PROBABILISTIC AIRSPACE CONGESTION MANAGEMENT

Stephen Zobell,[1] Craig Wanke,[2] and Lixia Song[3]
The MITRE Corporation, McLean, Virginia

1 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.

An example of a weather-related congestion situation is shown in Figure 1. A moving line of thunderstorms is reducing airspace capacity in the Kansas City Air Route Traffic Control Center (denoted "ZKC"). Sectors are labeled by facility and sector number (e.g., ZKC32 means ZKC center, sector 32) and altitude range in flight levels (e.g., 240-319 indicates 24,000 to 31,900 feet above ground level.) Due to uncertainties in predicting the storm motion and intensity, the future capacity of the impacted sectors is uncertain. In addition, predicted traffic demand is subject to several uncertainties, which will be discussed below. Several flights in this example have filed flight plans to skirt the southern edge of the storm, and these sectors are now heavily congested, as indicated by the red and yellow alert overlays.

Such congestion problems typically involve hundreds of flights, and affect multiple en route ATC facilities and airspace users (e.g., airlines). It is difficult for human decision-makers to develop effective, coordinated solutions to such problems in real time, even if predictions are good. Thus, predefined large-scale strategies, such as "National Playbook" routes [FAA 2005], have become standard means of addressing congestion.

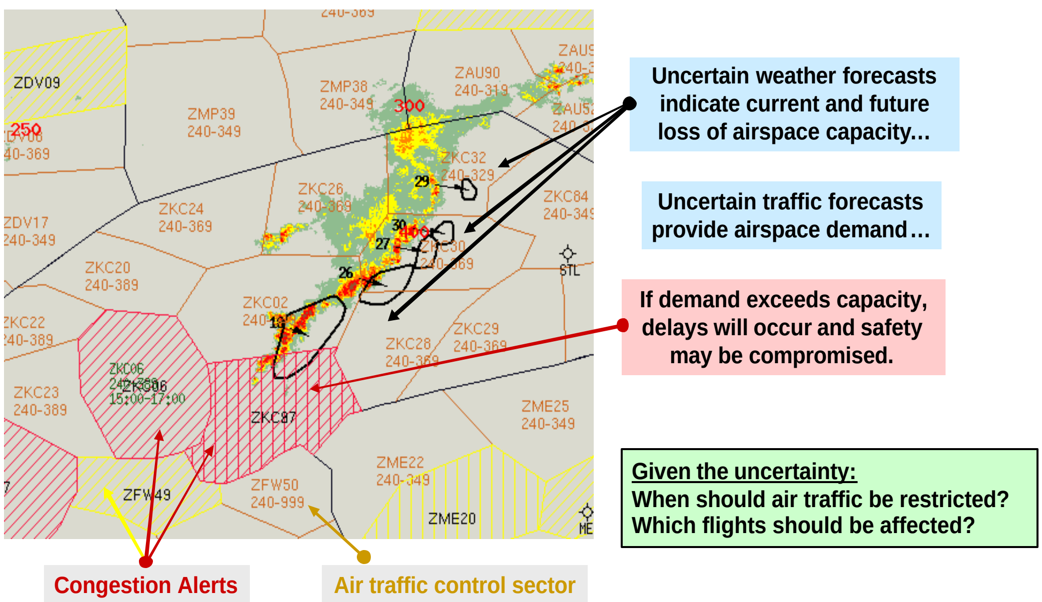

Figure 1. Weather-Related Airspace Congestion

[1] Lead Engineer, Mailstop N450.
[2] Senior Principal Engineer, Mailstop N450.
[3] Senior Engineer, Mailstop N450.

*Stephen M. Zobell, The MITRE Corporation's Center for Advanced Aviation System Development, 7515 Colshire Drive, McLean VA, 22102-7508; e-mail: szobell@mitre.org

The combination of prediction uncertainty and large-scale solution strategies leads to highly conservative decision-making. Such decisions may produce unnecessary delays, and may be taken at inappropriate times based on the actual accuracy of prediction data. Automation support is needed to help identify congestion problems and appropriate, efficient resolution strategies in the presence of prediction uncertainty.

2 BACKGROUND: CONGESTION PREDICTION UNCERTAINTY

2.1 Present-Day En Route Congestion Alerting

In the NAS, the Enhanced Traffic Management System (ETMS) [Volpe 2002] provides demand predictions for most NAS sectors in 15-minute bins, for prediction look-ahead times (LAT) of several hours. This information is available for particular sectors by user request, or collected on a *Center Monitor* (CM) display as illustrated in Figure 2. The CM is a user-configurable display showing alerts for some or all of the sectors in a single Air Route Traffic Control Center (ARTCC).

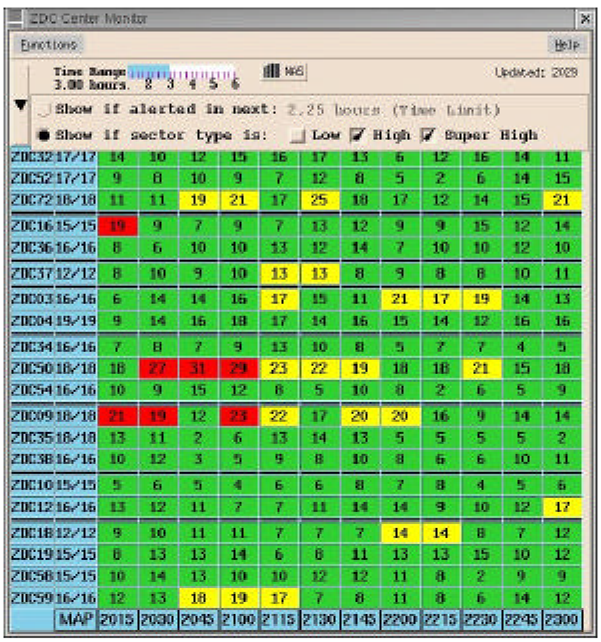

Figure 2. ETMS Sector Count Monitor Display

Each cell in the CM matrix represents a 15-minute period, and the number in the cell represents the maximum predicted traffic count for any single minute within that 15-minute span. This value is often referred to as maximum instantaneous aircraft count (IAC) or simply "peak count" for the interval. The horizontal axis indicates increasing LAT (corresponding to 2015 to 2300 Coordinated Universal Time (UTC), in this case). Each matrix row represents predictions for a single sector (e.g. ZDC50). Next to the sector name are two sector alert thresholds (e.g. "18/18"), although currently, only one is used. This threshold is called the Monitor/Alert Parameter (MAP) and is compared to the peak count to determine whether a sector should be alerted. When the peak count is predicted to exceed the MAP for a sector, the corresponding box is colored yellow or red. Red alerts indicate that, of the aircraft involved in the peak count, enough are already airborne to exceed the MAP even if pre-departure flights are not counted. Otherwise, the alert will be yellow.

