A spectral approach towards analyzing air traffic network disruptions

Thirteenth USA/Europe Air Traffic Management Research and Development Seminar (ATM2019)

A Spectral Approach Towards Analyzing Air Traffic Network Disruptions

Max Z. Li, Karthik Gopalakrishnan, Hamsa Balakrishnan
Department of Aeronautics and Astronautics
Massachusetts Institute of Technology
Cambridge, MA, USA
{maxli, karthikg, hamsa}@mit.edu

Kristyn Pantoja
Department of Statistics
Texas A&M University, College Station, TX, USA
kristynp@stat.tamu.edu

***Abstract*—The networked nature of the air transportation system leads to systemwide delays and cancellations as a result of disruptions at an airport. A comprehensive analysis of system performance requires understanding the inherent interdependencies between various airports, in order to characterize off-nominal disruptions and to aid in recovery. In this work, we apply Graph Signal Processing (GSP) techniques to the analysis of flight delay networks, yielding two novel contributions: (1) We use the notion of the total variation (TV) of a graph signal in order to identify and quantify unexpected distributions of delays across airports; and (2) we present a spectral eigendecomposition analysis of airport disruption and delay networks. We investigate and characterize different patterns of delay distribution based on the relationship between TV and total delay, using 10 years worth of operational data from major US airports. We show that attributes of the resultant eigenvector modes and energy contributions are useful metrics to characterize specific disruptions caused by events such as nor'easters, Atlantic hurricanes, and equipment outages at airports.**

Keywords: Graph Signal Processing; graph Fourier analysis; spectral methods; weather disruptions; airport operations; flight delays

I. Introduction

The scale of air transport operations conducted within the US National Airspace (NAS), along with the interconnectedness inherent within the NAS contribute to the challenges in characterizing the resilience and predictability of the system [1]. Obtaining a better understanding of these attributes are of particular importance in the context of disruptive events that affect the NAS, leading to systemic penalties such as congestions, delays, and flight cancellations. The importance of this line of research has been highlighted both domestically as well as internationally; the Federal Aviation Administration (FAA) and EUROCONTROL have identified the need for increased resiliency (how tolerant, or robust, airport-airspace operations are to operational perturbations) and predictability (reduction in variability and uncertainty) within the air transportation system.

Weather disruptions were responsible for 53% of all delays experienced in the NAS in 2017 [2]. One particular meteorological phenomenon we examine in this work is the *nor'easter*, a type of winter storm that predominantly affect the Mid-Atlantic and New England region of the US. Nor'easters significantly disrupt air traffic operations; the nor'easter that occurred 22-24 January 2016 resulted in the cancellation of over 11,000 flights, and was estimated to lead to financial losses of $100 to $120 million USD to the US aviation industry [3]. Airport and airline equipment outages also lead to delays and cancellations. We examine two particular cases involving a power outage at Hartsfield-Jackson Atlanta International Airport (ATL) and a fire incident at the Chicago Air Route Traffic Control Center (ZAU); both resulted in hundreds to thousands of cancelled flights throughout the NAS [4, 5].

Furthermore, since each aircraft typically operates more than one flight per day, it is well-known that both air traffic and crew-induced delays tend to propagate throughout the system (Section II-A). Moreover, combined with the increasing trend of airline hub strengthening at a handful of major airports, the share of traffic – and consequently, delays – is being handled by a smaller subset of increasingly critical and interconnected hub airport nodes. Due to the underlying networked structure of the NAS, any analysis of airport disruptions that do not consider the connectivities and interactions between airports may result in myopic findings. To this end, we propose using *Graph Signal Processing* (GSP) techniques (Section II-B) to analyze the impact and dynamic signatures of different categories of disruptions at major US airports. We present insights into the delay impact of different types of disruptive events (*e.g.* Atlantic hurricanes) and the nature of recovery by analyzing the spectrum of the airport delay signals on the underlying graph network.

A. Motivation: Delays in a network

High delays within the NAS are known to have significant economic impacts. When analyzing system performance involving multiple airports and routes, delay can be measured in several ways. Since the airport is typically the bottleneck of the system and the location where delay statistics are recorded, airport-centric delay metrics are commonly used. For instance,

This work was supported in part by the US National Science Foundation.

some classic metrics to quantify the performance of an airport use the average delay of departing flights, or the total delay of all the departing flights. However, due to the strong underlying network connectivity, some airport delays are highly correlated to other. This correlation leads to patterns of delay propagation that are familiar to airport operators, airlines, and air traffic controllers.

Within the NAS, system-wide disruptions can take on two forms: (1) The obvious problematic scenario when delays are high across all airports; (2) a subtler scenario when delays are high, but distributed only across a limited set of airports. The former can be identified via classical metrics, but the latter is more difficult to discern. There are also important operational implications of the second scenario. Consider an example involving Boston Logan International Airport (BOS) and LaGuardia Airport (LGA) in New York. Historically, delays at BOS and LGA are strongly correlated due to factors such as their geographical proximity and connectedness via traffic flows. This means that when delays are high at one of the two airports, it is likely that the other airport will be in a state of high delays as well. Operationally, this relationship is significant and may warrant air traffic controllers to coordinate Traffic Management Initiatives such as Ground Delay Programs and En Route Separation Programs to alleviate the delays at both BOS and LGA, providing spatial and temporal buffers for recovery. However, if there are high delays at only one of the two airports, it would constitute an unexpected delay pattern, potentially requiring a different approach towards recovery at the network-level.

We would like to define the terms *off-nominal*, *expected*, and *unexpected* in the context of our work. Continuing the example above with BOS and LGA, envision the following two scenarios: (1) BOS and LGA both have 90 minutes of delay; (2) BOS has 10 minutes of delay whereas LGA has 170 minutes. In both scenarios, a singular or a series of *off-nominal* events (*e.g.* pop-up thunderstorms, terminal area congestion) resulted in delays at BOS and LGA. However, the first scenario is an *expected* distribution of delays whereas the second scenario is an *unexpected* one given historically high correlations of delays at BOS and LGA. In short, the magnitude of delays determines whether the system is nominal or off-nominal, whereas the distribution of delays determines whether the system is expected or unexpected. Operationally, the two scenarios can have very different impacts, require different TMIs, and necessitate different airline recovery tactics. This example shows precisely why classical metrics fail in these circumstances. To address this knowledge gap, our work focuses on how graph-theoretic and spectral methods can provide insights to these unexpected and off-nominal events.

A priori knowledge that delays at two airports are *not* strongly correlated can also be important operationally: The lack of delay correlations provide a natural decomposition of network control decisions that may advocate for operationally independent recovery plans for these two airports. An example in this case would be Phoenix Sky Harbor International Airport (PHX) and Detroit Metropolitan Wayne County Airport (DTW), an airport pair that we found to have very low delay correlations. It is important to identify and quantify such historical delay dynamics in order to improve operational response and recovery planning. The main challenge lies in identifying the expected performance with a network perspective, where delays at multiple airports simultaneously affect each other. Another challenge is defining what expected patterns are, keeping in mind that even high delays, if present in a consistent pattern, could be considered expected. Finally, an open question in airline network operations is identifying fundamental differences in delay dynamics caused by one type of event (*e.g.*, a nor'easter) as opposed to another (*e.g.*, an Atlantic hurricane), even when the resultant total delays are the same. Our goal is a data-driven analysis at the network level in order to (1) identify off-nominal days during which unexpected delay dynamics occurred, (2) discuss the operational implications, and (3) demonstrate the value of such an analysis using case studies.

