

J12.4 TRANSLATING WEATHER INTO TRAFFIC FLOW MANAGEMENT IMPACTS FOR NEXTGEN

Tenny Lindholm *, Robert Sharman
STAR Institute, Boulder, CO

Jimmy Krozel, Victor Klimenko, Shubh Krishna, Nathan Downs
Metron Aviation, Inc., Herndon, VA

Joseph S. B. Mitchell
State University of New York, Stony Brook, NY

1. INTRODUCTION

The uncertainties of weather, and its uncertain impact on traffic flow management (TFM), significantly affect the overall performance of the national airspace system (NAS). The NextGen vision looks toward resolving weather impacts through more precise weather forecast information integrated into Decision Support Tools (DSTs) and weather information assimilated into NextGen decision loops for air traffic management (ATM).

The first step toward mitigating weather impact is to define the impact itself—that is, by type of hazard and the preferred pilot, airline, and aircraft response to that hazard. This has been pursued to date with greatest emphasis for convective weather hazards. However, we present weather impact models for turbulence and in-flight icing to gain further insight into a more general weather hazard model for NextGen. Our model includes convective, turbulence, icing, and other hazard types in a Weather Impact Interaction Grid (WIIG). In particular, we describe:

- Incorporation of a Convective Weather Avoidance Model (CWAM) as described in the literature for en route and in the terminal area,
- Incorporation of a Clear Air Turbulence (CAT) Weather Avoidance Model (CATWAM) for en route, based on actual observed pilot and aircraft responses to Graphical Turbulence Guidance (GTG) forecast data,
- Incorporation of in-flight Icing Weather Avoidance Model (IWAM) for expected pilot and aircraft responses in the terminal airspace,
- An important classification of weather hazards as either hard or soft constraints:
 - Hard constraints are formed by weather hazards through which no aircraft can safely fly (e.g., severe convection, turbulence or icing).
 - Soft constraints are formed by weather hazards through which some pilots or airlines decide to fly, while others do not (e.g., moderate turbulence or icing).
 - Mathematical modeling of deterministic and probabilistic weather forecasts, as well as the modeling of deterministic and probabilistic weather impacts.

Second, through modeling the weather hazard and simulation of future, NextGen traffic flows within NASA's Future ATM Concepts Evaluation Tool (FACET), we draw conclusions about the systemic affects of the individual preferred responses. Through these simulations, we begin to define TFM strategies that will help mitigate both individual and NAS-wide TFM impacts in the NextGen environment. Results from the simulations demonstrate the NAS impacts of imprecise weather forecasts and uncertain aircraft responses. We conclude the paper by suggesting how such a simulation capability can help develop requirements for more precise weather forecasts for NextGen.

2. NEXTGEN CONCEPT OF USE

A NextGen Concept of Use guides our study of weather translation to ATM impact (Lindholm, 2009). The concept is consistent with the NextGen Concept of Operations (JPDO, 2007), and how weather is integrated into future NAS trajectory-based operation (TBO). NextGen is foundationally based on TBO and Trajectory Management (TM). They apply any time an air vehicle is in a 4-Dimensional (4D) environment; that is, from liftoff to touchdown, as described by 4D Trajectories (4DTs).

The Air Navigation Service Provider (ANSP) will handle both individual 4DTs and aggregate flows representing many aircraft trajectories. For the management of convection, turbulence, and icing constraints, flexible 4D route definitions allow traffic flows to be shifted as necessary around regions of airspace constrained by weather hazards to enable more effective weather avoidance and manage demand into and out of the arrival/departure environment. Capabilities for managing airspace structures include a common mechanism for implementing and disseminating information on the current airspace configuration to ensure that all aircraft meet the performance requirements for any airspace they enter. The Concept of Use defines a 4D Flow Constrained Area (FCA) to specify that special requirements apply for aircraft entering into the FCA and that flows are designed to optimize the FCA capacity. Furthermore, users must specify their individual flight limitations and preferences as inputs to flight planning and execution, and flight operators may dynamically update these features. With this input, the ANSP can support 4DTs tailored to individual flight preferences. Figure 1 illustrates the concept of a FCA and how it fits into the NAS when weather presents a constraint to normal traffic flow.

* Corresponding author address: Tenny A. Lindholm, 3125 Sterling Circle, Unit 107, Boulder CO 80301; email: lindholm@starinst.org.

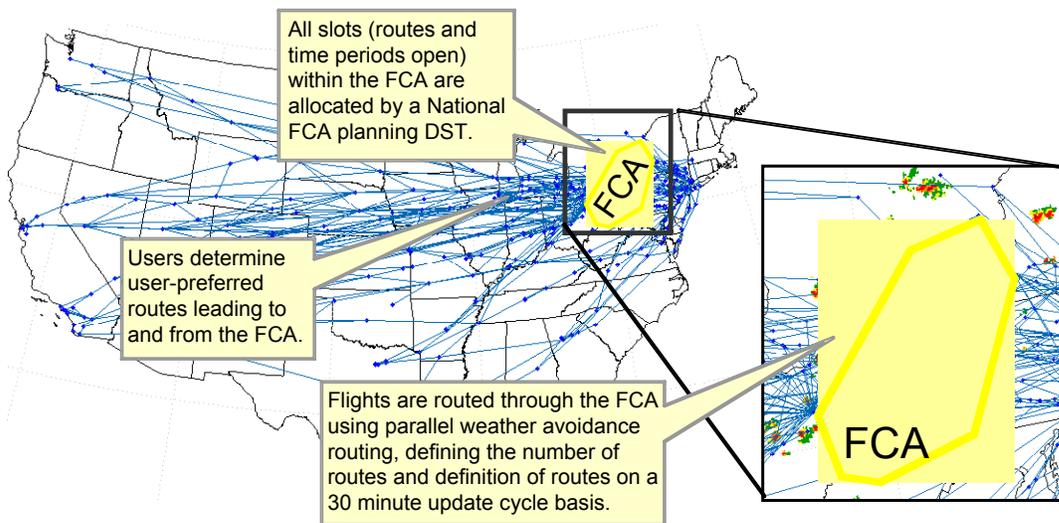


Figure 1: Routes within the FCA are defined by an FCA planning DST; routes outside the FCA are determined by routing preferences of the user.

3. CONVECTIVE WEATHER AVOIDANCE MODEL (CWAM)

Because convective weather is the leading cause of delays in the NAS, much of the previous research has focused on its impact to TFM. To understand the criteria pilots use to avoid convective weather, research (Rhoda, 1999) has shown that pilots generally avoid NWS Level 3 and higher weather cells (reflectivity greater than 41 dBZ). Later research (DeLaura, 2006) shows that altitude of cloud tops in severe storms is also an important factor that pilots consider in determining which storm cells to avoid. Figure 2 illustrates the relationship between the level of convection (expressed in terms of Vertical Integrated Liquid (VIL)) and the echo top relative to whether the pilot deviates or continues. This weather translation model has been referred to as the Convective Weather Avoidance Model (CWAM) [CRD07]. CWAM produces Weather Avoidance Fields (WAFs) (Figure 3) which can be used by DSTs to determine the impact of convective weather on ATM.

In order to determine the impacts of convective weather on en route operations, airspace is portioned into passable and impassable regions. CWAM calculates WAFs as a function of observed and/or forecast weather, to determine 2D or 3D grids whose grid points are assigned either a probability of deviation or a binary deviation decision value (0 or 1).

4. HARD AND SOFT CONSTRAINTS AND THE WEATHER IMPACT INTERACTION GRID (WIIG)

The Weather Impact Interaction Grid (WIIG) may assist NextGen planners involved with the integration of weather hazards into ATM planning. The WIIG provides the framework for defining the impact of aircraft class and capabilities against varying weather phenomena, which in turn guides the development of the most efficient risk management strategies.

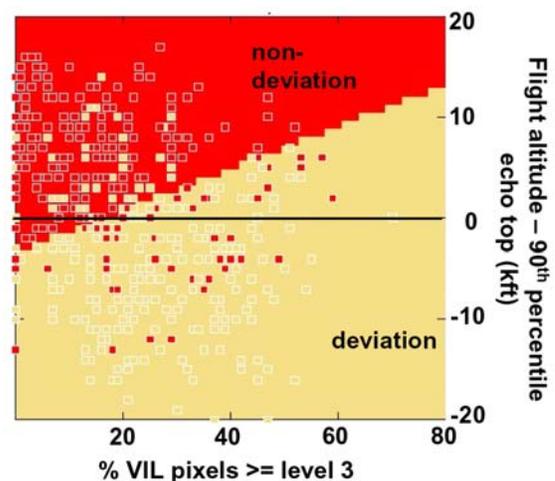


Figure 2: Convective Weather Avoidance Model (CWAM); source: MIT Lincoln Laboratory.

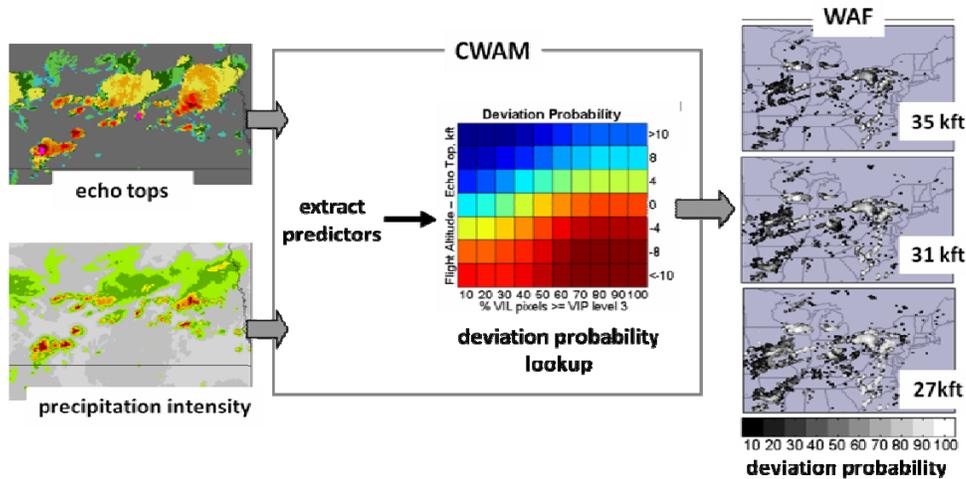


Figure 3: CWAM implementation to create Weather Avoidance Fields (WAFs); source: MIT Lincoln Laboratory.

Not all aircraft respond to weather hazards in the same way. Therefore, the WIIG has two dimensions:

- Class of Aircraft and
- Weather Hazard Type.

Two classes of aircraft are considered relative to the response to hard and soft constraints:

- Class 1 aircraft are those that avoid both hard and soft constraints.
- Class 2 aircraft are those that avoid hard constraints but are still willing to fly through soft constraints.

More aircraft classes may be defined to account for response strategies to weather hazard types.

Aircraft class is determined by a number of factors and not just type of operation. For example, a passenger-carrying jetliner may respond as a Class 1 aircraft under one airline's policy and a Class 2 under another set of rules. General Aviation (GA) aircraft will operate as either Class 1 or 2 depending on user preferences. Factors defining aircraft class include

- Rules of operation (VFR, IFR, Part 91, Part 121 or 135)
- Aircraft equipment and performance limitations
- Airline policies
- Personal pilot preference.

For example, Table 1 defines the constraint type based on whether an aircraft is certified for flight into icing conditions. Class 1 avoids hard and soft constraints; Class 2 avoids hard constraints but can fly through moderate or less severe icing (although aircraft flight manuals specify a time limit by aircraft type). Severe icing (SIGMET icing in particular) is always a hard constraint.