The MAP value is set to represent a traffic level high enough to be of concern to the traffic manager. The nominal value can be manually changed to reflect the impact of weather or other adverse conditions, though it can only have a single value; it cannot have different values at different LAT. It is not strictly accurate to refer to the MAP as a sector *capacity*, since there are many factors involved in sector workload beyond the number of aircraft present [Kopardekar, 2002][Masalonis, 2003]. However, it is an easily-understood abstraction of workload for alerting purposes. Thus, a simple definition of congestion is when demand exceeds the MAP.

2.2 Prediction Uncertainty

These sector traffic predictions are key TFM decision aids. Traffic managers use the alerts to identify areas of potential en route congestion, and by studying the flights and traffic flows involved, to identify candidate solutions such as reroute initiatives or MIT restrictions. Also, proposed TFM decision support systems [Wanke, 2004] make direct use of these predictions when, for example, predicting the impact of a proposed reroute initiative. However, the usefulness of these predictions is a function of their accuracy. At the long prediction timeframes associated with strategic TFM decision-making, the predictions may not be very accurate.

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 magnitude and characteristics of these uncertainties have been measured for several months of NAS operations [Wanke, 2003], and a computational model has been developed [Wanke, 2005].

While traffic managers know that sector demand predictions are uncertain, they have very little information to use in quantifying that uncertainty and taking account for it when making decisions. ETMS sector load predictions include a crude estimate of uncertainty, in that alerts are differentiated into "red" and "yellow" based on whether or not all the aircraft involved are airborne. This is based on the assumption that departure time uncertainty is the largest source of uncertainty in the predictions. While this is useful for prioritizing traffic situations, in that the traffic manager would be justified in looking at red alerts before yellow alerts, it says little about the actual magnitude of uncertainty. For example, does a yellow alert of 3 aircraft over the MAP mean that there is an 80 percent chance that demand will exceed the MAP, or a 20 percent chance? Clearly, the answer to this question should influence the traffic management decision.

Airspace capacity itself is harder to quantify. As noted earlier, the MAP is not intended to represent capacity directly. 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 the presence of convective weather.

2.3 Applying Uncertainty to Decision-Making

One way to factor in prediction uncertainty is to present probabilities directly or indirectly on traffic management decision support displays, relying on the skill of the traffic manager (and some procedural guidance) to use such information appropriately. Research is underway to understand the human factors issues in this area. Masalonis, et al have developed candidate visualization methods for probabilistic sector demand information, in research that is directly linked to the work presented here. A simple application is to replace the current point estimates (as used in Figure 2) with an estimate of known statistical properties, such as the median of the probability distribution. This would automatically compensate for biases in the predictions, without requiring the traffic manager to absorb any new information. For example, ETMS predictions at longer LAT are more frequently too low than too high, since they cannot include flights that have not yet filed plans. An example of such a display will be shown later.

Probabilistic predictions can also be used by decision support automation. Given detailed knowledge of demand and capacity prediction error distributions, standard decision analysis techniques can be applied to improve decision-making. This will be discussed further below.

3 Overview of Probabilistic Congestion Management

A notional probabilistic forecast of congestion is shown in Figure 3. If the uncertainties in traffic demand and capacity predictions are known and quantifiable, then a probability of congestion can be calculated from the demand and capacity uncertainty 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.

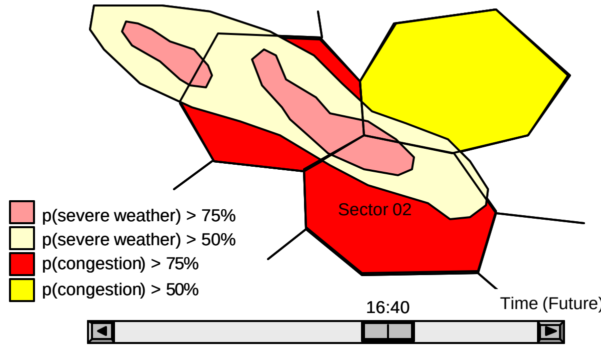

Figure 3. Probabilistic Congestion Forecast.

In addition, the probability of severe weather coverage is shown, based on a probabilistic severe weather forecast. Such products are now becoming available [Megenhardt, 2004] [Wolfson, 2004].

Figure 4 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. The capacity is well-known at short LAT, 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.

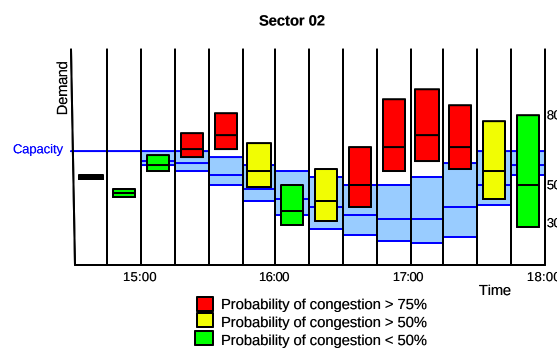

Figure 4. Detailed Congestion Predictions for a Singel Sector

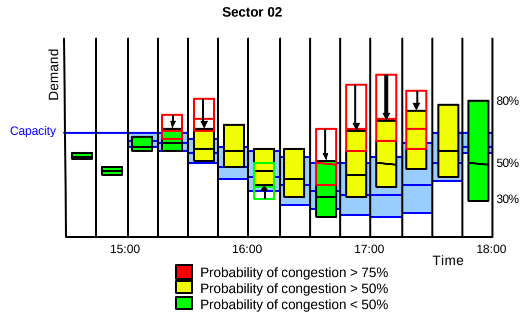

Figure 5. Reducing Congestion risk to an Acceptable Level

Given this representation, the congestion management goal can be expressed as a target level of probability, or "congestion risk." The traffic manager may, for example, decide that action should be taken to reduce the congestion risk below 50 percent for the next three hours. This is shown in Figure 5. The resolution problem, then, is to determine a set of actions that achieves this risk reduction while minimizing impact on airspace users. Ideally, this would be done by flight-specific adjustments, such as individual reroutes or ground delays, rather than actions on entire traffic flows.