II. Literature Review

A. Delays and disruptions

Due to the structure of air transportation networks, local airport-level disruptions may spread to other airports and persist in the system. The complex dynamics of delays are primarily due to aircraft rotations, wherein a single aircraft flies multiple trips in a given day between airport pairs. Consequently, the delays at two airports are not independent, and depend on the traffic between the airports, airline scheduling and crew rotation policies, geographic proximity, airline networks, and air traffic control procedures. Several approaches, including queueing theory [6], network models [7], simulations [8], and machine learning methods [9], have been employed to understand delay dynamics.

Localized and NAS-wide disruptions cause delays and flight cancellations requiring airlines to adjust their flight schedules, re-assign crews, and accommodate stranded passengers. Typically, off-nominal events in the NAS are identified by looking at delay statistics at airports or on specific routes. Recently, clustering has been proposed as a tool for network-level identification of similar delay patterns [10, 11]. However, these network-level identification methodologies do not distinguish nor decompose the observed delay patterns into components that occur often (and therefore, would be expected to occur), and ones that are more unusual (the unexpected components). This paper addresses this open question.

B. Graph Signal Processing (GSP) and graph Fourier decomposition

A foundational tool in classical signal processing is the *Fourier transform*, an integral transform that decomposes a time-based signal into its components in the frequency domain. The Fourier transform is used in a wide selection of applications ranging from signal de-noising and reconstruction to circuit analysis. Fourier transforms have been previously used in aviation applications, including the modeling and prediction of aircraft trajectories on final approach [12], the projection of

large trajectory data sets into feature spaces that can be easily clustered and compressed [13], and the analysis of airport capacity and delay profiles [14, 15]. This body of prior work also considered the relationship between individual airports in terms of delay propagation [15] . However, the underlying graph structure of the NAS – a critical component contributing to the behavior of delays and inefficiencies – cannot be taken into account by the standard Fourier transform.

Given the natural occurrence of graphs and graph signals in aviation, we use GSP techniques to analyze airport disruptions due to weather and other outages. Our work builds off of previous applications of spectral and graph wavelet analysis on traffic congestion in networks [16, 17]. Continuous and spectral graph wavelet transforms have also been used to study air traffic flows [18, 19], albeit at the resolution of Air Route Traffic Control Centers (ARTCCs) and with limited interpretability of spectral results. We extend these approaches by focusing on origin-destination airports. We focus on delays rather than traffic, and provide theoretical results that map spectral analysis to operational interpretations.

III. Methodology

We first present the graph-theoretic notations and basic definitions, followed by Proposition 1 in which we show the relationship between total variation and the graph spectral decomposition. We then formally define the notions of signal smoothness across a graph, setting the stage for Proposition 2 which relates correlation networks to smoothness and unexpected events. We provide explicit probabilistic bounds on unexpected events given our setup. Propositions 1 and 2 are the core of our theoretical contributions, and enable the interpretability of our results when we apply GSP techniques to airport delay networks. Due to the limitations of space, we omit the proofs of these propositions in this paper; we refer interested readers to [20] for proofs and additional theoretical work on identifying outliers in graph signals.

A graph G is defined by a set of nodes $\mathscr{N}$ containing $|\mathscr{N}| = N$ nodes, and edges $\mathscr{E}$ where an edge $e_{i,j} \in \mathscr{E}$ connects a pair of nodes $(n_i, n_j) \in \mathscr{N} \times \mathscr{N}$. The weights on the edges are represented by an adjacency matrix $A \in \mathbb{R}^{N\times N}$, where the element a_{ij} is the weight on the edge from node n_i to node n_j. We consider only undirected graphs, which implies that $A = A^\intercal$. The *degree* of a node n_i is given by

$$\deg(n_i) = \sum_{j=1}^{N} a_{ij}. \tag{1}$$

The *degree matrix* D is a diagonal matrix with elements $d_{ii} = \deg(i)$. Given a graph G along with its adjacency matrix A and degree matrix D, we can define the *combinatorial graph Laplacian*, $L = D - A$. The graph Laplacian L is a real symmetric matrix with a full set of orthonormal eigenvectors. We denote the Laplacian eigenvectors by $v_i \in \mathbb{R}^{N\times 1}$ for $i \in \{1,...,N\}$. By orthogonality, we also have that $v_i^\intercal v_j = \delta_{ij}$, where $\delta_{ij} = 1$ if $i = j$, and 0 otherwise.

The corresponding eigenvalues λ_i are given by solutions to $Lv_i = \lambda_i v_i$. The eigenvalues are sorted such that $\lambda_1 \le \lambda_2 \le \ldots \le \lambda_N$. By construction, the rows of the Laplacian sum to zero, indicating that $\lambda_1 = 0$ is always an eigenvalue. The multiplicity of λ_1 is the number of connected components of the graph. Thus, for a fully connected graph, we would have $0 = \lambda_1 < \lambda_2 \le \lambda_3 \le \ldots \le \lambda_N$.

The set of eigenvectors $\{v_1, \ldots, v_N\}$ form the *graph Fourier basis* for any graph signal $x \in \mathbb{R}^{N\times 1}$ on the nodes of the graph; we will refer to x as a *nodal graph signal*, or simply a graph signal with the understanding that it is supported on the nodes of the graph. The *graph Fourier transform* of a graph signal x is given by $(\alpha_1, \ldots, \alpha_N)$, where each α_i can be computed as

$$\alpha_i = x^\intercal v_i. \tag{2}$$

Intuitively, the graph Fourier transform (also known as *spectral decomposition*) quantifies the "contribution" of each eigenvector in the graph Fourier basis to x.

The *total energy* of the signal x is defined as $\|x\|^2$, which is also equal to $\sum_i \alpha_i^2$. The energy spectrum of a signal, $\left(\alpha_1^2, ..., \alpha_N^2\right)$, quantifies the distribution of the energy across different graph eigenvectors. α_i^2 is also referred to as the energy contribution of the i^{th} mode to the graph signal x.

Definition 1 (Total variation (TV))**.** *The total variation (TV) of a graph signal x on the graph $G = (\mathscr{N}, \mathscr{E}, A)$ is given by:*

$$TV = x^\intercal L x = \frac{1}{2}\sum_{i\neq j} a_{ij}(x_i - x_j)^2. \tag{3}$$

A graph signal with lower TV is said to be smoother relative to another graph signal with a higher TV.

More precise definitions of *graph signal smoothness* are made in Definitions 2 and 3.

A. Relationship between TV and spectral decomposition

Proposition 1. *The following two statements are equivalent:*

(i) *Total variation (TV) of a signal x is $x^\intercal L x$*

(ii) *Let $Lv_i = \lambda_i v_i$, for $i = 1, ..., N$, $\langle v_i, v_j \rangle = \delta_{ij}$, and $x = \sum_i \alpha_i v_i$, for some scalars α_i. Then the TV of x is $\sum_i \alpha_i^2 \lambda_i$.*

B. Correlations, smoothness, and unexpected events

Given two nodes $n_i, n_j \in \mathscr{N}$ connected by an edge $e_{i,j} = \{n_i, n_j\} \in \mathscr{E}$, observe k samples of the paired data $\left\{\left(x_{n_i}^{(1)}, x_{n_j}^{(1)}\right), ..., \left(x_{n_i}^{(k)}, x_{n_j}^{(k)}\right)\right\}$ where $x_{n_{\{i,j\}}}^{(\ell)}$ is the ℓ^{th} observation of the pair of random nodal signals on nodes n_i and n_j, for $\ell = 1, ..., k$. Let a pair of observations $\left(x_{n_i}^{(\ell)}, x_{n_j}^{(\ell)}\right)$ be drawn from the bivariate Gaussian distribution given by

$$\begin{pmatrix} X_{n_i}^{(\ell)} \\ X_{n_j}^{(\ell)} \end{pmatrix} \sim \mathcal{N}\left[\begin{pmatrix}\mu \\ \mu\end{pmatrix}, \quad \sigma^2 \begin{pmatrix} 1 & \rho_{X_{n_i}^{(\ell)}, X_{n_j}^{(\ell)}} \\ \rho_{X_{n_i}^{(\ell)}, X_{n_j}^{(\ell)}} & 1 \end{pmatrix}\right]. \tag{4}$$