The concept is illustrated in Figure 4, where three Classes of aircraft and four types of weather events are depicted. Example routes (dark blue, light blue, magenta) are shown to illustrate possible ways Class 1, 2, and 3 aircraft, respectively, can transverse the region

of weather events of Types 1-4. All Classes of aircraft must avoid the hard (red) constraints, so no viable routes are allowed to pass through such a constraint.

Severity	Constraint Type for Uncertified Aircraft	Constraint Type for Certified Aircraft
Light Icing	Hard Constraint	No Constraint
Moderate Icing	Hard Constraint	Soft Constraint
Severe Icing	Hard Constraint	Hard Constraint

Table 1: Weather constraint types for different levels of icing

However, each Class of aircraft has a different threshold for the soft constraints (shades of grey). Consider four types of weather events. Types 1 through 3 are soft constraints, and Type 4 is a hard constraint. We have already noted that all Classes of aircraft must avoid Type 4; this is indicated with an x symbol in the WIIG. The soft constraint of Type 3 (dark grey) must be avoided by Class 3 aircraft, but may be traversed by Class 1 or 2 aircraft (if they benefit from doing so). The soft constraint of Type 2 (medium grey) must be avoided by Class 1 and 2 aircraft, but may be traversed by Class 3 aircraft (if there is a benefit in doing so). Also, the soft constraint of Type 1 (light grey) must be avoided by Class 2 and 3 aircraft, but may be traversed by Class 1 aircraft (if there is a benefit to do so). Given this WIIG established by the preferences of the pilots and airlines, and certain regulatory requirements, the problem is to identify a set of routes crossing the airspace such that the demand is satisfied for all Classes of aircraft given the constraints imposed in the WIIG.

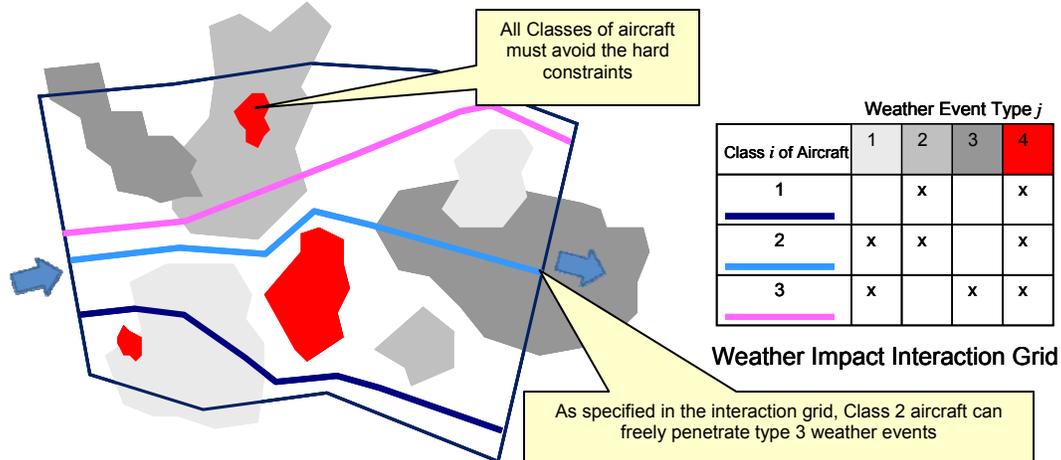


Figure 4: Capacity computation for three Classes of aircraft (1, 2, 3) among hard (red) and soft constraints (shades of grey) of multiple types (1, 2, or 3).

The WIIG models all known weather hazards as hard and soft constraints. Some considerations for modeling the constraints are:

- Hard constraints close airspace; soft constraints have varying degrees of “softness” based on user preferences or aircraft limitations.
- Known hard constraints such as 4D SIGMETs only impact TFM during their associated valid times and for the affected airspace. By definition, SIGMETs close airspace. Today, SIGMETs are valid for four hours unless rescinded or revised by another SIGMET. In NextGen, SIGMETs will be more dynamic in 4D space and time similar to the GTG and Current Icing Product (CIP). The GTG and CIP are gridded, 4D, high spatial and temporal resolution weather hazard products for turbulence and in-flight icing (FAA, 2008). A SIGMET in NextGen may be valid from a specific start time to a specific finish time, defined to the minute or with 5 minute accuracy, and for geometric size constraints defined to some NextGen design constraint. Figure 5 shows how GTG severity contours are translated to hard and soft constraints, red indicating severe turbulence or greater. CIP translations are done in a similar way.
- Soft constraints will be probabilistic and 4D. The behavior of the pilot is also uncertain. What is most relevant (and perhaps more easily modeled and operationally implemented) is the conditional probability of response impact x given a deterministic forecast of y . This situation can include the case where a forecast predicts no hazard and one actually is experienced. For example, what is the probability an aircraft will broadcast an intent to descend when the pilot is presented an area of moderate turbulence just ahead at its flight level? Or, what is the probability that a certain flight number/carrier will strategically plan around a forecast weather hazard? An

alternative is to treat soft weather constraints as probabilistic and then the relevant question is: Given a 40% probability of moderate turbulence, what is the probability of [response A, B, C, ...]?

- A collection of weather forecasting tools for convection, turbulence, icing, etc. would imbed weather hazard forecast and uncertainty information into the 4D weather cube for airline and ANSP use as a common definition of the weather hazard.

5. MATHEMATICAL MODEL DEVELOPMENT

5.1 Deterministic Case

The software solution related to deterministic ATM-weather impact examples is described next. The mathematical theory related to this work has been previously documented in Krozel, 2008.

A potential solution methodology is the Hard-Soft Constraint Solver (HSCS). Input to the HSCS algorithm includes the region R , a FCA as described above, and a source/sink pair of edges of R where the flow begins at the source and ends at the sink. The HSCS determines if it is possible to route I lanes of Class-1 aircraft and J lanes of Class-2 aircraft, with each lane entering/exiting the region R at the source/sink edge. The lanes for Class-1 (resp., Class-2) aircraft are of width w_1 (resp., w_2), where the lane width parameter includes the effect of both Required Navigation Performance (RNP) and of horizontal separation between containment zones. Figure 6 is an example HSCS output from processing a CIP icing map. Lanes for Class-1 aircraft must avoid hard constraints (red or magenta in Figure 6); lanes for Class-2 aircraft must avoid hard constraints but must also avoid soft constraints (blue in Figure 6).

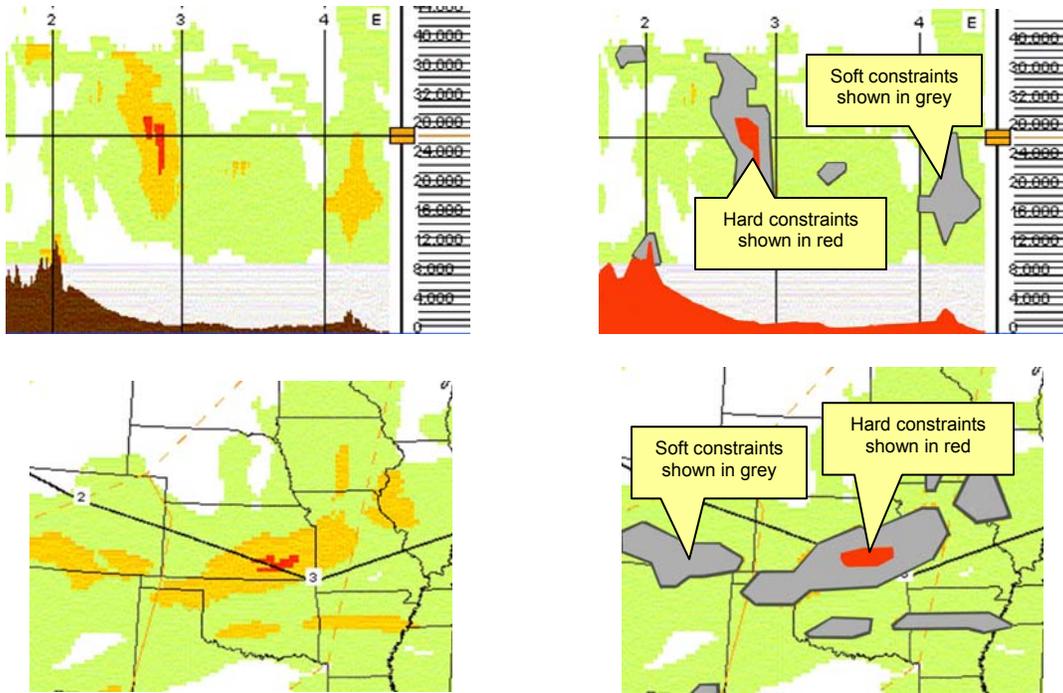


Figure 5: Example of translation of GTG severity contours to hard and soft constraints, vertical cross-section on the top and horizontal cross-section on the bottom.

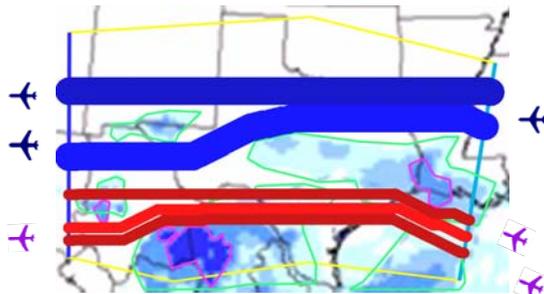


Figure 6: Example application of HSCS for planning East-West flow through icing.

5.2 Probabilistic Case

The theory and software related to probabilistic ATM-weather impact examples is described next. The HSCS determines if I lanes of Class 1 and J lanes of Class 2 aircraft can be routed through the region of interest, R , under the forecasted set of weather constraints. In reality, the weather constraints are not known in advance (and may not even be known with precision in a nowcast). The HSCO mathematical model allows for probabilistic weather forecasts in which there is an explicit stochastic model of uncertainty in the weather prediction. Then, the HSCS determines from the probabilistic forecast a set of probabilities associated with the routability of lanes of Class 1 and Class 2 aircraft.

Specifically, our model is currently formulated with an ensemble-based specification of weather forecast uncertainty. Instead of using a single forecast, F , that

specifies a set of hard and soft constraints, in the ensemble model there is a set $F = \{F_1, F_2, \dots, F_K\}$ of K forecasts, each consisting of a set of hard and soft constraints. Associated with forecast F_i is a probability, p_i , which estimates our belief that F_i will be the observed weather. Without any prior bias, the probabilities p_i are assumed to be the same ($1/K$), representing the *uniform prior case*. (This is the default case.) For any finite set F of forecasts that constitute a discrete ensemble, the probability that any one forecast is the actual observed weather is zero, since the set of all possible weather maps is a continuum, with an uncountably infinite set of possibilities, any one of which has probability zero. The discrete set F , together with the probabilities p_i , serves as an approximation of the continuum, giving a discrete set of sample instances of weather. Abstractly, think of each member of the ensemble F as representing a small neighborhood in the continuum of all weather instances. These neighborhoods may not represent a covering of the space of all possible instances; in general, they will not, and there will be some small probability that the observed weather is not close to *any* of ensemble members.

We have developed a simple method to generate random instances of weather maps (with hard and soft constraints) from a given seed forecast, F . First, we do not have access to rich ensembles of weather forecasts for the phenomena (turbulence, icing) in our experiments (although this may be available in NextGen), and, second, even if given discrete ensemble sets of forecasts, our random generation method allows us to generate instances of weather constraints in a

(continuous) neighborhood of any one ensemble forecast, F_i , since we utilize continuous probability distributions to generate a set of constraints that is “close” to those of the seed forecast.

