Cyber- Threat Assessment for the Air Traffic Management System: A Network Controls Approach

Sandip Roy, Washington State University
Banavar Sridhar, National Aeronautics and Space Administration

1. Introduction

Air transportation networks are being disrupted with increasing frequency by failures in their cyber- (computing, communication, control) systems [1-3]. Whether these cyber- failures arise due to deliberate attacks human errors, or equipment failure, they can have far-reaching impact on the performance of the air traffic control and management systems. For instance, a computer failure in the Washington DC Air Route Traffic Control Center (ZDC) on August 15, 2015, caused nearly complete closure of the Center's airspace for several hours. This closure had a propagative impact across the United States National Airspace System, causing changed congestion patterns and requiring placement of a suite of traffic management initiatives to address the capacity reduction and congestion. A snapshot of traffic on that day clearly shows the closure of the ZDC airspace and the resulting congestion at its boundary, which required augmented traffic management at multiple locations. Cyber- events also have important ramifications for private stakeholders, particularly the airlines. During the last few months, computer-system issues have caused several airlines' fleets to be grounded for significant periods of time: these include United Airlines (twice), LOT Polish Airlines, and American Airlines. Delays and regional stoppages due to cyber- events are even more common, and may have myriad causes (e.g., failure of the Department of Homeland Security systems needed for security check of passengers, see [3]).

The growing frequency of cyber- disruptions in the air transportation system reflects a much broader trend in the modern society: cyber- failures and threats are becoming increasingly pervasive, varied, and impactful. In consequence, an intense effort is underway to develop secure and resilient cyber- systems that can protect against, detect, and remove threats, see e.g. [4] and its many citations. The outcomes of this wide effort on cyber- security are applicable to the air transportation infrastructure, and indeed security solutions are being implemented in the current system [5,6]. While these security solutions are important, they only provide a piecemeal solution. Particular computers or communication channels are protected from particular attacks, without a holistic view of the air transportation infrastructure. On the other hand, the above-listed incidents highlight that a holistic approach is needed, for several reasons. First, the air transportation infrastructure is a large scale cyber-physical system with multiple stakeholders and diverse legacy assets. It is impractical to protect every cyber- asset from known and unknown disruptions, and instead a strategic view of security is needed. Second, disruptions to the cyber- system can incur complex propagative impacts across the air transportation network, including its physical and human assets. Also, these implications of cyber- events are exacerbated or modulated by other disruptions and operational specifics, e.g. severe weather, operator fatigue or error, etc.   These characteristics motivate a holistic and strategic perspective on protecting the air transportation infrastructure from cyber- events.

The analysis of cyber- threats to the air traffic system is also inextricably tied to the integration of new autonomy into the airspace [24]. The replacement of human operators with cyber functions leaves the network open to new cyber threats, which must be modeled and managed. Paradoxically, the mitigation of cyber events in the airspace will also likely require additional autonomy, given the fast time scale and myriad pathways of cyber-attacks which must be managed. The assessment of new vulnerabilities upon integration of new autonomy is also a key motivation for a holistic perspective on cyber threats.

Very recently, a few research efforts have begun to consider on the holistic implications of cyber-threats and other man-made disruptions on the air traffic system, among other infrastructure networks [7-9]. The long-term goals of these research efforts are to:

1) Analyze the NAS-wide impacts of cyber- failures and attacks, as well as other disruptions (e.g. space vehicle operations, integration of unmanned vehicles, human operator error/fatigue).

2) Identify critical vulnerabilities in the cyber- network, such as software components, communication links, or data sources whose failure may incur wide impact; and, likewise, identify critical physical-world vulnerabilities.

3) Develop protection schemes for these critical vulnerabilities.

These holistic analysis and design tasks are challenging. Analysis of NAS-wide impacts requires appropriate models that capture the traffic network, cyber system, and human assets in a way that is tractable yet descriptive. The interfaced model is necessarily high-dimensional, and detailed analyses of wide-area transient responses are needed for threat evaluation. Hence, effective simulation and formal analysis techniques are needed.   Identification and protection of critical vulnerabilities is even more complicated, since it requires comparing and evaluating a large pool of potential threats over variable weather and traffic states.

In this paper, we introduce a *layered network modeling framework* for assessing cyber- threats to the air traffic management system, and advocate for a *network control theory approach* for threat assessment using the model. The described network-theoretic approach builds on a growing literature on graph- and network- theoretic approaches to air traffic management [11,25,26]. Relative to this literature, the main innovation here is to study the spatiotemporal impacts of disruptions from a graph-theory and network-controls perspective, and to develop models and analyses for meshed cyber, traffic, and weather dynamics. The research described here also connects to a broader effort to evaluate threats to cyber-physical systems from a network-controls perspective, but achieves a keener analysis focused on the specific models used in air traffic management (e.g., [19,27-29]).

Specifically, a flow- and queueing- centered modeling paradigm is proposed (Section 2). To assist in identification of vulnerabilities, a linearized approximation is also considered (Section 3). Several exploratory analyses are conducted, that show how the proposed state-space modeling frameworks can allow threat assessment from a control theory perspective (Section 4).

2. Layered Network Model

A network model with three layers is envisioned: 1) a traffic layer which captures air traffic at the resolution of major flows and also major controls (e.g. traffic management initiatives such as GDPs and AFPs); 2) a cyber- layer that abstractly represents the information flow among stakeholders (airline dispatch offices, Centers, ATCSCC) required for operations, and the impacts of this information flow on traffic; and 3) a weather layer that tracks forecasted severe weather impacts on traffic and capacities. The model as whole comprises a multi-layer nonlinear flow and queueing network model, which has structured interfaces between the layers (Figure 1).
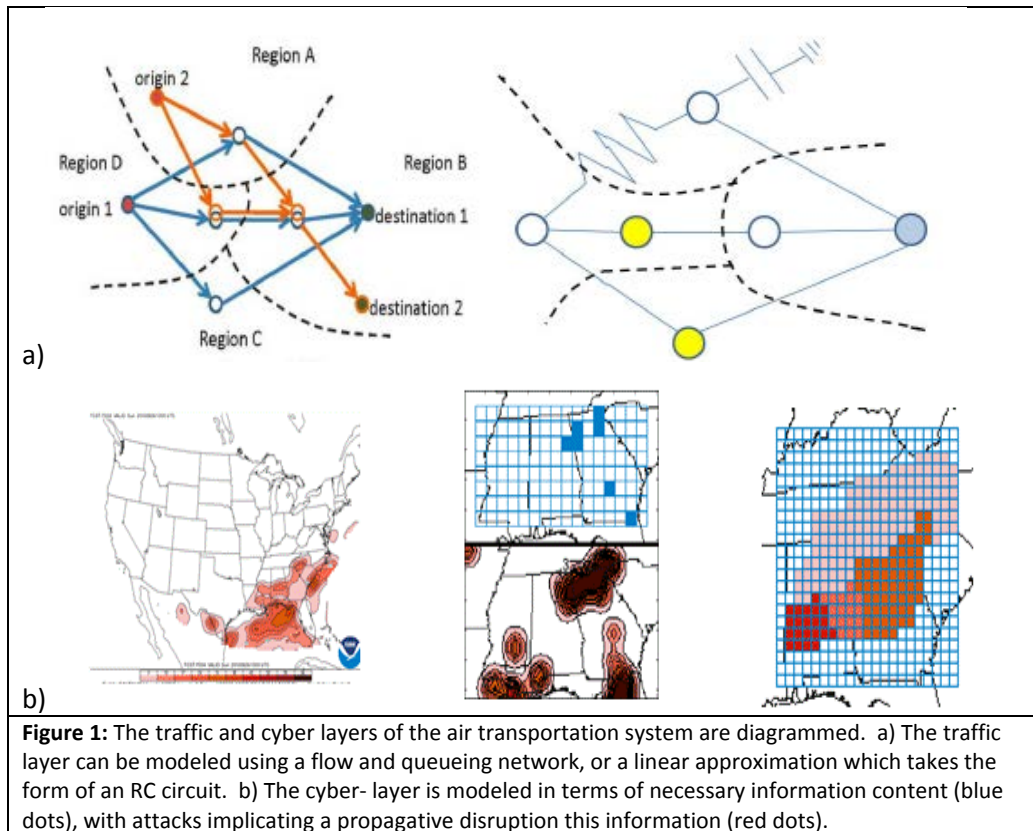
The models for traffic flow and weather, and their interface, were developed in our previous studies [10-12]. These studies are part of a research thrust on flow-level or Eulerian modeling of air traffic (e.g.,