The Pearson correlation coefficient for random variables $X_{n_i}^{(\ell)}$ and $X_{n_j}^{(\ell)}$ is $\rho_{X_{n_i}^{(\ell)}, X_{n_j}^{(\ell)}}$. We will now formalize the elements a_{ij} of the adjacency matrix A by defining them explicitly in terms of a *weight map*. Let $w : \mathscr{E} \to [0,1] : e_{i,j} \mapsto w(e_{i,j}) = a_{ij} = r_{x_{n_i}, x_{n_j}}^{(k)}$ be the weight map that maps an edge to the

projected *sample* Pearson correlation coefficient for a signal on node n_i and n_j computed from k prior data observations $\left\{\left(x_{n_i}^{(1)}, x_{n_j}^{(1)}\right), \ldots, \left(x_{n_i}^{(k)}, x_{n_j}^{(k)}\right)\right\}$ as

$$r_{x_{n_i},x_{n_j}}^{(k)} = \max\left(0, \frac{\sum_{\ell=1}^{k}\left(x_{n_i}^{(\ell)} - \hat{\mu}_i\right)\left(x_{n_j}^{(\ell)} - \hat{\mu}_j\right)}{\sqrt{\sum_{\ell=1}^{k}\left(x_{n_i}^{(\ell)} - \hat{\mu}_i\right)^2}\sqrt{\sum_{\ell=1}^{k}\left(x_{n_j}^{(\ell)} - \hat{\mu}_j\right)^2}}\right). \quad (5)$$

We use the notations $\hat{\mu}_i, \hat{\mu}_j$ and $\hat{\sigma}_i^2, \hat{\sigma}_j^2$ to represent the in-sample mean and variance. Furthermore, $r_{x_{n_i},x_{n_j}}^{(k)}$ is an estimate (formally, a *statistic*) for the population Pearson correlation coefficient $\rho_{X_{n_i}^{(\ell)},X_{n_j}^{(\ell)}}$. We also formalize the notion of *smoothness* by relating it to the TV for graph signals:

Definition 2 (Absolute smoothness)**.** *A signal $\left(x_{n_i}, x_{n_j}\right)$ on two nodes n_i and n_j is said to be smooth with respect to the edge $e_{i,j}$ with weight $w(e_{i,j}) = r_{x_{n_i},x_{n_j}}^{(k)}$ if the contribution to the TV of the graph signal approaches 0, i.e. $r_{x_{n_i},x_{n_j}}^{(k)}\left(x_{n_i} - x_{n_j}\right)^2 \to 0$.*

Since the pair of signals $\left(x_{n_i}, x_{n_j}\right)$ is being drawn from the joint distribution given in Equation (4), we modify the definition of smoothness to a weaker, probabilistic version. This can be thought of as the *data-driven* definition of smoothness:

Definition 3 (Probabilistic smoothness)**.** *A signal $\left(x_{n_i}, X_{n_j}|\left(X_{n_i} = x_{n_i}\right)\right)$ on two nodes n_i and n_j at the ℓ^{th} observation is said to be smooth in probability with respect to the edge $e_{i,j}$ with weight $w(e_{i,j}) = r_{x_{n_i},x_{n_j}}^{(k)}$ if, given observation x_{n_i}, the contribution to the TV of the graph signal converges to 0 in probability, i.e. $r_{x_{n_i},x_{n_j}}^{(k)}\left(x_{n_i} - X_{n_j}|\left(X_{n_i} = x_{n_i}\right)\right)^2 \xrightarrow{p} 0$. The notation $\xrightarrow{p}$ indicates a convergence in probability.*

The notation $X_{n_j}|\left(X_{n_i} = x_{n_i}\right)$ indicates that X_{n_j} is being drawn from the conditional distribution of X_{n_j} given an observation x_{n_i}. Probabilistic smoothness can be interpreted as follows: It is highly probable that the contribution to the TV of the graph signal approaches 0. Our main proposition relating smoothness and nodal signal correlation is as follows:

Proposition 2 (Smoothness and nodal signal correlation)**.** *The following remarks relating graph signal smoothness via TV and nodal signal correlation hold true:*

(i) *Given two weakly correlated nodal signals ($r_{x_{n_i},x_{n_j}}^{(k)} \to 0$), the graph signal supported by nodes n_i, n_j and edge $e_{i,j}$ is absolutely smooth.*

(ii.a) *Given two strongly correlated nodal signals ($r_{x_{n_i},x_{n_j}}^{(k)} \to 1$) and an observation x_{n_i}, the graph signals supported by nodes n_i, n_j and edge $e_{i,j}$ is smooth in probability if and only if $\left\|x_{n_i} - X_{n_j}|\left(X_{n_i} = x_{n_i}\right)\right\|_1 \xrightarrow{p} 0$.*

(ii.b) *Given two strongly correlated nodal signals ($r_{x_{n_i},x_{n_j}}^{(k)} \to 1$), the graph signals supported by nodes n_i and n_j, and edge $e_{i,j}$ is not smooth with probability $1-p$, where p is defined as*

$$p = \Pr\left(x_{n_i} - \mathscr{S} \leq X_{n_j} \leq x_{n_i} + \mathscr{S}|X_{n_i} = x_{n_i}\right) \quad (6)$$

where $\mathscr{S} = t_{p,k-2}\hat{\sigma}_j\sqrt{1 + \frac{1}{k} + \frac{\left(x_{n_i} - \hat{\mu}_i\right)^2}{(k-1)\hat{\sigma}_i^2}}$ is half of the width of the $(100p)\%$ prediction interval for $X_{n_j}|\left(X_{n_i} = x_{n_i}\right)$, with $t_{p,k-2}$ as the $(100p)^{th}$-quantile of a t-distribution with $(k-2)$ degrees of freedom.

IV. Data processing and setup

We consider the top airports (FAA Core 30) in terms of passenger and airline activity in the US for our analysis. We extract the total hourly departure and arrival delays at each of these 30 airports for the period 2008-2017. The signal vector for each day x is obtained by summing over the hourly departure and arrival delays from 0000Z to 2300Z.

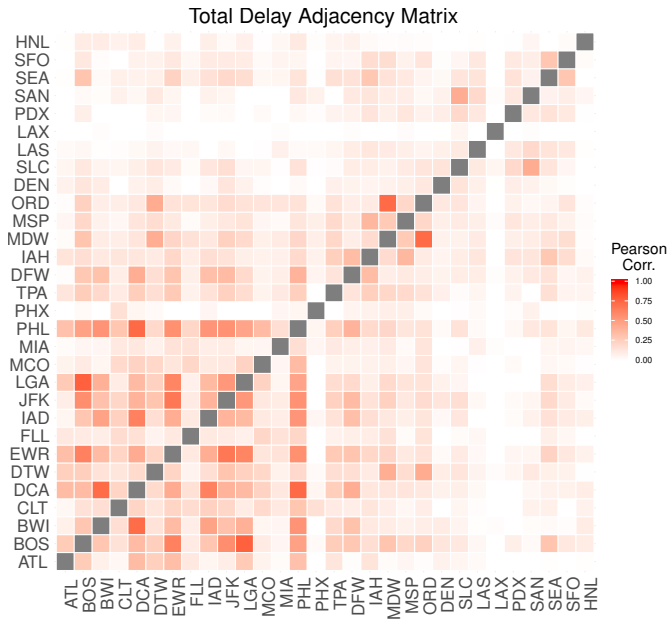

Fig. 1: Heat map of the adjacency matrix displaying the delay correlation between the top 30 airports, computed from Equation (5).