Our random generation method is as follows. Suppose that the seed forecast F consists of a set, $\{P_1, P_2, \dots\}$ of polygons corresponding to hard constraints and a set $\{Q_1, Q_2, \dots\}$ of polygons corresponding to soft constraints. Our goal is to generate a forecast, F' , that is close to F . We describe the method for generating a set of hard constraints that are close to the seed set $\{P_1, P_2, \dots\}$; similarly, we generate a set of soft constraints that are close to the seed set $\{Q_1, Q_2, \dots\}$. Together, these hard and soft constraints constitute the forecast F' .

For a user-specified (or randomly generated, based on a user-specified range of values) integer parameter, m_P , we generate m_P points, a_j , at random (e.g., according to a uniform distribution) within the union of the polygons P_i (the set of hard constraints). In order to allow for the possibility of having some hard constraints of F' in regions far from the polygons P_i , we allow the randomly generated points a small probability, p_P , of falling outside the union of the polygons P_i . Then, for each random point a_j , we randomly generate a polygon, $P(a_j)$ (or a disk) centered at a_j . The size/radius of $P(a_j)$ is generated randomly according to a distribution; e.g., we use a uniform distribution over a user-specified interval, (r_1, r_2) , for $0 \leq r_1 \leq r_2$. The size/radius of $P(a_j)$ controls the closeness of the forecast F' to the seed forecast F . The number m_P and the probability p_P also influence the closeness. If m_P is huge (approaching infinity), and the size/radius of $P(a_j)$ and the probability p_P are tiny (approaching zero), then the union of the randomly generated polygons $P(a_j)$ is very nearly the same as the union of the hard constraints P_i of the seed forecast.

Figure 7 shows an example of a seed forecast F given first as a map (the composite GTG forecast), then with hard/soft constraints extracted, based on thresholding at two different intensities. Also, six instances of a randomly generated forecast, F' , are shown based on our method. Here, the points a_j are uniformly randomly distributed within the hard/soft constraints, and the random polygons $P(a_j)$ are generated as random quadrilaterals as follows: for a specified pair of radii, $r_1 \leq r_2$, a point at a uniformly random distance $r \in (r_1, r_2)$ from the center point a_j is generated in each of the four quadrants with respect to a_j ; these four random points specify the vertices of the random quadrilateral. The hard constraints of ensemble member F' are shown as red polygons (quadrilaterals), while the soft constraints are blue-green.

In order to do probabilistic reasoning with the ensemble forecast, we run the deterministic version of

the HSCS algorithm for each ensemble member F_i . For a given permutation of Class 1 and Class 2 air lanes (e.g., (1,1,2), indicating a south to north ordering of a Class 1, then Class 1, then Class 2), the HSCS algorithm computes whether or not the permutation is routable for forecast F_i . Then, the probability that the permutation is routable is computed as the sum of the probabilities p_i associated with routable instances F_i . The probability that I lanes of Class 1 and J lanes of Class 2 are routable (in *any* order) through the airspace region R is obtained by taking the probability associated with the maximum-probability permutation having I Class 1 lanes and J Class 2 lanes. This probability represents the probabilistic capacity estimate, based on the demand that is expected to use the airspace.

6. Turbulence Impact Analysis

Quantification of the impact of severe turbulence on the NAS resources is a complex problem because of the following difficulties:

- **Accurate Forecasting of Turbulence.** CAT cannot be seen and does not accompany any visible phenomenon (like a thunderstorm in the case of CIT) and thus the stakeholders have to rely on turbulence forecasts (that are imprecise by nature and based on scientific models with known or unknown deficiencies) and/or PIREPs (that are currently subjective and sporadic).
- **Varying Spatial and Temporal Extent of a Turbulence Event.** A single turbulence event can be contained within a single sector at one or two flight levels or may affect large volumes of airspace – tens of NAS sectors and flight levels at once and to varying degrees (also changing with time). Besides, different sector/altitude levels have different air traffic densities and thus different overall importance to the NAS.
- **Understanding the Factors that affect the Pilot Response to Turbulence.** Multiple factors need to be identified, understood, quantified (if possible) and taken into account for building a turbulence impact model.
- **Generalization of Pilot Responses to the NAS.** The turbulence impact model built from a limited number of turbulence situations in certain parts of the NAS at certain times of the year may or may not be projected into the impact on the whole NAS for the whole year. This is an open research question.

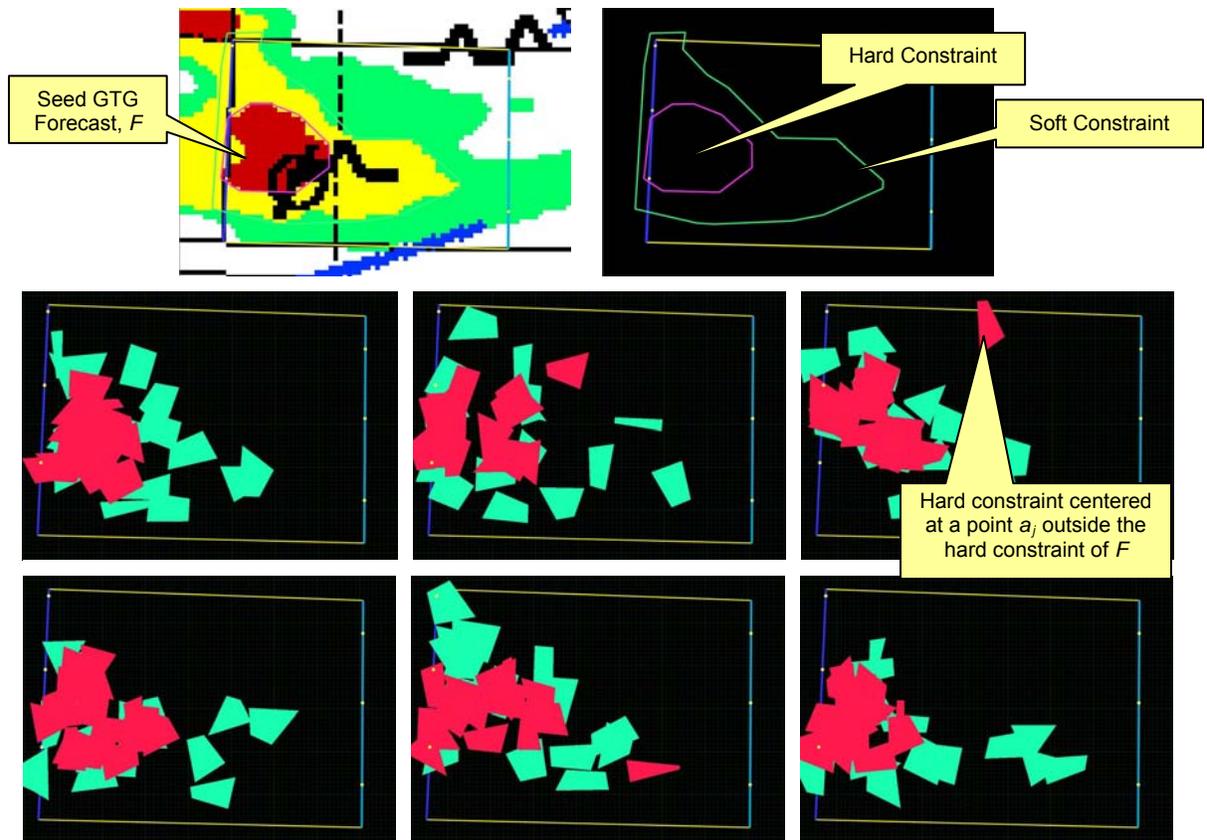


Figure 7: The seed forecast is thresholded to identify hard and soft constraints (one hard, one soft), and then used to randomly generate ensemble forecasts that are “close” to the given seed forecast.

6.1 Method of Analysis

We begin our analysis by (a) using the probabilistic severe or greater (SoG) GTG (Sharman, 2006) maps to quantify the probability of a severe turbulence encounter in a given sector/flight level and then (b) building a quantitative relationship between the observed air traffic behavior in this sector/flight level and these probability figures. In other words, knowing the probability of a SoG turbulence encounter for every sector/flight level along the flight path for every flight in the selected group, we accumulate the statistics of the type, magnitude and frequency of the pilots’ turbulence-avoidance maneuvers – and the ensuing changes in air traffic density – as a function of the probability of severe turbulence encounter that caused them.

We express the probability of a severe turbulence encounter in a given NAS sector at a given flight level for a given GTG cutoff (green, yellow or red) as the percentage of the sector area covered by the corresponding GTG color and “higher” colors. For example, the probability figures for flight level FL390 of ZDV3500 sector (Figure 8) are 7% for the “red” GTG cutoff, 28% for the “yellow”, and 60% for the “green”.

This way of describing turbulence in a given sector/flight level allows us to establish clear relationships between

turbulence and its impact on the NAS and provides useful ballpark impact estimates that can be done in a real-life turbulence situation.

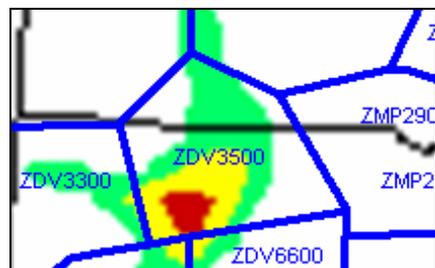


Figure 8: A section of the probabilistic SoG GTG map for 1800Z, Jan. 24, 2007, FL390.

Our analysis is based on two severe turbulence events occurring on Oct. 28, 2006 in Central Eastern US, and Jan. 24, 2007 in the Colorado Rockies. In both cases, we limit our scope to the 1500Z–1900Z timeframe, and the 24,000 ft – 45,000 ft altitude band in the first case and the 30,000 ft – 45,000 ft in the second. To collect flight data, in each case we set up a rectangular area (magenta in Figure 9) that

encompassed all the red/yellow/green high-probability-of-turbulence regions on probabilistic SoG GTG maps.

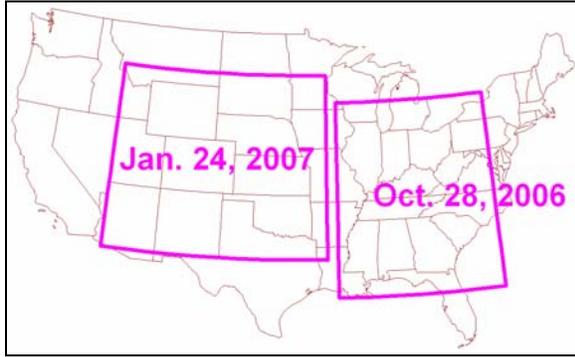


Figure 9: Geographical areas covered by the turbulence events studied in the present analysis.

Even though our analysis is based on just two turbulence events, we believe that the very scale of them provides us with enough data to draw conclusions about the impact of severe turbulence on the NAS in general, as these two events covered 8,200+ flights flying through several hundred NAS sectors with various forecasted turbulence probabilities over the course of several hours..

Since some sectors cover just a few flight levels (e.g., FL330 through FL360), we combined all sectors into vertical stacks, each spanning the whole altitude band of interest (FL240 to FL450 for the first turbulence event, and FL300 to FL450 for the second) and consisting of one, two or three sectors on top of one another. For each stack/flight level, we calculated the percentages of area covered by the red, yellow, or green color on probabilistic GTG maps - i.e., the percentage of area of a given flight level in a given sector for three different thresholds for agreement among GTG indices predicting severe turbulence. As previously discussed, we treat these numbers as representing the relative probabilities of encountering severe turbulence in a given sector/flight level.

