Temporal Assessment of Resilience in the US Airline Network

Neel Rakholia, Poorvi Bhargava, and Noam Habot

December 12, 2016

Abstract

In this paper, we propose a unique temporal assessment of resiliency within the US airline network over the course of 2015. Nodes, representing airports, and edges, representing flights, were systematically removed from the network. As these networks broke down, resiliency was measured by the number of nodes that remained in the largest weakly connected component and by the average longest shortest path. As nodes were removed, we discovered that the network was less resilient on holidays such as Thanksgiving and New Year's Eve. We also observed that major weather storms during this time frame disrupted flight routes, causing a significant drop in resilience. Additionally, in our analysis of major airline carriers, Southwest Airlines appeared to have the network most resistant to disruption. Thus, our research can aid airports and airlines in better reacting to and designing infrastructure for adversity.

1 Introduction

Airline networks are important agents for bridging economical, geographical, political, and social boundaries. Analysis of their structure gives significant information about migration patterns, tourism, economics, and business. Moreover, the topology of the network also helps predict its behavior under adversity (closure of airspace, cancellations, delays, etc.). In this project, we intend to analyze the resilience of the US airport network over the course of 2015. Resilience of an airport network is determined by how important a given single airport, set of airports, flight, or set of flights are within the bigger picture of the entire network. We are interested in which factors affect resiliency over the duration of the year, including airlines, days of the week, passenger flux, time of the year, etc. Although much research has been conducted around modeling a static airport network, less attention has been placed on how these networks differ over time. We plan to take the work of Verma et al.[5] as an initial framework and extend its results by appending temporal flight data.

2 Motivation and Problem Statement

Airport network resiliency analysis is very important in understanding the deep backbone of the US flight network and observing how robust it is to sudden adversity, which leads to immediate changes in its components and structure. For example, certain weather conditions might prompt officials to shut down specific airports for a given period of time, or to cancel a flight that would have needed to fly through the storm. Another example is a terrorist attack on any scale - one that may shut down a single airport or one that may shut down an entire region of airports. This paper aims to show the effects of such airport closures or flight cancellations if these or similar events were to happen.

Another important calculation and consideration that we tend to in this paper is the resiliency of the airport network for a given airline (6.2.2). This analysis can prove to be meaningful to an airline that is trying to maximize their profits by maintaining a network that is well connected and resistant to unintended changes.

Furthermore, this analysis can assist consultants, business executives, or any other frequent flyers in picking a single airline that is most worth an investment of their loyalty. Our results guide a prospective frequent flyer into understanding which airline is best connected, reaches a wider variety of destinations, and has the highest resiliency to inevitable airport closures or flight cancellations.

3 Related Work

Verma et al.[5] examined the resilience of World Airline Network (WAN). They measured resilience of a static airline network by examining the fraction of nodes that remain in the largest weakly connected component (LWCC), upon removal of edges or nodes. Several conventional (removal of nodes) and unconventional (removal of edges) removal measures were used. Nodes were removed by highest degree whereas edges were removed by passenger flux. Passenger flux estimates the number of passengers traveling a particular route. Verma et al.[5] used their resilience analysis to propose a novel hierarchical structure for WAN.

Guimera and Amaral[4] analyzed past research for models to describe the WAN. They highlight that airports with the largest degree centralities are not usually the cities with the largest betweenness centrality. Moreover, there are critical airports (those with high betweenness centrality) that are not necessarily hubs within the network that act as major weak spots within the network. Their analysis motivated our choice of removal measures.

Fleurquin et al.[3] presented a model for quantifying, predicting, and analyzing time delay propagation in airline networks. Their model reproduces the delay propagation patterns by looking closely into three specific factors that influence flight delays: aircraft rotation, flight connectivity, and airport congestion. The model described in their paper accurately reproduced delay propagation within the system and lead to a few interesting results. However we did not heavily utilize these results in our paper.

Wuellner et al.[6] analyzed the resilience of airline networks with respect to the seven major airline carriers within the United States. They discovered that Southwest airlines had the most resilient network with targeted node removal and random edge removal. They suggested that the structural differences of the networks of these airline carriers attributed to the difference in resiliencies.

4 Data

The Bureau of Transport Statistics (BTS)[2] under the US Department of Transportation harbors detailed information about flights within the US in its Airline On-Time Performance Database. The database describes scheduled flights connecting a multitude of different airports operated by various airlines. The attributes of the full dataset extend beyond the scope of our research problem. We focused on information about the origin, destination, flight number, departure time and date, airline carrier, cancellations, delays, and flight distance. This data not only allows us to keep a record of which airports are connected to other airports, but also details delays experienced on each routes and along which airlines.

Variable	Magnitude
Flights	5,819,079
Airports	323
Airlines	15

Table 1: Summary of BTS Dataset

One challenge in the data collection process for this dataset was that the data was only downloadable by month. In addition, the datasets were so massive and contained extraneous information that loading in graphs from each CSV file at runtime would take an excessive amount of time. Thus, we created a script that took the entire month's flight data and generated multiple binary files. Each of these binary files represented a single day's flight network, including the desired nodes, edges, and attributes. We ran this script for all twelve months and thus generated 365 binary files. These binary files allowed us to quickly load in graph for each day into subsequent scripts for analysis.

