

PROBABILISTIC CONGESTION MANAGEMENT

*Craig Wanke, Lixia Song, Stephen Zobell, Daniel Greenbaum, and Sandeep Mulgund
The MITRE Corporation, cwanke@mitre.org, McLean, VA*

Abstract

A new approach to airspace congestion management is presented, in which uncertainties in predicting traffic levels and airspace capacity are quantified and considered in developing congestion resolutions. Such a “probabilistic” approach has the potential to greatly improve decision-making. Probabilistic models for predicting traffic demand are presented, with results, and an initial model for predicting airspace capacity is outlined. Candidate displays and human factors issues for use of probabilistic information are discussed, and a probabilistic decision-making framework for en route congestion management is presented.

Introduction

In the U.S. National Airspace System (NAS), en route Traffic Flow Management (TFM) is the function which balances air traffic demand against available airspace capacity, to ensure a safe and expeditious flow of aircraft. A variety of flow control actions, such as weather avoidance routes, miles-in-trail (MIT) flow restrictions, and ground delay programs (GDPs) are used to achieve this. Planning these actions requires predictions of both traffic demand and airspace (en route sector) capacity. Since TFM decisions are typically made 30 minutes to several hours in advance of anticipated congestion, these predictions are subject to significant uncertainty. However, the magnitude of this uncertainty is not known, presented, or understood. As a result, decisions are often overly conservative, and may be taken at inappropriate times based on the actual accuracy of prediction data.

Background

Traffic demand uncertainties arise from many sources. Flight schedules undergo constant changes in response to daily events, and such changes often occur between the time of demand prediction and the time for which demand is predicted. These include flight cancellations, departure time changes, and initiation of previously unscheduled flights. This latter category is increasing in the U.S., as air taxi and “executive jet” operations become more prevalent.

There are uncertainties in wind forecasting and aircraft performance modeling, and unforeseen changes in flight route and cruising altitude due to weather and air traffic control (ATC) intervention. The aggregate effect of uncertainties has been measured for several months of NAS operations. [1]

Airspace capacity itself is harder to quantify. The current Enhanced Traffic Management System (ETMS) [2] provides a congestion alerting function which uses peak one-minute aircraft count as a sector congestion alerting criterion (the “Monitor/Alert Parameter,” or MAP). This is not meant to be a measure of airspace capacity, but rather a threshold which, when exceeded by predicted demand, alerts traffic managers to examine the sector for potential congestion. The actual capacity of a sector is dependent on the complexity of the traffic flows within, as well as the presence or absence of hazardous weather. Thus, predicting sector capacity is also subject to uncertainty, particularly in predicting convective weather.

Research Focus

A notional probabilistic forecast of congestion is shown in Figure 1. If the uncertainties in traffic demand and capacity predictions are known and quantifiable, then a probability of congestion can be calculated by convolving the two distributions. “Congestion”, in this case, is simply defined as when demand exceeds capacity. The red and yellow codes indicate the probability of congestion for the three sectors shown.

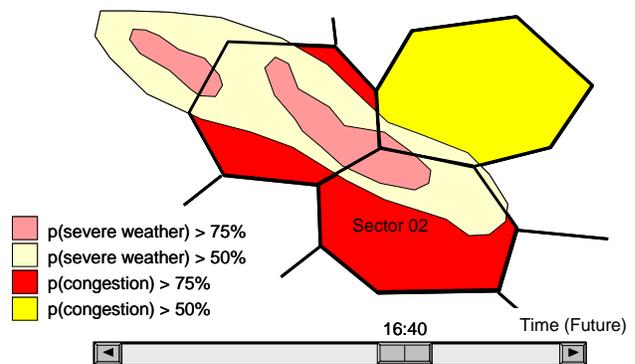


Figure 1. Probabilistic Congestion Forecast

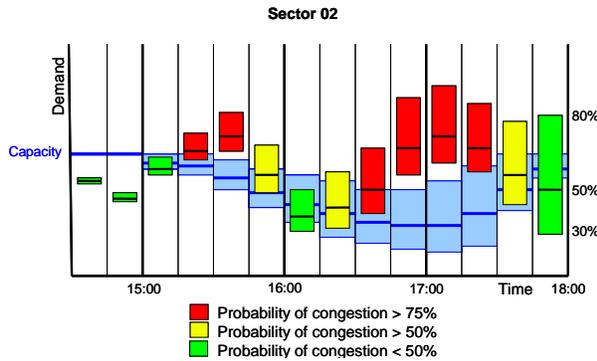


Figure 2. Detailed Congestion Predictions for a Single Sector

In addition, the probability of severe weather coverage is shown, based on a probabilistic severe weather forecast. Such products are now becoming available, and will be operational in the next one or two years. [3,4]

Figure 2 represents the same situation in a time-series plot. In this case, the demand and capacity prediction distributions for Sector 02 are broken out and shown as ranges for a series of 15 minute intervals. The blue line represents the 50th percentile prediction of airspace capacity. At short prediction look-ahead times (LAT), the capacity is well-known, since weather is not predicted to impact the area until, at earliest, 30 minutes into the future (1500). At greater LAT, the weather is expected to reduce capacity, and the spread in possible values of that capacity reflects the uncertainty in the future position, size, and intensity of the weather.

