AIAA 2016-3761

AIAA Aviation
13-17 June 2016, Washington, D.C.
16th AIAA Aviation Technology, Integration, and Operations Conference

Multi-Scale Data Mining for Air Transportation System Diagnostics

Emily Clemons, Richard DeLaura, Yan Glina, Richard Jordan, Alex Proschitsky & Tom G. Reynolds[*]
MIT Lincoln Laboratory, Lexington, MA 02420, USA

Jacob Avery, Hamsa Balakrishnan, Cal Brooks, Mayara Conde Rocha Murca, Karthik Gopalakrishnan & R. John Hansman[†]
Massachusetts Institute of Technology, Cambridge, MA 02139, USA

I. Introduction

THE US National Airspace System (NAS) is remarkably safe, but there continue to be challenges of reduced efficiency and robustness, particularly when issues such as inclement weather or high traffic volume impact the system. As system demand and technology levels have increased, so too has the availability of air transport system data, including information on the current or future state of the system (e.g., weather, demand and capacity), aircraft data (e.g., surveillance tracks, flight plans) and airline data (e.g., schedules, expected push-back times and preferred routes) for the approximately 30,000 commercial flights operating per day in the NAS. Despite the prevalence of such data, the development of insights on system behavior has remained a challenge. Air transportation data sets present many of the challenges traditionally associated with "Big Data" problems, such as, large **volume** (e.g., over 1.0 TB of weather data and 1.2 GB of flight track/plan data per day), large **velocity** (flight data and weather radar updates that stream continually) and large **variety** (different data types, wide range of spatial and temporal scales, etc.). They are drawn from a range of sources (manual and automated), are often inconsistent, and are prone to errors (or even missing data), presenting additional challenges of **variability** and **veracity**.

The fundamental premise of this research is that air transport system performance can be improved through the application of "Big Data" techniques at a range of spatial and temporal scales of relevance to different types of NAS operations. In general, current alerting and decision support capabilities used by air transport system stakeholders (e.g., air traffic control, airlines, regulators, etc.) do not take full advantage of all the available data. The absence of scalable techniques that can process Big Data and convert them to actionable insight over a range of spatial and temporal resolutions is a key reason for this gap. It is anticipated that by overcoming these challenges, significant improvements can be made in the way the NAS is characterized, predicted and controlled, eventually leading to improved alerting and decision support tools, with associated system performance improvements and/or pointers to novel operating paradigms for future research.

The next section outlines the range of spatial and temporal scales of relevance in NAS operations and some of the opportunities for application of Big Data techniques to improve inefficiencies at each scale. A generic Big Data Analysis Framework is then proposed which contains the key steps in transitioning from raw system data to enabled insights at a given operational scale. The utility of this framework is then demonstrated through case studies at the three main scales of interest: tactical ATC operations, strategic ATC operations and airline network planning. Conclusions and recommendations for future work are given in the final sections.

II. NAS Inefficiency Identification

The main NAS inefficiency areas are shown in Table 1 covering a range of spatial and temporal scales with various data availabilities illustrated in Figure 1. These scales are relevant to different stakeholders in the NAS, including federal regulators, Air Traffic Management (ATM), Air Traffic Control (ATC), airports, airlines and pilots. At the strategic level, issues are observed over yearly/decadal and global scales, with drivers such as world

---

[*] All with Air Traffic Control Systems Group.
[†] All with Department of Aeronautics and Astronautics.

events, fuel price variations, airline competition and regulator policies. Inefficiencies are mainly driven by imbalances in gross systemic demand compared to capacity. Opportunities potentially afforded by Big Data include guidance to regulators and airlines regarding long-term investment and policy decisions. The strategic and tactical ATC scales cover weekly through minute temporal scales and national through local spatial scales, with demand and capacity imbalances at each scale being the main driver of inefficiencies (although the causes can be very different). The stakeholders who could benefit from the appropriate application of Big Data to mitigate these inefficiencies vary from airline schedulers/flight planners and FAA air traffic managers at the strategic ATC scale, to tactical controllers and pilots at the tactical ATC scale.

Table 1. NAS inefficiency summary.

Research Scale	Typical Temporal *& Spatial* Aggregation	Potential Drivers of System Behavior	Main NAS Inefficiencies	Opportunities for Big Data Application
Strategic Planning	Decadal/Annual/ Seasonal/Monthly *Global*	World events (e.g., 9/11) Fuel price Airline competition FAA policy	Limited airspace/ airport capacity Airline over-scheduling	Govt & Airline investment & Policy planning guidance
Strategic ATC	Monthly/Weekly/ Daily *National*	Macro demand (e.g., airline schedules) Macro capacity (e.g., large scale weather)	US-wide propagation of delays & cancellations	ATM and airline improved NAS-wide operational planning
Tactical ATC	Daily/Hourly *Regional*	Meso demand (e.g., airline flight plans) Meso capacity (e.g., regional weather)	Regional demand/capacity imbalances	Improved flight planning and traffic flow management
	Hourly/ Quarter Hourly *Local*	Micro demand (e.g., aircraft trajectories) Micro capacity (e.g., local weather)	Local demand/capacity imbalances	ATC & pilot improved tactical operations

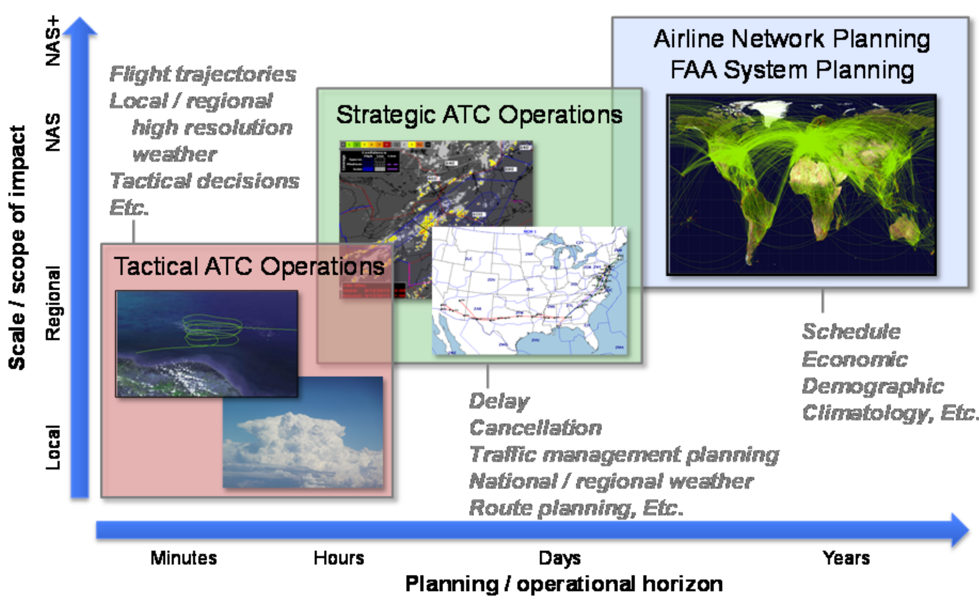

Figure 1. Relevant NAS spatial & temporal scales.

III. Big Data Analysis Framework

In order to explore the potential application of Big Data techniques to address the NAS inefficiencies identified above, a generic Big Data Analysis Framework has been created, as shown in Figure 2. "Raw" system data, such as demand (e.g., flight plans), capacity (e.g. maximum airspace & airport throughputs), system data (e.g., flight trajectories) and constraints (e.g., weather, traffic management initiatives), are reduced and aggregated into compact mathematical representations which enable more efficient data handling and analysis. From these representations, descriptive metrics for system characterization and identification of patterns of system behavior can be derived. These descriptive metrics provide searchable metadata that guide the partitioning of large-scale and heterogeneous datasets into operationally significant subsets that can be compared or mined further to generate insights, identify inefficiencies and formulate potential solutions. In this way, the size of the data reduces from the left to the right in the diagram, but the level of insight should increase if executed properly. The utility of this general framework is demonstrated in a range of relevant analysis case studies discussed at a high level in the next sections, but greater detail can be found in companion papers [1-5].

