to strengthen the segments of a road network in preparation for a potential earthquake disruption, with the explicit objective to preserve OD connectivity for emergency responders. They consider a hypothetical earthquake in Istanbul for a road network consisting of 25 nodes and 30 arcs. The network model does not include road segment flow capacities or congestion, and they restrict attention to five OD pairs representing connectivity to the districts with the highest expected number of injured people for the earthquake scenarios under consideration. They pose a two-stage stochastic program in which the first stage involves seismic retrofit of bridges or viaducts to reduce their failure probability, and the second stage reveals the residual connectivity and path distances following the disaster. They solve the problem using a Monte Carlo sampling procedure.

Chen and Miller-Hooks (2012) study the resilience of an intermodal freight transport network in which recovery activities can be taken following a disaster. Specifically, they define resilience as the postdisaster expected fraction of demand that can be delivered in a specified recovery time and budget. Arc traversal times and capacities are random variables (realized following the disruption), however these can be improved by undertaking one of several recovery activities. Their model does not include prevent decisions. They formulate a stochastic mixed-integer program that maximizes the expected fraction of deliveries that can be satisfied. They consider the example of an 8-node, 12-arc network faced with six possible scenarios involving both natural and man-made disruptions, and with the possibility of six recovery activities, each of which improves a subset of network arcs. They solve their problem using a combination of Benders decomposition, column generation, and Monte Carlo simulation.

Miller-Hooks, Zhang, and Faturechi (2012) extend this work by considering a two-stage problem in which preparedness decisions are made in the first stage and recovery actions are made in the second stage. Again, resilience is defined as the fraction of demand that can be satisfied postdisaster. They propose a stochastic two-stage integer program in which the aim is to determine the optional portion of the budget to spend on preparedness and the amount of the budget to reserve for postdisaster recovery given an uncertain future network state. The probability of each disaster scenario is assumed to be known in advance, and the optimal investment plan is the one that yields the maximum expected network resilience. They study an

intermodal freight network with eight nodes, 24 one-way rail arcs, and 22 bidirectional virtual highway arcs. They consider two possible preparedness activities and five potential recovery activities. They use Monte Carlo estimation to generate 100 realizations for each of five disaster scenarios, and they solve the overall problem using an L-shaped decomposition technique.

Faturechi and Miller-Hooks (2014a) generalize the previous concept to an optimization problem to maximize the expected performance of an infrastructure system given first-stage preparedness actions and second-stage recourse actions. Here, “performance” could be any one of several measures (e.g., robustness, flexibility, recovery, resilience), subject to constraints on service guarantees and budget limitations. They formulate a two-stage stochastic program that is, in general, both nonlinear and nonconvex, and propose various decomposition techniques to solve it. As an illustrative example, they consider the previous multimodal shipping network, with the performance measured in terms of OD flows and under the same five disaster scenarios.

Faturechi and Miller-Hooks (2014c) apply this generalized framework to a three-stage stochastic program for the travel time resilience of roadway networks. System performance is the total travel time to serve a given OD demand. The travel times and capacities on individual roadway links are random variables that depend on the occurrence of a disaster. Disasters are characterized in terms of type, location, and consequences (imposed on link travel times and capacities), and the problem as a whole evolves in three stages. In the first stage, the possible disaster type, location, and consequences are known probabilistically. Decision variables in this first stage include capacity expansion and retrofit (protection of links). In the second stage, the disaster event type and location are known, but consequences are probabilistic. Decision variables in the second stage are preparedness actions that affect the efficiency of postdisaster response actions. In the third stage, the disaster event has occurred and consequences are known. Decisions at this stage include response actions and postresponse traffic flows. System users are assumed to select their routes unselfishly, and traffic demand is assumed to be unchanged following the disaster. Link travel times are estimated using the BPR function for congestion and approximated as a piecewise-linear function. It is assumed that a UE exists before the disaster and once recovery has happened, and a partial user equilibrium (PUE)—in which only affected users reconsider their routes—adequately captures driver route-choice postdisaster. They present a progressive hedging algorithm to solve the problem, noting that Benders decomposition is too computationally intensive. They illustrate the technique with a

small example involving six nodes and 16 links, involving four OD pairs of traffic, and considering three disaster scenarios (earthquake, flood, and malicious acts). They solve this example using GAMS (GAMS Development Corp 2013) on a personal computer, obtaining optimal solutions in three to four hours.

2.2. Sequential Games for Infrastructure Defense Analysis

We use a sequential-game model called a DAD model for assessing and improving system resilience against disruptions. The model, first defined in Brown et al. (2006), is a three-stage Stackelberg game (e.g., Luh, Chang, and Chang 1984) whose name originates with our assumption that there is a decision maker corresponding to each stage of the game. We associate the first decision with a *defender* of an infrastructure system, i.e., the individual or group that is charged with improving the operational resilience of the system. A defender could be a system operator, owner-company executives, or a group of policy makers. First-stage decisions suggest structural changes for the system of interest, such as adding a new system component that contributes new or redundant capability, modifying an existing component (e.g., increasing its capacity), or protecting a component to reduce the impact of any damage to that component, subject to budgetary constraints. We refer to these decisions collectively as the *defense options*, and any set of defense options that can be selected simultaneously is a *defense*.

We associate the second decision with an *attacker* representing a hypothetical, intelligent adversary who uses a plausible set of resources and capabilities to carry out a worst case *attack* (a deliberate, disruptive event) to damage or destroy components of the modified system; the attacker can observe the consequence of any first-stage defense decision and, accordingly, the resulting changes made to the operational system. Throughout, “worst case” means that system functionality is maximally reduced or, equivalently, the system operating cost is maximized. We are *not* trying to predict the decision of any specific attacker; rather, we are using the notion of an omniscient, rational attacker as a means to discover operational vulnerabilities. In practice, an attacker may not be smart enough to act optimally or may have motives that lead him to attack in a less-than-optimal fashion in the view of the defender and operator. Trying to build a model of such a situation would force us to represent an attacker’s intent, his intelligence, or even his modeling skills. Such models are simply impossible to validate (see Brown and Cox 2011a, b), and we believe it is more prudent to assume a worst case.

We associate the third decision with an *operator* (or “defender-as-operator”) who selects among a set of feasible operational decisions to minimize the operating cost of the defended, and subsequently attacked,

system. The notion of an operator is certainly appropriate for systems like an electric power grid controlled by an independent system operator (“ISO”; see O’Neill et al. 2006), but for simplicity we also apply the term to system optimization or equilibrium models that represent the interactions of the system’s users, perhaps influenced by centralized or decentralized controls.

When working to minimize the operating cost of the attacked system we assume the system operator is indifferent to the source of disruptive events. In practice, real infrastructure owners and operators must routinely contend with disruptive events caused by Mother Nature, accidents, failures, and sometimes even deliberate acts; in each of these cases the operator’s first priority is maintaining the best possible system performance for the degraded system. In reality, a worst-case disruption might be caused by a combination of deliberate and/or nondeliberate events that leads to a “perfect storm,” and initiating these events might be beyond the capability of any single specific attacker. The intent of modeling the attacker is to identify the disruptive events that lead to worst-case system performance, so that the operator understands in advance what is possible (even if history or intelligence suggests such combinations of events are unlikely) and to prepare appropriately for those possibilities.

As described in Alderson et al. (2013), the DAD formulation builds on a class of *system interdiction models* that have their roots in military planning problems; see also Danskin (1967); Wollmer (1964, 1968); McMasters and Mustin (1970); Ghare, Montgomery, and Turner (1971); Fulkerson and Harding (1977); Golden (1977); and Wood (1993). For applications and additional background, we refer the reader to Wood (2011); Cappanera and Scaparra (2011); Alderson et al. (2011); Scaparra and Church (2012); Dimitrov and Morton (2013); Alguacil, Delgado, and Arroyo (2014); and Yuan, Zhao, and Zeng (2014). These models have been applied to a variety of infrastructure domains, including electric power grids (Salmerón, Wood, and Baldick 2009), facility-location problems (Church and Scaparra 2006), supply chain networks (Snyder et al. 2006), and telecommunications networks (Murray, Matisziw, and Grubescic 2007).

Notably, Bell et al. (2008, p. 1897) apply the DAD concept to identify a risk-averse routing strategy for the traversal of a single OD path across the road network in London, under the assumption that “link travel cost is directly influenced by link status and is not a function of link flow.” That is, the traversal cost of an individual disrupted road segment is assumed to be a scalar multiple of that under normal conditions, where these costs are precomputed (i.e., traffic congestion resulting from rerouting after the failure of individual road segments is not explicitly considered). Alderson et al. (2011) describe a DAD analysis of a road network that

is similar to the one in this paper, but apply a simplified operating model and test only a small, hypothetical example.

2.3. Contributions of the Current Work

Building on the recent tutorial by Alderson, Brown, and Carlyle (2014), the current paper addresses the challenge of assessing and improving regional roadway resilience using the DAD perspective. We present an operator model that reflects traffic routing to satisfy OD travel demands according to a system optimization that minimizes overall travel time. Travel times follow the BPR function for congestion, and the model allows for individual travelers to stay at home if their overall travel time is excessive. The modeled impact of a disruption on an individual road segment is different for highway segments (for which alternate surface routing might be available) and bridges or tunnels (where alternate surface routes are unlikely to exist). We present a multiperiod model that reflects nonuniform recovery times for individual roads, bridges, and tunnels following a disruption. We solve for the worst-case loss of one or more highway, bridge, or tunnel segments and identify the defensive investments that mitigate such losses in the best possible way.

We apply this model to a full-scale regional highway network in the United States using recent census data on commuters and a road-network model based on parameters endorsed by the relevant regional transportation agency. Our representation has 91 nodes, 133 edges, 266 arcs, and 8,190 OD pairs. We solve the complete trilevel DAD model using nested, generalized Benders decomposition.

We believe this paper presents the most complete description of a realistic application of DAD available in the literature, and that it can serve as a template for future studies seeking to improve the resilience of an infrastructure system to worst-case losses.

3. DAD for a Highway Network

We seek a method to identify a budget-limited design that modifies the regional highway infrastructure system to maximally improve that system’s operational resilience to worst-case losses. We use a hypothetical attacker to discover worst-case losses. The model has several key components.

System operation and performance measure. Given estimated demands for point-to-point travel during periods of peak demand, over a specified time horizon, the operating decisions for the system (specifically, the routes of individual travelers) are computed for each morning commute period. We use a standard model of traffic flow, extended appropriately, and measure performance as total vehicle-hours of travel (vht), plus certain penalties. We assume that the morning commute

period is much like the evening commute period and that nonpeak demand for travel is relatively low and will be accommodated with relatively little difficulty. As in Petersen (1975) and other papers, we call this cost *delay* (or “system delay”) although it means total travel time, as opposed to that quantity less nominal, total travel time under ideal conditions (which would be called “incremental delay”).

As in other recent work, we assume that users choose their routes unselfishly. That is, we use an SO model of traffic assignment that minimizes the total system delay given available routes for OD traffic. Modern metropolitan highway systems are thoroughly instrumented with traffic sensors and displays of traffic conditions, obstructions, and estimated transit times to destinations. Real-time traffic information and automated routing from services such as Google Maps (Google 2016) and Waze (2016) are widely used. Collectively, with the advent of driverless vehicles, SO models are becoming more realistic representations of traffic. Moreover, we are interested in assessing the capability of a network to respond to disruptive events, so a best-case response from an SO model is appropriate as an upper bound on what is possible.

Attacks. An *attack* is a shock that damages one or more *targets*, each target being a road segment, bridge, or tunnel. For simplicity, we ignore the possibility of damaging highway interchanges, so every target is an *edge* $[i, j]$ connecting two *nodes* i and j in a network.

We are interested in the resilience of an infrastructure system not to a random shock, but to one that yields the “worst” possible consequence. When evaluating the operational resilience of an infrastructure system, it is prudent to ask the question, “How bad could things be?” One way to answer this question is to consider the vulnerability of system performance to the targeted loss of sets of components.

In what follows, we consider the role of a hypothetical, intelligent adversary called the *attacker* who deliberately chooses component losses for the system. The goal of our hypothetical attacker is to discover the set(s) of components that, if lost, hurt the performance of the system in the worst possible way. We refer to these sets as the *worst-case component losses*, or simply *worst-case losses*. The importance of understanding worst-case losses follows from several observations.

- The performance of infrastructure systems is typically more vulnerable to the loss of targeted sets of components than to a random set of components. This is because of the specific way in which system components work together to achieve performance, and it is nontrivial to determine which sets of components are “most vital” (see Alderson et al. 2013, for a discussion). It is often the case that the simultaneous loss of a small number of seemingly insignificant components can be devastating to system performance.