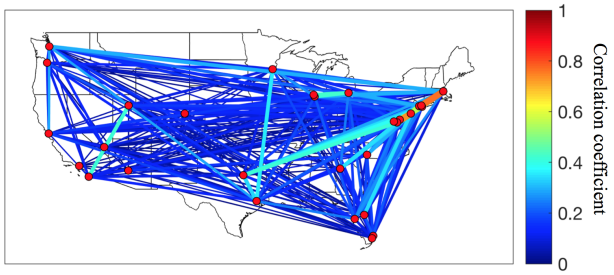

Fig. 2: Network with the top 30 airports as nodes and delay correlation as edge weights. Higher edge weights are also shown with a wider line.

We construct a graph using these 30 airports as nodes. The edge weight between any two nodes of the graph is the correlation coefficient between the delays at the two airports. We use FAA ASPM data for September 2018 to obtain the edge weights via Equation (5). The set of nodes

and edge weights forms the airport-delay correlation graph, on whose nodes we observe an airport performance signal x for each day. Figure 1 displays the value of the correlation coefficient between different airports. The same information is presented with geographical context in Figure 2, which has the 30 airports as nodes, and the edge weight representing the correlation.

The highest correlations are observed between the airports in the Northeast region. This means that the delays in these airports tend to move in the same direction. High traffic between airports in the Northeast region is one contributing factor to high correlations. Another instance of a high correlation pair of airport is ORD and MDW. While there is no traffic between these two airports (they are less than 20 miles apart), they almost always have the same weather impact. However, geographical proximity alone does not lead to similar delay trends, as the network connectivity of the two airports may be completely different. For instance, San Diego (SAN) and Los Angeles (LAX) are about 120 miles apart, and yet have a much lower correlation than LAX and Salt Lake City (SLC), which are 750 miles apart.

V. Implications of Propositions 1 and 2 for NAS Delay Networks

In this section, we aim to provide intuition for our theoretical results by interpreting TV and the spectral decomposition in the context of airport delays. The two main ideas we wish to convey are: (1) TV is a useful metric to help identify when a certain spatial distribution of delays across the air transportation network is unexpected; (2) Unexpected delays can be analyzed and interpreted by examining its spectral signature. Proposition 2 quantifies the likelihood of observing a certain spatial distribution of delays using the TV as a metric, whereas Proposition 1 relates the TV of the airport delays to the graph spectrum and identifies the specific patterns that are unexpected.

The key setup that enables our analysis is the edge weight being the correlation of delays between two airports. If delays at two highly-correlated airports are very close to each other, then $r^{(k)}_{x_{n_i},x_{n_j}}\left(x_{n_i}-x_{n_j}\right)^2$ is small, since $\left(x_{n_i}-x_{n_j}\right)$ is small (case (ii.a) of Proposition 2). If an airport pair is uncorrelated, then $r^{(k)}_{x_{n_i},x_{n_j}}$ is small, so the product $r^{(k)}_{x_{n_i},x_{n_j}}\left(x_{n_i}-x_{n_j}\right)^2$ is small regardless of the values of x_{n_i} and x_{n_j} (case (i) of Proposition 2). Note that the term $r^{(k)}_{x_{n_i},x_{n_j}}\left(x_{n_i}-x_{n_j}\right)^2$ is the mutual pairwise contribution of airport nodes n_i and n_j to the TV.

Recall that since TV is the sum of the terms $r^{(k)}_{x_{n_i},x_{n_j}}\left(x_{n_i}-x_{n_j}\right)^2$ over all airport node pairs n_i and n_j, if all airport node pairs have delays that follow expected trends (*i.e.*, the airport delays on a particular day follow historical correlations), then the TV is small. On the other hand, when the airport delays at historically highly-correlated airports do not follow typical correlations on a particular day, the signal would have a high TV. The spatial distribution of delays across airports is said to be *smooth* if the TV is small (Definitions 2 and 3). Thus, spatial delay distributions at airports that follow historical trends tend to be smooth.

v_i	ARTCC(s)	Trend 1	Trend 2
v_1	*Constant*	*Constant*	*Constant*
v_2	ZAB, ZLA, ZJX, ZMA, ZOA, ZSE, ZTL	MCO, FLL, CLT	SFO, SAN, ABQ LAS, PDX
v_3	ZJX, ZLA, ZOA, ZTL	CLT	MCO, SFO, LAX
v_4	ZAB, ZJX, ZLA ZMA, ZTL	FLL, ABQ	CLT, MCO, LAX
v_5	ZAU, ZBW, ZDC, ZDV, ZFW, ZJX, ZLA, ZLC, ZMA, ZNY, ZOA, ZOB ZSE, ZTL	PDX, LAS, SAN DEN, SLC, SFO	*All other airports in the ARTCCs* $(*)$
v_6	ZDV, ZLA, ZMA ZOA, ZSE	DEN	LAS, PDX, MIA SFO
v_7	ZDV, ZLA, ZSE	PDX	LAS, SAN, DEN
v_8	ZBW, ZDC, ZLA ZMA, ZNY, ZSE	PDX, FLL, LAX	*Same as* $(*)$
v_9	ZJX, ZLA, ZLC ZMA, ZOA, ZSE, ZTL	MIA, SAN, SLC	LAS, SFO, ATL MCO
v_{10}	ZDV, ZJX, ZLA, ZMA, ZOA, ZSE, ZTL	SFO, SEA, MIA DEN	*Same as* $(*)$
v_{11}	ZBW, ZDC, ZDV ZFW, ZHU, ZJX ZLA, ZMA, ZMP ZNY, ZOA, ZOB, ZTL	FLL, MCO, SFO SAN, DEN	*Same as* $(*)$
v_{12}	ZJX, ZMA, ZTL	ATL, FLL	MCO, CLT
v_{13}	ZAU, ZDC, ZFW, ZHU, ZJX, ZLA, ZLC, ZMA, ZMP, ZNY, ZOA, ZOB, ZSE, ZTL	ATL, MCO, SFO SAN	*Same as* $(*)$
v_{14}	ZAU, ZDV, ZHU ZLA, ZMA, ZOA	LAS, DEN	MIA, IAH, ORD SFO
v_{15}	ZAU, ZBW, ZDC, ZFW, ZHU, ZJX, ZMP, ZNY, ZOB, ZSE, ZTL	CLT, BWI, IAD SEA	*Same as* $(*)$
v_{16}	ZDV, ZLA, ZMA ZSE, ZTL	ATL, DEN	MIA, LAS, SEA
v_{17}	ZAU, ZBW, ZDC, ZFW, ZHU, ZJX, ZMP, ZNY, ZOB, ZSE, ZTL	MSP, CLT, SEA ATL, DFW	*Same as* $(*)$
v_{18}	ZAU, ZBW, ZDC, ZFW, ZHU, ZJX, ZMP, ZNY, ZOB, ZSE, ZTL	ORD, MDW, DTW SEA, CLT, ATL	*Same as* $(*)$
v_{19}	ZBW, ZDC, ZFW, ZJX, ZLC, ZMP, ZNY, ZTL	TPA, CLT	*Same as* $(*)$
v_{20}	ZAU, ZBW, ZDC ZFW, ZHU, ZJX ZMP, ZNY, ZSE	IAH, DFW, ORD	*Same as* $(*)$
v_{21}	ZDC, ZDV, ZFW, ZLA, ZJX, ZMA, ZOB, ZTL	ATL, LAS, TPA MIA	DFW, DEN, DTW IAD
v_{22}	ZAU, ZBW, ZDC ZFW, ZHU, ZNY	DFW, MDW	*Same as* $(*)$
v_{23}	ZAU, ZBW, ZDC ZNY, ZOB	MDW	*Same as* $(*)$
v_{24}	ZAU, ZBW, ZDC ZNY, ZOB	BWI, MDW	*Same as* $(*)$
v_{25}	ZBW, ZDC, ZNY	IAD	*Same as* $(*)$
v_{26}	ZBW, ZDC, ZNY	LGA, IAD	JFK, EWR, PHL DCA
v_{27}	ZBW, ZDC, ZNY	DCA, PHL	IAD, BWI, JFK
v_{28}	ZBW, ZDC, ZNY	BOS, EWR	JFK, LGA, DCA
v_{29}	ZBW, ZDC, ZNY	EWR	BOS
v_{30}	ZBW, ZDC, ZNY	PHL	*Same as* $(*)$

