

Network Model to Address Capacity/Demand Imbalances in the National Airspace System

Tim Myers*

Metron Aviation, Inc. Dulles, VA, 20166

David Kierstead, Ph.D.†

Daniel H. Wagner, Associates Vienna, VA, 22180

August 2008

An imbalance between capacity and demand in one region of the National Airspace System (NAS) will impact other regions through complicated response functions that cannot easily be modeled or understood using traditional, top-down approaches. Understanding how the NAS reacts to capacity/demand imbalances, such as weather disruptions, is vital to the optimal routing and scheduling of air traffic. This paper summarizes research completed to support the development of new Traffic Flow Management concepts, tools, and modeling capabilities that address local, regional, and national flow problems. We present a time-varying, capacitated, multi-commodity Network Flow Model (NetFM) representation of the NAS. We describe the development and application of NetFM as a tool for estimating the unavoidable cost of a NAS event, identifying the optimal solution for mitigating various NAS constraints, and measuring the impact of spatial and temporal uncertainties in weather forecasts to aid in risk management.

I. Introduction

En route constraints, such as weather disruptions, can trigger an imbalance between capacity and demand in one region of the National Airspace System (NAS) with impacts that propagate throughout the system. What is the optimal approach to dealing with a given capacity/demand imbalance so as to minimize the total cost to the system? Does the ideal approach involve spatial control, temporal control, or a combination? How do we measure the unavoidable cost of a given NAS event so that the impact of various events can be quantified on the same scale?

Surveys of related modeling capabilities [1][2] yield significant contributions from Odnoi [3] and Bertsimas [4][5] to the exploration of the traffic flow management problem (TFMP) for optimizing the scheduling and routing of aircraft in the NAS. Additional capabilities include the Boeing National Flow Model, which is a queuing model incorporating airports and Air Route Traffic Control Centers (ARTCCs) as queues in a discrete network [6][7]. There is also the Railroad Network Analysis System for modeling the flow of commodities over the nation's rail infrastructure [8].

This paper covers the development of an aggregate Network Flow Model (NetFM) and its application to answering questions concerning capacity/demand imbalances in the en route environment. We present a brief discussion of existing minimum-cost network models along with a description of the time-varying, multi-commodity approach implemented by this research team. Nodes in NetFM represent en route sectors, arcs are the pathways between sectors, and markets represent clusters of airports supplying demand to the network. Operational data were used to derive parameters required for building the network. This paper presents the methodologies used for representing clusters of airports as markets and an empirical approach for defining en route sector capacity. NetFM is applied as a tool for comparing traffic flows in the NAS during normal and constrained conditions, exploring the tradeoff between rerouting and ground delay, and quantifying the impact of weather forecast uncertainties.

* Senior Analyst, Research and Analysis Department, 45300 Catalina Court, Suite 101, AIAA Member

† Senior Scientist, 450 Maple Avenue, East, Suite 206

II. Modeling

This section describes NetFM and its application to answering questions related to capacity/demand imbalances in the NAS.

While imbalances most commonly result from capacity reductions due to adverse weather, they can also arise from excessive demand. Indeed, a capacity reduction in one portion of the NAS may lead to increased demand in other portions as flights route around the affected area, thus creating new imbalances which may themselves propagate. We developed NetFM to answer such questions as:

- What is the inherent cost of a given capacity reduction?
- What is the optimal combination of rerouting and ground delay for mitigating a given capacity reduction?
- How does a local capacity/demand imbalance affect the rest of the NAS?

Our intent was not to produce a detailed model of how individual flights or carriers should respond to a given imbalance scenario. Instead we studied a reasonably optimal, system-wide response in terms of aggregate flows. Thus NetFM may be viewed as a tool for producing and evaluating strategic plans while glossing over the tactical implementation. NetFM is a very high-level model that treats the traffic as continuous flows rather than discrete flights. Similarly, instead of including all the airports, we aggregate all origins and destinations into approximately 40 geographic *markets*. Thus, for example, one market might capture all of the airports in the Chicago metropolitan area.

NetFM is a multi-commodity network model in which: the nodes represent regions of capacitated airspace; the arcs represent routes between regions, with associated flight costs; and the commodities represent traffic to the various destinations. The details of NetFM are discussed in the remainder of this section.

We first include a general discussion of network-flow problems and how they relate to traffic flows in the NAS. We then describe precisely how we formulated NetFM and computed its parameters. This section then covers the algorithms for computing optimal network flows, and describes the details of the algorithm (based on a linear-programming formulation and a freely available software package).

A. Minimum-Cost Network Flows

We begin our discussion of network flow models with a relatively simple example, involving a single *commodity* being transported over a network. This is analogous to traffic flowing to a single destination. We then extend the problem to multi-commodity networks, which are analogous to simultaneous flows between multiple origin/destination pairs. Finally, time-varying networks are considered which can be used to model dynamic demand and capacities. See [9] for an extensive treatment of network flow problems and methods for their solution.

Single-Commodity Networks

A *network* consists of a set, $N = \{N_i\}$, of *nodes*, and a set, $A = \{A_{ij}\}$, of *directed arcs*, where arc A_{ij} connects node N_i to node N_j . We call N_i the *tail* of A_{ij} and N_j the *head*. There need not be an arc connecting each pair of nodes, and in fact in most transportation networks, each node is directly connected to only a limited number of its neighbors. Flow can only take place along arcs or across nodes. Thus direct flow from N_i to N_j is possible only if the arc A_{ij} exists in A . However, indirect flow may be possible along a *path*, say (N_i, N_k, N_j) , if both A_{ik} and A_{kj} exist. Arcs are *directed*, meaning that flow along A_{ij} may take place only from the tail to the head (N_i to N_j), however if the reverse arc, A_{ji} , also exists then direct flow may be possible from N_j to N_i . Thus bidirectional arcs may be represented by a pair of directed arcs. Flow enters the network through a designated set of nodes called *sources*, and leaves through designated *sinks*. All other nodes, which can neither introduce nor remove flow, are called *transshipment* nodes.

Figure 1 illustrates a simple, single-commodity network consisting of: four sources, one sink, two transshipment nodes, and ten arcs. In order to emphasize the relationship with air traffic we have labeled the sources as origins, the sink as the destination, and the transshipment nodes as en-route airspace.

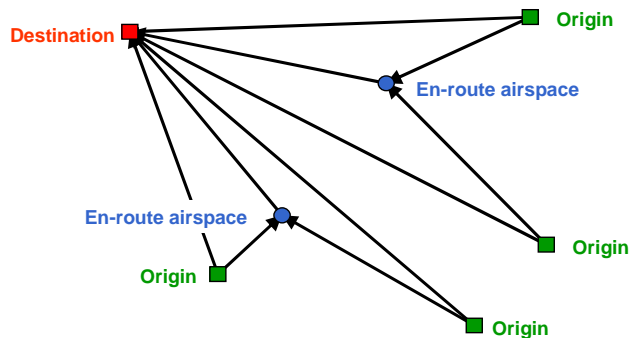


Figure 1. A single-commodity network.

since a transshipment node neither adds nor removes flow). In the case of a source (or sink) it is an arbitrary decision whether to count the flow into (or out of) the network toward the capacity constraint. We have chosen not to, so the capacity constraint at a source or sink applies only to flow that actually passes through the node. That flow is equal to the inflow of a source or the outflow of a sink. (The issue does not arise in **Figure 1**, but would if one of the transshipment nodes were also a source.)

The objective of the *single-commodity network flow problem* is to find the set of flow values for each arc that yields the minimum transportation cost for the entire network, satisfies the given demand for source-sink pairs, and is subject to the capacity constraints at the nodes. In a multi-commodity NetFM, commodities represent specific destinations so there is only one sink per commodity that accommodates the total demand.

We note in passing that there are a multitude of equivalent formulations of the network flow problem. For example, arcs can have capacities and nodes can have costs. Any of these formulations may be transformed to any other via suitable manipulations. For example, arc constraints may be modeled by introducing an artificial (capacitated) node in the middle of each arc. We have chosen this particular formulation because it seems to correspond well with our intuitive notion of the NAS. We also note that our particular solution accommodates most variations without requiring transformations of the underlying network.

In NetFM, transshipment nodes correspond to volumes of airspace with associated capacities, arcs represent flight paths with associated costs (e.g., distance), and sources and sinks represent *markets* which are surrogates for centers of demand (see B. Formulating the NetFM and Setting its Parameters for further details). **Figure 2** illustrates the solution for a single commodity NetFM with 36 origins (green squares) to a single destination (red square). The gray lines represent arcs, arc intersections represent the nodes, and the blue lines depict the flows over the arcs.

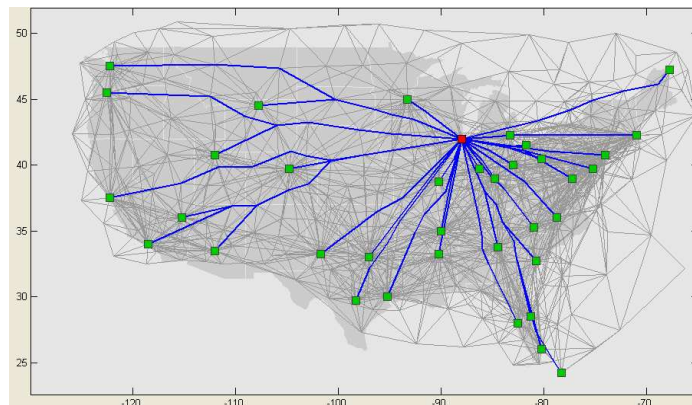


Figure 2. Flows to a single destination.

Multi-Commodity Networks

In a multi-commodity network, demand and flows are not homogeneous. Each unit of demand is associated with a particular commodity, and that commodity must reach an identified subset of the sinks. However, all flow occurs within the same underlying network, and in particular the nodes' capacity constraints apply to the sum of all flows for all the commodities. In our application a commodity is the collection of flights to a given destination. Thus each

* It is more common in the literature to think of the sinks as having demand which must be satisfied from the sources, but in our context it seems more natural to think of the demand originating at the sources and being satisfied at the sinks. The two approaches are entirely equivalent.

origin will have different demand for each commodity, corresponding to the amount of traffic departing for that destination.

In order to model a multi-commodity network, we must track the flows of each commodity separately. Whether a given node is a source, sink, or transshipment node now depends on the commodity in question. If a node is a transshipment node for a given commodity, then the sum of inflows of that commodity must equal the sum of outflows of that *same* commodity. The standard way to model a multi-commodity network is to replicate the underlying network once for each commodity, and then apply the capacity constraint at each node to the sum of all flows of all commodities through that node.

Figure 3 illustrates the concept for a two-commodity network. Here we have augmented the single-commodity network of **Figure 1** with a second commodity consisting of flights destined to Destination 2. Destination 1 has become an origin (source) for the second commodity. Reverse arcs have also been added so that all of the arcs are effectively bidirectional.




Figure 3. A two-commodity network.

Figure 4 illustrates a more realistic representation of the NAS as a 37-commodity network. There are 37 red squares, each representing a market which is a destination (sink) for one commodity and an origin (source) for all the others. The blue lines represent flow between market pairs.