The second dataset that we augmented to this data was the Federal Aviation Administration (FAA) Passenger Boarding (Enplanement) database[1]. It contains information about total passenger traffic at all US airports in 2015. We used this information to compute passenger flux – a measure of flight importance – for each edge in the network.

5 Methodology

5.1 Building Graphs

We created 365 separate directed multi-graphs that contained all flights for any given day throughout the year of 2015. The original dataset contained lots of information for each flight, including but not limited to: time of departure, flight distance, time at gate, time taxiing, etc. Therefore, we needed a way to map the edge to this flight information. Thus, we used airline ID and flight number as attributes for each edge. Airline ID is a unique string identifier for each airline obtained from the dataset. Flight number is an identifier distinguishing different flights from the same airline in a given day. It is generally unique except for a special case where a given airplane flies from airport X to airport Y and then either returns to airport X or continues to airport Z right after. This phenomenon, however, does not affect our analysis of the network because for a given origin airport on a given airline per day, the flight number is unique.

5.2 Network Statistics

Once we created the graphs, we computed statistics for each graph and averaged across all 365 of them (Table 2). We also constructed a picture of the airline network on January 1, 2015 and its degree distribution on the same day (Figure 1).

Total Number of Graphs	365
Mean Number of Nodes	305
Mean Number of Edges	15943
Mean Number of Out Edges	52
Mean Number of In Edges	52
Mean of Max Number of Out Edges	1039.7
Mean of Max Number of In Edges	1039.9

Table 2: Mean Statistics across all Graphs

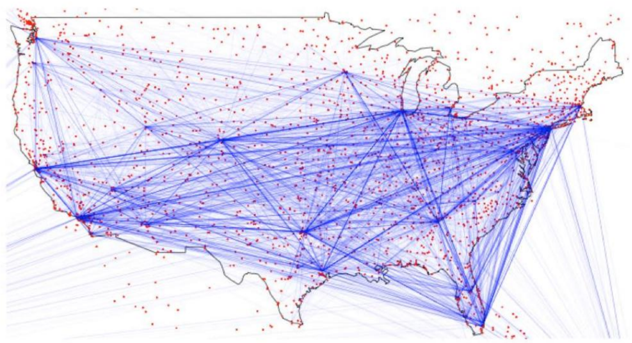

Figure 1: Graphical representation of the U.S. Airline Network on January 1st, 2015

5.3 Measuring Resilience

We measured the resilience of a graph in two different ways.

5.3.1 Resiliency Index: Fraction of Nodes in Largest Connected Component

We created a resiliency index based on the definition of Verma et al.[5]. They quantify the resilience of a network by the fraction of nodes that remain in the largest weakly connected component upon removal of edges or nodes from the network. In a real world situation, this could mean that an airport was shut down or a particular flight route was cancelled. One approach to visualizing this quantitatively is to plot the fraction of nodes that remain in the connected component, $S(q)$, as a function of q, the fraction of nodes or edges removed. This can be seen in Figure 2. We defined the area under this curve to be the resiliency index for a given day. To calculate the integral of the plot, we used the trapezoidal numerical integration, which approximates the region under the function as a trapezoid:

$$\int_a^b f(x)\,dx \approx (b-a)\left[\frac{f(a)+f(b)}{2}\right]$$

We used a unit spacing of $\frac{1}{(number\ of\ airports\ in\ graph)} \approx \frac{1}{300}$

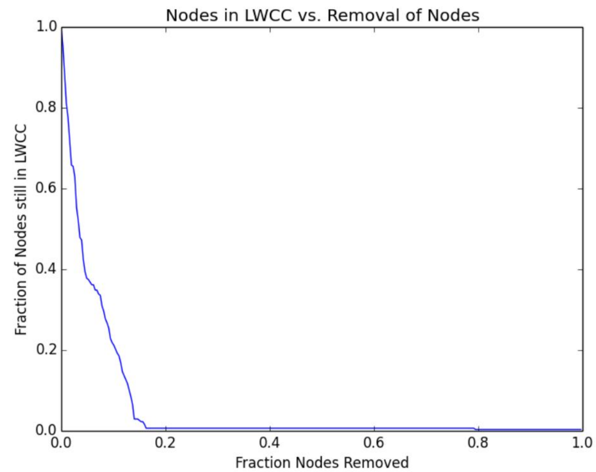

Figure 2: Example of resilience measured through removal by betweenness centrality

Note that another way to define the index from this graph is the x-intercept of the curve. This shows how quickly the graph breaks down based on how many nodes are removed from the graph. We did not use this definition.

5.3.2 Eccentricity Index: Longest Shortest Path

The longest shortest path, also known as Eccentricity, is defined as the distance between a given node X and the furthest node away from X. The distance to the furthest node from X is measured using the shortest path technique, which returns the number of nodes required to traverse the path between X and the furthest node. The longest shortest path of a given node is also called the eccentricity of that node.

For a graph corresponding to a given day in the year, we can compute a value called the Eccentricity Index. The process of computing the eccentricity index requires an iterative technique that iterates N times, where N is the number of total nodes in the graph. Within each iteration, we compute the average eccentricities of all the nodes in that iteration. At the end of each iteration, we remove the current highest centrality node, based on either degree, closeness, or betweenness centrality.