- Infrastructure systems do face threats of deliberate attack (e.g., the April 2013 attacks on the Pacific Gas and Electric Metcalf electric substation in San Jose, California), and in general these systems have not been designed to be resilient to deliberate threats. For example, electric power transformers are often constructed from high-quality ceramics that provide a mean time between random failure of hundreds of years. However, these same components can be easily destroyed by a high-powered rifle.

- It is the losses of system components that matter, more so than the process that leads to those losses. In practice, infrastructure owners and operators are often agnostic to the source of a disruption—accident, failure, natural disaster, or attack—once one or more components are lost, the operator’s focus is on doing the best she can to maintain function with whatever is left of the system (see Alderson, Brown, and Carlyle 2015, for a discussion).

- Strictly speaking, the worst-case situation for an infrastructure system would be the simultaneous loss of *all* system components. However, in practice our interest is understanding the worst among a specific set of possible losses; thus, our notion of worst-case losses is conditioned on a given set of possibilities. Specifically, we introduce the notion of the *attacker’s resource(s)*, which simply reflects a plausible limit on the number and type of targets that a coordinated attack might damage, explained in detail below.

In summary, the use of a hypothetical attacker provides a mechanism that is convenient, both conceptually and mathematically, for discovering worst-case losses. We emphasize that our goal is the discovery of worst-case losses and their implication for the operational resilience of the infrastructure system, not the intent, values, or behavior of any particular attacker (see Keeney 2007, for an example of the considerable effort already devoted to this).

Repair epochs. Following an attack that damages one or more targets, we assume that traffic patterns will adjust in accordance to our model of traffic flow to satisfy the original OD demands as best as possible given available road segments. However, in general we should expect that the amount of time required to restore the functionality of different components will differ, perhaps significantly. Our interest is in understanding the performance of the system, not just in the immediate aftermath of the attack but *as the system recovers*. To capture this, our model includes *repair epochs*, discrete intervals of time (perhaps different in duration) that collectively reflect the points in time where individual components will be restored. Thus, our representation of system operation is a multi-period model that reflects the schedule of repairs in the days, weeks, or months following an attack. At each point along the way, we assume that traffic adapts to

take advantage of restored roadway segments as they become available.

System designs. The designs of interest apply to an existing system, and cover budget-limited defenses of potential targets, e.g., structural hardening of a tunnel, or access control via checkpoints to a bridge. A more general model might consider adding lanes to road segments, building new road segments or bridges, etc. (see Alderson, Brown, and Carlyle 2014, for a description of the mathematics to support general design options), but this paper only covers designs that can be represented as *defenses* of edges.

3.1. Full Formulation

This section presents a full DAD model for a regional highway network; we label this “**T-DAD**” for “transportation DAD model.” The three stages of **T-DAD** are difficult to see, so we clarify by showing how (a) assuming a null defense and no attacks, the model specializes to the underlying operational model, (b) adding a representation of optimized attacks to the operational model results in a two-stage Attacker-Defender (AD) model, and (c) adding a representation of optimized defenses to the AD model produces the full model **T-DAD**.

Sets and Indices:

- $i, j, p, n \in N$ nodes in a highway network;
- $[i, j] \in E$ undirected edge in a highway network ($i < j$);
- $(i, j) \in A$ directed arc in a highway network; if $[i, j] \in E$, then $(i, j) \in A$ and $(j, i) \in A$;
- $t \in T$ repair epoch in planning horizon;
- $d \in D$ defense options;
- $d \in D_{ij} \subseteq D$ defense options available for edge $[i, j] \in E$ (e.g., do nothing, add checkpoints);
- $d_0 \in D_{ij} \subseteq D$ null defense option that leaves edge $[i, j] \in E$ unchanged (i.e., undefended).

Data (units):

- b_{pj} for each commute period, b_{pp} denotes the total supply rate for vehicles originating at node p , and $-b_{pj}$, $j \neq p$, denotes the demand rate at j for vehicles originating at p (vehicles per hour, “vph”);
- u_{ijt} capacity of arc $(i, j) \in A$ during repair epoch t (vph);
- q_{ijt} penalty for traversing targeted arc $(i, j) \in A$ during epoch t (hours);
- q_{\max} maximum tolerable commute duration (hours);
- n_t number of commute periods during epoch t (periods);
- h duration of commute period (hours);
- h_t $h \times n_t$, total commute time t (hours);

- a_{ij} cost to target edge $[i, j] \in E$ as part of an attack (attack dollars);
- \bar{a} total attack resource (attack dollars);
- u_{ijdt} under defense option d in epoch t , the baseline capacity of arc $(i, j) \in A$ (vph);
- u'_{pijdt} upper bound on flow rate originating from p and traversing $(i, j) \in A$ under defense option d in epoch t (vph); in practice, $u'_{pijdt} = \min\{2u_{ijdt}, b_{pp}\}$;
- q_{ijdt} if arc $(i, j) \in A$ is targeted by an attack, the penalty (if any) for traversing that arc under defense option d during epoch t (hours);
- $F^0_{ijdt}(\cdot)$ travel time as a function of traffic flow rate on arc $(i, j) \in A$, if that arc is not targeted by an attack under defense option d during epoch t (hours);
- $F^1_{ijdt}(\cdot)$ travel time on arc $(i, j) \in A$ that is targeted by an attack under defense option d during epoch t as a function of traffic flow rate (hours);
- c_{ijd} cost to defend edge $[i, j] \in E$ using defense option $d \in D_{ij}$ (defense dollars);
- \bar{c} total defense budget (defense dollars).

Decision Variables (units, when appropriate):

- w_{ijd} 1 if edge $[i, j] \in E$ is defended using defense option $d \in D_{ij}$, and 0 otherwise;
- x_{ij}, x_{ji} 1 if edge $[i, j] \in E$ is attacked, and 0 otherwise (Note: Defining x_{ji} simplifies notation here, but $x_{ji} \equiv x_{ij}$ for all $[i, j] \in E$, and is substituted out in practice.);
- y'_{pijdt} for each commute period during epoch t , the flow rate for traffic that originates from p and that traverses arc $(i, j) \in A$ under defense option $d \in D_{ij}$ (vph);
- y''_{ijdt} for each commute period during epoch t , the total flow rate for traffic on arc $(i, j) \in A$ under defense option $d \in D_{ij}$ (vph);
- y'''_{pjdt} for each commute period during epoch t , the rate for unmet demand (“dropped demand”) having origin p and destination j (vph).

Formulation of T-DAD(w, x, y)

$$z^*_{\text{DAD}} \equiv \min_{\mathbf{w}} \max_{\mathbf{x}} \min_{\mathbf{y}', \mathbf{y}'', \mathbf{y}'''} \left\{ \sum_{\substack{(i,j) \in A, \\ d \in D_{ij}, t \in T}} h_t (y''_{ijdt} F^0_{ijdt}(y''_{ijdt}) (1 - x_{ij}) + y''_{ijdt} F^1_{ijdt}(y''_{ijdt}) x_{ij}) + \sum_{\substack{(i,j) \in A, \\ d \in D_{ij}, t \in T}} h_t q_{ijdt} y''_{ijdt} x_{ij} + \sum_{\substack{p, j \in N, p \neq j, \\ t \in T}} h_t q_{\max} y'''_{pjdt} \right\} \quad (1)$$

$$\text{s.t.} \quad \sum_{\substack{j | (n, j) \in A, \\ d \in D_{nj}}} y'_{pnjdt} - \sum_{\substack{i | (i, n) \in A, \\ d \in D_{in}}} y'_{pindt} - y'''_{pnt} = b_{pn} \quad \forall p, n \in N, p \neq n, t \in T, \quad (2)$$

$$\sum_{p \in N-j} y'_{pijdt} - y''_{ijdt} = 0 \quad \forall (i, j) \in A, d \in D_{ij}, t \in T, \quad (3)$$

$$\sum_{p \in N} y'_{pijdt} \leq u_{ijdt} w_{ijd} \quad \forall [i, j] \in E, d \in D_{ij}, t \in T, \quad (4)$$

$$\sum_{p \in N} y'_{pjidt} \leq u_{jidt} w_{ijd} \quad \forall [i, j] \in E, d \in D_{ij}, t \in T, \quad (5)$$

$$\sum_{[i, j] \in E} a_{ij} x_{ij} \leq \bar{a}, \quad (6)$$

$$\sum_{\substack{[i, j] \in E, \\ d \in D_{ij}}} c_{ijd} w_{ijd} \leq \bar{c}, \quad (7)$$

$$\sum_{d \in D_{ij}} w_{ijd} = 1 \quad \forall [i, j] \in E, \quad (8)$$

$$x_{ij} \in \{0, 1\} \quad \forall [i, j] \in E, \quad (9)$$

$$x_{ji} \equiv x_{ij} \quad \forall [i, j] \in E, \quad (10)$$

$$w_{ijd} \in \{0, 1\} \quad \forall [i, j] \in E, d \in D_{ij}, \quad (11)$$

$$\mathbf{y}' \geq \mathbf{0}, \quad \mathbf{y}'' \geq \mathbf{0}, \quad \mathbf{y}''' \geq \mathbf{0}. \quad (12)$$

Our notation **T-DAD**($\mathbf{w}, \mathbf{x}, \mathbf{y}$) emphasizes the sequence of decisions associated with defenses \mathbf{w} , attacks \mathbf{x} , and operation \mathbf{y} , and allows us to specify fixed values for some of these decision variables. The rationale behind the form for this trilevel model is presented in the tutorial by Alderson, Brown, and Carlyle (2014). In the remainder of this section, we explain **T-DAD** by extracting and discussing its “submodels.”

3.2. The Operational Model

The need to track defense options d and repair epochs t makes **T-DAD** complicated. To understand the underlying operational model, consider the following (temporary) simplifications:

1. $T = \{t_0\}$, i.e., we consider only a single repair epoch;
2. $n_{t_0} = 1$, i.e., epoch t_0 contains only a single commute period;
3. $w_{ijd_0} \equiv 1$ for all $[i, j] \in E$, i.e., only the null defense option is relevant;
4. $x_{ij} \equiv 0$ for all $[i, j] \in E$, i.e., no attack is carried out.

Then, letting $\hat{\mathbf{w}}_0$ denote the vector of null defenses and $\mathbf{0}$ denote the vector of zeroes to represent the case of no attacks, **T-DAD** simplifies to

T-DAD($\hat{\mathbf{w}}_0, \mathbf{0}, \mathbf{y}$):

$$\min_{\mathbf{y}, \mathbf{y}'', \mathbf{y}'''} \left\{ \sum_{(i, j) \in A} h_{t_0} y''_{ijd_0 t_0} F^0_{ijd_0 t_0}(y''_{ijd_0 t_0}) + \sum_{p, j \in N, p \neq j} h_{t_0} q_{\max} y'''_{pj t_0} \right\} \quad (13)$$

s.t. (2)–(5), (12).

With fixed defenses and zero attacks, we drop the respective min and max operators for the outer two problems. The first term in the objective function (13) evaluates system delay (vht) by summing, for each arc (i, j) in the network, [hours in the commute period] \times

[vehicle flow rate on (i, j) during that period] \times [the travel time on (i, j) given the flow rate]. In practice, we use BPR travel times for $F^0_{ijdt}(\cdot)$ —Appendix A describes the generic travel-time function $F_{ij}(y''_{ij})$ and the implementation of the total travel-time function $y''_{ij} F_{ij}(y''_{ij})$ using a piecewise-linear approximation—but the framework is general and could accommodate other travel time models.

To allow for the extreme conditions that might occur after a severe shock—for instance, all unattacked paths between two nodes might actually be eliminated—we modify the standard model slightly by incorporating a second term in (13) that represents a penalty for “dropped demand” (i.e., commuters who elect to stay at home rather than travel). Specifically, we incorporate elastic flow-balance constraints that allow for demand to go unmet given that a sufficiently high penalty is paid. In our case study to follow (see Section 4), we set the penalty for this dropped demand to six hours per commuter who stays home, assuming that a commuter will choose to stay at home if the best available commute time exceeds three quarters of the length of a standard workday. (The nominal, inelastic model assumes that constant demand rates for point-to-point travel are presented to the system over a four-hour period, but does not limit observed travel times. This “six-hour safety valve” could be seen as an alternate OD path with a fixed length, and so we refer to our model as an “enhanced” model that has an implicit six-hour bound on travel times because of these alternate paths.)

For a single period of time, and for each possible origin of traffic, each constraint (2) maintains “balance of vehicle flow rate” at a node, adjusted for dropped demand. Each constraint (3) accumulates the total flow rate on an arc in each time period, so that the objective function may be expressed succinctly; these variables may be substituted out, however. Constraints (4) and (5) fix certain variables to 0 or these constraints are assumed slack and can thus be omitted. Stipulations (12) enforce nonnegativity of decision variables.