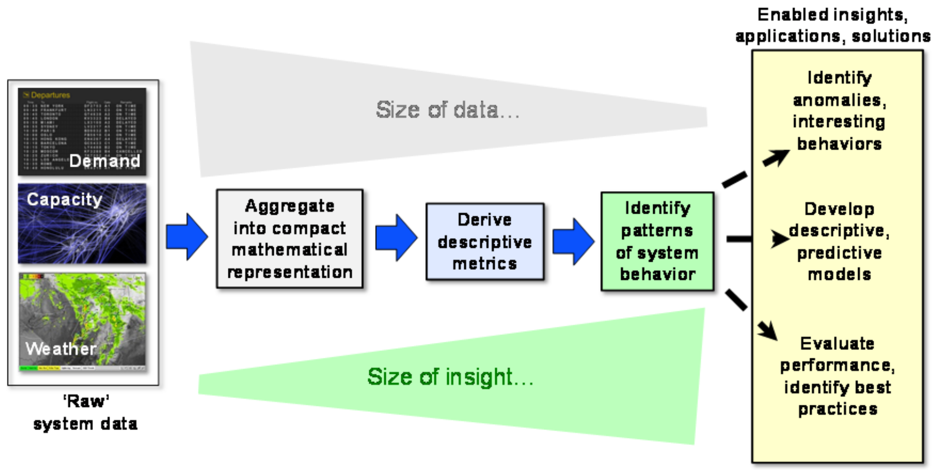

Figure 2. General anatomy of Big Data Analysis Framework.

IV. Tactical ATC Operations Analysis

Tactical ATC Operations may be divided into two functions: Tactical Air Traffic Control (Tactical ATC) and Tactical Traffic Management (Tactical TM). The primary goal of Tactical ATC is to maintain the safe separation of aircraft, which may be difficult as demand for an airspace resource (e.g. sector) exceeds its capacity. The task of Tactical TM is to allocate demand to local airspace resources efficiently while ensuring that capacities are not exceeded, even during off-nominal conditions. To date this work has addressed the needs of Tactical TM. A summary of the results is presented here; a more detailed description can be found in [3].

Several efforts have been made to develop analyses to identify traffic flows [6-9], sector capacity [10], flow capacity [11], and route availability [12] for post-event efficiency assessment, monitoring and alerting, and real time decision support for Tactical ATC Operations. However, none of these efforts address the need to estimate system throughput achievable under different operating conditions (e.g., demand, weather impacts, etc.) using different Tactical TM strategies. This work aimed to provide foundational methods to predict system throughput and identify effective Tactical TM strategies under different operational conditions. An initial case study was performed for New York Metro operations (John F. Kennedy International Airport (JFK), Newark Liberty International Airport (EWR), and LaGuardia Airport (LGA)). The methods were also applied to San Francisco (SFO), Dallas-Fort Worth (DFW), and Chicago O'Hare (ORD) airport operations as proof-of-concept of generalizability.

The framework adaptation for this scale is presented in Figure 3 and is composed of three sequential modules. Arrival and departure flows were defined by clustering Aircraft Situation Display to Industry (ASDI) flight tracks, which provide one minute updates of aircraft position, to learn typical patterns of operation. A trajectory classification scheme assigned flight trajectories to traffic flows and identified non-conforming trajectories. Finally, Tactical TM strategies were characterized as operational patterns by clustering hours with similar flows. An exploratory analysis examined the relationship between resource use, non-conformance, operational patterns and adverse weather. Convective weather impacts were characterized using the Route Availability Planning Tool (RAPT) [12], and terminal weather impacts (surface wind, ceiling, and visibility) were derived from hourly ASPM airport weather reports. The dataset included 1000 days of operations from 2013–2015.

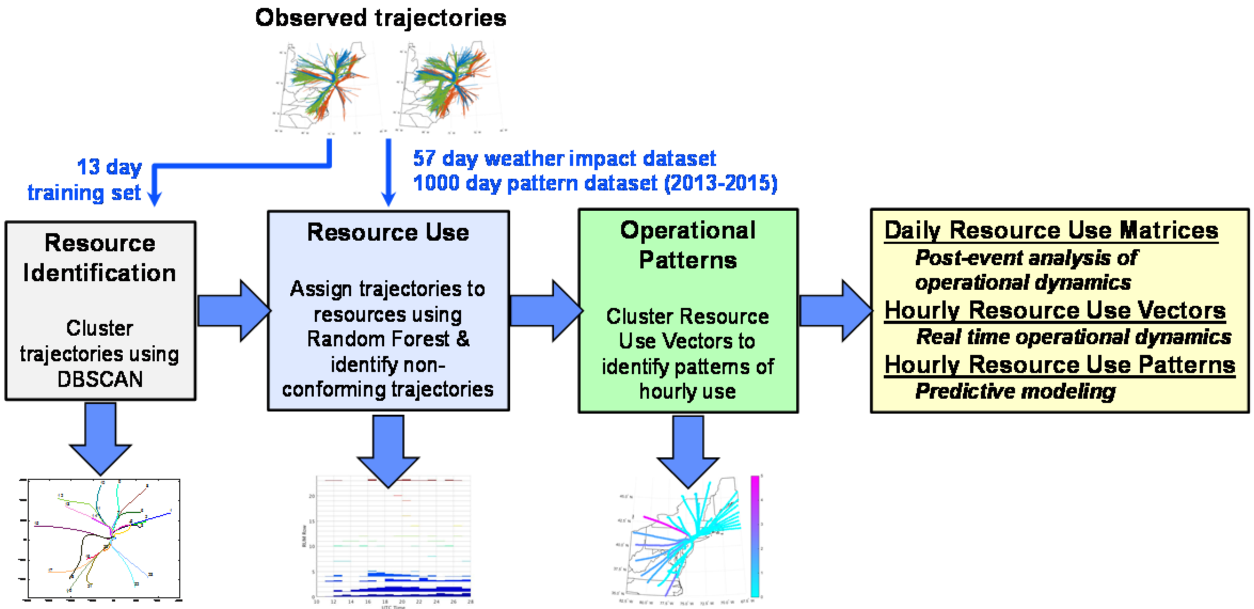

Figure 3. Tactical ATC operations big data analysis framework.

Resource Identification**:** A traffic flow definition clustering analysis was performed on a representative set of thirteen days selected from the dataset. Clustering analysis seeks to discover the structure of the observed data (flight trajectories) by identifying groups of observations with similar characteristics (traffic flows). A significant portion of the dataset, usually labeled as 'outliers', may not be captured by the clustering algorithm. Outliers, or "non-conforming" trajectories, represent flights that deviated in some significant way from the main traffic flow structure. Non-conforming trajectories may be operationally significant: they can increase the workload on air traffic controllers by not following the nominal procedures and hence adversely impact system throughput. They may also indicate dynamic use of airspace that increases throughput in the presence of constraints. Because of the importance of non-conforming trajectories and the spatial nature of the dataset, an algorithm called Density-Based Spatial Clustering of Applications with Noise (DBSCAN) [13] was chosen for this analysis. DBSCAN offers the ability to trade off the spatial spread of trajectories within clusters against the degree of non-conformance, making it well suited to flow identification. Cluster validity indices (Davies-Bouldin Index and Silhouette Index [3, 14]) were used to guide the selection of DBSCAN parameters. Figure 4 shows the set of arrival and departure flows defined by the method and a detailed view of EWR traffic flows, including the observed frequency of use for each flow.