The green, red, and yellow boxes reflect probabilistic demand predictions. The heights of the boxes reflect the uncertainty, increasing in size with increasing LAT. The bottom, midline, and top of the boxes represent the 30th, 50th, and 80th percentile of the predicted demand distribution, in this example. The boxes are color-coded to reflect the probability that the actual demand exceeds the actual capacity.

To date, current and proposed TFM decision support systems have ignored prediction uncertainties. As a result, they tend to produce conservative, overly strict problem resolution actions. This work is focused on displays, algorithms, and decision-making techniques which explicitly acknowledge and account for prediction uncertainty. Effective TFM decision-making in the presence of uncertainty, or “probabilistic TFM”, should have the following characteristics:

- Rather than attempting to resolve all possible congestion, incremental actions are taken to

keep congestion risk at a tolerable level, while retaining flexibility to take further action as the situation becomes more certain.

- Predicted congestion areas are continually re-evaluated for further control action.
- NAS users are informed of predicted congestion, so they can proactively reduce schedule risk if desired (e.g., by replanning flights through less congested airspace.)
- Probabilistic congestion predictions are presented to traffic managers and NAS users in an intuitive way, to maintain good situation awareness.

Several new techniques and technologies are required to provide probabilistic TFM decision support. First, prediction uncertainty must be known and quantifiable. Second, a metric is needed for rating the goodness of candidate solutions. Third, decision-making algorithms are needed to develop congestion solutions, given the prediction uncertainty and goodness metric. And finally, there are significant human factors issues to be resolved due to the combination of information uncertainty and complex automated processes.

Probabilistic Congestion Prediction

Figure 2 illustrates the desired goal of probabilistic predictions, specifically, statistical distributions of demand and capacity for the near future, which can be used to determine the likelihood of congestion. The precise form and characteristics of these predictions depends on the metric used for sector congestion alerting. If peak count is used, then the demand distribution can be expressed as a histogram of possible peak count outcomes. Since this is the current NAS practice, an initial uncertainty model was developed for peak count predictions.

Predicting Aggregate Traffic Demand

A comprehensive set of statistics on sector peak count prediction uncertainty was compiled in past research [1]. In that work, three variables were identified as having strong effects on demand uncertainty distributions: LAT, predicted peak count (N), and the primary sector traffic type. N is the sum of proposed (P) and active (N-P) flights, and further research has determined that the proportion of active flights is also an important variable.

Based on this work, a parameterized statistical model was developed. This model is able to forecast

probabilistic distributions of sector demand (peak counts) as a function of the four variables described above. These distributions can be used to drive probabilistic demand displays, such as that shown in Figure 2, as well as probabilistic congestion resolution algorithms.

Sample Data Collection

The peak sector count predictions in 15-minute intervals were collected from a TFM decision support prototype, using traffic prediction algorithms similar to those used in the ETMS. Every 15 minutes, predictions of active and proposed aircraft with LAT of 0 to 4 hours are recorded. The modeling data is based on 171 days, between 1 Jan 2004 and 30 Jun 2004, for the 754 NAS sectors with established MAP values. The sample data includes approximately 185 million predictions of 12 million actual 15 minute peak sector counts.

Each sector is assigned to one of four traffic categories based on predominant traffic type: (A) arrival, (D) departure, (E) en route, and (M) mixed [1]. Of the 754 sectors categorized, 47 are arrival sectors, 49 are departure sectors, 406 are en route sectors, and 252 are mixed traffic sectors.

Sample data is arranged to group predicted counts with the sector counts that actually occurred. The following is an example record of sample data:

LAT = 45 minutes
 N = 8 flights (predicted)
 P = 6 flights (predicted)
 Sector Type = Mixed
 Actual = 11 flights

In this case a prediction was made for a Mixed sector that the peak sector count will be 8 flights in the 15 minute time interval that begins 45 minutes in the future, and 6 of those flights are proposed. When 45 minutes passed, the actual peak count was 11.

Although 6 months of sample data were used to develop the model, in many cases usable data does not exist for all situations we want to model. Analysis suggests that flow control actions are sometimes taken to avoid congestion when predicted peak counts are greater than MAP - 5, and since the model should reflect the traffic counts that would occur in the absence of TFM action, these predictions were removed from the sample data [1]. As a result, the model must be able to extrapolate predictions for larger values of N than the sample data contains.

Parameterized Modeling

A Chi-Square test was used to find the best distribution to fit the sample data. The best fit is a

Binomial distribution for most cases, but a Poisson distribution is better for the cases where variance is larger than the mean, such as for very small N.