6.2 Impact on Air Traffic Density

Now, we analyze the reduction in air traffic through a given 3D region of airspace due to severe turbulence (encountered and/or anticipated).

For our study, a 3D region of airspace is defined by a NAS sector area and a group of 1, 2, 3, etc. odd or even flight levels inside the sector between the minimum and maximum altitudes of the altitude band under consideration. For each sector and each altitude level between the lowest and the highest, the archived Enhanced Traffic Management System (ETMS) data were used to calculate average instantaneous *flight level occupancies* (the average number of aircraft in a given flight level at an arbitrary moment of time within the specified hour) for each hour. For 3D volumes containing more than one flight level, the flight level occupancies for all flight levels were added together.

The same parameters were also calculated for these 3D volumes and time intervals for the same day of

the week one week before and one week after the turbulence event to provide a comparison to a “no turbulence” baseline. (GTG maps for those days were obtained to ensure the usability of those days for this purpose).

The goal of this part of analysis is to link the changes in the density of air traffic through a given sector/flight level to the forecasted probability of a severe turbulence encounter in this sector/flight level and other parameters, such as volume, time.

The change in air traffic through a given 3D volume,

$$\Delta N = (N_{\text{turbulence}} - N_{\text{no turbulence}}) / N_{\text{no turbulence}} \quad (1)$$

is dependent on the following four parameters:

(a) threshold for GTG index agreement, a_{GTG} - i.e., the minimum percentage of GTG indices that should agree on predicting severe turbulence (these thresholds are depicted by red, yellow or green on GTG maps),

(b) relative probability p_{turb} of a severe turbulence encounter in a given sector/flight level - expressed as the percentage of a flight level area covered by a given GTG color (for multi-flight-level volumes it is averaged over all flight levels in a group),

(c) vertical size n_{level} of a 3D volume of interest (1, 2, 3, etc. even or odd flight levels)

(d) time interval Δt used to average the turbulence probability and the flight level occupancies in a given 3D volume (1, 2, 3 or 4 hours)

The change in air traffic density due to severe turbulence is a multi-dimensional surface:

$$\Delta N = \Delta N(a_{\text{GTG}}, p_{\text{turb}}, n_{\text{level}}, \Delta t) \quad (2)$$

An exponential-decay-based mathematical model could be used to fit the experimental data:

$$\Delta N = A(a_{\text{GTG}}, n_{\text{level}}, \Delta t) \cdot \exp(-B(a_{\text{GTG}}, n_{\text{level}}, \Delta t) \cdot p_{\text{turb}}) - C(a_{\text{GTG}}, n_{\text{level}}, \Delta t) \quad (3)$$

Currently, this model is still in development.

Here we present various “2D slices” of this dependence for different values of a_{GTG} , n_{level} , and Δt , emphasizing the most important features of the turbulence impact on air traffic density.

6.3 Impact over One Hour

Here we study how air traffic density changes based on a single probabilistic GTG hourly forecast. Plots of the changes in air traffic density (relative to the “no turbulence” baseline day) are generated as a function of the relative probability of a severe turbulence encounter for a single flight level inside a NAS sector (Figure 10).

Each curve is labeled by a corresponding threshold for GTG diagnostic agreement, the number of flight levels in an airspace volume under consideration, and the time interval used to average the turbulence probability and flight level occupancy data. Each (x,y) point represents a median of all y-axis values whose x-values fell into $x \pm 2\%$ range (that is, 4%-wide bin). The

number of y-axis values ranged from ten to several thousand.

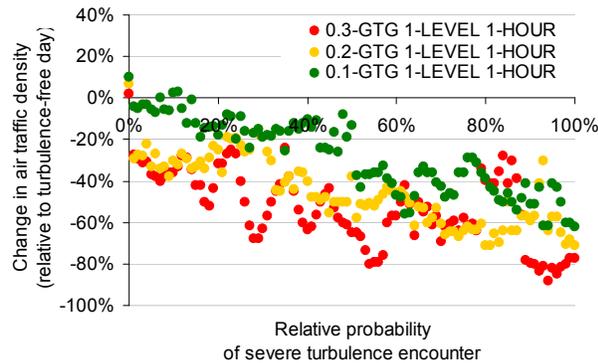


Figure 10: Air traffic density change in a flight level versus severe turbulence probability

Several observations can be made:

(a) The more definite the probabilistic severe turbulence forecast becomes (that is, the higher the GTG index agreement), the faster the air traffic decreases with the increase of the turbulence in a sector/flight level.

(b) With no turbulence forecasted in a given sector/flight level (the leftmost point of the above curves), its occupancy is typically higher compared to a turbulence-free baseline day. This is most likely due to air traffic being re-routed from the neighboring turbulent sectors. Furthermore, this increase in air traffic is getting higher as the GTG diagnostics' agreement decreases from "red-or-higher" to "yellow-or-higher" to "green-or-higher" (this is investigated further later).

(c) The data has an offset at $x=0$. This is to be expected, since the overwhelming majority of the data points taken at 0% turbulence probability for a given

GTG agreement threshold are from the turbulence-free sectors – that experience no decrease in air traffic – while even the smallest non-zero turbulence probabilities (1% and higher) indicate the presence of turbulence and thus immediately make a portion of air traffic to climb/descend away from this flight level.

For the 2-flight-level, 4-flight-level, and 8-flight-level volumes, we see the decrease in the air traffic volume does not go as low as it does for a single flight level. The larger the vertical extent of a 3D volume is, the slower the overall occupancy decreases with the turbulence probability increase. This is due to the air traffic escaping into flight levels above or below the volume under consideration. For multi-flight-level 3D volumes, flights resorting to these altitude-adjustment maneuvers are more likely to end up within the same 3D volume as before – thus slowing down the rate of the occupancy decrease with turbulence. For 3D volumes of the largest vertical size (e.g., 8+ even or odd flight levels), a reasonable assumption can be made that only a small percentage of air traffic escapes these sectors in the vertical direction (that is, descending below 30,000 ft or climbing above 45,000 ft) – so for higher turbulence probabilities horizontal re-routing is a main source of the air traffic population loss in these sectors.

6.4 Impact over Several Hours

For turbulence probabilities and sector/flight level occupancies averaged over 2, 3, and 4 hours, there are only minor differences among these cases as compared to 1-hour-averaged data – mostly in the low-turbulence-probability section of the plots. This initial difference can be explained by the assumption that for longer time periods used for averaging, there is a higher probability that there were periods of zero or low turbulence (included in the averaging) that were acceptable to air traffic. Overall, larger intervals used for averaging have the effect on the data similar to larger 3D volumes.

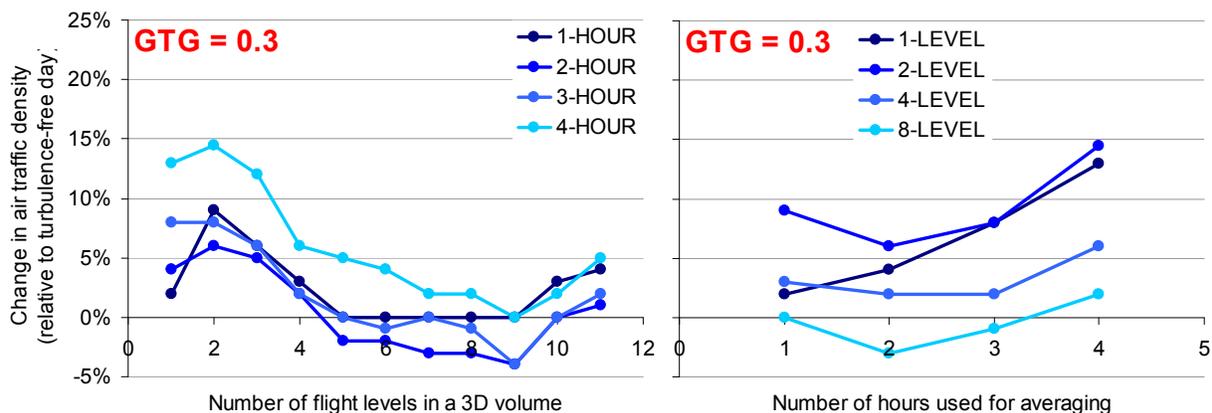


Figure 11: Air traffic density change in a 3D volume versus time and the number of flight levels when no turbulence is present.

6.5 No-Turbulence Situation

The zero-turbulence-probability part (leftmost point) of Figure 10 deserves attention, as it gives us the information about the amount of air traffic re-routed away from turbulence into neighboring sectors and the resulting increase in the sector controller workload.

Figure 11 (left) shows a non-trivial dependence of the amount of re-routed air traffic on the vertical size of a 3D volume of interest. Qualitatively, this dependence is expected, since the same number of aircraft entering/exiting a volume of airspace will result in progressively smaller relative change in the number of aircraft present in this volume as the size of the volume increases. Quantitatively, this result will be further investigated within our model. Figure 11 (right) suggests that the longer a 3D volume of airspace is able to maintain zero probability of turbulence (i.e., below a given GTG agreement threshold), the more trustworthy it becomes and thus a larger percentage of air traffic is rerouted there from the neighboring turbulent sectors. (This increase seems to happen in a non-linear fashion).

When we look at similar plots with decreasing GTG agreement with its indices (specifically, .2 and .1), there is a significant dependence on the threshold used for the GTG diagnostics' agreement: the higher the threshold is (i.e., the fewer GTG indices agree that the turbulence will be severe), the less truly turbulence-free the GTG map regions below the threshold become - thus resulting in less re-routed traffic. The leftmost point on the upper plots suggests that for the GTG agreement threshold slightly above 0.3 (red) the air traffic re-routed into the 3D volume from a more turbulent sector and the traffic re-routed from this volume will cancel each other.

6.6 Pilot/Airline Behavior Model for Turbulence

Here, we present how turbulence-avoidance maneuvers lead to air traffic density changes.

The flight track/flight plan amendment information was acquired for all flights that posted track hits inside the selected experiment areas and specified altitude bands in the 1500Z – 1900Z time frame on Oct. 28, 2006, and Jan. 24, 2007. For two turbulence events combined, this group included 8,200+ flights. A flight plan amendment represents a change in the flight route, altitude, or both. Combining the flight track with the flight plan amendment information allows us to see if/how the sector route and cruising altitude tactically chosen by a flight correlate with the relative probability of a severe turbulence encounter predicted by GTG for the sector/flight level an aircraft is about to enter.

For 8,200+ flights flown through 300+ NAS sectors spanning up to 22 flight levels from FL240 to FL450 over the course of 4 hours, we recorded about 26,500+ flight encounters with sectors/flight levels characterized by various probabilities of severe turbulence.

This statistic was analyzed to answer the following question: for a sector with a given relative probability of a severe turbulence encounter of a given threshold ("green-or-greater", "yellow-or-greater", or "red-or-greater" for probabilistic GTG maps), what is the

distribution of the air traffic responses among no response, altitude adjustment, sector route adjustment (re-route), altitude adjustment and re-route, and other? Table 2 summarizes this distribution.

Turbulence Situation Response	No Turbulence (baseline)	High probability of Severe Turbulence in next sector along a route [GTG threshold = 0.3 (red)]
Fly through	56%	35%
Adjust sector route	10%	6%
Adjust altitude	8%	18%
Adjust sector route and altitude	6%	22%
Adjust something else	20%	19%

Table 2: Distribution of turbulence-avoidance maneuvers for all air traffic combined.

