BUILDING A STOCHASTIC TERMINAL AIRSPACE CAPACITY FORECAST FROM CONVECTIVE WEATHER FORECASTS

Diana Michalek * and Hamsa Balakrishnan
Massachusetts Institute of Technology, Cambridge, Massachusetts

1. INTRODUCTION

The global commercial aviation industry generated $485 billion dollars in 2007, after a steady increase in revenue over the last decade. In the United States alone, this activity included 660 million passenger enplanements and the transport of over 11 million tons of freight (IATA, 2008; Bureau of Transportation Statistics, 2008a). This increase in demand for air travel has coincided with increased delays in the National Airspace System (NAS). According to Joint Economic Commitee Majority Staff (2008) estimates, the annual cost of domestic air travel delays to the U.S. economy was $41 billion in 2007. This figure includes $19 billion in costs to airlines, and $12 billion in costs to passengers. What's more, 76.9% of NAS delays in 2007 were weather-related, an increase over previous years (Bureau of Transportation Statistics, 2008b). This situation will only get worse, as demand for airspace is projected to climb further in the coming years, and it will become increasingly important to operate the airspace system efficiently even in the face of storms (IATA, 2007).

Researchers have been working on the challenge of minimizing delays for several decades, mostly in the field of Air Traffic Flow Management (ATFM). ATFM is the process of making strategic decisions a few hours ahead of the time of operations, to balance the demand for aircraft operations with the capacity of the NAS. The capacity of airspace is affected by the presence of hazardous weather, since aircraft must avoid unstable regions and are often forced to deviate from their planned trajectories. Historically, ATFM models have taken expected capacity to be a known input, and have only focused on optimizing routing decisions based on this input. Under clear weather conditions, these expected capacities are fairly stable and tend to reflect reality. However, under stormy weather conditions, capacity is highly variable, and the use of estimated capacity for planning is unrealistic. More recently, researchers have developed stochastic models for TFM, which take as input a probabilistic weather forecast or a probability distribution for capacity.

The missing piece has been a realistic model of airspace capacity under hazardous weather conditions. Such a model is critical for decreasing weather-related delays, yet it has been a challenging problem due to the chaotic nature of weather and to the unique requirements of ATFM forecasts. More concretely, forecasts for ATFM must be reasonably precise and fine-grained. Knowing that there is a 30% chance of rain in Boston today, for instance, does not help to determine if there will be a route open from the east into Boston Logan Airport at 5PM, or if flights should incur delay on the ground and avoid entrance into Boston between 5 and 9PM.

In this paper we present the first steps towards creating a stochastic forecast of terminal capacity, by using currently available weather forecasts. We focus on terminal airspace, or airspace surrounding an airport, since it is the bottleneck of aircraft operations. We consider a scenario in which short-term routing decisions must be made based on available forecast data. We find that the quality of a forecast can be measured in terms of how likely it is for a trajectory through forecast weather to be blocked by true weather conditions. This leads to a model for route robustness, which correlates features of a particular trajectory with blockage. We present the steps taken in creating this model, as well results using weather forecast data for Hartfield-Jackson Atlanta International Airport (ATL) terminal area from the summer of 2007. In addition, we show how route robustness can be mapped into a probabilistic forecast of terminal airspace capacity. Although this work focuses on the terminal area, the methodology developed can be modified to model en-route airspace.

The structure of this paper is as follows. Section 2 gives an overview of relevant TFM models and existing weather forecast research and opera-

* *Corresponding author address:* Diana Michalek, 77 Massachusetts Avenue, Bldg. E40-149, Cambridge, MA 02139. E-mail: dianam@mit.edu

tional products. Section 3 introduces the Lincoln Lab Convective Weather Forecast, which is used in this work. Section 4 describes the issues with trying to model capacity by studying errors in weather forecasts at individual pixels. Section 5 introduces a trajectory-based approach to modeling capacity, and contains the main contributions of this paper. Finally, Section 6 contains conclusions and future work.

2. RELATED WORK

In this section we present Traffic Flow Management models whose objectives include minimizing delays. These models take as input airspace capacity or weather conditions. In addition, we briefly describe currently available forecasts of airspace capacity and weather. This discussion will highlight the gap between TFM input requirements and the state of weather forecasting.

2.1 ATFM Models

Early work in TFM involved large-scale integer programming models which determined how to route a set of aircraft from their planned departure locations to their planned destinations while minimizing the cost of delays. Capacity was the major system constraint and limited how many aircraft could simultaneously occupy a region of airspace. Models involving both a static set of capacities for each sector of airspace, as well as multiple capacity scenarios, each associated with a probability of occurrence, were developed (Bertsimas and Patterson, 2000; Bertsimas and Odoni, 1997).

More recently there has been work on algorithms to efficiently synthesize routes through regions of airspace affected by convective weather. This work requires fine-grained and time-varying weather forecast data as static weather input, and focuses on synthesizing short and flyable routes which do not get too close to regions of airspace impacted by weather (Prete and Mitchell, 2004; Krozel et al., 2006). However, the weather forecasts are treated as ground truth, and routes are not evaluated against actual weather scenarios.

Finally, the Route Availability Planning Tool (RAPT) product uses Lincoln Lab Convective Weather Forecasts to model jet route blockage deterministically. The product is used operationally to help controllers determine if aircraft can take off (DeLaura and Allan, 2003).

2.2 Capacity Estimation

Some recent work has begun to look at the problem of creating stochastic models of capacity. Krozel et al. (2007) considers the problem of estimating the capacity of a sector of en-route airspace by computing a theoretical capacity given weather in the region. This is done through the application of continuous maximum flow theory. However, this work relies on static weather forecasts and does not incorporate uncertainty intervals or any measure of forecast accuracy. This research is taken a step further by Mitchell et al. (2006), which considers weather forecasts accompanied by a region of uncertainty. However, the uncertainty profiles are randomly generated.

We expand on this line of research by using maximum flow theory to help identify bottlenecks in regions of airspace, and to help synthesize routes which are then validated against observed weather.

