



A Wavelet Analysis Approach for Categorizing Air Traffic Behavior

Michael Drew*

*University of California, Santa Cruz at NASA Ames Research Center,
 Moffett Field, CA 94035-0001, USA*

and

Kapil Sheth[†]

NASA Ames Research Center, Moffett Field, CA 94035-0001, USA

In this paper two frequency domain techniques are applied to air traffic analysis. The Continuous Wavelet Transform (CWT), like the Fourier Transform, is shown to identify changes in historical traffic patterns caused by Traffic Management Initiatives (TMIs) and weather with the added benefit of detecting when in time those changes take place. Next, with the expectation that it could detect anomalies in the network and indicate the extent to which they affect traffic flows, the Spectral Graph Wavelet Transform (SGWT) is applied to a center based graph model of air traffic. When applied to simulations based on historical flight plans, it identified the traffic flows between centers that have the greatest impact on either neighboring flows, or flows between centers many centers away. Like the CWT, however, it can be difficult to interpret SGWT results and relate them to simulations where major TMIs are implemented, and more research may be warranted in this area. These frequency analysis techniques can detect off-nominal air traffic behavior, but due to the nature of air traffic time series data, so far they prove difficult to apply in a way that provides significant insight or specific identification of traffic patterns.

I. Introduction

WHEN characterizing the state of air traffic in the National Airspace System (NAS) for historical or real-time analyses, most data are evaluated in the time domain. These data are useful for understanding and analyzing the performance of the NAS, and are the basis for the majority of the work in this field. Several methods of modeling, predicting, and optimizing the performance of the NAS using these metrics have been proposed. For example, Sridhar et al. in Ref. 1 construct a linear time-varying model of aggregate traffic flow using flight counts as the state variable. In Ref. 2, an autoregressive model for predicting sector demand is presented that uses historical data and takes convective weather conditions into account.

Recent work by Drew and Sheth in Ref. 3 demonstrates that insight into NAS operations can be found by transforming time domain traffic count data into the frequency domain. Specific traffic flows (e.g flights to and from a major regional airport) are shown to exhibit a periodic signature that can be readily identified in the frequency domain by using the Discrete Fourier Transform (DFT). Furthermore, some disruptions in regional traffic counts (like those caused by weather, or by controller operations) are shown to be more clearly identified in the frequency domain than in the time domain.

One of the drawbacks to the DFT is that once transformed into the frequency domain, all notion of time is lost. Thus, if a disruption is noticed in the frequency domain data, it is impossible to determine when the disruption occurred. The desire to resolve data into both its frequency and time domain content is what gives rise to wavelet analysis. Though it is impossible to have perfect resolution in both domains

*Research Analyst, Systems Modeling and Optimization Branch, MS 210-8; michael.c.drew@nasa.gov.

[†]Aerospace Research Engineer, Systems Modeling and Optimization Branch, MS 210-15, AIAA Associate Fellow; kapil.sheth@nasa.gov.

simultaneously, the wavelet analysis provides a notion of both the frequency and time domain content of a signal.

Primarily a tool used for signal filtering and compression, wavelet analysis has also found uses in other domains. In Ref. 4 for example, wavelets are used to analyze the audio signal of a mechanical watch. By matching known modal frequencies provided by a finite element analysis to the 2D time and frequency results of the wavelet analysis, malfunctions of the watch's movement can be detected. In Ref. 5 the authors present a guide for analyzing time series data by performing a wavelet analysis on ocean temperature data. In Ref. 6, the authors apply wavelets to the problem of detecting disease outbreaks from time series disease counts data sets. In Refs. 7,8,9 the authors apply wavelets to graph models of data and road traffic analysis, and wavelets are found to be capable of characterizing the network performance and of identifying failures in traffic links.

In order to better analyze NAS traffic and identify disruptions (potentially prior to becoming difficult to manage) the authors have begun investigating the use of wavelet analysis on air traffic data. In this paper two techniques of wavelet analysis are performed. The first approach is to apply the wavelet transform in a traditional fashion to airspace traffic count time series data. The second approach is to explore the use of wavelets on a network graph model of the national airspace. Simulations of several Traffic Management Initiatives (TMIs) are analyzed using these methods and compared to their nominal traffic flow behavior.

This paper is organized as follows: Section II begins with a brief description of both the Continuous Wavelet Transform (CWT), and the Spectral Graph Wavelet Transform (SGWT). Section III contains the results of the CWT analysis, and Section IV contains those of the SGWT analysis along with a description of the center based graph network traffic model. Concluding remarks are in Section V.

II. Background

Much has been written about wavelet transforms in all their varied forms and applications. In this section, only a broad conceptual overview is presented to familiarize the reader with the general idea of wavelet transforms and their application in this context.

A. Continuous Wavelet Transform

Unlike the Fourier transform, which correlates a signal expressed in time (time series data) with infinitely long sines and cosines, the wavelet transform correlates the time domain signal with scaled and shifted versions of a 'mother wavelet' function $\psi(t)$. There are many choices for a mother function, but certain conditions must be met. For instance, it must have a zero mean and be localized in time and frequency. An example is the Morlet wavelet shown in figure 1.

The Continuous Wavelet Transform (CWT), which is used in this work, computes the inner product of a signal $x(t)$ with the scaled and shifted version of the mother wavelet. It is defined as follows:⁴

$$C(s, \tau) = \int_{-\infty}^{\infty} x(t) \frac{1}{\sqrt{s}} \psi^* \left(\frac{t - \tau}{s} \right) dt, \quad (1)$$

where $\psi^*(t)$ is the complex conjugate of the mother wavelet, s is the scaling factor, and τ is the shift in time. The shift in time τ provides temporal resolution, while the shift in scale s provides an approximation of frequency resolution. High frequency content within the signal will have high correlation with small scaling factors, whereas large scaling factors, which stretch the mother wavelet, will capture low frequency events in $x(t)$. In practice, the scale term can be converted into 'psuedo-frequency'—or, as presented herein for ease of interpretation, 'psuedo-period'.

An example of the CWT applied to air traffic is shown in figure 2. Here, one day of air traffic counts in Cleveland Air Route Traffic Control Center (ARTCC) is shown in black and the CWT coefficients are




Figure 1: The Morlet mother wavelet function

shown as colored contours. In this plot, the Morlet wavelet is used. The warmer and cooler colors represent peaks and valleys respectively at their indicated periodicity. For instance, as found in Ref. 3, flights within Cleveland Center going in and out of Detroit Metropolitan Wayne County Airport are responsible for a periodicity in the data of about 110 minutes. This is validated by the CWT with red and blue contours shown centered at 110 minutes. Unlike the Fourier transform results in Ref. 3, the CWT plot has the added benefit of clearly displaying when that periodicity is present—in this case it begins at about 12:00 UTC and persists until about 21:00 UTC where a trend of higher periodicity emerges.