In this concept, not all possible congestion is resolved at once. Congestion probabilities are continuously monitored, and incremental actions are taken to keep congestion risk at a tolerable level, affecting only a few flights at a time. This helps to avoid unnecessary or premature actions.

Another advantage of the probabilistic representation is that *developing* problem areas (those approaching the risk threshold) can be identified easily, and airspace users can be alerted before actions are taken. This allows users to proactively re-plan critical flights, or to provide preferred alternative routes should action be required.

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. This paper describes research into these topics, including a scenario in which a prototype system has been used to develop a congestion solution.

4 INITIAL PROBABALISTIC CONGESTION MANAGEMENT: AN EXAMPLE

The operational concept described here is significantly different from today's congestion management procedures. Therefore, a two-phase approach has been defined. In the first phase, traffic demand is computed probabilistically, but the sector alerting criterion (MAP) remains as defined today. The traffic manager identifies when congestion must be addressed, chooses which options are available to resolve it, and initiates the computation which generates the resolution maneuvers. This concept has been implemented in a real-time decision support prototype for testing and evaluation. An example of congestion management using this concept is given below, using illustrations from the prototype. The second phase capabilities will be described in Section 5.

4.1 Congestion Problem Identification

This example centers on a relatively small-scale congestion problem, illustrated in Figure 6. Slow-moving, intense thunderstorms with high cloud tops are impacting a high-altitude sector, reducing its effective capacity. Selected routes through this sector are highlighted in blue. Two things are apparent from the situation geometry. First, the traffic pattern includes several jet airways (J35, J80, J105/J181) and several crossing points. Second, at least one of these airways (J80) is obstructed by the storms. Based on this situation, the traffic manager decides to reduce the MAP value for this sector, normally 16, to 10.* The traffic manager then observes the predicted traffic situation on the Center Monitor (Figure 7).

The CM in Figure 7 is a probabilistic version of that shown in Figure 2. In this version, the median predicted peak count for each sector is shown, thus compensating for prediction biases. Also, the red and yellow alert colors are based directly on probabilities of congestion. On the probabilistic CM, a red alert indicates a higher than 75 percent probability that the actual peak traffic demand will exceed the MAP. A yellow alert indicates a 50 percent to 75 percent probability. These probabilities are user-selectable parameters.

In Figure 7, it is clear that congestion is highly likely in sector 84, and for a period of 90 minutes (1415Z to 1545Z), the median predicted peak values exceed the adjusted MAP by between 3 and 8 aircraft. Based on this prediction, in combination with other factors (such as the operational importance of sector 84's location) the traffic managers decide to initiate congestion management planning.

4.2 Resolution Strategy Definition

The traffic managers begin by defining sector 84 as a Congestion Resolution Area (CRA). This is the area in which flights are eligible for congestion management actions. The automation identifies 97 flights which will enter the CRA during the 1415-1600Z time period. Next, the traffic managers define a surrounding region as the Congestion Management Area (CMA). As the automation develops a resolution strategy for the CRA, it will also avoid producing any new congestion in the CMA. Also, there is an active military operations area (Red Hills) within the CMA, and the traffic managers identify this as an area to be completely avoided when resolution maneuvers are developed. The result of this process is shown in Figure 8.

---

* This is a contrived example, and does not reflect any real operational event. Nor does it represent the judgment of a qualified traffic manager.