2.3 Convective Weather Forecasts

There are several aviation weather forecasts available for the United States. In general, weather forecasts take the form of a grid, where each grid cell, or pixel, corresponds to a 2-dimensional section of airspace. Each pixel is associated with a value indicating the severity level of weather at that point.

MIT Lincoln Laboratory's Convective Weather Forecast product is a state-of-the-art 0-2 hour forecast, used throughout the United States to aid air traffic control (Wolfson et al., 2004). The forecast is static, meaning that each pixel contains one deterministic value indicating weather severity, with no additional estimate of the likelihood that the forecast is correct, or a distribution over possible severity. Specific details about the forecast are provided in Section 3.

Due to randomness in the weather and the resulting inaccuracy of weather forecasts, creating a plan for routes is not realistic using static forecasts alone. Indeed, flying through a region of airspace that turns out to be stormy would compromise safety. This has lead to research into developing probabilistic weather forecasts for aviation. NCWF-2 is one such forecast developed by the National Center for Atmospheric Research, which at each pixel gives a probability p that the pixel will contain convective weather. Initial validation of the forecasts show that these values of p have significant errors associated with them, and tend to be large overestimates of true values (Seseske et al., 2006).

Another probabilistic weather forecast that has been proposed is based on polygons (Sheth et al.,

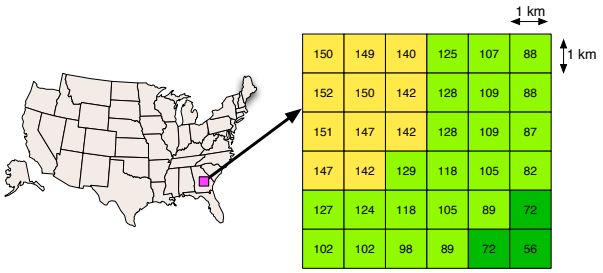

Figure 1. Sample Lincoln Lab Convective Weather Forecast near ATL.

2006). In this model, a weather cell is represented by a polygon, and the probability that weather will occur at a point in the polygon decreases with increased distance from the center. This structure is then used to estimate flight lengths and deviation delays. However, the model is not validated against behavior of weather, and for the polygons to have much meaning, they might have to be very large and hence not useful for fine-grained ATFM.

3. LINCOLN LAB CONVECTIVE WEATHER FORECAST

In order to create a realistic model of terminal area capacity, it is necessary to start with some available weather forecast. In this section we present the details of Lincoln Lab's Convective Weather Forecast (CWF), which we use in modeling terminal capacity.

The 0-2 hour CWF consists of a grid of 1km x 1km pixels covering a region of the NAS (Wolfson et al., 2004). The region will contain the entire continental United States as of 2008, though we only use data from Atlanta's Hartsfield-Jackson Airport (ATL) for the purposes of this work. Each pixel contains a value of Vertically Integrated Liquid (VIL), indicated by an integer value in the range $[0, 255]$. Figure 1 shows a sample forecast for ATL. A VIL value above a certain threshold (133, in practice) corresponds to weather of severity level 3 or higher, which is commonly considered hazardous weather that pilots are thought to avoid. A forecast has a horizon every 5 minutes between 5 and 120 minutes, and is updated every 5 minutes. In other words, at time T_0, forecasts are available for time $T_0 + 5, T_0 + 10, T_0 + 15, \ldots, T_0 + 120$. The forecast data is accompanied by observed VIL values for the same region of airspace at that time, allowing for validation of the quality of the forecast.

This static forecast is useful in obtaining a general idea of what weather will look like, and is used in various decision support tools by air traffic controllers and airlines. Lincoln Lab, as well as other forecast products, provide daily statistics such as rates of false positives, false negatives, and a skill score, but these vary daily and by storm. However, no large-scale historical validation of accuracy has been performed.

The natural first step in modeling airspace capacity is to see if historical trends in forecast skill can be correlated with capacity. In effect, it seems desirable to compute the conditional distribution of actual weather given a forecast. This approach is further discussed in the next section.

4. A PIXEL-BASED MODEL

One possible approach to developing a stochastic model of airspace capacity using deterministic forecast data is to determine a distribution for the probability that hazardous weather occurs at pixel x given the forecast value at pixel x. This probabilistic forecast can then be mapped into a distribution for capacity. This section discusses our attempts in this direction, and the associated issues.

Let $w_x(t)$ be the observed weather at pixel x at time t. Let $f_x(t, \tau)$ be the τ-minute weather forecast for pixel x at time t. In other words, $f_x(t, \tau)$ is the forecast made at time $t - \tau$ for time t. We would like to know the conditional distribution $\Pr(w_x(t) = v \mid f_x(t, \tau) = u)$. Note that this distribution would likely be independent of the pixel x, but might depend on the geographical area.

After trying to approximate this distribution using historical weather data from 8 days of stormy weather conditions surrounding ATL, we found that the deterministic forecasts have large errors when measured using this metric, making their use for ATFM potentially problematic.

Figure 2 shows a sample of results from this approach. The two figures contain a histogram of the VIL that actually occurred after a VIL of level 3 (VIL in the range [133, 162]) was predicted, for forecast horizons of 30 and 60 minutes. Although both plots look like they have a roughly Gaussian curve, even the 30-minute forecast results in non-hazardous weather more than half of the time. This would translate to wasted capacity if the static forecast were followed.

Several conclusions can be reached from these results. First, it is possible that the forecasts get the general weather trends right, but are incorrect in the position of the weather cells. This turns out to be a major shortcoming of the pixel-based approach. A