After completing all of the iterations, we obtain a vector that corresponds to the average eccentricity, where the i^{th} value corresponds to the average eccentricity with $n-i$ nodes in the graph. To compute the eccentricity index of the entire graph, i.e. for that day, we compute the area under the curve using the trapezoidal function, as explained in 5.3.2.

5.4 Removal Measures

In the physical world, removing a node from the network represents a situation in which an entire airport is forced to shut down and no flights can take off from or land at the specified airport. We began by analyzing this phenomenon. We systematically removing nodes from the network until no nodes were left. In a real world scenario, however, flight, or edge, cancellations occur more frequently as compared to multiple airport-wide shutdowns. Thus, a scenario with systematic edge removal can more accurately depict the overall topology of the network under adversity, as opposed to systematic removal of nodes. Additionally, removing particular edges allows us to analyze the effect of adversity at a more granular level.

5.4.1 Node Removal

Verma et al. analyzed the resilience of their network by successively removing the highest degree nodes from the graph. In a similar fashion, we analyzed the resilience of each graph representing each day in the year. However, instead of only removing nodes based on degree, we also considered two other measures:

1. **Highest degree centrality node:** compute the sum of out degree and in degree for each node, sort on descending order, and remove nodes until no nodes remain.
2. **Highest closeness centrality node:** compute the closeness centrality for each node, sort on descending order, and remove nodes until no nodes remain.
3. **Highest betweenness centrality node:** compute the betweenness centrality for each node, sort on descending order, and remove nodes until no nodes remain.

5.4.2 Edge Removal

In addition to removing nodes, we also explored the effect of edge removal based on the following methods:

1. **Number of passengers traveling over a particular route (passenger flux)**
   To remove edges from the network, we needed a quantifiable measure of their importance. Verma et. al. [5] defined passenger flux, an estimate of the number of passengers flying a particular route between two cities, as a metric for flight importance. We followed a similar scheme. We defined passenger flux p_{ij} as:

$$p_{ij} = p_i \frac{k_j}{\sum\limits_{z \in neigh(i)} k_z}$$

   Here p_{ij} is the value of the flux for a flight from i to j. p_i is the total number of passengers visiting (arriving, departing, or transiting) airport i. k_j is the out-degree of node j. Qualitatively, this measures the importance of a flight route based on the estimated number of passengers flying that route. We perform this computation for all edges in the network. We then follow two removal schemes: 1) Sort the passenger fluxes in descending order and successively remove edges with the highest passenger flux first. 2) Sort the passenger fluxes in ascending order and successively remove edges with the lowest passenger flux first (see Figures 3, 4).

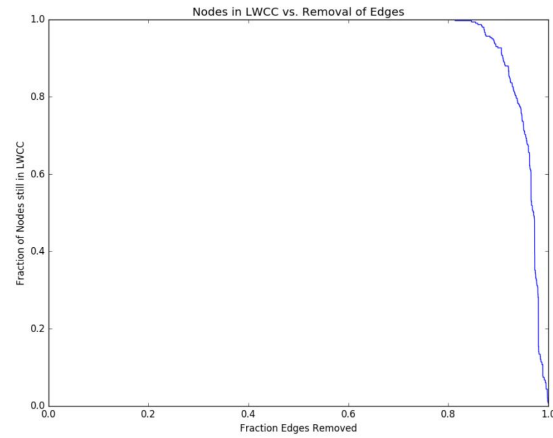

Figure 3: High to low passenger flux

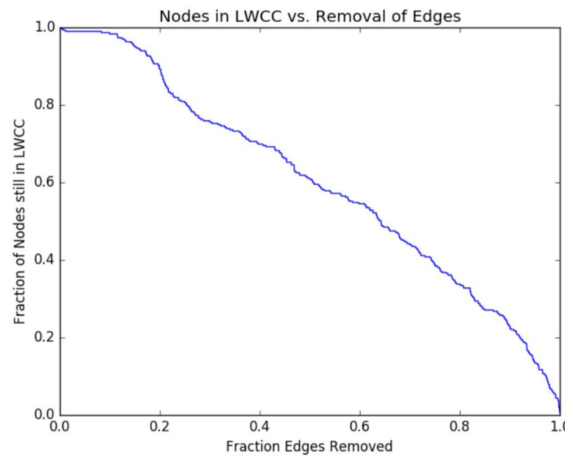

Figure 4: Low to high passenger flux

As can be seen in the plots above, removing edges on high passenger flux did not significantly affect the resiliency of the graph until 80% of the total edges were removed. Removing edges based on low passenger flux, however, did affect the resiliency of the graph drastically, confirming the results of Verma et. al. [5]. This is probably attributed to the fact that removing edges with low passenger flux essentially disconnects less popular airports first, reducing the number of nodes in LWCC at a faster rate. From a practical standpoint, removing edges with low passenger flux is of lower importance.

2. **Highest Betweenness centrality edge**
   Thus we explored another method of edge removal - removing edges by betweenness centrality in descending order. Instead of using passenger flux as a measure of how important an edge was, the betweenness centrality for each edge, or the number of shortest paths passing the edge, was computed. The edges were then sorted in descending order of betweenness and removed sequentially until no edges remained in the graph. The resilience of the network under highest betweenness centrality edge removal was measured using the resiliency index as previously defined.

