Operational Perspectives Into the Resilience of the U.S. Air Transportation Network Against Intelligent Attacks

Karl H. Thompson and Huy T. Tran

***Abstract*— Protecting critical infrastructures has been deemed a natural security imperative by the U.S. Government for the past 20 years. The resilience of transportation infrastructures is of particular importance, as they serve a crucial rule in economic development and citizen mobility throughout the world. This paper presents a defender–attacker–defender model to analyze the potential impacts of intelligent attacks and worst case disruptions on the U.S. air transportation network, as well as possible protection steps that could be taken to minimize the negative outcomes of such disruptions. Furthermore, to analyze the effects of intermodal connections on the resilience of the air network, a second layer representing a hypothetical bus network is added to the model and studied. We use these models, supported by publicly available data, to identify routes likely to be attacked by intelligent adversaries and those critical to the resilient operation of the air network in such scenarios. We also demonstrate the potential benefits of intermodal linkages toward maintaining network operations and identify promising research directions for this type of integrated and intelligent transportation system.**

***Index Terms*— Transportation, networks, optimization, air, resilience, intelligent, attacker, defender, intermodal, multimodal.**

I. Introduction

PROTECTING critical infrastructures is of utmost importance to national security. Lewis [1] traces the development of the US Government's approach to critical infrastructure protection (CIP) over the past few decades. He notes that it has evolved over the years from initial awareness of the problem due to the emerging need to counter terrorism to what it is today: a strategy characterized by risk-informed policy-making guided by practical infrastructure assessments.

Infrastructure resilience in the context of CIP is defined according to Presidential Policy Directive (PPD) 21 [2] as "the ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions." The directive goes further to assert that this definition encompasses disruptions resulting from not only natural incidents, but also human-caused accidents and malicious attacks. Faturechi and Miller-Hooks [3] conducted a review of literature on the assessment of transportation infrastructure performance in disasters, both natural and man-made. The authors determined that while resilience is widely regarded as one of seven key performance characteristics, along with risk, vulnerability, survivability, reliability, flexibility and robustness, it is unique in two aspects: its inherent property of accounting for the strengths and weaknesses measured by other characteristics, and its encapsulation of the benefits of a system's flexibility which enables it to adapt to post-disruption circumstances. These aspects position resilience as one of the more important system characteristics with regard to CIP.

Transportation infrastructure has been studied extensively over the past few decades. In the early nineties, Winston [4] underlined problems facing both ground and air transportation and emphasized the significant potential benefits that efficient infrastructure policy would have. Since then, many authors have continued to study the problem, and renew the call to establish effective policy positions to maintain and protect different transportation systems [5]–[7]. Air transportation in particular, as a subset of the greater transportation infrastructure, is of valuable interest in the context of resilience analysis due to its substantial scale and economic role. Comprising approximately 19,700 airports, heliports, and sea plane bases [8] connected by hundreds of thousands of individual routes, analysis of the US air transportation network plays a major role in understanding the resilience of the overall transportation system. Also of interest are the interdependencies between the air network and other transportation modes, and how they affect its own resilience.

This paper develops a defender-attacker-defender model for exploring and improving the resilience of transportation networks against attacks by an intelligent adversary, through quantification of potential disruption impacts and identification of optimal mitigation actions and critical routes for resilient operation. The developed model is applied to the US air transportation network, based on publicly available data from the Bureau of Transportation Statistics that show average traffic, ticket cost, and distance across routes in the air network. We also extend this model to consider intermodal connections between air and ground transportation infrastructures. Our results demonstrate potential benefits of intelligent and integrated transportation systems towards overall resilience in various disruption scenarios.

Manuscript received July 24, 2018; revised February 10, 2019; accepted March 31, 2019. Date of publication April 15, 2019; date of current version March 27, 2020. The Associate Editor for this article was W. Jin. *(Corresponding author: Karl H. Thompson.)*

The authors are with the Department of Aerospace Engineering, University of Illinois at Urbana–Champaign, Urbana, IL 61801 USA (e-mail: karlht2@illinois.edu; huytran1@illinois.edu).

Digital Object Identifier 10.1109/TITS.2019.2909177

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/

II. Background

A. Critical Infrastructures

Much work has been dedicated to the study of critical infrastructures in recent years in an effort to understand their vulnerabilities and interdependencies and, by extension, the optimal options for their protection. In its 1997 report, the President's Commission on Critical Infrastructure Protection [9] found that while there was no indication of an immediate threat, a wide range of infrastructure assets remained highly susceptible to unsophisticated attacks, both by foreign and domestic actors. The report also highlighted the role growing complexity and interconnectivity among infrastructure systems could inadvertently play towards exacerbating the effects of an outage or attack on even a minor network component. Extending this analysis, Rinaldi *et al.* [10] explored the various connections and dependencies between infrastructure systems, identified a number of common infrastructure characteristics that are central to interdependency studies, and added that further analysis of these connections is crucial to optimal infrastructure operation and defense. Zimmerman [11] theorized that decisions made throughout infrastructure project phases from planning to operation and maintenance could cause many unintended vulnerabilities due to limited understanding of these complexities, and therefore developed a method to track and share a catalog of infrastructure interdependencies between decision makers to mitigate this problem.

Researchers have taken different approaches in search of ways to reinforce and protect critical infrastructures. Brown *et al.* [12] introduced the application of new bi-level and tri-level optimization models to improve the resilience of various infrastructure networks such as the US Strategic Petroleum Reserve and Border Patrol against terrorist attacks, utilizing open-source historical data sets pertaining to these networks. Scaparra and Church [13] developed an approach for solving the r-interdiction median problem with fortification (RIMF) that can be utilized to form an effective defensive strategy for a given infrastructure system that increases its resilience in the face of an attack. Murray *et al.* [14] established a spatial optimization mathematical model called the flow interdiction model (FIM) that is able to assess the worst-case effect on network flow given a specified interdiction scenario. While these three approaches are similar, they do differ in a few ways. Most notably, the first two favor a multi-level formulation for their representation of the interdiction problem, while the third favors a single-level formulation. Primary interdiction metrics used by these methods are arc length, weighted distance, and network flow. Building on these methods, our analysis utilizes operational cost, defined as a function of network flow and routing expense, as its main interdiction metric. Details of our approach follow in Section 3.

B. Transportation Networks

Transportation infrastructure has been a key subject of resilience studies recently due to its elevated importance to the economy and industry supply chain. Authors have therefore attempted to devise methods to assess and improve the resilience of different transportation systems to a wide range of real and artificial disturbances. Cox *et al.* [15] presented a metric for assessing operational resilience of transportation systems based on the economic resilience metric first introduced by Rose [16], and applied it to a case study of the London 2007 bus and subway terror attacks. In a different direction, Alderson *et al.* [17] analyzed the resilience of the San Francisco regional highway system accounting for peak travel periods and nonlinear traffic congestions using a sequential tri-level program.

Resilience of air transportation systems, on the other hand, has been studied less extensively. This may be attributed to their inherent large scales and dependence on consolidated regional operators. Nevertheless, some researchers have attempted to tackle the problem of modeling and assessing the resilience of air transportation networks using specific regions as case studies. For example, Cardillo *et al.* [18] examined the dynamics of air transport in Europe using a multi-level modeling approach, and analyzed the resilience of the network against random flight failures. Dunn and Wilkinson [19] also analyzed air transportation resilience from a network theoretic point of view, and suggested two methods for increasing resilience: the first involves adaptively modifying the network route structure in response to disruptions, while the second proposes permanently changing route structure in favor of an optimized network topology. In a different direction, Dray *et al.* [20] and Marzuoli *et al.* [21], [22] investigated the effectiveness of using ground transportation, in coordination with air operations, to alleviate airline costs and passenger delays in the aftermath of airport-wide disruptions in Europe and the United States, respectively. Their results suggest that both metrics could indeed be reduced, and the air network resilience ultimately improved, with the use of ground transportation to optimally reroute stranded passengers within the connected networks. These works, however, do not thoroughly consider optimal disruptions caused by intelligent adversaries, nor the ways in which their effects could be mitigated through network defense plans.