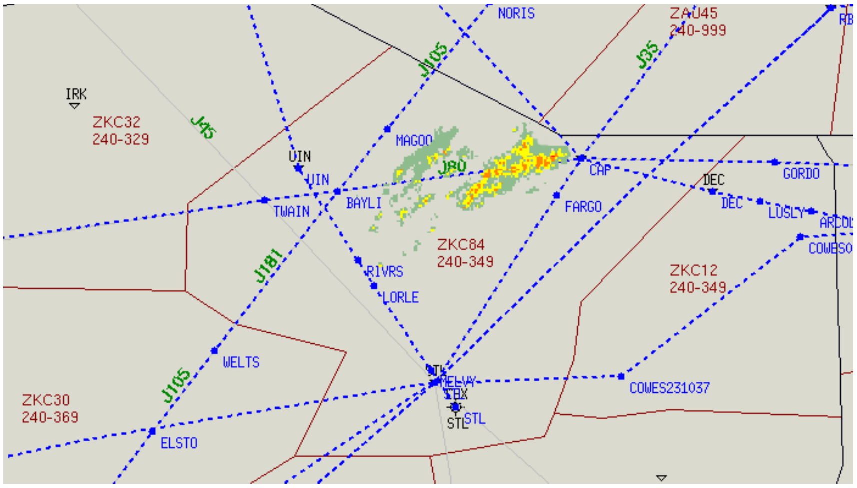

Figure 6. Flow – and Weather Induced Complexity in ZKC84

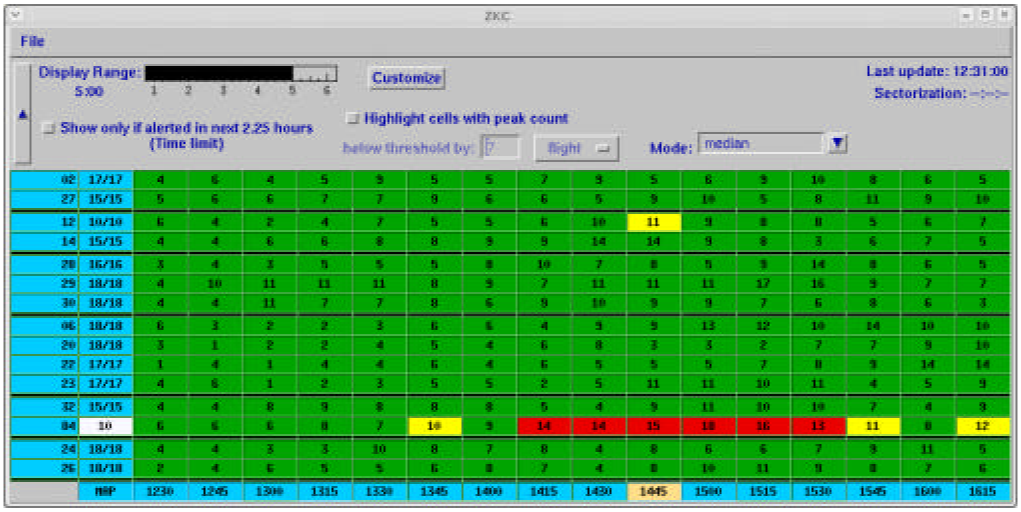

Figure 7. Predicted Congestion in Kansas City ARTCC (ZKC) Sector 84

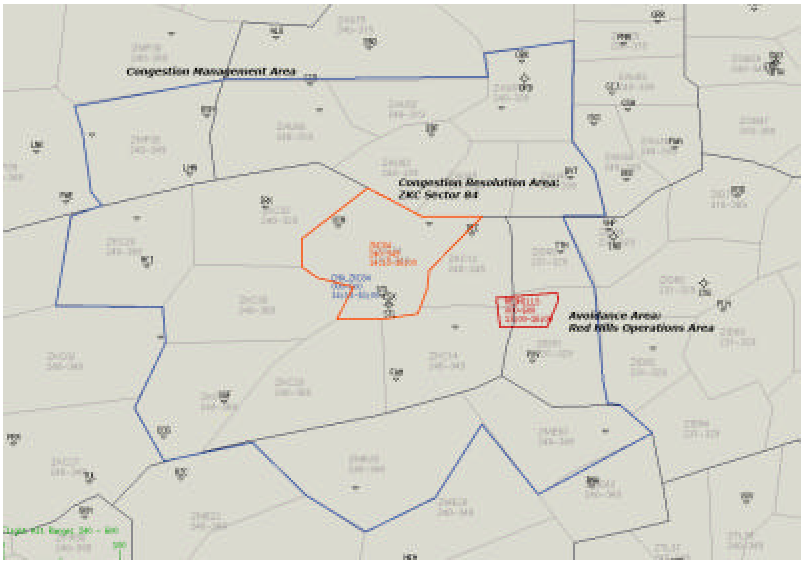

Figure 8. Congestion Problem Definition: Management Area, Resolution Area, and Avoidance Area

At this stage, a planning advisory is disseminated to NAS facilities and users, noting that a traffic management initiative is likely to be implemented. It includes the areas and timeframe of likely action. This enables NAS users to submit preferences as to how they would like their flights handled, should an initiative be required. In this example, it is assumed that some users would submit desirable alternate routes. Users may also simply file new flight plans to avoid the potentially congested area.

For planning the strategy details, the automation presents the list of flights identified as entering the CRA during the time period of interest (Figure 9). Since Sector 84 overlays Lambert-St. Louis International Airport (STL), the traffic managers modify this list by exempting 10 flights departing or arriving at STL. The result is a list of 4 active (airborne) and 83 inactive (pre-departure) candidate flights for resolution maneuvers. This list includes basic plan data for each flight, as well as preferences submitted by the flight operator.

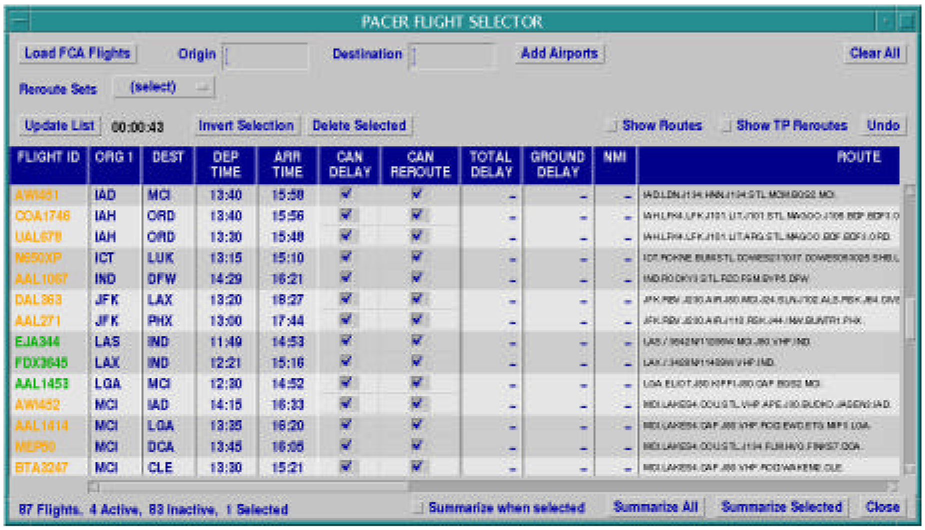

Figure 9. Non-STL Flights Predicted to Enter ZKC84 During the Congested Period

Once the flight list is established, the details of the resolution strategy can be developed. Figure 10 shows the probabilistic congestion resolution planning window from the prototype. The current prototype generates resolution maneuvers comprised of ground delays and reroutes, though other strategies are possible. Parameters governing the possible values of ground delay are shown at left; in this case, the traffic manager has established a maximum value of 20 minutes of delay per flight.

Reroutes can be generated from a variety of sources. The traffic manager can manually define corridors around the weather, or can ask the automation to choose routes from databases of predefined or historically-flown routes. The prototype currently uses the FAA Coded Departure Route (CDR) database as a source of predefined routes; an operational capability could have these and many other route options available. These route databases are keyed by origin-destination pair. In this case, the traffic manager has selected the predefined route database, and that only inactive flights will be rerouted. Also, any user-submitted alternate routes are automatically considered in the resolution process.