3.3. An Attacker-Defender Model

Next, keeping $\mathbf{w} \equiv \hat{\mathbf{w}}_0$, $T = \{t_0\}$, and $n_{t_0} = 1$, we allow \mathbf{x} to vary

$$\mathbf{x} \in X \equiv \{x_{ij} \mid \forall (i, j) \in A \mid (6), (9), (10)\}. \quad (14)$$

This results in the AD model **T-DAD**($\hat{\mathbf{w}}_0, \mathbf{x}, \mathbf{y}$)

$$\max_{\mathbf{x} \in X} \min_{\mathbf{y}, \mathbf{y}'', \mathbf{y}'''} \left\{ \sum_{(i, j) \in A} (h_{t_0} y''_{ijd_0 t_0} F^0_{ijd_0 t_0}(y''_{ijd_0 t_0})(1 - x_{ij}) + h_{t_0} y''_{ijd_0 t_0} F^1_{ijd_0 t_0}(y''_{ijd_0 t_0}) x_{ij}) + \sum_{(i, j) \in A} h_{t_0} q_{ijd_0 t_0} y''_{ijd_0 t_0} x_{ij} + \sum_{p, j \in N, p \neq j} h_{t_0} q_{\max} y'''_{pj t_0} \right\} \quad (15)$$

s.t. (2)–(5), (12).

With fixed defenses we drop the outermost min operator. The solution to this model identifies a worst-case attack against the nominal system. The variables $\mathbf{x} \in X$ primarily control which “version” of the travel-time function applies: $F_{ijd_0t_0}^0$ if $x_{ij} = 0$, and $F_{ijd_0t_0}^1$ if $x_{ij} = 1$. They also impose an additive penalty $q_{ijd_0t_0}$ to the travel time for arc (i, j) if $x_{ij} = 1$.

These two mechanisms are used to increase the travel time on an attacked arc in the following manner. As described in Appendix A, we use a piecewise-linear approximation to the BPR function for both F_{ijdt}^0 and F_{ijdt}^1 . However, we parameterize function F_{ijdt}^1 so that it represents a 50% reduction in lane capacity and a 20% reduction in free-flow speed when compared with F_{ijdt}^0 . Thus by construction, we have $F_{ijd_0t_0}^1(y''_{ijd_0t_0}) > F_{ijd_0t_0}^0(y''_{ijd_0t_0})$ for $y''_{ijd_0t_0} > 0$. Moreover, an attacked arc is not only slower but also more sensitive to congestion. The idea is that a road segment might lie adjacent to surface streets onto which traffic could be rerouted locally, at some penalty, in the event of an attack. Thus the attacked segment is still “passable” but expensive from a travel-time perspective.

Alternatively, we recognize that some segments, such as a bridge, might be impossible to traverse or bypass if targeted in an attack. In such cases, we use the additive penalty q_{ijdt} , which in practice we set to q_{\max} (six hours), making it effectively infinite because it will always be less costly for a traveler to stay at home (be dropped) than to attempt to traverse an attacked bridge. From this point on, we refer to all such impassable edges, bridges, or otherwise, as *bridge-like*; all others are *standard edges*.

Thus, attacks do not restrict flow via capacity constraints (4) and (5). Rather, we rely entirely on this form of *cost-based interdiction* because it avoids computational complications if the incumbent solution to system operation suddenly becomes infeasible as a result of an attack (see Alderson, Brown, and Carlyle 2014, for a discussion).

3.4. The Full DAD Model

Next, we release \mathbf{w} to

$$\mathbf{w} \in W \equiv \{w_{ij} \forall [i, j] \in E \mid (7), (8), (11)\}, \quad (16)$$

and obtain the three-stage model **T-DAD** = **T-DAD**(\mathbf{w} , \mathbf{x} , \mathbf{y}). In effect, the variables $w_{ij\hat{d}}$ control which versions of the arcs (i, j) and (j, i) the operator must use, and the variables x_{ij} determine which travel-time functions apply. For example, if $w_{ij\hat{d}} = 1$, then only the variables $y_{ij\hat{d}t}$ and $y_{ji\hat{d}t}$ can take on positive values, and thus only the functions $F_{ij\hat{d}t}^0(y_{ij\hat{d}t})$, $F_{ij\hat{d}t}^1(y_{ij\hat{d}t})$, and $F_{ji\hat{d}t}^1(y_{ji\hat{d}t})$ are relevant. If defense option $\hat{d} \in D_{ij}$ completely protects $[i, j]$ from attack, then $F_{ij\hat{d}t}^1(\cdot) = F_{ij\hat{d}t}^0(\cdot)$ and $F_{ji\hat{d}t}^1(\cdot) = F_{ji\hat{d}t}^0(\cdot)$. If not, then we

expect that $F_{ij\hat{d}t}^1(y_{ij\hat{d}t}) > F_{ij\hat{d}t}^0(y_{ij\hat{d}t})$ for any $y_{ij\hat{d}t} > 0$ as before.

Each defense option d has its own set of data. Each edge has a special null defense (d_0) that represents “doing nothing,” i.e., leaving the arc in its original condition. The data associated with this option are nonzero; the corresponding arcs have nominal capacity, travel time, etc. However, if we choose a defense option that means “do something special to defend this edge” then the attack penalty will be lower, but the cost might increase, or the capacity might drop, etc. We handle the dependence between the flow rate and the attack through the travel-time functions F_{ijdt}^0 and F_{ijdt}^1 . Constraints (8) ensure that a single defense option is chosen for each edge.

The full model allows for the possibility of an attack that causes structural damage to edges, and that the time to repair such damage can vary by edge, so repair epochs $t \in T$ now come into play. Model **T-DAD**(\mathbf{w} , \mathbf{x} , \mathbf{y}) is instrumented so that the repairs to individual edges are independent and follow a predetermined schedule. Specifically, if an attack would not be repaired by epoch t then, in general, $F_{ijd_0t}^1(y''_{ijd_0t}) > F_{ijd_0t}^0(y''_{ijd_0t})$ for $y''_{ijd_0t} > 0$; otherwise $F_{ijd_0t}^1(y''_{ijd_0t}) = F_{ijd_0t}^0(y''_{ijd_0t})$ for all $y''_{ijd_0t} \geq 0$. If an arc has been completely defended by defense option $\hat{d} \in D_{ij}$, then $F_{ij\hat{d}t}^1(\cdot) = F_{ij\hat{d}t}^0(\cdot)$ and $F_{ji\hat{d}t}^1(\cdot) = F_{ji\hat{d}t}^0(\cdot)$ for all $t \in T$.

Remark 1. Commensurate restoration cost could be added to total system operating cost, if desired. If we estimate defense option d for edge $[i, j]$ incurs restoration cost c_{ijdt}^R during repair epoch t , then the following term, linearized and scaled, could be added to the objective function (1) for **T-DAD**:

$$\sum_{[i, j] \in E, d \in D_{ij}} c_{ijdt}^R w_{ijdt} x_{ij}. \quad (17)$$

Converting units of dollars into believable units of vehicle-hours traveled, or vice versa, might be difficult, but the dollar cost of congestion could be estimated (e.g., Weisbrod, Vary, and Treyz 2003) and those could be combined with the dollar cost of restoration.

3.5. Solving T-DAD

We have explored a number of techniques for solving **T-DAD**. All are based on converting the inner AD model into a mixed-integer program using the technique described by Cormican, Morton, and Wood (1999). Given that the AD “subproblems” can be solved by standard LP-based branch and bound, and given the modest number of defenses and attacks that might arise in our situation, **T-DAD** can be solved in a number of ways, including the following:

1. Applying the decomposition method described in Alderson et al. (2011).

2. Viewing **T-DAD** as a simple DA problem, and applying the “covering decomposition” described by Israeli and Wood (2002) for a network-interdiction problem.

3. Applying the implicit-enumeration technique described by Alguacil, Delgado, and Arroyo (2014) for defending an electric power grid.

We combine methods 1 and 2 to solve the case-study problems described in Section 4. The algebraic modeling system GAMS (GAMS Development Corp 2013) generates models and implements the decomposition algorithms, with master problems and subproblems solved by CPLEX, version 12.6.0 (IBM 2014). The operator’s problem has more than 200,000 variables and 40,000 constraints and takes more than a second to solve on a Windows-based Lenovo W530 laptop computer with 32 GB RAM and using a single thread.

As an example of empirical computational efficiency, we note that the 15 DAD models in Table 9 require a total of 6.5 hours of computation time. By contrast, the results in Table 9 would require solving the operator model approximately 1.17×10^{14} times to achieve the same results by exhaustive enumeration. If we could solve the operator model in one millisecond, this would still take more than 3,700 years.

4. A Case Study: The Major Highway Bridges, Roads, and Tunnels of the San Francisco Bay Area

We present an illustrative case study that investigates the resilience to worst-case losses of a regional highway system, specifically, the major roads, bridges, and tunnels of the San Francisco (SF) Bay Area. Several caveats not already mentioned apply: (a) data modeling was carried out in early 2013, so data reflect the infrastructure system at that time (for example, the Caldecott Tunnel is modeled as having three operating bores, rather than the four it now has); (b) we do not account for how mass transit (buses, the “BART” train system, ferries, etc.) might alleviate some traffic congestion after an attack; (c) we do not account for the possibility of converting nominally unidirectional roadways into bidirectional roadways to alleviate congestion after an attack; (d) we ignore random events such as traffic accidents and emergency lane closures; (e) we ignore any error induced by the piecewise-linear approximating functions used in the operating model; and (f) we implement the attack-resource constraint (6) as a cardinality constraint on the number of edges attacked rather than as a (set of) general knapsack constraint(s).

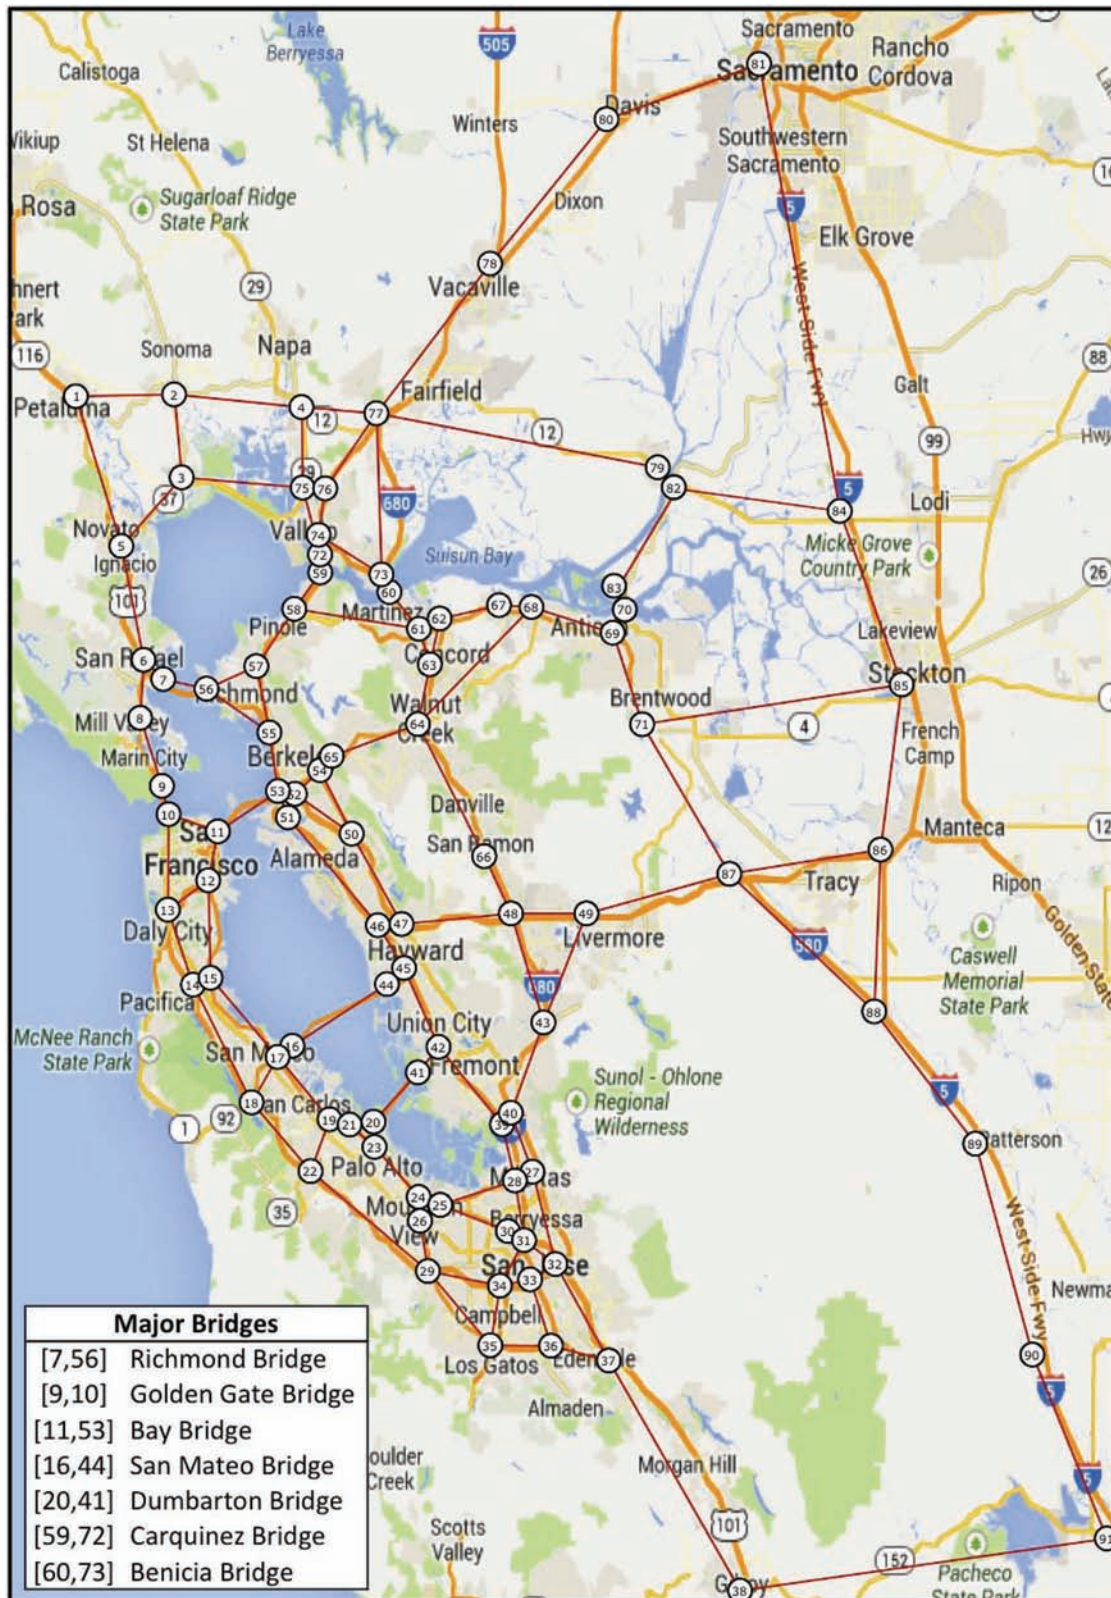
Our operational transportation model covers the San Francisco Bay Area with a network of 91 nodes and 133 edges; see Figure 1. This study models the morning commute period, between 5 A.M. and 9 A.M., as one