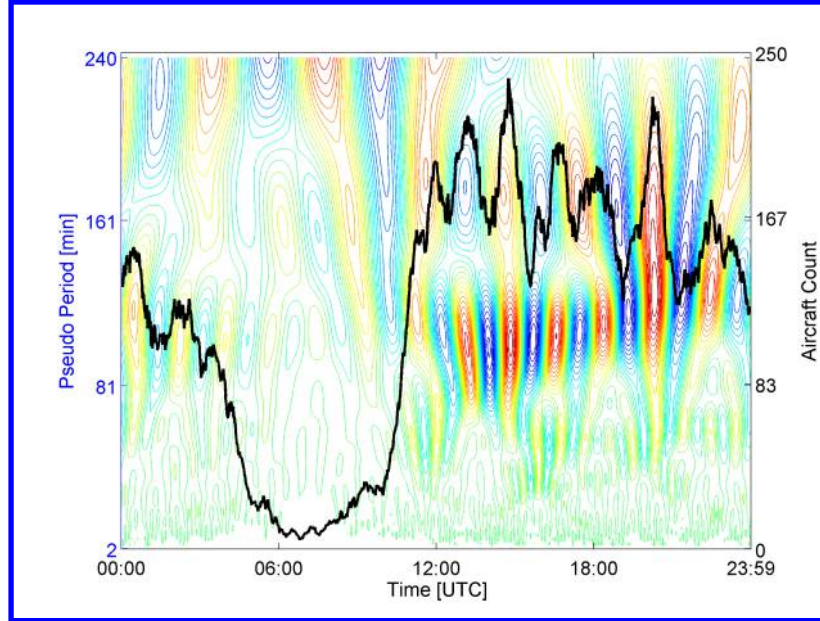


Figure 2: Aircraft count in Cleveland Air Route Traffic Control Center (ZOB), and contours of the corresponding continuous wavelet transform.

B. Spectral Graph Wavelet Transform

Recently, work has been done in applying traditional wavelet techniques, which operate on a temporal data signal $x(t)$, to the signals on a network line graph model. The motivation is that the so-called Spectral Line Graph Wavelet Transform could reveal and detect the impact of disruptive traffic events in the context of its impact on the overall traffic network. Formal derivations of the SGWT can be found in Ref. 8,9. Here, only a conceptual overview is presented.

Given a weighted connected graph $G = \{V, E, w\}$ with a finite set of vertices V , $|V| = N$, a set of edges E , and weight function w that resides on the edges, an $N \times N$ adjacency matrix A can be defined with entries $a_{m,n}$ such that

$$a_{m,n} = \begin{cases} w(e), & \text{if } e \in E \text{ connects vertices } m \text{ and } n \\ 0, & \text{otherwise} \end{cases}$$

The edge weights in this work are set to one, however, less trivial weighting models could be easily employed.

At each vertex of the graph resides a graph signal $f : V \rightarrow \mathbb{R}^N$, consisting of sampled data associated with that network node. For example, the m^{th} component of f could represent the traffic volume (or speed, etc.) at node m .

For any simple graph G , the Laplacian matrix L can be defined as $L = D - A$, where D is the degree matrix having only diagonal elements d_m equal to the sum of weights of all edges connected to vertex m . A is the adjacency matrix defined above. The Laplacian matrix has many useful properties in graph theory.¹⁰ The eigenvalues and eigenvectors of L found by

$$L \chi_l = \lambda_l \chi_l$$

for $l = 0, 1, \dots, N - 1$ produce eigenvectors χ_l . Noting the analogy with the exponential eigenfunctions of the classical Fourier transform, it can be shown that the graph Fourier transform \hat{f} of the function f on the vertices of G can be defined as:⁹

$$\hat{f}(l) = \sum_{n=1}^N \chi_l^*(n) f(n). \quad (2)$$

The graph wavelet coefficients defined at each vertex n and scale s are then found with

$$W_f(s, n) = \sum_{l=0}^{N-1} g(s\lambda_l) \hat{f}(l) \chi_l(n). \quad (3)$$

Here g is a wavelet generating kernel analogous to the mother wavelet function found in the CWT. Recall that for the CWT, the mother wavelet provides localization in time and frequency depending on how it is stretched and scaled. The SGWT with its kernel function g provides localization in frequency (of signals whose domain is the graph vertices—not time), and graph space. Therefore, as the mother wavelet of the CWT is selected to provide the best resolution tailored for the application, the wavelet generating kernel g must also be selected to provide an appropriate trade-off between localization in the frequency and graph space domains. For this work the wavelet generating kernel function is a cubic spline with parameters matching those in the examples of Ref. 9 so that sufficient resolution on the vertices of the graph can be achieved. It must be pointed out that despite the discrete nature of the graph vertex domain, g is a continuous function, and s can be any positive real number. The domain of g must span the spectral range of the graph (given by the largest eigenvalue of L). Furthermore, similar to how the mother wavelet of the CWT is scaled to find correlation with a signal at different pseudo-frequencies, the SGWT kernel function g is scaled by s to find correlation with the signal residing on the graph vertices. An example of this function is shown in figure 3 for various scale values.




Figure 3: Wavelet generating kernel function $g(s \cdot x)$ with various scale values.

To be clear, the SGWT is very different from the CWT, which produces a view of time series data revealing the temporal location of its frequency components. The independent variables s (for periodicity) and τ (for time shift) of Eq. (1) allow the creation of plots like those of figure 2. In contrast, note that the SGWT equation of Eq. (3) has independent variables s for network scale, and n for vertex location. There is no knowledge of time (though successive calculations can be done for each time sample of the data), and here scale is related to how events are correlated across the network. Rather than identifying the periodicity of a signal as it moves through the time domain, the SGWT identifies the periodicity of a signal as it exists on one graph vertex relative to another. Lower scale values capture how events in the network vary over large vertex ranges (low frequency, or high periodicity within the network). Higher scale values (high frequency) highlight events that have a greater correlation with their more immediate neighbors in the network. An example plot of SGWT coefficients is shown in figure 4 based on one day of NAS Center graph network traffic. Results




Figure 4: SGWT results of NAS network for August 22, 2014. Scale $s = 0.78$.

for individual vertices are shown as horizontal rows. Note that though the SGWT does not operate on time-domain signals, it can be computed sequentially on the network for each sample in time. Thus, for this plot it was calculated 1440 times using a fixed scale of $s = 0.78$. Again, warmer colors indicate higher values and cooler colors lower. A more thorough discussion of how to interpret SGWT plots will be presented in Section IV.

III. Continuous Wavelet Transform Analysis Results

The results of this work are divided into two parts. In this section a traditional wavelet analysis is performed on air traffic data using the CWT. This begins by recording historical sector counts sampled every minute using the Future Air Traffic Management (ATM) Concepts Evaluation Tool (FACET).¹¹ Specific air traffic streams, like flights of a particular altitude range, can be filtered out if desired and recorded separately. Simulations of various Traffic Management Initiatives (TMIs) can also be conducted in FACET and compared to simulated baseline traffic in order to observe their effects in the wavelet results. The CWT contour plots are created using the ‘cwtft’ function of MATLAB’s wavelet toolbox.¹² This function uses an FFT algorithm to calculate the CWT coefficients using the desired mother wavelet.

A. Wavelet Choice

As mentioned, computing the CWT involves a mother wavelet with which the time series data is compared. Depending on the wavelet chosen, the resulting wavelet plot will exhibit a trade off between time and frequency resolution. Thus, the first step in the analysis is to choose the appropriate mother wavelet for the type of analysis involved. Wavelet analysis seeks a compromise between time and frequency resolution, and some wavelets can be biased towards one of those domains. In figure 5, two different wavelets are applied to the same time domain signal. In figure 5a, the Morlet-12 wavelet is used. Here 12 is the base frequency of the mother wavelet. With a high base frequency, it is broad in the frequency domain, but narrow in the time domain. In figure 5b, the DOG-2 (Derivative of Gaussian) wavelet is used. It has broad time domain signal, but narrow in frequency. Thus, the periodic signature at 110-minutes is sharply resolved in figure 5a in terms of its periodicity, but it is not as clear when in time it starts to show up in the signal. (Some time between 8:00 and 12:00 UTC.) In figure 5b, the DOG-2 wavelet shows poor resolution in frequency, since it is difficult to pinpoint the periodicity of the same signature. On the other hand, it is clear that it does not start before 12:00 UTC. A wavelet that strikes a good compromise between time and frequency domain resolution for air traffic is the Morlet wavelet with the standard base frequency of 6 (Morlet-6). The results of this wavelet applied to the same traffic data are shown in figure 2, which depicts much sharper warm and cold colors.