Several observations can be made:

1. First, the baseline: on a turbulence-free day, 56% of all flights enter the next sector along their route without any adjustments; 10% will amend their sector route; 8% will adjust their altitude; 6% will make some combination of altitude and sector route adjustment; and 20% will resort to some other type of amendment (jet route, waypoints, etc.) while staying on the same sector route and at the same altitude.

2. On a turbulent day, the behavior of traffic depends on the relative probability of a severe turbulence encounter (for a set GTG threshold): predictably, the percentage of all air traffic that will fly through a turbulent sector without any response, decreases with the increase of the likelihood of a severe turbulence encounter in that sector, while the percentage of air traffic that will adjust their altitude (with or without the sector route) increases. The percentages of flights that will only amend the sector route or something other than sector route and altitude stay roughly the same – thereby suggesting that the altitude adjustment is indeed the favored response to the encountered or anticipated turbulence.

3. For high likelihoods of a severe turbulence encounter, the percentages of the flights that display various kinds of responses level off. For the "red-or-higher" GTG threshold, about 35% of flights will fly through without any adjustments, 6% will amend their sector route; 18% will adjust their altitude; 22% will make some combination of altitude and sector route adjustment; and 19% will resort to some other type of amendment (jet route, waypoints, etc.) while staying on the same sector route and at the same altitude.

Further study of aircraft response patterns indicate:

1. Descending is generally favored over climbing.

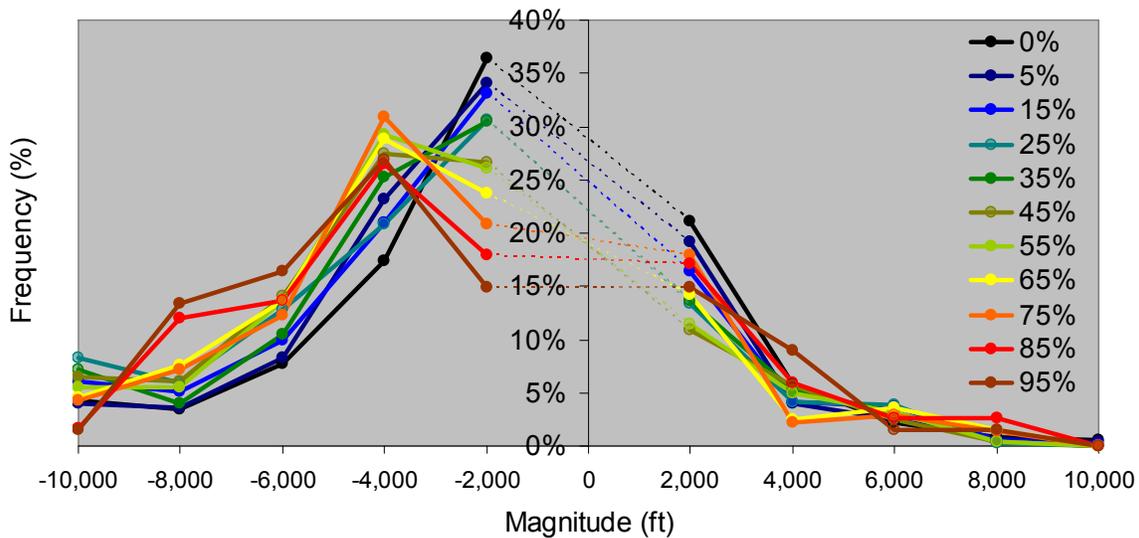


Figure 12: Magnitude distribution for altitude-changing turbulence-avoidance maneuver

2. Heavier aircraft in general and freight in particular seem to favor re-routing over altitude adjustment. Lighter aircraft (like General Aviation) strongly favor altitude adjustment over re-routing.

6.7 Altitude Distribution

Next, we analyze the magnitude distribution of climbing/descending maneuvers as a function of the probability of a severe turbulence encounter in an upcoming sector. Figure 12 data are labeled according to the probability of a severe turbulence encounter for the “red-or-higher” (≥ 0.3) threshold for GTG diagnostic agreement. As the probability of severe turbulence becomes higher, the distribution shifts towards descending and becomes wider. For the highest probabilities, the 4,000-ft descending maneuver dominates, while the likelihood of a 8,000-ft descending maneuver is almost the same as for a 2,000-ft climbing maneuver. These data demonstrate that when there is a substantial probability of severe turbulence in a given sector/flight level, a large portion of air traffic also avoids the next flight level below and above.

6.8 Impact on Flight Time

Finally, we discuss the impact of turbulence avoidance on time – both on the ground and in-flight. The former has to do with strategic delays and cancellations, the latter with strategic and tactical re-routing, and diversions. We found that re-routing, not unexpectedly, was the only significant time-related impact. While altitude-changes do not, in general, lead to any substantial en route delays, re-routing often does – which, in turn, has obvious implications for the overall cost of fuel burned during the flight.

7. In-flight Icing Impact Analysis

In-flight icing is a significant safety hazard to any aircraft exposed to atmospheric conditions conducive to airframe and powerplant (induction, carburetor, fan/propeller) icing. These conditions are generally present when the outside air temperature is below 8 degrees Celsius with visible moisture. Although low-altitude, low-performance aircraft are most susceptible to the affects of icing from a safety viewpoint, NAS traffic flow impact is also significant in the terminal arrival/departure phase of flight. Safety is the primary issue; avoidance because of safety considerations creates a NAS impact. This impact is severe enough to warrant further study as it relates to NAS performance.

7.1 Analysis Approach

Our analysis primarily focused on 12 SIGMETs issued during the winter seasons of 2005-06, 2006-07 and 2007-08 across the NAS. The analysis of SIGMETs is particularly helpful because they identify geographic and temporal bounds in which to focus the analysis. SIGMETs are issued for non-convective icing when the intensity has the potential to be severe and the threat exists within an area larger than 3,000 square miles. These SIGMETs are also initiated – and verified – by PIREPs. SIGMET information, once issued, is disseminated to pilots and airlines, and decisions pertaining to avoidance of the SIGMET may be made by the pilot(s) and airline(s). These avoidance decisions may include delay, cancellations, or a change in the planned route or altitude. While these 12 SIGMETs are only a subset of the 78 SIGMETs identified over these three seasons, further study has shown they are representative of the total set, and subsequently provide

an opportunity to conduct a detailed analysis on various NAS performance metrics.

The dynamic and complex nature of the NAS poses challenges to determining causes of delays, cancellations, etc., because there are so many possible reasons that flights can be affected – reasons which have nothing to do with icing. Since there are a variety of random and independent factors contributing to impacts, this analysis determined whether the severe icing SIGMETs had a statistically significant effect on the impact, beyond what can be explained by day-to-day random, independent changes in the NAS. Statistical methods were established to check for statistical significance.

7.2 SIGMET Effect on Delays

Delays are caused by a multitude of factors – as discussed previously – and can propagate through the NAS both spatially and temporally. This poses a challenge in the analysis of cause and effect relationships for delays and severe in-flight icing. The analysis, therefore, was focused on determining whether severe icing – in the form of SIGMETs – caused an increase in delays which was statistically significant. In order to conduct statistical hypothesis testing, there must first be an adequate amount of sampled data. For each SIGMET, flights departing to, or arriving from, airports affected by the SIGMET were identified, and statistics were gathered on the *percentage* of those flights arriving 15 or more minutes late to their destination. Delays were categorized as departure or arrival delays; SIGMETs were classified as ground-level or above-ground. These statistics are collected for the hours of SIGMET activity, for each day of the same month in which the SIGMET is active, and for each airport affected by the SIGMET. For example: For the SIGMET that occurred on Dec. 13, 2007 18Z-22Z, statistics on delayed flights departing to, or arriving from, *each* of the five airports affected on the SIGMET day (LGA, JFK, EWR, BWI, PHL) are gathered for Dec. 1, 2007 18Z-22Z, Dec. 2, 2007 18Z-22Z, and so on – until all days of the month are collected. Doing this results in 31 data points (for each of 31 days in December 2007), for each of the five affected airports, and thus a total of 155 data points. Five of these data points – represented by the flight delay percentages occurring on the day and hours the SIGMET was active for each of the five airports – are separated from those initial set of 155 data points. The end result is two empirical samples: one sample representing the percentage of flights delayed when the SIGMET was inactive for each of the five airports; the second sample representing that same metric, but for when the SIGMET was active. This same approach was then taken for each of the SIGMETs in the sample subset, and all of these data were combined, such that the result is two sampled empirical data sets which still

represent the percentages of flights delayed when the SIGMET is inactive and active, respectively, but for *all* SIGMETs within our subset.

Choosing which hypothesis test to use depends on the underlying distribution – either observed or assumed – of the sampled data. In this case, our samples are collectively taken from widely disparate spatial and temporal regions of the NAS and, thus, the assumption of normality cannot be made because each spatial and temporal region possesses a unique variability. Non-parametric hypothesis tests do not require the assumption of normality for the sampled data sets and subsequently makes it a desirable option for this analysis. A two-sample Kolmogorov-Smirnov (K-S) test (Lehmann, 2005) is a type of non-parametric hypothesis test which compares two empirical samples, each represented with its Empirical Cumulative Distribution Function (ECDF), making it suitable for this analysis.

Figure 13 shows example ECDFs of the two samples, arrival and departure delays for the cases where the SIGMET touches the ground. Red and blue curves represent the percentage of flights delayed when the SIGMET is active and inactive, respectively. Because SIGMETs are a relatively rare occurrence, the ECDFs representing flight delay percentages when SIGMETs are inactive (the blue curve) can be thought of as resembling the background population. Therefore, they involve a denser amount of data and appear as a smoother curve. Conversely, the relative rarity of active SIGMETs (red curve) results in much fewer data points, causing the curve to look coarse. In the more numerous cases where the SIGMET is above ground level, the relationship is still significant but not as strong. This analysis shows that the existence of a SIGMET does indeed significantly affect arrival and departure rates. It should also be noted that even though there are only three ground-level SIGMETs in our subset, they affected a total of nine airports, which is why there so many “steps” in the ground-level SIGMET plots than actual number of ground-level SIGMETs – delays are analyzed independently for each affected airport. The test statistic, D , used by the K-S test represents the maximum distance between the ECDFs of the two samples. Because the value of D represents the maximum separation between the two ECDFs, larger values of D tend to be correlated with higher levels of statistical significance. More importantly, the p -value (P) represents the probability that the difference between the two samples was caused by random variation. The statistical significance threshold used in this analyses is 0.05, meaning that a p -value lower than 0.05 is deemed a statistically significant result. The very low p -value resulting from the K-S test shows a strong, statistically significant relationship between SIGMETs and an increase in flight delays.