C. Sequential Game Models

Attacker-defender (AD) models are an approach to analyzing critical infrastructure resilience with explicit consideration of intelligent adversaries and optimal recovery actions. AD models, and their extension defender-attacker-defender (DAD) models, are multi-level, deterministic Stackelberg games that mimic the dynamics of a real system with operators and defenders, as well as possible attackers. They were first introduced as infrastructure resilience assessment tools by Brown *et al.* [12], building on the works of Golden [23] and Wood [24]. The fundamental premise of these models is the simulation of three opposing agents (two in the case of AD models). The first agent is a 'defender' (described by Brown as the system operator or user) who seeks to minimize the network's operational cost through the efficient routing of system resources. The second is an 'attacker' that seeks to inflict the most damage on system components given a limited offensive budget in order to drive the operational cost up.

And finally, the third agent is a second 'defender' that aims to mitigate the attacker's actions through the defense of system components given a limited defensive budget.

Since their introduction, DAD models have been used in a variety of applications. Ding *et al.* [25] developed a DAD model that represents a power grid defense problem with variable attacks and load types, and demonstrated results for a bulk power reliability test system. Costa *et al.* [26] also examined the power grid fortification problem; however, they formulated it as a two-stage optimization program where the defender's decisions directly influence the attacker's uncertainty set, impeding their ability to target certain network components. In the field of networked systems, Rao *et al.* [27] considered the problem of defending connected classes of infrastructures from physical and cyber attacks that target their communications. They formulated the problem as a bi-level AD game, where the costs and benefits of offensive and defensive actions are defined using sum-form and product-form functions.

This paper examines our development of a DAD model for the air network resilience problem. Sections 3 and 4 detail limitations of existing DAD models for our problem, our methodology to address these issues, and results from application of our approach to the US air transportation network based on publicly available data.

III. Methods

Previous efforts have demonstrated the suitability and effectiveness of DAD models towards assessing the resilience of large-scale infrastructure systems. However, one common theme in many of these studies is the focus on physical elements, such as electrical transmission lines, transformers, highway segments, etc., as the subject of both attacks and reinforcements. Our analysis, which examines the long-term resilience of the commercial US air transportation network, is approaching the problem from a different angle. In our view, air routes, despite being virtual components of the air network, are as critical to the resilience of the overall system as the physical airports that support them. The premise for this notion is twofold. First, due to commercial and competitive considerations, airline operators rarely modify or cancel key routes in the network that drive most of the passenger traffic. This allows us to assume that all air routes can be considered as fixed assets, such as roads and bridges, within a reasonable time period without a significant loss of accuracy. Second, while it undoubtedly carries a much more significant impact, an airport-wide disruption is far less likely to happen than a disruption to a specific route. Indeed, adverse weather conditions, human disturbances during flights, and other factors make route-specific disruptions a more common occurrence, and therefore, an appropriate subject for air transportation resilience analysis. Furthermore, many implementations of DAD models allow traffic over the entire network to be reoptimized after an attack. However, we believe that re-optimization of the entire network is unrealistic for air traffic operations, where traffic is not continuous but discrete. Instead, our approach assumes that traffic and supply and demand on individual air routes are not affected by disruptions on other routes, and therefore only affected passengers are rerouted after an attack.

Our DAD model takes a tri-level formulation. The system operator (i.e. the first 'defender') aims to reroute all passengers affected by attacker actions to their destinations in a cost-effective manner, based on supply and demand. The attacker agent seeks to maximize the operator's operational costs by targeting and disrupting key air routes. Finally, the defender agent attempts to mitigate potential damage to the network resulting from the attacker's actions by fortifying a number of routes in a way that makes them invulnerable to attack. The network supply and demand are extracted directly from the Airline Origin and Destination Market Survey (DB1BMarket) dataset [28], issued by the Bureau of Transportation Statistics (BTS). The dataset contains a 10% sample of airline tickets sold by reporting carriers from Q4 2012 to Q3 2017, showing the origin and destination airports, ticket price, number of passengers, and the flying distance for each entry. In our model, the number of passengers and average ticket price on each route taken from the dataset are used as proxies for the network supply and demand at each airport (node) and the cost for traversing a route (arc), respectively.

This formulation models worst-case air route disruptions. These disruptions could be due to intentional, man-made events, such as terrorist attacks, or the result of natural disasters, such as hurricanes or wildfires. We present the formulation specifically in the context of intelligent attacks under the assumption that an intelligent attacker rationally optimizes its attack to inflict maximum damage on network operations. In contrast, non-intelligent attacks, like severe weather disruptions, are not guaranteed to cause the maximum amount of damage to the network, and are better modelled as random variables. We focus on intelligent and intentional attacks within this study to understand impacts of potential worst-case disruptions. Additionally, we find that analyzing disruptions from a route-centric, rather than an airport-centric, perspective is beneficial as it opens the door to finding potential new measures to improve the air network's long-term resilience. Such measures could include establishing new modes of transportation with the objective of augmenting critical air routes, as well as the identification of ground transportation routes that could be utilized for passenger rerouting after a disruption. Finally, we note that our formulation could be modified in the future to model more detailed temporal effects of attacks, which in turn would enable the analysis of specific times or days of interest to system planners.

A. Problem Formulation

Our resilience assessment problem is formulated as a nonlinear optimization program with the following form and constraints:

$$\min_{D} \max_{A} \min_{P,R} \sum_{[i,j]\in E} \left[t_{ij} + h_{ij} A_{ij} \left(1 - D_{ij}\right)\right] P_{ij} + \sum_{n,q\in N} l_{qn} R_{qn} \tag{1}$$

$$\text{s.t.} \sum_{[i,j]\in E} A_{ij} \leq AB \quad \forall [i,j] \in E \tag{2}$$

$$\sum_{[i,j]\in E} D_{ij} \leq DB \quad \forall [i,j] \in E \tag{3}$$

TABLE I
DEFINITIONS OF MODEL INPUT DATA AND DECISION VARIABLES

Variable	Definition
t_{ij}	Mean airfare on a specified route $[i,j]$
h_{ij}	Penalty for traversing an attacked or non-existent route $[i,j]$
A_{ij}	Binary variable; 1 if route $[i,j] \in E$ is chosen to be attacked, 0 otherwise
D_{ij}	Binary variable; 1 if route $[i,j] \in E$ is chosen to be defended, 0 otherwise
P_{ij}	Integer variable: total number of passengers travelling on route $[i,j] \in E$
l_{qn}	Penalty of non-routed passengers at airport $q \in N$ who were originally on an itinerary ending at airport $n \in N$. In practice: $l_{qn} < h_{ij}\ \forall\, n,q \in N,\ [i,j] \in E$
R_{qn}	Integer variable: number of non-routed passengers at airport $q \in N$ who were originally on an itinerary ending at airport $n \in N$
AB	Budget constraint on the number of allowed simultaneous attacks (integer)
DB	Budget constraint on the number of allowed simultaneous defenses (integer)
P_{qij}	Integer variable: number of passengers travelling on route $[i,j] \in E$ who originated at airport $q \in N$
m_{qn}	Total demand rate at airport $n \in N$ for passengers originating at airport $q \in N$. Total supply rate of passengers originating at airport q is represented by m_{qq}
u_{ij}	Upper bound on total number of passengers on a specified route $[i,j] \in E$

$$\sum_{(n,j)\in Z} P_{qnj} - \sum_{(i,n)\in Z} P_{qin} - R_{qn} \le m_{qn} \quad \forall n,q \in N \tag{4}$$

$$\sum_{q\in N} P_{qij} - P_{ij} = 0 \quad \forall [i,j] \in E \tag{5}$$

$$A_{ij} \in \{0,1\} \quad \forall [i,j] \in E \tag{6}$$

$$D_{ij} \in \{0,1\} \quad \forall [i,j] \in E \tag{7}$$

$$R_{qn} \ge 0 \quad \forall n,q \in N \tag{8}$$

$$0 \le P_{ij} + P_{ji} \le u_{ij} \quad \forall [i,j] \in E \tag{9}$$

where $[i,j] \in E$ refers to an undirected route between airports i and j, $(i,j) \in Z$ refers to a directed route between i and j, and the indices $n,q \in N$ represent airport nodes. The set E is defined as the set of all undirected routes, Z is the set of all directed routes, and N is the set of all nodes.