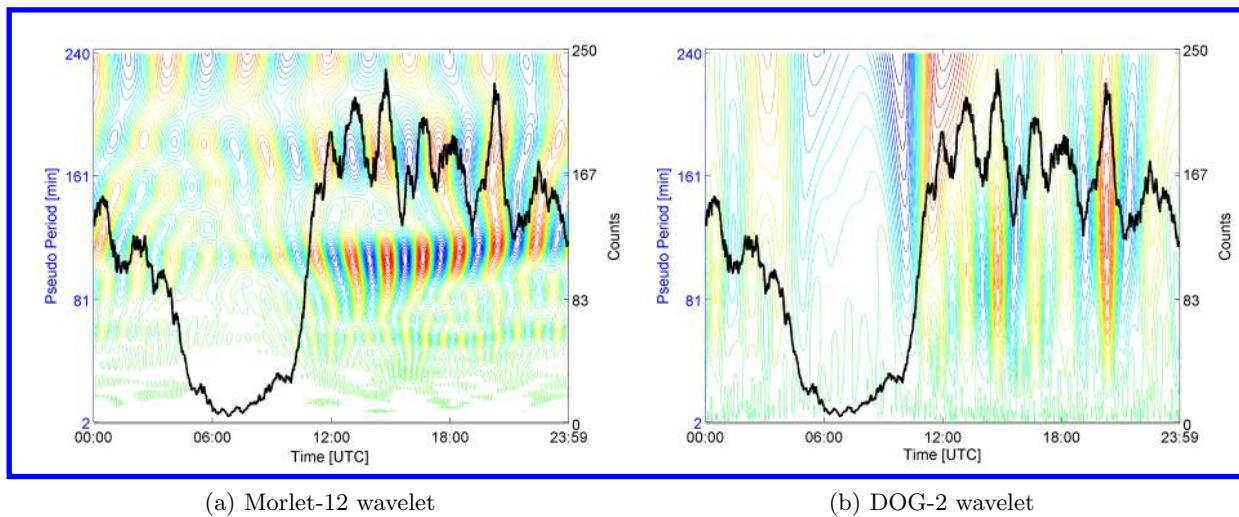


Figure 5: A comparison of wavelet analysis between two wavelets on the same air traffic data.

B. Historical Traffic Analysis

Noting that known traffic flow signatures (like the Detroit Metropolitan Wayne County Airport flights in figure 2) can be observed with wavelet analysis, the next step is to apply it to several days of air traffic to see if off-nominal traffic patterns can be detected. Cleveland Center (ZOB) is one of the most congested ARTCCs in the country. Focusing on flights above flight level 250 in ZOB during July 2014, the wavelet coefficients for a 24-hour period of air traffic counts are calculated for each day. Then, seven mean days of coefficients are calculated for each day of the week from the month. Figure 12 in the Appendix shows the difference between every July day's wavelet coefficients and the mean coefficients computed for that day of the week. Select dates are shown in figure 6 below. For example, note the dark red and blue contours shown toward the end of July 3 and the beginning of July 4 shown in figures 6a and 6b. This indicates that the wavelet coefficients for these two days are significantly different from those of the average Thursday and Friday of this month. Thus, plots with more reds and blues (representing high peaks and low valleys) indicate possible off-nominal traffic conditions.

Figure 6 demonstrates that the wavelet analysis highlights differences in traffic flow characteristics that are difficult to detect in the original time domain data. Several days stand out as being significantly different in the wavelet domain. During the July 3-4 time period there was a large system of convective weather that moved through the Cleveland and New York regions. Similarly, July 14 (figure 6c) also shows a large difference from mean toward the end of the day. Weather records show that there were severe thunderstorms west of Cleveland Center near Chicago, as well as within the center itself. Operational data indicate that due to heavy volume, the playbook route CAN-BRAVO-EAST¹³ was put into effect from 17:00 UTC to 2:00 UTC the following day. July 23, which also exhibits major wavelet domain contrast in figure 6d, also witnessed severe convective weather in the Midwest and the Northeast, resulting in a CHICA-ROUTE-OUT playbook to be put in place from 17:00 to 23:59 UTC. Again, the wavelet analysis detects this off-nominal traffic by a signature in the frequency domain. Unlike the Fourier analysis, however, it also indicates the time at which this traffic anomaly took place.