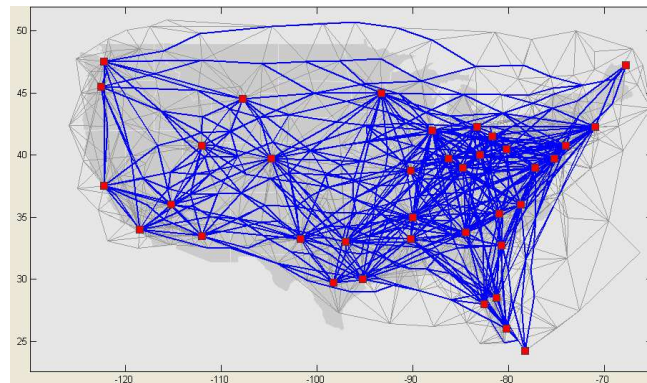


Figure 4. A 37-commodity network.

Time-Varying Networks

The multi-commodity network described above models a static network, in which capacity and demand remain constant over time. This assumption implies that the only choice is over which routes to send the various commodities. In the NAS, conditions change and there is also the option to delay flights on the ground until the capacity/demand balance improves. In the current NetFM, we approximate the option to incur ground delay by considering a sequence of static networks, one for each of a set of discrete time steps, with different capacity and demand profiles at each time step. As illustrated in **Figure 5**, we then connect the origins of the networks for successive time steps with one-directional *delay arcs*, represented by red lines. We assign a cost to each delay arc to reflect the cost of a ground delay of one time step. A flow through the augmented network consists of traversing zero or more delay arcs until reaching the desired time step, and then traversing that network to the appropriate destination. Of course, as with the underlying networks, there must be a separate delay arc for each commodity. Although we have not implemented it in the current NetFM, it




Figure 5. A multi-time-step network with delay arcs.

would be just as easy to account for the possibility of airborne holding by introducing delay arcs between successive versions of the transshipment nodes. The airborne-delay arcs would presumably incur greater costs than the ground-delay arcs.

We refer to this multi-time-step network as *semi-dynamic* because it captures some of the dynamic nature of the NAS, but imperfectly. What the model fails to capture is that temporal variation is reflected in spatial variation. For all its limitations, the semi-dynamic model can provide useful predictions for relatively localized events. For example, if we wish to predict the effect of a weather event near Cleveland at 1500Z we should set demand profiles based on departure times of flights that would be near Cleveland at about 1500Z. Thus we should set Los Angeles to New York demand based on typical traffic departing Los Angeles at about 1100Z, but set New York to Los Angeles demand based on typical traffic departing New York at about 1400Z.

Fully-Dynamic Networks

It is possible to develop a fully dynamic network model. Such a model would have to ensure consistency between time and distance. For example, we might restrict the nodes to a grid of hex cells and set the time steps to reflect flight time between adjacent nodes as shown in **Figure 6**. The model would remain imperfect as it would artificially increase the modeled distance between points that were not connected by straight lines on the grid. However, the errors would be moderate and could be controlled to be no more than +/-8%. There are other, more complicated, approaches that could reduce modeling errors to arbitrarily small values, but at the cost of increased computational complexity. While results presented in this paper are based on the semi-dynamic model, we are in the process of implementing the fully-dynamic model and anticipate results being available in the near future.

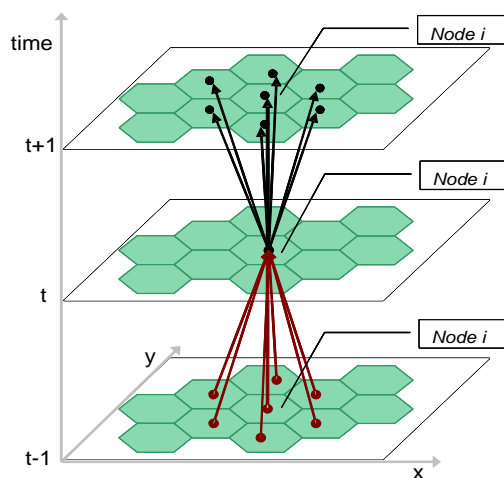


Figure 6. A fully-dynamic implementation of NetFM will consist of a grid of hex cells to reflect equal flight time between adjacent nodes.

B. Formulating the NetFM and Setting its Parameters

We developed methodologies for computing the network parameters based on operationally-observed data in an attempt to create a model that would behave and respond similarly to the actual NAS. These methodologies are described in this section.

Defining Markets

Demand in the NAS originates and terminates at thousands of airports. One approach to creating a network model could be to create a multi-commodity network with individual airports represented by independent nodes. Due to the high number of airports, such a network could become computationally very expensive to solve. As an alternative, we cluster groups of airports into markets and treat each market as single sink/source node. For example, all the airports in the New York City metropolitan area are grouped into one market airport which is then treated as one commodity in NetFM. This assumption seemed reasonable given our main focus was on modeling en route events, as opposed to the details of flows in the terminal airspace which is the topic of related research [10].

Markets were defined by identifying spatial concentrations of arrival demand based on Enhanced Traffic Management System (ETMS) arrival data between July 7, 2005 and August 31, 2005 through the following steps:

1. Arrival counts for each airport were matched with the airport coordinates to populate a spatial grid throughout the NAS. This arrival count grid was then processed using a 2-D Gaussian smoother having standard deviation of about 30 nmi to create a smooth arrival density grid as shown in **Figure 7**.
2. Find all peaks in the smoothed arrival density and denote these as candidate markets.
3. Associate each grid cell with its closest candidate market and denote the sum of arrival densities from all cells associated with each candidate market as the candidate market's weight.
4. Define major markets as the top 35 candidates by weight. These major markets account for 65% of all arrivals.
5. In order to capture the remaining 35% of arrival demand, define orphan peaks as all peaks in the arrival density not associated with a major market.
6. Create an orphan density grid by assigning a 1 to grid cells containing an orphan peak, 0 otherwise.
7. Smooth the orphan density (1 or 0) with a 2-D Gaussian smoother using a standard deviation of approximately 150 nmi. The smoothed orphan density appears in **Figure 8**.
8. Add any peak in the smoothed orphan density to the list of major markets:
35 major markets + 10 orphan markets = 45 markets
9. Filter the set of markets to those inside the NAS. Total number of markets reduces to 37 markets.

NetFM does not differentiate between “major” and “orphan” markets. These designations are used only in reference to way market airports are identified. **Figure 9** shows the final set of 37 markets used in NetFM.

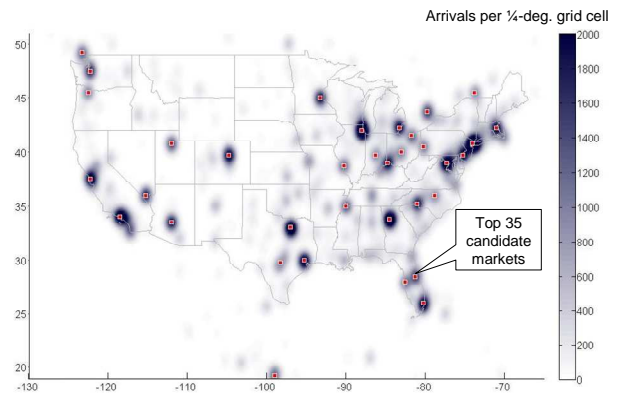


Figure 7. Smoothed arrival density used to identify airport markets. Major markets defined by top 35 candidate markets

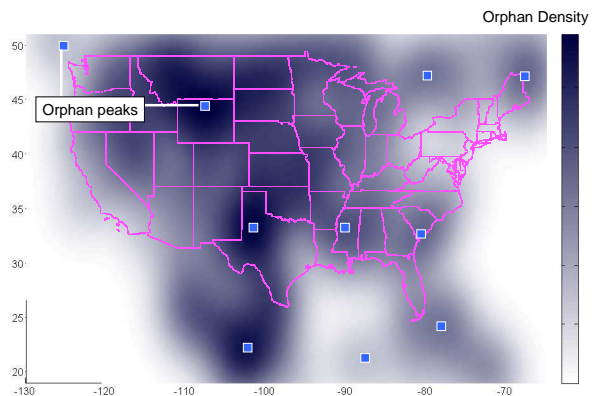


Figure 8. Smoothed orphan density used to identify non-major markets and to account for 100% of the arrival demand

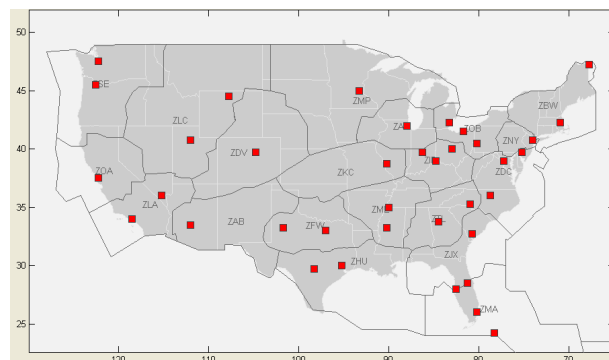


Figure 9. Location of 37 NetFM airport markets

Defining Nodes

Nodes for NetFM were defined from the set of operational en route sectors having both a floor at or below flight level 240 (FL240 or about 24,000 feet) and ceiling above FL240. This is equivalent to identifying all sectors that intersect with a horizontal slice taking through the 3-D sector airspace at FL240. Node locations were presented in NetFM by the centroid of each sector. Using this logic, 353 NetFM nodes were defined to represent individual en route sectors. **Figure 10** shows the boundaries of these sectors at the FL240 slice along with the locations of the nodes representing each sector.

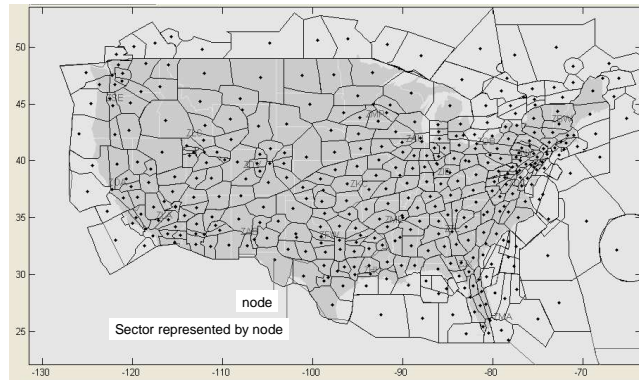


Figure 10. 353 NetFM nodes and the sectors they represent

Estimating Node Capacity

En route constraints such as weather can lead to an imbalance between demand and capacity, which may require airborne delay, Miles-In-Trail (MIT), or other Traffic Flow Management (TFM) initiatives. Unlike demand, capacity cannot be easily measured. Although Monitor Alert Parameters (MAP) have been published to establish the maximum instantaneous sector occupancy, no metric exists for defining the capacity per unit time (e.g. total flights per hour) that can be handled by a given piece of airspace. In this section, we describe an empirical approach to defining a capacity metric for a region of airspace which is then applied to capacitating the nodes in NetFM. From this point on, we will refer to “sectors” and “nodes” interchangeably in reference to the node representation of 3-D sectors in NetFM.