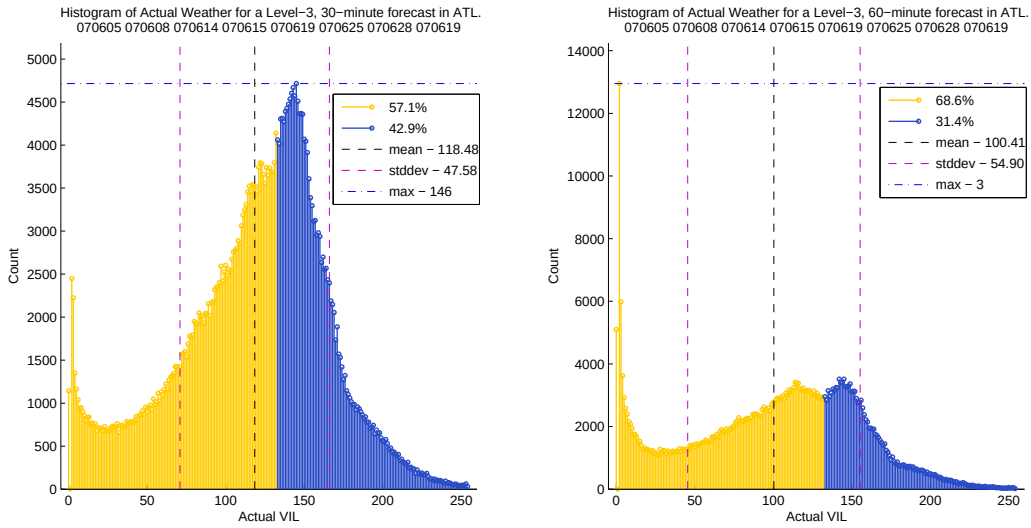

Figure 2. Plot of true VIL when Level 3 VIL (in the range [133, 162]) is forecast. Results for 30-minute and 60-minute forecast are on the left and right, respectively.

10km x 10km storm cell forecast, for example, might be displaced by 10km to the west when observed, resulting in no correct pixel predictions. For planning purposes, however, this forecast is quite good because moving an aircraft's planned trajectory 10km east might be acceptable.

Along similar lines, different storm types have different implications for ATFM. A storm consisting of many small sparsely distributed cells of weather activity (called a popcorn storm), might have very low forecast accuracy in pixel-by-pixel comparisons, because it is hard to predict the exact location of a small cell. However, there may be routes through the non-stormy sections between the popcorn cells, resulting in no practical loss in capacity.

These observations suggest that a route-centric approach may be a better way to think about weather forecasts used for TFM. Identifying persistent routes through weather might identify opportunities for increased capacity even in the presence of storms and of inaccurate forecasts. The remainder of this paper outlines how we validate static forecasts and correlate path characteristics with capacity.

5. ROUTE-BASED CAPACITY FORECAST

The main contribution of this paper is a route-based approach to modeling terminal airspace capacity. In this section we first formalize the terminal area capacity problem. Then we describe our data-driven approach to measuring the robustness of a trajectory through airspace. The high-level steps taken are: 1) Compute the theoretical capacity of a region R of airspace using an algorithm for continuous maximum flow developed in Mitchell (1988) and Mitchell et al. (2006). The resulting minimum cut contains the bottlenecks for flow in R. 2) Use the minimum cut to generate paths through R by solving set of integer programs. 3) For each generated path p, solve another linear program to find a path p' in the neighborhood of p, in the observed airspace. 4) Using data from steps 2 and 3, build a model for the probability that a route will be available by correlating blockage with various features of a path in the forecast space, such as mean distance to weather. 5) Use this model for path robustness to determine a probability distribution for the capacity of R.

5.1 Problem Formulation

Consider the following version of the terminal airspace capacity problem. The input is a terminal area, defined by two concentric circles: an outer circle C_O of radius R, and an inner circle C_I of radius r. The outer circle C_O represents the points at which aircraft first enter the terminal airspace, and R is typically about 40 nautical miles, or 75 km. The inner circle C_I represents the point at which aircraft start their final descent into the airport. Figure 3 illustrates the model of the terminal area described.

We are interested in determining the answers to the following questions:

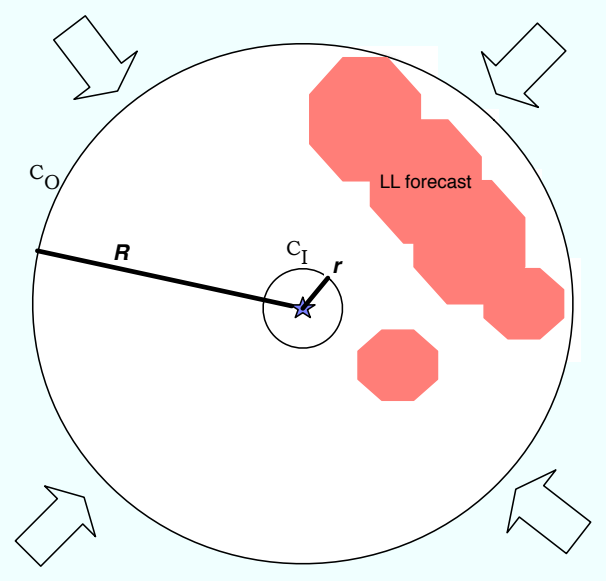

Figure 3. Model of terminal area flow. Flow comes in through the outer circle of radius R, C_O, and into the inner circle C_I. The red region within the terminal region represents a weather hazard according to the static weather forecast for some time t.

1. If aircraft arrive at C_0 t minutes from the current time, how many will be able to get through to C_I? This is the question of theoretical capacity.

2. If aircraft are routed along trajectories that are clear according to the t-minute weather forecast, what is the probability that that these trajectories will be clear in the weather that actually materializes?

This paper aims to answer the second question. We believe this question is an important one for air traffic control for several reasons. First, it aims to capture trends in how true weather cells differ from predictions. And second, it better takes into account the realities of scheduling aircraft routes. Theoretical capacity might suggest that N aircraft will be able to enter airspace over the next hour, but may not indicate the possibility (which is critical for planning) that these aircraft must necessarily arrive from the West, for instance. Furthermore, it is possible for a forecast of theoretical capacity to exactly match the true theoretical capacity, yet require that aircraft use trajectories that are very far from original plans.

5.2 Dynamic Weather Grid

Since we assume that aircraft move inwards towards the airport from the outer circle, we can splice together weather data for time instants t that increase from the outer to inner circle. This effectively creates a dynamic weather grid.