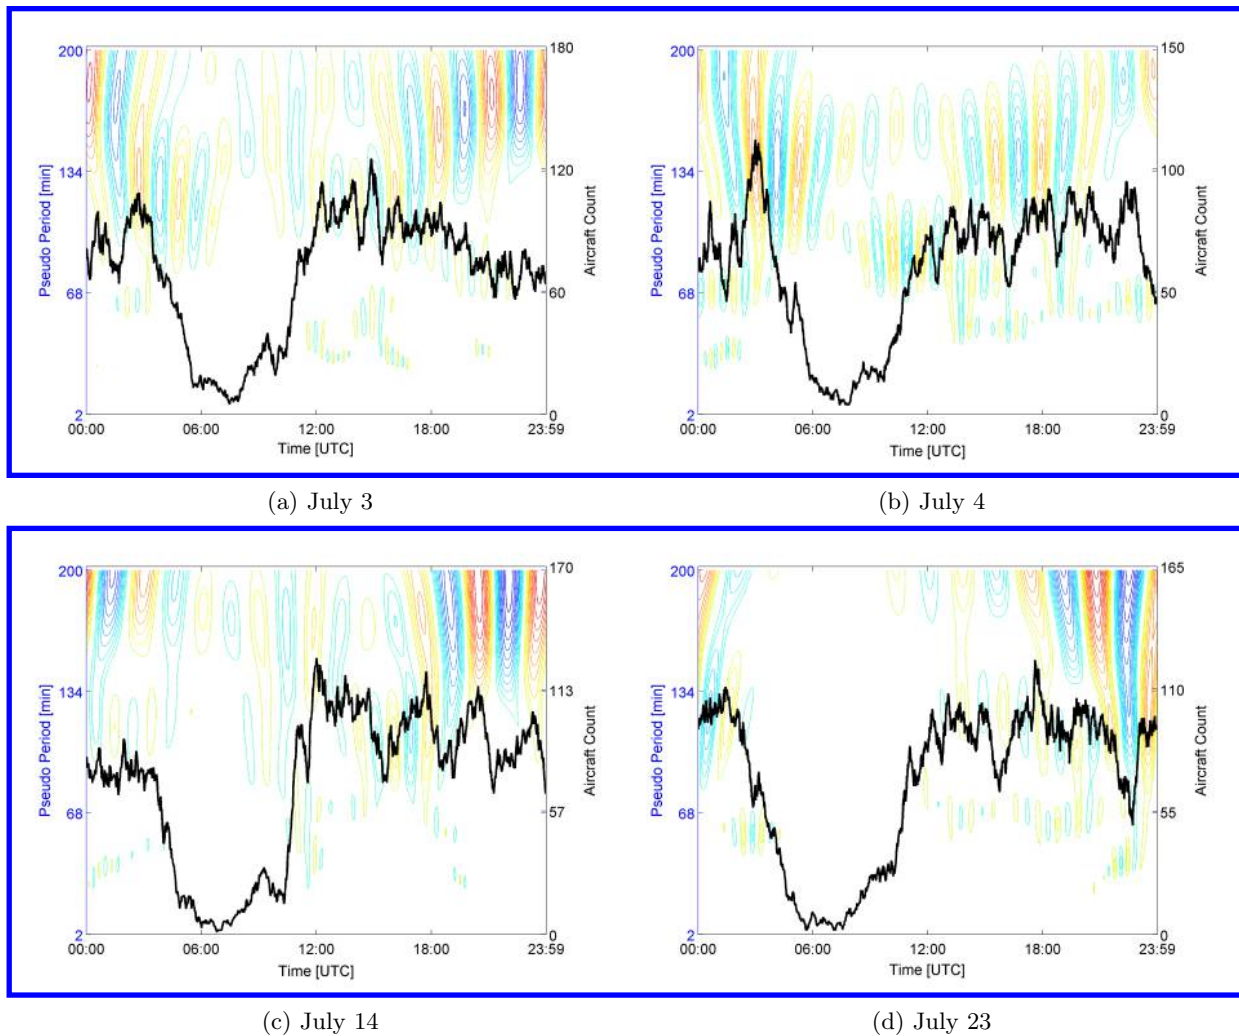
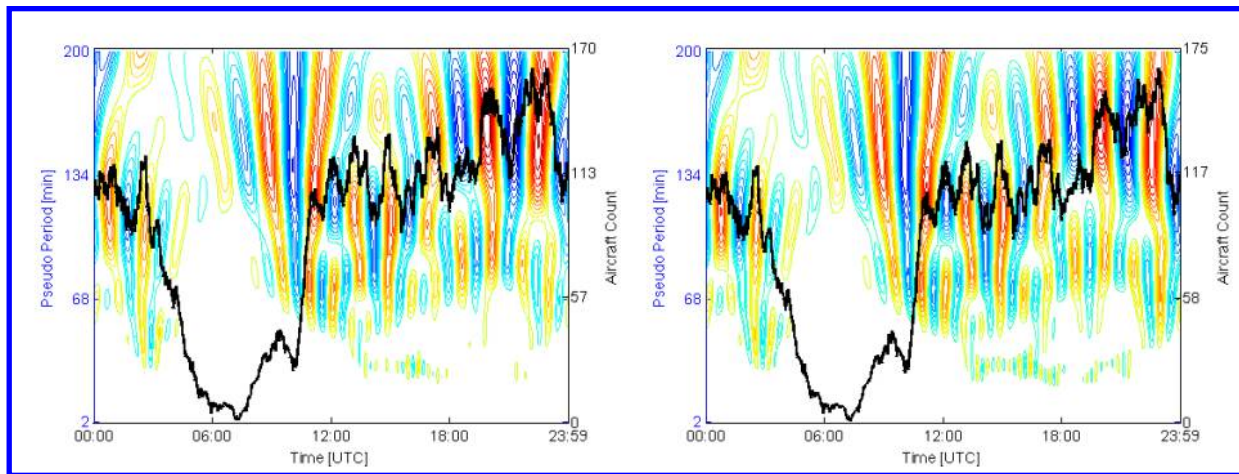


Figure 6: Traffic counts and wavelet analysis for flights above FL250 in Cleveland Center on select July 2014 dates. Wavelet contours shown are the difference between that day and the mean of that day of the week's wavelet contours. See figure 12 for the entire month.

C. Simulated TMI Detection

In figure 12, July 17 is shown as having high volume but near-nominal behavior in terms of the CWT analysis (indicated by the lack of dark red and blue contours). Hoping that the CWT could detect the presence of a single TMI within a center, two simulations of July 17, 2014 historical traffic were conducted in FACET, and the CWT was applied to the resulting ZOB traffic counts. In the first simulation, traffic was flown without any TMIs applied. In the second, the Green Bay (GRB) playbook route was implemented from 12:00 - 20:00 UTC. This playbook routes east-bound flights into Cleveland Center. Also, 30 Miles in Trail (MIT) metering was applied through the DJB fix (near Cleveland). The resulting counts and contours for both simulations are shown in figure 7. Note that these plots are displaying the CWT contours as directly applied to the counts. They are not the difference from a mean day's contours as in figure 12.

In the previous subsection, the CWT was able to detect off-nominal historical traffic due to a combination of factors involving weather and various TMIs. However, it appears difficult to observe a difference in the CWT results between figure 7a and figure 7b. Thus, it seems that this application of the CWT is not capable of detecting a single TMI event, and it is not yet clear how to distinguish or identify one TMI or traffic pattern from another. Individual traffic streams, like those traveling to and from a specific airport, can occasionally be identified due to their operational patterns (as was seen in figure 2). However, historical and



(a) Baseline, no TMIs implemented.

(b) Green Bay playbook implemented from 12:00 - 20:00 UTC with 30 MIT.

Figure 7: Aircraft count and CWT contours for flights above FL 250 in ZOB. Simulated traffic based on historical flights from July 17, 2014.

simulated analyses like these have yet to reveal a reliable signature for many traffic flows—either nominal or off-nominal. This may not be possible due to the relatively small data volume and finite nature of traffic time series data. Hence, the focus was moved to analyzing traffic behavior with the SGWT method described above to look at reroute disturbances.

IV. Spectral Graph Wavelet Transform Analysis

Using a network model as a traffic network realization is not a new technique. For the purposes of this work, a network graph of continental Air Route Traffic Control Centers is created such that the vertices of the graph are the centers, with edges connecting each center vertex to all contiguous neighboring center vertices. That is, each pair of control centers that share a boundary are connected with a graph edge. This graph network is depicted in figure 8, where vertices are shown with black circles and edges are shown in blue. Note that neither the locations of the vertices, nor the edge lengths as shown is indicative of weights or any other mathematical attribute of the model. The vertices are shown located at the centroid of each control center, but this is only for convenient graphical depiction. All edge weights are set to one in this work. This graph has 20 vertices and 45 edges.

Both historical and simulated traffic can be processed and converted into this network flow model by recording the number of flights that are traveling from one center to another at each time sample (1-minute). This is recorded for all 45 edges. For instance, if there are 15 flights traveling from Chicago Center to Cleveland Center, and 12 flights traveling in the reverse direction, the value of the signal recorded at the edge connecting both centers will be 27 at that time step. Flights that are departing and arriving within the same center, as well as flights transitioning to international airspace are not counted.

The traffic flow data that exist in this network graph model resides on the edge of the graph, but the SGWT analysis deals with signals on the vertices f . This leads to defining a ‘line graph’ based on the original graph model. In graph theory, given a graph G , each vertex of the line graph $L(G)$ represents an edge of G , and the line graph edges exist between those vertices of $L(G)$ if the edges they represent in G share a common vertex in G .¹⁴ The resulting line graph edges of the NAS center network graph are shown in figure 8 in dark red, with vertices shown as green triangles at the midpoints of each of the original graph’s edges. It consists of 45 vertices and 176 edges. The line graph is only used for computational purposes in the SGWT. For the remainder of this paper, the original center network graph with traffic flow along its edges is an acceptable way of visualizing and describing the model.