Resource Use**:** The Random Forest algorithm [15] was used to assign flight trajectories to the identified resources. Random Forest is an ensemble-based classification method that relies on classification and regression trees (CART) built from bootstrap samples of the training data. The decision trees are computed by taking into account a random sample of the predictors at each node split. Each observation is then passed through each tree and its class is determined using a majority rule (highest percentage of tree votes in the ensemble). Five-fold cross validation was performed on the training data for tuning the parameters of the classification scheme. In order to detect non-conforming flight trajectories, we used a Conformal Prediction based method [16]. For any time scale, the classification scheme makes it possible to determine the patterns of traffic flow use and the level of trajectory non-conformance. The results are aggregated into hourly "Resource Use Vectors" (RUV) and daily "Resource Use Matrices" (RUM) that show the combination of resources used during each hour of the day. The RUM and RUV are very compact representations of tactical flows that characterize Tactical Operations. RUM and RUV are also suitable inputs to the clustering analysis used to identify operational patterns. The trajectory assignment algorithm was tested on a 52 day dataset that included fair weather, convection, and adverse terminal weather (Instrument Flight Rule conditions and/or strong winds). For each hour of operations, trajectories were assigned to arrival and departure flows, and hourly RUV, daily RUM, and the number of non-conforming trajectories was calculated. Non-conforming trajectories were far more prevalent on weather impacted days than on fair weather days (see Table 1).

Several forms of non-conformance were identified: direct-to routings, dynamic tactical reroutes to avoid convective weather, and holding and vectoring (observed most often on days with significant terminal weather impacts).

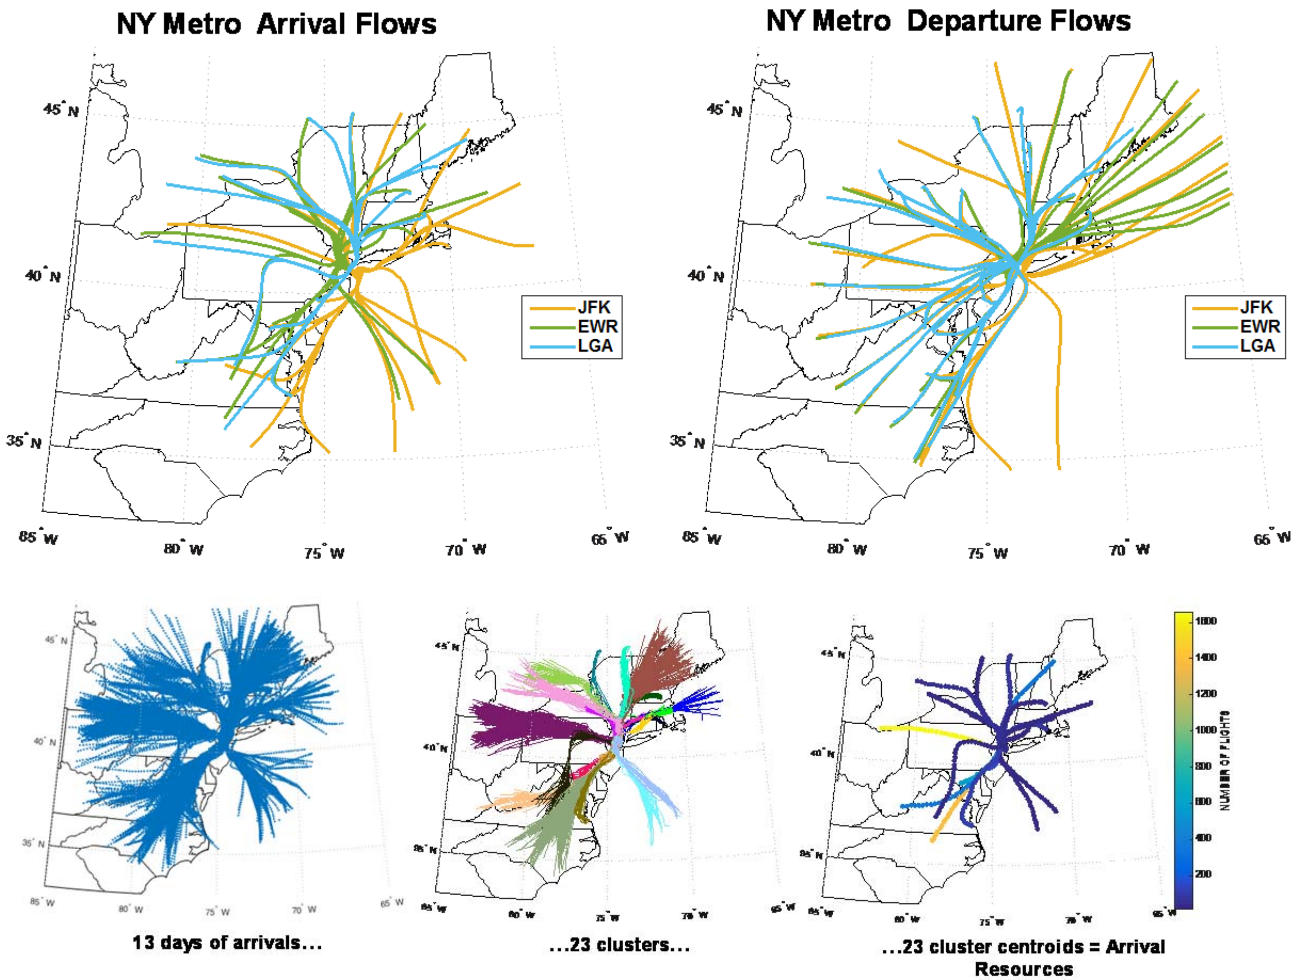

Figure 4. NY Metro arrival & departure flows (top); EWR trajectory clustering for flow identification (bottom).

Table 1. Daily percentage of non-conforming trajectories.

	Fair Weather (27 days)	Convective Impacts (16 days)	Terminal Impacts (9 days)
Pct. non-conforming	4.2	11.7	14.6
Std. deviation	2.4	5.2	8.7

Figure 5 compares daily operations from two days using the RUM: one with fair weather (Oct. 8, 2014) and the other with convection (July 14, 2015). The aggregate flow assignments (center panel) identify tactical reroutes and show an increase in non-conforming trajectories on the convective day. The RUM (left panel) identify the periods of the day when reroutes were most prevalent and show that reroutes and non-conforming trajectories were most observed during convective impacts.