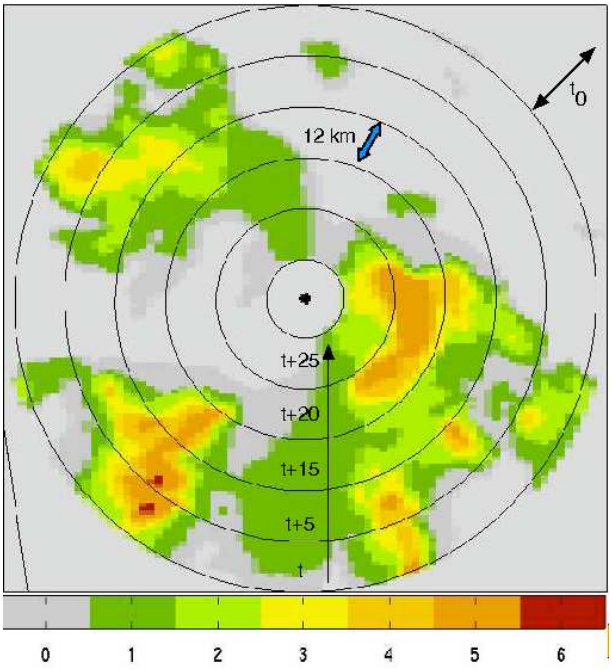

Figure 4. Sample forecast region, created by splicing together consecutive 5-minute forecasts. This is for a 30-minute time horizon on July 29, 2007, where aircraft reach the outer circle at time 21:00. The regions with Level 3 and higher weather become obstacles in the forecast grid.

Figure 4 illustrates how this is done, showing a sample aircraft coming in, for reference. We assume aircraft arrive at C_O at time t, with a t_0- minute time horizon. This captures all planning that can occur at time $t - t_0$.

5.3 Synthesizing Paths

Given a t_0-minute time horizon, we can model paths through airspace which are likely to stay open. Our approach is a data-driven approach to exploring route robustness. This section outlines the steps taken in creating a dataset of paths through weather-constrained airspace.

Theoretical Capacity

First we compute the theoretical capacity in a dynamic forecast grid. This computation identifies the bottlenecks for flow in the region. To compute the theoretical capacity, we follow the developments

on continuous maximum flow extended to the case of airspace in Mitchell et al. (2006), Krozel et al. (2007), and Mitchell (1988). This work presents a polynomial-time algorithm for computing the maximum flow through a polygon with holes, from a set of source edges to a set of sink edges. In our case, the polygon represents the terminal airspace, the holes represent weather, and C_0 and C_I are sets of sources and sinks, respectively. The algorithm involves the creation of a discrete *critical graph*, where a shortest path through this graph gives the cost of the minimum cut through the continuous region, and is also equal to the maximum flow.

The minimum cut gives the bottleneck for flow through the airspace, and all trajectories would necessarily have to pass through this bottleneck region.

Finding Paths in the Forecast Grid

We identify potential aircraft trajectories through the forecast grid by solving the following modified shortest path problem as an integer program:

Given a snapshot of the weather forecast for the terminal region taken at time $t - t_0$, construct a directed graph $G(\mathcal{N}, \mathcal{A})$ such that the set of nodes $\mathcal{N}$ contains all pixels free of weather hazards, and each set of adjacent nodes form an arc $a \in \mathcal{A}$ as long as the arc moves towards the center. At time t, a unit of flow is sent from a set of source nodes $\mathcal{S} = \{s_1, \ldots, s_l\} \subseteq \mathcal{N}$ to a set of sink nodes $\mathcal{T} = \{t_1, t_2, \ldots t_m\} \subseteq \mathcal{N}$, and must pass through a node $K \in \mathcal{N}$. K could correspond to points of interest, such as the midpoint of the minimum cut or arrival gates. For simplicity, we use a standard transformation and introduce a supersource $\bar{\mathcal{S}}$ and a supersink $\bar{\mathcal{T}}$, and route one unit of flow between the two through the source nodes and sink nodes (Ahuja et al., 1993). Define $\mathrm{NX}(i, j)$ to be the node $k \in \mathcal{N}$ which constitutes a straight next arc if (i, j) is used. In other words, nodes i, j, k form a straight line in the weather grid, pointing towards the center. The objective is to find the minimum cost flow f such that out of all minimum cost flows, f has the minimum number of turns.

This problem is modeled by the IP below, and can be solved with different sets of sources and sinks and nodes K to generate a large set of candidate paths for a given weather forecast scenario.

$x_{ij} :=$ flow on arc $(i, j) \in \mathcal{A}$

$z_{ij} := 1$ if $(i, j) \in \mathcal{A}$ is a turn, 0 otherwise.

$$\min \sum_{(i,j)\in\mathcal{A}} c_{ij}x_{ij} + \lambda \sum_{(i,j)\in\mathcal{A}} z_{ij}$$

$$\text{s.t.} \quad \sum_{\substack{j\in\mathcal{N}:\\(i,j)\in\mathcal{A}}} x_{ij} - \sum_{\substack{j\in\mathcal{N}:\\(j,i)\in\mathcal{A}}} x_{ji} = b_i \quad \forall i \in \mathcal{N} \qquad (1)$$

$$\sum_{\substack{j\in N:\\(K,j)\in A}} x_{Kj} = 1 \qquad (2)$$

$$z_{ij} \geq x_{ij} - \sum_{\substack{k\in NX(i,j):\\(j,k)\in A}} x_{jk} \quad \forall (i,j) \in \mathcal{A} \qquad (3)$$

$$x \in \{0,1\}^n \qquad (4)$$

$$z \in \{0,1\}^n \qquad (5)$$

Constraints (1) are the flow balance constraints, with $b_i := -1$ for a supersource $\bar{\mathcal{S}}$, $b_i := +1$ for a supersink $\bar{\mathcal{T}}$, and $b_i := 0$ for all other nodes i in $\mathcal{N}$. Constraint (2) ensures that the path goes through node K. These first two constraints along with the non-negativity constraint (4) finds shortest-paths through the airspace region. However, these constraints alone do not limit the number of turns in the resulting trajectory. It is however desirable that aircraft trajectories have a limited number of turns for simplicity. Constraints (3) serve to minimize the number of turns in the path without changing the path length. All arcs that follow (i, j), *except* (j, k) for $k = \mathrm{NX}(i, j)$, pay a penalty in the objective function. We set λ to be sufficiently small (less than the maximum length of any path) to ensure that a longer route but with fewer turns is never chosen. Finally, x and z are binary variables because a single path cannot be split up, and the existence of a turn is a binary quality.

