NAS GENOMICS: NEW TECHNIQUES AND INITIAL RESULTS FOR SYSTEM-LEVEL UNDERSTANDING OF NAS BEHAVIOR

Mark Klopfenstein, Terence R. Thompson, John DiFelici, Ray Jakobovits, Jesse Clayton, Brian Capozzi, Andrew Ryan, and Bruce Ware, Metron Aviation, Herndon, VA, USA
James Wetherly, FAA/AUA-740, Washington, D.C., USA

Abstract

The NAS Genome Project seeks to improve our understanding of the behavior of the NAS as a complex system. We report here on initial results from three of the current Genome research areas.

Information Flow and Demand Prediction - We have analyzed the magnitude and content of demand predictions in several data sets and have determined that:

1. There is an error in the way ETMS processes altitudes in historical filed flight plans to create scheduled route (FS) messages, causing some flights to be projected into the wrong sectors.
2. Flight-plan filing time varies significantly by airline, and therefore sector demand predictability is a function of the particular mix of users comprising the demand.
3. The identification of individual flights that will use a sector is quite poor several hours in advance of sector entry.

System-wide NAS Mental Model – We are engaged in systematic description of major NAS elements and their interactions using system dynamics and the Vensim software environment. We have developed initial representations of five major classes of elements of the mental model: command/control, airspace, surface, environment, and flight.

NAS Network Graph - We have developed algorithms to transform very large numbers of the flight trajectories actually used in the NAS into a mathematical graph, that is, as nodes and links with attached attributes describing the traffic that flows through and along them under varying conditions. This enables application of powerful mathematical and computational tools to analysis and control of traffic in this network.

Introduction

The NAS is a complex system, and its users are increasingly demanding improvements in predictability, flexibility and equity. Our challenge as researchers is to better understand the behavior of the system so that we may, in turn, propose enhancements to the NAS. We have recognized that beyond the near-term benefits accruing from improvements in TFM data exchange and collaboration, there needs to be a revolutionary change in the manner in which we perform decision making in an uncertain environment. This requires more fundamental research into the behavior of the NAS, and it is this that is the focus of the NAS Genome Project.

NAS Genome research is tailored to include the investigation and documentation of the daily, monthly, and yearly operation of the NAS from the viewpoint of complex system behavior. It develops in-depth longitudinal data sets that enable NAS operations to be characterized in a statistically significant fashion, including explicit treatment of uncertainty. It also develops means of visualizing and manipulating these data sets in order to facilitate the discovery process. In addition, it leads to understanding and managing the many forms of uncertainty that pervade the NAS (e.g., demand levels/locations, capacities of various NAS elements, weather characteristics), and forging a clear understanding of equitable allocation of limited resources throughout the NAS. Most fundamentally, it develops understanding and lessons learned concerning the behavior of the NAS as a system and as a set of interacting sub-systems.

The principal current thrusts of the Genome effort are in the areas of system analysis (both bottom-up and top-down), data visualization, and data management. More specifically:

1. NAS Behavior Analysis. The goal of this work is to develop understanding of NAS behavior from the sub-system level upward. This includes detailed analysis of information flows, routing decisions, TFM control actions, etc.
2. NAS Mental Model. The goal of this work is to develop a functional description of key processes and data flows throughout the NAS and its subsystems at a level that is appropriate for comprehension of system and sub-system behaviors under varying circumstances.
3. NAS Behavior Visualization. The goal of this work is to augment traditional data-visualization techniques with new approaches and to apply these to long-term, statistically significant volumes of NAS data.

Data management for a system as complex as the NAS is a significant effort in its own right, but is beyond the scope of this paper.

Methods and Results

We discuss here recent results from each of the major areas of NAS Genome research:

1. Analysis of information flows related to demand prediction;
2. Development of the NAS mental model; and
3. Development of radar-based network (graph) representations of NAS traffic flows.

Information Flow and Demand Prediction

To support effective TFM the FAA needs to have a good idea of the expected demand on individual NAS elements several hours into the future. Unfortunately, airlines and other NAS users typically file detailed flight plans only 30 to 90 minutes in advance. Thus, for each regularly scheduled flight, ETMS uses a collection of algorithms to generate a “scheduled route” based on historical data to use for demand predictions. This scheduled route is replaced with information from the filed flight plan once the user submits it to the FAA, and then it is eventually replaced with actual data as the flight becomes active.