of the two periods of highest demand on the system. We consider only the morning commute to simplify exposition and to reduce the data-processing workload in this example. The road network represents all of Alameda, Contra Costa, Marin, San Francisco, San Mateo, and Santa Clara counties, and parts of Napa, San Joaquin, Solano, and Sonoma counties. Commuter data cover the populations of those counties as well as the populations of the peripheral counties of Merced, Sacramento, San Benito, and Santa Cruz.

For simplicity, all bridges are treated as single edges, even though some bridges have parallel spans that might be viewed as separate targets. Similarly, all bores of the Caldecott Tunnel are treated as a single edge. (For the morning commute, the tunnel has four lanes westbound and two lanes eastbound.) We do not separately model the arcs for the Waldo Tunnel on U.S. Highway 101 (US 101) north of the Golden Gate Bridge (edge [8,9]), because the tunnel has a single bore with fixed lanes that do not differ from the highway.

We extract highway-network data from Google Earth imagery and compute travel distances between nodes using the Google Maps calculator (Google 2016). Data on the commuting population derive from (a) the Quarterly Workforce Indicators for 2009 (U.S. Census Bureau 2009b), (b) Table 32226C from the Census Transportation Planning Products for 2009 (U.S. Census Bureau 2009a), and (c) Table 32106 from the same source. Data source (a) identifies the home and work locations, by census tract, for about 90% of the population in any region. Data source (b) provides an estimate for the fraction of commuters who commute by car during the morning commute period of 5:00 A.M. to 9:00 A.M., and data source (c) provides an estimate of the number of passengers per vehicle depending on origin and destination counties. We ignore truck traffic, assuming this to be light during commute hours (Hallenbeck et al. 1997), and make no attempt to adjust for the workers not covered by the Quarterly Workforce Indicators (e.g., federal government employees). With certain exceptions, a commuter living in one census tract and working in another contributes a fraction of a vehicle (based on carpooling rates) to the demand for traffic between the two nodes closest to the population centroids of those two tracts. If these are the same node, we assume that the trip for that commuter is short enough that a highway is not used, and therefore does not contribute to traffic demand. Our final estimates have 834,200 vehicles traveling through the road network during the morning commute period when there is no major traffic accident or any significant damage to any of the infrastructure in the area.

Figure 1. (Color online) The Roads and Bridges of the San Francisco Bay Area



Notes. A network of major highway and seven bridge edges connects 91 nodes, numbered for identification. Each node can be both an origin and a destination for vehicular traffic. For simplicity, the edges representing roads, bridges, and tunnels are drawn as straight connections, and some node locations have been perturbed to avoid overlaps that hide smaller features.

Source. Background image from Google Maps.

4.1. Nominal Morning-Commute Traffic in the SF Bay Road Network

In the absence of an attack, vehicles may use any of the edges. The baseline solution of $\mathbf{T-DAD}(\hat{\mathbf{w}}_0, \mathbf{0}, \mathbf{y})$, which represents the flows in the nominal network, calculates an optimal delay of 0.797×10^6 vht for a single morning-commute period (approximately 57 minutes per vehicle), and no demand is dropped. (See the data extracts in Appendix B.) We cannot illustrate the individual routes followed between each OD pair, but Figure 2 does provide a graphical indication of traffic density during the morning commute for the undamaged road network.

Table 1 lists the 30 most-congested arcs identified in the baseline results, ranked by delay. Our validation work shows model-estimated traffic congestion patterns closely match observed patterns. For example, a report by the Metropolitan Transportation Commission (2013) lists the top 10 traffic “hotspots” in the SF Bay Area for 2013, ranked by daily (weekday) added delay measured in vht. This list covers both morning and evening commutes, and one hotspot can actually correspond to multiple arcs in our model, or multiple hotspots can correspond to a subset of one of our network’s arcs.

All 10 top hotspots on this 2013 list correspond to arcs within the first 24 entries in Table 1. Our data have approximately twice the number of arcs indicated (graphically) in the 2013 report, so our top 24 correspond quite well to the top 10 in that report. The vehicle-hours reported here are 10%–30% lower than those reported in the 2013 report, but we do not account for accidents, temporarily closed lanes, repairs, and construction that occur from time to time in a real road network.

4.2. Basic Modeling of Attacks and Defenses

Initial testing considers only two defense options: $d_0 \in D_{ij}$ means “do not defend” and $d_1 \in D_{ij}$ means “defend.” A defended edge becomes invulnerable to attack, i.e., we set the penalty $q_{ij,d,t} = 0$.

We assume that an edge corresponding to a bridge, or bridge-like edge, cannot be traversed or bypassed if targeted in an attack, so the corresponding travel-time function is modified by adding a penalty of six hours, which is effectively infinite, to the nominal travel time as discussed for Equation (15). We have also identified several edges in the SF Bay Area that are not bridges but, because of their location, allow little rerouting of the traffic they carry onto local streets. In particular, we model the Caldecott Tunnel (edge [54, 65] in Figure 1), the section of I-680 connecting Fremont to Sunol (edge [40, 43]), and SR 37 connecting Sears Point to Vallejo (edge [3, 75]) as bridge-like, in that an attack on any one of these edges renders it unusable.

To demonstrate attackers with differing technological capabilities in $\mathbf{T-DAD}$, we consider two extreme

cases. We assume a *low-capability attacker* can only remove an edge from service for a single commute period, perhaps by releasing a hazardous material on a roadway or simply blocking traffic (e.g., protesters in Berkeley, California recently did exactly that; see Pogash and Grady 2014).

We assume a *high-capability attacker* can damage edges structurally, resulting in repair epochs summarized in Table 2. Specifically, “simple” highway segments (i.e., standard edges) require only one month to repair, but a bridge could take years to repair, depending on its construction. We model this situation using repair epochs $t = 1, \dots, 4$ such that the most easily repaired edges are repaired after n_1 commute periods, the second easiest after $n_1 + n_2$ commute periods, etc. The disruption caused by an attack now depends on the repair times and the corresponding repair epochs. The first epoch has a duration of 1 month (22 commute periods), the second has a duration of 11 months (242 commute periods), and the last two each have durations of 24 months (528 commute periods). The total time horizon under consideration in this paper is therefore five years (1,320 commute periods).


4.3. The Most-Damaging Attacks to the Nominal Network

We assume no defenses, and consider a hypothetical intelligent adversary who can mount an attack targeting one or more edges in the San Francisco Bay Area highway network. The adversary seeks to maximize the network’s operating cost (system delay).

4.3.1. Single-Period Losses from a Low-Capability Attacker. To get a better sense of the relative importance of segments individually and in combination, we begin by simply identifying the 10 most damaging single-edge attacks ($a_{ij} = 1$ for all $[i, j] \in E$, $\bar{a} = 1$) and the 10 most damaging two-edge attacks ($a_{ij} = 1$ for all $[i, j] \in E$, $\bar{a} = 2$) for a low-capability attacker who can cause losses for only a single epoch. Because the number of attacks is small, it may be possible to produce these results by enumerating solutions to the operational model, but we actually solve $\mathbf{T-DAD}(\hat{\mathbf{w}}_0, \mathbf{x}, \mathbf{y})$ for the most damaging attack, add a solution-elimination constraint (e.g., Alderson et al. 2011), solve for the second most damaging attack, and so on. The key purpose of investigating these scenarios is to demonstrate the varying superadditive effects of lost edges.

Table 3 presents the 10 most-damaging single-edge losses for a low-capability attacker. Eight of the top 10 single-edge losses are bridge-like edges. Losing the Bay Bridge connecting San Francisco and Oakland (edge [11, 53]) increases the system delay by nearly 29%; the bridge-like edge [40, 43], part of I-680, is second worst, with a 24% increase; the Golden Gate Bridge (edge [9, 10]), the cultural icon of the San Francisco Bay, ranks

Figure 2. (Color online) Traffic Flows in the San Francisco Bay Area



Notes. (A) The baseline flows during a normal commute period. Here, "saturation" is the ratio of traffic flow (vph) to arc capacity (vph). For each edge $[i, j]$, the diagram indicates the saturation and capacity on the arc with the highest saturation, (i, j) or (j, i) . (B) Traffic flows in response to an attack targeting the Golden Gate Bridge and Bay Bridge. (C) In response to an attack on those bridges, the number of "delivered" and "dropped" vehicles per commute period, by origin.

Table 1. Arcs Associated with the Top 30 Most-Congested Road Segments, Bridges, or Tunnels in the San Francisco Bay Area Under Morning-Commute Conditions

Rank	Road, bridge, or tunnel	From	To	Lanes	Speed limit (mph)	Speed (mph)	Observed delay (vht)
1	I-680	63_ConcordS	64_WalnutCreek	5	65	19.9	7,057.7
2	I-680	43_Sunol	40_FremontSE	4	65	26.8	6,202.1
3	Bay Bridge	53_BayBridgeE	11_BayBridgeW	5	65	25.9	6,071.8
4	I-80	55_Albany	53_BayBridgeE	5	65	21.7	5,649.8
5	I-880	46_SanLeandro	45_Hayward	5	65	21.4	5,264.8
6	US 101	38_Gilroy	37_CoyoteValley	4	65	41.0	5,184.7
7	US 101	05_Novato	06_SanRafael	4	65	30.6	5,070.3
8	I-580	87_Tracy	49_Livermore	4	65	35.6	4,972.1
9	I-80	58_Hercules	57_Richmond	4	65	25.7	4,497.7
10	SR 24	64_WalnutCreek	65_CaldecottTunnE	4	65	32.6	4,009.3
11	SR 4	69_AntiochOakley	68_Pittsburgh	2	65	21.7	3,889.2
12	I-580	48_Dublin	47_CastroValley	4	65	35.5	3,761.8
13	I-80	57_Richmond	55_Albany	4	65	26.9	3,468.6
14	US 101	01_Petaluma	05_Novato	3	65	35.6	3,418.3
15	I-280	29_Cupertino	22_Woodside	4	65	40.7	3,405.5
16	US 101	24_MountainViewN	23_PaloAlto	4	65	30.6	2,858.8
17	I-5	86_Manteca	85_Stockton	2	65	33.7	2,786.9
18	I-880	39_FremontSW	42_Fremont	4	65	36.7	2,782.4
19	I-580	49_Livermore	48_Dublin	4	65	35.6	2,486.1
20	I-880	45_Hayward	46_SanLeandro	5	65	31.9	2,399.6
21	US 101	12_SanFranciscoS	11_BayBridgeW	6	65	34.8	2,329.2
22	I-680	40_FremontSE	43_Sunol	4	65	40.7	2,244.5
23	US 101	08_MillValley	09_GGBridgeN	4	65	38.0	1,959.3
24	US 101	32_SanJoseE	31_SanJoseN	4	65	30.6	1,887.9
25	I-80	78_Vacaville	80_Davis	3	65	45.7	1,867.5
26	I-880	46_SanLeandro	51_Oakland	4	65	47.5	1,664.1
27	US 101	06_SanRafael	08_MillValley	4	65	35.6	1,657.4
28	SR 152	91_LosBanos	38_Gilroy	2	55	45.8	1,650.9
29	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS	4	50	22.8	1,511.9
30	Caldecott Tunn.	65_CaldecottTunnE	54_CaldecottTunnW	4	55	21.0	1,485.9

Note. For each arc, we list its name, two end points, number of lanes, speed limit (mph), calculated vehicle speed (mph), and observed delay (in vht for a single commute period).