The first step in exploring this analysis is to run a historical simulation in FACET. Historically-scheduled flights are flown according to their filed flight plans without any controller intervention or changes due to

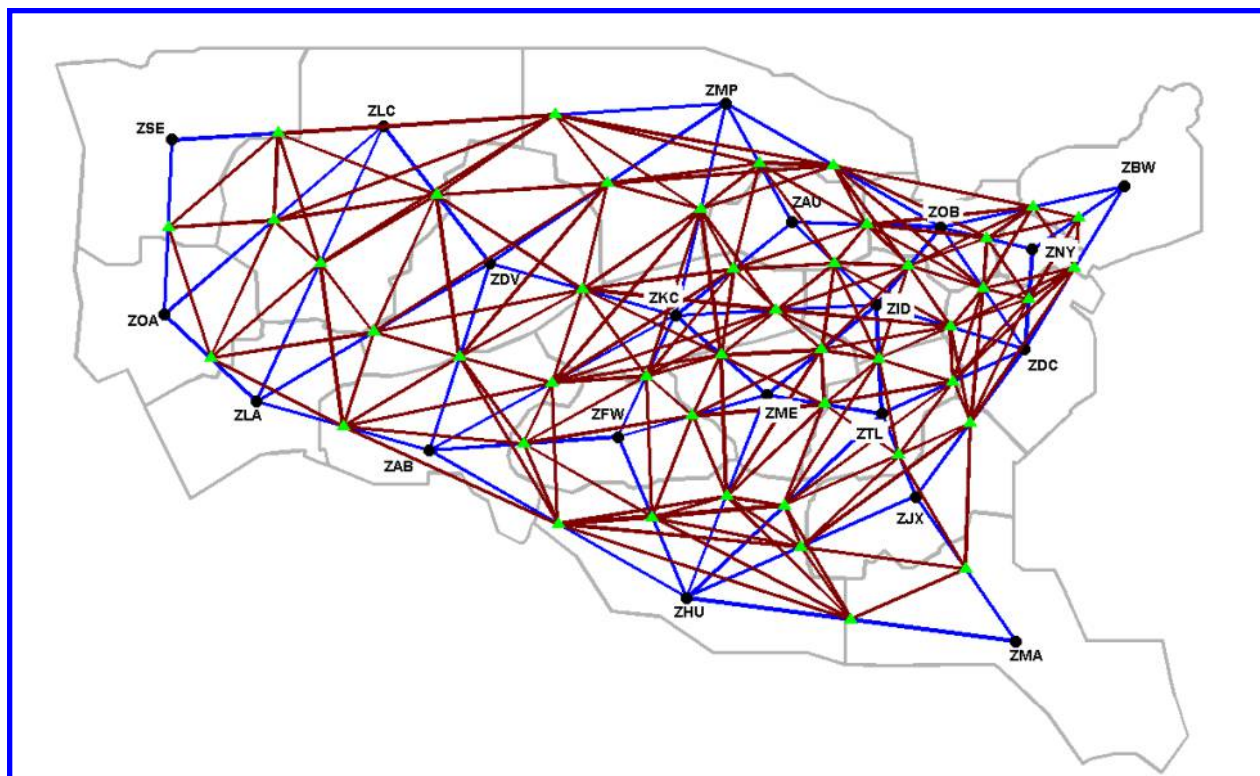


Figure 8: NAS Center network graph G shown in blue, and line graph $L(G)$ shown in green.

weather, etc. Choosing August 22, 2014, which was a high volume low delay day, the recorded traffic flow is converted into the network model structure. Then the SGWT coefficients are calculated for a range of scale values for all 45 center-to-center edges. This is performed using MATLAB¹² and the SGWT Toolbox from Ref. 15. The results are shown in figure 13 in the Appendix. It can be difficult to interpret these results, but from figure 13, it is clear that the network has the highest coefficients for a scale range at roughly 0.08 to 1.0 due to the warmer colors distributed in this range of the plots.

Recall that in this context the range of scale values is based on the Laplacian matrix of the graph, with larger values related to the high frequency (on the vertex domain) content of the signals across the network, and vice versa for smaller scale values. For example, if a warm color is depicted at a low scale value (low frequency = long period) at a specific link location, it indicates that the traffic at that link has a strong cause and effect relation to traffic at a relatively long distance across the network. Disruptions in traffic on such links would likely have a large effect on traffic a long distance away. On the other hand, if a link shows warm colors for higher scale values, its traffic is more strongly correlated to closer neighboring links in the network. Disruptions on those links would likely affect closer neighboring links.

A more illustrative view of these results can be found by calculating the coefficients at a specific scale and plotting the results at their corresponding edges in the center network. Such a depiction is shown in figure 9, which shows the SGWT results for a scale value of 0.08. Each color band shown on every graph edge represents the 24-hour period of August 22, and can be thought of as representing a slice at 0.08 through each of the plots in figure 13. Accordingly, the colors tend to be warmer in the first quarter and last half of each band where volume is highest. At the low scale value of 0.08, the results depict the correlation between edges that are farther away in the network (low frequency across the network). In other words, warmer colors on links in this plot indicate flows that have a stronger cause/effect relation with traffic at a farther distance across the network. Figure 9 reveals that flows with the greatest long-distance impact include the northerly flows of Salt Lake (ZOB) \leftrightarrow Denver (ZDV), Denver (ZDV) \leftrightarrow Minneapolis (ZMP), and Minneapolis (ZMP) \leftrightarrow Chicago (ZAU). Of note on the East Coast, is the New York (ZNY) \leftrightarrow Washington DC (ZDC) flow, which clearly has a greater correlation to faraway flows than other edges in the region.

Figure 10 shows the same analysis at a scale of 1.0, which depicts a more localized correlation of traffic flows. Warmer colored bands indicate times and links where traffic is better correlated to closely neigh-