Instead of modeling the parameters of the Binomial or Poisson distribution directly, the mean and standard deviation (SD) of the sector peak count prediction were modeled. These mean and SD values were then used to calculate the parameters for the Binomial or Poisson distribution. Modeling mean and SD has several advantages. First, the parameters for both distributions can be calculated from mean and SD, so separate fitting calculations are not needed for the different distribution types. Second, mean and SD are easily interpreted to determine if the models were producing intuitive results. Third, the model should accurately predict mean, because it may be used as the best estimate for future sector counts. Finally, it became apparent that mean and SD had linear relationships to N and P for most of the data space. This made it easier to extend the model to larger values of N-P and P where sample data was not available.

Figure 3 shows the relationship of the mean of peak sector count predictions for en route sectors as the function of N-P and P, with LAT = 60 minutes in the sample data. It shows that the relationship of mean to N and P is approximately linear. Graphs of the relationship of mean to P and N-P for different LAT all show an approximately linear relationship. SD as a function of P and N-P is also approximately linear for a given LAT.

Figure 4.1 shows mean plotted as a function of LAT and P (for cases where N=P). This plot does not show a strong linear relationship of mean to LAT. A 15 minute change in LAT has a large effect for small LAT, but for larger LAT has a smaller effect, and for LAT larger than 2 hours, changes in LAT have

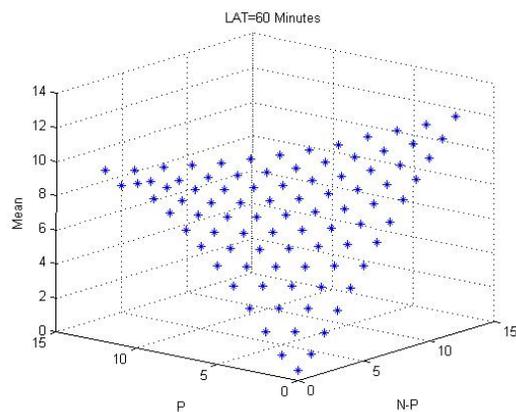


Figure 3. En route Sector Peak Count Prediction Mean as Function of P, N-P

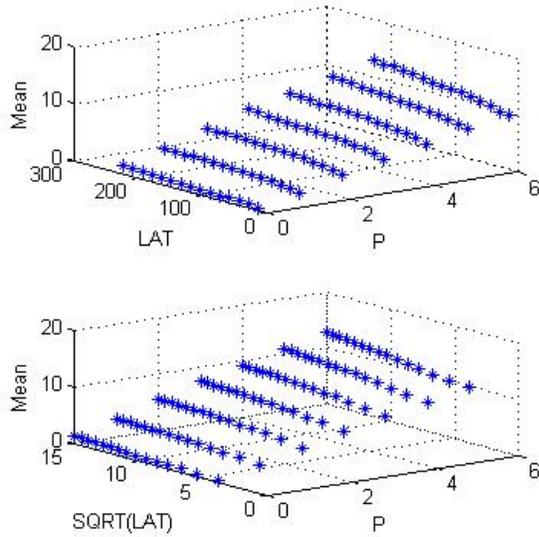


Figure 4. En route Sector Peak Count Prediction Mean as Function of P, LAT

almost no effect. Several non-linear functions were tested to see if they model the linear relationship of $f(LAT)$ to mean and SD. The square root of LAT is a close fit. That is, mean has a strong linear relationship with the square root of LAT, as shown in Figure 4.2.

The following linear function based on N-P, P, and \sqrt{LAT} models the mean and SD of the sample data:

$$Mean = a_1 \times (N - P) + b_1 \times P + c_1 \times \sqrt{LAT} + d_1$$

$$SD = a_2 \times (N - P) + b_2 \times P + c_2 \times \sqrt{LAT} + d_2$$

Further sample data analysis identified regions of the space spanned by the variables where a single linear model was not sufficiently accurate. Piecewise linear models were developed to model these regions. These exception regions are:

- LAT = 15 minutes (smallest modeled LAT),
- LAT > 120 minutes,
- Small N ($N < 5$) and LAT > 15 and sector type = departure.

LAT > 120 minutes is modeled the same as LAT = 120, since LAT has little effect beyond 120 minutes. The equations used in the exception regions are simple variations of the general equations.

The parameters a , b , c , d are found by minimizing the sum of squared errors between the prediction from the model and the mean and SD from the sample data. Table 1 shows the parameters of the

model for the mean and SD equations for En Route sectors with $15 < LAT \leq 120$ minutes.

Table 1. Model Coefficients for En Route Sectors with 15 min < LAT ≤ 120 min

	a	b	c	d
Mean	0.9570	0.9197	0.2440	-1.1724
SD	0.0878	0.2483	0.1244	0.4558

Some interesting conclusions can be derived from these coefficients. Active flights predicted to enter the sector are very good predictors of mean (96%), but proposed flights are slightly less accurate (92%). When LAT=120 minutes, the model is predicting an additional 1.5 ($0.2440 \times \sqrt{120} - 1.1724$) flights beyond what is predicted by the active and proposed flights. This reflects the presence of unexpected flights in the actual traffic; some of these flights may have not filed flight plans when the prediction was made. The c and d coefficients are adjusting for these “pop-up” flights. For SD, proposed flights are a much larger contribution to uncertainty than active flights ($b > a$), and LAT is also a contributing factor.