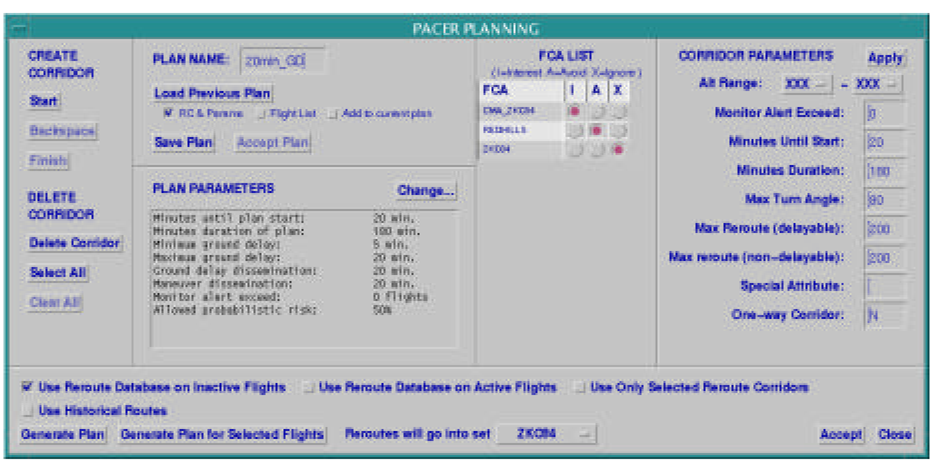

Figure 10. Congestion Resolution Planning Interface

The automated resolution algorithm attempts to resolve congestion by proposing flight-specific ground delays and reroutes for flights. The optimal solution that both solves the congestion and assigns the very minimum delays to flights is very difficult to calculate in a reasonable amount of time. Instead, a faster approach is used that prioritizes flights based on earliest time of arrival to the congestion area. The automation first removes all flights that are eligible to maneuver from the traffic and then adds flights back one at a time in this priority order. As each flight is added back, the automation determines if the flight will increase the sector congestion probability beyond the threshold value. If the flight does not cause congestion problems, it is added back without changes to the flight's route or schedule. If the flight causes congestion problems, then the algorithm searches the reroute or ground delay options for a maneuver that avoids congestion and also causes the smallest arrival delay for the flight. The resolution results are proposed flight-specific reroutes and ground delays for a select set of flights.

This is an initial algorithm, and the results are highly dependent on the priority order in which flights are considered for maneuvering. Ordering by predicted congestion area entry time is one of many possible ways to prioritize flights, and further research is needed to determine the "best" priority scheme for NAS operations. Also, resolutions are developed without explicit consideration for future adjustments. It may be desirable to choose maneuvers such that some level of future flexibility is maintained. For example, avoid delaying flights that do not depart until considerably later, so that they can be delayed in future if the situation worsens.

4.3 Resolution Impact Assessment

Once the calculation is complete, the traffic manager can examine the structure and predicted impact of the resolution. A summary of the plan is shown in Figure 11. The automation has assigned 10 reroutes, averaging an increased flight distance of 88 nautical miles (nmi), and 41 ground delays, averaging 11 minutes per flight. The total number of delayed flights is 50, since one of the reroutes actually had a shorter flying time than the originally-filed plan for that flight. Overall, the resolution plan would incur 582 minutes of delay. This is significant, but probably considerably less than if a large scale reroute plan, miles-in-trail restrictions, or ground delay programs were run to solve this weather problem.

* The center monitor in Figure 13 is configured to show only the high-altitude sectors (i.e. not low- or super-high-altitude sectors) in ZKC, but the resolution did not create new alerts in any other sectors either.

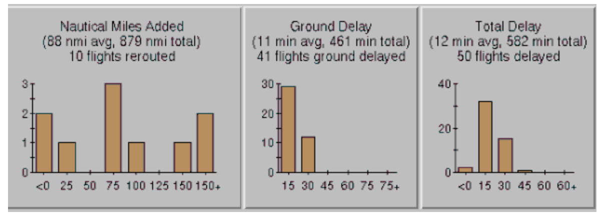

Figure 11. Resolution Plan Summary

Details of the resolution plan are also available (Figure 12). The assigned ground delay and/or reroute for each flight is shown, and if rerouted, the calculated increase or decrease in flight distance and flying time is shown. The portion of the list shown here includes two ground-delayed flights and three rerouted flights. Note that if a flight is rerouted, and has provided a preferred alternate route, that route will be chosen.

Finally, the traffic managers can determine whether the resolution plan is acceptable for managing the congestion. The automation predicts the change in sector traffic loads due to the resolution plan and presents it on a modified Center Monitor display (Figure 13). Given the resolution constraints – a maximum ground delay of 20 minutes, and routes from the Coded Departure Route database – the system was not able to reduce all congestion probabilities below 50 percent. However, it left median counts of only one or two over the MAP, which is likely to be an acceptable solution, and did not create new congestion elsewhere.* Note that heavy blue outlines indicate where the median peak count increased due to the resolution, and light blue outlines indicate a decrease.

The reroutes generated by this plan are chosen to be as short (in flying time) as possible, given the congestion and avoidance constraints. An example is shown in Figure 14. This flight was assigned a CDR which took it clear of both ZKC84 and the Red Hills operations area, with a relatively small increase in flying time.

At this point, the plan may be deemed acceptable, or the parameters can be changed and a new plan produced. Recall that a maximum of 20 minutes of ground delay was chosen for this plan. If this number were to be changed to 30, then the results are considerably different. As shown in Figure 15, this plan produces fewer and shorter reroutes, with correspondingly less airborne delay (28 minutes), but considerably more total delay (758 minutes). If the decision makers would like to avoid reroutes and/or airborne delay, then this may be a better plan.

Note, however, that the congestion results of a ground-delay-heavy strategy are quite different (Figure 16). Using larger ground delays, the congestion is completely resolved in the timeframe of interest (1415-1545), but since more flights are being pushed later, there is a bubble of high-probability congestion at 1545-1615. This may be acceptable, since it is three hours in the future and can be resolved later, but is certainly not as satisfactory as the previous solution.