Table 2. Repair Epochs Following a High-Capability Attack

Type	Road, bridge, or tunnel	Edge	Repair epochs			
			t_1 mos. 0–1	t_2 mos. 2–12	t_3 mos. 13–36	t_4 mos. 37–60
I	Golden Gate Bridge	[9, 10]	0	0	0	0
II	Bay Bridge	[11, 53]	0	0	0	1
	Carquinez Bridge	[59, 72]				
III	Benicia Bridge	[60, 73]	0	0	1	1
	Dumbarton Bridge	[20, 41]				
	Nejedly Bridge	[70, 83]				
	Richmond Bridge	[7, 56]				
	Rio Vista Bridge	[79, 82]				
	San Mateo Bridge	[16, 44]				
	IV	SR 37	[3, 5]	0	1	1
SR 37		[3, 75]				
SR 84		[43, 49]				
SR 152		[38, 91]				
Caldecott Tunnel		[54, 65]				
I-680		[40, 43]				
Kirker Pass Road		[64, 68]				
V	Standard edges	Other	R	1	1	1

Notes. For edge $[i, j]$ in epoch t , “0” indicates that $y''_{ijt} \equiv 0$, that is, the arc is unavailable for use; “1” indicates that arc $[i, j]$ is fully operational and that the travel-time function is defined with nominal parameters; “R” indicates that lane capacities u'_{ij} and u''_{ij} are reduced by 50% and free-flow speeds s_{ij} and s_{ji} are reduced by 20%. (See Appendix A.1 for a description of the travel-time function.)

Table 3. The Most Disruptive Single-Edge Losses in the San Francisco Bay Area, and Three Others, Assuming a Low-Capability Attacker

Rank	Edge lost			Sys. delay (vht × 10 ⁻⁶)	Delay increase (%)
	Road, bridge, or tunnel	Junction 1	Junction 2		
1	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.025	28.6
2	I-680	40_FremontSE	43_Sunol	0.985	23.6
3	Golden Gate Bridge	09_GGBridgeN	10_GGBridgeS	0.933	17.1
4	Caldecott Tunnel	65_CaldecottTunnE	54_CaldecottTunnW	0.916	14.9
5	Carquinez Bridge	72_CarquinezN	59_CarquinezS	0.858	7.7
6	I-880	45_Hayward	46_SanLeandro	0.857	7.6
7	Benicia Bridge	73_Benicia	60_Martinez	0.853	7.1
8	I-680	63_ConcordS	64_WalnutCreek	0.851	6.8
9	I-80	53_BayBridgeE	55_Albany	0.847	6.3
10	Richmond Bridge	07_RichmondBridgeW	56_RichmondBridgeE	0.844	5.9
17	SR 37	03_SearsPoint	75_VallejoW	0.831	4.3
18	San Mateo Bridge	16_SanMateoBridgeW	44_SanMateoBridgeE	0.828	3.9
24	US 101	08_MillValley	09_GGBridgeN	0.825	3.5

Notes. An attacked edge is lost from service for only a single commute period, and thus delay is measured only for one such period. “Delay increase” is measured from the nominal delay of 0.797×10^6 vht for a single commute period. The last three entries in the table enable computation of “synergy” in Table 4.

only third, with a 17% increase; and the Caldecott Tunnel (edge [54, 65]), which feeds the Bay Bridge, ranks fourth, with a 15% increase. The single-edge losses ranked fifth through 10th (and beyond) are significantly less disruptive.

Table 4 presents the 10 most-damaging two-edge losses for a low-capability attacker. The most disruptive of these combines the top two single-edge losses (the Bay Bridge and a segment of I-680); it results in a system delay of 60.8% or 1.281×10^6 vht. All of the top 10 most disruptive two-edge losses involve one of the two most disruptive single-edge losses. We observe, however, that for each of these pairs, the loss of two edges in combination is more disruptive (by 3%–32%) than would be predicted by the sum of the single-edge losses based solely on the data in Table 3. For example, the San Mateo Bridge (edge [16, 44]) ranks as the 18th most disruptive single-edge loss, but paired with the Bay Bridge it ranks as part of the fifth most disruptive two-edge loss. This can be understood intuitively because the Bay Bridge and the San Mateo Bridge are components of substitute paths that connect the east and west sides of the bay.

We end this section by solving T-DAD(\hat{w}_0, x, y) to identify the most damaging attacks for $\bar{a} = 1, \dots, 5$, for a low-capability attacker (see Table 5). We observe that these worst-case attacks represent monotonically increasing sets.

4.3.2. Multiperiod Losses from a High-Capability Attacker. We repeat the previous analyses for the case of a high-capability attacker who can damage components for multiple periods, as specified in Table 2.

Table 6 presents the 10 most-damaging single-edge losses for a high-capability attacker. The effect of repair

times now becomes apparent: bridge-like edges that are not bridges (e.g., edge [40, 43] on I-680) drop in importance compared to preceding results, because they are repaired quickly, and true bridges rise in importance because they will take years to repair.

Table 7 presents the 10 most-damaging two-edge losses for a high-capability attacker. Here, each of the top 10 entries consist solely of components with longer repair times.

Table 8 lists the worst possible attacks by a high-capability attacker on the undefended system, with the attacker’s resource ranging from zero to five edges.

We observe that edges that are road segments such as I-680, although bridge-like in terms of their penalty cost in the short term, do not show up as a worst-case loss when considering actual bridges having longer repair times. The five worst single-edge losses are all bridges. The length of the repair time alone is not enough to determine the worst-case attacks; however, losing the Bay Bridge for three years is *more* costly in travel delays than losing the Golden Gate Bridge for five years (with the longest repair time).

4.4. Improving Operational Resilience

Our ultimate goal is to determine how to reconfigure or defend edges in a road network to improve that network’s operational resilience. For simplicity, we consider only the case of a highly capable adversary.

Table 9 presents the best defenses and the resulting worst-case losses for T-DAD with the heterogeneous repair times in Table 2. We observe that defending the Bay Bridge is the best way to mitigate the worst-case loss of one or two other edges. However, the best defense against the loss of three to five edges protects the Golden Gate Bridge. In our experience, this kind

Table 4. The 10 Most Disruptive Two-Edge Losses, Assuming a Low-Capability Attacker

Rank	Edges lost			Sys. delay (vht × 10 ⁻⁶)	Delay increase (%)	Synergy (%)
	Road, bridge, or tunnel	Junction 1	Junction 2			
1	Bay Bridge I-680	11_BayBridgeW 40_FremontSE	53_BayBridgeE 43_Sunol	1.281	60.8	16.3
2	Golden Gate Br. Bay Bridge	09_GGBridgeN 11_BayBridgeW	10_GGBridgeS 53_BayBridgeE	1.254	57.4	25.5
3	Bay Bridge I-880	11_BayBridgeW 45_Hayward	53_BayBridgeE 46_SanLeandro	1.183	48.4	34.0
4	Richmond Bridge Bay Bridge	07_RichmondBridgeW 11_BayBridgeW	56_RichmondBridgeE 53_BayBridgeE	1.152	44.5	29.1
5	Bay Bridge San Mateo Bridge	11_BayBridgeW 16_SanMateoBridgeW	53_BayBridgeE 44_SanMateoBridgeE	1.137	42.7	31.3
6	Golden Gate Br. I-680	09_GGBridgeN 40_FremontSE	10_GGBridgeS 43_Sunol	1.134	42.3	4.0
7	I-680 Caldecott Tunn.	40_FremontSE 65_CaldecottTunnE	43_Sunol 54_CaldecottTunnW	1.126	41.3	7.2
8	I-680 I-880	40_FremontSE 45_Hayward	43_Sunol 46_SanLeandro	1.110	39.3	26.2
9	US 101 Bay Bridge	08_MillValley 11_BayBridgeW	09_GGBridgeN 53_BayBridgeE	1.083	35.9	11.7
10	SR 37 Bay Bridge	03_SearsPoint 11_BayBridgeW	75_VallejoW 53_BayBridgeE	1.081	35.6	8.4

Notes. See the caption of Table 3 for a description of the results except for “Synergy.” Synergy represents as a percentage the incremental cost that the two-edge loss incurs above the sum of the costs of the constituent one-edge losses. (For example, the synergy of function $g(a, b)$ over $f(a) + f(b)$, given baseline $g(0, 0) = f(0)$, is $100\% \times (\Delta g(a, b) - (\Delta f(a) + \Delta f(b))) / (\Delta f(a) + \Delta f(b)) - 100\%$, where $\Delta g(a, b) = g(a, b) - g(0, 0)$, $\Delta f(a) = f(a) - f(0)$, and $\Delta f(b) = f(b) - f(0)$.)

Table 5. Worst-Case Single-Period Losses for the Undefended System, as Discovered by a Low-Capability Attacker

Attacker resource	Edges lost			Sys. delay (vht × 10 ⁻⁶)	Delay increase (%)
	Road, bridge, or tunnel	Junction 1	Junction 2		
0	—	—	—	0.797	—
1	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.025	28.6
2	Bay Bridge I-680	11_BayBridgeW 43_Sunol	53_BayBridgeE 40_FremontSE	1.281	60.7
3	Bay Bridge I-680	11_BayBridgeW 43_Sunol	53_BayBridgeE 40_FremontSE	1.521	90.8
4	Golden Gate Br. Bay Bridge I-680	09_GGBridgeN 11_BayBridgeW 43_Sunol	10_GGBridgeS 53_BayBridgeE 40_FremontSE	1.735	117.7
5	Golden Gate Br. I-880 Bay Bridge I-680 Golden Gate Br. I-880 Richmond Bridge	09_GGBridgeN 46_SanLeandro 11_BayBridgeW 43_Sunol 09_GGBridgeN 46_SanLeandro 07_RichmondBridgeW	10_GGBridgeS 45_Hayward 53_BayBridgeE 40_FremontSE 10_GGBridgeS 45_Hayward 56_RichmondBridgeE	1.813	127.5

Notes. The table reports the baseline system operating cost and the worst losses for one to five edges lost. For each attack, the table lists the operating cost and the percentage increase over the baseline cost. All solutions are optimal.

of nonmonotonicity often appears in the solutions to AD and DAD problems. This means that prioritized (ranked) target lists, for either defense or attack, may provide poor guidance to decision makers, even if created using a complicated scoring scheme (e.g., Leung, Lambert, and Mosenthal 2004, Apostolakis and Lemon 2005, Tranchita, Hadsaid, and Torres 2006); see Alderson et al. (2013) for a discussion.

Figure 3 shows the impact of optimal defenses of up to three edges on the worst-case delay that an adversary can inflict on the system, as a function of an increasing number of losses.

The case of no defenses in Figure 3 provides a baseline for measuring the operational resilience of the highway infrastructure to various levels of attack: with no defense in place, an adversary could inflict consid-

Table 6. The 10 Most Disruptive Single-Edge Losses, Assuming a High-Capability Attacker

Rank	Road, bridge, or tunnel	Edge lost		Sys. delay (vht × 10 ⁻⁹)	Delay increase (%)
		Junction 1	Junction 2		
1	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.232	17.1
2	Golden Gate Bridge	09_GGBridgeN	10_GGBridgeS	1.231	17.0
3	Carquinez Bridge	59_CarquinezS	72_CarquinezN	1.093	3.9
4	Benicia Bridge	73_Benicia	60_Martinez	1.067	1.4
5	Richmond Bridge	07_RichmondBridgeW	56_RichmondBridgeE	1.064	1.1
6	San Mateo Bridge	16_SanMateoBridgeW	44_SanMateoBridgeE	1.060	0.8
7	Dumbarton Bridge	20_DumbartonBridgeW	41_DumbartonBridgeE	1.058	0.6
8	I-680	40_FremontSE	43_Sunol	1.056	0.4
9	Caldecott Tunnel	65_CaldecottTunnE	54_CaldecottTunnW	1.054	0.2
10	Rio Vista Bridge	79_RioVistaBrN	82_RioVistaBrS	1.053	0.1

Notes. Table 2 specifies repair times, which can cover as many as 1,320 commute periods (five years). Consequently, “Delay increase” is measured with respect to the nominal delay of 1.052×10^9 vht over 1,320 commute periods.