Probabilistic Demand Display

This model has been used to prototype and display probabilistic sector demand predictions with real traffic. One of these displays, shown in Figure 5, shows confidence intervals for predicted future sector counts for a single sector in look-ahead increments of 15 minutes. Each box shows sector count prediction confidence intervals from 25% to 75%, and the line in the middle of each box is the predicted mean. The color code represents the probability the sector count will exceed the MAP (MAP is 18 for this sector) within the interval. Yellow indicates a greater than 50% probability of exceeding the MAP, and red

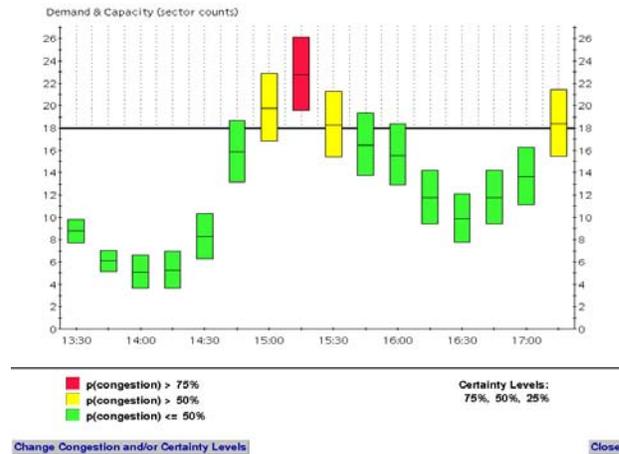


Figure 5. Sample Predicted Peak Sector Count Display Using the Aggregate Model

indicates a greater than 75% probability. These ranges for the confidence interval boxes and the color thresholds for probability of exceeding the MAP can be dynamically adjusted.

This model is useful for showing uncertainty in sector count predictions, and can provide more accurate results by correcting several sources of bias, such as correcting for the number of unexpected flights that typically occur in this type of sector.

Predicting Situation-Specific Demand

There are limitations to the aggregate approach. For example, it was necessary to treat predictions for all sectors of a particular type (e.g., “arrival”) as having equivalent statistical properties. This is a strong simplification, which hides some airspace-specific uncertainty features. It also makes it difficult to choose the optimal flights to move for a given problem, since the model treats all flights as statistically identical. In reality, different flights have different uncertainty characteristics. Finally, improved sector capacity metrics are sought. These metrics require prediction of traffic flow patterns, and the uncertainty in such predictions cannot be obtained through the aggregate method.

To address these needs, an advanced uncertainty modeling technique was developed. This method fuses empirically-observed uncertainty characteristics with Monte Carlo simulation, and is described in detail in Reference 5.

In this simulation (Figure 6), prediction uncertainties are broken down into five categories:

1. Flights that are predicted, but do not actually operate (“no shows”)
2. Flights that operate, but were not predicted at all (“pop-ups”)
3. Departure time prediction errors

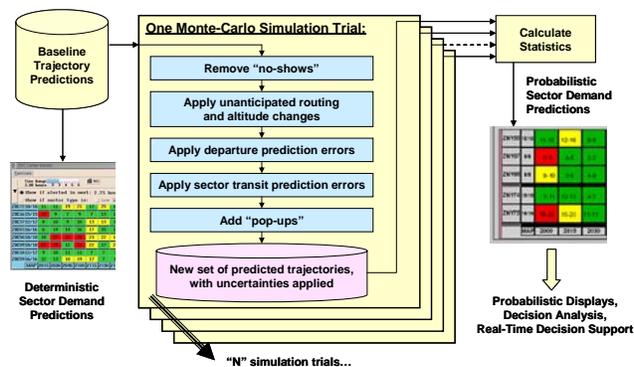


Figure 6. Monte Carlo Simulation for Prediction Uncertainty Analysis

4. Future route or altitude changes
5. Sector transit time prediction errors

Statistical models have been developed for each of these error types, based on extensive analysis of recorded traffic data and trajectory predictions. Given these models, prediction uncertainties can be calculated as follows:

1. From the baseline trajectory predictions, compute the deterministic sector demand, as in present-day ETMS Monitor/Alert predictions.
2. Conduct N Monte Carlo trials, each generating an alternate “actual” set of trajectories, for computation of probabilistic sector demand. The distributions are applied one-by-one, as illustrated in the center of Figure 6.
3. Calculate sector demand metrics from the modified trajectory set (e.g., peak counts).
4. Once the N trials are complete, calculate desired demand statistics, by sector. Flight-specific statistics can also be computed.

This model has been implemented, using statistics from a single month of recorded flight data and predictions. An example result is shown in Figure 7. In this case, a 75 min LAT prediction is shown. The most likely peak demand value is 10 aircraft, but the distribution shows significant likelihoods for peak counts between 8 and 12, inclusive.