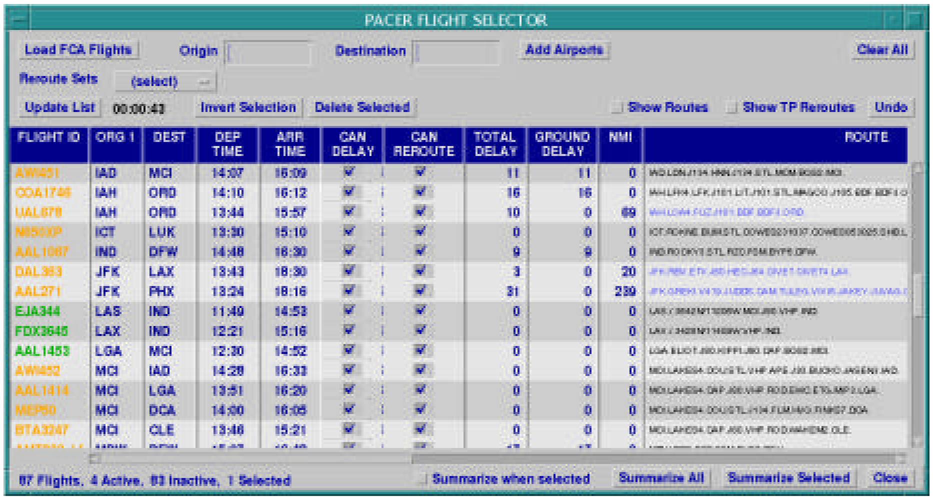

Figure 12. Resolution Plan Details

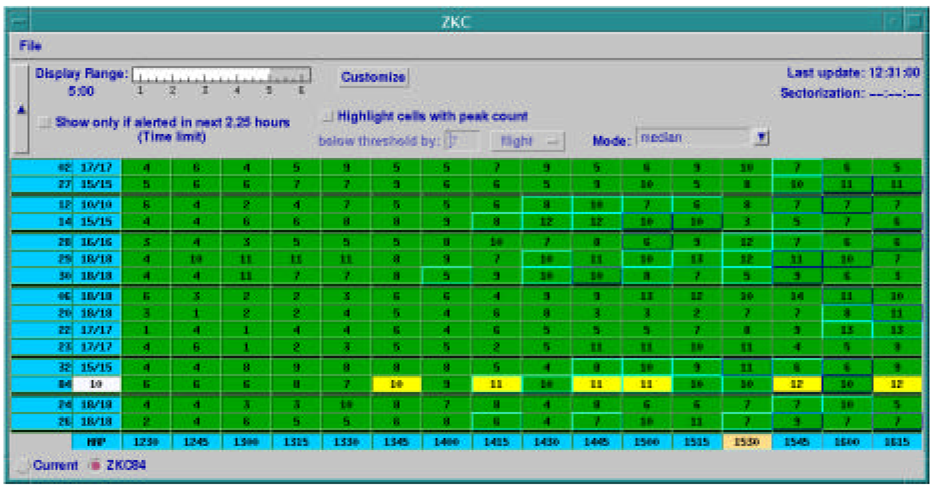

Figure 13. Resolution Plan Impact on Congestion

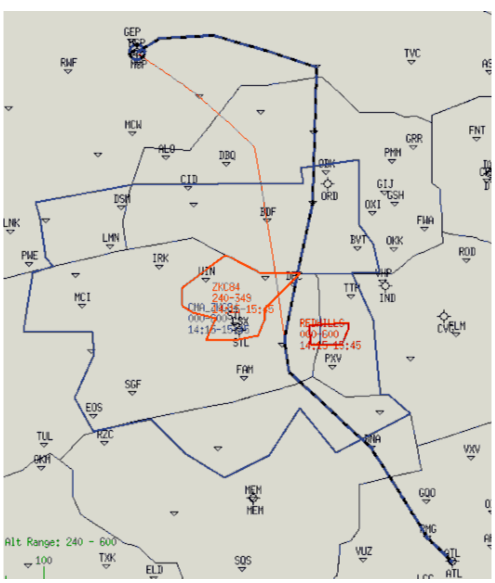

Figure 14. Congestion Avoidance Route

Plan adjustments such as this take less than a minute to construct and analyze using the prototype, and thus might be done collaboratively (e.g., between FAA facilities, or involving NAS users) in an operational setting.