TABLE I
Description of eigenvectors, delay trends (trends 1 and 2 move in opposite directions), and the impacted airports and ARTCCs. Figure 5 provides a visualization of v_{24} and v_{29}.

As an illustrative example, take a pair of airports with perfectly correlated delays (*i.e.* their correlation coefficient is 1). The contribution to the TV from this pair of airports is always 0, since any new delays at both airports would be equal to each other with probability 1. However, in reality, such correlations may be close to, but not equal to, 1. Thus, given delays at one airport, the delays at the other cannot be

deterministically known and instead follow some probability distribution. Case (ii.b) of Proposition 2 characterizes this distribution and specifies a prediction interval for the second airport's delay (with probability p). For airport pairs with a higher correlation, the interval width would be smaller. In other words, the difference between the delays at both airports would be bounded by an interval of width $2 \times \mathscr{S}$ (Equation 6 in case (ii.b) of Proposition 2) with probability p. The contribution of this airport pair to the TV would be small with probability p. On the other hand, with probability $1-p$, the airport pair would have an unexpected distribution of delays, resulting in a larger contribution to the TV.

In practice, there are many airport pairs with a correlation coefficient between 0 and 1. In such cases, they will contribute to the TV whenever their delays are not the same. Thus, it is important to remember that even an expected distribution of delays would result in a non-zero TV. Furthermore, with increasing delays at airports, the TV is expected to grow quadratically (Equation 3). Therefore, it is important to only draw comparisons of TVs between days with similar quantities of delay. With these considerations, it is possible to conclude whether or not a distribution of delays is unexpected (*i.e.* not smooth) with respect to historical patterns.

While Proposition 2 allows for an airport-pair-specific decomposition of TV and interpretation of unexpected events, Proposition 1 provides a more global interpretation via spectral decomposition (Equation 2). The graph Fourier decomposition of airport delays into *eigenvector modes* provides a set of "templates" or spatial delay patterns (see Table I) in the form of eigenvectors that can be used to reconstruct any particular day's delay patterns across our airport graph. These eigenvector modes represent a set of airport delay patterns each with varying abilities to contribute to the TV (Proposition 1). Furthermore, the contribution of each eigenvector mode to the TV is proportional to its eigenvalue. Due to case (ii.b) of Proposition 2, delay patterns with higher TV are less likely to occur. Proposition 1 tells us that having more energy at higher eigenvector modes leads to a higher TV. Conversely, when higher TVs are observed, there is more energy in the higher eigenvector modes. This duality is highlighted by our case studies in Table II.

In Figure 3, we plot the total delay and corresponding TV for each day in the US NAS from 2008 to 2017. This plot enables us to visualize and interpret Proposition 2. The red line is the best quadratic fit for the data, and the dashed red line shows the 99% prediction interval for the quadratic fit. Recall that TV is expected to grow quadratically with total delay due to Equation 3. However, days with a significantly higher or lower TV for a given level of total delay are unexpected; the delays on those days were distributed across airports in a manner not seen previously. In the 10-year period we analyzed, 291 days (7.9%) fell outside of the 99% prediction interval.

It is worth noting that these days with unexpected spatial distribution of delays are *not* the same as days with high total delays. Furthermore, unexpected days (in terms of TV) may or may not also have high delay; one attribute does not imply the other, and vice versa. Specifically, what characterizes unexpected days is the fact that a very unusual and unexpected spatial distribution of delays were observed across the airports within the network.

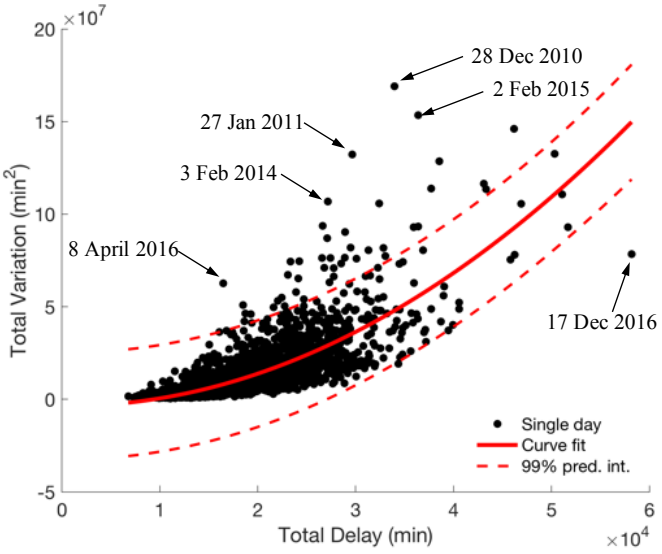

Fig. 3: TV versus Delay plot for 2008-2017.

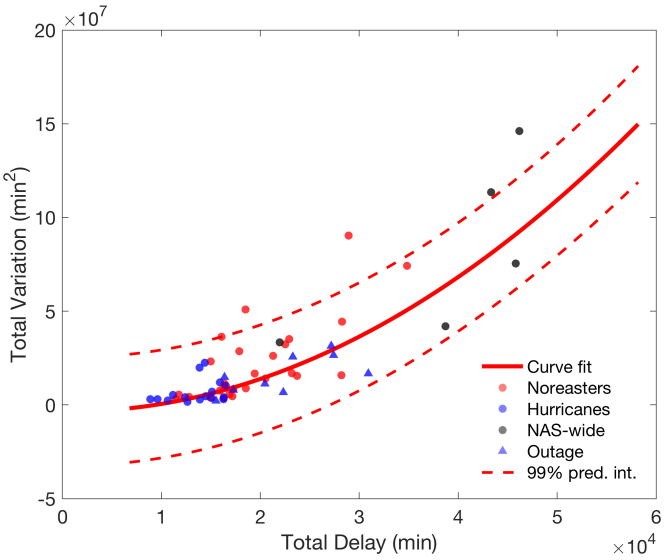

Fig. 4: TV versus Delay plot for specific types of disruptions (nor'easters, hurricanes, airport equipment outages, and NAS-wide disruptions, in 2008-2017.

VI. Analysis of Specific Off-Nominal Events

We chose four categories of off-nominal events that have the potential to cause severe disruptions: Atlantic hurricanes, nor'easters, airline-specific outages, and NAS-wide days when the entire system experiences significant delays, classified by [10]. Each event is recorded in Table II on a per-day basis. For

each day of a specific off-nominal event, we compute the total delay, the TV, and the graph Fourier decomposition. We also compute the energy contribution of each of the 30 eigenvector modes.