[13,14]). The main innovation in this work is to develop a network model for the cyber (communication, computing, decision-support) architecture of the airspace system, and to interface it with the traffic layer so as to form the full layered-network model. The cyber- layer may be viewed abstractly as modeling the availability and flow of information required for the air traffic system to function. Attacks and failures serve to corrupt information or prevent information flow: a percolation-type model for the propagative impact of attacks is envisioned here [15].

Since the traffic and weather models were described in previous work, we only briefly overview them here, and refer the reader to the literature for details. The model for the traffic and management layer that we use here was introduced in [11]. The model falls within the broad class of flow- and queueing-models, or Eulerian models, that represent aggregate flow densities or traffic counts rather than individual aircraft positions [11,13,14,20]. The particular model considered here represents traffic at the resolution of inter-Sector flow densities in an *area of interest* with high congestion or severe weather, and at a lower resolution outside the area of interest. Specifically, traffic is modeled using overlaid flow networks for different origin-destination (OD) pairs, see Figure 1a. Structured queueing elements are used to represent traffic management initiatives such as ground-delay programs, airspace flow programs, miles-in-trail or minutes-in-trail. Queues also are used to model intrinsic capacity restrictions on airspace resources (e.g., Sector capacities, arrival and departure rate constraints). Demand is modeled as having a deterministic component which represents scheduled traffic, and a stochastic component which reflects schedule uncertainty and pop-up traffic [21]. Resource capacities are modulated by forecasted weather dynamics, see discussion on the weather layer below. Model parameters – including the flow-network structure, demand profiles, possible traffic management initiatives, and nominal capacities – are obtained from archived data along with day-of-operations data. The queueing model has been evaluated for several historical bad-weather days, and has been shown to provide adequate forecasting of traffic characteristics. The model has also proved effective for tuning of traffic management initiatives, so as to optimize multi-objective cost metrics [22,23]. We refer the reader to [11] for a mathematical formulation of the model.

Weather significantly modulates *en route* and terminal area air traffic, and hence modeling the traffic management system requires modeling of forecast weather. At the strategic decision-making horizon, weather is subject is significant uncertainty, and hence appropriate statistical forecasts of weather are needed. The described queueing-network model, in particular, requires estimates of en route (Sector) capacities as well as airport arrival and departure rates, both of which depend on weather; a weather layer is envisioned in our modeling framework to generate these capacities. Although statistical weather forecasting tools are available in the public domain (e.g. ensemble forecasts [30]), these tools often do not output weather data at the proper resolution for traffic management, and also do not capture the regional-scale variabilities and uncertainties in forecast weather. In our previous work, we have used a stochastic automaton network known as the influence model [31] to represent the spatiotemporal progression of severe convective weather, so as to forecast *en route* capacity impacts [32-34]. The main idea is to parameterize the influence model to statistically match public-domain forecasts at snapshot times, whereupon the model can be run and analyzed to get interpolated forecasts at desired resolutions and to capture small-scale variabilities in weather evolution. The convective coverage predicted by the model can then be translated to a reduction in the en route capacity. Meanwhile, airport capacity trajectories can be obtained from local wind, ceiling, and convection variables obtained from ensemble forecasts (or derived influence models), or alternately from terminal aerodrome forecasts (e.g., [35-37]). The weather layer of the proposed model includes the spatiotemporal models for weather evolution, and their translation to airport and en route

capacities, see Figure 1b. The weather layer is interfaced with the traffic layer in that it sets en route and terminal area capacities.

Figure 1: The traffic and cyber layers of the air transportation system are diagrammed. a) The traffic layer can be modeled using a flow and queueing network, or a linear approximation which takes the form of an RC circuit. b) The cyber- layer is modeled in terms of necessary information content (blue dots), with attacks implicating a propagative disruption this information (red dots).

Since the focus of this work is on evaluating cyber threats to the air traffic system, the cyber- system which is used for air traffic control and management is explicitly modeled in this work. The modern air traffic system uses numerous networked cyber assets for traffic control and management. Relevant to this effort, control of aircraft for collision avoidance is undertaken by human controllers located in about 20 regional offices, known as Air Route Traffic Control Centers (ARTCCs or Centers), which are each responsible for a partition of the United States' airspace. At each Center, a small group of controllers (typically 3-5) are assigned to each Sector in the Center's airspace, and are responsible for the control of aircraft in the Sector. The controllers for each Center rely on a number of cyber- systems, including radar displays of aircraft and weather, collision alert tools, and computer systems that provide directives from traffic managers.   In similar fashion, controllers for the Terminal Radar Approach Control facilities (TRACONs) associated with major airports, as well as airport-control tower personnel, have numerous cyber tools which provide radar data, filed flight plans, and relevant weather data. Meanwhile, wider-area and longer-term traffic management is undertaken via coordination of traffic managers at the regional offices, the central command center (Air Traffic Control Strategic Command Center or ATCSCC), and major commercial airlines. The personnel involved in traffic management also use numerous cyber tools, including weather and traffic data sources, telephone as well as web-based communication, simulators, etc.