The input data and decision variables included in Equations (1-9) are defined in Table I. The variable t_{ij} represents the per-passenger operational cost incurred on the system operator for utilizing route ij. Variables h_{ij} and l_{qn} are arbitrary high-value penalty figures designed to incentivize the system operator to only utilize functional routes and to reroute as many passengers as possible, respectively. Both variable values are arbitrary, but for computational stability purposes, h_{ij} should be, at a minimum, an order of magnitude higher than l_{qn}. The variable u_{ij} represents the upper limit of traffic on a given route at any time, and is taken to be 190% of the nominal traffic value. Finally, variable m_{qn} constitutes the supply and demand of passengers with respect to each pair of airports q and n.

The objective function, shown in Equation (1), is a three-level problem. The first level represents the system operator attempting to minimize operational costs incurred by airlines through rerouting of passengers affected by disrupted routes. This process is represented mathematically through the multiplication of the ticket cost variable t_{ij} and the number of route travelers variable P_{ij} on all air routes. In the last term of the equation, the variable R_{qn} exists as a deterrent to prevent the network operator from failing to route travelers to their destinations through the use of the penalty variable l_{qn}. The second level of the problem represents the system attacker aiming to maximize that same operational cost to the system operator through the disruption of a set of air routes, which in turn incurs additional rerouting costs. This process is achieved mathematically by the attacker agent selecting one or more routes in A_{ij} that then get multiplied by the penalty variable h_{ij}, effectively making it impossible for the first-level operator to utilize disrupted routes to transport passengers. The third level then represents the system defender attempting to mitigate the effects of the attacks by 'defending' one or more air routes in a way that makes them invulnerable to attack. This objective is mathematically attained by the $\left(1 - D_{ij}\right)$ term, which effectively prevents attacks on defended routes.

Eight constraints are included in the optimization problem, represented by Equations (2-9). Constraints (2) and (3) ensure that the number of attacked and defended air routes fall within their allocated, respective budgets. Constraints (6) and (7) define the acceptable values for an air route to have in the arrays A_{ij} and D_{ij} as the integers zero (for non-attacked and non-defended routes) and one (for attacked and defended routes). Constraints (4) and (5) ensure that original passenger supply and demand is satisfied by making certain at each airport, for each airport pair, that the sum of the number of passengers arriving minus the number of passengers departing and the number of passengers who are stranded is less than or equal to the original data-driven passenger supply and demand value. Constraint (8) defines the variable R_{qn}, representing the number of unrouted passengers at each airport for each airport pair, as a positive integer. Finally, Constraint (9) sets an upper limit for the number of passengers travelling on each route at any given time according to the values of the variable u_{ij}.

In our representation of the airline operational cost as the product of ticket cost and passenger numbers, we assume that this product serves as a good proxy for the true operational cost incurred on the airlines. For example, we assume this cost accounts for wages for crew, ground, and management teams among other expenses, due to the proportionality of these costs to flight sizes.

B. Decomposition Algorithm

Our formulation of the problem, represented by Equations (1-9), is solved computationally by converting it into a mixed-integer, non-linear optimization program in the form:

$$z = \min_{d\in D} \max_{a\in A} \min_{p\in P, r\in R} f(a,d,p,r), \tag{10}$$

$$\begin{aligned} f(a,d,p,r) &= \sum\nolimits_{[i,j]\in E} \left[t_{ij} + h_{ij}a_{ij}\left(1-d_{ij}\right)\right] p_{ij} \\ &\quad + \left[t_{ji} + h_{ji}a_{ji}\left(1-d_{ji}\right)\right] p_{ji} \\ &\quad + \sum\nolimits_{n,q\in N} l_{qn}r_{qn} + l_{nq}r_{nq} \end{aligned} \tag{11}$$

$$A = \left\{ a \in \{0,1\}^{|E|} | \sum_{a \in A_{ij}} a_{ij} \leq AB \forall [i,j] \in E \right\} \quad (12)$$

$$D = \left\{ d \in \{0,1\}^{|E|} | \sum_{d \in D_{ij}} d_{ij} \leq DB \forall [i,j] \in E \right\} \quad (13)$$

$$P = \left\{ p \in R^{|N||E|} | \sum_{(i,j) \in Z} p_{nij} - p_{nji} \leq m_{nn} + r_{nn} \forall n \in N \right\} \quad (14)$$

$$R = \left\{ r \in R^{|N||E|} | r_{qn} + r_{nq} \geq 0 \forall n, q \in N \right\} \quad (15)$$

Such a non-linear program, however, is often difficult to compute directly. To reach more efficient computation, researchers have historically resorted to decomposition [29], [30] and implicit enumeration techniques [31], among other approaches. For our model, we apply elements of the decomposition technique described by Alderson *et al.* [32]. The full decomposition algorithm works as follows:

1. *Initialization*:
   a) Input DB1BMarket data and optimality tolerance $\varepsilon > 0$.
   b) Set the lower bound (LB) to $-\infty$ and the upper bound (UB) to ∞.
   c) Set loop counter K to 1.
   d) Set initial defense plan $D_{ij}^{(0)} = 0$.
2. *Attack Subproblem*:
   a) Given defense plan $D_{ij}^{(K)}$, compute attack plan $A_{ij}^{(K)}$ and rerouted traffic $P_{ij}^{(K)}$ such that $z_{AS}^{UP} - z_{AS}^{LO} \leq \varepsilon z_{AS}^{LO}$.
   b) If $z_{AS}^{UP} <$ UB, update UB to z_{AS}^{UP}.
   c) If any previous attack is repeated, add the temporary constraint $\sum A_{ij}^{(K)} - A_{ij}^{(K-1)} \geq 1$.
   d) If UB LB $\leq \varepsilon$ LB, end the program.
3. *Defense Subproblem*:
   a) Compute defense plan $D_{ij}^{(K+1)}$, such that $z_{DS}^{UP} - z_{DS}^{LO} \leq \varepsilon z_{DS}^{LO}$.
   b) If $z_{DS}^{LO} >$ LB, update LB to z_{DS}^{LO}.
   c) If UB LB $\leq \varepsilon$ LB, end the program.
4. *Looping*:
   a) Update loop counter $K = K + 1$.
   b) Return to step 2a.

The terms z^{LO} and z^{UP} refer to the lower and upper bounds of the value of the respective subproblem's objective function at the termination of the algorithm. In step one, the program processes the DB1B input data and sets values to the various constants and variables described in the problem formulation as well as the optimality tolerance ε, loop counter K, and initial defense plan $D_{ij}^{(0)}$. Steps two and three describe the computation of the various iterations of the attack and defense plans, respectively, until an optimal solution is found and program is ended. In step two, to avoid the repetition of previous attacks, which could lead to cycling, we enforce that a new attack plan be found that contains at least one unique route not included in previous plans. Finally, if the program is not terminated in either step two or three, the loop counter is updated, and the two subproblems are re-solved until an optimal, or near optimal, objective value is reached.