While an initial version of the model is complete, it has not yet been validated. This can be the results can be approached in several ways. It is impossible to validate the model entirely with empirical data; predictions of high demand do not come true, since TFM or ATC actions are taken to prevent such outcomes. However, empirical data can

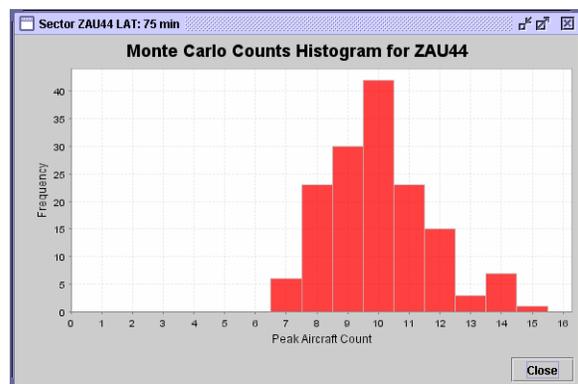


Figure 7. Sample Monte Carlo simulation results for peak aircraft count

be used in light traffic situations, and the aggregate models described in the previous section can be used to establish face validity of the results. Also, the Monte Carlo model assumes independence of the 5 uncertainty distributions, and the degree to which this assumption holds will be analytically determined.

Defining and Displaying Congestion

Assuming that uncertainties can be quantified, methods for displaying and using uncertainties are needed. An earlier study [6] presented candidate information requirements and visualization concepts for explicit representation of uncertainty in TFM decision support systems. These concepts were developed through a set of interviews with operational experts in the TFM domain. The degree of MAP exceedance (i.e., count of flights in the sector above MAP) and the range of alert start times were rated highly, as was the range in peak count estimates. The study results indicated a strong preference for a new color-coding scheme based on the probability of MAP exceedance, severity, and duration. These findings were used to generate candidate displays for visualization of results from the Monte Carlo simulation model described in the preceding section. Figure 8 presents one of these.

Rather than presenting a single predicted peak aircraft count for each 15-minute window, results are now shown as a maximum likelihood estimate for peak count along with user-specified confidence intervals. Color coding for each cell is now based on the *probability of congestion* for the 15-minute window of interest. This probability is estimated using the raw Monte Carlo results. For example, if the peak aircraft count exceeds the MAP in 50 out of

100 Monte Carlo evaluations, the probability of congestion is estimated as 0.5.

In the example display, the following rules are used to choose the colors:

Color = green if $p(\text{Congestion}) < 0.5$

Color = yellow if $0.5 \leq p(\text{Congestion}) < 0.75$

Color = red if $p(\text{Congestion}) \geq 0.75$

These thresholds were selected as an initial educated guess on sensible ranges, and are being refined through human-in-the-loop evaluation. There are several advantages of this scheme for representing uncertainty. Alerts are naturally prioritized by their probability of congestion, which makes it easier for traffic managers to identify which problems to address first. Also, the spread in the numerical ranges provides a natural, intuitive representation of the uncertainty in the predictions, allowing traffic managers to modulate their actions based on the quality of information available.

This display also allows the user to obtain histograms of the predicted count distributions. Double-clicking on a cell produces the display previously shown in Figure 7.

Predicting Sector Capacity

As noted earlier, complete probabilistic congestion predictions require both probabilistic demand predictions and probabilistic capacity predictions. Thus, a metric of sector capacity is needed that (1) is a good approximation of the amount of traffic that can be effectively handled in the sector, (2) can be predicted at look-ahead times of 30 minutes to several hours, and (3) can include the impact of convective weather on available capacity.

Capacity as a function of traffic flow pattern

The central hypothesis of this work is that sector capacity is a function of primary traffic flow patterns. Clustered traffic (flow) properties are more predictable and perturbation-resistant than individual flight characteristics. This idea is consistent with current NAS operations, in that flow pattern recognition has been identified as a primary component of traffic controllers' and managers' mental decision processes. Also, quantifying sector capacity as a function of traffic flow pattern provides a basis for capturing weather impact on sector capacity.



Figure 8: Probabilistic Demand View

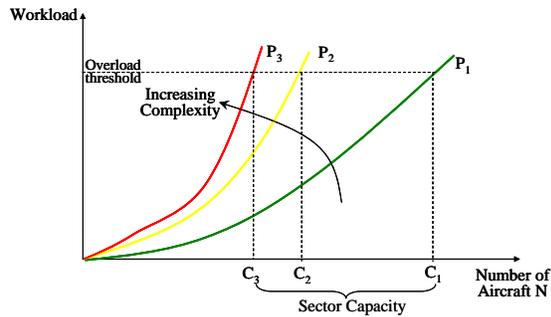


Figure 9. Relationship between Workload, Sector Capacity, Complexity, and Traffic Flow Pattern

Figure 9 shows the relationship between traffic flow patterns and sector capacity. Given a traffic flow pattern (for example, P_1), as the number of aircraft within the sector increase, controller workload increases monotonically. And given the same number of aircraft within a sector, controller workload is a function of traffic complexity, which is represented by different traffic flow patterns. For example, in Figure 9, the given sector has three normal traffic flow patterns (P_1 , P_2 , and P_3), representing three different levels of traffic complexity. When the controller workload reaches the threshold, the sector capacity is reached (C_1 for P_1). As shown in the figure, given the threshold of controller workload, the capacity is different (C_1 , C_2 and C_3) for different traffic flow patterns (P_1 , P_2 and P_3). Since P_3 is the most complex traffic flow pattern, the sector has least capacity when the traffic has that pattern.