In trying to balance demand with capacity the FAA routinely makes TFM decisions based on a combination of scheduled, filed, and actual flight information. Therefore, in order to understand the behavior of the NAS it is necessarily to understand this flow of information, and the accuracy of these demand predictions.

1. ETMS Scheduled Route (FS) Error

While investigating demand predictability we discovered an error in the way ETMS processes historical filed flight plans to create the scheduled route (FS) messages. This error causes ETMS to sometimes project flights into the wrong sectors, which is not updated until the airlines file their flight plans for a particular flight. The significance of looking at these projected sector lists is that they directly affect the accuracy of the sector demand lists and the ETMS Monitor Alert function.

Generally speaking the scheduled route including the cruise altitude is supposed to represent the route & altitude that the airline will most likely file based on what it has filed most often in the recent past. We found that the scheduled cruise altitude is often an altitude that was never actually filed by the airline (see Figure 1). For example, if an airline were to file the same route but half the time at an altitude of 37,000 ft and the other half of the time at 33,000 ft, then ETMS is would assign an altitude of 35,000 ft— the arithmetic mean—for the scheduled route even though the airline has never actually filed this altitude in the past (in some cases because this altitude is reserved for traffic coming the other direction).

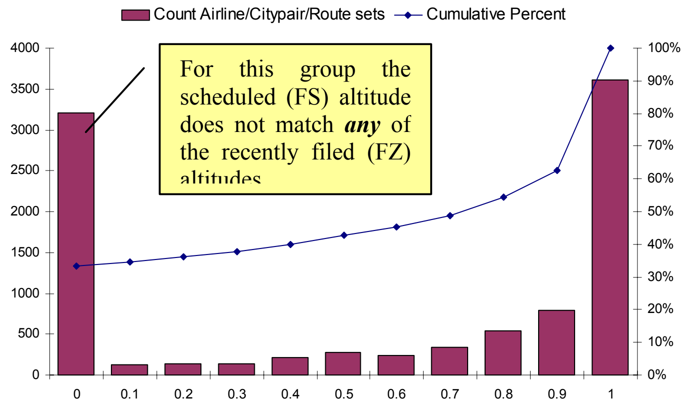

Figure 1: ETMS sets the FS altitude as the mean of the previous week's FZ altitudes for each airline/city pair/route combination. This figure shows the fraction of previous week's FZ altitudes that actually match the current FS altitude (FS date = 5/24/2002, FZ range = 5/16/2002 – 5/22/2002, observed change date = 5/23/2002 ~12:00).

To gauge the significance of this averaging error we looked at the cases where the FS altitude was different from the most commonly filed altitude (vice the arithmetic mean) to see if ETMS would have projected these flights to traverse a different list of sectors. For example, if the altitude difference was enough to cause the flight to move from the high to the super-high sectors. Using the same data set as used in Figure 1 we found the following:

- 4,240 out of 10,146 (42%) of the airline/city pair/route sets have FS altitudes different from previous week's FZ mode (most common) altitude
- 1,563 (15% of total) of these sets result in different sector projections
- In terms of flights, 4,230 out of 29,885 (14%) scheduled flights would have been predicted in different sectors if the mode (most common) of the recently filed altitudes was used as the scheduled altitude instead of the mean.

Based on our findings we have recommended that ETMS be modified to use the corresponding most commonly filed altitude when computing historical routes for FS messages, which the FAA is in the process of implementing.

2. Distribution of FZ Filing Times

In the previous section we saw how inaccuracies in the process that computes schedule flight information can negatively impact demand predictability until it is eventually replaced by filed flight plans. The next question then is, "When do the airlines file their flight plans?"

To answer this we compared the flight plan filing time to both the actual and the scheduled departure times. The data that we examined consisted of all flights from Chicago (ORD or MDW) to Newark (EWR) that flew between 4/9/2002 and 6/7/2002. After removing obvious outliers, the total sample size consisted of 1707 flights. The results of this comparison are shown in Figure 2. For each of the major airlines in this market we show the distribution of the filing times relative to scheduled departure times for all flights with the particular airline's contribution highlighted in dark blue.