Holistically, the cyber system acting in support of the air traffic system can be viewed as transmitting the information that is necessary for effective traffic management and control, see Figure 2a. This cyber

system comprises a mixture of specialized information transfers for the air traffic system and generic information gathering from the broader Internet (e.g., public-domain weather forecasts). The cyber systems used by traffic managers are very often networked to the broader web, whether for required data transmission or for convenience. To the best of our knowledge, cyber systems used in the airspace system use only standard protection technologies (e.g., standard firewalls and virus-checking software, limited or no encryption). The cyber system may be subject both to failures and to deliberate software and hardware attacks, and indeed both types of threats have been observed [1-3].
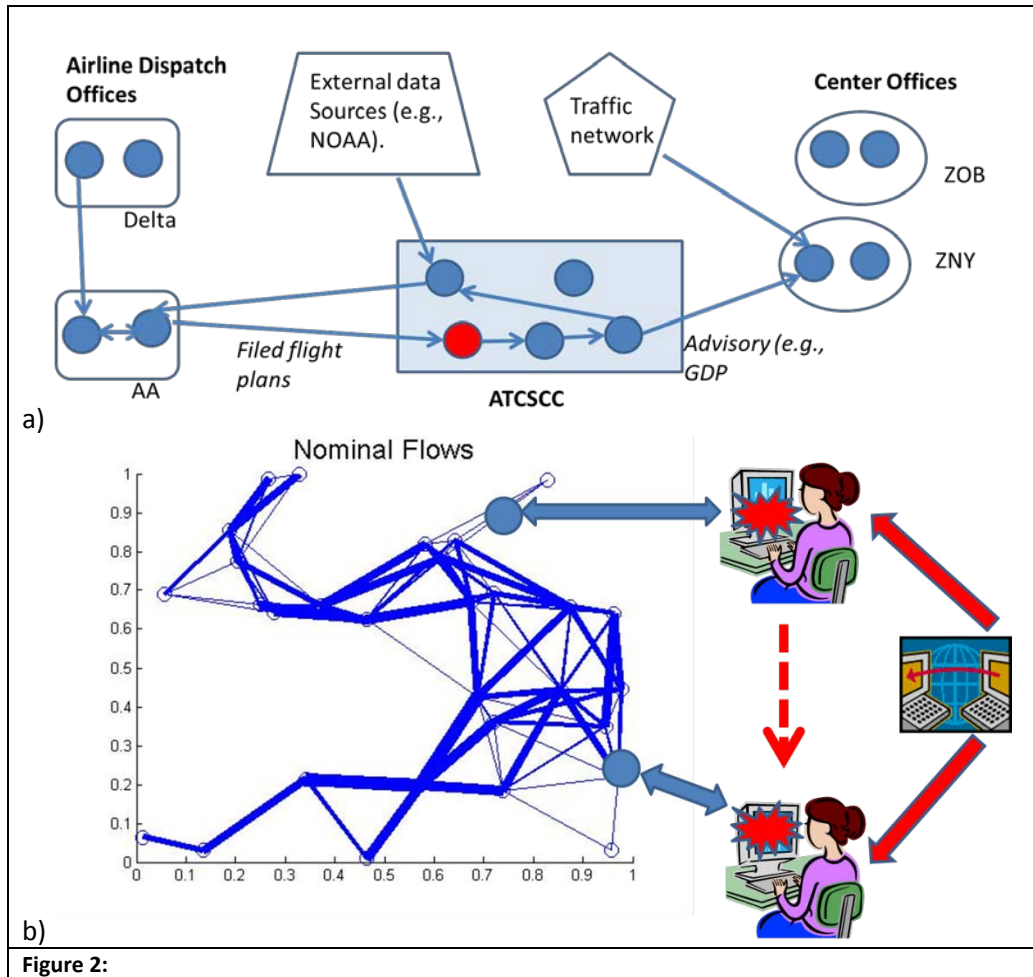
Figure 2:

In this work, we abstractly model the cyber system as a network of information resources that are necessary for control and management of traffic, see Figure 2b. Under nominal conditions, each piece of information is modeled as being present, which then allows control and management. The main purpose of our cyber-layer model is to represent disruptions to the needed information resources (which are the nodes in our network model). These disruptions then cause changes to the traffic network, which are modeled as the interface between the cyber and traffic layers (also see Figure 2b).

Formally, an information network with m nodes labelled $i = 1, \ldots, m$ is considered, which each node represents an information resource needed for traffic management and control (e.g., a radar screen used by the controllers for a particular Sector, or the flight manifest data that are needed by a Center's traffic managers, etc). Each node is modeled has having a nominal state `Normal' or `N', which indicates

that the information content is available and uncorrupted.   During a particular operational period of interest, each node may transition to a failed state (`Failure' or `F') which indicates that the information content associated with the node is unavailable, whether due to a failure or an attack. The state of node i during the period of interest is referred to as $x(i)$.

Two probabilistic models for failure are considered. In the simpler model, an attack or failure event is modeled as causing the state $x(i)$ of each network node to be `F' with probability $p(i)$, independently of all other nodes. This simple model for failures is descriptive of independent component failures, which cause individual information resources to become unavailable. The model also encompasses structured deterministic failure scenarios where the failure of a fixed set of information resources needs to be evaluated (e.g., during post-processing after a failure or event, or for common failure paradigms). The model further captures certain types of cyber-attacks, for which information resources are independently impacted. For instance, phishing attacks wherein an attacker sends an e-mail with a computer-virus file attached to a long list of recipients may be modeled in this way. In this scenario, personnel who are responsible for traffic control (e.g., controllers, managers, airline's dispatchers) each have some probability of independently receiving and opening the attack e-mail using a particular cyber system, hence causing failure of this cyber system for a period of time. Thus, a model where each cyber resource is independently disrupted with some probability is apt.

A second, more sophisticated model for cyber disruptions is also considered. This second envisioned model reflects that information flows among resources according to a specified network, and hence disruptions of information flow may be correlated. Specifically, the model captures that information disruptions may propagate through the cyber network. This type of propagative disruption in cyber systems has been studied widely, in the context of computer-virus spread, cascading failure modeling, and other contexts [38-40]. Numerous probabilistic models for propagation or spread have been proposed. Here, a stochastic percolation model for disruption propagation is considered. Specifically, first each node i in the network is modeled as probabilistically being infected (having a failed status) with some probability, say $p_0(i)$; this is the initial stage (stage *k=0)* of the infection. In further stages of the infection, each node that has just been infected has some probability of infecting further nodes. Specifically, at stage *k*, each node i that was infected at stage *k-1* infects any neighboring node *j* with probability $p_k(j,i)$, where the neighbors of a node are specified by the digraph Γ. The infection process continues until no new infections are produced. We notice that the percolation model generalizes the simple probabilistic-failure model, by capturing cascading impacts of failures in the information-flow network. The percolation model is useful when the failure of one cyber system implicates an impact on other information resources used in traffic control and management: for instance, the failure of systems which store flight manifests may simultaneously impact information resources at multiple Centers.

The cyber- layer of the model is interfaced with the traffic layer as follows. Each information resource is viewed as being necessary for operation of some airspace resources over a time period of interest– for instance, a major flow or jet route, a sector or group of sectors, or one airline's traffic.   Thus, information-resource failures modulate the associated traffic resources' parameters for their nominal values. Specifically, airspace resources such as Sector or flow capacities may be curtailed, demand patterns may be altered, traffic management initiatives parameters (e.g., rates, scope) may be modified, etc. Thus, the traffic models parameters and inputs are changed over an interval in reflection of the information-layer failures. More sophisticated interfaces between the cyber and traffic layers are also envisioned, which can capture deliberate attacks by sophisticated sentient adversaries. Specifically, to capture such attacks, we model failure of an information resource as permitting commandeering the information set by the adversary. In this case, the failure can allow for an arbitrary time-varying parametric or actuation signal that is set by the adversary.

In sum, a layered network model has been proposed, which includes: 1) a flow and queueing model for traffic and its management; 2) an influence model-based weather simulator; and 3) a percolation model for disruptions to information resources. The different layers are coupled in structured ways. Specifically, the weather and cyber layers both modulate parameters in the traffic layer.