Process of predicting sector capacity

In the initial analysis, directed sector transits were used to describe traffic flow patterns. Flights in a sector were grouped into flows based on the sector from which they entered and the sector into which they exited (a “sector transit triplet”). Figure 10 shows one observed traffic flow pattern for a single en route sector (12). Each flow is labeled by its transit triplet (e.g., flights in the purple flow entered sector 12 from sector 04, and exited to sector 19).

Two activities are required to characterize

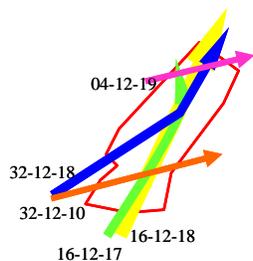


Figure 10. Sector Flow Pattern Example

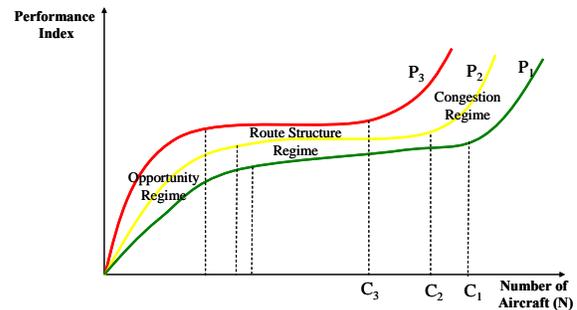


Figure 11. Sector Behavior Curves to Identify Sector Capacity

sector capacity according to the flow pattern hypothesis. First, a set of primary flow patterns for each sector of interest must be identified. Second, the sector capacity for each pattern must be established.

As far back as the 1970’s, techniques were developed to relate the traffic variables, route and sector geometry, and control procedures to an index that quantifies the workload required on the part of the air traffic control team [7]. This study asserts that workload or control difficulty is related to the frequency of occurrence of events that require the controller team to make decisions and take action, as well as to the time required to accomplish the tasks associated with those events. This study has been applied to assess sector capacities through setting the controller workload threshold (e.g., 70 percent of hourly task time) [8]. But what is the real operational workload threshold? A new methodology is proposed to assess sector capacity based on *observed* system performance transition behavior.

In an earlier study [9], a preliminary analysis was conducted to model transition in system behavior. Three regimes of system behavior were found: opportunity, route structure, and congestion. If such system behavior can be measured for a sector with a particular flow pattern, the transition point between the route structure and congestion regimes can be considered as an indicator that the controller team has reached the workload threshold. Thus, the corresponding number of aircraft within the sector would be the historical sector capacity for that flow pattern. Sector behavior curves as shown schematically in Figure 11 have been detected through observations of sector throughput. These curves support the theory that sector behavior is a function of sector count and traffic flow pattern. In these observations, the performance index was the average distance traversed by flights on the heaviest flow within the sector, though many other measures are possible. Figure 11 also shows that for more complex traffic flow patterns ($P_3 > P_1$), fewer flights

need to be present for congestion behavior to be observed, and thus the sector capacity is smaller ($C_3 < C_1$). Systematic sector studies, using this methodology, are underway.

After the primary set of traffic flow patterns and the corresponding sector capacities are identified for each sector, future sector capacity for a given LAT can be predicted through pattern recognition, by comparing the predicted traffic flow pattern with the primary set of traffic flow patterns.

Weather impact on sector capacity

When severe weather impacts the sector, probabilistic weather forecasts can be used to determine the probability of blockage for each flow within the sector, and the capacity of blocked flows is subtracted from the overall sector capacity. In this way, sector capacity in the presence of severe weather can be predicted probabilistically. Figure 12 shows an example of sector capacity distribution given the predicted traffic flow pattern and probabilistic severe weather forecast.

The key to building the mapping function between the probabilistic weather forecast and the sector capacity distribution is to identify the flow blockage distribution. Flow blockage distribution will be calculated based on probability of pilot rejection, given the location of the flow and the probabilistic weather distribution around the flow.

Numerous research areas remain open in this work. To identify the basic set of traffic flow patterns for each sector, more data analysis, field observations and interviews are necessary. More study is necessary to find the best performance index for identifying sector behaviors as shown in Figure 11. Finally, more research is needed to implement and validate the weather impact method described in Figure 12.