3. **By Airline** In addition to simply removing edges in a particular order, we studied the effect of (a) removing all of the edges of a particular airline and measuring the subsequent resiliency of the remaining graph and (b) removing edges of all airlines except a particular airline and measuring the resiliency of the remaining airline's network. After removing the flights of the airlines not being studied, edges were removed based on highest betweenness to analyze the resiliency of the remaining networks. This analysis allowed us to answer the questions of (a) Is the resilience of the entire US network dependent on just one or two main airlines? and (b) How resilient are each of the major airline networks? Is one significantly more resilient than the others?

5.5 Temporal Assessment

A key aspect of the entirety of these analyses is that they are executed not only for randomly chosen individual days, but also for the entire year. This allows us to spot trends and significant differences in the resiliency and eccentricity indices throughout the year, subject to both the node and edge removal measures that are described above. This extended analysis helps us determine how the network changes over the year, including whether it is affected by seasons, specific weeks, or specific days.

6 Results

6.1 Node Removal

6.1.1 Analysis by Year

Once we computed the resiliency and eccentricity indices for each graph, we plotted the indices for each day of the year to analyze how airline network resiliency fairs over the course of a year (Figures 5, 6, 7).

Observations and Discussion
Eccentricity index plot for betweenness centrality (Figure 7) and resiliency index graph for closeness centrality (Figure 6) showed that there was less fluctuation (low variance) in index values during the summer months (day 150 to day 200). Moreover, we observed that Thanksgiving (day 330), Christmas (day 359), and New Year Eve (day 365) had significantly lower resiliency index values compared to other days in the year. This was especially visible in graphs of resiliency index for degree and closeness centrality. A similar trend was also observed for the eccentricity index.

In an effort to further understand the reasons for these sudden significant changes, we loaded in the individual graphs for those specific holidays as well as three other random days throughout the year, and compared and contrasted the properties between these graphs. Some of the results can be see in Table 3. After applying two sample t-tests to some of the metrics in the table below, we noticed that although there was no significant difference in the number of nodes between the two groups ($p = 0.284$), there was a statistically significant difference between the number of edges between the two groups ($p = 0.0138 < 0.05 = \alpha$). This means that on these holidays, there was a statistically significantly lower number of flights throughout the U.S. airline network.

6.1.2 Analysis by day of the week

In addition, we also plotted both indices (resiliency and eccentricity) by day of the week to determine if there was a difference in resiliency by day, based on patterns in flight schedules (Figure 8).

Observations and Discussion
When we aggregated the data and computed the average resiliency index by day of the week (Fig. 5), we found that there was no statistical difference between each day for all three node removal methods. This was surprising as we expected airlines to have different flight schedules on each day of the week, resulting in different resiliency indices for each day.

When we aggregated the eccentricity index by day of the week, however, we did find a difference in the resiliency of the network by day. The resilience of the airline network was significantly lower on Thursdays and Sundays and higher on Mondays, Tuesdays, and Fridays. It would be interesting to see if, for example, this variation in eccentricity index was correlated or explained by variation in flight demand by day of the week.

This variation in resilience by day of the week is also apparent in the eccentricity index by day of the year plots. The eccentricity index seems to fluctuate between higher and lower values roughly every other day.

6.2 Edge Removal

Utilizing the previous methods for computing resiliency of a given day's graph according to the Edge Removal methods outlined in 5.4.2.1 (passenger flux, ascending and descending) and 5.4.2.2 (highest betweenness cen-

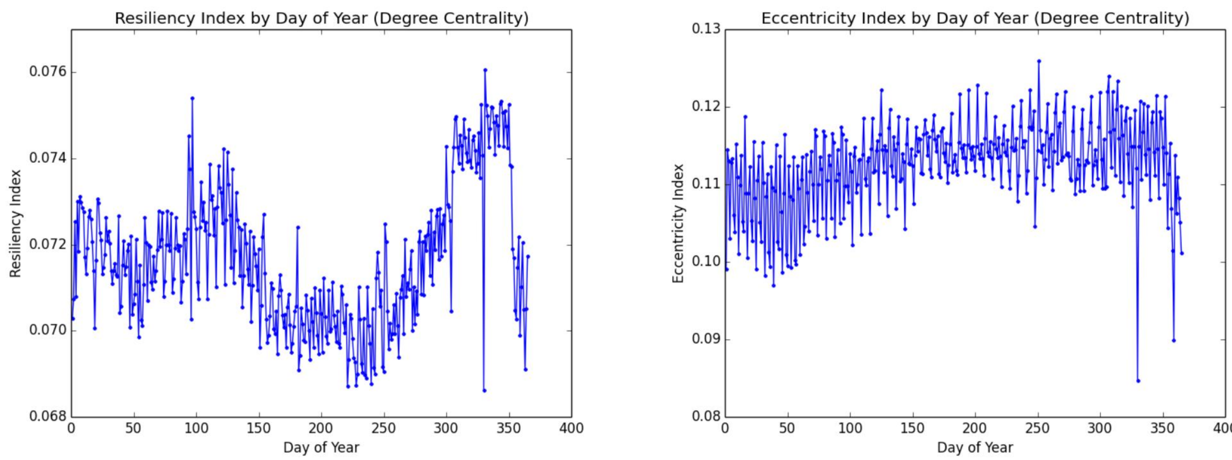

Figure 5: By Degree, Resiliency index (left), and Eccentricity index (right)