Figure 2 shows that some airlines such as COA and AAL appear to have automated processes for filing their flight plans. All of AAL's & COA's flight plans are filed at ~75 or ~42 minutes prior to scheduled departure, respectively. This suggests that these airlines employ an automated process to submit flight plans. Conversely, the distribution of filing times for UAL is both significantly earlier and spread over a broader range (most between 100 – 250 minutes prior to scheduled departure). This suggests that UAL is using a different process to submit flight plans; perhaps one where individual dispatchers may submit flight plans when ready. AMT appears to lie somewhere between the two with most of their flight plans being submitted at either 60 or 120 minutes prior to departure.

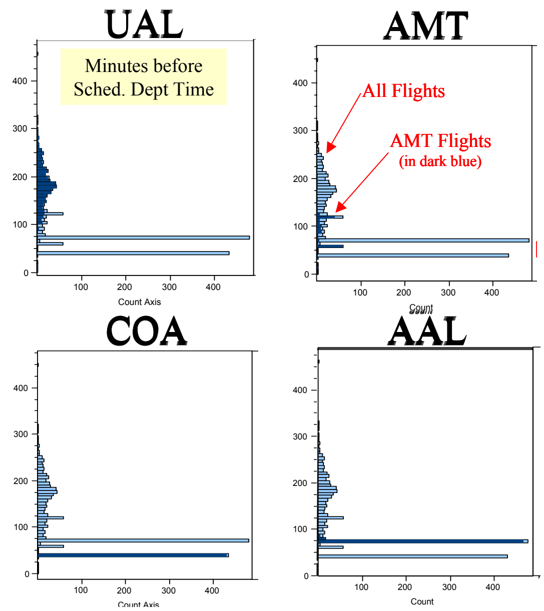

Figure 2: Distribution of FZ filing times

Although we do not show the comparison of flight plan filing times with actual departure times, the results are similar with the exception that the spikes are spread out slightly due to the variance of actual departure times relative to scheduled times.

From this analysis we see that the timing of filed flight plans varies significantly by airline. Because errors observed in scheduled routes are not currently updated until flight plans are filed, sector demand predictability is a function of particular mix of users comprising the demand.

3. Airspace Demand List Stability

Current FAA practices for managing the flow of aircraft into constrained airports include, among other things, the assigning of individualized delays to the projected to arrivals while still on the ground using a variety of algorithms that attempt to do so in an equitable manner [1]. Similar ideas for managing en route traffic congestion are also being developed in which en route TFM control actions with be applied to individual flights. One particular example is the Flow Constrained Area (FCA) procedures under development for use during the 2003 convective weather season. In this case, select flights projected to traverse a defined section of airspace during a time of predicted congestion will be identified by the FAA as needing to be rerouted or delayed by the airlines, so they do not enter the FCA during the time of interest. A more current example, is that TMU specialists will often "pull a flight list" to see the projected demand on a sector to support decisions concerning the need for miles-in-trail restrictions. Therefore, we need to answer the question, "How stable (and accurate) are the airspace demand lists that the FAA uses to manage en route traffic?"

Pulling an airspace demand list involves three variables, look ahead time, spatial extent, and time bin size. For this study, we looked at predictions up to five hours into the future. Spatial extent involves the size of the airspace in which we are interested in knowing the demand. The current Common Constraint Situation Display (CCSD) interface to ETMS, which we used to collect data, only allows for demand lists (for airspace) to be pulled for individual sectors. Future versions supporting flow-constrained areas (FCAs) will allow more flexible spatial configurations. However, we can simulate larger spatial extents by merging individual sector lists. Finally, the time bin size governs the time grouping of the list request. The default is one hour time bins, but this can be anywhere from 15 minutes to multiple hours. The FCA process is reliant on the accuracy of these flight lists, as is any attempt to perform sector (or area) based TFM control actions that affect individual flights. Keeping these instances in mind, we not only looked at the look ahead time component of list stability; we also investigated variation due to the spatial extent and time-bin size of the list request.

Using the CCSD interface to ETMS we set up an automated process to perform regular list requests for a set of forty "high interest" sectors in the NAS. Our initial set of data for March 2002 spanned a look-ahead time of six hours into the future and five hours into the past (which provides actual flight in sector data) binned into one-hour time bins. Data were then exported to a database where for each list request the following metrics were computed:

1. ***Correctly predicted***: Flights that were predicted and actually showed up in the sector during the specified time bin as a fraction of the total number of flights predicted (in that list request)

2. ***Pop*-ups**: Flights that were not predicted but showed up in the specified time bin.
3. ***Over-predicted***: Flights that were predicted but did not show up in the specified time bin.