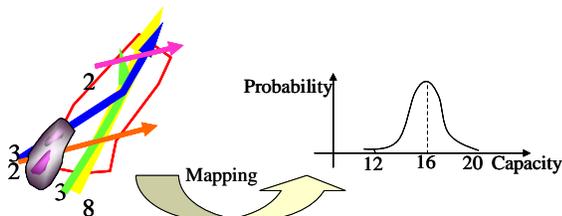


Figure 12. Example of Sector Capacity Distribution under Weather Impact

Probabilistic Decision Making

The idea of probabilistic decision making is deceptively elegant. Based on an observable utility metric, choose actions that maximize that utility. Other researchers have addressed parts of this

problem. Recent work includes methods for selecting routes to meet several congestion and equity constraints [10], and route selection while considering weather uncertainty [11]. To date, a technique that allows incremental decision-making in the presence of both demand and capacity prediction uncertainty has not been successfully demonstrated.

There are several challenges present in developing an algorithm which satisfies the goals set forth in the Introduction. First, the utility metric must reflect real policy goals and operational limitations, leading to acceptable solutions. Second, the algorithm must develop solutions that can be iteratively modified as conditions change; it is unacceptable, for example, to change the assigned route of a single flight several times. Third, the algorithm must identify both immediate actions and likely future actions, so that NAS users can proactively manage their own schedule risk.

Defining a Utility Metric

One approach is to base the utility metric explicitly on safety, by using a function of the probabilities that sectors will be congested, where congestion is defined as demand exceeding capacity. This approach is too simple, since it provides no discrimination between solutions which may be different in other important ways, such as imposed delay.

Thus, at a minimum, a mathematical representation is needed for the tradeoff between congestion risk and the impact on NAS users of reducing that risk. The relative value given to congestion risk vs. impact is an explicit statement of traffic management policy. A utility metric that penalizes congestion risk highly will produce actions that move many flights to achieve low congestion risk. A utility metric that penalizes delay highly will tend to leave flights alone and tolerate high congestion risk. The metric may also include values for other resolution features. For example, the distribution of delay across NAS user types (e.g., airlines, military, general aviation) can be explicitly captured in the utility metric, thus producing solutions that are inherently equitable, provided that equity can be suitably defined.

Since the demand and capacity are stochastic, the utility function will also be stochastic. Thus, inherent uncertainty in the predictions will limit the potentially realizable utility at a particular decision point. Decision-making performance can always be improved by reducing the inherent prediction uncertainty.

Solution Methods

Once risk and utility functions are defined, algorithms for developing congestion resolutions can be explored. While classical decision theory provides several alternative approaches to such a problem, no previous researchers have usefully applied it to a problem of this type. Our initial approach is to divide the problem into two parts. The first part involves deciding when to take action, based solely on congestion risk. One method would be to set a simple threshold of congestion risk; when the probability of a sector being congested exceeds a preset value, action is required. Similarly, a lower threshold could be used to identify regions where future action is likely, but not yet required. More complex techniques include partially observable Markov decision processes [12], which are designed explicitly for handling uncertainties in incremental decision making.

Once the decision to act is made, a resolution that maximizes (or nearly so) the utility function must be developed. Optimization techniques for problems of this type exist, such as genetic algorithms (GA) or linear programming variants, but can be computationally expensive. An initial GA prototype has been implemented, and while initial tests show promise, results are not yet available.

Roles of Humans and Automation

A crucial question that must be addressed in the design of any decision support system is the allocation of responsibilities between humans and computer-based automation. In an effort to assess the effects of human/machine interaction on task performance, Sheridan [13] identified a number of different levels of interaction between a human user and a machine depending upon the nature of the task at hand. He provided a breakdown of the *degrees of automation* applied to a system where human intervention is steadily decreased across a scale of automation. At the lowest level of the scale, the human is acting entirely unassisted by the computer, while at the highest level, the computer takes over complete responsibility for the task and ignores any human input. Each successive level in the scale calls for the automation taking an increasing role in choosing and implementing a solution option, with a corresponding decreasing role for the human in choosing and approving that option.

For some tasks it is acceptable to allow the automation component to go the highest level, while for others a much lower limit is preferred. The general tendency in decision support design has been

to automate what is easiest and leave tasks with more cognitive complexity to the human.

The concept for probabilistic congestion management involves a higher level of automation than any existing TFM decision support system. As discussed earlier in this paper, probabilistic estimation and decision-making are mathematically complex. It is impractical for traffic managers to make detailed, probabilistic decisions in real time without automation support. In the concept presented here, the traffic manager has controls to affect the general form and magnitude of the recommended congestion solution, but the automation is responsible for determining the detailed resolution actions. This raises a host of human factors issues, such as the ability of human decision-makers to maintain situation awareness. Also, can humans maintain a high level of trust in a complex, probabilistic system, when the decisions made may not be entirely consistent with their intuition and experience? These issues will eventually have to be addressed through human-in-the-loop evaluation.