3. Linearized Approximation and Control-Theory Perspective on Threat Assessment

To assist in threat assessment, a linearized approximation of the layered model is also considered, which preserves the network's topological structure but simplifies modeling of the flow processes and interfaces. The linearization of the traffic layer is developed using the techniques described in [10,16,31], which consider a stochastic linearization around a predicted operating conditions and achieve an equivalence with a linear circuit model. The linearization of the queueing network model represents a significant abstraction from the detailed operations of the air transportation system, but preliminary studies have shown that the linearized model can adequately capture wide-area ripples in traffic flow, and represent essential connections between the network's topology and disruption impacts [16]. Linear and jump-linear approximations of layered weather and traffic models have also been proposed, see e.g. [10,41]. These analyses draw on the moment-closure properties of the influence model [31] along with the stochastic linearization approaches used for the traffic model. In similar fashion, linear approximations for probabilistic infection models, such as those used for the cyber layer, are available in the literature [38-40] and can be used in our framework. The interfaces between the cyber layer and physical layer can similarly be approximated as linear using the arguments given in [31], or alternately cyber events can be modeled as actuating or modifying the traffic network's topology.

The linearized model is useful for threat assessment for several reasons. First, it enables rapid simulation and formal statistical analysis of wide-area impacts of disruptions, and hence potentially also simplify design of mitigation schemes. Second, the model strips away operational details and highlight the essential connection between the network's graph topology and disruption impacts, which can allow development of simple graph-theoretic rubrics for analysis and design. Third, because the linear model permits explicit mathematical expression in state-space form, it allows concrete description of the threat assessment problem in control-theory and (specifically) network-control-theory language. We develop this control-theoretic interpretation next.

We argue, specifically, that deliberate manipulation of the cyber network by a deliberate sentient adversary can be viewed as a *reachability* problem for the linearized model, while tolerance of natural disruptions or blunt attacks may be phrased as a *robustness* analysis. These reachability and robustness formulations can be addressed using a mixture of simulation and formal analysis, and also are a useful stepping stone for obtain topological insights into attacks. Formally, the linearized model with the traffic and cyber layer can be written as follows:

$$\begin{bmatrix} x_t[k+1] \\ x_c[k+1] \end{bmatrix} = \begin{bmatrix} G_{tt}(\Gamma_t) & G_{ct}(\Gamma_{ct}, u[k]) \\ G_{tc}(\Gamma_t) & G_{cc}(\Gamma_c, u[k]) \end{bmatrix} \begin{bmatrix} x_t[k] \\ x_c[k] \end{bmatrix} + \begin{bmatrix} B_t \\ B_c \end{bmatrix} u[k],$$ where $x_t[k]$ and $x_c[k]$ are states of the traffic network (flows, congestion levels) and cyber network (information resource failure statuses or probabilities) respectively, Γ specifies the networks' graph (including inter-network interactions and interfaces between the networks), and *u[k]* captures an attack or disruption initiation which may couple into the cyber or traffic dynamics additively via B_c or B_t, or may modify the network's graph. We note here that the weather layer has been suppressed to simplify the presentation, but can also straightforwardly be included using an extended state vector, if desired. For sophisticated cyber-attacks,

the reachability question of interest is to understand whether and how the input u*[k]* can be designed to manipulate the physical system's state away from its nominal trajectory. Meanwhile, for blunt attacks and natural disruptions, the robustness of the dynamics to stochastic or impulsive disruptions is of interest. Formal algebraic analyses of these problems can be completed by applying and building on control-theory techniques. These analyses then give a starting point for developing topological insights into threat impact, using recent results that tie network structure to reachability and robustness levels [17-19]. They also bring forth interesting new questions regarding reachability and robustness in networks with layered structures, as well as reachability using sign-definite inputs. We envision that these reachability and robustness analyses will also enable identification of critical vulnerabilities and eventually design of protection schemes, based on actuator-placement techniques in the controls literature.

A comprehensive general treatment of the posed reachability and robustness problems is beyond the scope of this paper. However, we pursue an exploratory analysis of one threat-assessment problem that can be approached using the modeling framework. Specifically, we consider the robustness or vulnerability assessment for cyber attacks/failures that cause restriction of traffic flows (local network topology changes), as measured by the overall impact on traffic that they implicate.

4. Threat Assessment: Exploratory Flow-Vulnerability Analysis

The proposed layered network model and its linearization are appealing for threat assessment because they permit evaluation of the propagative impacts of weather and cyber threats on traffic. Because the models represent weather, traffic, and cyber systems, they can support a variety of analyses on the potential impacts of cyber-attacks and failures as well as the interactions among technological and natural disruptions. Crucially, the model allows evaluation of the wide-area impact on air traffic of local disruptions. Here, we discuss one exploratory analysis using the model, focused on 1) identifying vulnerable traffic flows and 2) using the vulnerability analysis to gauge the impact profiles of cyber-attacks and to determine overall network robustness. This analysis is meant to illustrate how the nonlinear and linear modeling frameworks can be used in tandem with network-controls concepts for threat assessment.

The exploratory flow-vulnerability analysis is undertaken as follows. First, simulations of the linearized traffic-network model for constructed examples are used to conceptualize which flows are vulnerable (Section 4.1). Based on these simulations, a network-theoretic calculation of flow vulnerability is proposed and evaluated (Section 4.2). Using this calculation, a network-theoretic metric for overall robustness to cyber-attacks is proposed, and the metric is tested using the linearized layered model (Section 4.3). Finally, the developed network-theoretic calculations and metrics are applied to a realistic small-scale case study, and the results are tested using the detailed (nonlinear) layered model (Section 4.4).

4.1. What-If Analysis of Vulnerability: Simulations and Concepts

As a first step toward assessing the vulnerability of flows, simulations of the linearized traffic network model are undertaken to gain insight into the propagative impacts of disruptions. These simulations may be viewed as a ``what-if" analysis of the attack impact. That is, we seek to understand the

consequences of a significant-duration stoppage or constriction of different traffic flows, whatever the cause of the disruption. Since a what-if analysis is conducted, the cyber and weather layers of the model are suppressed for the simulations in this subsection, and only the propagative impact in the traffic layer is considered. The linearized model is used for the simulations in this section, so as to provide a pathway toward developing graph-theoretic and analytical insights into vulnerability.


Figure 3:

The linearized flow model is simulated on a constructed network with 30 waypoints, see Figure 3. The flow network was constructed by placing 30 vertices in the unit square which correspond to the 30 waypoints, and allowing flows between waypoints whose corresponding vertices are sufficiently close. Traffic to three destination airports from 10 origin points (which may represent either origin airports or points at which flows enter from outside the modeled region) is considered. The linearized flow model, which is analogous to an RC network model, is simulated for the 30-waypoint network. The nominal flow densities on the links are shown in Figure 3.