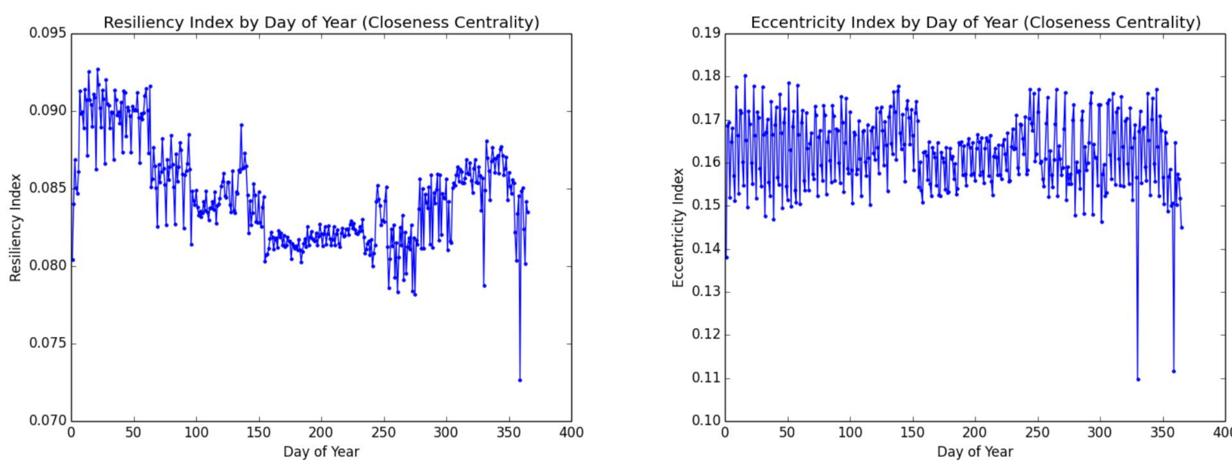

Figure 6: By Closeness, Resiliency index (left), and Eccentricity index (right)

trality edge), we plotted the corresponding resiliency values throughout 2015. The plots can be seen in Figure 9.

6.2.1 Analysis by Year

Observations and Discussion

The left most plot in Figure 9 shows the resiliency index as edges are removed in the order of highest to lowest passenger flux. Although the plot may give the first impression that there is high volatility, a closer look at the scale of the resiliency index in that plot shows that the range of the indices is low ($\sim$0.01), which implies that the observed fluctuations are due to noise as opposed to true variation. This means that removing edges sorted by high to low passenger flux does not significantly alter the resiliency indices.

On the other hand, the plot in the center does provide us with more conclusive results. It presents the resiliency index as edges are removed according to the order of lowest to highest passenger flux. The range of the resiliency index is around 0.05, which is 5 times more significant than the previous plot. We can see that in the scope of the entire year, the resiliency increases overall as the year progresses, with some volatility in between. We also observe the usual resiliency lows on Christmas and New Years Eve, but this plot also draws us to an interesting observation about the very low outlier on February 2nd, 2015 (Day 33). Upon further inspection into the circumstances of that day, there happened to be a major snow storm in the Northeastern United States which adversely affected resiliency of the entire US airport network.

The right most plot in this sequence shows the resiliency indices over the course of 2015 for when airports were iteratively removed from highest to lowest betweenness centralities. Other than the extreme low on Thanksgiving - of which we have already paid attention to from other measurements - no significant observations can be made from this plot. The resiliency indices seem consistently varying around the mean throughout the year.

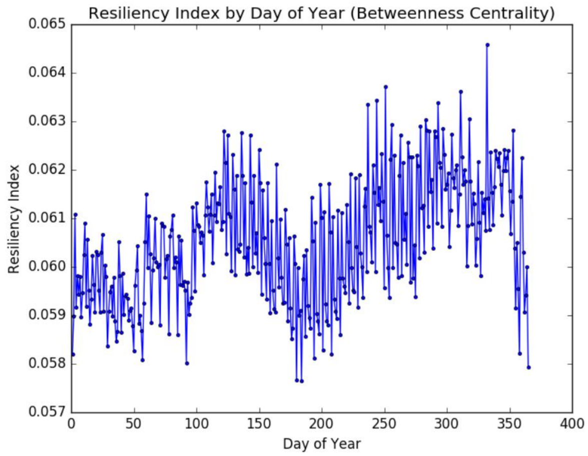

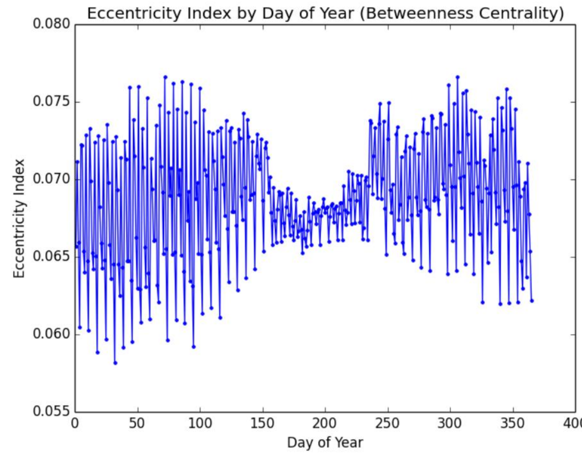

Figure 7: By Betweenness, Resiliency index (left), and Eccentricity index (right)