Sector capacities were estimated based on observe sector usage. Usage was measured within 3-D regions of airspace defined laterally by the boundaries of ETMS sectors at the FL240 slice (see **Figure 10**). The vertical boundaries of NetFM sectors were set to FL180 and FL600. This approach was chosen, versus using the altitude regimes defined in the ETMS grid files, so that the sector capacities in NetFM would represent actual en route usage within a consistent altitude range.

ETMS track (TZ) data were used from the eight-week period 7/7/2005-8/31/2005 to provide 4-D (time, latitude, longitude, altitude) trajectory data of actual flights. We then implemented a sector matching algorithm to identify and record the times when each flight entered each sector. Sector entries times were aggregated into 30-minute bins over the eight-week period. Initial results suggested that inaccuracies in the TZ altitude were producing unreliable sector entry counts due to the altitude-dependant nature of the sectors as illustrated in **Figure 11**. As a result, we implemented an altitude filtering-smoothing algorithm to produce more realistic altitude profiles, resulting in more accurate sector entry counts.

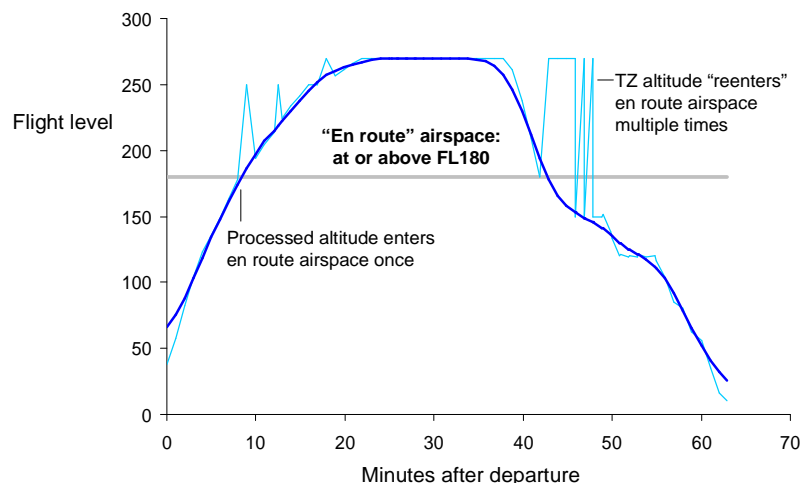


Figure 11. In the absence of TZ altitude data processing, a flight might be recorded as making repeated entries into the same sector due to poor data quality. A filtering-smoothing algorithm was implemented to define a more realistic altitude profile resulting in improved sector usage data.

The number of entries into each sector during each 30-minute time bin was then analyzed to identify a reasonable estimate for the maximum sector capacity. **Figure 12** shows the distribution of observed usage for the sector defined by the lateral boundary of ZAB15.

Notice the sharp climb in the usage curve starting just after the 95th percentile leading up to a maximum value of 95 flights per 30-minute time bin. The sharp climb indicates that entry rates for this sector seldom exceed about 60 flights per time bin. Analysis of usage curves for many sectors identified a common spike just after the 95th percentile. We therefore determined that the maximum usage (100th percentile) was not a reasonable representation of a sector's true capacity given the steep climb in the usage curve approaching the maximum. Instead, the 95th percentile usage was selected to define node capacity for the sectors represented in NetFM.

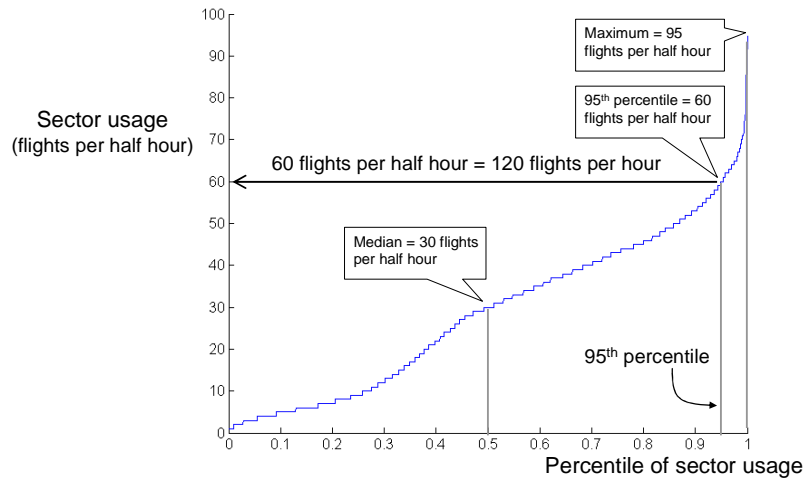


Figure 12. Inverse cumulative distribution of sector usage per 30-minute time bin. Entry counts for ZAB15. The 95th percentile usage of 60 flights, per time bin (120 flights per hour) was used to define the capacity of this sector in NetFM.

Figure 13 displays the computed node capacities for the 353 NetFM nodes. Shades of blue indicate relative hourly capacities with the maximum capacity of 168 flights per hour coming from ZDV19. The large relative area of ZDV19 (43,000 nmi² versus the average sector area of about 19,000 nmi²) contributes to it having the highest estimated capacity.

Another approach to visualizing sector capacity is to look at the capacity per unit area or “capacity density” as illustrated in **Figure 14**. To create this figure, we simply normalized the sector capacities in **Figure 13** by sector area then multiplied by 10,000 to define capacity density in terms of flights per 10,000 nmi² per hour. As one may expect, the highest and lowest capacity per unit area are found in the North East and oceanic/Canadian airspace, respectively. The median capacity density throughout the NAS is 55 flights per 10,000 nmi² per hour, which means for example that the average 10,000-nmi² sector has a maximum capacity of about 55 flights per hour. White shading indicates regions with capacity density below 33 flights per 10,000 nmi² per hour.

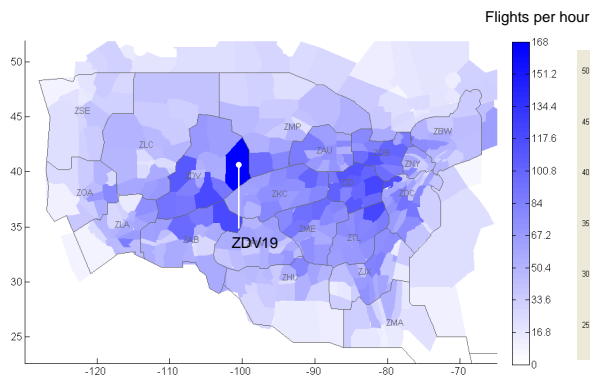


Figure 13. Sector capacity based on 95th percentile observed usage

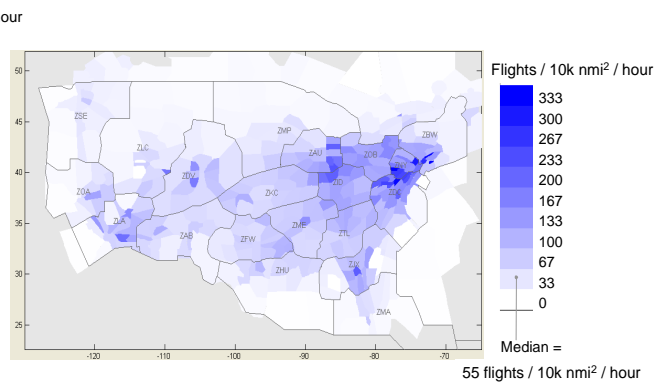


Figure 14. Sector capacity density

Defining Arcs

Arcs were defined in NetFM as pathways between nodes. These node-to-node transitions were defined from observed sector entry data using actual flight tracks during the eight-week period of 7/7/2005 through 8/31/2005. An arc was included if and only if that arc or its reverse arc was observed being used at least 10 times during the eight-week period. We put this filter into place for two reasons

1. To filter out anomalies; and
2. To ensure all NetFM arcs were bidirectional.

We also added arcs to connect airport markets to the nodes. This step was completed by identifying the market nearest to each departing flight and connecting that market with the first sector entered. The reverse was done on the arrival end to connect nodes to the arrival markets. The resulting NetFM network consists of 37 markets, 353 nodes, and 4,782 arcs as shown in **Figure 15**.

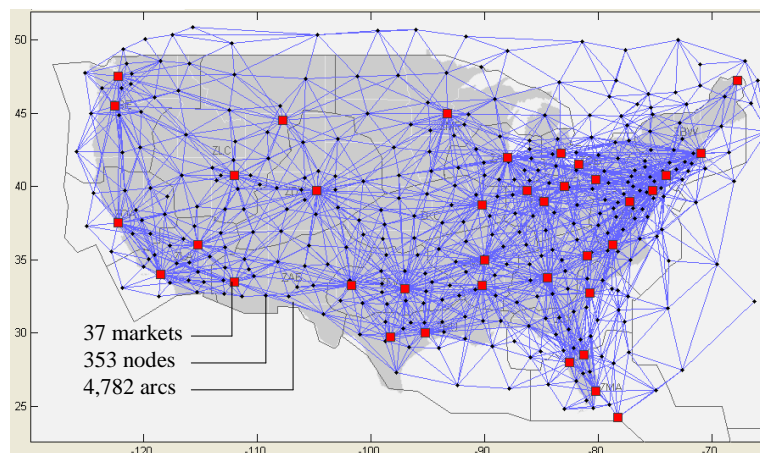


Figure 15. NetFM underlying network consisting of markets, nodes, and arcs. Arcs defined by observe market-node, node-node, and node-market transitions.

Defining Hourly Market Pair Demand

The number of flights in the NAS varies significantly throughout the day. In order to capture this temporal variability, we computed a set of baseline market demand parameters specific to each hour of the day. Quantifying the demand between each market pair is important because this is the amount of demand being transferred through the network.

We computed the hourly demand between each market pair by identifying the markets nearest the origin and destination airports for each flight. As with our definition for sector capacity, we computed market pair demand based on actual traffic in the FL180-FL600 altitude range. In particular, flights that never climbed above FL180 were not included as contributing to market pair demand. This altitude constraint was added so that both node capacity and market pair demand would represent the same set of flights; those flying between FL180 and FL600. The temporal component of demand; the hour in which the demand was contributing, was defined by the flight midpoint. Using this approach, the average demand per hour was computed during the eight-week period starting on 7/7/2005 and is plotted in **Figure 16**.




Figure 16. Average hourly total market pair demand

Eq. (1) is the total transportation cost function. Eqs. (2) are called the *flow conservation* constraints to ensure that the inflow and outflow to any node sum to the node's demand. If N_i is a source ($d_i > 0$), then its outflow must exceed its inflow by d_i since N_i adds d_i units of flow to the network. If N_i is a transshipment node ($d_i = 0$), then its outflow must equal its inflow since N_i neither adds nor removes flow from the network. In either case, the difference between outflow and inflow must be d_i . If N_i is a sink ($d_i < 0$), then its inflow must exceed its outflow by whatever quantity of flow N_i removes from the network, and that quantity may not exceed $|d_i| = -d_i$. After rearranging the terms we arrive at the inequality in Eq. (2)*. Eq. (3) represents the capacity constraints. As discussed in section (A. Minimum-Cost Network Flows), the definition of capacity requires that we apply the constraint to the inflow of a source and the outflow of a sink. We may apply the constraint to either the inflow or outflow of a transshipment node, but have chosen to apply it to the inflow. Finally, the non-negativity constraint in Eq. (4) merely ensures that flows conform to our notion of direction.