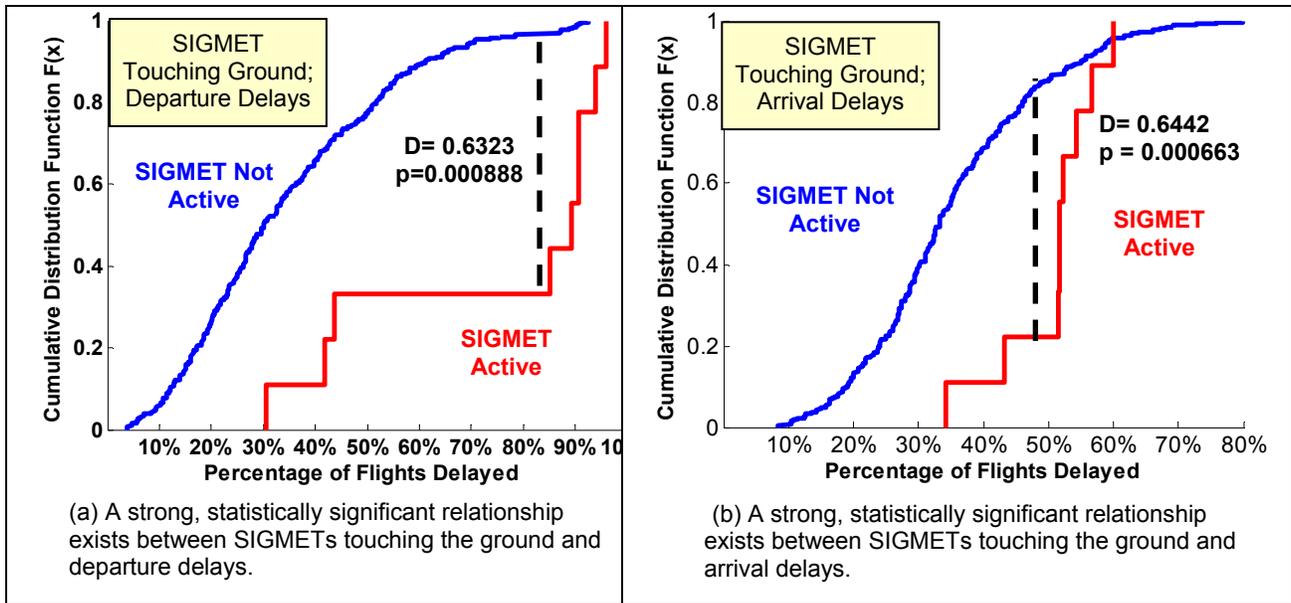
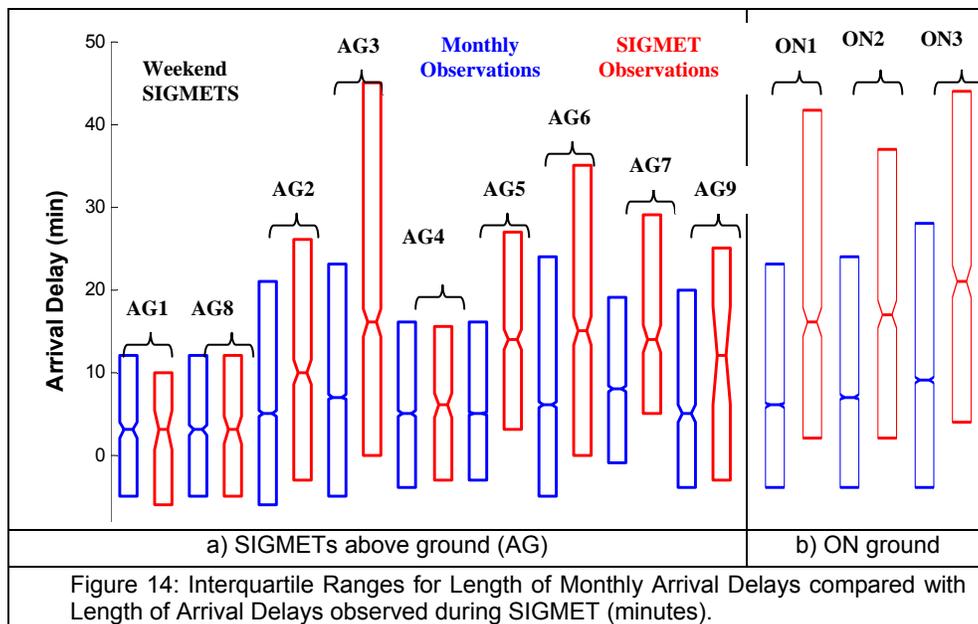


Figure 13: ECDFs comparing percentages of flights delayed at airports affected by SIGMET.



While the frequency of flight delays was just discussed, the magnitude, represented as the amount of time flights were delayed, was also analyzed. Arrival delay times for flights flying through the areal region of airspace covered by the SIGMET were analyzed and compared with arrival delay times for flights flying through the same spatial-temporal airspace over the entire month. This airspace-centric analysis of delay includes flights that flew through the SIGMET region but did not depart or arrive at any airports directly affected by the SIGMET. The magnitude of arrival time delays observed when the SIGMET was active, relative to observations in the same spatial-temporal region across

the entire month, is shown in Figure 14. Here, each SIGMET analyzed has a pair of boxes: the blue box represents interquartile (25-75 percentiles) ranges of flight delays observed for similar hours over the same month the SIGMET was issued; the red box represents interquartile ranges of flight delays observed only during the time and day the SIGMET was active. The notch represents the median. Figure 14 shows that, in 9 of the 12 cases, the amount of time flights were delayed when the SIGMET was active was significantly higher than what is typically observed within the same airspace during the same month.

7.3 SIGMET Effects on Cancellations

A statistical approach, similar to the one used to analyze flight delays, was also used to analyze flight cancellations. Cancellations, like delays, can be caused by many factors such as traffic management initiatives and poor weather; however, they occur much less frequently as compared to delays. This analysis first identifies major airports affected by the SIGMETs, either by being located underneath or in the vicinity of the SIGMET boundary. For each airport, the number of flight cancellations and scheduled arrivals/departures was identified, for the same hours of the day the SIGMET was active and for each day of the month the SIGMET was active. The number of departure and arrival cancellations for each day is divided by the number of scheduled departures and arrivals, respectively; the resulting metric is the percentage of scheduled departures and arrivals that were cancelled for each day of the month. Using the same methodology used for delays, data collected for the day of the SIGMET is separated and placed into a different sample set for comparison purposes during statistical hypothesis testing. For example, for the SIGMET that occurred on Dec. 13, 2007 18Z-22Z, five airports were identified as being affected: EWR, LGA, JFK, PHL, BWI. For each airport, data representing the number of departure cancellations, arrival cancellations, scheduled departures and scheduled arrivals occurring during 18Z-22Z was collected for each day in December. In all cases, a significant relationship exists between SIGMETs and departure flight cancellations, with the effects of SIGMETs touching ground being much more significant, as one would expect.

7.4 SIGMET Effects on AAR and ADR on Affected Airports

Delays and cancellations, collectively, can significantly reduce the AAR and ADR. The hourly ADR and AAR observations from major airports affected by the SIGMET, either by being underneath it or on the edge of the SIGMET boundary, were averaged for the day and hours the SIGMET was active. These observations were compared with AAR and ADR observed at similar hours and days across the entire month. For example, the averaged hourly ADR at Newark airport (EWR) when the SIGMET was active (December 13, 2007 18Z-22Z) was 8 flights per hour; this was then compared to ADR, taken as a monthly average, observed during similar hours (18Z-22Z) across all weekdays of the month. In this case, the monthly averaged hourly ADR for December, for the hours of 18Z-22Z, was 28 flights per hour at EWR. Therefore, the average of 8 flights per hour observed during the SIGMET was approximately 71.4% below the monthly average. The SIGMETs were separated into two categories: SIGMETs touching ground (Table 3), and SIGMETs above the ground (Table 4). Tables 3 and 4 summarize SIGMET effects on AAR and ADR for both categories.

When SIGMETs touch the ground, the AAR and ADR drastically go down. Nine instances involving eight major airports (with CVG involved in two of the instances) were affected by SIGMETs touching the ground. Eight out of nine airports observed a decrease in AAR and ADR. The one case that had a positive increase in AAR and ADR was a very small increase, close to zero.

SIGMET date	Airports affected	Change in ADR	Change in AAR
Dec 13 2007 18Z - 22Z	BWI	-11.6%	-0.7%
	EWR	-71.4%	-40.4%
	JFK	-26.4%	-22.2%
	LGA	-50.1%	-34.5%
	PHL	-21.5%	-21.3%
Feb 21 2008 21Z - 01Z	CVG	-7.9%	-28.0%
	IND	1.6%	0.1%
	STL	-51.7%	-83.2%
Feb 22 2008 01Z - 05Z	CVG	-42.8%	-0.3%

Table 3: SIGMETs touching ground: AAR and ADR monthly averages during the SIGMET.

SIGMET Date	SIGMET Altitude Range (ft)	Airports affected	Change in ADR	Change in AAR
Oct 20 2007 14Z-18Z	16000 - 23000	PDX	14.3%	0.5%
Dec 14 2005 19Z-23Z	9000 - 12000	MDW	6.6%	-0.3%
		ORD	-27.7%	-35.8%
Feb 7 2008 07Z-11Z	8000 - 14000	SEA	-32.3%	-6.5%
Dec 1-2, 2005 22Z-02Z	2000 - 5000	CVG	16.6%	0.5%
Jan 2 2007 00Z-04Z	8000 - 14000	PIT	-13.5%	-14.3%
Feb 28 2006 17Z-21Z	10000 - 20000	SLC	-11.7%	-16.0%
Nov 14 2006 01Z-05Z	10000 - 20000	SFO	-5.9%	-19.2%
Nov 14 2006 05Z-09Z	10000 - 20000	SFO	13.0%	22.2%
Feb 26 2007 18Z-22Z	6000 - 15000	SFO	-18.5%	-16.0%

Table 4: SIGMETs above ground: AAR and ADR monthly averages when SIGMET was active.

When SIGMETs are above the ground, the effects are mixed. Ten instances involving eight major airports (with SFO involved in two SIGMETs, one of which, on November 14, 2006, was extended an additional four hours) were affected by SIGMETs issued above the ground. Only four of those ten instances observed a decrease in ADR, while six of those ten observed a decrease in AAR. However, the relationship between SIGMETs above the ground and changes in ADR/AAR of affected airports is dependent on the operational procedures of the airports and the approach taken by air traffic controllers to route traffic under the SIGMET boundary if possible. There are some instances where a significant decrease in ADR/AAR was observed, but others instances where a significant increase was observed. These differences can be attributed to cases where the SIGMET was still in effect but either not verified or confirmed.

7.5 SIGMET Effects on Diversions

Flight diversions are defined as flights landing at an airport that is different than the arrival airport filed in the flight plan. Diversions, while rare, have a high impact for airlines because they may involve placing passengers on another plane, staffing issues, wasted fuel, and other unplanned use of resources. Analysis of planned vs. actual flight plan data suggest that diversions are a rare event, even when a SIGMET is active; however, the issuance of a SIGMET does show some increase in diversions, particularly when the SIGMETs touch the ground. (Figure 15, SIGMET above the ground).

7.6 SIGMET Effects on Airborne Circular Holding

Flights arriving to an airport affected by a SIGMET that is above ground may have to engage in airborne circular holding in order to navigate underneath the SIGMET during descent. Airborne holding is undesirable because of inefficient use of time and fuel. Figure 16 shows an example for SFO during a SIGMET case in November 2006. Here we show the percentage of flights engaging in airborne holding prior to arrival. An hourly average is given for the same month the SIGMET was issued—it is the percentage of flights which engaged in airborne

holding prior to arrival. This is compared to the day of the SIGMET. The hours of SIGMET activity are highlighted on the chart. As might be expected, not all SIGMETs resulted in airborne holding due to specific airport procedures or lack of confirmation of actual icing conditions. When holding does occur, however, the delays are large and generally widespread during the SIGMET period.