C. Application Problem

Our main data source is the DB1BMarket table within the Airline Origin and Destination Survey database that is maintained by the BTS. As the table only contains a 10% sample of tickets sold by reporting air carriers, a need arises for a way to extrapolate the data to reflect actual air traffic figures on the entire air network that would be impacted by a disruption on any given route. Our data, spanning a five-year period from Q4 2012 to Q3 2017, contains records for approximately 240.6 million passengers flying domestically. This figure translates to a daily average of 131.8 thousand passengers across all domestic US routes. The Federal Aviation Administration (FAA) estimates the daily number of passengers flying in and out of US airports to be approximately 2.587 million [33]. Moreover, the Department of Transportation, in its US International Air Passenger and Freight Statistics report [34], estimates that approximately 217.3 million passengers have travelled to and from the United States by air for the year-ended March 2017. Assuming mostly uniform travel patterns throughout the year, we can then calculate the average international travel rate to be 595.3 thousand passengers per day. We finally subtract the international travel rate from the total travel rate to obtain an average net domestic travel volume of 1.992 million daily passengers. Using linear interpolation, we multiply our DB1B data points by a factor of 1.992/0.1318=15.11 to obtain an approximate figure for daily traffic volume on all domestic US air routes, visualized in Figure 1.

Figure 1 shows the average ticket costs for domestic US air routes plotted versus their average estimated daily traffic volumes. The number of routes composing the study's air network, and that are represented in the graph, is 15,469. They were selected after eliminating all routes with fewer than 1,000 annual passengers from the raw dataset. The routes connect a total of 327 airports, all of which are commercial service airports. In the figure, we notice three notable groups of routes: a group comprising routes that are high in both average ticket costs and number of daily passengers such as JFK-LAX and JFK-SFO, a group comprising routes with relatively higher ticket costs but lower average passenger traffic such as PHL-SFO and EWR-HNL, and finally a group composed of routes with relatively higher average daily number of passengers but lower ticket costs such as LAX-SFO and LGA-ORD. In the next section, we examine the overall network dynamics, as well as the relative vulnerabilities of these groups.

IV. Results

We now examine some practical results pertaining to the resilience of the US air transportation network. First, we observe and interpret the application of the model to the air network for multiple cases representing different attacker and defender capabilities. This is achieved by varying the attack and defense budget parameters, effectively limiting the number of air routes the attacker and defender can target and defend, respectively. We then investigate the potential effect integration with other transportation modes could have on the resilience of the air network. To this end, a second layer is added to the model representing an operational bus network that has the same nodes as the air network but with a different topology representing the US highway network.

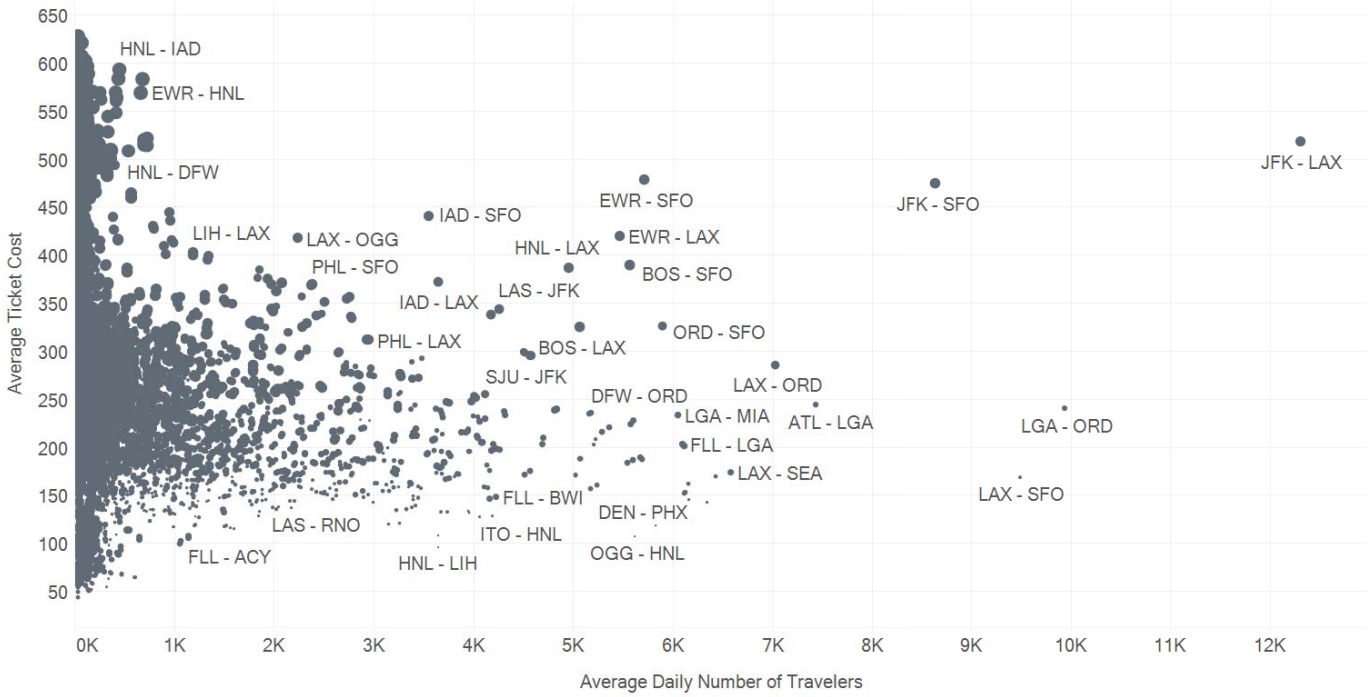

Fig. 1. Distribution of the number of daily travelers and ticket costs across all domestic US air routes. Circle sizes represent relative route lengths.

A. Operational Scenarios of Interest

As defender-attacker-defender models are intended to provide general insights to inform high-level policy decisions, it follows that the quality of these recommendations increases with the number of data points used as their underlying basis. In other words, the higher the number of cases an optimization model is used to analyze, the better overall idea we get of the implications of its results due to various uncertainties and assumptions made in the model. For our air network model, we chose to examine 25 cases representing attack and defense budgets ranging from one to five. We find that this is an adequate choice due to the longer run times incurred by additional run cases and the limited value added by them. The average computing time for each case is approximately 20 minutes on a quad-core second-generation Intel i7 processor. Before attempting to analyze the overall network behavior, however, we first examine three individual run cases of interest: a case with a larger attack budget than defense, a case with the opposite conditions, and a case with balanced budgets.

Figure 2 shows the resulting network dynamics for a case where the defender's budget is five times that of the attacker. We see that the defender used their budget to defend key routes connecting different regions, while the attacker made use of their limited resources to target a high-traffic, regional route. This result underscores the advantages of a higher defense budget with respect to protecting critical resources, and in effect, limiting the effects of a potential attack. Based on the inset close-up image of passenger rerouting operations following the attack in Figure 2, we observe that the system operator made use of multiple airports in Southeastern California to reroute passengers, with a majority of the traffic passing through the LAX and SAN airports.

We now consider the case where the attacker's budget outweighs the defender's, also with a margin of five to one, shown in Figure 3. We notice that the attacker is now able to more freely target long-range routes connecting the coasts. The result is a near six-fold increase in the number of disrupted passengers compared to the first case, as well as a 14% increase in the per-passenger rerouting cost incurred on the system operator. In practice, this result underscores the potential cost of failing to adequately appropriate funds for the protection of critical infrastructures in an active rather than a reactive manner.

We also explore the effect of balanced attacker and defender budgets, with each set to five. The resulting network behavior is shown in Figure 4. We notice that the attacker is still able to target multiple long-range routes. However, thanks to the added resources, we see that the defender is able to minimize the overall effect of the attacker's actions by protecting important routes in different regions of the network. The outcome is a 38% decrease in the number of disrupted travelers compared to the previous case, and a 9% decrease in the per-passenger rerouting cost. This result further cements our view that system operators and policy makers alike ought to keep track of potential threats to their critical systems, and devise appropriate defense strategies to protect them. Table II summarizes numbers of disrupted travelers and per-passenger rerouting costs for these three cases.