Table 7. The 10 Most Disruptive Two-Edge Losses, Assuming a High-Capability Attacker

Rank	Road, bridge, or tunnel	Edges lost		Sys. delay (vht × 10 ⁻⁹)	Delay increase (%)	Synergy (%)
		Junction 1	Junction 2			
1	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.486	41.3	20.9
	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS			
2	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS	1.267	20.4	12.5
	Richmond Bridge	07_RichmondBridgeW	56_RichmondBridgeE			
3	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.266	20.3	11.5
	Richmond Bridge	07_RichmondBridgeW	56_RichmondBridgeE			
4	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS	1.264	20.2	-3.6
	Carquinez Bridge	59_CarquinezS	72_CarquinezN			
5	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.261	19.9	11.1
	San Mateo Bridge	16_SanMateoBridgeW	44_SanMateoBridgeE			
6	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.253	18.7	-9.0
	Carquinez Bridge	59_CarquinezS	72_CarquinezN			
7	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS	1.249	18.4	1.5
	Benicia Bridge	73_Benicia	60_Martinez			
8	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.246	18.2	4.3
	Dumbarton Bridge	20_DumbartonBridgeW	41_DumbartonBridgeE			
9	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS	1.243	18.1	2.1
	San Mateo Bridge	16_SanMateoBridgeW	44_SanMateoBridgeE			
10	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.242	35.7	-2.6
	Benicia Bridge	73_Benicia	60_Martinez			

Notes. See the captions in Tables 3 and 4 for further explanation. Here, positive synergy means that the total increase in delay is greater than the sum of the increase from the delay resulting from each individual loss; negative synergy means the total is less than the sum. Two components in parallel tend to exhibit positive synergy, while components in series tend to exhibit negative synergy.

erable additional delay on the system and, as expected, this cost increases with an increasing number of lost edges.

However, when the defender can implement the optimal single-edge defense for each separate number of attacks, the resulting worst-case attacks produce significantly less disruption. With the optimal two-edge defense, the disruptions caused by worst-case attacks are less, and with the optimal three-edge defense, these are mitigated even more. This illustrates the improvement in resilience that comes with increasing optimal defensive investments.

Collectively, these *resilience curves* characterize the returns on investment a decision maker could achieve. For example, we see that the first bridge defended


reduces disruption the most for the worst-case attacks, and therefore offers the most significant improvement in operational resilience. Similar curves should apply to any infrastructure system for which an industry-standard operational model exists; see Alderson et al. (2013) and Alderson, Brown, and Carlyle (2015) for further discussion.

Although Figure 3 quantifies the benefit for each additional defended edge, it does not specify how many bridges should be defended. That decision belongs to the system defender or a policy maker, based on an assessment of the actual costs of implementing each defense plan and the decision-maker’s willingness to trade those costs for the benefits they provide.

Table 8. Worst-Case Attacks by a High-Capability Attacker Against the Undefended System

Attacker resource	Edges lost			Sys. delay (vht × 10 ⁻⁹)	Delay increase (%)
	Road, bridge, or tunnel	Junction 1	Junction 2		
0	—	—	—	1.052	—
1	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.232	17.1
2	Bay Bridge	11_BayBridgeW	53_BayBridgeE	1.486	41.3
3	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS	1.515	44.0
	Bay Bridge	11_BayBridgeW	53_BayBridgeE		
4	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS	1.549	47.2
	San Mateo Bridge	16_SanMateoBridgeW	44_SanMateoBridgeE		
	Bay Bridge	11_BayBridgeW	53_BayBridgeE		
	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS		
5	San Mateo Bridge	16_SanMateoBridgeW	44_SanMateoBridgeE	1.590	51.1
	Dumbarton Bridge	20_DumbartonBridgeW	41_DumbartonBridgeE		
	Bay Bridge	11_BayBridgeW	53_BayBridgeE		
	Golden Gate Br.	09_GGBridgeN	10_GGBridgeS		
	Carquinez Bridge	59_CarquinezS	72_CarquinezN		
	Richmond Bridge	07_RichmondBridgeW	56_RichmondBridgeE		
	Benicia Bridge	73_Benicia	60_Martinez		

Notes. See the caption of Table 5 for further explanation.

Figure 3. Resilience Curves That Characterize How Different Levels of Defense Mitigate Against Varying Attack Resource Given a Highly Capable Attacker

4.5. Multiple Defense Types

The tests above assume that a defended bridge becomes invulnerable, and could have been modeled with simpler binary variables, i.e., $w_{ij} = 1$ if edge $[i, j]$ is defended and $w_{ij} = 0$ otherwise. This section illustrates how multiple defense options might actually be used, focusing on bridge-like edges. In particular, assuming a high-capability attacker, for any bridge-like edge

$$w_{ijd_0} = 1 \text{ if } [i, j] \text{ is undefended,}$$

$$w_{ijd_1} = 1 \text{ if the latest vulnerable repair epoch for } [i, j] \text{ is made invulnerable,}$$

$$w_{ijd_2} = 1 \text{ if the latest and second-to-latest vulnerable repair epochs for } [i, j] \text{ are made invulnerable, etc.}$$

Of course, the requirement that $\sum_{d \in D_{ij}} w_{ijd} = 1$ still applies. These definitions would enable the modeling of progressive structural improvements to bridges, tunnels, and certain roadways (see Alderson, Brown, and Carlyle 2014, for additional discussion).

For the sake of brevity, we consider only a single case: In effect, one unit of defense resource can protect one repair epoch for a given edge, and we assume a total of 10 units of resource. Table 10 presents the results. In the case with a defense budget of 10 units, we observe that the optimal solution is to reduce the repair times of all bridges so that they are about the same. The reasons for this invite additional analysis.

4.6. Discussion

Our modeling of regional highway systems shares many features with the recent literature on transportation resilience. Nearly all of the papers in the recent literature use the BPR congestion model we use, and all but one of those reviewed are based on a capacitated multicommodity (OD pair) traffic model, minimizing some cost or expected cost. Most of them allow for rerouting of traffic following a disruptive event and some consider explicit recovery activities as part of this response. Some of these papers deal with successive decisions to first defend, fortify, make more redundant, add more infrastructure, etc., prior to the event itself.

However, the technique presented here differs from the literature in several important ways. First, nearly all of the papers in the literature deal with two-stage stochastic optimizations of preparations for a random (i.e., probabilistic) disaster followed by the operation of the randomly damaged traffic system. The use of probability distributions to represent damage from natural disasters is well established, however there is considerable concern about the use of probabilities to model the behavior of an intelligent, goal-oriented terrorist (National Research Council 2008, 2010); see also Alderson, Brown, and Carlyle (2015) for a detailed discussion.

Table 9. Best Defenses and Worst-Case Losses Against a High-Capability Attacker

Num. def.	Num. lost	Defended edges	Lost edges	Sys. delay (vht × 10 ⁻⁹)	Delay increase (%)
1	1	Bay Bridge [11, 53]	Golden Gate Bridge [9, 10]	1.231	17.1
1	2	Golden Gate Bridge [9, 10]	Bay Bridge [11, 53] Richmond Bridge [7, 56]	1.266	20.3
1	3	Golden Gate Bridge [9, 10]	Bay Bridge [11, 53] Carquinez Bridge [59, 72]	1.304	24.0
1	4	Golden Gate Bridge [9, 10]	Bay Bridge [11, 53] Richmond Bridge [7, 56] Carquinez Bridge [72, 59] Benicia Bridge [73, 60]	1.345	27.9
1	5	Golden Gate Bridge [9, 10]	Bay Bridge [11, 53] Richmond Bridge [7, 56] Carquinez Bridge [72, 59] Benicia Bridge [73, 60] San Mateo Bridge [16, 44]	1.376	30.8
2	1	Bay Bridge [11, 53] Golden Gate Bridge [9, 10]	Carquinez Bridge [72, 59]	1.093	3.9
2	2	Bay Bridge [11, 53] Golden Gate Bridge [9, 10]	Carquinez Bridge [72, 59] Benicia Bridge [73, 60]	1.145	8.8
2	3	Bay Bridge [11, 53] Golden Gate Bridge [9, 10]	Carquinez Bridge [59, 72] Benicia Bridge [73, 60] Richmond Bridge [7, 56]	1.172	11.4
2	4	Bay Bridge [11, 53] Golden Gate Bridge [9, 10]	Carquinez Bridge [72, 59] Benicia Bridge [73, 60] Richmond Bridge [7, 56] Nejedly Bridge [70, 83]	1.180	12.2
2	5	Bay Bridge [11, 53] Golden Gate Bridge [9, 10]	Carquinez Bridge [72, 59] Benicia Bridge [73, 60] Richmond Bridge [7, 56] San Mateo Bridge [16, 44] Dumbarton Bridge [20, 41]	1.198	13.9
3	1	Bay Bridge [11, 53] Golden Gate Bridge [9, 10] Carquinez Bridge [72, 59]	Benicia Bridge [73, 60]	1.067	1.4
3	2	Bay Bridge [11, 53] Golden Gate Bridge [9, 10] Carquinez Bridge [72, 59]	Carquinez Bridge [72, 59] Benicia Bridge [73, 60]	1.082	2.9
3	3	Bay Bridge [11, 53] Golden Gate Bridge [9, 10] Carquinez Bridge [72, 59]	Richmond Bridge [7, 56] San Mateo Bridge [16, 44] Dumbarton Bridge [20, 41]	1.104	4.9
3	4	Bay Bridge [11, 53] Golden Gate Bridge [9, 10] Carquinez Bridge [72, 59]	Richmond Bridge [7, 56] San Mateo Bridge [16, 44] Dumbarton Bridge [20, 41] Benicia Bridge [73, 60]	1.114	5.9
3	5	Bay Bridge [11, 53] Golden Gate Bridge [9, 10] Carquinez Bridge [72, 59]	Richmond Bridge [7, 56] San Mateo Bridge [16, 44] Dumbarton Bridge [20, 41] Benicia Bridge [73, 60] I-680 (Fremont–Sunol) [40, 43]	1.124	6.8

Notes. For each combination of number of defended edges (Num. defs.) and number of losses (Num. lost), the table reports the best edges to defend and worst loss(es) by solution rank, along with the resulting total system operating cost, and percentage increase over the baseline objective value of 1.052×10^9 . Percentage increases are computed assuming that solutions are optimal, although a relative optimality gap of up to 1% is allowed.

We have evaluated stochastic optimization, and Monte Carlo simulation of scenarios versus those that an intelligent adversary can plan. The results have shown uniformly that the worst case is *much worse* than typically revealed by Monte Carlo simulation (see Alderson et al. 2013 and Alderson, Brown, and Carlyle 2014 for a detailed discussion). Because our goal is understanding and mitigating against worst-case disruptions no matter how they result, we believe it is

prudent to begin with an assessment from our hypothetical intelligent adversary, provided it is practical to do so.

None of the existing work lets an attacker view all defensive preparations before committing to the most-damaging attack he can afford, and perhaps for practical reasons. This is a trilevel optimization, with decisions at each level, and as noted by some of the recent work, there is an exponential number of admissible

Table 10. Best Defenses Given 10 Units of Variable Defense Against an Attacker Who Can Strike Any Five Bridge-Like Edges

	Repair epochs				
	t_1	t_2	t_3	t_4	
Defended edges					
Golden Gate Bridge [9, 10]	D	D	D	D	D = Defended
Bay Bridge [11, 53]	V	D	D	—	— = Never vulnerable
Carquinez Bridge [59, 72]	V	V	D	—	V = Still vulnerable
Richmond Bridge [7, 56]	V	D	—	—	
San Mateo Bridge [16, 44]	V	D	—	—	
Benicia Bridge [60, 73]	V	D	—	—	
Attacked edges					
Bay Bridge [11, 53]	0	1	1	1	1 = Available
San Mateo Bridge [16, 44]	0	1	1	1	0 = Unavailable
I-680 (FremontSE–Sunol) [40, 43]	0	1	1	1	
Dumbarton Bridge [20, 41]	0	0	1	1	
Carquinez Bridge [59, 72]	0	0	1	1	

Note. The total cost of these attacks is 1.085×10^9 vht compared to a baseline of 1.052×10^9 vht.

attacks. So, a distinguishing difference we bring here is the trilevel successive decisions with full visibility of the sequential actions of opponents, and we can solve it for a realistic case study using generalized Benders decomposition. We demonstrate with this example that it is practical to solve problems at a regional scale, but considerable work remains to solve problems at a higher fidelity and/or a national scale.