We generate paths for many pairs of source and sink sets, each passing through each segment of the minimum cut.

Validating Paths in Observed Weather Grid

Given a set of paths through a region of airspace, we validate the paths against observed weather. We define a route u as *open* if there exists a corresponding route in the forecast grid which is within B km of u and does not pass through any actual weather hazards.

We synthesize open routes by solving an IP identical to the one defined in the previous section, except over a modified graph $\mathcal{G}'(\mathcal{N}', \mathcal{A}')$, and without the requirement of passing through node K. $\mathcal{G}'$ corresponds to the dynamic grid of *observed* weather in the neighborhood of u. $\mathcal{N}'$ contains all nodes in

the truth grid that are within B km of the forecast route u, and $\mathcal{A}'$ contains all pairs of adjacent nodes in $\mathcal{N}'$. The buffer B allows for slight perturbations in the path (on the order of several kilometers), which represents only a slight change from the original planned trajectory of u.

Figure 5 contains a few examples of routes synthesized in the forecast gird, side-by-side with the same routes validated against true weather. Weather scenarios corresponding to time horizons of 10, 30, and 60 are included to show the difference in pixel accuracy at different forecast horizons. The figure also shows two routes that become blocked, and one that remains open, subject to slight changes in the planned trajectory.

5.4 Creating a Dataset of Routes

This section describes details of steps taken to generate a dataset of trajectories through forecast weather with corresponding trajectories through true weather.

Selection of Weather Scenarios

A dataset was created containing routes for several weather scenarios during the 18 worst weather days in June and July 2007, ranked according to weather-related delays.

Though the Lincoln Lab CWF data as described is simply a matrix of integers in the range $[0, 255]$, the archives of this data are kept in a proprietary format, and each day of data takes several hours to extract and decompress, leaving 30 GB of uncompressed binary data. To identify the times of day which contain stormy weather, we extracted the forecasts for airspace surrounding ATL and watched visualizations of the entire two month time period to identify the stormiest hours. This resulted in an average of 4 weather scenarios per day, spaced 30 minutes apart, for a total of over 300 trajectories in forecast weather. Four datasets were created, corresponding to the 10, 30, 60, and 90 minute time horizons.

Dataset Details

Tables 1 contains overall statistics of route blockage for the datasets. The vast majority of synthesized routes end up being open in the true weather grid, even for a time horizon of 60 minutes, which had poor pixel-based accuracy. This suggests that subject to minor adjustments, planning at a 10-, 30-, 60-, and 90-minute time horizon is quite reasonable. This is encouraging, and shows that moving away from fixed jet routes gives us the flexibility in the airspace to allow that to happen.

Horizon t_0	Size (# Paths)	Viability rate (%)
10	356	87
30	369	87
60	335	80
90	323	74

Table 1. Overall dataset statistics. Each time horizon contains over 300 paths through different weather forecast scenarios, and the viability rate shows the rate that these routes are viable in observed weather.

For each path, eight features of interest were indentified and each feature was correlated with route blockage. The eight features, chosen for their possible correlation with route blockage, are listed below:

1 Mean VIL along path
2 Standard Deviation of VIL along path
3 Min distance to level 3+ weather along path
4 Mean distance to level 3+ weather along path
5 Number of turns in path
6 Theoretical capacity for weather scenario
7 Number of segments in the minimum cut
8 Length of path's minimum cut segment

5.5 Robust Routes

This section introduces a method for route robustness based on creating correlations of features with blockage.

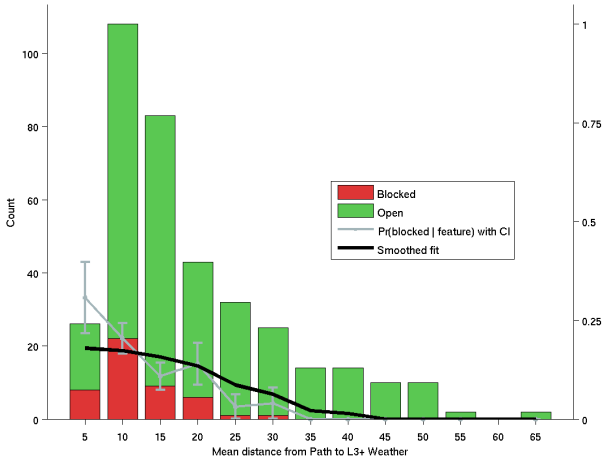

Figure 6. Individual features give conditional probabilities of route blockage. This figure shows this probability (black line) as well as raw data (histogram) for Feature 4 at a 30-minute time horizon.