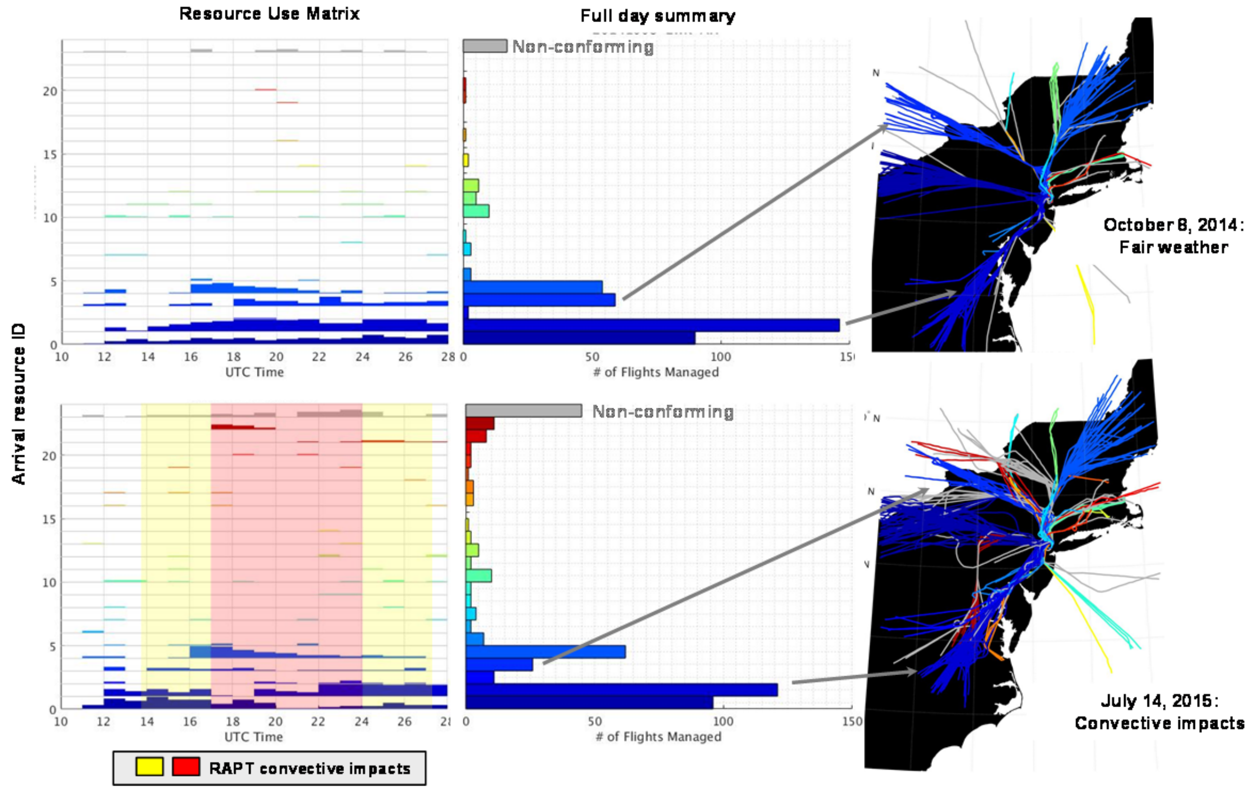

Figure 5. Comparison of two days of NY Metro tactical operations.

***Operational Patterns*:** Aggregate RUV for all three NY Metro airports were calculated for each hour of the 1000 days in the total dataset. Operational pattern recurrence was explored by a clustering algorithm that identified six major Resource Use Patterns (RUP). RUP clearly identified traffic flow patterns associated with different schedule pushes over the course of the day (e.g. morning departure push, early afternoon arrival peaks associated with international arrivals, late afternoon/early evening peak throughput). Figure 6 illustrates the average use of traffic flows for selected RUP. Figure 7 presents the average Metro throughput for each RUP and the observed probability of occurrence of RUP as a function of RAPT blockage. The probability of observing the High Non-conformance RUP increased markedly with moderate convective impacts; under severe convective impacts, the observed probabilities for both the High Non-conformance and Very Low Throughput RUP were elevated.

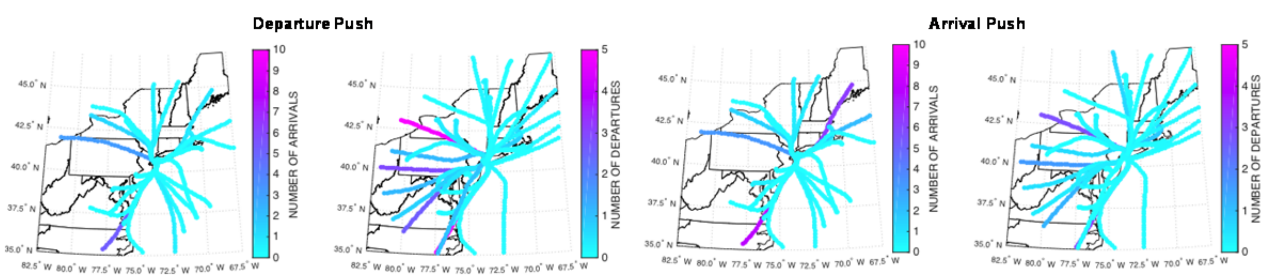

Figure 6. Comparison of average throughput for arrival and departure push Resource Use Patterns (RUPs).

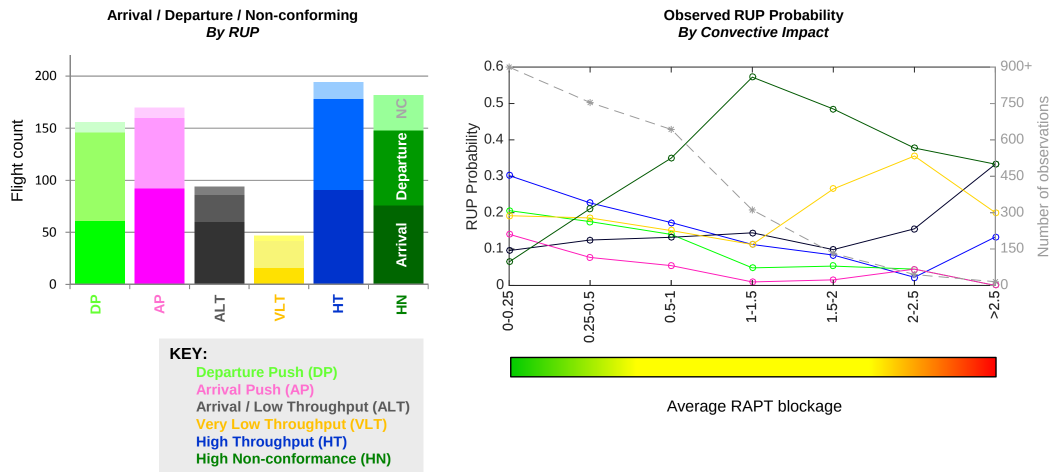

Figure 7. Average RUP throughput (left) and observed probability partitioned by convective impact (right).

In terms of the general data analysis framework introduced in Figure 2, this example illustrates how large amounts of raw data can be distilled into the compact RUP concept. These preliminary results suggest that the RUP concept permits terminal area behaviors to be classified into patterns that could form the foundation for a Metro system throughput prediction model. The intensity of flow use captured by the RUP appears related to patterns of demand and the presence of adverse weather constraints. Fine-grained detail of Tactical TM – for instance, the implementation of specific reroute patterns – may be captured by RUP derived from individual airport RUV and compared to less aggregated RAPT convective impact metrics to correlate reroute patterns to weather impacts. A RUP-derived model could identify operational patterns with sufficient throughput to accommodate expected demand and the feasibility of those patterns given forecast weather constraints.

V. Strategic ATC Operations Analysis

Strategic ATC operations typically cover broader air traffic management tasks with spatial scales of an air route traffic control center (ARTCC) size up to a national scale, and temporal scales of hours to days. Network representations of system behavior at this scale enable a range of standard and novel network metrics and methodologies to be explored. Based on these analyses, data-driven stochastic models were developed of the propagation of delay across the O-D pair network, and laid the groundwork for using the models for assessing system stability and resilience in the presence of delay-inducing disruptions. The key accomplishments and findings in these areas are described below. The adaptation of the analysis framework for analysis at this scale is shown in Figure 8.

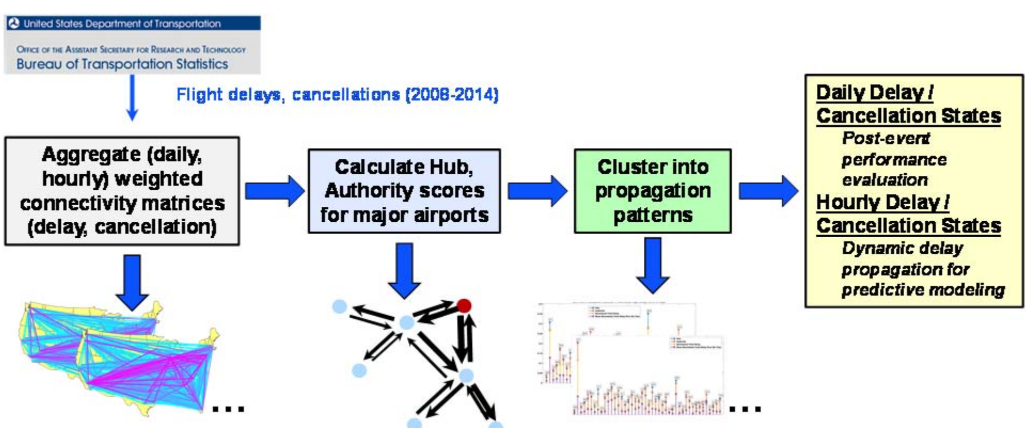

Figure 8. Strategic ATC operations big data analysis framework.