	Thanksgiving	Christmas	New Year's Eve	May 28	February 24	December 1
Number of Nodes	292	295	305	302	310	299
Number of Edges	10068	12633	13133	17052	15906	15823
Resiliency (Degree)	0.069	0.070	0.072	0.0715	0.0715	0.0747
Resiliency (Closeness)	0.0788	0.0727	0.0835	0.0829	0.0866	0.0843
Resiliency (Betweenness)	0.0614	0.0615	0.0579	0.0607	0.0590	0.0624

Table 3: Graph Statistics for Holidays vs. 3 Random Days of the Year (2015)

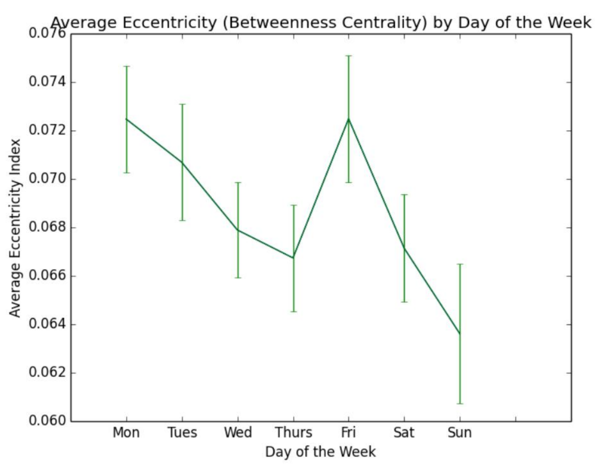

Figure 8: Average Eccentricity versus Day of the Year (for node removal by highest betweenness)

6.2.2 Analysis by Airline

Our analysis with respect to airline explored two additional questions: (1) How does the resiliency of the entire US network change with removal of a particular airline or group of airlines? and (2) What is the resiliency of an individual airline's network?

With regards to the first question, all edges pertaining to a particular airline, for example, those of Delta Airlines, were removed. This same analysis was conducted for most of the major airlines within the United States, including American Airlines, United Airlines, Southwest, etc. In addition, all flights from various combinations of these airlines were also removed. Surprisingly, in each analysis, the resiliency of the entire network was only affected by at most 1.5%, even if up to 20% of the edges of the entire network were removed from the graph. Thus, we determined that the resiliency of the entire network is not significantly affected by the removal of a single or group of a few of the major airlines.

To study the second question, we deleted all edges and nodes not pertaining to the airline being studied. Edges were subsequently removed (from the reduced netword) based on highest betweenness. Table 4 and Figure 10 below detail the resiliency index by day of the year for four major airlines within the United States: American Airlines, Delta Airlines, Southwest Airlines, and United Airlines.

Observations and Discussion

As summarized in Table 4, Southwest Airlines has the highest average resiliency index, suggesting that their network is relatively more well connected than the others. This hypothesis is confirmed by the results of Wuellner et al.[6], which outlines that Southwest's core network is composed of over half of its destinations and is "a dense

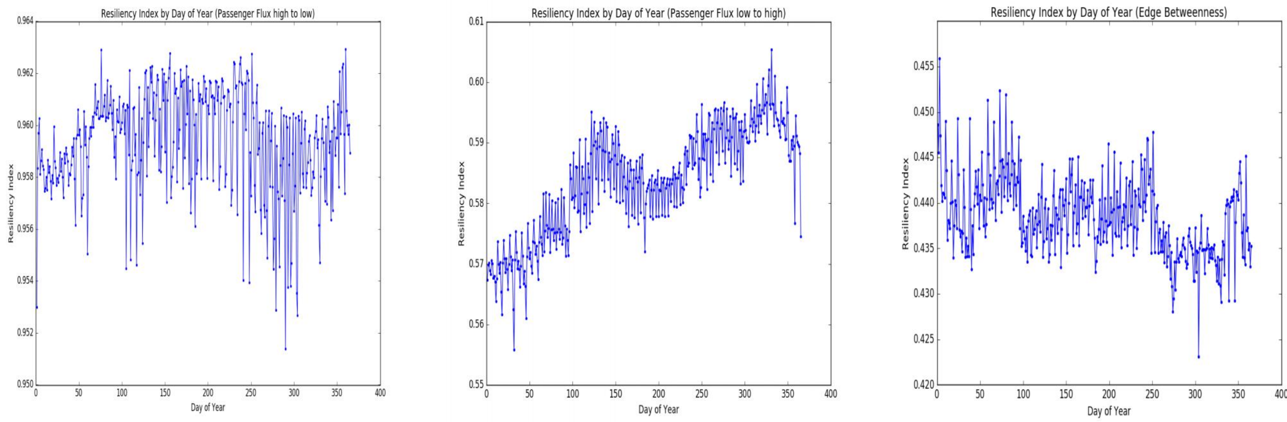

Figure 9: Resiliency index for high to low flux removal (left), Resiliency index for low to high flux removal (center), Resiliency index for edge removal by betweenness (right).