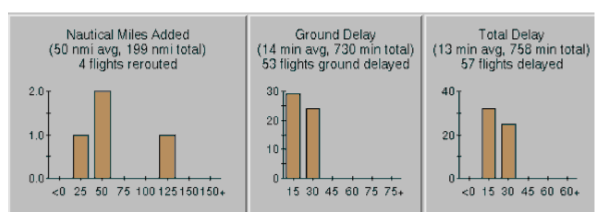

Figure 15. Resolution Plan for Maximum Ground Delay of 30 Minutes

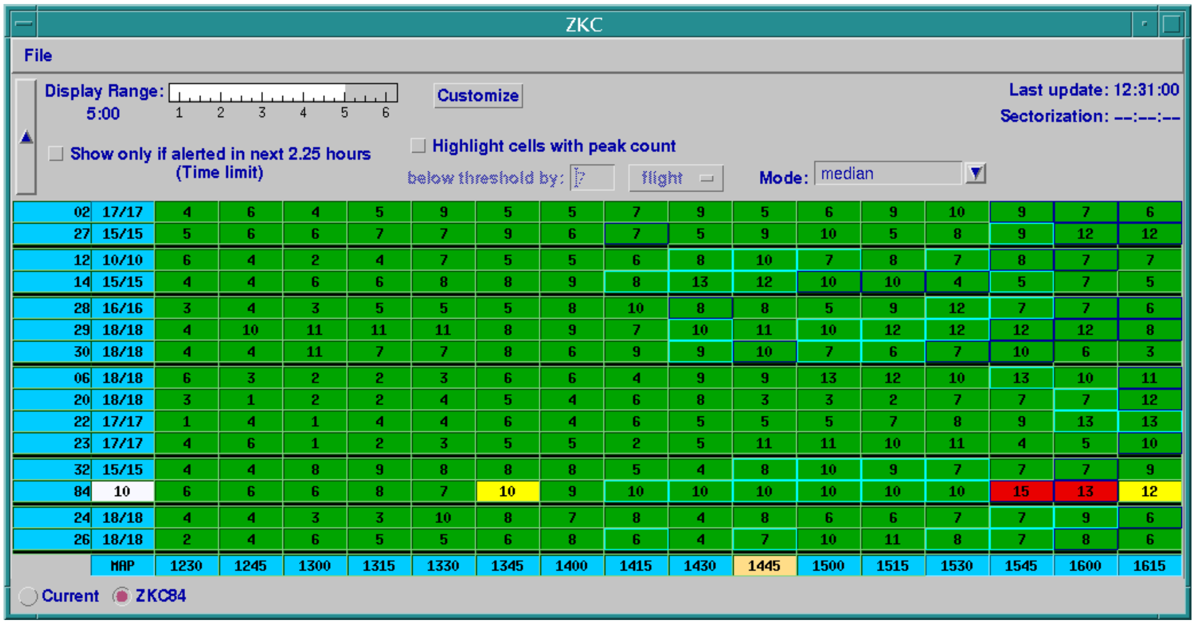

Figure 16. Congestion Impact of the Revised Plan

4.4 Coordination, Execution, and Monitoring

Once the traffic managers working on the plan decide that it is acceptable, it is coordinated with other en route facilities and the Air Traffic Control System Command Center (ATCSCC). Further adjustments are made, if needed to reach consensus. At this point an advisory is issued, and the plan is executed. Ground delays would be executed in a similar fashion to how Ground Delay Programs are executed today, by issuing Estimated Departure Clearance Times (EDCTs) for involved flights. Reroutes would also be disseminated automatically to the appropriate airspace users and FAA facilities. Such a system does not yet exist, but is planned for deployment in the next few years.

As time passes, traffic managers throughout the NAS monitor the execution and impact of the plan, and tactical adjustments are made if needed.

5 Advanced Probabilistic Congestion Management

The process illustrated above represents a major step forward in congestion management, and is expected to be significantly more effective than current procedures in resolving congestion without imposing excess delay. A benefits analysis is underway to estimate the size of this improvement. However, it imposes significant workload on traffic managers, and relies on manual problem recognition and coordination. The resolution plans generated, while expected to be smaller than "National Playbook" initiatives, will still involve large numbers of flights. The second phase of the probabilistic congestion management concept addresses these and other issues.

5.1 Probabilistic Sector Capacity Prediction

As previously discussed, 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. This would replace the manually-adjusted MAP that was used in the resolution example described earlier.

Considering that the clustered traffic (flow) properties are more predictable and perturbation-resistant than individual flight characteristics, an approach is being developed to predict the sector capacity as a function of primary traffic flow pattern. NAS sectors typically exhibit a small set of common traffic flow patterns, and different patterns represent different levels of traffic complexity. In higher complexity conditions, it takes fewer flights to generate high workload for the controller team, and thus the sector capacity is lower.

Figure 17 shows the relationship between traffic flow patterns and sector capacity. Given a traffic flow pattern (for example, P1), 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 17, the given sector has three normal traffic flow patterns (P1, P2, and P3), representing three different levels of traffic complexity. When the controller workload reaches the threshold, the sector capacity is reached (C1 for P1). As shown in the figure, given the threshold of controller workload, the capacity is different (C1, C2 and C3) for different traffic flow patterns (P1, P2 and P3). Since P3 is the most complex traffic flow pattern, the sector has least capacity when the traffic has that pattern.

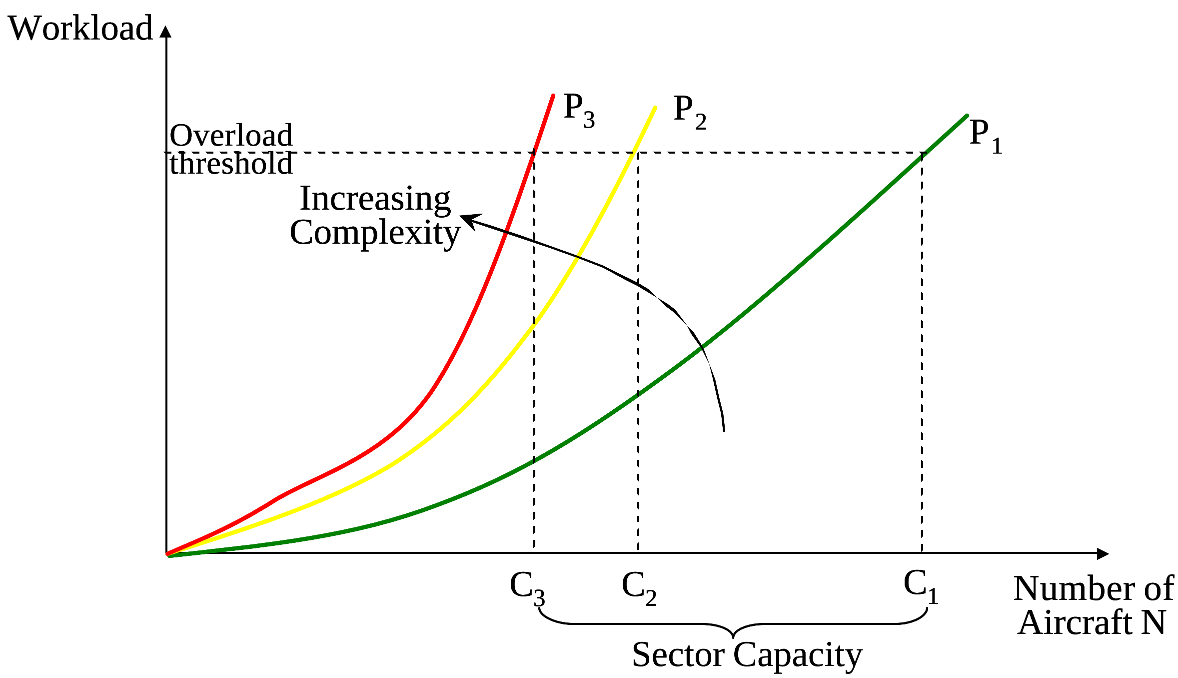

Figure 17. Relationship between Workload, Sector Capacity, Complexity, and Traffic Flow Pattern

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 18 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).