Figure 3 shows the *correctly predicted* metric as a function of look-ahead time broken out by individual sectors. As expected, it shows that as time nears zero, the fraction of the list that is a correctly predicted flight increases. However, when examining the magnitude we see that the overall stability of the list is quite poor. For example, on average we see that the sectors start out around 40% at five hours prior and only rise to approximated 70% at a look ahead of zero hour, which represents the prediction made at the top of the hour for that hour. Therefore, if a TFM control action were implemented at a time of two hours before based on a one of these list requests then approximately half of the flights in the list would incorrectly be acted (delayed, rerouted, etc.) upon.

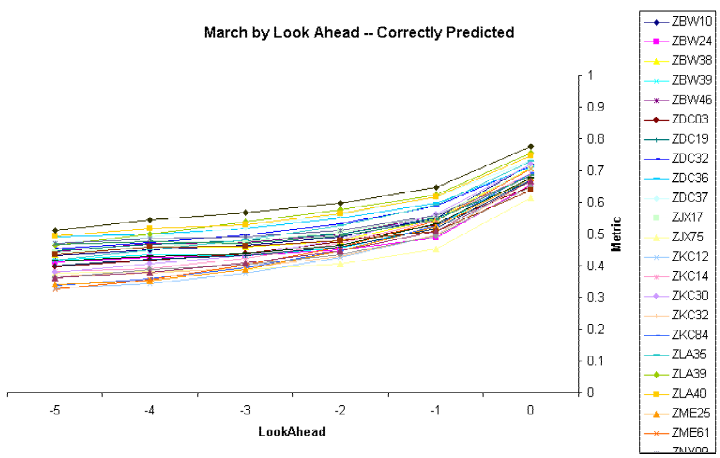

Figure 3: Correctly Predicted by Sector

Next we explored how predictability varies as a function of time bin. We expect that in choosing larger time bins that predictability will improve because some flights that are not correctly predicted may have entered the sector slightly before or after the predicted time bin. Thus, larger time bins should capture some of these flights. For example, a 24-hour time bin would capture all flights that went through the predicted sector as correctly predicted. A secondary goal given the poor results presented above is to discover how much predictability might be improved by expanding the time bin, to better support operational decision-making.

To create multi-hour time bins, we merged one-hour time bins together to create two-hour time bins. demonstrates the gains in using a two-hour time bin as opposed to one. The two-hour bin provides increased predictability as it accounts for flights that may have slipped, in time, from one bin to the other. In future research, we plan to fill in this figure with curves from a spectrum of time bins.

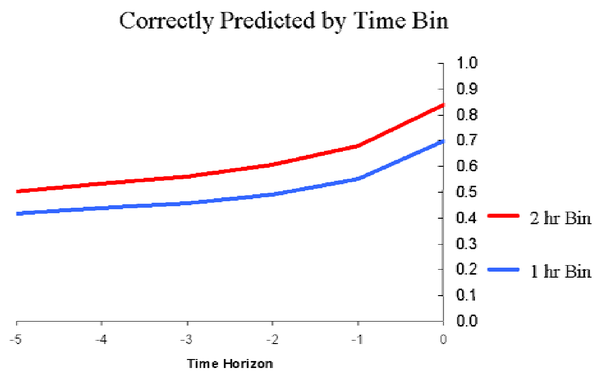

Figure 4: Correctly Predicted Two Hour Time Bin

Another vector of our research involved exploring how predictability varies as a function of the spatial extent of the airspace in which demand predictions are being made. Again the idea is that increases in the spatial extent may improve predictability by including flights that are not correctly predicted because they flew in nearby adjacent areas (i.e. horizontal or vertical sectors). This is directly relevant to the FEA/FCA process as it is often focused on looking at larger than sector areas. In attempting to mimic this scenario, we combined sector lists from two sectors to create a 'super-sector.' We then calculated the "Correctly Predicted" metric. shows some of our initial results for the combination of ZKC30 and ZKC32. Here, the combined results are noticeably better than ZKC30, but only slightly better than ZKC32.