For each day, we selected the most energy-contributing eigenvectors and constructed a set $\{v_1, ...v_i\}$ of eigenvectors, ordered by their decreasing energy contribution. Typically, the constant eigenvector v_1 contains the majority of the energy. Using a threshold of 80% of the total spectrum energy, we construct the minimal set of eigenvectors such that the total energy of that set of eigenvectors meets or exceeds the threshold. We selected the 80% energy threshold in order to retain the most important eigenvector modes while maintaining interpretability. Each row in Table II corresponds to a particular day of an off-nominal event; we give the eigenvector index $i \in \{1, ..., 30\}$ as well as its individual energy contribution as a percentage.

A. Hurricanes

While hurricanes are known to cause major disruptions for airlines and airports, our spectral analysis shows that hurricane-induced disruptions tend to result in spatial delay distributions with lower TV across our system. In Figure 4, we see that days during three major Atlantic hurricanes (Sandy in October 2012, Harvey in end of August/beginning of September 2017, and Irma in September 2017) cluster towards the lower end both in terms of total delay and TV. More relevant in the context of the interplays between TV, smoothness, and unexpected events, the TVs exhibited by hurricanes do not exceed the 99% prediction interval. This indicates that, at least for the three hurricane events we examine, the airport delay distributions across the system are not considered to be unexpected.

This noticeable lack of graph signal variation for hurricane-related disruptions could be explained by a combination of operational and airline policy-driven responses specific to hurricanes. Typically, the most probable impact areas for Atlantic hurricanes are limited to the Southern and occasionally Mid-Atlantic portions of the US. These projected impact areas are continuously updated as the hurricane approaches, allowing airlines to preemptively cancel and re-position flights and crews. For example, once Hurricane Sandy was projected to make landfall on the East Coast, more than 7,000 flights were canceled ahead of Hurricane Sandy's East Coast landfall on October 29, 2012 [21]. This action may have reduced the TV induced by Hurricane Sandy's impacts due to the non-activation of East Coast modes (eigenvectors v_{25} through v_{30}, see Table I). The highly-correlated airports on the East Coast did not experience unexpected differences in their total delays.

We can also observe some interesting temporal trends in both total delay and TV for disruptions caused by hurricanes. For all three hurricanes, the recovery phase (*i.e.* the last one or two days of the off-nominal event) is typically characterized by a drop in both total delay and TV. In the context of Proposition 2, this trend suggests that the recovery phase towards the end of a hurricane's impact is smooth and as expected. This distinguishes hurricanes from nor'easters, where the TV towards the end of a nor'easter may increase dramatically (*e.g.* the January 2016 nor'easter), even though the total delay continues to dissipate. In other words, our spectral analysis suggests that airline recovery responses towards hurricanes versus nor'easters may differ in nontrivial ways.

B. NAS-wide high delay days

Recall that these days were selected from clusters identified in [10] wherein the overall delays experienced by airports within the NAS were large. Even within this small selection of NAS-wide high delay days, there are some interesting observations. For example, the total delays on January 3 and January 5, 2014 are comparable (4.61×10^4 versus 4.58×10^4 minutes), but the TV exhibited within our delay graph on January 3 is almost twice as high (146.17×10^6 versus 75.39×10^6 minutes squared). The TV exhibited in the delay graph of the former day exceeds the 99% prediction interval of the TV versus total delay curve, whereas the latter day does not. In the context of Proposition 2, the delay graph exhibited on January 3 represents an unusual event, even if the total delay was comparable to the total delay two days later.

A large contributing factor resulting in the unusual delay graph on January 3 is the activation of eigenvector v_{26}, wherein delays at LGA and IAD are not trending with delays at JFK, EWR, PHL, and DCA, even though their delay signals are historically highly correlated. This is a highly energetic mode, providing a large contribution to the TV. Other highly energetic modes such as v_{23} and v_{24} – the latter of which was also activated on January 3 – capture delay trends at BWI and MDW moving out of sync with other major airports in ZAU, ZBW, ZDC, ZNY, and ZOB ARTCCs. In contrast, we only see the activation of a medium-energy eigenvector v_{17} on January 5 – the behavior of this mode summarized in Table I suggests that airlines with hubs at MSP, SEA and ATL, or at CLT and DFW may be operating in a way that causes the delays at these five airports to not trend with other major airports in the listed ARTCCs for v_{17}.

C. Nor'easters

As mentioned earlier, nor'easters are known for severely impacting aviation operations, particularly around the East Coast and Mid-Atlantic regions. The severity of these off-nominal events often forces airlines to preemptively cancel flights in order to lessen the propagative effects of more tactical delays and cancellations, as well as to better position aircraft and crews for a quicker recovery.

From Table II, we note that nor'easter-type off-nominal events tend to activate many eigenvectors that encapsulate very high variations, *i.e.* v_i's where $i \geq 20$. In particular, eigenvectors v_{24} and v_{29} (visualized in Figure 5) appear 7 and 5 times, respectively. We can further augment our analysis from a geographic perspective by observing which airport nodes activate in the two prevalent eigenvectors v_{24} and v_{29}. Recall that all East Coast and Mid-Atlantic airports tended to be strongly positively correlated in terms of total delay (Figures