A multi-commodity network. To extend our model to multiple commodities, we extend the indexing of the control variables (the flows) and the flow conservation constraints to include multiple commodities:

$$\begin{aligned} d_{ik} & : && \text{the demand for node } N_i \text{ for commodity } k. \\ f_{ijk} & : && \text{a control variable defining the flow of commodity } k \text{ across arc } A_{ij} \text{ if it exists, } 0 \\ & && \text{otherwise.} \end{aligned}$$

Thus a given node may be a source, sink, or transshipment node for one commodity, independent of its role for other commodities.

The formulation of the multi-commodity network flow problem is nearly identical to the single-commodity formulation (although the multi-commodity problem is much harder to solve).

$$\text{minimize} \quad C = \sum_{i,j,k} p_{ij} f_{ijk} \quad (5)$$

subject to:

$$(\forall N_i, k) \quad \sum_j f_{ijk} - \sum_j f_{jik} \begin{cases} = d_{ik} & \text{if } d_{ik} \geq 0, \\ \geq d_{ik} & \text{if } d_{ik} < 0, \end{cases} \quad (6)$$

$$(\forall N_i) \quad \sum_k F_{ik} \leq c_i, \quad \text{where } F_{ik} = \begin{cases} \sum_j f_{jik} & \text{if } d_{ik} \geq 0, \\ \sum_j f_{ijk} & \text{if } d_{ik} < 0, \end{cases} \quad (7)$$

$$(\forall A_{ijk}) \quad f_{ijk} \geq 0. \quad (8)$$

Note that the flow conservation constraints apply separately to each commodity, so each commodity may be thought of as traveling over its own, private network. However, the capacity constraints introduce interactions between these independent networks. The capacity constraints are somewhat more complicated than those for the single-commodity network (Eq. (3)) because nodes may play different roles in the various networks.

A semi-dynamic, multi-commodity network. Finally, we extend our model to a semi-dynamic network, say with time steps indexed by t . For simplicity and to conform to our application, we assume that there is a unique sink for each commodity (the same for all time steps). We first replicate the static network once for each time step, extending the indexing of all nodes, arcs, and flows to include the time step. Next, we add delay arcs, with associated costs, connecting each potential source node to its version in the next time step:

* If there happens to be a zero-cost path from one sink, S_0 , to another, S_l , then it is possible that in some optimal solutions S_0 would contribute flow which would be removed at S_l . Such flow would have no cost so the solutions would still be optimal for the problem as stated. In our case, that problem cannot arise since there is only one sink for each commodity. If the potential for such solutions exists, then one could add an additional constraint that the sum of the outflows of a sink not exceed the sum of the inflows.

- D_{ikt} = delay arc from node N_{it} to $N_{i(t+1)}$, defined only if time step $t+1$ exists and N_{is} is a source for commodity k at *some* time step s .
 π = the cost of delay for one time step, which for simplicity we assume to be the same for all delay arcs.
 ϕ_{ikt} = a control variable defining the flow of commodity k across delay arc D_{ikt} if it exists, 0 otherwise.

The reason that we have limited delay arcs to sources is that we are modeling ground delay, which must be taken at the origin.

The formulation of the semi-dynamic network flow problem is similar to the static formulation, except that the flow conservation constraints become somewhat more complicated at sources, as we shall explain below.

$$\text{minimize} \quad C = \sum_{i,j,k,t} p_{ij} f_{ijkt} + \pi \sum_{i,k,t} \phi_{ikt} \quad (9)$$

subject to:

$$(\forall N_i, k, t) \left(\phi_{ikt} + \sum_j f_{ijkt} \right) - \left(\phi_{ik(t-1)} + \sum_j f_{jik(t-1)} \right) \begin{cases} = d_{ikt} & \text{if } d_{ikt} \geq 0 \\ \geq d_{ikt} & \text{if } d_{ikt} < 0 \end{cases} \quad (10)$$

$$(\forall D_{ikt}) \quad \phi_{ikt} - \phi_{ik(t-1)} \leq d_{ikt},$$

$$(\forall N_i, t) \quad \sum_k F_{ikt} \leq c_{it}, \quad \text{where } F_{ikt} = \begin{cases} \sum_j f_{jik(t-1)} & \text{if } d_{ikt} \geq 0, \\ \sum_j f_{ijkt} & \text{if } d_{ikt} < 0, \end{cases} \quad (11)$$

$$(\forall A_{ij}, k, t) \quad f_{ijkt} \geq 0, \quad (12)$$

$$(\forall D_{ikt}) \quad \phi_{ikt} \geq 0.$$

The flow conservation constraints in Eq. (10) have become somewhat more complicated. First, we must account for flow delayed to the next time step as part of the node's outflow, and delayed flow from the previous time step as part of the inflow. Second, since we are modeling ground delay, we need to restrict flow over the delay arc to the sum of current demand from the source and previously delayed demand from the *same* source. In other words, the flow over D_{ikt} cannot exceed the flow over $D_{ik(t-1)}$ by more than the current demand. The second part of Eq. (10) imposes that constraint. It is implicit in this constraint that N_{ik} is a source (or possibly a transshipment node) for commodity k at time t since we assume that D_{ikt} is only defined for nodes that are sources for commodity k at some time and that there is a unique sink for that commodity over all times.

D. Computational Algorithms

There are a multitude of variations of the network-flow problem, of which we have touched on only a small fraction. These problems have been studied extensively, both because of their practical applications and mathematical significance [1][2][3][4][5][8][9]. Very efficient algorithms are known for many special cases, with computational complexities that are low-order polynomials of numbers of nodes and arcs. For example, if the maximum degree of the network (the maximum number of arcs with a given node as head or tail) is bounded by a constant, then the single-commodity problem previously described may be solved by an algorithm that scales as $(n \log(n))^2$, where n is the number of nodes. Unfortunately, the multi-commodity problem does not appear to admit such efficient algorithms. There are ways to exploit the network structure and partition the problem into a sequence of single-commodity problems, the parameters of which are the solutions of smaller LP problems. However, there does not appear to be a canonical, general-purpose algorithm for the multi-commodity problem.

We have chosen to treat NetFM as a standard LP problem. This approach is desirable due to the generality of the LP formulation that provides a flexibility to add any linear costs or constraints. Specialized algorithms, though perhaps

more efficient on the problems for which they were designed, tend to fail when any of the assumptions are altered in even minor ways. Additional benefit for LP formulation is the availability of open-source software packages like the GNU Linear Programming Kit (GLPK)* that solves moderately large-scale LP problems with reasonable efficiency. There are also commercial packages, CPLEX probably being the premier, available for solving very large LP problems very efficiently with the ability to exploit network structure where it exists. Both GLPK and CPLEX were used in solving NetFM problems described in this paper.

III. Applications of the Network Flow Model

This section presents results of applying NetFM to problems involving capacity/demand imbalances in the NAS. We begin by presenting the NetFM solution for normal operational conditions followed by a description of how NetFM can be used to model constraints resulting from changes in capacity and demand, such as a weather disruption. Scenarios involving constrained conditions discussed in this paper include reduced node capacity, arc flow constraints, and increased market pair demand. Comparisons between normal and constrained conditions are then examined to identify the impact of these modeled scenarios. Dynamic airspace constraints are discussed followed by a discussion of the “sphere of influence” concept for predicting and evaluating the spatial propagation of demand fluctuations resulting from airspace constraints. This section concludes with a description of how NetFM can be used to measure the impact of spatial and temporal uncertainties in weather forecasts.

A. Normal Operations

A fundamental application of NetFM is defining the minimum cost (distance) solution during normal operations. This solution is computed by applying the default values for node capacity and market pair demand as defined in previous sections. **Figure 18** shows the node utilization modeled during normal operations based on the market pair demand during the 2200Z hour. Node utilization refers to the ratio of flow over capacity for each node. Darker shades indicate nodes having flow at or near maximum capacity.

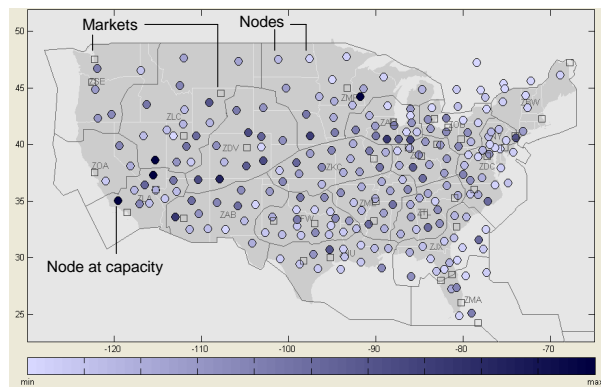


Figure 18. NetFM solution showing node utilization during normal 2200Z operations. Shading based on node usage relative to node capacity

NetFM also computes the flow through each arc as shown in **Figure 19**. Arcs are shaded based on their relative flow. As depicted in the figure, predominant east-west flows can be observed near the Los Angeles and San Francisco markets, moving east towards the Chicago market. Another predominant flow is visible on the Atlantic coast between the Miami and New York City markets.

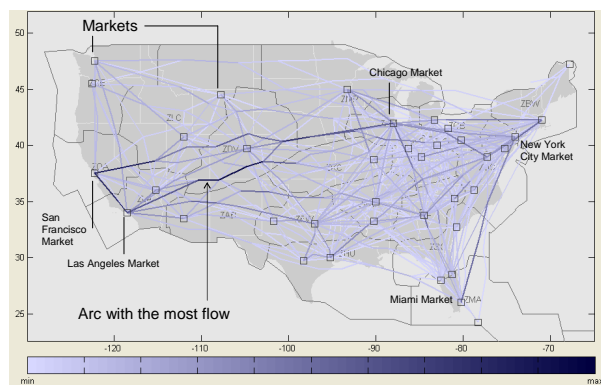


Figure 19. NetFM solution showing arc flow during normal 2200Z operations shaded relative to arc with the most flow

The map display in **Figure 20** combines node utilization and arc flows using the shading scheme described for **Figure 18** and **Figure 19**.

NetFM defines the minimum cost solution in terms of flow through each node and arc for a given set of demand and capacity conditions. The total cost of the normal operations solution based on 2200Z demand is computed to be 1,098,430 nautical miles, which is the minimum total distance required to satisfy all

* The GLPK package is a part of the GNU project, released under the aegis of GNU. GLPK is available free, under the terms of the GNU General Public License as published by the Free Software Foundation.

market pairs' demand, while not exceeding the maximum capacity of any nodes. Now that we have presented the NetFM solution during normal operations, the following section introduces the results of altering capacity and demand parameters as a means for modeling their impact on the NAS.

B. Airspace Constraints

The NetFM solution shown in **Figure 20** illustrates node utilization and arc flows during normal operations (i.e. clear weather day). A goal of this research was to model the high-level impact of a NAS event by comparing NetFM solutions during normal and constrained conditions.