B. Air Network Resilience Metrics

Two key metrics of interest to this study are the number of potentially impacted travelers for each of the attack scenarios, and how effective increasing the defense budget is at reducing that impact. Figure 5 shows the average number of daily

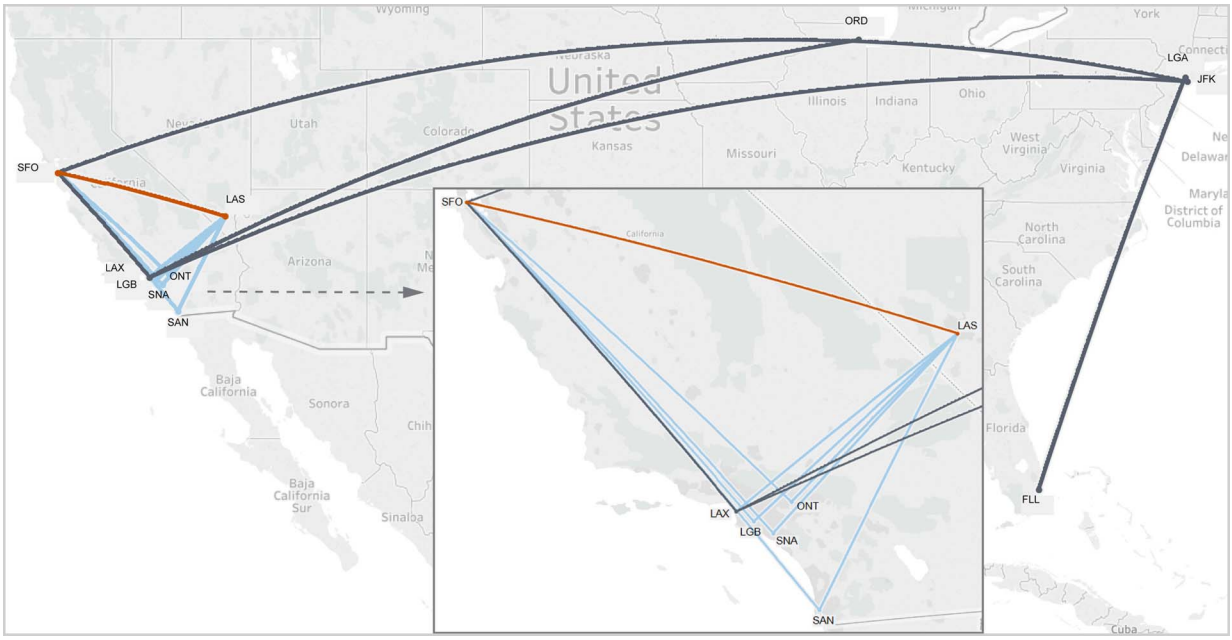

Fig. 2. Best defended (black) and attacked (red) routes for an Attack budget of one and a defense budget of five, with routes used for Optimal passenger rerouting (blue).

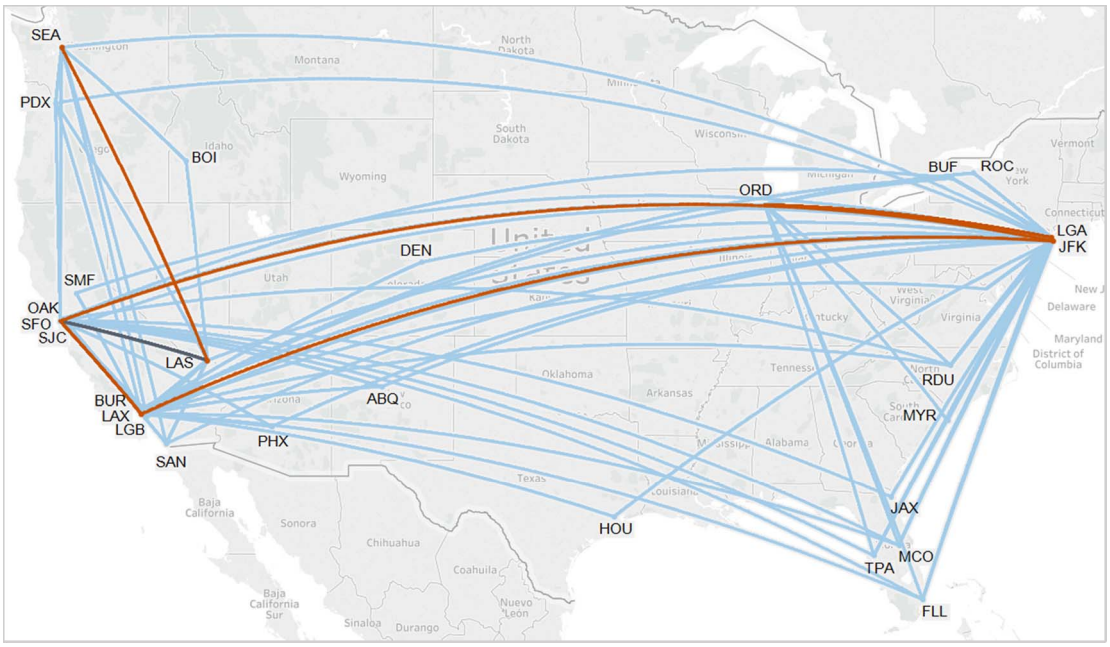

Fig. 3. Best defended (black) and attacked (red) routes for an attack budget of five and a defense budget of one, with routes used for optimal passenger rerouting (blue).

travelers that could potentially be disrupted given different attack and defense scenarios. As expected, we notice that as the attack budget grows, the number of affected travelers generally does as well, if the defense budget remains constant. Similarly, as the defense budget is increased, we notice a decrease in the number of disrupted travelers in some cases for the same attack budget. While the results are inconclusive, they are indicative of a trend of the potential disruption being proportional to the discrepancy between offensive and defensive resources. In practical terms, these results indicate that in order to minimize the effects of an attack on a given network, defensive capabilities should sufficiently surpass that of a potential attacker, or the estimated impact of a natural disaster.

Further examination of the results displayed in Figure 5 reveals that, in some cases, in order to minimize operational costs due to route disruptions, the system defender opts to utilize increased defense budgets to protect

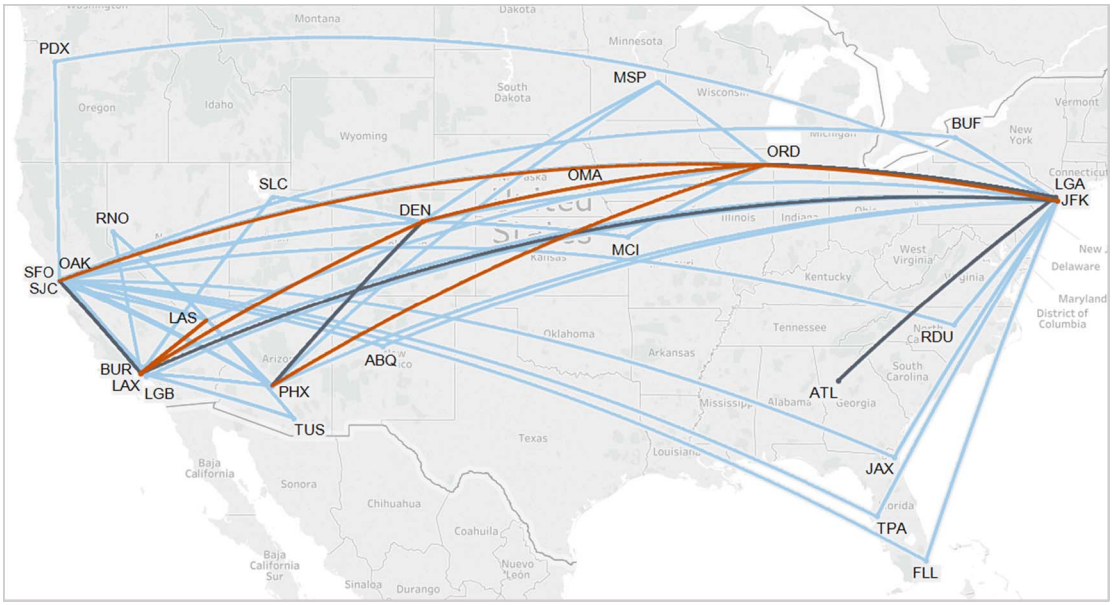

Fig. 4. Best defended (black) and attacked (red) routes for attack and defense budgets of five, with routes used for optimal passenger rerouting (blue).