***Unique characteristics of air traffic networks*:** There exists a rich literature on network models for a vast range of infrastructure systems, including power grids, communication systems, ground transportation systems, public transport, railroads, and air transportation. In addition, spreading processes (such as disease epidemics, rumor propagation and consensus formation) have also been studied for these systems. While the development of data-driven models for these systems is far more limited, fundamental issues that this work has addressed includes identification of what makes air traffic systems unique, and what air transport characteristics can be identified by leveraging Big Data analytics.

Different interactions between components of air transport systems can be identified in the data by network abstractions which treat airports as the nodes of the network and the links between them as the edges. It is then possible to develop networks where each edge is weighted by the delay between those airports (and the delay associated with each node is the total delay at that airport). In doing so, several distinguishing characteristics of air traffic networks emerge, key among them being: (1) Nodes can experience a wide range of delays that are best described by continuous variables – by contrast, most standard epidemic models assume that a node is in one of a small set of discrete-states; (2) Not all interactions between nodes are identical, that is, the networks are weighted – by contrast, most prior literature assumes unweighted networks in which nodes either interact or they do not; (3) Interactions between nodes are inherently asymmetric – by contrast, most prior network models assume undirected networks in which interactions are symmetric; (4) Network connectivity is time-varying – by contrast, most prior literature assumes static networks. For example, Figure 9 shows the network of delay connectivities between different airports in the US, at different periods of time. The links are colored based on their weights, and illustrates that the weights can vary significantly from one time step to another.

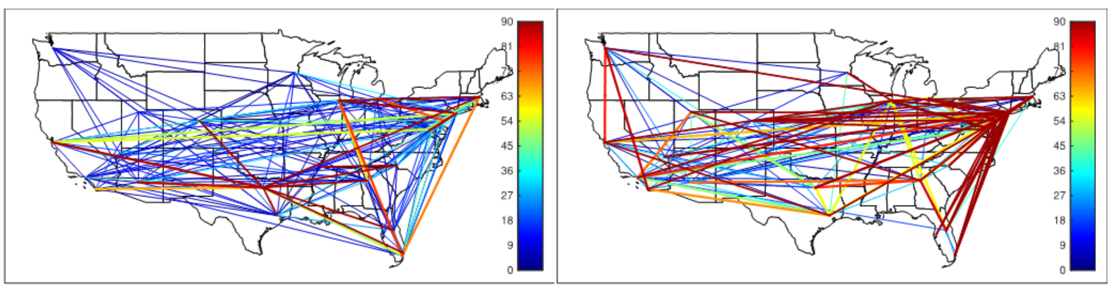

Figure 9. Network showing weighted connectivity of delays between airports at two different times. Links are colored by the average of the weights in the two directions for ease of visualization.

These distinctive characteristics of air traffic networks are important to identify and model using real data, since they drive system behavior. In other words, disruptions propagate differently in air transport systems *because* of these characteristics. Our analysis of these issues is presented in more detail in [1].

***Scalable analysis techniques*:** Big Data challenges are especially evident when exploring interactions between components of the air transport system. For example, when considering the 343 airports in the Bureau of Transportation Statistics (BTS) database, the resultant network has potentially ~120,000 (343^2) links, which means that there may be a very large number of parameters to estimate. The resulting challenge is two-fold: (1) The sheer size of the model and the required computation and storage; and (2) the amount of data that would be needed to reliably estimate such a large number of parameters (avoiding, for example, over-fitting). Motivated by these challenges, reduced-order and even sparse representations of the underlying networks were developed by leveraging the concepts of "hub" and "authority" scores (developed for evaluating relationships between webpages). A strong hub refers to an airport that has significant outbound or fan-out delays to strong authorities, airports that have significant inbound or fan-in delays from strong hubs. This approach is scalable on two counts: firstly, it is an $O(n)$ rather than an $O(n^2)$ representation of the network (that is, the number of model parameters is linear in the number of nodes), and secondly, it relies on the calculation of eigenvectors of matrices, a process that is computationally tractable and lends itself to parallelization. The resulting models are readily interpretable (owing in large part to the reduction of dimensionality) and reflect the unique characteristics of air traffic networks identified above, namely, that they are weighted and directed networks which can adopt a wide range of nodal states. Clustering based on the delays and hub/authority scores were used to identify characteristic delay patterns, which can also help model time-varying network structures [1].

***Extensible analysis techniques*:** An important advantage of the network characterization method identified above is that it is fundamentally agnostic to the source or field considered in weighting the networks. In other words, it could be applied to modeling delays (arrival, departure, ground or airborne), cancellations, traffic flows (operations) or capacity (weather impacts). Similarly, it could be used to analyze networks at vastly different temporal scales, from monthly, daily or even hourly levels to understand dynamics at those different scales [2]. For example, Figure 10 shows five of the 12 recurrent daily delay states identified by applying network analysis and clustering methodologies to BTS departure delay data from 2008-2014. The methodologies were able to uncover delay states that are largely geographically localized at major airports (e.g., the ATL-dominated HIGH Delay State), and also delay states that that are characterized by large delays spread across the network (e.g. the NAS-wide HIGH Delay State). The hourly delay states were analyzed in detail, and insight gained on the frequency with which they occur during different times of the year (i.e., dependence on season) and how they co-occur with characteristic NAS network cancellation states have led to valuable insights that can be leveraged in the identification of interesting case days and scenario generation. Our extensible method can be used to compare the effects of using different features (for example, cancellations instead of delays), different data sources (ASPM instead of BTS), different data types (a time-series of hourly states to characterize a day instead of a static aggregate network), etc. It also provides a pathway to building multi-layer, multi-timescale networks planned for future work.

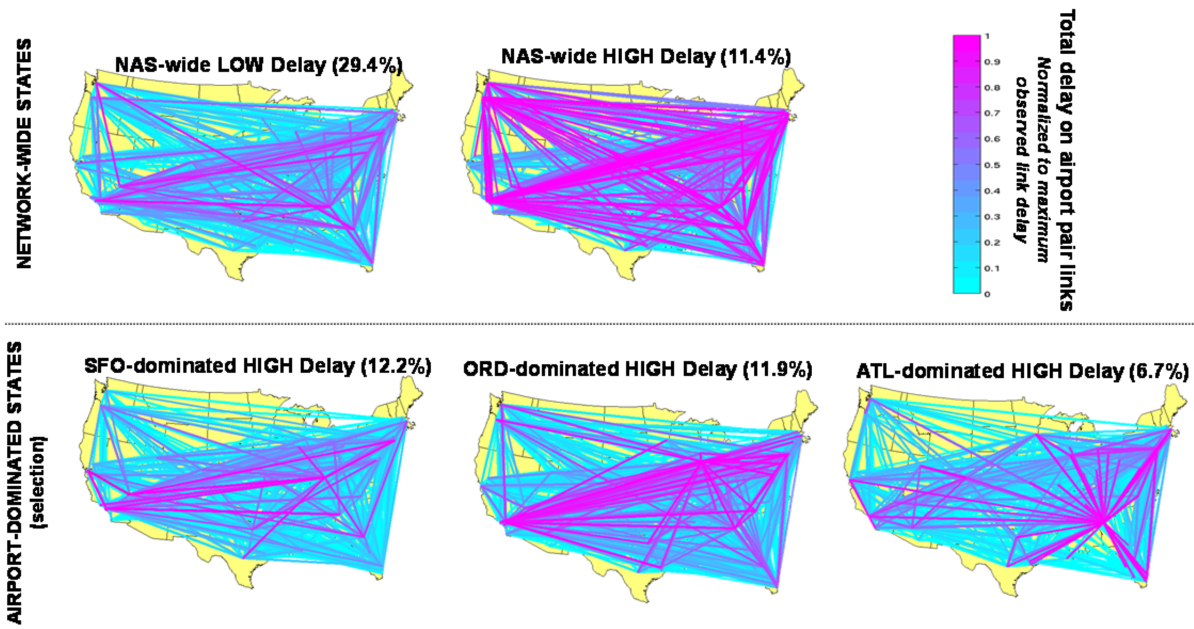

Figure 10. Selected recurrent daily delay states.