5. Conclusion

This paper studies the regional highway transportation system of the San Francisco Bay Area. Focusing on peak periods for commuter traffic, traffic patterns are computed from an enhanced traffic model that includes nonlinear congestion behavior and reflects actual travel demands as captured by U.S. Census demographic data. We have instrumented our model so that one can assess the consequence (measured as increased systemwide travel time) associated with the loss of a combination of one or more major road segments, bridges, and/or tunnels, each of which might require a different length of time to restore. Using this model, we systematically assess the operational resilience of the highway system to worst-case losses of segments, and we evaluate the extent to which damage to a small number of them can impact system performance.

We observe three features that are key to understanding why a major road segment, bridge, or tunnel is critical to the overall performance of this system. The first is the location of the segment within the overall network topology and in relation to the origins and destinations of commuter traffic. Bridges and tunnels, which tend to connect geographically isolated areas, often provide critical connectivity, and their loss can

dramatically increase travel times. This is the case for the Golden Gate Bridge and the Bay Bridge. However, we also observe that the loss of some seemingly less significant road segments (e.g., I-880, I-680), can significantly increase travel time for the system.

The second feature that makes loss of a major road segment, bridge, or tunnel critical to system performance is restoration time. Because bridges and tunnels tend to have longer repair times, their loss impacts performance for more commute periods. Indeed, the last numerical case shows an effective strategy for allocating defenses is to invest to decrease all restoration epochs so they are about the same in duration.

The final key insight is that, in general, the contribution to system function by a single component depends on the status of other components. That is, it is the combination of highway segments that when working (or lost) together has the greatest impact on system function. This type of dependence can only be uncovered when studying the system as a whole, and for assessing operational resilience this means that one needs to be concerned with the loss of sets of components. A primary question is, *Which sets of components are most worrisome?*

In this paper, we have concentrated on sets of components whose loss results in a worst-case consequence. Our experience in studying many types of infrastructure systems is that the loss of a relatively small number of components at specific locations can dramatically affect performance. One way of uncovering such sets is to take the perspective of an intelligent adversary (our attacker), even if in practice the “attack” itself comes in the form of an accident, failure, or just plain bad luck. The point is to uncover these critical sets in advance of a disruptive event.

Annually, the United States spends billions of dollars to build, maintain, and strengthen its infrastructure

systems. Given these huge expenditures, any return-on-investment analysis should quantify rigorously the resilience of such systems to natural disasters, attacks, and other disruptive events. For a well-understood system like the regional highway network studied in this paper, we have presented methods that can assess a system's resilience to component loss and even produce a maximally resilient, resource-constrained design. Future research will meld worst-case and probabilistic analysis to balance resilience to attacks with resilience to random events like natural disasters.

Acknowledgments

The basic ideas underlying this story were first presented at the Military Operations Research Society Symposium in June 2009. The authors would like to thank three anonymous reviewers for their comments, which helped to improve the manuscript.

Appendix A. Evaluating Travel Times

This appendix describes the function used to measure the total travel time associated with each arc in the road network, and the piecewise-linear approximation used for computational purposes.

A.1. A Travel-Time Function for Each Arc

We begin with a model of traffic congestion originally developed by the Bureau of Public Roads (1964, Traffic Assignment Manual) and fitted to traffic data from the Highway Research Board (1965, Highway Capacity Manual). The model, known as the BPR function, calculates the average travel speed on each highway arc (i, j) based on traffic-flow rate and empirical characteristics of the arc.

Definitions for parameters and variables:

- y''_{ij} traffic flow rate on (i, j) , i.e., the number of vehicles per hour (vph) that arrive at i and seeking to traverse (i, j) ;
- s_{ij} unobstructed travel speed on (i, j) (mph);
- l_{ij} length of (i, j) (miles);
- $t_{ij} = l_{ij}/s_{ij}$; unobstructed travel time on (i, j) (mph);
- m_{ij} number of lanes on (i, j) ;
- u'_{ij} capacity per lane on (i, j) (vph);
- $u_{ij} = m_{ij}u'_{ij}$; baseline capacity of (i, j) (vph);
- α, β, γ empirical parameters chosen to fit observed traffic flows; we use $\alpha = 0.20$, $\beta = 6$, and $\gamma = 4/3$ (California Metropolitan Transportation Commission 2012).

BPR function: The function $S_{ij}(\cdot)$ represents the traffic-flow speed on arc (i, j) , defined by

$$S_{ij}(y''_{ij}) = s_{ij} \left(1 + \alpha \left(\gamma \frac{y''_{ij}}{u_{ij}} \right)^\beta \right)^{-1}. \quad (\text{A.1})$$

Table B.1. Node Data, Including Node Name, and Location in Latitude (Decimal Degrees North of the Equator) and Longitude (Decimal Degrees West of the Prime Meridian)

Node name	Lat (°N)	Lon (°E)	Node name	Lat (°N)	Lon (°E)
01_Petaluma	38.2343	-122.6184	47_CastroValley	37.6901	-122.0978
02_Sonoma	38.2362	-122.4615	48_Dublin	37.7010	-121.9227
03_SearsPoint	38.1506	-122.4495	49_Livermore	37.7012	-121.8022
04_Napa	38.2232	-122.2578	50_OaklandE	37.7834	-122.1774

Then, given the relationships between time, speed, and arc length, the inverse of the BPR function estimates travel time for every vehicle traversing arc (i, j) , as exploited from Section 3 forward.

Travel-time function: The function $F_{ij}(\cdot)$ represents the travel time arc (i, j) , defined as

$$F_{ij}(y''_{ij}) = l_{ij}/S_{ij}(y''_{ij}) = t_{ij} \left(1 + \alpha \left(\gamma \frac{y''_{ij}}{u_{ij}} \right)^\beta \right). \quad (\text{A.2})$$

This function is relatively flat until $y''_{ij} \approx u_{ij}$, and then rises steeply.

A.2. A Piecewise-Linear Approximation for the Travel-Time Function

Let $F_{ij}(y''_{ij})$ denote a generic form of the various travel-time functions used to define T-DAD (i.e., $F_{ijdt}^1(y''_{ijdt})$ or $F_{ijdt}^1(y''_{ijdt})$), and define $\bar{f}_{ij}(y''_{ij}) = y''_{ij}F_{ij}(y''_{ij})$. In effect, the nonlinear function $\bar{f}_{ij}(y''_{ij})$ appears many times in T-DAD's objective function, but that is the only nonlinear function of concern.

To use linear programming and mixed-integer linear programming in a solution algorithm for T-DAD, $\bar{f}_{ij}(y''_{ij})$ always replaces $y''_{ij}F_{ij}(y''_{ij})$, in a primal or dualized form. In a standard fashion, we define the approximation through the solution of a linear program

$$\bar{f}_{ij}(y''_{ij}) \approx \min_g \sum_{r=1}^{\bar{r}} s_{ijr} g_{ijr} \quad (\text{A.3})$$

$$\text{s.t.} \quad \sum_{r=1}^{\bar{r}} g_{ijr} = y''_{ij}, \quad (\text{A.4})$$

$$0 \leq g_{ijr} \leq \lambda_{ijr}, \quad (\text{A.5})$$

where the variables g_{ijr} represent the traffic flow on arc (i, j) on each subsequent piece r in the piecewise linear representation of f_{ij} , $s_{ijr} \equiv \lambda_{ij}^{-1} [\bar{f}_{ij}(r\lambda_{ij}) - \bar{f}_{ij}((r-1)\lambda_{ij})]$ for all (i, j) , and r ; $\lambda_{ij} \equiv 2u_{ij}/\bar{r}$ for all (i, j) ; and testing has shown that $\bar{r} = 40$ yields an adequate approximation given other definitions. We note that these definitions imply that $y''_{ij} \leq 2u_{ij}$, but travel time on any arc (i, j) with $y''_{ij} = 2u_{ij}$ would exceed the total allowed travel time for a commute period, except in a few minor instances.

We note that this linearization would work for other nonlinear travel time functions, should there be a need to replace the BPR function with something else.

Appendix B. Case-Study Data

This appendix presents excerpts of the data tables used in our models. Table B.1 provides a complete list of nodes, and Table B.2 presents representative arc data.

Table B.1. (Continued)

Node name	Lat (°N)	Lon (°E)	Node name	Lat (°N)	Lon (°E)
05_Novato	38.0799	-122.5458	51_Oakland	37.8002	-122.2792
06_SanRafael	37.9624	-122.5103	52_OaklandN	37.8243	-122.2683
07_RichmondBridgeW	37.9427	-122.4787	53_BayBridgeE	37.8273	-122.2954
08_MillValley	37.9027	-122.5157	54_CaldecottTunnW	37.8487	-122.2280
09_GGBridgeN	37.8325	-122.4810	55_Albany	37.8877	-122.3091
10_GGBridgeS	37.8035	-122.4704	56_RichmondBridgeE	37.9328	-122.4095
11_BayBridgeW	37.7856	-122.3916	57_Richmond	37.9560	-122.3299
12_SanFranciscoS	37.7352	-122.4068	58_Hercules	38.0150	-122.2694
13_DalyCity	37.7050	-122.4715	59_CarquinezS	38.0526	-122.2276
14_SanBruno	37.6276	-122.4313	60_Martinez	38.0325	-122.1170
15_SFO	37.6346	-122.4032	61_ConcordNW	37.9946	-122.0698
16_SanMateoBridgeW	37.5647	-122.2722	62_ConcordNE	38.0051	-122.0367
17_SanMateo	37.5532	-122.2960	63_ConcordS	37.9580	-122.0525
18_CrystalSprings	37.5062	-122.3375	64_WalnutCreek	37.9563	-122.0716
19_RedwoodCity	37.4889	-122.2130	65_CaldecottTunnE	37.8634	-122.2092
20_DumbartonBridgeW	37.4864	-122.1426	66_SanRamon	37.7600	-121.9654
21_MenloPark	37.4834	-122.1807	67_Baypoint	38.0190	-121.9419
22_Woodside	37.4353	-122.2433	68_Pittsburgh	38.0170	-121.8900
23_PaloAlto	37.4604	-122.1409	69_AntiochOakley	37.9905	-121.7600
24_MountainViewN	37.4087	-122.0705	70_NejedlyBrS	38.0190	-121.7510
25_Sunnyvale	37.4006	-122.0357	71_Brentwood	37.8960	-121.7130
26_MountainViewS	37.3842	-122.0683	72_CarquinezN	38.0712	-122.2280
27_MilpitasE	37.4347	-121.8885	73_Benicia	38.0506	-122.1298
28_MilpitasW	37.4257	-121.9165	74_VallejoS	38.0915	-122.2305
29_Cupertino	37.3324	-122.0557	75_VallejoW	38.1400	-122.2560
30_SJC	37.3737	-121.9275	76_VallejoN	38.1393	-122.2200
31_SanJoseN	37.3642	-121.9019	77_Fairfield	38.2166	-122.1379
32_SanJoseE	37.3397	-121.8519	78_Vacaville	38.3710	-121.9560
33_SanJoseCentral	37.3238	-121.8922	79_RioVistaBrN	38.1610	-121.6880
34_SanJoseW	37.3173	-121.9403	80_Davis	38.5200	-121.7700
35_LosGatos	37.2557	-121.9559	81_Sacramento	38.5770	-121.5260
36_SanJoseS	37.2555	-121.8589	82_RioVistaBrS	38.1550	-121.6760
37_CoyoteValley	37.2402	-121.7664	83_NejedlyBrN	38.0300	-121.7510
38_Gilroy	37.0030	-121.5565	84_Lodi	38.1160	-121.3970
39_FremontSW	37.4836	-121.9369	85_Stockton	37.9370	-121.2980
40_FremontSE	37.4955	-121.9232	86_Manteca	37.7670	-121.3320
41_DumbartonBridgeE	37.5372	-122.0716	87_Tracy	37.7419	-121.5735
42_Fremont	37.5635	-122.0385	88_Vernalis	37.6000	-121.3420
43_Sunol	37.5887	-121.8709	89_Patterson	37.4637	-121.1804
44_SanMateoBridgeE	37.6282	-122.1210	90_Gustine	37.2462	-121.0890
45_Hayward	37.6448	-122.0939	91_LosBanos	37.0567	-120.9698
46_SanLeandro	37.6890	-122.1356			

Table B.2. Roadway Types, Associated Data, and Example for the Highway Network Arcs

Roadway type	FFS s_{ij} (mph)	Capacity u'_{ij} (veh./hr/lane)	Example		
			Arc	Road or bridge	Lanes
1_Freeway65	65	1,950	(05_Novato, 06_SanRafael)	US 101	4
2_GGBridge50	50	1,780	(09_GGBridgeN, 10_GGBridgeS)	Golden Gate Br.	4
3_Freeway55	55	1,850	(25_Sunnyvale, 26_Milpitas)	US 237	3
4_Rural55	55	1,530	(03_SearsPoint, 75_Novato)	SR 37	1
5_SigCoor30	30	950	(72_CarquinezNorth, 07_VallejoWest)	SR 29	2
6_LocalStreet25	25	900	(22_Woodside, 19_RedwoodCity)	SR 84	3
7_Collector45	45	1,500	(20_DumbartonBridgeW, 21_MenloPark)	Bayfront Expwy.	3

Notes. For each type, the table shows free-flow speed (FFS) in miles per hour (mph) and capacity in vehicles per hour per lane. (See Appendix A.1 for definitions.) For each roadway type, an example is given from the network specified in Figure 1, including the number of lanes modeled. Note that an arc of type “3_Freeway55” also represents the Caldecott Tunnel in computational tests.