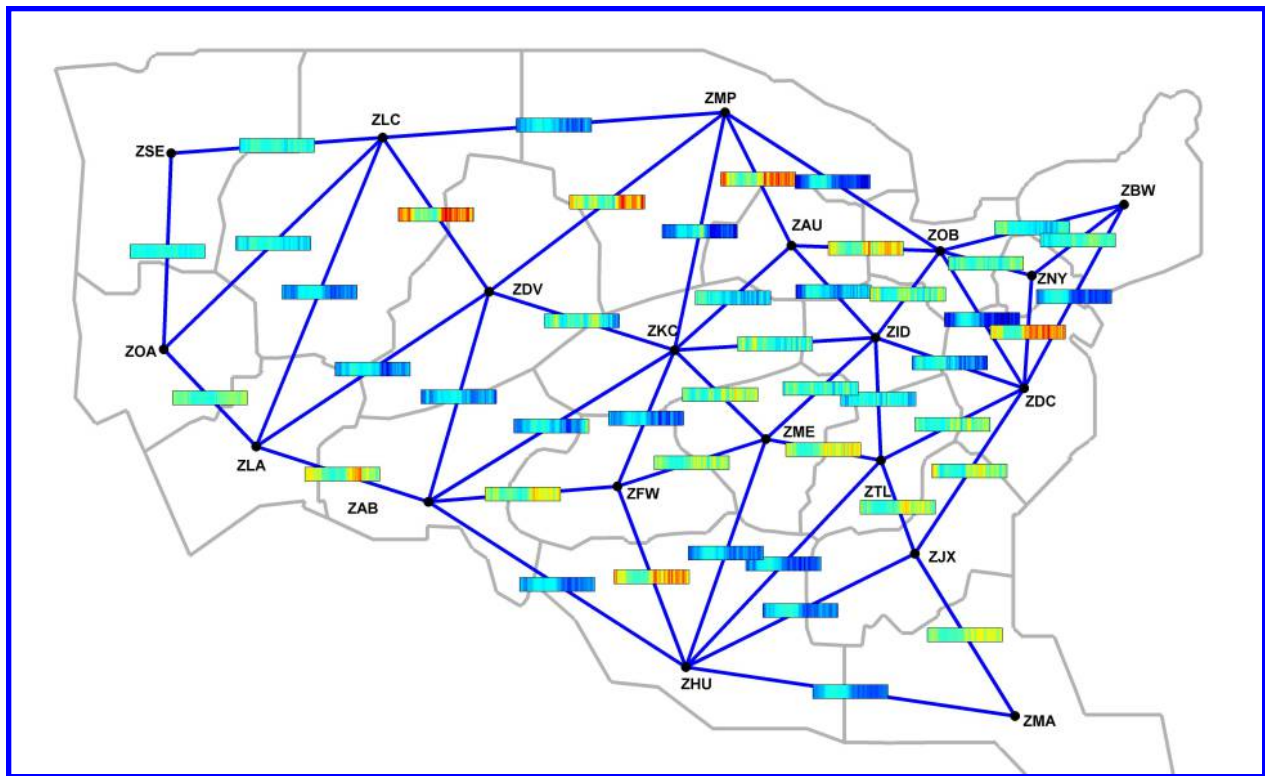


Figure 9: SGWT results for simulated traffic based on flights of August 22, 2014. Scale = 0.08

boring links. Here, most of the highlighted edges from figure 9 exhibit a more muted correlation to closely neighboring edges, except for the New York (ZNY) \leftrightarrow Washington DC (ZDC) flow. Also, the Miami (ZMA) \leftrightarrow Jacksonville (ZJX), and New York (ZNY) \leftrightarrow Boston (ZBW) flows show periods of high correlation to local traffic. These two plots, along with figure 13 suggest that traffic in the NAS, from a center network perspective, has a strong long distance cause/effect relationship. Thus, traffic events and disturbances in one center can have a large effect on traffic many centers away. This is shown in figure 13 where most of the links show warmer colors at the low scale side of the scale spectrum.

In order to determine how sensitive the SGWT technique is to TMIs, both the CAN-1-East and CAN-1-West playbook routes were implemented in simulation for the entire day of August 22. This is an unlikely scenario in real operations, since traffic blocking convective weather does not persist for 24-hours, but it is desired to see to what extent and for what time of the day the SGWT may show a distinction in results. Using a scale of 0.08, which will highlight low frequency (long period) signatures across the network, the results are shown in figure 11. Both CAN-1 playbooks route cross-country flights to the north and into Canadian airspace north of the Great Lakes. Comparing figure 11 to figure 9, many of the links show a significant difference when the playbook route is implemented. The most striking of which is the Minneapolis (ZMP) \leftrightarrow Chicago (ZAU) flow, which shows markedly less red and orange at peak time, indicating that when the playbook route is implemented, this link has less of a cause/effect correlation with traffic at distant edges. The same is true for the New York (ZNY) \leftrightarrow Washington DC (ZDC) link. Some of this difference is caused by a reduction in volume at these links, but when the same analysis is done comparing plots using a scale value of 1.0, less of a distinction is observed, indicating that the SGWT is capturing the effects of the playbook at a long periodicity of the signal across the center network graph.

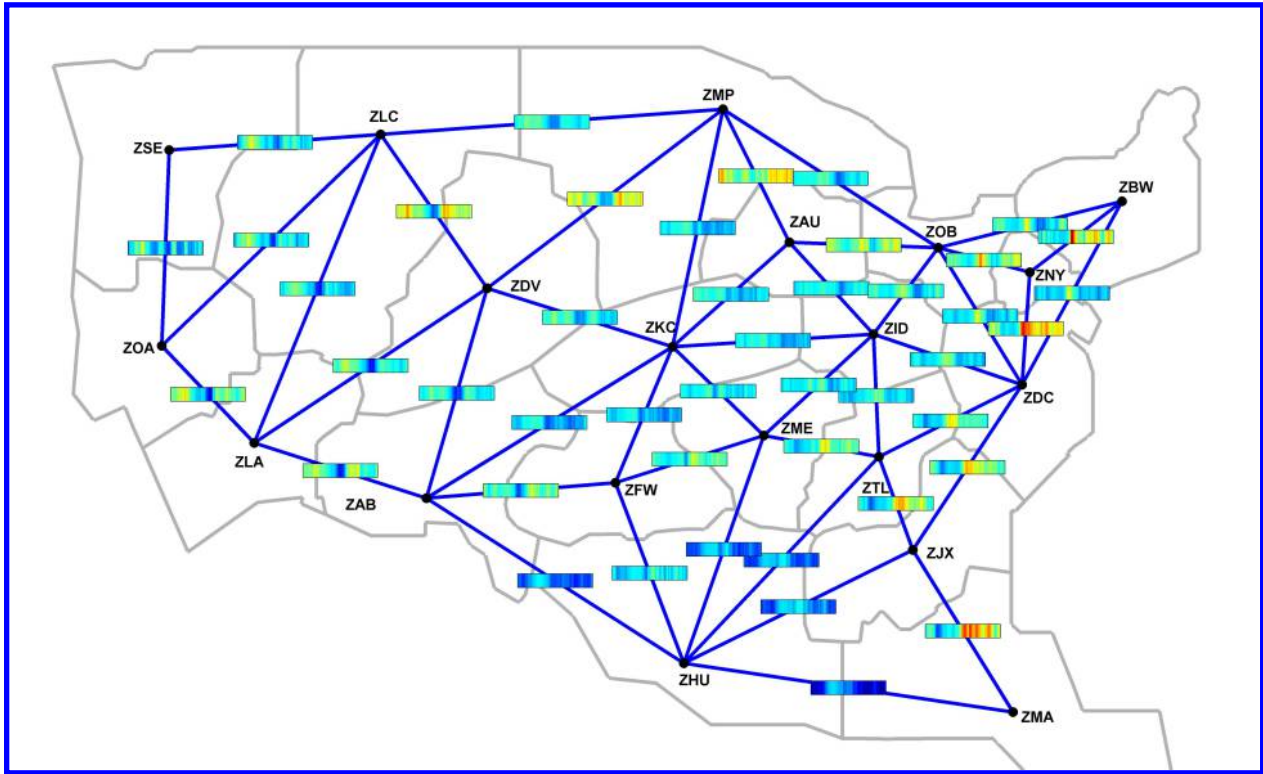


Figure 10: SGWT results for simulated traffic based on flights of August 22, 2014. Scale = 1.0