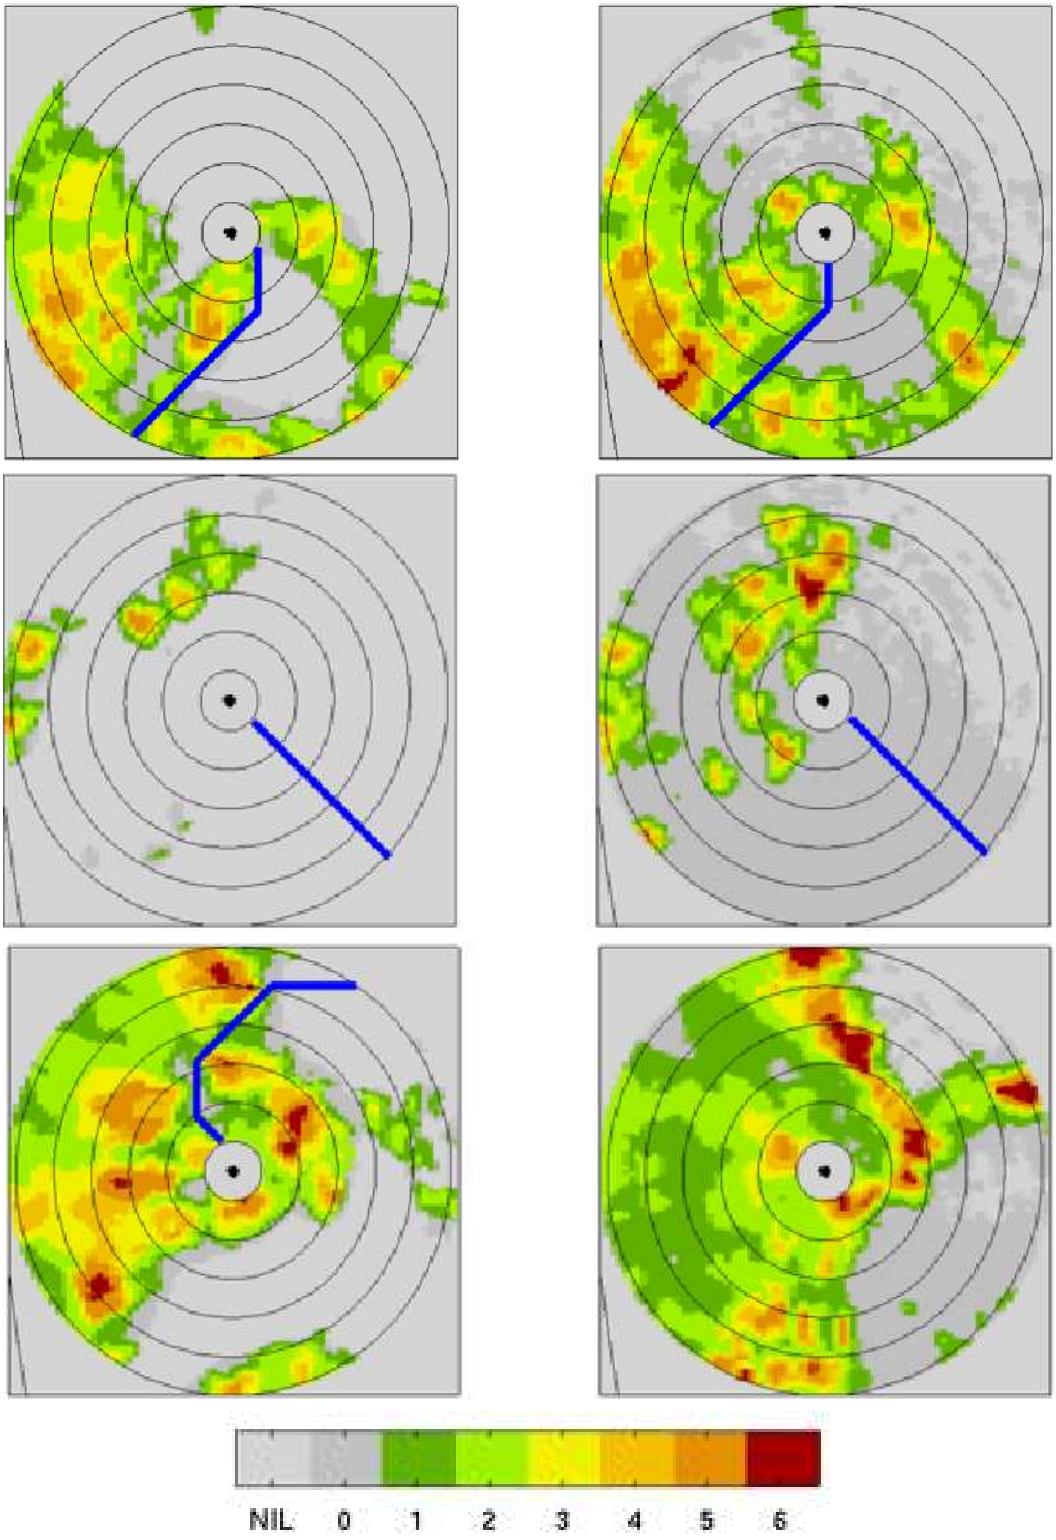

Figure 5. Sample routes in a dynamic forecast grid on are in the left-hand column, and corresponding routes in the actual weather are on the right-hand column. The top weather scenario is from June 8, 2007 at 1930hrs with 10-minute time horizon. The second is from June 28, 2007 at 2300hrs with 30-minute time horizon. The bottom is from June 8, 2007 at 2100hrs with 60-minute time horizon. Notice that the middle route is viable in the observed weather, while the other two are not.

Intuitively, if a route through airspace is open according to the forecast, but has an average VIL of just below Level 3, we expect it is more likely to be blocked in the actual weather than a similar route with average VIL at Level 0. Using this idea, we estimate the conditional probability that a route is blocked given its value for various features. The following equation is used to compute these conditional probabilities:

$$\mathrm{P}(\ u \text{ blocked } | f_i(u) = v) = \frac{\mathrm{P}(f_i(u) = v \ \&\ u \text{ blocked})}{\mathrm{P}(f_i(u) = v)} \tag{6}$$

$$= \frac{\#(\text{blocked} \ \&\ f_i = v)}{\#(f_i = v)} \tag{7}$$

where $f_i(u)$ is the value of Feature i for route u.

Since the denominator in equation 7 may be 0 in the case of continuous data, the feature values must be binned where necessary. Due to sampling error, these conditional probabilities contain some noise. To account for this noise, smoothing was used to create revised estimates of the desired conditional probabilities. The smoothed conditional probability $\mathrm{P}(\ u \text{ blocked } | f_i(u) = v)$ was computed by taking the average of 5 neighboring bins (bin v as well as 2 bins on each side of v), weighted by the number of paths in each bin.

Figure 6 shows a histogram of the raw data for Feature 4 in the 30-minute time horizon. The green (red) section of the bar shows how many routes having value v ended up being open (blocked) when the route was validated against true weather conditions. In addition, the light gray line overlaying the plot shows the raw conditional probabilities as computed by equation (6), along with confidence intervals. The black line corresponds to the smoothed conditional probability that a route is blocked given Feature 4.