Airline	Average Resiliency Index	Average Number of Nodes in Network
American	0.538	81.9
Delta	0.447	141.1
Southwest	0.720	86
United	0.514	74.2

Table 4: Graph Statistics for Four Major Airlines within the United States (2015)

mesh of interconnected high degree airports" [6]. In addition, Southwest relies on a Point-to-Point (PP) system as opposed to a 'Hub and Spoke' (HS) network preferred by the other major airlines. When designing the infrastructure of an airline network, there are various trade offs between choosing a HS versus PP model. For example, HS models minimize the number of connecting flights passengers must take, whereas PP models minimize the overall distance traveled [6]. This difference in topology and high degree of the core nodes attributes to the higher resilience of the Southwest network.

In addition, from the perspective of a frequent flyer trying to standardize on a particular airline, understanding the resilience of a network is important in determining how flexible an airline is to cancellations or shut downs. In other words, this describes how easily can a passenger be rerouted in adversity. In a well connected graph, such as Southwest's, there are relatively more routes linking two nodes.

Additionally, from a customer perspective, knowing the breadth of destinations of the airline is also relevant. Thus, we decided to report the average number of nodes over all days of the year for each of the four airline networks. Delta Airlines has the highest average number of nodes (141.1), indicating that it services a larger number of airports. Since edges were removed based on highest betweenness centrality, the less important and more remote airports were removed from the LWCC at a later point in the edge removal process. Thus, the presence of these less important airports should not have significantly affected the average resiliency index. Another approach to measuring the resiliency index that would have accounted for this issue would have been to analyze the network consisting of only the intersection of all airports serviced by the four major airlines.

Lastly, as can be seen in the plots in Figure 9, all airline networks seem to become more resilient as the year progresses, except Southwest Airlines, which seems to maintain its resiliency throughout the year.

7 Evaluation of Model

Few papers have explored the change in resilience of airline networks over time. Verma et. al. [5] analyzed the yearly resilience of the world airline network by (1) targeted removal of edges by degree and (2) targeted removal of edges by passenger flux. To evaluate the validity of our approach, we compared our resilience plots for a single day, January 1st, (Figure 2) to their plots for the whole year. We observed that removing nodes by degree for a single day yielded similar results. When Verma et. al. removed 20% of the nodes from the World Airline Network, they observed that very few nodes remained connected. We discovered the same results for the US Airline Network (Figure 2).

We also compared our edge removal plot for January 1 (Figures 3, 4) with Verma et. al.'s work. As men-