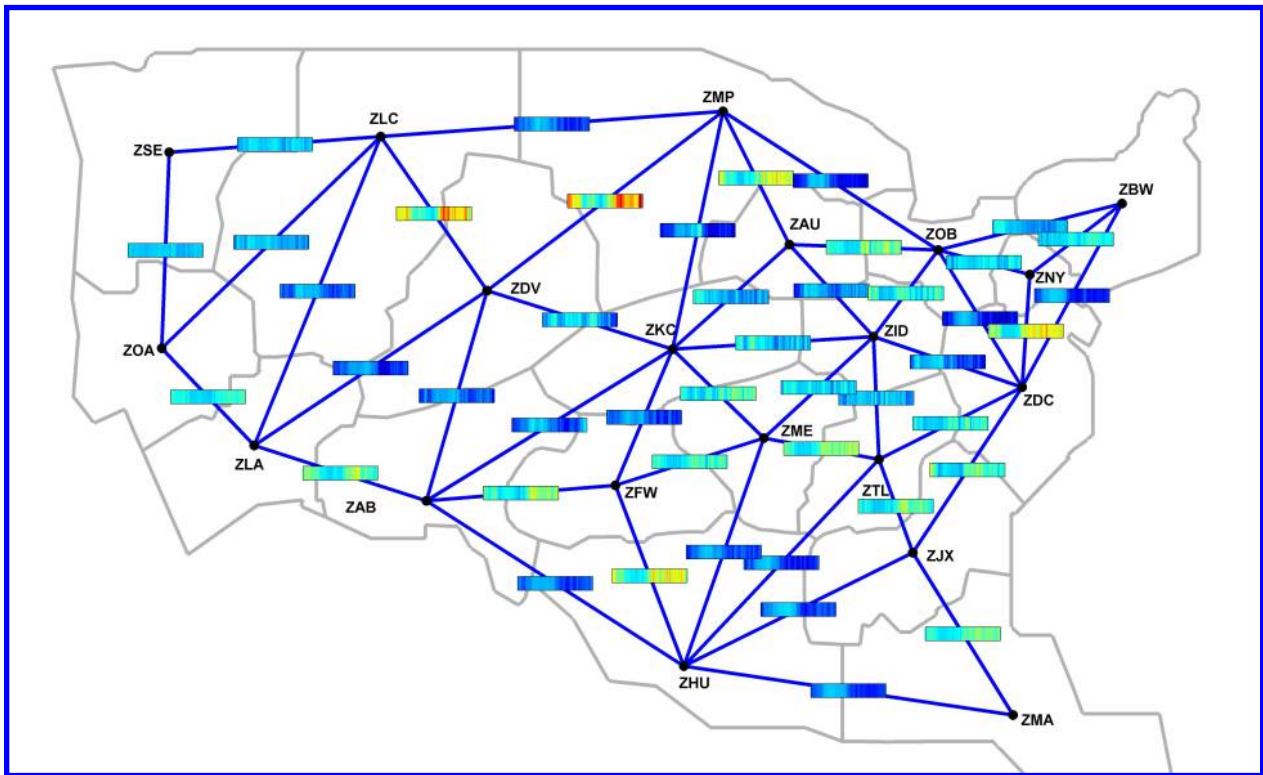


Figure 11: SGWT results for CAN-1-East and CAN-1-West implemented in simulation for August 22, 2014. Scale = 0.08. Compare to figure 9.

The previous figures reveal that the SGWT can highlight interesting characteristics of traffic flow in the NAS. It can also detect the presence of major playbook routes as well as provide some notion of the extent to which the playbook route affects the NAS as a whole. However, more work is required to determine if it can identify, and even predict the severity and extent of a disruption in the NAS due to weather or TMI actions. Unfortunately, it is not yet clear how, or even if this is possible, since it remains difficult to quantify and interpret the SGWT results. In most cases, differences between the baseline and TMI simulations are minimal. Prior applications of the SGWT include internet and highway traffic where the volume per graph edge is much greater, and disruptions have a more noticeable impact. Disruptions and delays in air traffic caused by TMIs, on the other hand, can be much more subtle and/or localized in time and space.

V. Concluding Remarks

In prior work, a frequency domain analysis was shown to provide insight into air traffic streams that was not possible in the time domain alone. In this paper, two additional frequency domain techniques were applied to air traffic. The Continuous Wavelet Transform was shown to provide similar insight that the Fourier transform did, but with the added benefit of preserving some time domain resolution. Though some off-nominal TMI- and weather-affected traffic was detected in the historical data using the CWT, no conspicuous signature could be found to reliably detect and identify specific patterns in the traffic. The CWT analysis of TMI simulations failed to reveal major differences when compared to baseline nominal traffic results. Then, the Spectral Graph Wavelet Transform was applied to a center-based network graph of the NAS. This alternate method was also applied to nominal and off-nominal traffic simulations. The SGWT highlighted some characteristics of traffic flow in the NAS, revealing the relative extent for which traffic events in one area affect the traffic in distant regions of the NAS. It also identified the location of traffic differences caused by a major TMI playbook route, and hinted at the extent (spatially) to which the rerouting affected the NAS as a whole.

However, it is still difficult to reliably identify and quantify off-nominal traffic conditions and their effects on the NAS using either transform. It is suspected that this is due to several factors. First, since traffic within a center consists of many traffic streams, the detection of a single affected stream can be very difficult in either the time or frequency domains. If the scope is narrowed to the sector level, the volume of the traffic data is reduced, and noise (caused by spatial and temporal flight variation) tends to dominate the results. Second, even within a large airspace, the number of fights is small compared to packets traversing a router, or cars on a stretch of highway. When disruptions in the NAS occur, planes do not simply stop moving, unlike internet or road traffic where flow rates routinely drop to zero at affected nodes. To the extent that the most extreme traffic events can be detected using either the CWT or SGWT, it is difficult to argue that similar patterns are not as clear in time domain data. Third, related to the problem of volume, is the relatively short time domain and finite nature of the data. With a sampling period of 1-minute, there are only 1440 points per day. This particularly affects the CWT analysis, which performs better with longer signal length. On this point, however, this work is not without merit. It may be discovered that by widening the time scope of the data to weeks, or months, such a frequency analysis may provide insight to long-term trends in the NAS, like airline strategy, seasonal and/or weather effects, etc.