The spatial impacts on network-wide traffic of two individual flow disruptions are shown in Figure 4. Specifically, Figure 4a shows the changes in steady-state flow densities due to blockage of a particular flow, while Figure 4b shows the transient response on a nearby flow. Similarly, Figures 4c and 4d show the steady-state and transient impacts due to the blockage of another flow. The simulations indicate that the most drastic steady-state changes occur on flows that are proximal to impacted flow, and

Figure 4:

particularly on routes that are alternatives of the blocked flow. The transient responses indicate a traveling-wave phenomenon, wherein alternative routes are quickly impacted, immediate downstream
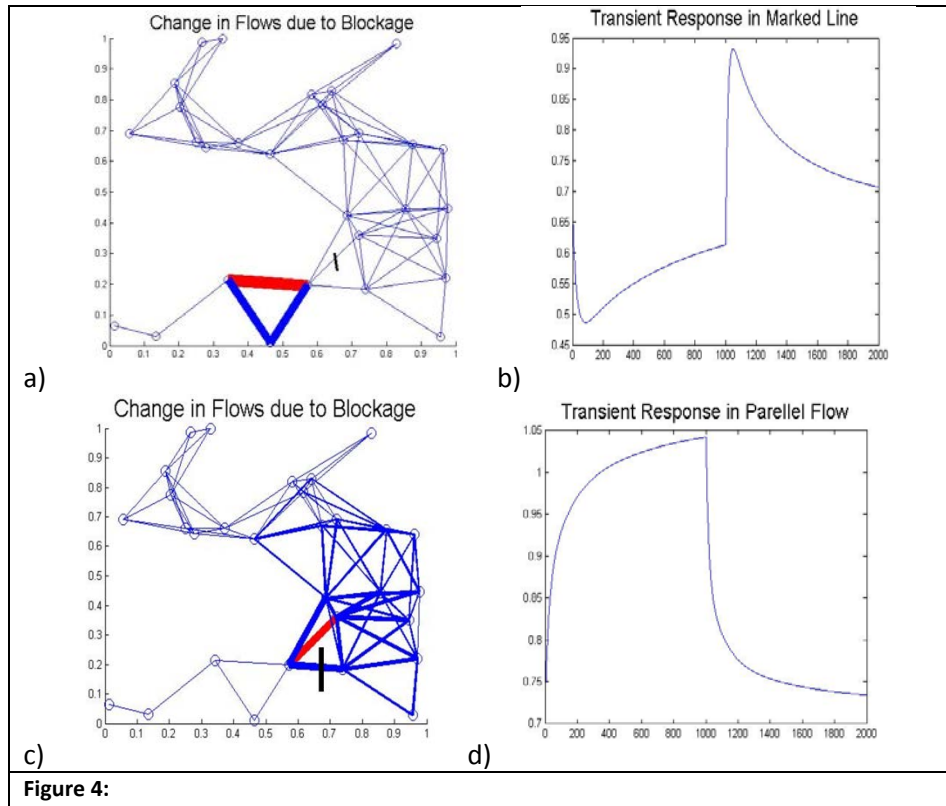
flows show a fast bimodal response (i.e., decrease followed by increase), and locations further away have a more limited and slower transient.

In Figure 5, the total impact of blocking each flow in a persistent way is compared. The impact of blocking a flow is measured as follows. For each other flow, the absolute deviation between the nominal and modified flow value is determined. These deviations are summed as a total measure of impact. Figure 5 shows that blocking a particular flow has large impact if: 1) the nominal traffic on the flow is large, and 2) there are few alternative routes for the flow traffic (equivalently, the flow is on a weak cut of the network's graph). Thus, the most vulnerable flows are large flows which offer few alternative routes.


Figure 5:

The comparison of blockage impact is shown for a second constructed network in Figure 6. Again, large flows which have few alternative routes are seen to cause disproportionate impact to network-wide traffic, or equivalently to be the most vulnerable. This example shows that the vulnerability may be particularly large if the alternative routes are long. It also highlights that routes on weak cuts tend to have large nominal flows since few alternatives are available.
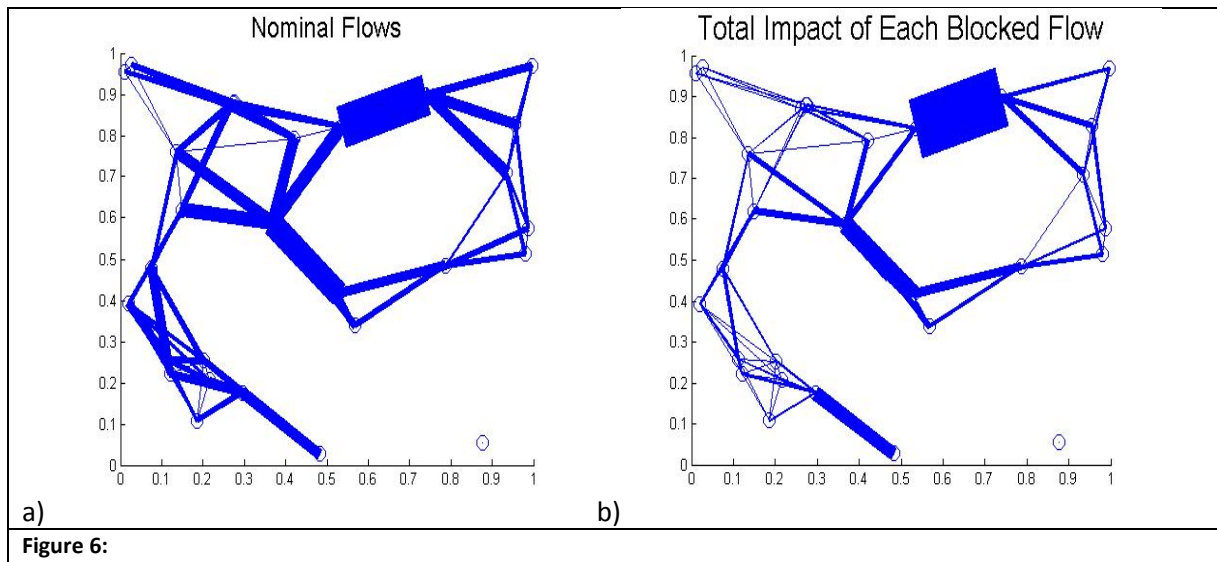

a)   b)

Figure 6:

4.2. Network-Theoretic Calculation of Flow Vulnerability

The insights on flow vulnerability developed in the previous subsection, together with network analysis/control concepts, suggest a simple graph-theoretic algorithm for identifying flows that are especially vulnerable to unknown disruptions like cyber attacks. The proposed algorithm draws on the *Laplacian matrix* of the network's graph, which has been widely used for network analysis and control. Specifically, the algorithm uses the fact that the eigenvector associated with the subdominant

eigenvalue of the Laplacian matrix can be used to measure whether each edge in a network is on a weak cut or not. By combining this information with nominal flow data, a good approximation for the vulnerability of each flow can be obtained, which turns out to be especially apt in finding highly-vulnerable flows.

Here is the algorithm:

1) Using the layered traffic and weather model (where a deterministic or average weather forecast is assumed), compute the nominal traffic flow between each pair of waypoints (i.e., on each route segment or link). Let us denote the flow by between waypoints i and j by f_{ij}.