Event	Date	1st	2nd	3rd	4th	5th	Total Delay ($\times 10^4 min$)	Total Variation ($\times 10^6 min^2$)
Hurricane	10/28/12	1 (83%)	–	–	–	–	1.51	7.01
	10/29/12	1 (77%)	18 (4%)	–	–	–	0.89	3.01
	10/30/12	1 (77%)	9 (4%)	–	–	–	1.11	5.23
	10/31/12	1 (86%)	–	–	–	–	1.24	4.04
	11/1/12	1 (87%)	–	–	–	–	1.45	4.57
Hurricane	8/24/17	1 (87%)	–	–	–	–	1.63	4.01
	8/25/17	1 (88%)	–	–	–	–	1.63	2.98
	8/26/17	1 (58%)	20 (35%)	–	–	–	1.39	19.84
	8/27/17	1 (74%)	20 (13%)	–	–	–	1.59	12.03
	8/28/17	1 (78%)	18 (4%)	–	–	–	1.65	10.21
	8/29/17	1 (91%)	–	–	–	–	1.50	3.76
	8/30/17	1 (91%)	–	–	–	–	1.26	1.66
Hurricane	9/9/17	1 (85%)	–	–	–	–	1.06	2.41
	9/10/17	1 (77%)	19 (4%)	–	–	–	0.96	2.99
	9/11/17	1 (47%)	13 (17%)	12 (16%)	11 (4%)	–	1.43	22.42
	9/12/17	1 (89%)	–	–	–	–	1.39	2.89
NAS-wide	1/2/14	1 (68%)	23 (9%)	17 (7%)	–	–	4.33	113.44
	1/3/14	1 (71%)	26 (6%)	24 (5%)	–	–	4.61	146.17
NAS-wide	1/5/14	1 (78%)	17 (4%)	–	–	–	4.58	75.39
	1/6/14	1 (79%)	5 (4%)	–	–	–	3.87	41.92
NAS-wide	6/17/15	1 (67%)	20 (21%)	–	–	–	2.19	33.39
Nor'easter	2/25/10	1 (69%)	24 (9%)	5 (5%)	–	–	2.29	35.08
	2/26/10	1 (60%)	24 (10%)	26 (5%)	5 (4%)	29 (4%)	2.89	90.30
	2/27/10	1 (87%)	–	–	–	–	1.69	5.67
Nor'easter	1/30/11	1 (87%)	–	–	–	–	1.14	2.95
	1/31/11	1 (74%)	6 (8%)	–	–	–	1.64	9.69
	2/1/11	1 (71%)	22 (10%)	–	–	–	2.82	44.38
	2/2/11	1 (72%)	24 (5%)	28 (4%)	–	–	2.25	32.25
	2/3/11	1 (82%)	–	–	–	–	1.85	8.84
Nor'easter	2/7/13	1 (80%)	–	–	–	–	1.59	7.66
	2/8/13	1 (86%)	–	–	–	–	1.73	8.51
	2/9/13	1 (65%)	26 (11%)	29 (7%)	–	–	1.50	23.17
	2/10/13	1 (75%)	15 (7%)	–	–	–	1.64	11.33
	2/11/13	1 (72%)	5 (5%)	27 (5%)	–	–	2.13	26.23
Nor'easter	2/11/14	1 (88%)	–	–	–	–	1.63	3.64
	2/12/14	1 (76%)	13 (6%)	–	–	–	2.31	16.87
	2/13/14	1 (70%)	24 (5%)	5 (5%)	–	–	3.48	74.23
	2/14/14	1 (86%)	–	–	–	–	2.82	15.87
Nor'easter	1/26/15	1 (74%)	5 (6%)	–	–	–	1.94	16.68
	1/27/15	1 (83%)	–	–	–	–	1.18	5.49
	1/28/15	1 (84%)	–	–	–	–	1.28	4.18
	1/29/15	1 (89%)	–	–	–	–	1.49	3.88
	1/30/15	1 (79%)	24 (7%)	–	–	–	2.05	14.15
Nor'easter	1/21/16	1 (90%)	–	–	–	–	1.72	4.55
	1/22/16	1 (80%)	–	–	–	–	2.37	15.45
	1/23/16	1 (58%)	29 (13%)	28 (8%)	12 (3%)	–	1.61	36.30
	1/24/16	1 (56%)	24 (11%)	29 (6%)	30 (6%)	25 (5%)	1.85	50.93
	1/25/16	1 (68%)	29 (12%)	24 (5%)	–	–	1.79	28.60
Outage	11/15/12	1 (93%)	–	–	–	–	1.55	2.18
Outage	9/26/14	1 (73%)	18 (10%)				2.33	25.71
Outage	9/17/15	1 (80%)	–	–	–	–	1.73	8.02
Outage	7/20/16	1 (84%)	–	–	–	–	2.05	11.26
	7/21/16	1 (77%)	18 (9%)	–	–	–	2.74	26.50
Outage	8/8/16	1 (70%)	13 (7%)	12 (5%)	–	–	2.72	31.35
	8/9/16	1 (87%)	–	–	–	–	2.23	6.69
Outage	1/22/17	1 (79%)	13 (8%)	–	–	–	3.09	16.66
Outage	12/17/17	1 (65%)	13 (12%)	12 (10%)	–	–	1.64	14.86

TABLE II

DIFFERENT OFF-NOMINAL EVENTS; COLUMNS "1ST" THROUGH "5TH" CONTAIN THE HIGHEST-CONTRIBUTING EIGENVECTORS AND THEIR ENERGY CONTRIBUTION, IN DESCENDING ORDER.

1 and 2), reflecting both the operational interconnectedness as well as the geographic proximity between these airports. Hence, the delay distribution is expected to be smooth (case (ii.a) of Proposition 2). However, we note that during the off-nominal conditions imposed by various nor'easters, the delays at BWI are not behaving in an expected manner when compared to other East Coast airports (eigenvector v_{24}). Similarly, the delays at BOS and EWR are not behaving in the expected pairwise manner, as represented by eigenvector v_{29}. Thus, our spectral analysis suggests that nor'easters, more so

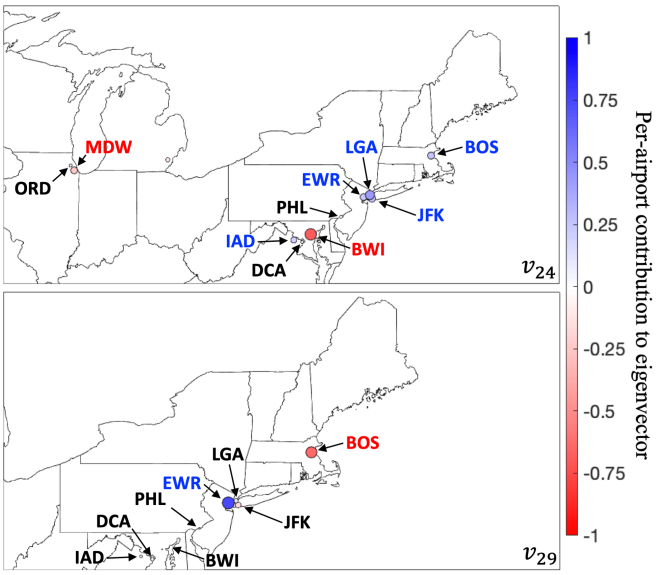

Fig. 5: Eigenvectors v_{24} (*top*) and v_{29} (*bottom*).

than other off-nominal events, drive the NAS into historically unexpected modes.

Even within nor'easter-type events, there are interesting differences in the temporal distribution of the high-energy contributing eigenvectors. For example, the February 2010 nor'easter peaks on February 26, driving the NAS into a very unexpected delay mode requiring five eigenvectors to characterize, three of which contain high levels of graph signal variance ($i \in \{24, 26, 29\}$). However, an abrupt recovery occurs between February 26 and 27, with the NAS settling back to a state where delay signals are smooth across the entire NAS (*i.e.* only v_1 is required to fulfill the 80% energy threshold) in just one day. This does not appear to be the case for the February 2013 nor the January 2016 nor'easter, where the recovery process was much less smooth in the context of graph signals. We hypothesize that this observation points to a non-trivial difference in recovery strategies beween different nor'easters – some strategies result in a more "even" network-wide recovery, whereas others result in recoveries that are less smooth and isolated. Future research that attempts to link different operational features in airline irregular operations (IROPs) recovery strategies with excited eigenvector modes would be of interest.

D. Airport- and airline-specific outages

Our spectral analysis shows that the airport- and airline-specific outages examined in Table II did not produce unexpected delay signals across our graph. Many of these off-nominal days (November 15, 2012; September 17, 2015; July 20, 2016; August 9, 2016) did not require additional eigenvectors past v_1 to pass the 80% energy threshold, indicating low levels of TV in the airport delays. Even though the total delay in the system was at a significant level for some of these days (*e.g.* 2.05×10^4 and 2.23×10^4 minutes of delay for July 20, 2016 and August 9, 2016, respectively), the TV was low and well within the 99% prediction interval that delineates an expected versus unexpected spatial delay distribution.

Furthermore, even when additional eigenvectors are required to achieve the 80% energy threshold, the eigenvectors tend to be less energetic, providing less contribution to the TV. For example, the two outages on August 8, 2016 and December 17, 2017 suffered by Delta Air Lines due to power outages at ATL identically triggered eigenvector modes v_{12} and v_{13}. Just from the index of the eigenvectors we can see that these are low-contributing modes to the TV. Both modes record delay signals at ATL trending opposite to delay signals at other airports within the NAS, but the delays at these other airports are weakly correlated to ATL (*i.e.* $r^{(k)}_{x_{n_i}, x_{n_j}} \to 0$), thus explaining the low contributions to the TV. If we examine Figures 1 or 2, we see that delay signals at ATL are more strongly correlated with East Coast airports such as BOS, DCA, EWR, LGA, and PHL, but these were not the airports whose delay signals were trending opposite to ATL's delay signals in eigenvector modes v_{12} and v_{13} (case (i) of Proposition 2). Recall that these eigenvector modes have airport-specific geographical interpretations; v_{12} and v_{13}, as described in Table I, implicate delays at ATL trending opposite to delays at other southeast airports such as MCO and CLT.