Similar correlations for a few other features at 10- and 60- minute time horizons are show in Figure 7 and Figure 8. The remaining features and time horizons are omitted in the interest of space constraints. Looking at these figures, there is correlation between features and blockage, and the correlation is especially strong for the shorter time horizon of 10 minutes. This is to be expected, since weather forecasts for a shorter time horizon are known to be more accurate. We see for mean VIL (Feature 1), standard deviation of VIL (Feature 2), and number of turns (Feature 5), that as the value of the feature increases, so does likelihood of blockage. These trends agree with intuition. Indeed, high VIL values along a path indicate the route moves through some weather-affected regions, which are more likely to show up as level 3 or higher in the true weather. Likewise, higher standard deviation of VIL indicates increased variability in weather conditions along the path, and hence higher likelihood that weather will materialize. And a high number of turns indicate that a route requires lots of maneuvering to avoid weather, so is more likely to be exposed to moving cells. For feature 3 (minimum distance to weather), we see that synthesized routes that are very close to weather end up being blocked more of the time, and that routes that are very far from any forecast weather stay viable. The same trend is found with average distance to weather (Feature 4). Finally, theoretical capacity (Feature 6) shows the general trend that increased theoretical capacity correlates with decreased likelihood of being blocked. The plots for Feature 7 were left out, because they showed very little correlation with blockage. This is probably because the number of segments in a minimum cut can occur both if the theoretical capacity is very high (there is 1 large cut segment, for instance), or very low (there is a very small cut segment). The high capacity case would make routes more likely to be viable, since they are less likely to be close to adverse weather cells, while the reverse is true for low capacity.

5.6 From Routes to Capacity

The route blockage model can be used to create a stochastic model of capacity. This section presents an initial version of such a model. For an airport with m arrival gates (in the case of ATL, $m = 4$), and for a given time horizon t_0, we can forecast capacity in the following way. First, four routes are synthesized through the forecast grid, each sourced from a different quadrant of the outer circle C_O. This can be done using the IP in Section 5.3. Next, for each synthesized route, let the overall viability rate for the given time horizon (listed in Tables 1) be the probability that the route will be blocked in the true weather grid. Let C be the clear-weather capacity of the airspace. Then the capacity of the the airspace can be forecast as $\frac{Ck}{m}$ with probability $\mathrm{Pr}(\text{ exactly } k \text{ of the synthesized routes are open})$, which will be a binomial distribution.

6. CONCLUSION

A route-based approach to modeling airspace capacity turns out to be fruitful for estimating airspace capacity using actual weather data. By analyzing stormy weather data from the summer of 2007, we were able to show that certain features of a candidate

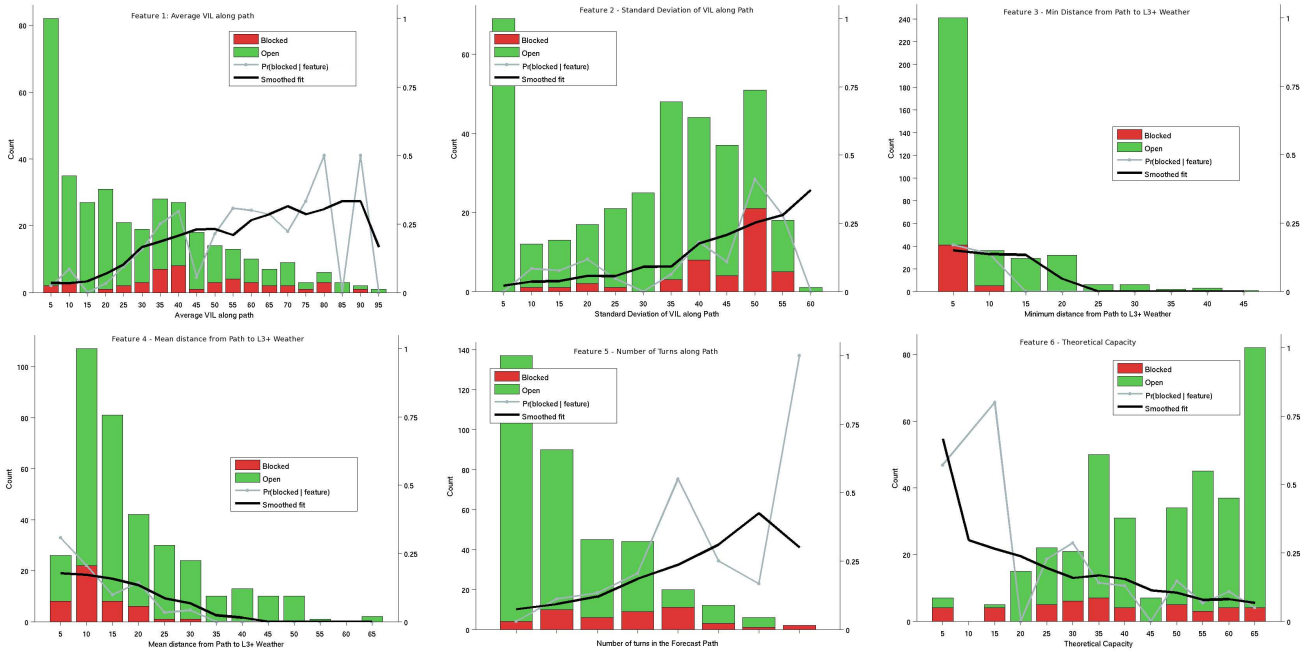

Figure 7. Plots showing how the values of the first 6 features correlate with route blockage at a 10-minute time horizon. The histogram showing the distribution of feature values and blockage is overlaid with the smoothed probability that a route is blocked given each feature value. Features 2, 3, and 4 correlate especially well.