References

- Alderson DL, Brown GG, Carlyle WM (2014) Assessing and improving operational resilience of critical infrastructures and other systems. Newman A, Leung J, eds. *Tutorials in Operations Research: Bridging Data and Decisions* (INFORMS, Catonsville, MD), 180–215.
- Alderson DL, Brown GG, Carlyle WM (2015) Operational models of infrastructure resilience. *Risk Anal.* 35(4):562–586.
- Alderson DL, Brown GG, Carlyle WM, Cox LA (2013) Sometimes there is no “most vital” arc: Assessing and improving the operational resilience of systems. *Military Oper. Res.* 18(1):21–37.
- Alderson DL, Brown GG, Carlyle WM, Wood RK (2011) Solving defender-attacker-defender models for infrastructure defense. Wood K, Dell R, eds. *Operations Research, Computing and Homeland Defense* (INFORMS, Hanover, MD), 28–49.
- Alguacil N, Delgadillo A, Arroyo JM (2014) A trilevel programming approach for electric grid defense planning. *Comput. Oper. Res.* 41:282–290.
- Apostolakis GE, Lemon DM (2005) A screening methodology for the identification and ranking of infrastructure vulnerabilities due to terrorism. *Risk Anal.* 25(2):361–376.
- Beckmann M (1967) On the theory of traffic flows in networks. *Traffic Quart.* 2:109–116.
- Beckmann M, McGuire CB, Winsten CB (1956) *Studies in the Economics of Transportation* (Yale University Press, New Haven, CT).
- Bell MGH, Kanturska U, Schmöcker JD, Fonzone A (2008) Attacker-defender models and road network vulnerability. *Philos. Trans. Royal Soc. A* 388:1893–1906.
- Boyce D, Bar-Gera H (2003) Validation of multiclass urban travel forecasting models combining origin-destination, mode, and route choices. *J. Regional Sci.* 43:517–540.
- Boyce DE, Mahmassani HS, Nagurney A (2005) A retrospective on Beckmann, McGuire and Winsten’s studies in the economics of transportation. *Papers Regional Sci.* 84(1):85–103.
- Brown GG, Cox LA (2011a) How probabilistic risk assessment can mislead terrorism risk analysis. *Risk Anal.* 31(2):196–204.
- Brown GG, Cox LA (2011b) Making terrorism risk analysis less harmful and more useful: Another try. *Risk Anal.* 31(2):193–195.
- Brown GG, Carlyle WM, Salmerón J, Wood K (2006) Defending critical infrastructure. *Interfaces* 36(6):530–544.
- Bureau of Public Roads (1964) Traffic assignment manual. U.S. Department of Commerce, Washington, DC.
- California Metropolitan Transportation Commission (2012) Initial examination of volume delay functions using PeMS data. http://mtcgs.mtc.ca.gov/foswiki/pub/Main/Documents/2012_03_06_RELEASE_Volume_delay_functions.pdf.
- Cappanera P, Scaparra MP (2011) Optimal allocation of protective resources in shortest-path networks. *Transportation Sci.* 45(1):64–80.
- Chen L, Miller-Hooks E (2012) Resilience: An indicator of recovery capability in intermodal freight transport. *Transportation Sci.* 46(1):109–123.
- Church RL, Scaparra MP (2006) Protecting critical assets: The r -interdiction median problem with fortification. *Geographical Anal.* 39(2):129–146.
- Cimellaro GP, Reinhorn AM, Bruneau M (2010) Framework for analytical quantification of disaster resilience. *Engrg. Structures* 32(11):3639–3649.
- Cormican KJ, Morton DP, Wood RK (1999) Stochastic network interdiction. *Oper. Res.* 46(2):184–197.
- Correa JR, Stier-Moses NE (2010) Wardrop equilibria. Cochran JJ, ed. *Wiley Encyclopedia of Operations Research and Management Science* (John Wiley, Hoboken, NJ), 1–12.
- Dafermos SC, Sparrow FT (1969) The traffic assignment problem for a general network. *J. Res. Natl. Bureau Standards* 73B:91–118.
- Danskin JW (1967) *The Theory of Max-Min* (Springer-Verlag, New York).
- Dimitrov NB, Morton DP (2013) Interdiction models and applications. Hermmann JW, ed. *Handbook of Operations Research for Homeland Security* (Springer, New York), 73–103.
- Dowling R, Singh GR, Cheng WW (1998) Accuracy and performance of improved speed-flow curves. *Transportation Res. Record* 1646:9–17.
- Fan Y, Liu C (2010) Solving stochastic transportation network protection problems using the progressive hedging-based method. *Networks Spatial Econom.* 10(2):193–208.
- Faturechi R, Miller-Hooks E (2014a) A mathematical framework for quantifying and optimizing protective actions for civil infrastructure systems. *Computer-Aided Civil Infrastructure Engrg.* 29(8):572–589.
- Faturechi R, Miller-Hooks E (2014b) Measuring the performance of transportation infrastructure systems in disasters: A comprehensive review. *J. Infrastructure Systems* 21(1):Article 04014025.
- Faturechi R, Miller-Hooks E (2014c) Travel time resilience of roadway networks under disaster. *Transportation Res. Part B: Methodological* 70:47–64.
- Florian M, Nguyen S (1976) An application and validation of equilibrium trip assignment methods. *Transportation Sci.* 10(4):374–389.
- Fulkerson DR, Harding GC (1977) Maximizing the minimum source-sink path subject to a budget constraint. *Math. Programming* 13:116–118.
- GAMS Development Corporation (2013) GAMS v.23.8.2. <http://www.gams.com>.
- Gazis DC (2002) *Traffic Theory*, Internat. Series Oper. Res. Management Sci., Vol. 50 (Springer Science and Business Media, New York).
- Geoffrion AM (1972) Generalized Benders decomposition. *J. Optim. Theory Appl.* 10(4):237–260.
- Ghare PM, Montgomery DC, Turner WC (1971) Optimal interdiction policy for a flow network. *Naval Res. Logist. Quart.* 18(1):37–45.
- Golden B (1977) A problem in network interdiction. *Naval Res. Logist. Quart.* 25:711–713.
- Google (2016) Google Maps. <http://maps.google.com>.
- Hallenbeck M, Rice M, Smith B, Cornell-Martinez C, Wilkinson J (1997) Vehicle volume distributions by classification. Technical Report FHWA-PL-97-025, Federal Highway Administration, U.S. Department of Transportation, Washington, DC.
- Highway Research Board (1965) Highway capacity manual. Special Report 87, Division of Engineering and Industrial Research, National Research Council, Washington, DC.
- Homeland Security Council (2007) National Strategy for Homeland Security. The White House, Washington, DC.
- IBM (2014) IBM ILOG CPLEX V12.6.0, User’s Manual for CPLEX. <http://pic.dhe.ibm.com/infocenter/cosinfoc/v12r6/index.jsp>.
- Israeli E, Wood RK (2002) Shortest-path network interdiction. *Networks* 40(3):97–111.
- Keeney RL (2007) Modeling values for anti-terrorism analysis. *Risk Anal.* 27(3):585–596.
- Leung M, Lambert JH, Mosenthal A (2004) A risk-based approach to setting priorities in protecting bridges against terrorist attacks. *Risk Anal.* 24(4):963–984.
- Liu C, Fan Y, Ordóñez F (2009) A two-stage stochastic programming model for transportation network protection. *Comput. Oper. Res.* 36(5):1582–1590.
- Luh PB, Chang S-C, Chang T-S (1984) Solutions and properties of multi-stage Stackelberg games. *Automatica* 20:251–256.
- McMasters AW, Mustin TM (1970) Optimal interdiction of a supply network. *Naval Res. Logist. Quart.* 17:261–268.
- Metropolitan Transportation Commission (2013) Bay Area freeway locations with most delay during commute hours, 2013. http://files.mtc.ca.gov/pdf/congestion/BayArea_Top-10_Congestion_Hotspots_2013.pdf.
- Miller-Hooks E, Zhang X, Faturechi R (2012) Measuring and maximizing resilience of freight transportation networks. *Comput. Oper. Res.* 39(7):1633–1643.
- Murray AT, Matisziw TC, Grubescic TH (2007) Critical network infrastructure analysis: Interdiction and system flow. *J. Geographical Systems* 9(2):103–117.

- National Research Council (2008) Department of Homeland Security bioterrorist risk assessment: A call for change. *Committee on Methodological Improvements to the Department of Homeland Security's Biological Agent Risk Analysis* (National Academies Press, Washington, DC).
- National Research Council (2010) Review of the Department of Homeland Security's approach to risk analysis. *Committee to Review the Department of Homeland Security's Approach to Risk Analysis* (National Academies Press, Washington, DC).
- O'Neill RP, Helman U, Hobbs BF, Baldick R (2006) Independent system operators in the United States: History, lessons learned, and prospects. Sioshansi F, Pfaffenberger W, eds. *Electricity Market Reform: An International Perspective* (Elsevier, Oxford, UK), 479–528.
- Peeta S, Salman FS, Gunnec D, Viswanath K (2010) Pre-disaster investment decisions for strengthening a highway network. *Comput. Oper. Res.* 37(10):1708–1719.
- Petersen ER (1975) A primal-dual traffic assignment algorithm. *Management Sci.* 22(1):87–95.
- Pogash C, Grady B (2014) Berkeley protesters block freeway over Garner and Brown killings. *New York Times* (December 9), <http://www.nytimes.com/2014/12/09/us/berkeley-sweeps-up-after-violent-protests-over-garner-and-brown-killings.html>.
- Rockafellar RT, Wets RJ-B (1991) Scenarios and policy aggregation in optimization under uncertainty. *Math. Oper. Res.* 16(1):119–147.
- Salmerón J, Wood K, Baldick R (2009) Worst-case interdiction analysis of large-scale electric power grids. *IEEE Trans. Power Systems* 24(1):96–104.
- Scaparra MP, Church R (2012) Protecting supply systems to mitigate potential disaster a model to fortify capacitated facilities. *Internat. Regional Sci. Rev.* 35(2):188–210.
- Snyder LV, Scaparra MP, Daskin MS, Church RL (2006) Planning for disruptions in supply chain networks. Johnson MP, Norman B, Secomandi N, eds. *Tutorials in Operations Research: Models, Methods, and Applications for Innovative Decision Making* (INFORMS, Hanover, MD), 234–257.
- The White House (2013) Presidential Policy Directive—Critical infrastructure security and resilience. Washington, DC. <https://obamawhitehouse.archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil>.
- Tranchita C, Hadjsaid N, Torres A (2006) Ranking contingency resulting from terrorism by utilization of the Bayesian networks. *Proc. IEEE MELECON, Malaga, Spain*, 964–967.
- U.S. Census Bureau (2009a) Census transportation planning products (CTPP). http://www.fhwa.dot.gov/planning/census_issues/ctpp/.
- U.S. Census Bureau (2009b) Longitudinal employer-household dynamics (LEHD). <http://lehd.ces.census.gov>.
- Van Slyke RM, Wets R (1969) L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM J. Appl. Math.* 17(4):638–663.
- Wardrop JG (1952) Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers, Part II* 1:325–378.
- Waze (2016) Waze. <http://www.waze.com>.
- Weisbrod G, Vary D, Treyz G (2003) Measuring economic costs of urban traffic congestion to business. *Transportation Res. Record: J. Transportation Res. Board* 1839(1):98–106.
- Williamson EB, Winget DG (2005) Risk management and design of critical bridges for terrorist attacks. *J. Bridge Engrg.* 10(1):96–106.
- Wollmer RD (1964) Removing arcs from a network. *Oper. Res.* 12(6):934–940.
- Wollmer RD (1968) Stochastic sensitivity analysis of maximum flow and shortest route networks. *Management Sci.* 14(9):551–564.
- Wood RK (1993) Deterministic network interdiction. *Math. Comput. Modeling* 17(2):1–18.
- Wood RK (2011) Bilevel network interdiction models: Formulations and solutions. Cochran JJ, ed. *Wiley Encyclopedia of Operations Research and Management Science* (John Wiley & Sons, Hoboken, NJ), 1–11.
- Yuan W, Zhao L, Zeng B (2014) Optimal power grid protection through a defender–attacker–defender model. *Reliability Engrg. System Safety* 121:83–89.