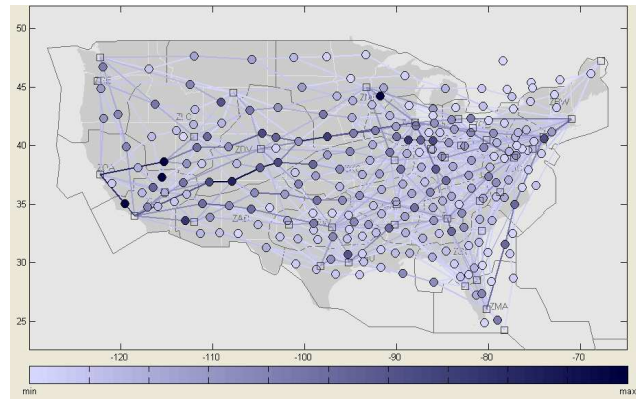


Figure 20. NetFM solution during normal 2200Z operations showing node utilization and arc flows

The NetFM representation of the NAS entails a certain degree of abstraction in that sectors, which are multi-dimensional regions of airspace are represented as points in space (i.e. nodes) at the intersection of arcs (pathways between sectors). This abstraction needs to be considered when applying constraints to specific portions of the network.

Three methods were explored for imposing airspace constraints on the network.

1. Reducing node capacity relative to maximum node capacity
2. Reducing node capacity relative to normal node flow
3. Reducing arc capacity relative to normal arc flow

We originally applied node capacity reductions relative to the maximum capacity of each node. Again, this maximum capacity was obtained from the 95th percentile of observed usage as described in section on Estimating Node Capacity. Since sectors rarely operate at such high usage levels (by definition only 5% of the time), a reduction in capacity of even 50% might not have a significant impact on normal levels of demand. An alternative approach which was found to be favorable is to record node and arc flows during normal operations as a baseline. Capacity reductions are then applied relative to this baseline of normal operations. For example, node i might have a maximum capacity of 60 flights per hour, yet during normal operations the node flow is only 20 flights per hour. Applying a 25% capacity reduction to this node would further reduce its capacity down to 15, rather than 45, flights per hour.

Figure 21 illustrates two approaches to implementing airspace constraints in NetFM. When placing a constraint on the nodes within a polygon **Figure 21 (b)**, only the arcs to and from constrained nodes are affected, leaving the potential for arcs that traverse the polygon to be left unimpeded. For example, three arcs are unaffected by the node flow constraint shown in **Figure 21 (b)** as they cross the airspace constraint (shaded region) without utilizing any restricted nodes. Placing the constraint on flows along arcs **Figure 21 (c)**, however, enables the isolation of all flows within a polygon boundary.




Figure 21. Normal operations with no arc constraints (a), arcs affected by capacity reduction placed on nodes within a polygon (b), and arcs affected by capacity reduction placed on arcs intersecting or within a polygon (c).

Whereas node flow constraints capture only the flow through nodes within a constrained region of airspace, arc flow constraints appear to be more realistic

in that they affect all flows within a region of airspace. We therefore chose method 3, arc capacity constraints relative to normal arc flow, as the preferred method for modeling airspace constraints such as weather in NetFM.

In the following example, NetFM was applied to evaluate the impact of an airspace constraint lasting one hour. The constraint, depicted as the polygon in **Figure 22**, may be thought of as a region of severe weather or a flow constrained area (FCA) within which normal arc flow totals 642 flights per hour. The modeled constraint restricted total arc flow within the polygon to 341 flights per hour, a 50% reduction.

The figure illustrates the modeled impact of this airspace constraint as the difference between solutions during constrained and normal operations. Red indicates nodes and arcs that had increased flow during constrained operations. Blue indicates reduced flow. Notice that the colors are shaded on a log scale to magnify subtle changes in flow.

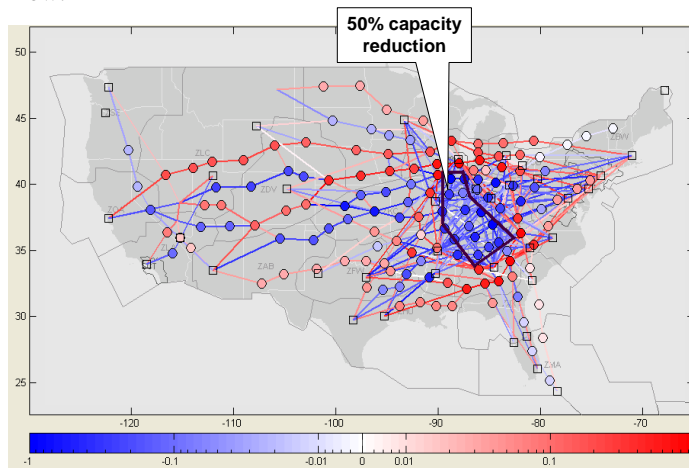


Figure 22. 50% reduction in arc capacity relative to normal arc flow

There is an additional network cost associated with the reduction in node capacity within the polygon. The added cost is due to some flows being forced to take longer paths between market pairs in order to satisfy the reduced capacity within the polygon. Comparative results are summarized in **Table 1**, which shows the additional mileage and dollar cost associated with the modeled airspace constraint.

Dollar costs for the normal and constrained solutions throughout this paper were computed assuming

- 400-knot average groundspeed
- \$60 per minute of flight time*.

Table 1. Impact of the 50% arc capacity reduction shown in Figure 22 lasting 60 minutes. Assumes \$60 per minute of airborne operations.

		Normal	Constrained	Change	Percent Change	
Within polygon	Total arc flow (flts/hr)	682	341	-341	-50.00%	Constraint on arc flows relative to normal flow
Entire network	Cost of solution (nmi)	1,098,430	1,109,576	11,146	1.01%	
	Cost of solution (\$)	9,885,870	9,986,184	100,314	1.01%	

Impact of airspace constraint in mileage and dollars

These assumptions are used only for converting NetFM outputs, which are given in units of distance, to units of cost. An estimate of the dollar cost associated with the one-hour constraint modeled above using these conversion parameters is roughly \$100,000.

Up until this point, NetFM applications discussed have involved a single time step. The results of these single time step models reflect the steady state solution, which assumes demand and capacity parameters are constant through time. The operational equivalent to these results would be an airspace constraint (such as a cell of thunderstorms) forever having a fixed position in space. In practice however, capacity/demand imbalances have a finite duration.

* This estimate was based on data available in 2006. Recent increases in fuel prices have undoubtedly impacted airborne operating costs. NetFM solution costs listed in dollars can be scaled up (or down) based on the ratio of current airborne operating costs to the baseline cost of \$60 per minute used throughout this paper.

This gives controllers and airlines the option of routing around a constraint or waiting on the ground until the constraint clears. In the next section, we discuss how ground delay was incorporated in NetFM through the modeling of multiple time steps.

D. Including the Option for Ground Delay

When an en route capacity constraint such as an Airspace Flow Program (AFP) or weather event occurs in the operational world, flights have the option to fly around the constraint – requiring additional flight time, or to wait on the ground until the constraint dissipates. By implementing a time-varying network as discussed in the section on Time-Varying Networks, NetFM begins to answer the following types of questions for a given NAS event:

- At what point does ground delay become an essential part of the optimal solution?
- Which market pairs should take the option for ground delay?

The cost associated with taking the option for ground delay π (introduced in Mathematical Formulation) is related to average groundspeed, the duration of the ground delay, and the cost ratio of ground versus air operations. The “delay arc cost”, is computed as:

$$\pi = v \times d \times r , \tag{13}$$

where v represents the average groundspeed of airborne flows in nautical miles per hour (knots), d denotes the duration of the constraint in hours, and the cost ratio of ground versus airborne operations is denoted by r . Assuming an average groundspeed of 400 knots, duration of half an hour, and cost ratio of ground to airborne delay of $\frac{1}{2}$:

$$\pi = 400kts \times 0.5hrs \times 0.5 = 100nmi .$$

In this example, waiting on the ground for half an hour has the same cost as flying an extra 100 nautical miles. Factors making the option for ground delay more attractive include:

- **Slower assumed groundspeed** Flying the extra distance to go around a constraint would take twice as long at 200 knots than it would at 400 knots, making the option to wait it out on the ground more attractive
- **Shorter constraint duration** It is more appealing to wait out a constraint lasting only 30 minutes than a constraint lasting an hour
- **Lower ratio of ground:air operating costs** If the cost of ground operations fell from $\frac{1}{2}$ to $\frac{1}{4}$ that of air operations, then it would be less costly to wait on the ground. Ground delay would also become a more attractive option if increases in fuel prices made air operations more than twice as costly as ground operations.

This tradeoff between taking ground delay and flying around a constraint is capture in NetFM by Eq. (13).

The following two-time step examples model the impact of airspace constraints of increasing size to identify the point at which ground delay becomes a part of the optimal solution. During the first time step $t(1)$, arc flows within the polygon are limited to 75% of their normal usage – a 25% reduction. Arc flows return to the unconstrained state during the second time step $t(2)$. Each time step in these examples represents 30 minutes, which simulates a weather event lasting half an hour and then clearing. We modeled three circular constraints of increasing size, as shown in **Figure 23**, having radii 100-, 200-, and 300-nmi, respectively. The three maps on the left show the impact of the constraint during the first time step by comparing node utilization and arc flow to the normal operations solution. Regions of increased flow appear in red, decreased flow in blue. Regions of increased flow are visible in red just outside the boundary of each constraint. This is equivalent to flights being pushed out into the neighboring airspace.