Operationally, it may be the case that airport- and airline-specific outages cause severe disruptions, but these disruptions are mostly localized to one particular airport node. In addition, for the cases of airport- and airline-specific outages that we examine, the airport node that was affected did not have a strong *a priori* positive correlation in terms of delay with other airports experiencing fluctuations in their delay signals. In other words, whatever delay signal fluctuations were occurring between the affected airport and other airports within the graph were not significant contributions to the TV, due to the weak correlation coefficient edge weights.

VII. Summary and Future Work

We analyze the networked behavior of airport delays within the US NAS using GSP techniques in order to take into account the underlying airport graph structure and delay correlations. In doing so, we compute the TV of our delay signals across our graph, and examine their eigenvector (spectral) components. Underpining our spectral analysis is our theoretical work linking together nodal signal correlation, the TV (smoothness) of a graph signal, and unexpected events.

Motivated by the knowledge gap in characterizing nominal or off-nominal disruptions with expected or unexpected airport delay dynamics at a network-level, our work resulted in two primary contributions. The first contribution – a systematic method, grounded in theory, to identify days with unexpected delay distributions at a network scale – has far-reaching implications for airlines and traffic managers. Through Proposition 2, we are able to show that, given historical delay correlations, a new delay signal induced by events such as hurricanes, nor'easters, and airport outages could be classified with explicit probabilistic bounds as *expected* or *unexpected*. Enabled by GSP techniques, this classification inherently accounts for

both historical trends and underlying connectivities, allowing for insightful, network-level interpretations of delay dynamics. Our second contribution stems from our case studies in Section VI, leading to the observation of fundamental differences in how delays begin, evolve, and dissipate between different categories of disruptive, off-nominal events (*e.g.* nor'easters versus Atlantic hurricanes).

In future work, we intend to explore the temporal dynamics of TV as a disruptive event unfolds; this direction already looks promising from our case study observations, hinting at deeper connections between TV and airport recovery procedures. We also will expand our case study to include all days classified as *unexpected*, enabling us to analyze other types of disruptions besides nor'easters, hurricanes, airport outages, and NAS-wide high delay days. Finally, we will compute airline-specific graph Laplacians and delay signals, leading to further, more detailed operational insights.

References

[1] A. Cook, H. A. Blom, F. Lillo, R. N. Mantegna, S. Micciche, D. Rivas, R. Vazquez, and M. Zanin, "Applying complexity science to air traffic management," *Journal of Air Transport Management*, vol. 42, 2015.

[2] US Department of Transportation, "Bureau of Transportation Statistics," 2018.

[3] US Department of Commerce, "The Historic Nor'easter of January 2016," 2016.

[4] A. Madhani, B. Mutzabaugh, and W. Spain, "FAA contractor charged with fire that halted flights," 2014.

[5] K. Sanders, D. Douglas, and K. Rosenblatt, "Power back on at Atlanta airport, but hundreds of flights still canceled," 2017.

[6] N. Pyrgiotis, K. M. Malone, and A. Odoni, "Modelling delay propagation within an airport network," *Transportation Research Part C: Emerging Technologies*, vol. 27, pp. 60–75, 2013.

[7] K. Gopalakrishnan, H. Balakrishnan, and R. Jordan, "Deconstructing Delay Dynamics: An air traffic network example," in *International Conference on Research in Air Transportation (ICRAT)*, 2016.

[8] S. Ahmadbeygi, A. Cohn, and M. Lapp, "Decreasing airline delay propagation by re-allocating scheduled slack," *IIE Transactions*, vol. 42, no. 7, pp. 478–489, 2010.

[9] Y. J. Kim, S. Choi, S. Briceno, and D. Mavris, "A deep learning approach to flight delay prediction," in *2016 IEEE/AIAA DASC*, 2016.

[10] K. Gopalakrishnan, H. Balakrishnan, and R. Jordan, "Clusters and Communities in Air Traffic Delay Networks," in *American Control Conference*, July 2016.

[11] J. J. Rebollo and H. Balakrishnan, "Characterization and prediction of air traffic delays," *Transportation Research Part C*, pp. 231–241, 2014.

[12] C. Gong and A. Sadovsky, "A final approach trajectory model for current operations," in *AIAA ATIO Conferences*, Sep 2010.

[13] R. Annoni and C. H. Q. Forster, "Analysis of aircraft trajectories using Fourier descriptors and kernel density estimation," in *2012 15th International IEEE Conference on Intelligent Transportation Systems*, Sept 2012.

[14] G. Yablonsky, R. Steckel, D. Constales, J. Farnan, D. Lercel, and M. Patankar, "Flight delay performance at Hartsfield-Jackson Atlanta International Airport," *Journal of Airline and Airport Management*, vol. 4, no. 1, pp. 78–95, 2014.

[15] T. Diana, "Do market-concentrated airports propagate more delays than less concentrated ones? A case study of selected U.S. airports," *Journal of Air Transport Management*, vol. 15, no. 6, pp. 280 – 286, 2009.

[16] M. Crovella and E. Kolaczyk, "Graph wavelets for spatial traffic analysis," in *IEEE INFOCOM 2003. Twenty-second Annual Joint Conference of the IEEE Computer and Communications Societies*, vol. 3, March 2003, pp. 1848–1857 vol.3.

[17] D. M. Mohan, M. T. Asif, N. Mitrovic, J. Dauwels, and P. Jaillet, "Wavelets on graphs with application to transportation networks," in *17th International IEEE Conference on Intelligent Transportation Systems (ITSC)*, Oct 2014, pp. 1707–1712.

[18] M. Drew and K. Sheth, *A Frequency Analysis Approach for Categorizing Air Traffic Behavior*. 14th AIAA ATIO Conference, Jun 2014.

[19] ——, *A Wavelet Analysis Approach for Categorizing Air Traffic Behavior*. AIAA ATIO Conference, Jun 2015.

[20] K. Gopalakrishnan, M. Z. Li, and H. Balakrishnan, "Identification of outliers in graph signals," *under review.*

[21] J. Freed and S. Bomkamp, "Hurricane Sandy Flight Cancellations: Northeast Airports Shut Down, Travellers Could Be Stranded For Days," October 2012.

Author Biographies

Max Z. Li is a PhD Candidate in the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology. His research interests include air traffic flow management, aviation systems modeling, and applied mathematics, particularly geometric and topological methods. He is a recipient of the 2018 FAA RAISE award, and a NSF Graduate Research Fellowship.

Karthik Gopalakrishnan is a PhD Candidate in the Department of Aeronautics and Astronautics at the Massachusetts Institute of Technology. His research interests include network dynamics, control theory, and modeling of air transportation systems.

Kristyn Pantoja is a PhD Candidate in the Department of Statistics at the Texas A&M University. Her research interests include machine learning, statistical education, and design of experiments.

Hamsa Balakrishnan is the Associate Department Head and Associate Professor of Aeronautics and Astronautics at the Massachusetts Institute of Technology. Her research interests are in the design, analysis, and implementation of control and optimization algorithms for large-scale cyber-physical infrastructures, with an emphasis on air transportation systems. She is a recipient of the AACC Donald P. Eckman Award (2014), the AIAA Lawrence Sperry Award (2012), the CNA Award for Operational Analysis (2012), an NSF CAREER Award (2008), and several best paper awards in ICRAT and ATM Seminars.