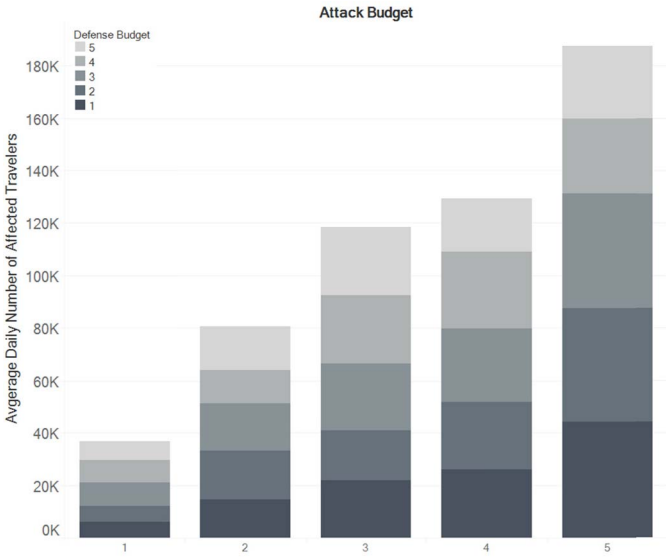

Fig. 5. Estimated number of daily affected passengers for attack and defense budgets ranging from one to five.

TABLE II
SUMMARY OF THE THREE TEST CASES

Attack budget	Defense budget	Number of disrupted passengers	Per-passenger rerouting cost ($)
1	5	7,246	713.26
5	1	44,076	813.61
5	5	27,608	744.37

certain critical routes, even at the cost of increasing the overall number of disrupted passengers. This observation highlights a key difference between passenger-centric, and our operator-centric view of the air network passenger rerouting problem. In a passenger-centric model, the number of impacted travelers is generally expected to increase as the attack budget increases for a constant defense budget, and vice versa. However, in our model, due to the incorporation of average route ticket cost as a proxy for the operational cost, and the use of this operational cost along with route traffic to formulate the attacker, defender, and operator actions, we reach a non-monotonic network behavior that attempts to attain optimal choices on a case-by-case basis.

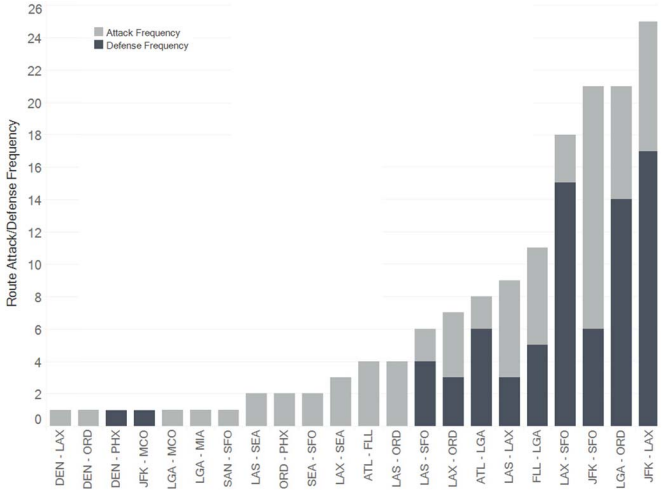

Fig. 6. Most frequently attacked and defended routes based on twenty-five test cases representing attack and defense budgets ranging from one to five.

We also examine the routes that are most attacked and defended in our 25 test cases to characterize the relative importance of individual routes in the network. Figure 6 shows the frequency with which different routes were chosen by the model to be attacked or defended. Inspecting the top ten routes, we notice that they all contain some combination of the east coast's JFK and LGA and west coast's LAX, LAS and SFO airports, either interchangeably, or with other key airports. This result indicates that connectivity between the two coasts is potentially a more important factor in shaping the resilience

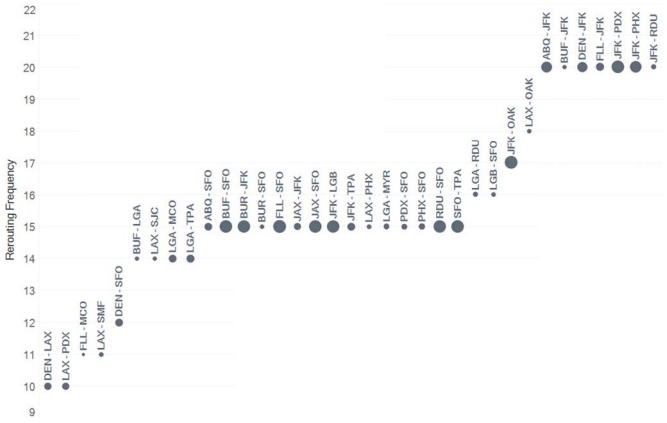

Fig. 7. Most frequently used routes for rerouting operations following one or more attacks. Circle sizes represent relative route lengths.

of the US air network than other criteria such as sheer traffic volume.

Looking back at the distribution of ticket prices and passenger traffic in Figure 1, we notice that the four rightmost routes – namely, JFK-LAX, JFK-SFO, LGA-ORD, and LAX-SFO – also dominate the attack and defense frequency chart. Conversely, routes such as EWR-SFO and BOS-SFO do not appear in the attacked nor the defended routes at all, despite both also having high daily passenger traffic and average ticket costs. This observation ties back to our original point that airport characteristics, beyond traffic and ticket cost, play an important role in the selection of routes for attack and defense. Specifically, the more connected an airport is to the rest of the network, the more valuable each of its routes are, from a resilience point of view, compared to other routes of comparable statistics but with less connected airports.

Relatedly, an analysis of the most frequently used routes by the network operator for rerouting air passengers in the aftermath of an attack is conducted, the results of which are shown in Figure 7. We again notice the prevalence of routes containing at least one key east or west airport. This is to be expected as a natural result of the more frequent disruption of routes serviced by these airports. Nevertheless, their relatively higher significance to the resilience of the air network in general is an area that could benefit from future research.

C. Bus Network Integration Case Study

In order to fully investigate the different elements that impact the resilience of a complex infrastructure system such as the US air transportation network, it is necessary to also consider its connections to and interdependencies with other transportation networks. As a first step, the bus network is a prime candidate for such analysis because of its ubiquitous availability and considerable popularity for low to mid-range travel. The Highway Performance Monitoring System (HPMS) geospatial database [35], a component of the National Transportation Atlas Database (NTAD) database issued by BTS, contains nation-wide data on the performance, characteristics, and conditions of all public roads in the US. Using this data, we construct a bus network for preliminary exploration of this problem, with road segments from HPMS serving as network edges and airport locations from DB1B serving as network nodes. We integrate this bus network into our optimization program described in Section III.A by allowing the network operator to reroute passenger over air or bus routes connecting their original origin and destination pairs. In order to model the temporal component of rerouting operations, we use BTS's Airline On-Time Performance Database [36] to define an average flight time for each origin-destination airport pair. To derive a similar metric for the bus network, we calculate the ground travel time between each airport pair by summing the lengths of the HPMS road segments along the shortest path divided by segment speed limits. For road segments with missing associated HPMS speed limits, we use an assumed speed limit of 65 mph.