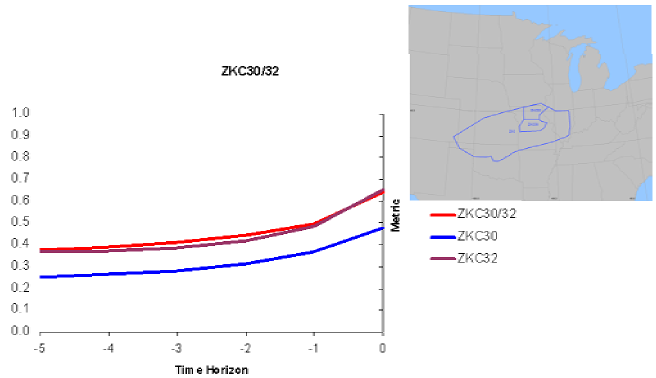

Figure 5: Correctly Predicted for Combined Sectors

There is little increase in the predictability for the combined airspace on ZKC30 and ZKC32. Our hypothesis is that these two sectors are mostly independent such that flights do not often "leak" over to the other sector. We expect that we will get different results when other sector combinations are examined particularly those involving sectors that are stacked on top of each other, which is part of our continued research.

NAS Mental Model

The NAS Mental Model seeks to support system-wide understanding of the operation of the NAS by systematic description of major NAS elements and their interactions. Although some elements of the NAS have been described in detail [2], there apparently exits no unified description of NAS elements, their various instantiations throughout the NAS, and their dynamic, feedback-laden interactions during daily operation. Lack of a common Mental Model makes analysis of current system operation very difficult, and makes exploration of operational alternatives equally difficult.

The principal elements of our approach are as follows:

- Description of NAS elements in terms of inputs, processing, and outputs;
- Organization of NAS elements into groupings that reflect major units and functional associations;
- Electronic representation of model elements for sharing with the community of NAS experts;
- Incremental evolution of model elements; and
- Emulation of some aspects of system behavior.

A much simplified sketch of the Mental Model concept is shown in Figure 6. It is expected that understanding of NAS behaviors will emerge as both the depth and the breadth of the Mental Model increase.

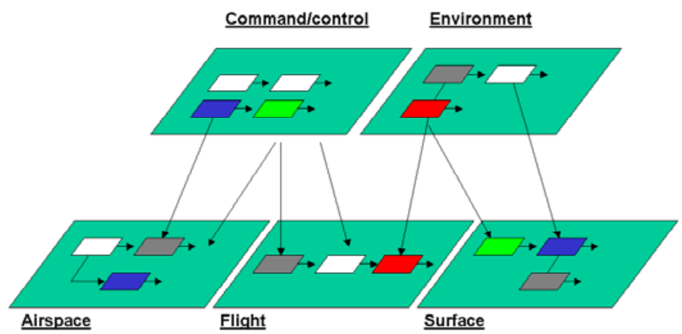

Figure 6: NAS Mental Model Concept

We have selected the Vensim simulation environment for development of the Mental Model because it offers a low-cost and efficient means of supporting the required description and emulation. Vensim is based on the discipline of system dynamics [3] and has been utilized in similar contexts [4] to describe and emulate complex systems.

We have achieved a partial representation of major NAS elements, organized into five principal layers or groups:

- *Command/Control* – Functional elements of the NAS primarily involved in the scheduling, routing, and control of flights;
- *Airspace* – Structural elements of the NAS forming an overlay of boundaries reflecting control responsibilities;
- *Surface* – Functional and structural elements of the NAS involved in the servicing and movement of flights on the airport surface; and

- *Environment* – Features of the environment that affect flight planning, operation, and control.
- *Flight* – Elements related to the operation of individual flights.

These groups serve only as a convenient means of organizing model elements into some initial major classes. As the Mental Model evolves, elements can be shifted between groups and different groups can be created without losing the detailed interactions between elements.

Within these groups, thus far we have achieved the following capabilities:

1. *Models for ARTCCs, sectors, and airports at different levels of detail*. For ARTCCs, flows are defined between upstream and downstream ARTCCs and MITs are modeled. For sectors, individual flows into and out of the sector are modeled, as are airborne holding, MIT, and travel times. For airports we have both a high-level source/sink representation, and a detailed version containing several types of runway holding, taxi transit time, and gate availability.
2. *Integration of a detailed traffic-flow model into the larger mental model.* We have generalized the detailed model (described below), and instantiated it at several specific locations in the NAS by defining specific flow parameters and linkages (transit times, neighboring sectors, etc.). See Figure 7.
3. *Creation of linkages from the command/control layer to the airspace layer*. These linkages define how control actions, such as MIT restrictions, affect resource usage in terms of the flows affected and the associated maximum flow rates.