***Understanding the importance of capturing temporal patterns*:** An important aspect of reliably modeling the air transport system is the ability to capture the temporal dynamics. Doing so enables a better understanding of system stability and resilience to disruptions. The characteristic delay network patterns identified above have been used to develop stochastic dynamic models of delay propagation based on BTS data [2, 13].These models capture the unique characteristics of air traffic delay networks identified above, including stochastic temporal switching between delay states. Conditions for the stability of such systems (i.e., the decay of delays) have been developed [13]. Findings demonstrate the stabilizing influence of temporal patterns in transitions between different delay states. This is the first step toward validating models of delay dynamics, since it is found that when these temporal patterns are ignored, the resultant models predict a growth in delay over the course of an entire day (even though it is observed that delays always decay very late in the day). Figure 11 shows the expected amplification of total delay at the Core 30 airports as a function of time, compared to the delay at 4:00 AM EST. It is seen that maximum amplification factor occurs at around 6:00 or 7:00 P.M. ET and is 5.7 times higher relative to the total delay at the start of the day. Modeling the propagation, amplification and decay of delay, as well as characterizing switching between delay states and how it impacts system stability, is a key first step in ultimately modeling and understanding the efficacy of various delay mitigation strategies (such as Traffic Management Initiatives (TMIs)), and the investigation of alternate control strategies [13] such as TBO concepts.

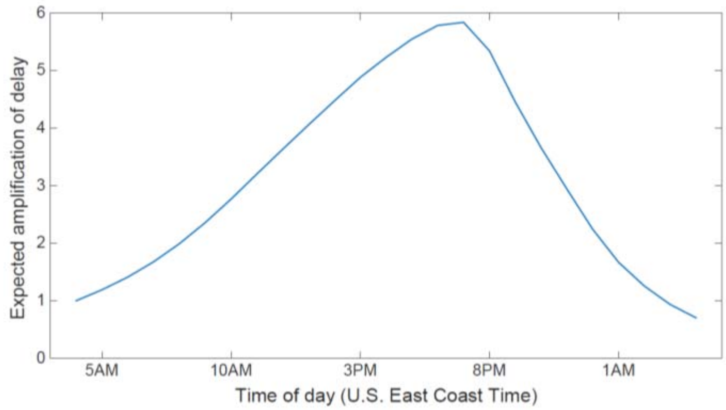

Figure 11. Upper bound on expected amplification of departure delay relative to 4:00AM ET, by time-of-day.

VI. Airline Network Planning Analysis

Analysis at the airline network planning time scale aggregates network and schedule data in order to observe structural changes in the NAS over multiple years. These structural characteristics are likely to impact system behaviors discussed in the previous sections, for example how quickly delay propagates through the NAS given the networks and frequencies serviced by the airlines. Analysis at this scale is less mature than the other scales just discussed, but some example insights are discussed here, and in more detail in [5]. By correlating network-wide descriptive metrics to exogenous factors such as airline mergers/bankruptcies, fuel price and economic activity, it may be possible to identify driving factors in NAS structure which in turn influences behaviors at the other levels. The ability to detect structural changes in the network also allows an analyst to correlate these changes to airlines' internal drivers such as passenger fare and load factor. Current work has used both standard and novel metrics and associated visualizations. Figure 11 shows the analysis framework at this airline network planning level.

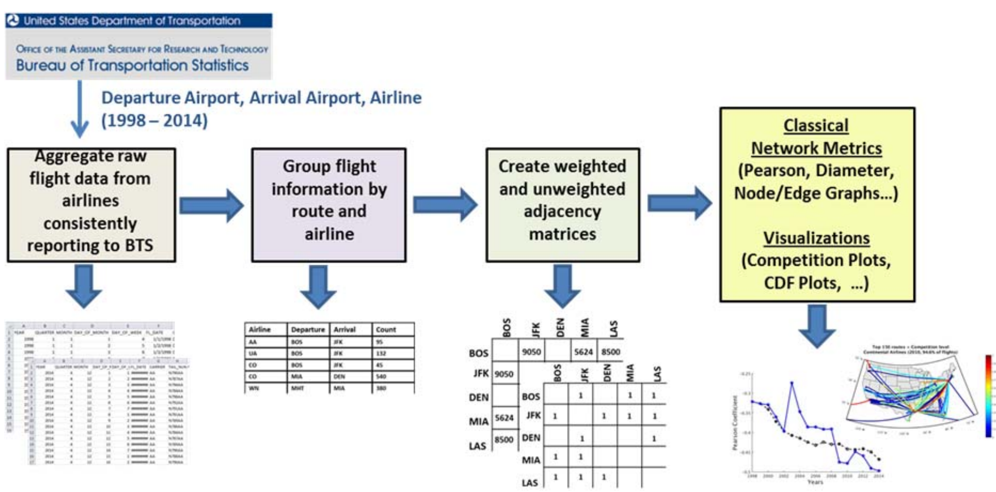

Figure 11. Airline network planning big data analysis framework.

Initial analysis has been performed on the aggregate data for a select group of airlines which consistently reported to BTS over the years of the study. Metrics were established from the network, with airports represented as nodes and with links corresponding to existing routes between airport pairs. Classical network metrics such as the Pearson coefficient [17, 18], Gini indices [19] and network diameter [17] have been applied alongside more fundamental metrics such as average stage length, load factor and the amount of competition along different routes. For example, Figure 12 shows the flight-count weighted and unweighted metrics for both the Pearson coefficient and the average stage length for the NAS-wide network. A steady decline in both the unweighted and weighted Pearson coefficients is observed over time, demonstrating a shift towards greater hub-and-spoke (or, in other words, away from point-to-point) structural characteristics. The average stage length metrics are also seen to climb. These tendencies impact how delay propagates through the network which would then be observed at the other analysis scales. For example, delay grows more rapidly when constraints such as bad weather occur at a key hub in a hub-and-spoke network, and Traffic Management Initiatives such as Ground Delay Programs take longer to have an effect when stage lengths in the network are longer.

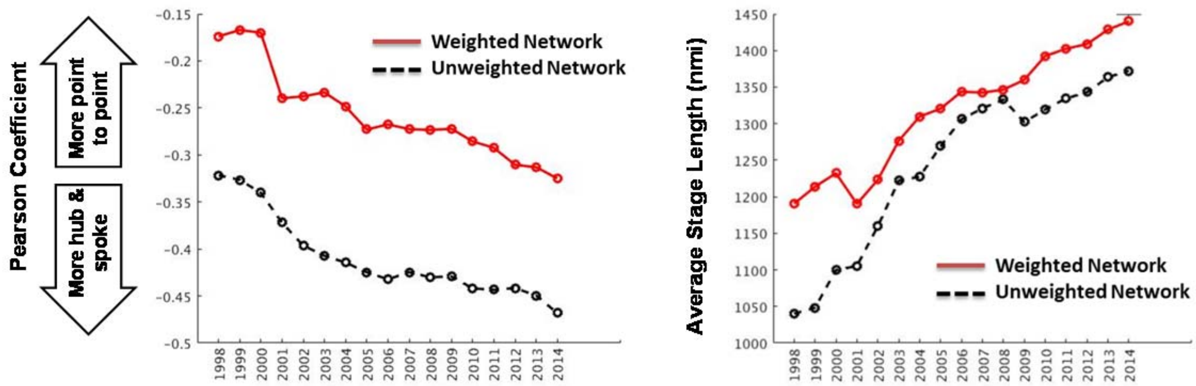

Figure 12. Pearson coefficients (left) and Average stage length metrics (right).