2) Construct the *Laplacian matrix L* for the traffic network. The Laplacian L is a square matrix, and for us has dimension equal to the number of waypoints in the traffic network. The entries in L are filled in as follows. The entry at row i and column j is set to *-1* if there is a flow (link) between waypoint i and waypoint j under the known operational conditions, and is set to zero otherwise. The diagonal entries are chosen so that each row sums to *0*. We notice that the Laplacian exactly encapsulates the topolog o the flow network. Alternately, if a linearization of the traffic model has been done, the state matrix of the linearized traffic model can be used in lieu of the Laplacian matrix. It is easy to check that the state matrix will have the form of a Laplacian or *grounded Laplacian* matrix.

3) Provided that the network topology is connected, the Laplacian matrix has a single eigenvalue at *0*, while the remaining eigenvalues are real and strictly positive. The next step of the algorithm is to find the smallest nonzero eigenvalue and corresponding eigenvector v.

4) For each pair of waypoints i and j, please find $C_{ij} = |f_{ij}(v_i - v_j)|$.

5) The scalar C_{ij} is a measure of impact extent due to a blockage of the flow between waypoints i and j, or equivalently is a measure of the vulnerability of the flow (link). Thus, to identify vulnerable flows, we rank the flows according to C_{ij}, and identify those that are above a threshold or largest.

a)   b)

Figure 7:

The algorithm for identifying vulnerable flows has been applied to the traffic-network examples introduced in the previous section. The algorithm is effective in finding the most vulnerable flows. For instance, the vulnerability metric C_{ij} and actual total flow disruption are compared in Figure 7, for the

second example network. The metric identifies the four worst vulnerabilities, in order of impact. Also, other high vulnerability flows are highlighted by the metric. Likewise, the three most vulnerable flows are found, in order of impact, for the first example network (not shown).

For the linearized flow model, a formal analysis can be used to show that the vulnerability metric effectively identifies vulnerable flows, particularly in the case where the smallest nonzero eigenvalue is significantly smaller than the remaining eigenvalues. This spectral gap is present, for instance, for many networks with planar graphs such as the typical graph of the air traffic network. A detailed presentation of the formal analysis is left to future work.

4.3. Global Vulnerability Analysis for Cyber Events

The flow-vulnerability analysis developed in the previous two subsections enables assessment of the global susceptibility of the traffic system to cyber- events which impact flows. Here, the flow vulnerability analysis is used to define global metrics for cyber- event vulnerability, and simulations are used to evaluate the metrics and also to assess the extent of impact of cyber- attacks, for the two examples.

The two probabilistic models for cyber events described in Section 2 are considered here. Whichever cyber model is used, information disruptions in the cyber system are assumed to block flows in the traffic network. Specifically, each information resource in the failed state is assumed to cause the blockage of one or a group of flows in the traffic system. For this model, each flow in the traffic network has a probability of blockage due to cyber events, which can be computed via formal analysis, approximation, or simulation of the percolation model. We use the notation q_{ij} for the probability of blockage of the flow between waypoints *i* and *j*.

A metric for the global vulnerability to cyber events can be defined by considering blockage probabilities and impacts. A natural metric is $V = \sum_{i,j} q_{ij} C_{ij} = \sum_{i,j} q_{ij} C_{ij} = \sum_{i,j} q_{ij} |f_{ij}(v_i - v_j)|$, which weights the flow-vulnerabilities by the blockage probabilities and then sums them. Alternately, if a probabilistic description of cyber events is unavailable or all flow blockages are equally like, then the metric $U = \sum_{i,j} C_{ij}$ is natural.

A specific cyber- disruption model has been simulated for the two example traffic networks. For the simulation, a single information resource is assumed to be associated with each flow. Each information resource is assumed to be failed with equal probability and independently during the period of interest (due to a phishing attack, for instance). Failure of the information resource is assumed to lead to persistent blockage of the associated flow. We note that the cyber-layer model is overly simplistic: in practice, information resource failures would likely disrupt a group of flows (e.g., all of the flows within a Sector). However, even the simple independent-failure model yields complex impacts on traffic flow patterns, which allow a basic evaluation of the performance metrics.

The global vulnerability metrics for the two networks have been computed. The metric value for Example 2 is larger than that for Example 1 by a factor of about 1.5. To further evaluate cyber-event impacts, Monte Carlo simulations of the layered traffic and cyber network models are undertaken, for two different information-resource failure probabilities (corresponding to ``more virulent" and ``less virulent" attacks). Histograms of the total flow disruption, as a fraction of the nominal total flow, are shown (see Figure 8). Indeed, the mean impact is significantly larger for the second example as compared to the first, primarily because of the significant probability for very impactful disruptions.

Interestingly, the shape of the histogram changes significantly between the less virulent and more virulent attacks and between the two examples, suggesting that the occurrence of larger but rarer events is a key concern in evaluating security to cyber- attacks.

Figure 8:
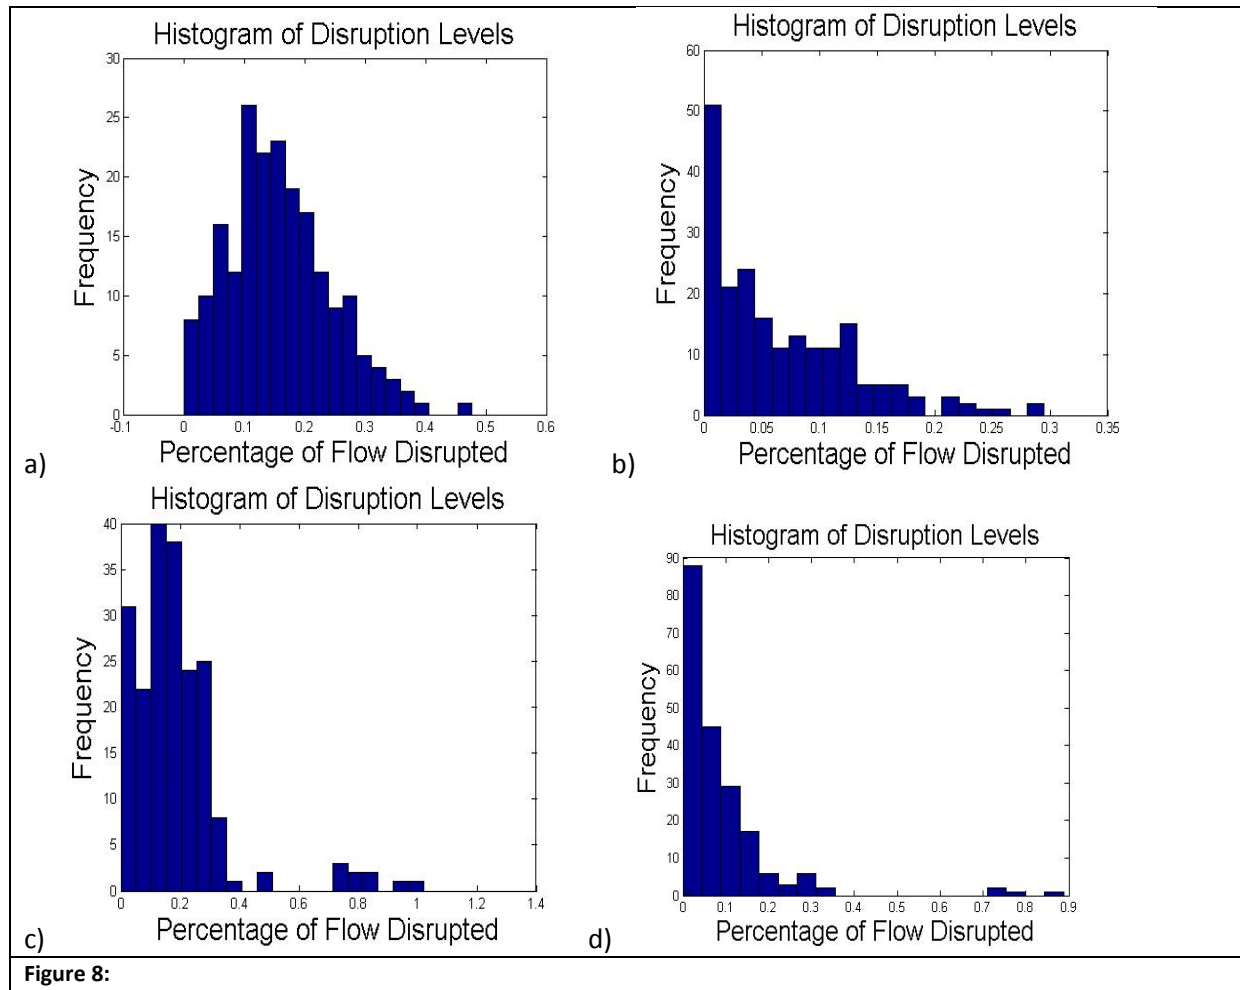
The simplicity of the flow-vulnerability calculation supports design of cyber- management strategies to mitigate event impacts. As one simple example, additional protection of cyber systems can be enabled for the information resources coupled with vulnerable flows. For the second example network, we have considered the cyber vulnerability when protection resources are provided for the most vulnerable flow, so that the probability of blockage of this flow becomes negligible. Monte Carlo simulations have also been completed for the protected network. The resulting histogram of impacts is shown in Figure 9. The additional protection eliminates rare very high impact events, while other characteristics of the impact profile remain similar.