We can now demonstrate propagation and feedback effects of control actions (e.g., passing back of MIT restrictions from one center/sector to an upstream center/sector).

We now have a tested process by which detailed models from other elements of the NAS community may be abstracted, cloned, and integrated into the Mental Model. This process will allow us to bring in expertise from other modeling efforts and save such knowledge in the growing Mental Model. We have attained an initial environment that utilizes Vensim's ability to view model dynamics from different perspectives (known as "views" in Vensim) to support better system-level understanding of the NAS.

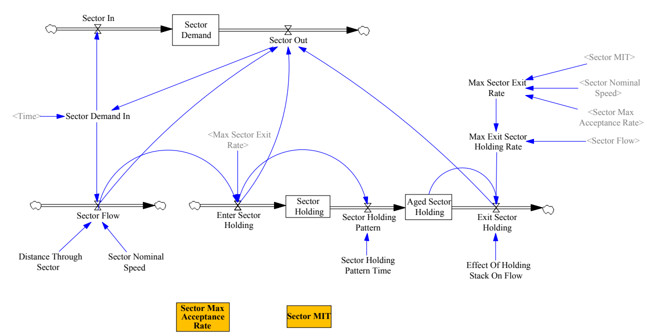

Figure 7. Top-level View of Airspace Layer and Generalized Sector Model Building Block

The detailed model mentioned above addresses a sub-set of the NAS between ORD and EWR. The resulting model represents west-to-east traffic flows through and among the 14 sectors in this "tube", and an initial command-and-control layer with east-to-west propagation of flow controls.

In order to both aid in the construction of the model and validate outputs produced by the model, we selected a single day (30 May 02) as the basis for an initial model database. We obtained the daily logs for that date, and extracted flight tracks for all flights passing through any of the tube sectors. By examining that data, it was clear that the model should include, at a minimum, all flights with destinations in the New York area (EWR, LGA, JFK, TEB, MMU, HPN, and ISP). Based on sector demand statistics, that traffic constituted approximately 25% of all tube traffic (by contrast, the ORD/MDW to EWR traffic comprised only about 4%) on that date. That traffic is the basis for flow-control actions throughout the day, and is much more representative of flows within the tube than just the ORD/MDW to EWR traffic.

Figure 8 shows the flow elements within the tube for all the NYC-bound traffic. Note that all traffic lands in ZNYJF, but flights that do not flow between adjacent tube sectors are indicated as exogenous flows whenever they leave, enter, or reenter a particular tube sector.

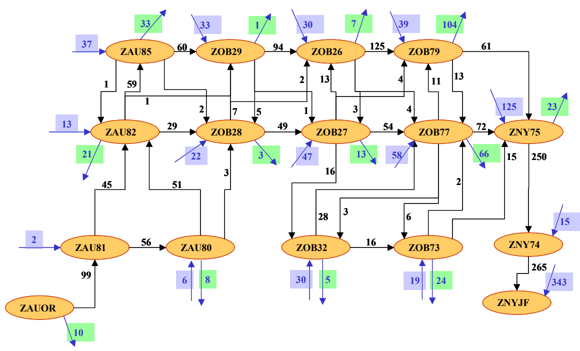

Figure 8. Traffic Flows into 7 NYC Airports Through 14 Sectors (30 May 2002)

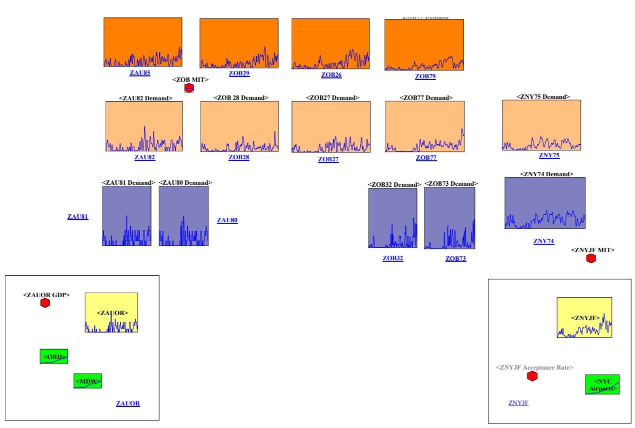

Figure 9. Top-level View of 14-sector Model

The top level "view" of the model is illustrated in Figure 9. This view contains a schematic layout of the tube elements with hyperlinks to underlying sub-models of each sector. The model can be operated in an interactive sensitivity-analysis mode wherein sector demands are computed and displayed as control parameters (e.g., sector capacity, MIT) are varied.