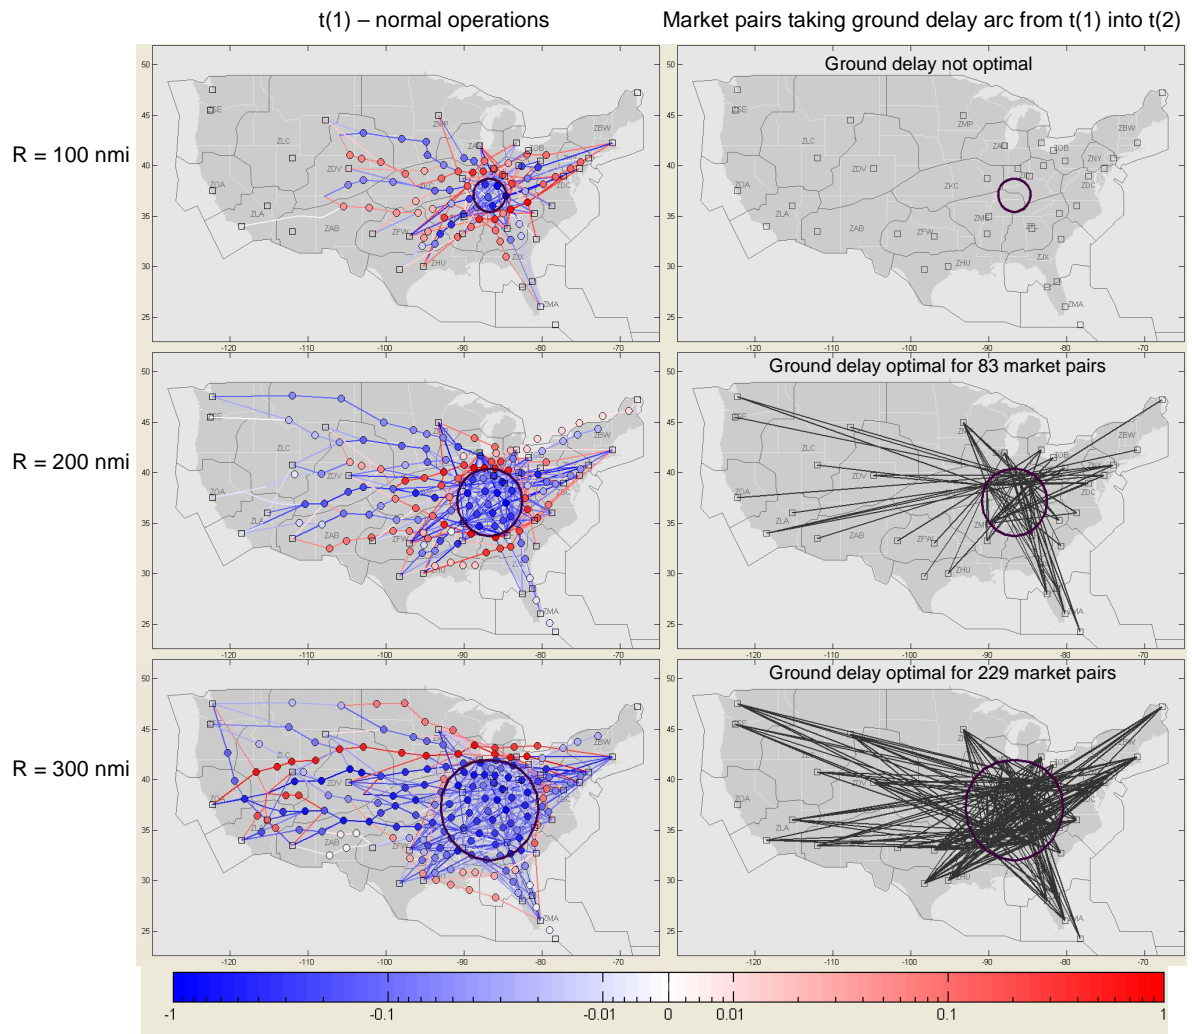


Figure 23. NetFM solution for multi time step scenario involving a 25% reduction in arc flow with polygons of radius 100-, 200-, and 300-nautical miles, respectively, applied during t(1) with full capacity restored during t(2)

Maps in the right hand column show which market pairs use ground delay in the optimal solution. When the airspace constraint is relatively small (e.g. 100 nmi), ground delay is not necessary and reroutes alone are suitable for addressing the en route flow restriction. As the constraint grows to 200 nautical miles in radius, the problem can no longer be handled spatially while maintaining a minimum cost. The optimal solution now requires the implementation of both reroutes and ground delay. In this way, NetFM may eventually be used as a tool to identify which type of TFM approach, or combinations of approaches, will best deal with a given capacity/demand imbalance. These results are summarized in **Table 2**.

Table 2. Summary of solutions displayed in Figure 23 where each time step equals 30 minutes. Assumes \$60 per minute of airborne operations and ground-to-airborne cost ratio of 1/2.

Radius (nmi)	costs (\$) relative to normal operations			Total	Market pairs taking ground delay arcs
	t(1) arc	t(1) ground delay	t(2) arc		
100	3,506	0	0	3,506	0
200	-147,002	29,349	155,763	38,111	83
300	-445,176	73,823	451,170	79,817	229

Arc costs during the 100-nmi constraint are elevated during t(1) by about \$3,500 relative to the normal operations solution. This added cost represents flights willing to travel the extra distance to go around the relatively small constraint rather than take ground delay. The constraint disappears during t(2) and capacity returns to normal. Since the ground delay arc from t(1) to t(2) was never used, demand during t(2) is identical to that during normal operations resulting in no added arc cost during the second time step. The total modeled cost of the 100-nmi constraint is just over \$3,500.

As the radius of the constraint increases to 200 nmi, arc costs during t(1) relative to normal operations are reduced by \$147,002. This cost reduction is due to a reduction in flow through the network arcs as demand begins taking the ground delay option instead. The demand that took ground delay during t(1) is then added to the preexisting demand during t(2). The result of the added demand during t(2) is a significant increase in arc costs of \$155,763 relative to normal operations. The total cost of the 200-nmi constraint is therefore \$38,111, which is the sum of the arc costs during each time step and ground delay costs from t(1) to t(2). The cost of the two time step scenario involving the 300-nmi constraint is \$79,817.

The results in **Table 2** show how NetFM can be used as a tool for measuring the impact of different size capacity/demand imbalance constraints as well as for identifying the point at which a constraint is so large that ground delay becomes an essential part of the optimal solution. As previously described in this section, the amount of ground delay included in the optimal solution depends on the trade off between routing around a constraint (taking additional flight time) versus waiting out the constraint on the ground. This trade off is defined by parameter π in Eq. (13). As a result, NetFM can also be used as a tool for modeling how a system wide TFM strategy might be impacted by various conditions such as rising fuel costs. Physical constraints such as limits on hold fuel can also be incorporated into NetFM.

E. Dynamic Constraints

The impact of static airspace constraints was examined in the previous section. Real constraints, in particular weather disruptions, change through time. NetFM enables the modeling of airspace constraints that are dynamic in position, size, and intensity. This section describes the solution to a problem involving a dynamic airspace constraint. The example consists of four time steps, each representing 30 minutes. During time steps 1-3, an airspace constraint imposes a 25% capacity reduction relative to arc flow during normal operations. The position and shape of the constraint change during each time step as depicted in **Figure 24**, eventually clearing during the fourth and final time step.

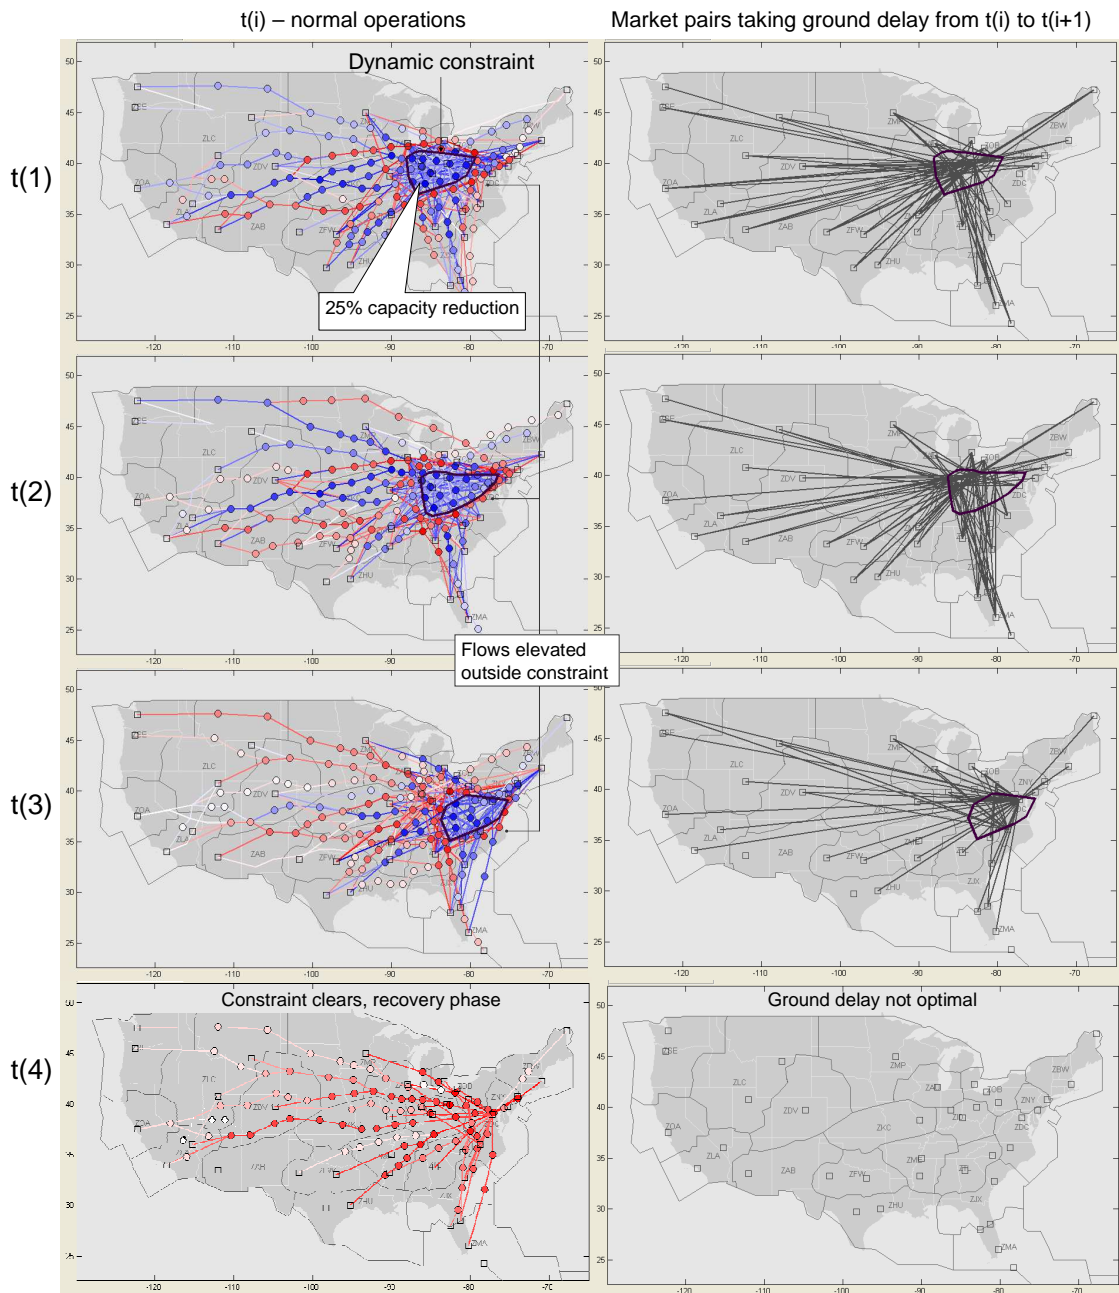


Figure 24. Impact of a dynamic airspace constraint during a four-time step NetFM simulation. Changes in arc and node flow relative to normal operations appear on the left. Ground delay market pairs during each time step are shown on the right.

NetFM determines the optimal combination of rerouting and ground delay so as to minimize the total cost during this four-time step problem. Maps on the left show arc and node flows relative to normal operations. Again, with blue indicating reduce flow and red showing increases. Maps on the right identify market pairs taking ground delay during each time step. These results are summarized in **Table 3**.

Table 3. NetFM solution to the four-time step problem involving dynamic capacity/demand imbalance constraints illustrated in Figure 24. Assumes \$60 per minute of airborne operations and ground-to-airborne cost ratio of 1/2.

Cost (\$) relative to normal operations				
Time Step	Arc	Ground Delay	Total	Market pairs taking ground delay arcs
1	-92,291	22,131	-70,160	94
2	-98,474	40,334	-58,140	94
3	94,667	22,437	117,104	44
4	124,902	0	124,902	0
Total	28,805	84,902	113,706	

Flows inside the constraint polygon are reduced, indicated by the blue shading. This reduction in flow is due to a reduction in the “airborne” demand in the network as a result of ground delay being taken. For example, the table shows that arc costs during t(1) are reduced by \$92,291 while ground delay costs add \$22,131. Ground delay is optimal for 94 market pairs during time steps 1 and 2. Only 44 market pairs would use ground delay during the third time step as the end of the constraint draws nearer. As the constraint clears at t(4), demand taking ground delay during previous time steps is added to the existing demand at t(4). The result is a surge in arc flow as illustrated in red in **Figure 24**. This four-time step event represents two hours of real time and has a total cost of about \$113,000.