To define an approximate measure of the operational cost associated with the two transportation modes, without explicit pricing data for bus transportation, we use the following formulation:

$$operational\ cost = ticket\ price \times (idle\ time + travel\ time)$$

where the bus ticket price is assumed to be 30% of that for air travel for sub-300-mile trips, 60% for sub-800-mile trips, and 100% for over-800-mile trips, where air travel costs are based on DB1B data. Idle times for both air and bus modes are estimated to be: 60, 60, and 90 minutes for air and 15, 30, and 60 minutes for bus transport for sub-300-mile, sub-800-mile, and over-800-mile trips, respectively.

The bus network is then added as a second layer to the air network in our mathematical model. To determine the practical implications of this change, we re-run the case with an attack budget of five and a defense budget of one, given the multi-modal network structure. Figure 8 shows the impacted routes for the integrated network. We observe that the attacker focused their effort on disrupting mid and long-range routes connecting key east and west coast airports to those in the inland. The defender, to counteract this plan, opted to protect inter-coast connectivity by defending the JFK-LAX route. In this case, the number of daily affected passengers is estimated to be 36,881, and the average rerouting cost incurred on the system operator per passenger is $724. Comparing this result to that of the air model alone in Table II indicates that an integration between the air and bus networks has the potential to not only reduce the number of affected passengers in the case of a disruption to air operations, but also lower the overall passenger rerouting cost resulting from such a disruption.

Figure 9 shows the most attacked and defended air routes across 25 run cases of the integrated model. In contrast to the same analysis of the air network alone, shown in Figure 6, two differences stand out between the two. First, we observe that the JFK-LAX route was chosen to be protected in all 25 run cases, compared to only 17 in the first analysis. This result is to be expected, and it can be attributed to the large annual traffic volume on the route and the lack of direct ground transportation options that could alleviate passenger rerouting burden in the event of a disruption. Second, we note that the

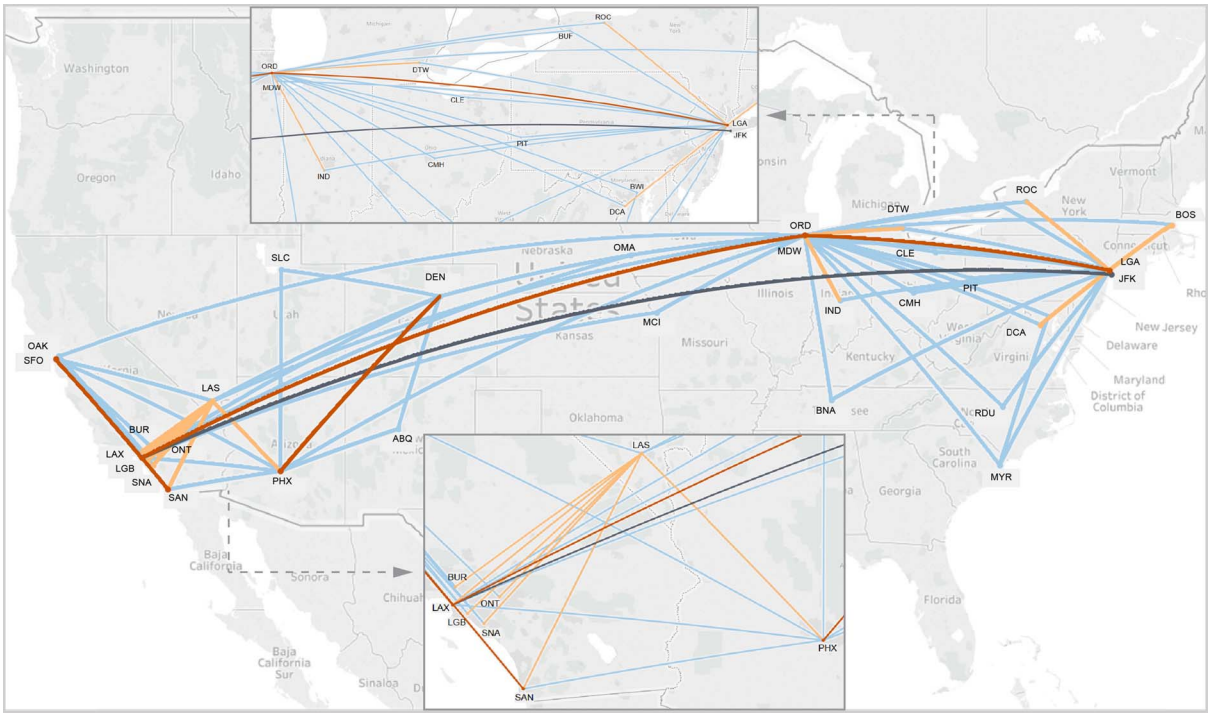

Fig. 8. Best defended (black) and attacked (red) routes for an attack budget of five and a defense budget of one, with air (blue) and bus (orange) routes used for optimal passenger rerouting.

Fig. 9. Most frequently attacked and defended air routes in the integrated air-bus network. Plotted are twenty-five test cases representing attack and defense budgets ranging from one to five.

JFK-SFO route has been targeted 60% less frequently than in the first analysis. This metric could be a result of the persistent protection of the JFK-LAX route and the existence of a variety of transportation options between LAX and SFO that rendered an attack on the route less effective.

These preliminary results highlight the potential advantages an integration between the different transportation modes could have on the overall transportation system's resilience and reliability. This effort would require cooperation between federal, state and local authorities, as well as private contractors, airlines, and regional transportation companies in appropriating the required funding and planning the needed infrastructural changes that would facilitate such an integration. New high-speed, low-cost transportation modes would also likely be needed to fill the gap in transportation range and accessibility between the long-range-but-high-cost air travel and low-cost-but-short-range bus travel.

V. Conclusion

This paper presents the use of a three-stage DAD optimization model to analyze the resilience of the US air transportation network against an intelligent adversary. The model approximates the dynamics between three opposing agents: an operator that seeks to minimize the network's operational cost through optimal passenger rerouting, an attacker that aims to maximize that cost by disrupting a number of network routes, and a defender that is tasked with mitigating the actions of the attacker by protecting key routes. The network topology and behavioral characteristics used in the model are derived from operational data issued by the Bureau of Transportation Statistics. An exploratory case study that examines potential benefits of inter-modal connections on the resilience of the air network is also considered. We find that in order to maintain a resilient air network, appropriate defensive resources ought to be dedicated to augmenting routes maintaining connectivity between the two coasts. We also find that multi-modal integration shows promise in reducing the impact of worst-case attacks on the air network, as measured through the number of disrupted passengers and the cost associated with rerouting passengers. We conclude that multi-modal integration between different transportation modes is potentially a highly valuable endeavor for the public and private sectors to pursue.

REFERENCES

[1] T. G. Lewis, *Critical Infrastructure Protection in Homeland Security: Defending a Networked Nation*. Hoboken, NJ, USA: Wiley, 2015.

[2] (2013). *Presidential Policy Directive—Critical Infrastructure Security and Resilience, Office of the Press Secretary*. [Online]. Available: https://obamawhitehouse. archives.gov/the-press-office/2013/02/12/presidential-policy-directive-critical-infrastructure-security-and-resil

[3] R. Faturechi and E. Miller-Hooks, "Measuring the performance of transportation infrastructure systems in disasters: A comprehensive review," *J. Infrastruct. Syst.*, vol. 21, no. 1, Mar. 2015, Art. no. 04014025.

[4] C. Winston, "Efficient transportation infrastructure policy," *J. Econ. Perspect.*, vol. 5, no. 1, pp. 113–127 , Mar. 1991.

[5] K. Button, "Infrastructure investment, endogenous growth and economic convergence," *Ann. Regional Sci.*, vol. 32, no. 1, pp. 145–162, Feb. 1998.

[6] K. A. Small, "Economics and urban transportation policy in the United States," *Regional Sci. Urban Econ.*, vol. 27, no. 6, pp. 671–691, Nov. 1997.