NAS Network Graph

There currently exists no representation of the flight trajectories actually used in the NAS as a mathematical graph, that is, as nodes and links with attached attributes describing the traffic that flows through and along them. Static depictions of airways, navigational aids, and fixes do exist, but these are not graphs in the mathematical sense, nor are they representations that show the dynamic nature of traffic flows. Lack of this NAS network graph hampers research into the behavior and control of traffic in this network, because tools from areas such as graph theory, computational geometry, and autonomous agents [5] cannot be applied. This element of the NAS Genome Project seeks to build a robust, dynamic, and mathematically approachable network graph for the NAS.

A two-stage algorithmic process constructs the network graph from raw ETMS or radar track data. The first stage groups the raw track data into unique flows and generates, for each flow, a geometric average track, called a "backbone". A brief summary of the first stage process for backbone construction is as follows:

- Assemble the raw tracks and filter any necessary track data (altitudes, etc.).
- Perform an initial track grouping using arrival and departure city-pair information.
- Sub-divide each city-pair group using initial and final bearing data.
- Use a spatial-difference clustering algorithm to further sub-divide each group into unique flows.
- Finally, construct a representative track (a backbone) for each flow.

Figure 10 shows typical results from the backbone construction algorithm utilizing 60 days of ETMS data from April and May of 2002.

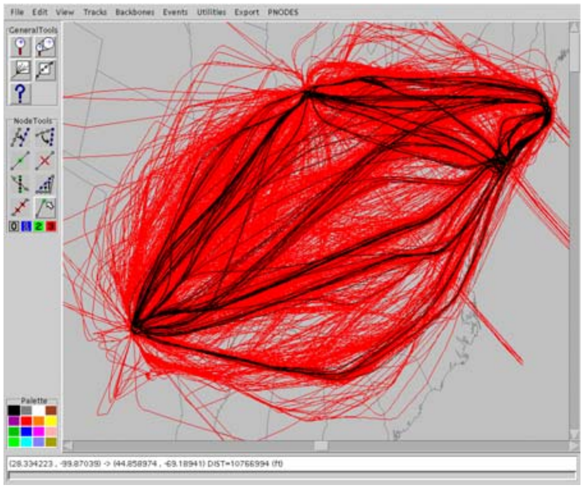

Figure 10 - Backbones For ORD/EWR, DFW/ORD, and BOS/EWR Traffic (radar in red, backbones in black; April-May 2002)

The second stage of network graph construction transforms the backbone tracks into a network comprised of nodes and directed links. A brief summary of the process for this stage is as follows:

- Find two tracks having common sections "close" together. Two track sections are considered "close" if the shortest distance from any point on either section to the other section does not exceed a predefined value.
- Merge the common track sections together to create a link and place a node at each end of the link.
- When all track pairs have been examined, adjust the network bys removing very short links, removing small triangles, removing any duplicate links, smoothing out any track kinks, and consolidating chained links.

Figure 11 shows typical results from the application of the network construction algorithms to a fully connected, twelve city-pair collection of ETMS tracks from April and May of 2002. The twelve airports are ATL, BOS, DEN, DFW, EWR, IAD, LAX, MCO, ORD, PHL, SFO, and STL.

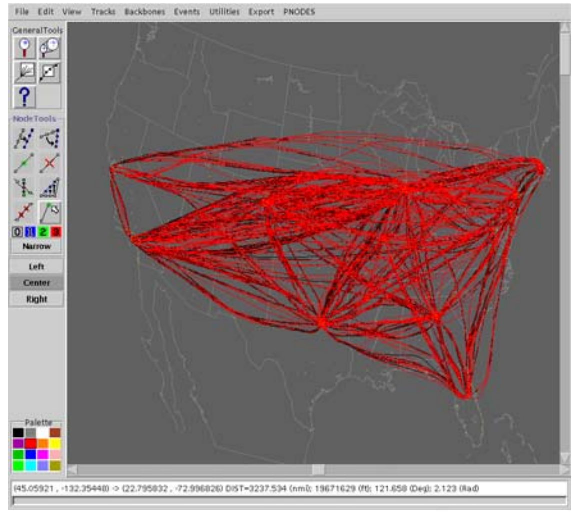

Figure 11 – Fully connected, twelve city-pair network graph (network in red, backbones in black; April-May 2002).