F. Measuring the Sphere of Influence

A capacity/demand imbalance such as the scenario modeled in the previous example (section E) can impact demand patterns far from the source of the constraint. Demand that normally would have flown through the region is now either ground delayed or pushed out into neighboring regions. The extent to which demand is spatially displaced during an airspace constraint is referred to here as the “sphere of influence”. This section describes how NetFM can be used to predict and measure the sphere of influence for a given capacity/demand imbalance.

We investigated an actual NAS event involving an en route capacity/demand imbalance, modeled the constraint in NetFM, and finally compared the NetFM prediction to the actual sphere of influence as shown in **Figure 25**. The actual event being modeled was June 14, 2005, a day on which multiple Ground Delay Program (GDP) in support of Severe Weather Avoidance Plan (SWAP) initiatives were being implemented to deal with en route weather. The map on the right compares actual traffic volume (demand) during the 1800-0200Z hours on June 14, 2005 to normal demand levels for that time range during June and July 2005.

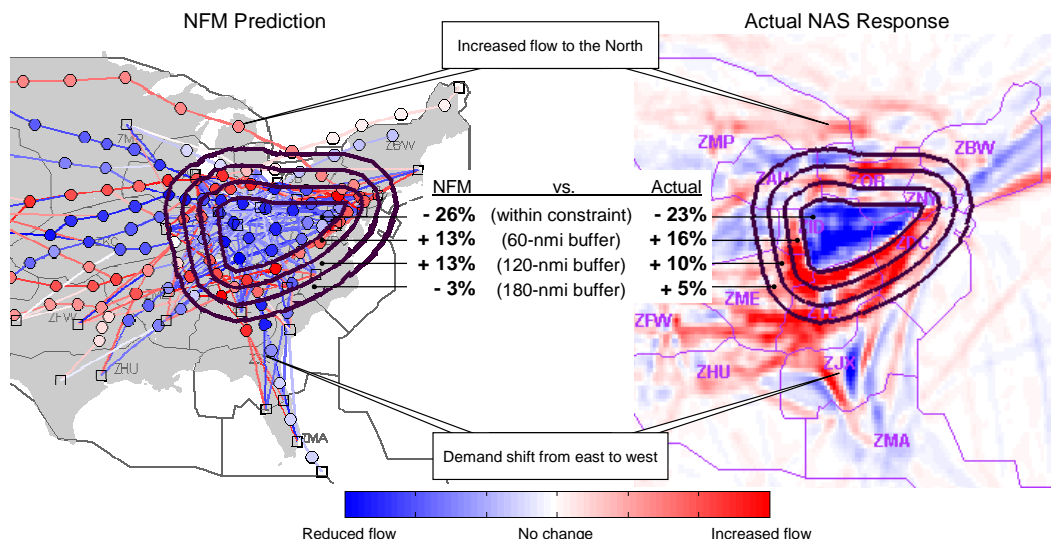


Figure 25. The sphere of influence of an actual airspace constraint is measured by comparing node usage within buffer zones relative to normal operations

Going with convention, blue indicates regions of relatively low traffic volume while red shows where volume is elevated. A polygon was then drawn to enclose the large region of reduced traffic volume in the center of the map. Traffic volume within the polygon was down 23% relative to normal usage. The polygon indicates the location of the constraint. This constraint was then modeled in NetFM as a 23% reduction in arc flow relative to normal operations within the same geographic polygon. The NetFM-predicted impact of the constraint was then compared to the actual NAS response observed on June 14, 2005.

Buffer zones were drawn around the constraint at 60-, 120-, and 180-nautical miles, respectively. Demand within each buffer zone was then compare with normal usage to identify the sphere of influence of the constraint. Applying the 23% reduction in arc flow within the modeled constraint resulted in a 26% reduction in node usage (demand) within the polygon. Just outside the constraint in the NetFM prediction, there is a 13% increase in demand which matches well with the 16% increase observed in the NAS. This increase in demand is due to flights being moved out of the constrained region and into adjacent airspace. The 120-nautical mile buffer zone shows a similar effect. At 180-nautical miles out however, the NetFM prediction suggests a slight reduction in demand signifying the sphere of influence in the optimal solution extends only to 120 nautical miles. The sphere of influence observed in the actual NAS response is about the same and tails off to a 5% increase in demand within the 180-nautical mile buffer zone. Other similarities between the NetFM prediction and actual NAS response can be observed. For example, elevated demand is observed in both maps over northern Minneapolis ARTCC (ZMP). Also, demand over central Jacksonville ARTCC (ZJX) appears to shift from east to west.

Numerous similarities appear when comparing the behavior of the NetFM optimal solution to the operational solution. This observation suggests we are on the right track to creating a working high-level model of how the NAS will respond to predicted constraints. Insight into the system-wide impact of localized events, such as weather disruptions, can be gained through applying NetFM to predict and evaluate the sphere of influence associated with capacity/demand imbalances.

G. Measuring the Impact of Spatial and Temporal Uncertainties in Weather Forecasts to Aid in Risk Management

TFM and airline planners deal with risks associated with predicted airspace constraints such as convective weather. There are two components to risk: likelihood and impact as shown in **Figure 26**. The likelihood of convective weather is an area of ongoing research pursued in studies concerning various weather forecast products [12][13][14][15][16][17][18][19][20]. Effective risk management also requires knowledge of the impact associated with each possible outcome. The impact of convective weather is less understood as it requires an understanding of the complex relationship between demand levels, airspace capacity, and network structure.




Figure 26. Risk table adapted from [11]

In this section, we demonstrate the application of NetFM for measuring the potential NAS-wide impact of possible weather scenarios taking into account spatial and temporal uncertainties in weather forecasts.

Weather Impact as a Function of Location

To go beyond modeling the impact of a single constraint, we configured NetFM to compute the impact of the same constraint placed at various locations throughout the NAS. This analysis was conducted using NetFM to solve a two-time step problem within each of the 65 five-degree by five-degree (300-nmi by ~230-nmi) grid cells depicted in **Figure 27**. The first time step involved a 25% reduction in arc flow with the cell. Arc flows were then set to their unconstrained state during the second time step. Each problem was solved independently based on normal 2200Z demand levels. Grid cell size was chosen to represent the approximate size of an actual weather disruption.

The figure shows the relative impact of the constraint at locations throughout the NAS as measured by NetFM. Red cells indicate regions where the impact is highest with cells 37, 25, and 49 being the top three. The figure makes clear that a constraint, such as weather, in the DC metro area (cell 37) has a greater impact than the same constraint located over Montana (cells 4 and 5).

These NetFM results agree well with our intuition regarding areas of the NAS that are sensitive to disruption. The importance of this NetFM application is that it identifies and quantifies the network impact of constraints, which could be modeled to have any shape, size, duration, or severity, on an equal scale throughout the NAS. This approach could be used to normalize for weather severity when comparing the operational costs of different TFM initiatives such as GDP SWAP, AFP, and MIT.

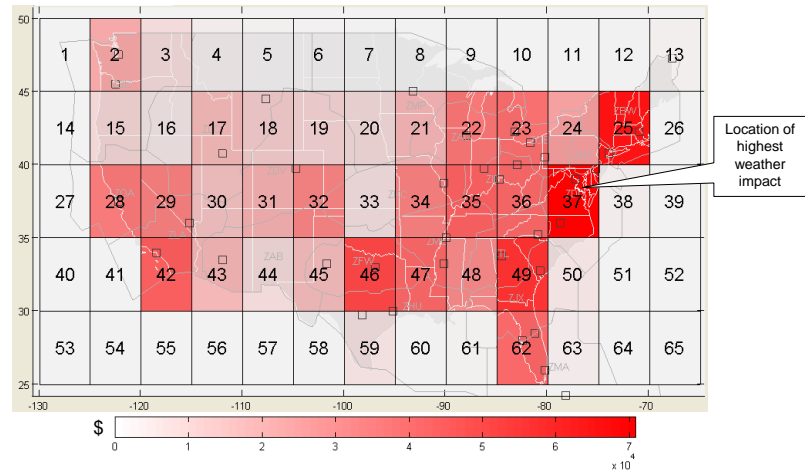


Figure 27. Network sensitivity as a function of constraint location based on modeling a 25% reduction in arc flow within each grid cell. The most sensitive regions appear in red.

Table 4 lists 20 cells with the highest cost increase along with additional parameters for each cell. The total cost in each cell before and after imposing an airspace constraint is a measure of the sensitivity of each cell to the modeled airspace constraint. Cells with the highest cost are those most sensitive to an imbalance between capacity and demand.

Table 4. Top 20 cells most sensitive to constraints based on 2200Z demand. Assumes \$60 per minute of airborne operations and ground-to-airborne cost ratio of 1/2.

Cell	Cost relative to normal operations			Ground Delay Market Pairs	Markets Inside Cell	Arcs In Cell	Arc Flow In Cell
	Reroutes (\$)	Ground delay (\$)	Total (\$)				
37	13,599	57,114	70,713	60	3	197	910
25	936	63,990	64,917	41	2	134	589
49	10,476	44,964	55,431	43	2	152	761
46	14,787	34,326	49,113	32	2	124	618
42	846	41,877	42,723	13	1	61	411
35	24,345	17,154	41,499	29	2	153	992
62	207	39,807	40,014	53	3	93	364
36	30,753	9,000	39,762	19	3	185	1186
47	15,633	23,949	39,582	36	2	121	670
23	23,733	14,049	37,782	25	4	152	830
22	28,782	6,849	35,622	5	1	114	873
28	279	34,245	34,524	22	1	52	304
29	12,483	20,097	32,580	31	1	88	851
34	14,967	17,163	32,130	26	2	104	792
32	17,649	13,248	30,897	10	1	78	701
48	10,107	19,278	29,394	24	1	112	548
24	21,519	1,989	23,499	6	0	183	690
2	0	21,231	21,231	24	2	54	189
31	12,519	6,012	18,531	5	0	55	730
21	17,046	90	17,127	1	1	64	596

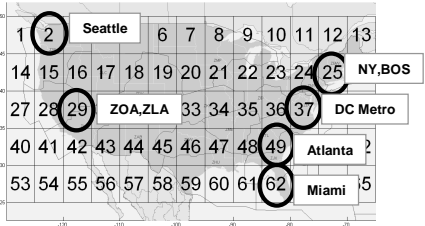
The table lists the added cost associated with the constraint being located in each cell. The most sensitive grid cell (by total cost) is number 37, which contains three market airports: Philadelphia, DC Metro, and Charlotte, NC. The impact of the 25% reduction in arc flow lasting 30 minutes is a combination of rerouting and ground delay costs. Reroutes and ground delay in cell 37 cost \$13,599 and \$57,114, respectively with a total cost impact of \$70,713. The number of ground delay market pairs listed in the table is the number of market pairs for which taking 30 minutes of ground delay is more optimal than the added distance to fly around the constraint. The number of arcs in each cell and arc flow in each cell are provided as a measure of the flow impacted by the constraint. For example, there are 197 arcs inside cell 37. These arcs normally carry 910 flights per hour. Under the constraint, each arc in cell 37 was limited to 75% of its normal flow, meaning total arc flow was reduced from 910 to 683 flights per hour.