4.4. Realistic Small-Scale Case Study

A preliminary application of the proposed vulnerability assessment to a more realistic case study is described. This case study is adapted from a case study described in [7], which is concerned impact of

long-duration non-weather disruptions on air traffic in a Center's airspace. The case study in [7] itself extends a study of weather impact on air traffic using the queueing- model for traffic [11]. The case study is roughly inspired by high-altitude traffic characteristics in Atlanta Center (ZTL), however the case study is constructed and much distanced from real operations in ZTL.
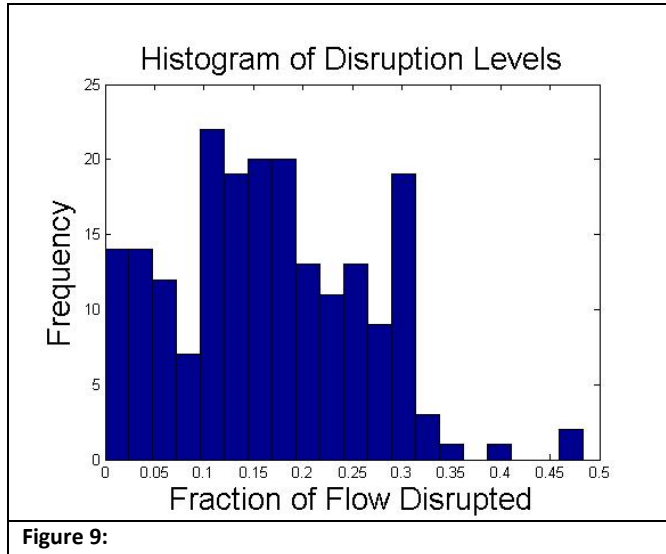

Figure 9:

The case study is concerned with traffic for four origin-designation pairs that traverses a region with six high-altitude Sectors. The demand densities, routing fractions, and capacities are illustrated in Figure 10. Limited traffic management initiatives are also nominally in place in the example, see [11] for details. The detailed queueing-network model has been simulated for the case study, to obtain flow dynamics, backlog and delay statistics, and other performance metrics. Under nominal conditions, a total delay of about 330 minutes is imposed on all aircraft using the network; thus, the airspace is relatively congested under nominal conditions, to the point of causing minor delays.

The eigenvector flow-vulnerability analysis algorithm has been applied to the case study network. The algorithm identifies flows 10-14, 1-4, and 2-10 as the most vulnerable. Conceptually, the flows identified as most vulnerable are reasonable choices. In each case, blockage of the flow requires significant re-routing of traffic via sparse alternate choices. Additionally, these are large flows whose disruption may cause significant impact. Simulations of the queueing model verify that, indeed, blockage of these flows cause large impact in terms of increased delay and backlog, and significant modification of flow patterns. !!!Need to expand on last sentence!!!
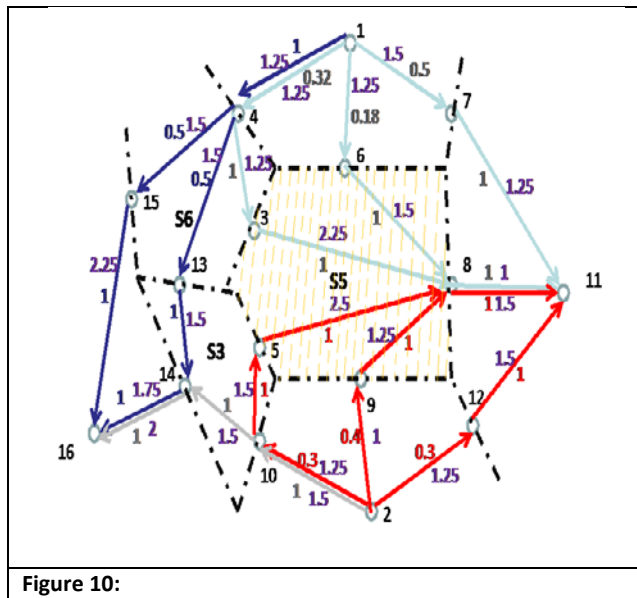

Figure 10:

References

[1] S. Gallagher, ``Computer systems outage grounded American Airlines at major hubs," http://arstechnica.com/information-technology/2015/09/computer-systems-outage-grounds-american-airlines-at-major-hubs/ .

[2] "Computer glitch delays hundreds of flights in New York, DC," https://www.rt.com/usa/312554-flights-delay-faa-glitch/ .

[3] ``Computer glitch caused delays at NY's Kennedy Airport," Associated Press, http://www.journalnow.com/news/nation_world_ap/computer-glitch-caused-delays-at-ny-s-kennedy-airport/article_5741a109-85e4-5ad3-a2ed-72987167e3f9.html

[4] M. Bishop, "What is computer security?." *Security & Privacy, IEEE* 1.1 (2003): 67-69.

[5] R. Baheti and H. Gill, "Cyber-physical systems." *The Impact of Control Technology* 12 (2011): 161-166.

[6] R. S. Huang, H. M. Yang, and H. G. Wu. "Enabling Confidentiality for ADS-B Broadcast Messages Based on Format-Preserving Encryption." *Applied Mechanics and Materials*, vol. 543, pp. 2032-2035. 2014.