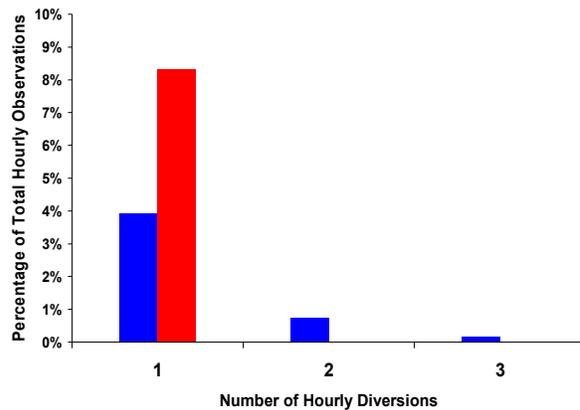


Figure 15: Diversions for SIGMETs above ground.

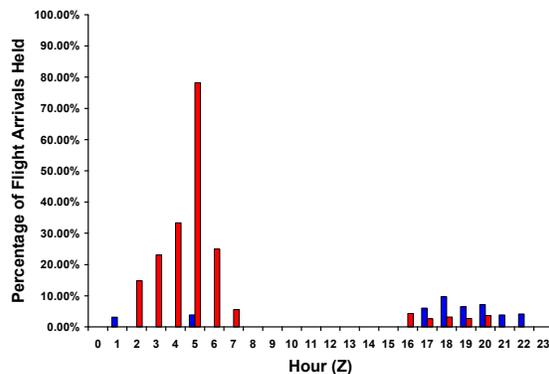


Figure 16: Percentage of arrival flights involved in airborne circular holding (by hour).

8. Cost Analysis

Our cost analysis includes costs associated with:

- Altitude Maneuvers
- Reroute Maneuvers
- Ground Delays
- Airborne Delays
- Cancellations
- Diversions

The primary cost of the altitude and reroute maneuvers is the cost of extra fuel consumed to perform these maneuvers. In the case of large rerouting maneuvers caused by either large turbulence areas, or large regions of severe icing, time, i.e. delay costs are also analyzed. We calculate fuel burn cost for altitude and reroute maneuvers using the exact Base of Aircraft Data (BADA) equations with BADA coefficients from BADA 3.6 files, comparing results with the fuel burn calculated for the baseline steady level flight without any maneuvers. We then analyze cost of ground and airborne delays, as well as the cost of cancellations and diversions.

The BADA model is the industry standard for fuel burn calculations. The BADA model is based on a Total Energy Model (TEM) (Rutowski, 1954). A point model is used for the aircraft, with forces acting on that point - thrust T , drag D , lift L , and gravitational force W , in combination with a total energy conservation law (Nuic, 2005, Eurocontrol 2004).

The Breguet range and endurance equation could also be used to calculate fuel burn. Both equations are derivable from the fundamental equations of motion governed by Newton's Laws, but are not derivable from each other. The BADA equations have an advantage over the Breguet equations because they take in account a larger number of significant parameters. This is why the BADA model, with its ongoing improvements, is the industry standard for fuel burn calculations (Trani 2009, Ssamula 2006). As part of our in-depth development of an impact cost model, these equations were adapted to our context and computations used to validate other more heuristic and literature-based cost models as described here. Full development of the

BADA equations is contained in the final NASA report, which can be accessed through the Corresponding Author.

8.1 *Toward Simplifying the Cost Model, or Reducing Complexity*

Generally, one can categorize airline costs in different ways. Cook et al (Cook 2004) divide costs in two categories: tactical and strategic costs. Strategic costs are incorporated into the Airline Operating schedule in advance, and are well described as unit costs, like those published regularly by Air Transport Association (ATA 2008). Unit costs are given in fixed amount of dollars per block minute, regardless of the duration of delay. Tactical costs are more appropriately described as marginal costs, which can vary from minute to minute, and can be triggered at different times during the delay duration, as illustrated in Figure 17. Essentially, marginal costs are those which give rise to the non-linearity of the delay cost function as a function of time. Delays caused by severe icing belong to daily tactical delays, and when they happen, constitute tactical delays, which, then, have significant effect on other aircraft causing their reactionary delays due to delay propagation. We illustrate the non-linearity of delay propagation in Figure 18.

For primary tactical delays the following are the most important cost components:

- Fuel burn costs (every delay results in an additional fuel burn cost, except if on the ground with engines off)
- Maintenance costs
- Flight and cabin crew salaries and expenses
- Depreciation, rentals and leases of flight equipment
- Handling agent penalties
- Airport charges
- Costs of passenger delay to airlines (hard and soft)
- Cancellations
- Diversions

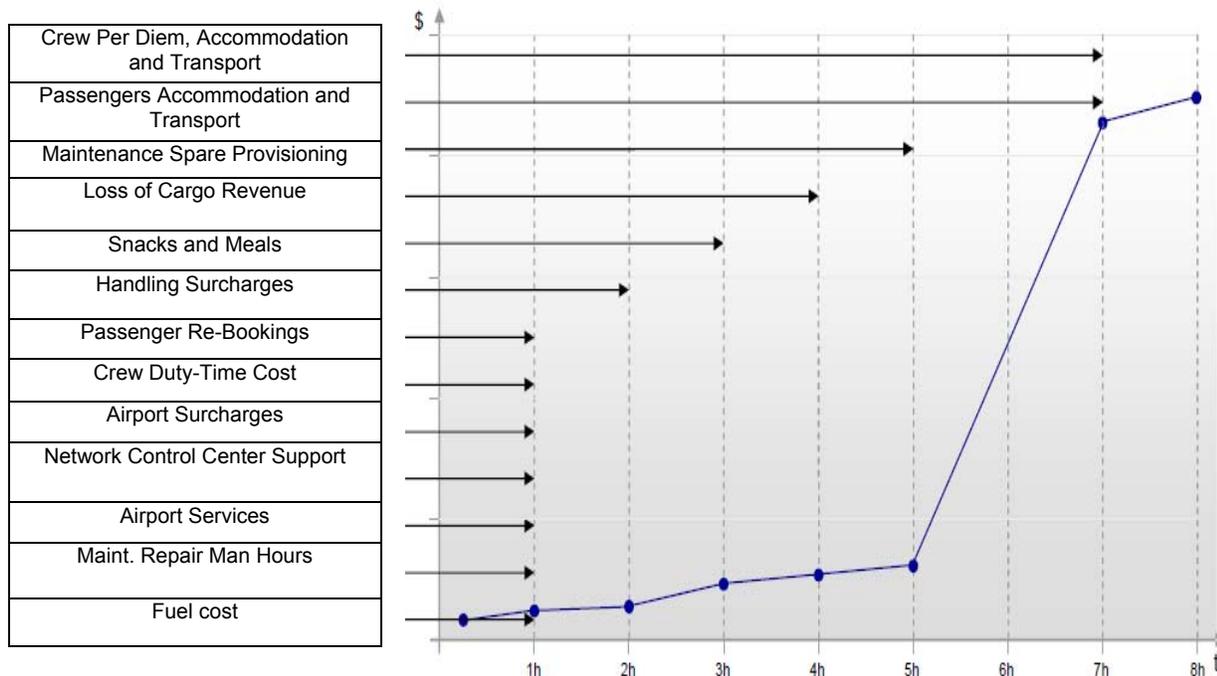


Figure 17: Delay costs as they are triggered.

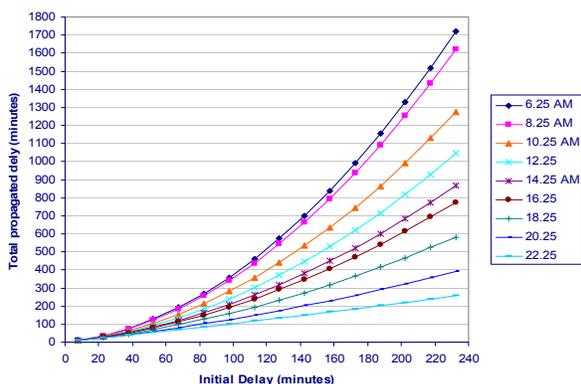


Figure 18: Total propagated delay as a function of the initial delay (Beatty 1998).

Some of these cost components are in effect only for airborne situations, like fuel burn, and direct operating block-minute costs. Some are present in both ground and airborne delay situations. DOCs can be further categorized into sub-categories, which can have different impact on the total DOCs.

Major cost components are:

- Fuel burn costs,
- Passengers costs (hard and soft),
- Crew (flight and cabin) costs,
- Maintenance costs.

However, even if we restrict ourselves to just these four major cost components, the problem is still quite complex (Cook 2008).

In addition, flight cancellation and diversion costs are a function of many variables. For example, costs

associated with flight diversions can be between \$10,000 and \$100,000. The presence of severe in flight icing, for instance, can cause ground delays where an aircraft departure is delayed, or airborne delays where the aircraft is rerouted to avoid a region of icing. Prolonged ground delays may result in flight cancellations, and severe icing conditions can cause a flight diversion from the destination airport to a diversion airport or back to the origin airport. Our analysis included all of these types of flight disruptions.

Ground delays due to terminal area weather constraints are not typically isolated events that concern just one individual aircraft, but they propagate to many upstream flights, and affect the airline operating schedule as well. Thus, we analyzed the costs incurred by propagation delays.

Several parameters contribute to the overall increased operating cost when delays happen. The delay cost varies with airline and by aircraft type. Furthermore, the operating costs consist of DOCs and Indirect Operating Costs (IOCs), each having a set of elements. The distinction between the two is not always clear (Cook 2004). Even the categorization of DOCs and IOCs is not unique. For instance, the ATA provides the update on cost of delays and gives the list of DOCs as dollar amount per block minute in Table 5 (ATA 2009).

Passenger related costs, which are IOCs, are usually treated separately from DOCs. These costs can vary from airport to airport, since they include passenger accommodations and ground transportation, which vary in price by location.

DOC 12 Months Ending Sept 2008	\$ Per Block min	Δ% vs CY 2007	Annual Delay Costs (\$ millions)
Fuel	\$39.35	41%	\$5,431
Crew - Pilots/Flight Attendants	\$13.08	3%	\$1,805
Maintenance	\$10.09	5%	\$1,398
Aircraft Ownership	\$7.72	0%	\$1,066
Other	\$1.88	-28%	\$259
Total DOCs	\$72.13	19%	\$9,954

Table 5: ATA direct operating costs of taxi plus airborne delays.

The complexity of the general problem of modeling the delay cost function with such a large number of independent parameters, results in a lack of consensus about how to approach this problem. There is a variety of approaches in the literature, and each is based on restricting assumptions. No general delay cost model is widely accepted. Furthermore, due to confidentiality, airlines' crew and maintenance costs are generally not available.

Our goal is to give an estimate of the cost of each minute of ground and airborne delay, as well as the cost of the flight cancellations and diversions. Of course, we expect that the cost will vary with time, i.e., that it will not be a constant for each minute.

8.2.1 Cost of Ground Delays and Cancellations

Generally, the cost function of the ground delays has to have some global characteristics as a function of time, regardless of the complexity of its detailed behavior. We make the following assumptions, illustrated in Figure 19.

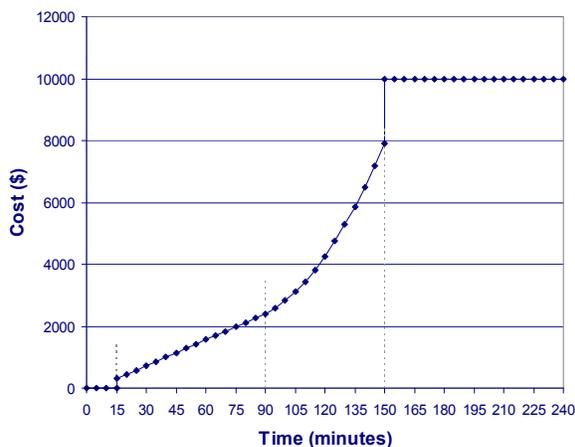


Figure 19: A prototypical nonlinear ground delay cost curve.