Appendix

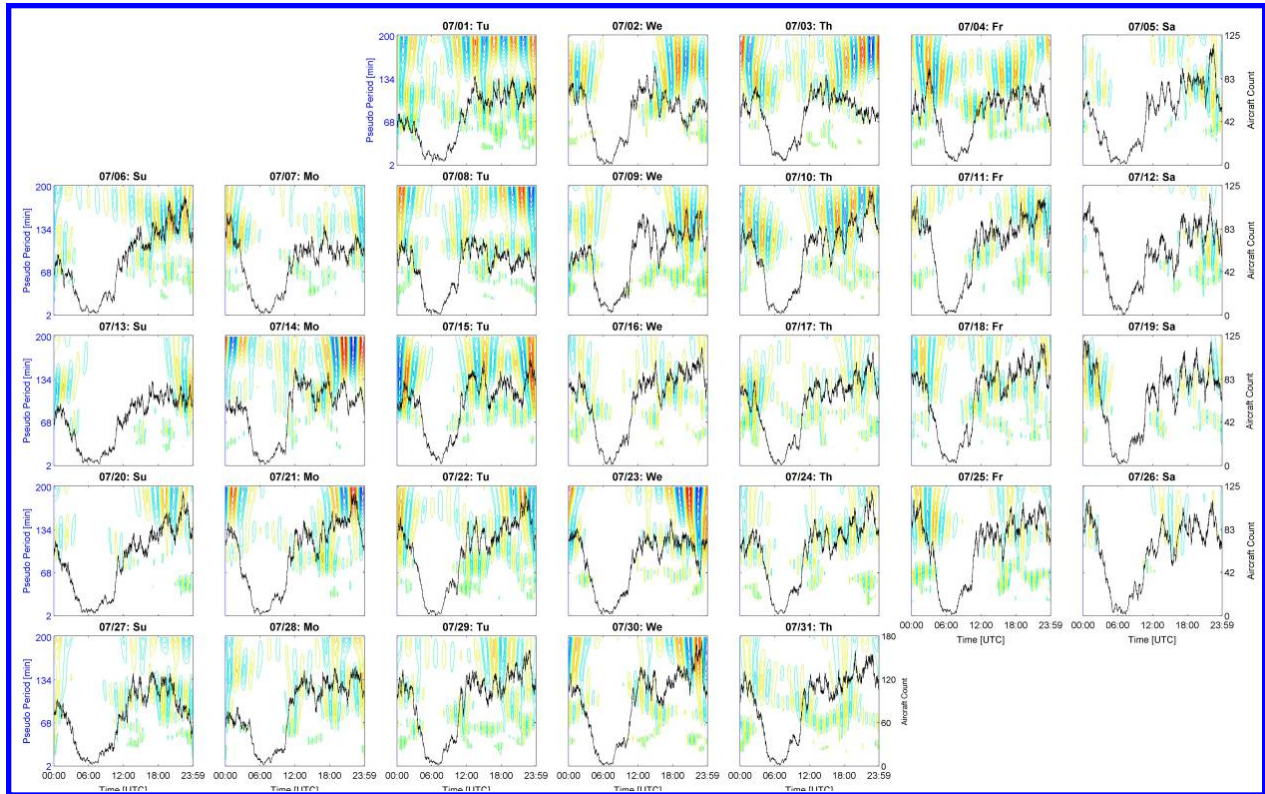


Figure 12: Traffic counts and wavelet analysis for flights above FL250 in Cleveland Center, July 2014. Wavelet contours shown are the difference between that day and the mean of that day of the week's wavelet contours.

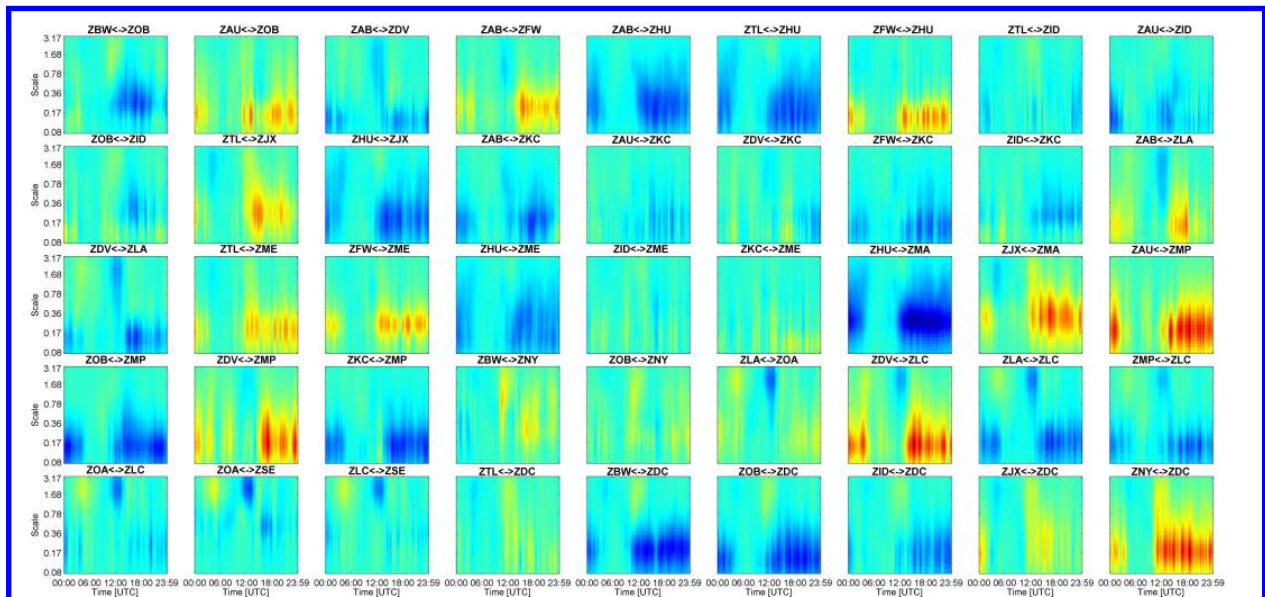


Figure 13: Results of the SGWT analysis on simulated air traffic for August 22, 2014 for every center-to-center flow. Graph scale is on the vertical axis, and time on the horizontal. Warmer colors represent higher coefficients (more traffic correlation to other signals on the graph at that scale).

References

- ¹Sridhar, B., Soni, T., Sheth, K., and Chatterji, G. B., "Aggregate Flow Model for Air-Traffic Management," *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 4, July - August 2006.
- ²Chen, N. Y. and Sridhar, B., "Weather-Weighted Periodic Auto Regressive Models for Sector Demand Prediction," *AIAA Guidance, Navigation, and Control Conference*, Chicago, IL, 10 - 13 August 2009.
- ³Drew, M. and Sheth, K., "A Frequency Analysis Approach for Categorizing Air Traffic Behavior," *14th AIAA Aviation Technology Integration and Operations Conference*, Atlanta, GA, 16 - 20 June 2014.
- ⁴Su, S. and Du, R., "Signature Analysis of Mechanical Watch by Reassigned Scalogram," *Proceedings of World Congress on Engineering*, Vol. I, London, UK, 2 - 4 July 2007.
- ⁵Torrence, C. and Compo, G. P., "A Practical Guide to Wavelet Analysis," *Bulletin of the American Meteorological Society*, Vol. 79, No. 1, 1998, pp. 61-78.
- ⁶Zhang, J., Tsui, F.-C., Wagner, M. M., and Hogan, W. R., "Detection of Outbreaks from Time Series Data Using Wavelet Transform," *American Medical Informatics Association Symposium Proceedings*, 2003, pp. 748 - 752.
- ⁷Crovella, M. and Kolaczyk, E., "Graph Wavelets for Spatial Traffic Analysis," *22nd Annual Joint Conference of the IEEE Computer and Communications Societies*, Vol. 3, IEEE, 2003.
- ⁸Mohan, D. M., Asif, M. Y., Dauwels, J., and Jaillet, P., "Wavelets on Graphs with Application to Transportation Networks," *17th International IEEE Annual Conference on Intelligent Transportation Systems*, August 2014.
- ⁹Hammond, D. K., Vandergheynst, P., and Gribonval, R., "Wavelets on graphs via spectral graph theory," *Applied and Computational Harmonic Analysis*, Vol. 30, No. 2, 2011, pp. 129 - 150.
- ¹⁰Wikipedia, "Laplacian matrix," http://en.wikipedia.org/wiki/Laplacian_matrix, 2015, [Online; accessed 23-April-2015].
- ¹¹Bilimoria, K., Sridhar, B., Chatterji, G., Sheth, K. S., and Grabbe, S., "FACET: Future ATM Concepts Evaluation Tool," *Air Traffic Control Quarterly*, Vol. 9, No. 1, 2001, pp. 1-20.
- ¹²MATLAB, *version 8.3.0.532 (R2014a)*, The MathWorks Inc., Natick, Massachusetts, 2014.
- ¹³FAA, "Air Traffic Control System Command Center National Severe Weather Playbook," <http://www.fly.faa.gov/Operations/playbook/current/current.pdf>, 2014, [Online; accessed 14-Nov-2014].
- ¹⁴Wikipedia, "Line graph," http://en.wikipedia.org/wiki/Line_graph, 2015, [Online; accessed 23-April-2015].
- ¹⁵Hammond, D. K., "The spectral graph wavelets toolbox," <http://wiki.epfl.ch/sgwt>, 2010, [Online; accessed 23-April-2015].