[7] S. Roy, Y. Wan, and J. Xie, "Proactive and reactive management of non-weather capacity disruption events in the National Airspace System: a flow modeling and design approach," in *Proceedings of the 15th AIAA Aviation Technology, Integration, and Operations Conference*, *AIAA Aviation* Dallas TX, June 2015.

[8] M. C. Aubert et al, "Toward the development of a low-altitude air traffic control paradigm for networks of small, autonomous unmanned aerial vehicles," *Aviation SciTech,* Jan. 2015.

[9] T. J. Colvin, and J. J. Alonso, "Near-elimination of airspace disruption from commercial space traffic using compact envelopes." *AIAA SPACE 2015 Conference and Exposition*, 2015.

[10] R. Dhal and S. Roy, "Layered moment-linear network models as tools for strategic air traffic flow management." *Proceedings of the 2012 AIAA Guidance, Navigation, and Control Conference*, June 2012.

[11] Y. Wan et al, "Dynamic queuing network model for flow contingency management." *Intelligent Transportation Systems, IEEE Transactions on* 14.3 (2013): 1380-1392.

[12] M. Xue et al, "Using stochastic, dynamic weather-impact models in strategic traffic flow management." *Proceedings of 91st American Meteorological Society Annual Meeting*, Seattle, WA, Dec. 2011.

[13] S. Roy, B. Sridhar, and G. C. Verghese. "An aggregate dynamic stochastic model for air traffic control." In *Proceedings of the 5th USA/Europe ATM 2003 R&D Seminar, Budapest, Hungary*. 2003.

[14] D. Sun et al, "Eulerian trilogy." *AIAA Guidance, Navigation, and Control Conference and Exhibit*, 2006.

[15] P.-Y. Chen, S.-M. Cheng, and K.-C. Chen. "Smart attacks in smart grid communication networks." *Communications Magazine, IEEE* 50, no. 8 (2012): 24-29.

[16] S. Roy and Y. Wan, "Geographical weather-impact sourcing: analytical and data-driven approaches," in *Proceedings of AIAA SciTech,* 2014.

[17] R. Dhal and S. Roy, "Vulnerability of continuous-time network synchronization processes: A minimum energy perspective." *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on.* IEEE, 2013.

[18] J. Abad Torres et al. "Local open-and closed-loop manipulation of multi-agent networks." *Proceedings of the ACM/IEEE Sixth International Conference on Cyber-Physical Systems*. ACM, 2015.

[19] F. Pasqualetti, S. Zampieri, and F. Bullo. "Controllability metrics, limitations and algorithms for complex networks." *Control of Network Systems, IEEE Transactions on* 1.1 (2014): 40-52.

[20] Wan, Yan, and Sandip Roy. "A scalable methodology for evaluating and designing coordinated air-traffic flow management strategies under uncertainty." *Intelligent Transportation Systems, IEEE Transactions on* 9.4 (2008): 644-656.

[21] Wanke, Craig, et al. "Modeling air traffic demand for a real-time queuing network model of the national airspace system." *AIAA Modeling, Simulation and Technologies Conference, Minneapolis, MN*. 2012.

[22] Wanke, Craig, and Christine Taylor. "Exploring Design Trade-offs for Strategic Flow Planning." *AIAA Aviation* (2013): 12-14.

[23] Taylor, Christine, et al. "Designing Traffic Flow Management Strategies Under Uncertainty." *FAA/Eurocontrol Air Traffic Management Research and Development Forum and Exhibit* (2015).

[24] B. Sridhar and P. Kopardekar, ``Toward autonomous aviation operations: what can we learn from other areas of automation,'' to appear in *2016 AIAA Aviation Forum, Washington DC* (2016).

[25] Martinez, Stephane, et al. "A weighted-graph approach for dynamic airspace configuration." *Proceedings of the AIAA Conference on Guidance, Navigation, and Control (GNC). American Institute of Aeronautics and Astronautics* (2007).

[26] K. Gopalkrishnan, H. Balakrishnan, and R. Jordan, ``Clusters and communities in air traffic delay networks,'' to appear in the *2016 American Control Conference, Boston, MA* (2016).

[27] Dhal, Rahul, and Sandip Roy. "Vulnerability of continuous-time network synchronization processes: A minimum energy perspective." *Decision and Control (CDC), 2013 IEEE 52nd Annual Conference on*. IEEE, 2013.

[28] Sundaram, Shreyas, and Christoforos N. Hadjicostis. "Distributed function calculation via linear iterative strategies in the presence of malicious agents." *Automatic Control, IEEE Transactions on* 56.7 (2011): 1495-1508.

[29] Roy, Sandip, Mengran Xue, and Sajal K. Das. "Security and discoverability of spread dynamics in cyber-physical networks." *Parallel and Distributed Systems, IEEE Transactions on* 23.9 (2012): 1694-1707.

[30] Steiner, Matthias, and J. Krozel. "Translation of ensemble-based weather forecasts into probabilistic air traffic capacity impact." *Digital Avionics Systems Conference, 2009. DASC'09. IEEE/AIAA 28th*. IEEE, 2009.

[31] Asavathiratham, Chalee, et al. "The influence model." *Control Systems, IEEE* 21.6 (2001): 52-64.

[32] Xue, Mengran, et al. "Using stochastic, dynamic weather-impact models in strategic traffic flow management." *Proceedings of 91st American Meteorological Society Annual Meeting*. 2011.

[33] Roy, Sandip, et al. "A stochastic network model for uncertain spatiotemporal weather impact at the strategic time horizon." *Proceedings of AIAA Aviation Technology, Integration, and Operations Conference*. 2010.

[34] Xue, Mengran, et al. "Refinement and Enhancement of an Influence-Model-based Weather-Impact Simulator." *Proceedings 2012 AIAA Modeling and Simulation Technologies Conference*. 2012.

[35] Dhal, Rahul, et al. "An Operations-Structured Model for Strategic Prediction of Airport Arrival Rate and Departure Rate Futures." *2014 Aviation Technology, Integration, and Operations Conference*. 2014.

[36] Buxi, Gurkaran, and Mark Hansen. "Generating probabilistic capacity profiles from weather forecast: A design-of-experiment approach." *Proc. of USA/Europe Air Traffic Management Research & Development Seminar*. 2011.

[37] Ramanujam, Varun, and Hamsa Balakrishnan. "Estimation of maximum-likelihood discrete-choice models of the runway configuration selection process." *American Control Conference (ACC), 2011*. IEEE, 2011.

[38] Cohen, Reuven, Shlomo Havlin, and Daniel Ben-Avraham. "Efficient immunization strategies for computer networks and populations." *Physical review letters* 91.24 (2003): 247901.

[39] Wan, Yan, Sandip Roy, and Ali Saberi. "Designing spatially heterogeneous strategies for control of virus spread." *Systems Biology, IET* 2.4 (2008): 184-201.

[40] Saito, Kazumi, Ryohei Nakano, and Masahiro Kimura. "Prediction of information diffusion probabilities for independent cascade model." *Knowledge-based intelligent information and engineering systems*. Springer Berlin Heidelberg, 2008.

[41] Xie, Junfei, et al. "A Jump-Linear Model Based Sensitivity Study for Optimal Air Traffic Flow Management under Weather Uncertainty."