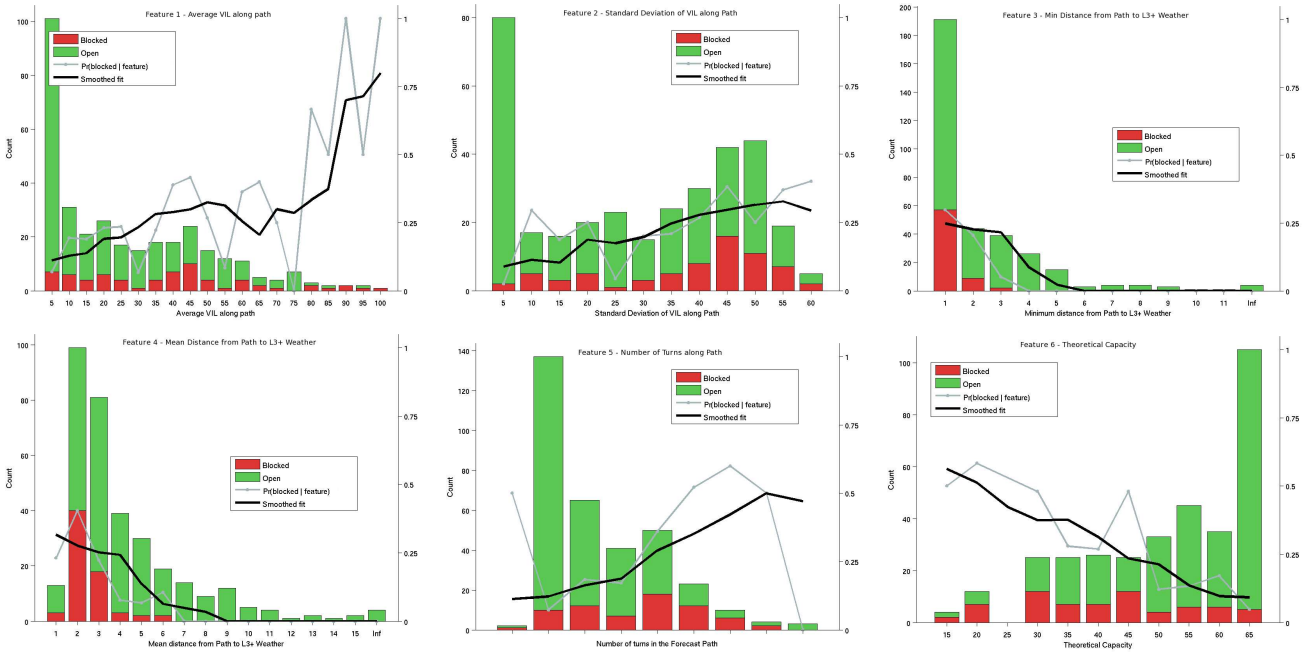

Figure 8. Plots showing how the values of the first 6 features correlate with route blockage, for 60-minute time horizon.

trajectory through airspace correlate with blockage in true weather. We found that VIL levels along the trajectory and in neighboring regions, along with the complexity of the route (number of turns), to be good individual predictors for blockage. Furthermore, we introduced an integer program for synthesizing turn-constrained routes through forecast regions, which turned out to produce routes that tended to be viable in the true weather grid, subject to slight shifts in position. This is promising for short-term route planning under uncertain weather constraints. Future work includes the incorporation of departures, additional weather features such as echo tops, and ultimately, using these probabilistic capacity forecasts for ATFM.

7. Acknowledgments

The authors are grateful to Marilyn Wolfson, Rich De Laura and Mike Robinson at MIT Lincoln Lab for help with CIWS data and fruitful discussions.

References

Ahuja, R. K., T. L. Magnanti, and J. Orlin, 1993: *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall.

Bertsimas, D. and A. Odoni, 1997: A critical survey of optimization models for tactical and strategic aspects of air traffic flow management. Tech. rep., NASA Ames Research Center.

Bertsimas, D. and S. S. Patterson, 2000: The traffic flow management rerouting problem in air traffic control: A dynamic network flow approach. *Transportation Science*.

Bureau of Transportation Statistics, 2008a: `http://www.transtats.bts.gov/`.

Bureau of Transportation Statistics, 2008b: Understanding the reporting of causes of flight delays and cancellations. `http://www.bts.gov/help/aviation/html/understanding.html`, Retrieved 19 April 2008.

DeLaura, R. and S. Allan, 2003: Route selection decision support in convective weather: A case study. *USA/Europe Air Traffic Management R&D Seminar, Budapest*.

IATA, 2007: *Passenger and Freight Forecasts 2007 to 2011*. `http://www.iata.org/NR/rdonlyres/E0EEDB73-EA00-494E-9408-2B83AFF33A7D/0/traffic_forecast_2007_2011.pdf`.

IATA, 2008: *Industry Statistics Factsheet*. `http://www.iata.org/NR/rdonlyres/65A3F2C3-3656-4243-8396-3D77CB31C2FC/0/factsheet_industry_stats_apr20082.pdf`, Retrieved 19 April 2008.

Joint Economic Commitee Majority Staff, 2008: Your flight has been delayed again. Tech. rep.

Krozel, J., C. Lee, and J. S. B. Mitchell, 2006: Turn-constrained route planning for avoiding hazardous weather. *Air Traffic Control Quarterly*, **14 (2)**, 159–182.

Krozel, J., J. S. B. Mitchell, V. Polishchuk, and J. Prete, 2007: Capacity estimation for airspaces with convective weather constraints. *AIAA Guidance, Navigation and Control Conference and Exhibit, Hilton Head, South Carolina*.

Mitchell, J. S., 1988: On maximum flows in polyhedral domains. *SCG '88: Proceedings of the fourth annual symposium on Computational geometry*.

Mitchell, J. S. B., V. Polishchuk, and J. Krozel, 2006: Airspace throughput analysis considering stochastic weather. *AIAA Guidance, Navigation, and Control Conference, Keystone, Colorado*.

Prete, J. and J. S. B. Mitchell, 2004: Safe routing of multiple aircraft flows in the presence of time-varying weather data. *AIAA Guidance, Navigation, and Control Conference and Exhibit, Providence, Rhode Island*.

Seseske, S. A., M. P. Kay, S. Madine, J. E. Hart, J. L. Mahoney, and B. Brown, 2006: *2006: Quality Assessment Report: National Convective Weather Forecast 2 (NCWF-2)*. Submitted to FAA Aviation Weather Technology Transfer (AWTT) Technical Review Panel.

Sheth, K., B. Sridhar, and D. Mulfinger, 2006: Impact of probabilistic weather forecasts on flight routing decisions. *AIAA Aviation Technology, Integration and Operations Conference, Wichita, KS*.

Wolfson, M., et al., 2004: Tactical 0-2 hour convective weather forecasts for FAA. *11th Conference on Aviation, Range and Aerospace Meteorology, Hyannis, MA*.