Some noticeable changes in the metrics are observed around 2001. A number of events occurred in the 2001 timeframe which could be driving the observed behaviors, including a merger between two big airlines (American Airlines and Trans World Airlines (TWA)) and the terrorist attacks of 9/11. This event caused a large spike in the number of cancellations as well, a metric which is intrinsically tied to the number of total flights – fewer scheduled flights mean a smaller number of total possible cancellations. The left side of Figure 13 highlights this behavior, in addition to revealing the trend of fewer flights and cancellations beginning after 2001.

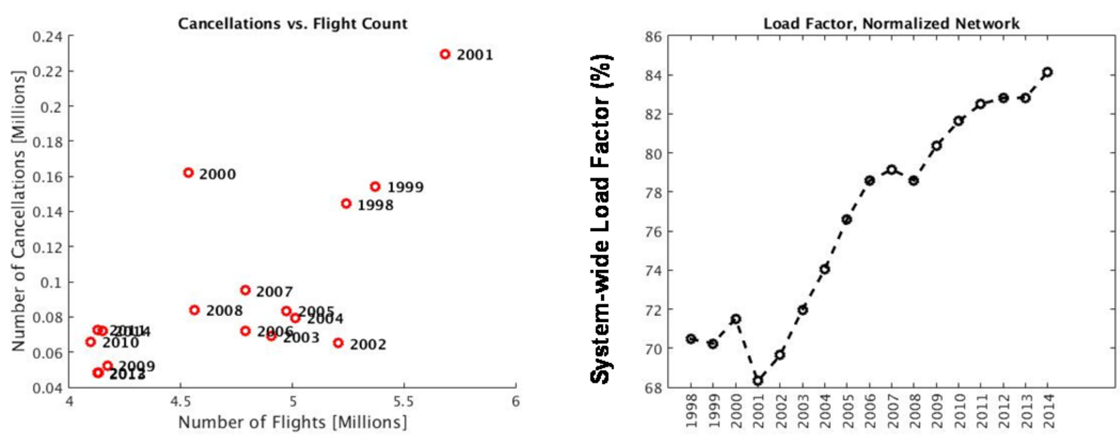

Figure 13. Cancellation count versus total flight count (left) and Load factor (right).

It can be seen that two distinct cancellation groups exist; after 2001, there is a significant drop-off in flight count and thus cancellations. This behavior is compounded after the financial crisis from 2007-2008, where there are even fewer available flights from the airlines present in the normalized network. This decrease is not so much a response to an actual demand in the system as to a greater tendency to optimize the available seats in departing aircraft by flying at a higher load factor. The right side of Figure 13 illustrates this trend, showing the increase in load factor over the years, rising steadily since 2002. This strategy of flying aircraft at higher load factor leads to fewer flight options for passengers and greater impacts of cancellations.

Another factor that impacts network dynamics is the route network served by individual airlines. Figure 14 below highlights an example case of the analysis of United Airline's network before and after the 2012 merger with Continental Airlines. These plots show the top 50 routes of the airline before and after the merger, color-coded by the fraction of flights on each route that is serviced by United Airlines. The competition plots show not only high levels of competition along United's main routes, which disappear after the merger, but a large shift in the main route network serviced by the airline. Note that despite this enormous change in United Airline's route network, many of the aggregate NAS metrics shown in previous figures indicate the overall network remains relatively stable, suggesting the NAS is quite robust to changes made to individual elements of the network.

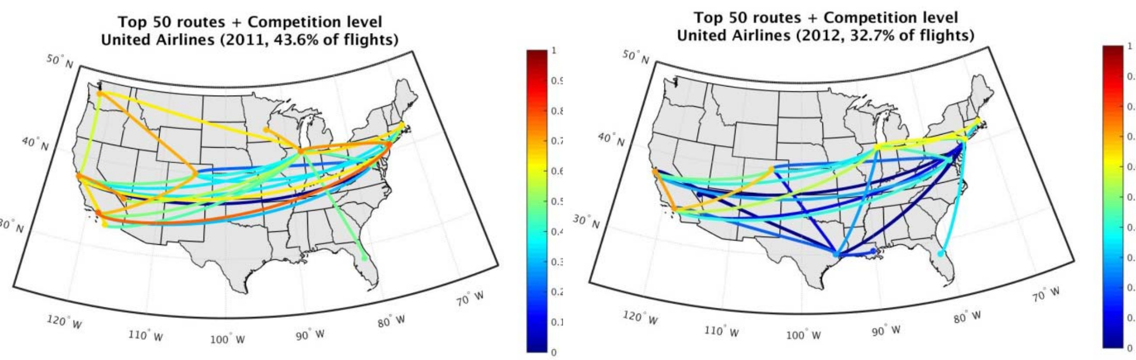

Figure 14. Example merger and airline competition impacts.

VII. Conclusions and Planned Next Steps

Big Data techniques hold significant potential for improving the robustness and efficiency of the air transportation system. This paper has highlighted some of the main NAS inefficiencies and resulting Big Data methodologies that can eventually lead to improved mitigation opportunities across a range of temporal and spatial scales relevant to different stakeholders. While the methods and analyses described above show great promise towards the possibility of reducing inefficiencies in the NAS, much work remains to be done. Recommended areas for follow-on efforts include extending and refining the work in several key areas:

- ***Transition to validated predictive capabilities*:** current efforts have focused on novel *diagnostic* characterizations of the NAS in compact and analyzable forms. Future work is planned to focus on the construction and validation of *predictive* models that could enable novel NAS control strategies and more effective use of existing control strategies by:
    - Incorporation of additional data sources that provide information about observed and anticipated NAS control decisions, resource constraints and system response.
    - Creation of compact metrics that characterize control, constraints and outcomes.
    - Development and validation of new analytics and predictive models to correlate decisions and constraints to system response and performance.
    - Development of concepts of operation for the application of predictive models.
- ***Development of big data approaches for multilayer, multi-timescale networks*:** The focus in the current work was the development of a methodology that focused on a single factor at a time, such as delays, cancellations or operations. However, in practice, air traffic networks are multilayered, carry multiple commodities, and nodes and links may interact in multiple dimensions. Future work is planned to extend the work to the analysis of multilayer, multi-scale networks. Figure 15 illustrates the schematic of a multilayered network comprising of an operation-weighted network layer, a cancellation-weighted network layer, and a delay-weighted network layer. Links can exist not just between the same layer at two different times (such as delays at two different times), but also different layers (such as delays at one time and cancellations at another). The figure also shows examples of how time-scales can differ: For example, delays at time t can influence delays at time t as well as those at time t + 1, depending on the nodes; similarly, they can also impact cancellations at t + 1 and t + 2 (or even later in time) along some of the links.