The resulting network files can be displayed in order to visualize link volume and directionality. Volume is represented by the thickness of a link, with a thicker link representing more flight tracks and a thinner link representing fewer flight tracks. Directionality is depicted by drawing the first half of a link in blue (the portion that leaves a node) and drawing the second half of the link in red (the portion that enters a node). Figure 12 shows a network graph constructed from BOS to EWR traffic along with the associated ETMS tracks.

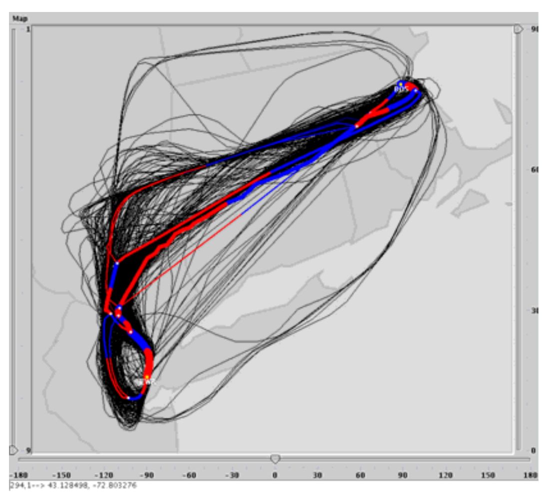

Figure 12 – Directional BOS to EWR network graph with underlying ETMS tracks (network in red/blue, backbones in black; April-May 2002)

Note that the outlying ETMS tracks in Figure 16 represent less than 1% of the traffic in the data set. Under the control of aggregation sensitivity parameters in the network-construction algorithms, they are associated with other flows in the figure shown. When aggregation sensitivity is set to very high levels, these tracks generate flows in the outlying regions.

Note that the network-generation algorithms keep track of all the flight tracks going into each backbone, and then propagate this information to the network links and nodes. These connections between the network elements (links and nodes) and the underlying flight tracks provide the ability to perform queries upon the network, and obtain information such as: aircraft ids, flight times, aircraft equipment types, aircraft carriers, departure/arrival times, and all other information provided by the ETMS or other data streams.

Conclusions and Next Steps

This work has shown that:

- Improvements in information flows in the NAS are essential for better prediction of demand;
- System-wide functional description via a NAS mental model is a key enabler of better understanding and control of system-wide behaviors; and
- Development of NAS network graph representations is feasible and opens the door for application of new analysis and control techniques.

Next steps in this work include the following:

- Comparison of airspace and airport demand predictability, and more detailed investigation of the underlying information flows for airspace demand;
- Extensions of depth and bread in the functional descriptions, information flows, and processing dynamics within the NAS Mental Model; and
- Extension of the NAS Network Graph to larger sections of the NAS, and characterization of resulting graphs according to various operational conditions.

References

[1] Chang, K., et al, 2001, "Enhancements to the FAA Ground-Delay Program Under Collaborative Decision Making", *INFORMS Interfaces*, 31(1): 57-76.

[2] FAA, 1991, *NAS System Requirements Specification, NAS-SR-1000, Change 10*, Washington, D.C.

[3] Forrester, J.W., 1961, *System Dynamics*, MIT Press.

[4] Sterman, J. D., 2000, *Business Dynamics: Systems Thinking and Modeling for a Complex World*, McGraw-Hill/Irwin.

[5] Lucertini, M., Adacher, L., and Meloni, C., 2002, "A Decentralized Approach to Real-time Air Traffic Flow Management", in *Proceedings ATM 2002 Advanced Workshop*, Capri, Italy, Galati and Zellweger, editors.

Keywords

National Airspace System (NAS), Complex system behavior, Information flow, Demand prediction, System dynamics, Network graph

Biographies

Mark Klopfenstein, Terence R. Thompson, John DiFelici, Ray Jakobovits, Jesse Clayton, Brian Capozzi, Andrew Ryan, and Bruce Ware are all members of the NAS Genome Project at Metron Aviation in Herndon, VA, USA.

James Wetherly is Traffic Flow Management Research & Strategic Planning Product Team Lead for FAA/AUA-740. He graduated from Pennsylvania State University in 1987 (B.S.E.E.), and received his M.S. in operations research from George Washington University in 1997.