- 1) Most researchers agree that the initial DOC for short ground delays, up to 15 minutes, is zero. This parameter of zero cost at initial time, t_0 , can be varied, depending on the available data for the particular airlines and/or aircraft types.

- 2) We assume that after the initial 15 minutes there is a discontinuity in cost.
- 3) If ground delay is prolonged, it will eventually end in the flight cancellation, since airlines are not likely to endure delay costs which surpass the cancellation cost. Thus, the cost function is naturally capped off at the cost of cancellation, at some delay t_c . We discuss the value of t_c later.
- 4) For delays between 15 and t_c minutes, the cost curve is very complex and irregular, since various cost components enter at various times as the delay increases. At present, there is no good model to describe this curve in detail. However, global behavior of the cost curve can be described first as linear, between t_0 and t_1 , and superlinear, between t_1 and t_c . (we approximate the superlinear part with another linear segment of different slope).
- 5) We assume that the increasing cost of a prolonged delay does not need to reach the cost of the cancellation for the flight to be cancelled. That means that our cost curve will have the first derivative discontinuity at time t_c . For times greater than t_c the cost function will cap off at the cost of the cancellation.
- 6) Finally, we assume that the ground delay cost per minute is roughly $\frac{1}{2}$ of the ATA value of airborne delay cost per minute, reduced by the amount of ownership costs, which belong to strategic costs and are out of scope of this model.

As illustrated in Figure 19, our assumptions pose certain limitations. Free parameters are t_0 – the time when the cost becomes non-zero ($t_0=15$ min, in our case); t_1 – the time which separates linear from non-linear part of the cost curve ($t_1=90$ min, in our case); t_c – the cut off time when delay ends with flight cancellation ($t_c=150$ min, in our case); the slope of the linear segment of the curve, which we define as $\frac{1}{2}$ of the ATA airborne cost per block minute value reduced by the ownership costs, and which can also vary from flight to flight; and finally, the cost of cancellation which caps off the delay cost curve.

8.2.2 Cost of Airborne Delays and Diversions

For the airborne delays costs we use the ATA data for the delay cost per block minute. Assume that there will be no zero cost time interval for airborne delays, due

to the fuel burn cost component when the aircraft is airborne. Otherwise, assume that the general shape of the cost function is similar to the cost function for ground delays (prototypically, Figure 19), except that the cap off cost value will be now equal to the cost of the aircraft diversion (Hanowski 2008). Further, we assume again the first derivative discontinuity at t_d , for similar reasons stated previously for the ground cancellation case. For example, an aircraft delayed due to an icing SIGMET will divert before the cost of the airborne delay reaches the actual cost of the diversion.

We assume a cost \$28,000 per diversion as in (Shavell 2001) for an A-319-type aircraft. However, this can vary in the literature from \$15,000 for smaller domestic aircraft to as high as \$893,000 for a jumbo aircraft on an international flight (Holmes 2009, Young 2009). This requires further investigation over wide range of aircraft types—a global cost for diversion is not available. We illustrate the general airborne delay cost models for the A-319 and B-777 aircraft types in Figures 20 and 21.

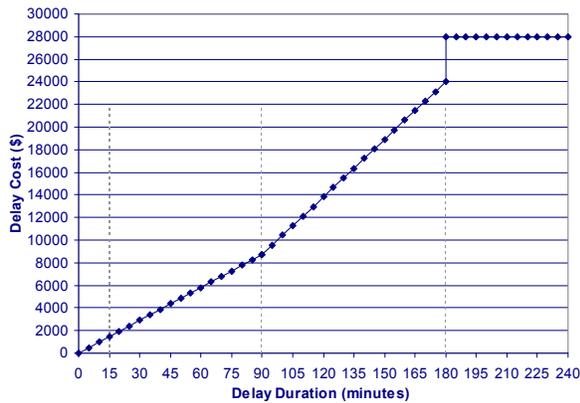


Figure 20: Delay cost function for airborne delays – 50 seats, for A319-100.

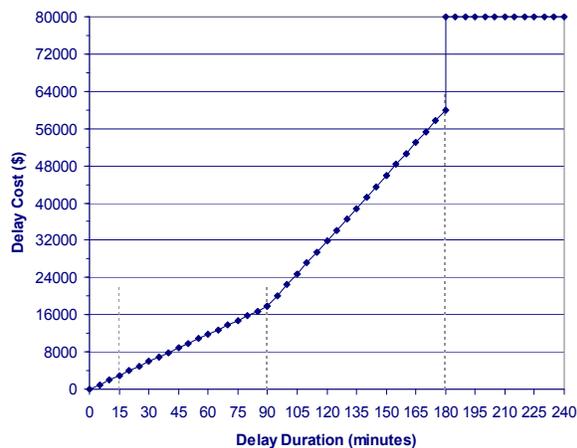


Figure 21: Delay cost function for airborne delays, for 218 seats for 777-200 Aircraft (\$80,000 diversion cost).

The cost models suffer, of course, from many simplifications. However, those simplifications are dictated by extremely large scattering of data over

different airlines, aircraft types, and flights, and, perhaps, more importantly, by the lack of data. Several more realistic cost models are presented in the NASA report that account for some of the scattering of data caused by reality. However, the general formulations presented here are sufficient for our objectives, where we want to compare cost impacts due to differing levels of weather integration and spatial/temporal resolution.

9. Conclusion

This paper discusses the translation of weather information into Traffic Flow Management (TFM) impact for Clear Air Turbulence (CAT) and in flight icing weather constraints. Data for CAT and in flight icing potential were studied at various flight levels for en route airspace in the National Airspace System (NAS). By analyzing aircraft flying in the NAS, we see that some aircraft fly through weather constraints where others do not, and the magnitude of the deviations varies on parameters investigated in this paper. Multi-dimensional mathematical models developed in this paper demonstrate the TFM impacts on a wide range of scenarios and conditions in the NAS, and when generalized, may be useful for studying the TFM constraints for the Next Generation Air Transportation System (NextGen).

ACKNOWLEDGEMENTS

This research was funded by NASA Ames Research Center under NRA contract NNA07BB36C for the NextGen Air Traffic Management (ATM) – Airspace Project – Subtopic 15: Translation of Weather Information to Traffic Flow Management Impact. The authors appreciate the frequent inputs from our contract monitor, William Chan, from NASA Ames Research Center. Finally, we appreciate the financial support of the sponsor of the research, NextGen Project Manager, Dr. Paramal Kopardekar. Prof. J. Mitchell is partially supported by the National Science Foundation (CCF-0528209, CCF-0729019)..

REFERENCES

- Air Transport Association (ATA), 2009: Costs of Delay, Available at: <http://www.airlines.org/economics/cost+of+delays/>, June, 2009
- Beatty, R. R. Hsu, L. Berry, J. Rome, 1998: Preliminary Evaluation of Flight Delay Propagation through an Airline Schedule", *2nd USA/EUROPE Air Traffic Management Research and Development Seminar*, Orlando, FL, Dec., 1998.
- Cook, A.J. G. Tanner, S. Anderson, 2004: Evaluating the true cost to airlines of one minute of airborne or ground delay, *Technical Report, Transport Studies Group, University of Westminster, May, 2004.*
- Cook, A. G. Tanner, 2008: Innovative cooperative actions of R&D in EUROCONTROL Programme CARE INO III: dynamic cost indexing: Airline costs of delayed passengers and how to estimate full network

- delay costs, *Technical Report, Transport Studies Group, Univ. of Westminster, Nov., 2008.*
- DeLaura, R. J. Evans, 2006: An Exploratory Study of Modeling Enroute Pilot Convective Storm Flight Deviation Behavior, *Proc. of the 12th Conf. on Aviation, Range, and Aerospace Meteorology, Atlanta, GA, 2006.*
- Eurocontrol Experimental Center, 2004: User Manual for the Base of Aircraft Data (BADA), Revision 3.6, Brussels, Sept. 2004.
- FAA, 2008: Federal Aviation Regulations / Aeronautical Information Manual, ASA Pub., Newcastle, WA, 2008.
- Hanowsky, M., 2008: A Model to Design a Stochastic and Dynamic Ground Delay Program subject to Non-linear Cost Functions, Ph.D. Thesis, Massachusetts Institute of Technology, Engineering Systems Division, Cambridge, MA, Dec., 2008.
- Holmes, M., 2009: Just How Much Does That Cost, Anyway? An Analysis of the Financial Costs and Benefits of the 'No-Fly' List", *Homeland Security Affairs*, Vol. V, No. 1, Jan., 2009.
- Joint Planning and Development Office (JPDO), 2007: Concept of Operations for the Next Generation Air Transportation System, Version 2.0, Washington, DC, June 13, 2007.
- Krozel, J. J.S.B. Mitchell, V. Klimenko, T. Lindholm, J. Prete, N. Downs, and S. Krishna, 2008: Translation of Weather Information to Traffic Flow Management (TFM) Impact: Weather Translation Model for TFM Decisions including FACET Demonstrations, Tech. Report 34N0707-002-R0, Metron Aviation, Inc., Oct., 2008.
- Lindholm, T. J. Krozel, J. S. B. Mitchell, 2009: Concept of Operations for Addressing Types of En Route Hazardous Weather Constraints in NextGen. *AIAA Aviation Technology, Integration, and Operations Conf.*, Hilton Head SC, 2009.
- Lehmann, E. J. Romano, 2005: Testing Statistical Hypotheses, Edition 3, Springer Pub., New York, NY, 2005.
- Nuic, A. C. Poinso, M. Iagaru, E. Gallo, F. Navarro, C. Querejeta, 2005: Advanced Aircraft Performance Modeling for ATM: Enhancements to the BADA Model, *24th Digital Avionics System Conf.*, Washington D.C., Oct. 30 – Nov. 2, 2005.
- Rhoda, D. M. Pawlak, 1999: The Thunderstorm Penetration / Deviation Decision in the Terminal Area, *American Meteorological Society, 8th Conf. on Aviation, Range, and Aerospace Meteorology*, Dallas, TX, pp. 308-312, Jan., 1999.
- Rutowski, E. S., 1954: Energy Approach to the General Aircraft Performance Problem, *Journal of the Aeronautical Sciences*, Vol. 21, No. 3, March, 1954.
- Sharman, R. S. Tebaldi, S., G. Wiener, and J. Wolf, 2006: An Integrated Approach to Mid- and Upper-Level Turbulence Forecasting, *American Meteorological Society, Weather and Forecasting*, Vol. 21, No. 3, 2006.
- Shavell, Z., 2000: Effects of Schedule Disruptions on the Economics of Airline Operations" *USA/Europe 3rd Air Traffic Management Seminar*, Napoli, Italy, April, 2000.
- Ssamula, B. Del Mistro, A. T. Visser, 2006: Using an operating cost model to analyse the selection of aircraft type on short-haul routes, *Journal of the South African Institution of Civil Engineering*, Vol. 48 No. 2, p. 2–9, Paper 579, Aug., 2006.
- Trani, A., 2009: Virginia Polytechnic Institute and State University, Blacksburg, VA, Personal Communication, March 27, 2009.
- Young, S., 2009: Ohio State University, Columbus, OH, Personal Communication, June 2009.