Weather Impact as a Function of Location and Time

Weather impacts were measured as a function of location and time culminating in the results presented in **Table 5**. The table lists the impact of the same 25%-reduction constraint occurring at various locations and times in the NAS. For example, the constraint occurring in the New York / Boston area (cell 25) during the 1800Z hour would have an impact of 1,141 equivalent air minutes. However, if the same constraint (such as a weather system) arrived two hours later, then the impact would increase 35% to 1,539 equivalent air minutes. In this way, NetFM can be used to quantify the impact of temporal uncertainty.

Table 5. Weather impact based on a 25% flow reduction lasting 30 minutes at select times and locations

Cell	Region	Impact (equivalent air minutes)				
		0800Z	1600Z	1800Z	2000Z	2200Z
2	Seattle	20	317	504	352	404
25	NY / BOS	2	1,376	1,141	1,539	1,237
29	ZOA / ZLA	116	753	645	?	584
37	DC Metro	2	1,832	1,350	1,495	1,347
49	Atlanta	2	2,064	1,053	?	1,593
62	Miami	12	1,683	936	1,181	762



Impact of **temporal** uncertainty

Impact of **spatial** uncertainty

NetFM can also quantify spatial uncertainty. If weather predicted for 2000Z in the Atlanta area (cell 49) drifts south into the Miami area (cell 62), then the impact will be reduced by 26% from 1,593 to 1,181 equivalent air minutes.

Relative Impact Matrix

Another application of these results involves constructing a relative impact matrix (RIM) as shown in **Table 6** where the impacts of potential weather outcomes are defined relative to the impact of the expected outcome. The expected outcome could be interpreted from an existing weather forecast or may eventually refer to the expected value based on multiple scenarios given in an ensemble forecast [18].

Table 6. Relative impact matrix for weather predicted for 1800Z in the Atlanta area

Spatial Deviation	Cell	Relative Impact						
		Temporal Deviation (hours)						
		- 6	- 4	- 2	none	+ 2	+ 4	+ 6
none	49	1.23	1.77	1.96	1.00	1.51	1.00	0.77
N	36	0.84	1.52	1.14	0.82	1.10	0.72	0.73
NE	37	1.48	1.83	1.74	1.28	1.42	1.28	1.11
E	50	0.36	0.18	0.28	0.20	0.15	0.08	0.14
SE	63	0.09	0.16	0.18	0.14	0.07	0.04	0.02
E	62	1.33	1.39	1.60	0.89	1.12	0.72	0.66
SW	61	0.02	0.02	0.04	0.02	0.04	0.01	0.01
W	48	0.31	0.56	0.55	0.26	0.88	0.53	0.49
NW	35	0.56	0.87	0.67	0.32	1.08	0.75	0.76

When the actual weather matches the expected weather in both location and time, then the spatial and temporal deviations are none and the relative impact is 1.00. If the weather arrives two hours early, however, then the table shows there is a 96% increase in the impact (1.96). Alternatively, the relative impact diminishes to 82% (0.82) if the weather drifts, say, north.

It is understood that temporal deviations in weather forecasts seldom reach six hours. However, this type of table could be adapted to provide useful information to TFM planners or other decision makers when managing risks associated with uncertain weather. In particular, this type of display could help define the impact axis of a risk table (**Figure 26**), giving decision makers the ability to weigh the potential impacts of possible weather outcomes when selecting the most appropriate strategy in the face of uncertain weather.

IV. Conclusion and Next Steps

Many questions arise in TFM planning when dealing with capacity/demand imbalances in the NAS. For example, what is the optimal strategy for scheduling and routing traffic around such constraints? At what point does the extent of a constraint require the use of both rerouting and ground delays? What is the unavoidable cost of a particular constraint? We have presented in this paper a semi-dynamic, time-varying, multi-commodity network flow model representation of the NAS that begins to answer these types of questions. This model is unique in that its underlying architecture of nodes, arcs, and markets was derived from operationally observed conditions. Furthermore, node capacities were set based on observed sector usage rates in an effort to create a high-level model that responds to demand and capacity constraints in ways that appear similar to the operational NAS.

The application of NetFM was demonstrated for multiple scenarios including quantifying the inherent cost of NAS events, predicting the extent to which demand fluctuations propagate outwards from an airspace constraint, and measuring the impact of spatial and temporal uncertainties in weather forecasts.

V. Future Research

NetFM currently models demand as a continuous flow from source to sink, similar to how electrical current flows through a circuit. In order to more accurately model the interactions between dynamic demand and capacity, a focus of future research is on the implementation of a fully-dynamic network flow model in which demand passes through a grid of evenly-space hex cells.

We are also working on algorithms for disaggregating the flow-based solutions produced by NetFM into flight-specific solutions. In this way, we will be able to use NetFM to optimize a schedule of individual aircraft based on dynamic demand and capacity constraints.

A major step towards readying NetFM for operational use will be translating the routing solutions defined in the NetFM architecture of nodes and arcs to the operational routing structure of jet routes, airways, standard instrument departures (SIDs), coded departure routes (CDRs), standard terminal arrival routes (STARs), and alike.

The existing architecture of nodes and arcs now represents a single layer of en route sectors. Addition layers could be added to the network to represent flows at different altitudes. An appropriate cost could be assigned to arcs that transition between altitude layers to reflect pilot preference to maintain a fixed altitude.

We are also pursuing continued comparisons between the NetFM and operational solutions to extend validation and to pinpoint problem areas in the NAS to help focus future investments and estimate potential returns. For example, identify ideal locations for installing Automatic Dependent Surveillance – Broadcast (ADS-B) ground stations to provide enhanced capacity corridors with the best system effects.

The application of NetFM for measuring the impact of uncertainty in weather forecasts could be extended to take into account weather probability. This research would help to identify regions of the NAS most likely to experience high-cost constraints.

VI. Acknowledgements

This research was funded by Ved Sud of the FAA under the Results National Contract Vehicle's SOW Task Area 1: Airspace Congestion Management and by Kareena Nair of the FAA under the Collaborative Air Traffic Management Technologies (CATMT) Concept Engineering and Development (CED) Program. Special thanks to Mark Klopfenstein of Metron Aviation for his guidance during this research.

VII. References

- [1] Bertsimas D., Odoni, A., "A critical survey of optimization models for tactical and strategic aspects of air traffic flow management," NASA/CR-97-206409, 1997.

- [2] Ball, M. O., Barnhart, C., Nemhauser, G., Odoni, A., "Air Transportation: Irregular Operations and Control," *Handbook in Operations Research and Management Science*, Vol. 14, C. Barnhart and G. Laporte (eds.), pp. 1-73, Elsevier, 2007.
- [3] Odoni, A., "The Flow Management Problem in Air Traffic Control," *Flow Control of Congested Networks*, A.R. Odoni, L. Bianco and G. Szego, 1987, pp. 269-288, Springer-Verlag, Berlin.
- [4] Bertsimas, D., Stock, S., "Air Traffic Flow Management Problem with Enroute Capacities," *Operations Research*, Vol. 46, No.3, 1998, pp. 406-422.
- [5] Bertsimas, D., Stock, S., "The Traffic Flow Management Rerouting Problem in Air Traffic Control: A Dynamic Network Flow Approach," *Transportation Science*, Vol. 34, No.3, 2000, pp. 239-255.
- [6] Schwartz, Albert, "FAA/EUROCONTROL Cooperative R&D Action Plan 5 and Action Plan 9: Traffic Flow Management in Fast-time Simulation: Current and Future Capabilities", 2006
- [7] van Tulder, P. Berge, M. Repetto, B. Haraldsdottir, A. Moerdyk, D., "Airline Schedule Recovery in Collaborative Flow Management and Weather Forecast Uncertainty", *Digital Avionics Systems Conference*, 2004.
- [8] MacLaren, Jon, "Network Optimization Models", *National Infrastructure Simulation and Analysis Center*, 2005.
http://www.sandia.gov/mission/homeland/factsheets/nisac/RNAS_factsheet.pdf#search=%22NISAC%20RNAS%22 [retrieved September 1, 2006].
- [9] Ahuja, R. K., Magnanti, T. L. and Orlin, J. B., "Network Flows", Prentice-Hall, New Jersey, 1993.
- [10] Krozel, J., Penny, S., Prete, J. M., and Mitchell, J. S. B., "Automated route generation for avoiding deterministic weather in transition airspace," *Journal of Guidance, Control, and Dynamics*, 30(1), 2007, pp. 144-153.
- [11] "NAS System Engineering Manual", Version 3.1, Section 4.10, June 6, 2006,
<http://www.faa.gov/asd/SystemEngineering/SEM3.1/Section%204.10.pdf>, [retrieved June 15, 2007].
- [12] Benjamin, S. G., Brown, J. M., Brundage, K. J., Devenyi, D., Grell, G. A., Kim, D., Schwartz, B. E., Smirnova, T. G., Smith, T. L., Weygandt, S. S., and Manikin, G. S., "RUC20 - The 20-km version of the Rapid Update Cycle." *NWS Technical Procedures Bulletin*, 490, 2002.
- [13] Evans, J. E., and Ducot, E. R., "The Integrated Terminal Weather System (ITWS)." *Lincoln Laboratory Journal*, 7(2), 1994, pp. 449-474.
- [14] Evans, J. E., and Ducot, E. R., "Corridor Integrated Weather System." *Lincoln Laboratory Journal*, 16(1), 2006, pp. 59-80.
- [15] Kay, M. P., Mahoney, J. L., and Hart, J. E., "An analysis of CCFP forecast performance for the 2005 convective season." *12th Conference on Aviation, Range and Aerospace Meteorology*, January 29-February 2, 2006, Atlanta, GA, 2006.
- [16] Krozel, J., Murphy, J. T., Mueller, C., Penny, S., Smith, P., and Spencer, A., "Weather data and forecast product literature survey." Technical Report, Metron Aviation, Inc., Herndon, VA, 2007.
- [17] Myers, T., Khorrami, B., Cross, C., Kierstead, D., "Airspace Congestion Management Research", Technical Report for FAA 32F0907-014-R0, Metron Aviation, Inc., Herndon, VA, 2007.
- [18] Steiner, M., "Ensemble Forecasts." *2nd ATM-Weather Impact Workshop*, Boulder, CO, March 3-4, 2008.
- [19] Wolfson, M. M., Forman, B. E., Calden, K. T., Dupree, W. J., Johnson, R. J., Boldi, R. A., Wilson, C. A., Bieringer, P. E., Mann, E. B., and Morgan, J. P., "Tactical 0-2 hour convective weather forecasts for FAA." *11th Conference on Aviation, Range and Aerospace Meteorology*, Hyannis, MA, 2004.
- [20] Wolfson, M. M., and Clark, D. A., "Advanced aviation weather forecasts." *Lincoln Laboratory Journal*, 16(1), 2006, pp. 31-58.