Conclusion

An operational concept and supporting research for probabilistic congestion management in en route airspace has been presented. Probabilistic decision-making has the potential for great improvements in system performance, since by explicitly accounting for prediction uncertainties, ineffective and unnecessary flow actions can be reduced. Probabilistic decision-making leads to incremental corrective actions which maintain an acceptable risk of congestion, rather than large scale flow controls which may affect more flights than needed.

To implement such a system, traffic demand and sector capacity prediction uncertainties must be quantified. Two models for demand prediction uncertainty were discussed, and a proposed model for sector capacity prediction uncertainty was outlined. A candidate probabilistic congestion alerting display was also introduced. Many open issues remain, especially in the human factors area. Research continues, with the eventual goal of developing a practical probabilistic congestion management decision-support system.

References

- [1] Wanke, C., et. al., "Measuring Uncertainty in Airspace Demand Predictions for Traffic Flow Management Applications," *2003 AIAA Guidance*,

Navigation, and Control Conference and Exhibit, Austin, TX, 11-14 August 2003.

[2] *Enhanced Traffic Management System (ETMS) Reference Manual, Version 7.5*, Report No. VNTSC-DTS56-TMS-004, Volpe National Transportation Systems Center, U.S. Department of Transportation, November 2002.

[3] Megenhardt, D. L., et al., "NCWF-2 Probabilistic Forecasts," *11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA, 4-8 October 2004.

[4] Wolfson, M. M., et al., "Tactical 0-2 Hour Convective Weather Forecasts For FAA," *11th Conference on Aviation, Range, and Aerospace Meteorology*, Hyannis, MA, 4-8 October 2004.

[5] Wanke, C. R., et al., "Modeling Traffic Prediction Uncertainty for Traffic Management Decision Support," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Providence, RI, 16-19 August 2004.

[6] Masalonis, A., et. al., "Using Probabilistic Demand Predictions for Traffic Flow Management Decision Support," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Providence, RI, 16-19 August 2004.

[7] Schmidt, D. K., "On Modeling ATC Work Load and Sector Capacity," *Journal of Aircraft*, Vol. 13, No. 7, pp. 531-537, 1975.

[8] EUROCONTROL, 2003, *Pessimistic Sector Capacity Estimation*, EUROCONTROL Experimental Centre Note 21/03, EUROCONTROL Experimental Centre, Brétigny-Sur-Orge, France.

[9] Histon, J., and R. J. Hansman, "The Impact of Cognitive Complexity Limits on ATC," Technical Interchange Meeting, McLean, VA, 2004.

[10] Sherali, H. D., et. al. "An Airspace Planning Model for Selecting Flight-plans Under Workload, Safety, and Equity Considerations," *Transportation Science*, Vol. 36, Issue 4, pp. 378-397, Nov. 2002.

[11] Nilim, A, et. al., "Robust Dynamic Routing of Aircraft Under Uncertainty," *21st Digital Avionics Systems Conference*, Berkeley, CA, 2002.

[12] Kaelbling, L. P., et al., "Planning and acting in partially observable stochastic domains." *Artificial Intelligence*, Volume 101, pp. 99-134, 1998.

[13] Sheridan, T. B., *Telerobotics, Automation, and Human Supervisory Control*, The MIT Press, 1992.

Notice

The contents of this material reflect the views of the authors and/or Director of the Center for Advanced Aviation System Development. Neither the FAA nor the Department of Transportation makes any warranty or guarantee, or promise, expressed or implied, concerning the content or accuracy of the views expressed herein.

Keywords

Traffic flow management, decision support, air traffic prediction, probabilistic decision-making, en route congestion

Author Biographies

Craig R. Wanke is a Senior Principal Engineer at MITRE, specializing in simulation and modeling. He has worked for the past 15 years on decision support systems for cockpit and air traffic management applications. He holds S.B, S.M., and Ph.D. degrees in aeronautical engineering from the Massachusetts Institute of Technology.

Lixia Song is a Senior Modeling and Simulation Engineer at MITRE. Her research includes decision support systems analysis and integration, estimation and control, optimization. She holds S.B and S.M. in Solid Mechanics from Hunan University, S.M. in Aerospace Engineering from University of Minnesota, and Ph.D. degree in aeronautical engineering from the Massachusetts Institute of Technology.

Sandeep Mulgund is a Lead Information Systems Engineer at MITRE. His research activities focus on modeling, simulation, visualization, and high performance computing. He holds a Ph.D. in Mechanical and Aerospace Engineering from Princeton University.

Dan Greenbaum is a Lead Software Systems Engineer at MITRE. His work includes developing software to model air traffic and related human decision making. He holds an M.S. in Computer Science from the George Washington University (pending), an M.E. in Systems Engineering from the University of Virginia, a J.D. from Cornell Law School, and a B.A. from the University of Pennsylvania.

Stephen Zobell is a Lead Engineer at MITRE, specializing in software systems and simulations. He has worked in air traffic management decision support system development for 20 years. He holds an M.S. in Computer Science from the George Washington University.