[7] D. Zeng, S. S. Chawathe, H. Huang, and F. Y. Wang, "Protecting transportation infrastructure," *IEEE Intell. Syst.*, vol. 22, no. 5, pp. 8–11, Sep. 2007.

[8] (2018). *Transportation Systems Sector, Department of Homeland Security*. [Online]. Available: https://www.dhs.gov/transportation-systems-sector

[9] *Critical Foundations: Protecting Americas Infrastructures: The Report of the Presidents Commission on Critical Infrastructure Protection*, Washington, DC, USA: Commission, 1997.

[10] S. M. Rinaldi, J. P. Peerenboom, and T. K. Kelly, "Identifying, understanding, and analyzing critical infrastructure interdependencies," *IEEE Control Syst.*, vol. 21, no. 6, pp. 11–25, Dec. 2001.

[11] R. Zimmerman, "Decision-making and the vulnerability of interdependent critical infrastructure," in *Proc. IEEE Int. Conf. Syst., Man Cybern.*, Oct. 2004, pp. 1–10.

[12] G. Brown, M. Carlyle, J. Salmerón, and K. Wood, "Defending critical infrastructure," *Interfaces*, vol. 36, no. 6, pp. 530–544, Nov. 2006.

[13] M. P. Scaparra and R. L. Church, "A bilevel mixed-integer program for critical infrastructure protection planning," *Comput. Oper. Res.*, vol. 35, no. 6, pp. 1905–1923, 2008.

[14] A. T. Murray, T. C. Matisziw, and T. H. Grubesic, "Critical network infrastructure analysis: Interdiction and system flow," *J. Geographical Syst.*, vol. 9, no. 2, pp. 103–117, Jun. 2007.

[15] A. Cox, F. Prager, and A. Rose, "Transportation security and the role of resilience: A foundation for operational metrics," *Transport Policy*, vol. 18, no. 2, pp. 307–317, Mar. 2011.

[16] A. Rose, "Economic resilience to natural and man-made disasters: Multidisciplinary origins and contextual dimensions," *Environ. Hazards*, vol. 7, no. 4, pp. 383–398, 2007.

[17] D. L. Alderson, G. G. Brown, W. M. Carlyle, and R. K. Wood, "Assessing and improving the operational resilience of a large highway infrastructure system to worst-case losses," *Transp. Sci.*, vol. 52, no. 4, pp. 1012–1034, 2017.

[18] A. Cardillo, M. Zanin, J. Gómez-Gardeñes, M. Romance, A. J. G. D. Amo, and S. Boccaletti, "Modeling the multi-layer nature of the european air transport network: Resilience and passengers re-scheduling under random failures," *Eur. Phys. J. Special Topics*, vol. 215, no. 1, pp. 23–33, Feb. 2013.

[19] S. Dunn and S. M. Wilkinson, "Increasing the resilience of air traffic networks using a network graph theory approach," *Transp. Res. E, Logistics Transp. Rev.*, vol. 90, pp. 39–50, Jun. 2016.

[20] L. Dray, I. Laplace, A. Marzuoli, E. Féron, and A. Evans, "Using ground transportation for aviation system disruption alleviation," *J. Air Transp.*, vol. 25, no. 3, pp. 95–107, Jul. 2017.

[21] A. Marzuoli *et al.*, "Improving disruption management with multimodal collaborative decision-making: A case study of the Asiana crash and lessons learned," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 10, pp. 2699–2717, Oct. 2016.

[22] A. Marzuoli *et al.*, "Multimodal impact analysis of an airside catastrophic event: A case study of the asiana crash," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 2, pp. 587–604, Feb. 2016.

[23] B. Golden, "A problem in network interdiction," *Naval Res. Logistics Quart.*, vol. 25, no. 4, pp. 711–713, Dec. 1978.

[24] R. K. Wood, "Deterministic network interdiction," *Math. Comput. Model.*, vol. 17, no. 2, pp. 1–18, Jan. 1993.

[25] T. Ding, L. Yao, and F. Li, "A multi-uncertainty-set based two-stage robust optimization to defender-attacker-defender model for power system protection," *Rel. Eng. Syst. Saf.*, vol. 169, pp. 179–186, Jan. 2018.

[26] A. Costa, D. Georgiadis, T. S. Ng, and M. Sim, "An optimization model for power grid fortification to maximize attack immunity," *Int. J. Electr. Power Energy Syst.*, vol. 99, pp. 594–602, May May 2018.

[27] N. Rao, C. Ma, K. Hausken, F. He, D. Yau, and J. Zhuang, "Defense strategies for asymmetric networked systems with discrete components," *Sensors*, vol. 18, no. 5, p. 1421, Mar. 2018.

[28] (2017). *Bureau of Transportation Statistics, 'Airline Origin and Destination Survey*. [Online]. Available: https://www.transtats.bts.gov/ DatabaseInfo.asp?DB_ID=125

[29] F. Margot, M. Queyranne, and Y. Wang, "Decompositions, network flows, and a precedence constrained single-machine scheduling problem," *Oper. Res.*, vol. 51, no. 6, pp. 981–992, 2003.

[30] E. Israeli and R. K. Wood, "Shortest-path network interdiction," *Networks*, vol. 40, no. 2, pp. 97–111, Jan. 2002.

[31] N. Alguacil, A. Delgadillo, and J. M. Arroyo, "A trilevel programming approach for electric grid defense planning," *Comput. Oper. Res.*, vol. 41, pp. 282–290, Jan. 2014.

[32] D. L. Alderson, G. G. Brown, W. M. Carlyle, and R. K. Wood, "Solving defender-attacker-defender models for infrastructure defense," in *Proc. 12th INFORMS Comput. Soc. Conf.*, Jan. 2011, pp. 36–46.

[33] (2018). *Air Traffic By The Numbers, Federal Aviation Administration*. [Online]. Available: https://www.faa.gov/air_traffic/by_the_numbers

[34] *U.S. International Air Passenger and Freight Statistics*. US Department of Transportation, Washington, DC, USA, 2017.

[35] Bureau of Transportation Statistics. (2018). *National Transportation Atlas Database*. [Online]. Available: https://www.bts.gov/geography/geospatial-portal/NTAD-direct-download

[36] Bureau of Transportation Statistics. (2019). *Airline On-Time Performance Data*. [Online]. Available: https://www.transtats.bts.gov/Tables.asp?DB_ID=120

Karl H. Thompson received the Bachelor of Aerospace Engineering and Mechanics degree from the University of Minnesota at Twin Cities in 2016. He is currently pursuing the master's degree with the Department of Aerospace Engineering, University of Illinois at Urbana–Champaign. In 2017, he joined the Department of Aerospace Engineering, University of Illinois at Urbana–Champaign, as a Research Assistant. Since 2017, he has been doing research in the design and optimization of aerospace systems and complex networks. He is currently a Graduate Student in aerospace engineering at the University of Illinois at Urbana–Champaign. His current research interests are centered on analyzing the resilience of the U.S. air transportation network, and exploring the role of an inter-connected, multi-modal transportation network may play to enhance that resilience.

Huy T. Tran received the B.S. degree in mechanical engineering from North Carolina State University, Raleigh, NC, USA, in 2008, the M.S. degree in mechanical engineering from the University of Wisconsin–Madison, Madison, WI, USA, in 2010, and the M.S. and Ph.D. degrees in aerospace engineering from the Georgia Institute of Technology, Atlanta, GA, USA, in 2014 and 2015, respectively. He is currently a Research Assistant Professor with the Aerospace Engineering Department, University of Illinois at Urbana–Champaign, Urbana, IL, USA. His current research focuses on designing resilient systems through a combination of machine learning, reinforcement learning, network science, optimization, social media analytics, and agent-based modeling. Current application areas include air traffic management, transportation systems and other critical infrastructures, and mission planning for autonomous systems. He is a member of the American Institute of Aeronautics and Astronautics (AIAA) and the Institute of Electrical and Electronics Engineers (IEEE).