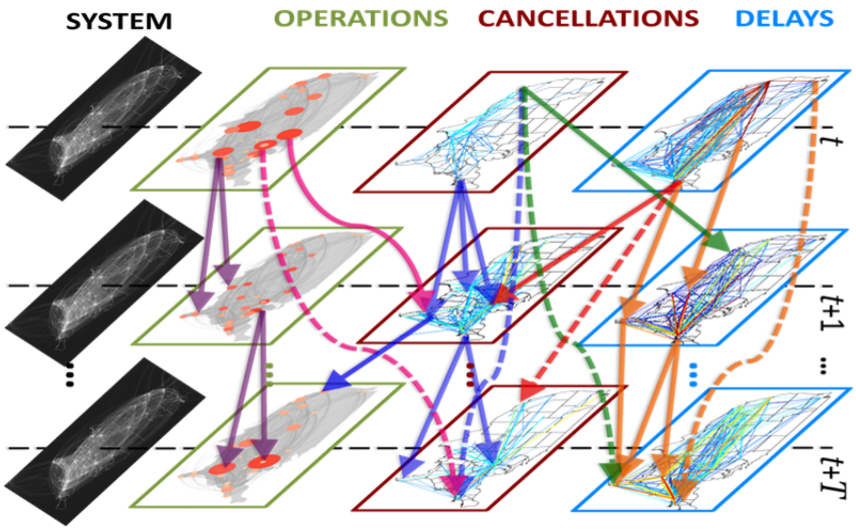

Figure 15. Multilayer, multi-timescale network perspective.

- ***Development of innovative visualization techniques*:** The analytics approaches developed in this work often result in compact representations of large data volumes. An effective visualization environment can enable researchers to identify possible correlations between operational strategies, constraints and system performance present in the data that can form the foundation of predictive modeling and can simplify the identification and preparation of interesting scenarios for concept development and simulation. The current work has created an initial concept and requirements for a Visual Environment for Data Analysis (VEDA) illustrated in Figure 16. The architecture is composed of three main layers: the batch layer, speed layer, and serving layer. The batch layer is responsible for pre-computing the vast quantities of historical data. The speed layer provides low latency computation for real-time streaming data products such as surveillance, weather and Traffic Management Initiative (TMI) data. The serving layer integrates the computation results from the batch and speed layers, and supports ad-hoc queries on the combined pre-computed dataset. VEDA elements are planned to be implemented in future work.

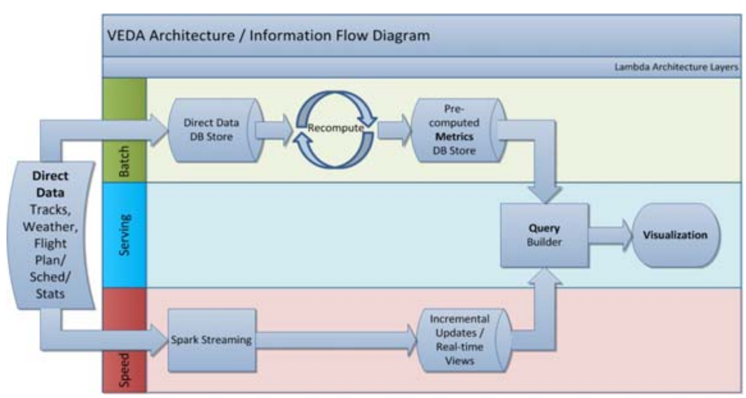

Figure 16. Visual Environment for Data Analysis (VEDA) concept.

Acknowledgments

This work was funded by NASA under the Leading Edge Aeronautics Research for NASA (LEARN) program. Many thanks to program manager Koushik Datta and technical monitors Sarah D'Souza and Michael Bloem from NASA Ames Research Center. Thanks also to the many NASA Ames Research Center program managers and researchers who have shared Big Data and domain application needs and insights.

Disclaimer

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited. This work was sponsored by NASA under the Leading Edge Aeronautics Research for NASA (LEARN) program. Opinions, interpretations, conclusions, and recommendations are those of the authors and are not necessarily endorsed by the United States Government.

References

[1] K. Gopalakrishnan, H. Balakrishnan & R. Jordan, "Clusters and Communities in Air Traffic Delay Networks", *IEEE American Control Conference*, Boston, MA, 6-8 July 2016.

[2] K. Gopalakrishnan, H. Balakrishnan & R. Jordan, "Deconstructing Delay Dynamics: An Air Traffic Network Example", *7th International Conference on Research in Air Transportation*, Philadelphia, PA (submitted), 20-24 June 2016.

[3] M. Conde Rocha Murca, R. DeLaura, R J. Hansman, R. Jordan, T. Reynolds & H. Balakrishnan, "Trajectory Clustering and Classification for Characterization of Air Traffic Flows", *AIAA Aviation Technology, Integration, and Operations Conference*, Washington DC, 13-17 June 2016.

[4] Y. Glina, E. Clemons, R. DeLaura, R. Jordan, A. Proschitsky & T. Reynolds, "A Visual Environment for Data Analysis and Air Traffic System Evaluation", *AIAA/IEEE Integrated Communications, Navigation and Surveillance Conference,* Herndon, VA, 19-21 April 2016.

[5] E. Clemons, Y. Glina, R. Jordan & T. Reynolds, "Airline Network and Competition Characterization using Big Data Approaches", *35th Digital Aviation Systems Conference,* Sacramento, CA, 25-29 September 2016.

[6] A. Eckstein, "Automated Flight Track Taxonomy for Measuring Benefits from Performance Based Navigation", *Integrated Communications, Navigation and Surveillance Conference*, Arlington, VA, 2009.

[7] M. Gariel, A.N. Srivastava & E. Feron, "Trajectory Clustering and an Application to Airspace Monitoring", *IEEE Transactions on Intelligent Transportation Systems*, 12(4):1511–1524, 2011.

[8] F. Rehm, "Clustering of Flight Tracks", *AIAA Infotech@Aerospace Conference*, Atlanta, GA, 2010.

[9] M. Enriquez, "Identifying Temporally Persistent Flows in the Terminal Airspace via Spectral Clustering", *10th USA/Europe ATM Research & Development Seminar*, Chicago, IL, 2013.

[10] L. Song, D. Greenbaum & C. Wanke, "The Impact of Severe Weather on Sector Capacity", *8th USA/Europe Air Traffic Management Research and Development Seminar (ATM2009)*, Napa, CA, 2009.

[11] M. Matthews, J. Venuti & R. DeLaura, "Translating Convective Weather Forecasts into Strategic Traffic Management Decision Aids", *5th Aviation, Range, and Aerospace Meteorology Special Symposium*, New Orleans, LA, 2016.

[12] M. Robinson, R. DeLaura & N. Underhill, "The Route Availability Planning Tool (RAPT): Evaluation of Departure Management Decision Support in New York During the 2008 Convective Weather Season", *8th USA/Europe Air Traffic Management Research and Development Seminar (ATM2009)*, Napa, CA, 2009.

[13] K. Gopalakrishnan, H. Balakrishnan & R. Jordan, "Stability of Networked Systems with Switching Topologies", *IEEE Conference on Decision and Control (CDC) 2016* (submitted).

[14] M. Ester, H.-P. Kriegel, J. Sander & X. Xu, "A Density-based Algorithm for Discovering Clusters in Large Spatial Databases with Noise," *KDD-96: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, pp. 226—231, 1996.

[15] L. Breiman, "Random Forests," *Machine Learning* Vol. 45, No. 1, pp. 5-32, 2001.

[16] G. Shaffer & V. Vovk, "A Tutorial on Conformal Prediction", *Journal of Machine Learning Research,* Vol. 9, pp. 371-421, 2008.

[17] M. E. J. Newman, "The structure and function of complex networks", *SIAM Review*, Vol. 45, No. 2, pp. 167—256, 2003.

[18] Bounova, G. A., "Topological Evolution of Networks: Case Studies in the US Airlines and Language Wikipedias", Ph. D. thesis, Massachusetts Institute of Technology, https://dspace.mit.edu/handle/1721.1/62965, 2009.

[19] J.C. Martin and A. Voltes-Dorta, "A note on how to measure hubbing practices in airline networks", *Transportation Research, Part E*, 45, 250—254, 2009.