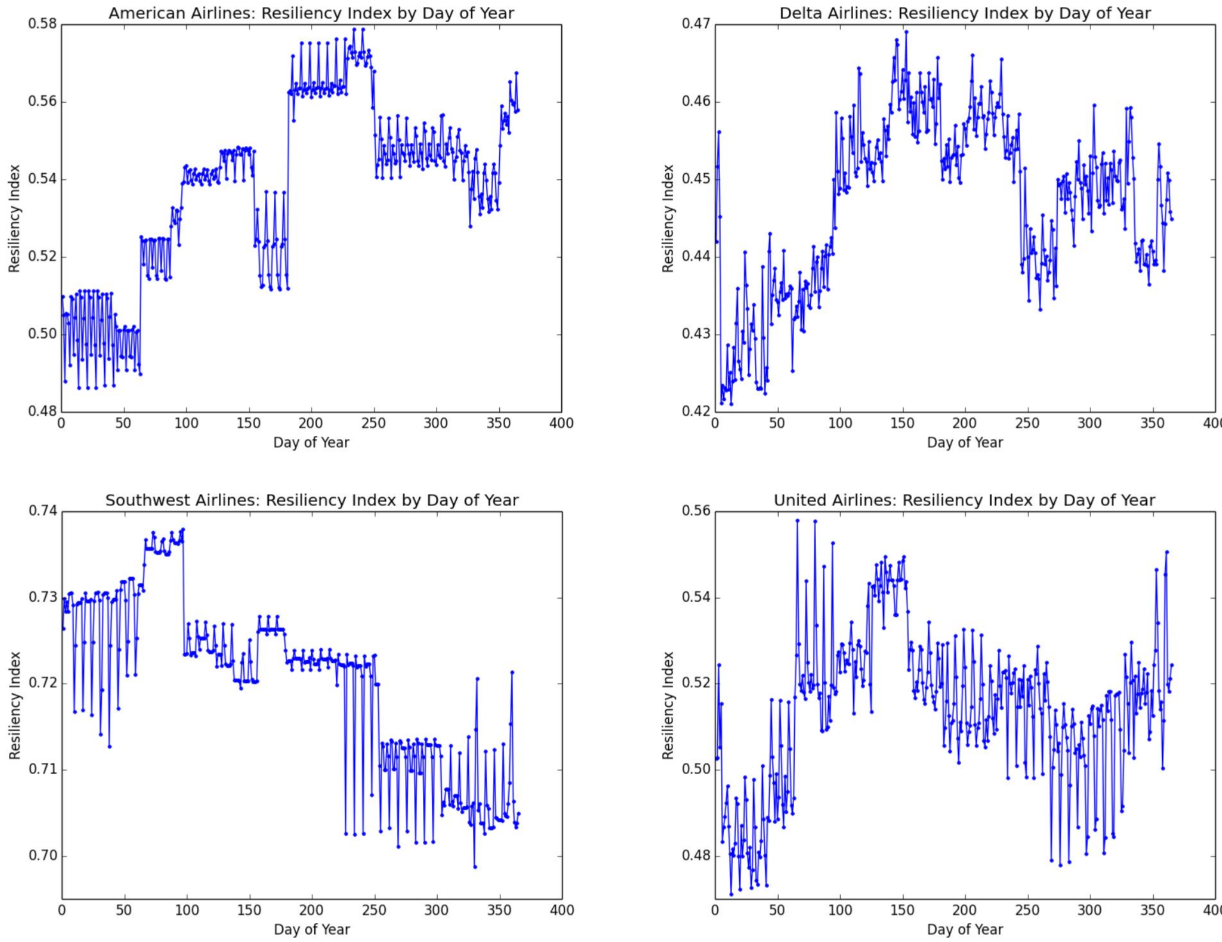

Figure 10: Resiliency comparison between the four major airline carriers

tioned previously, we noticed that removing edges with a high passenger flux first did not affect the network that much. This corroborated the findings of Verma et al [5]. However, the drop in the fraction of nodes in LWCC was not as drastic for our network. Verma et. al. observed that removing 40% of edges with the highest flux resulted in 25% drop in nodes connected. Whereas, we observed that removing 40% of edges resulted in no drop in connectivity. We hypothesized that this was a result of the differences in the networks being analyzed (US Airline Network vs World Airline Network).

Wuellner et. al. [6] used node removal by betweenness to measure the resilience of specific airline carrier networks. They observed that over the course of the year, Southwest Airlines' network was the most resilient. We also observed this for the two different removal measures we used: removal by passenger flux and removal by edge betweenness.

While our results reflect expected behavior, they do not necessarily suggest that our approach is ideal. Resiliency index, eccentricity index, and the various removal measures do not take into account factors like the difference in total number of nodes and edges between graphs of different days.

8 Conclusion

In this paper, we propose a unique temporal assessment of resiliency within the US airline network over the course of 2015. We defined our own resiliency and eccentricity indices to quantify its total resiliency for a given day. We used these measures to identify the days of the year on which the network is least resilient as well as identifying which US airline carrier's network is most resilient to adversity.

When we conducted targeted node removal, we discovered that the network was less resilient on both Thanksgiving and New Year's Eve. There were also no significant differences in the resiliency index between the days of the week. However, we did notice significant differences in the eccentricity index.

When we conducted edge removal, we observed that the network was least resilient on days such as February 2, 2015 on which flights were disrupted due to a snow storm in the Northeast. In addition, our results indicated that Southwest is the airline with the highest resiliency network. Therefore, a frequent flyer looking to minimize disruption of their flights should most likely pick Southwest Airlines over any of the other three airlines in our analysis. However, if their priority is an airline with a wider breadth of destination airports, the customers should standardize on Delta Airlines.

9 Future Work

With additional time, there are a few more analyses that we would perform on the network that would provide us with additional observations. More precisely, we would have liked to compute the eccentricity index under targeted edge removal and see how that affects the total resiliency of the network.

We would also like to further analyze the structure of each individual airline's network by examining parameters such as the average number of flights per day. The average degree of each node in en airline's network would provide us with metrics of how important a given airport is to that airline. Within the analysis of the different airlines' networks, it would also be interesting to pairwise analyze two networks when the problem is constrained to the set of intersecting airports between the two networks. This would eliminate the bias in the comparison of two airlines in which one of the airlines has more destination airports in its repertoire than the other airline.

References

[1] `https://www.faa.gov/airports/planning_capacity/passenger_allcargo_stats/passenger/`. Accessed: 2016-11-18.

[2] Airline on-time performance database. `http://www.transtats.bts.gov/DL_SelectFields.asp?Table_ID=236`. Accessed: 2016-10-19.

[3] Pablo Fleurquin, José J. Ramasco, and Victor M. Eguiluz. Systemic delay propagation in the us airport network. *Scientific Reports*, 3, 01 2013.

[4] R. Guimerá and L. A. N. Amaral. Modeling the world-wide airport network. *The European Physical Journal B*, 38(2):381–385, 2004.

[5] T. Verma, N. A. M. Araújo, and H. J. Herrmann. Revealing the structure of the world airline network. *Scientific Reports*, 4, Jul 2014. Article.

[6] Daniel R. Wuellner, Soumen Roy, and Raissa M. D'Souza. Resilience and rewiring of the passenger airline networks in the united states. *Phys. Rev. E*, 82:056101, Nov 2010.

10 Collaboration

We all contributed equally to the project. We used GitHub to work on the coding parts and ShareLatex to work on the writing requirements. Contact info:

1. Neel Rakholia: neelvr@stanford.edu
2. Poorvi Bhargava: poorvib@stanford.edu
3. Noam Habot: nhabot@stanford.edu