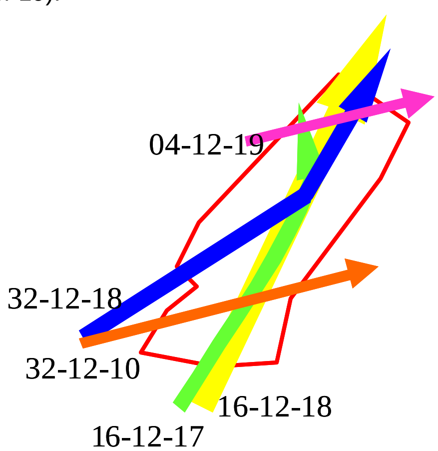

Figure 18. Sector Flow Pattern Example

Two activities are required to characterize 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 [Schmidt, 1975]. 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) [Eurocontrol, 2003]. 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, 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 19 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

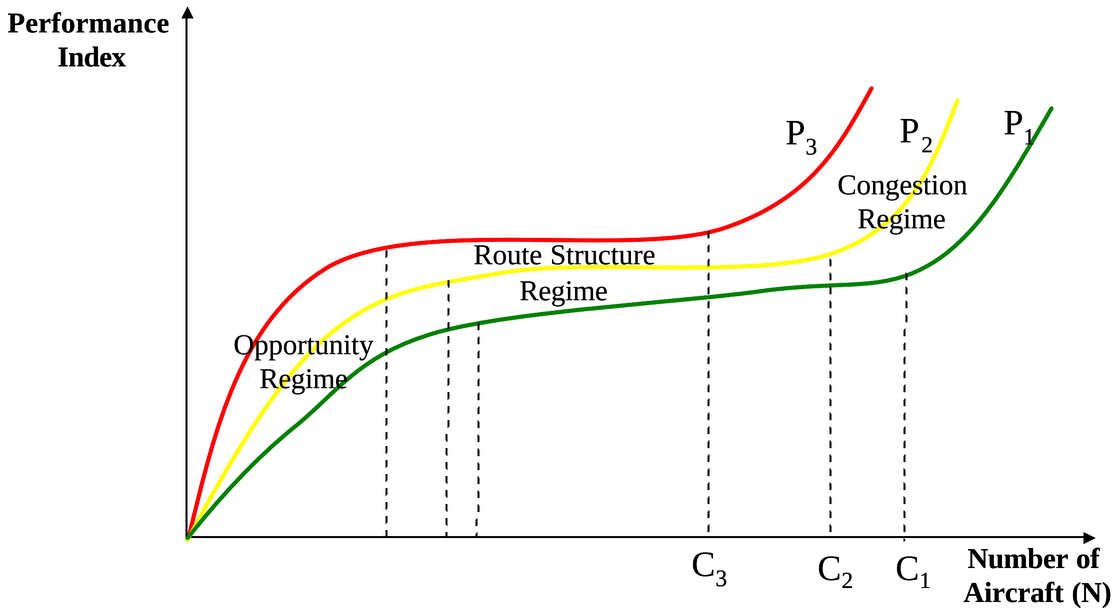

Figure 19. Sector Behavior Curves to Identify Sector Capacity

the sector, though many other measures are possible. Figure 19 also shows that for more complex traffic flow patterns (P3 > P1), fewer flights need to be present for congestion behavior to be observed, and thus the sector capacity is smaller (C3 < C1). 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 taffic flow pattern with the primary set of traffic flow patterns.

This idea is consistent with current NAS operations, in that traffic managers use flow pattern recognition, in addition to the predicted aircraft count, in evaluating the potential for sector congestion. Also, quantifying sector capacity as a function of traffic flow pattern provides a basis for capturing 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 20 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 19. Finally, more research is needed to implement and validate the weather impact method described in Figure 20.

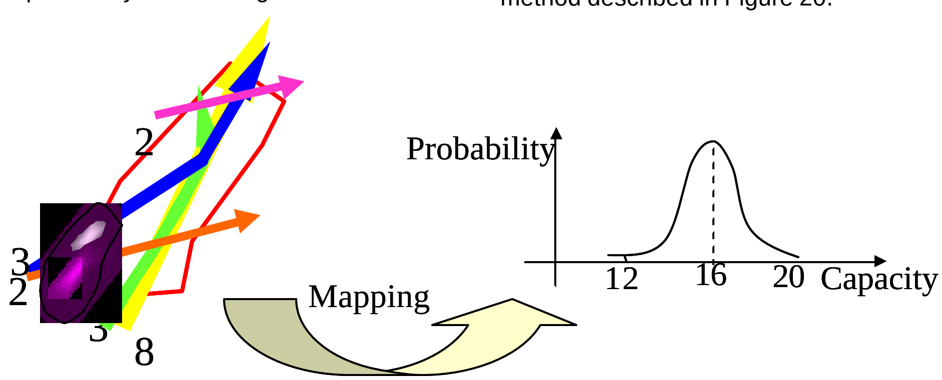

Figure 20. Example of Sector Capacity Distribution under Weather Impact

5.2 Automated Problem Identification

Given operationally-acceptable, probabilistic predictions of demand and capacity, then the probability of future congestion can be reliably used for automated problem identification. When congestion risk exceeds a certain threshold, then traffic managers and NAS users can be alerted that a potential problem exists. For NAS users, this would provide an early indicator of possible TFM action, and allow them to take action consistent with their business needs. If a flight is of critical importance, perhaps due to having many connecting passengers on board, then it might make sense to re-plan the flight and avoid a potential congestion problem.

To help accomplish this, the automation system should be capable of providing real-time feedback on potential congestion problems that a flight might face. Figure 21 illustrates this idea. The originally-filed route for this flight penetrates high-probability congestion areas, and passes near a convective weather system. The alternate route is slightly longer, but travels through less congested airspace, and is thus unlikely to incur future delay. A flight planning system could use this feedback to support dispatchers in deciding when and how to re-plan flights.

If congestion risk becomes sufficiently high, the system will suggest that action be taken to manage the risk. This suggestion would be in the form of a small set of resolution maneuvers, based on pre-established criteria for congestion risk, delay cost, and acceptable maneuver types. The set would be small, since unless there is a radical change in NAS conditions (e.g., a sudden loss of capacity due to equipment failure), the predicted congestion should increase slowly, and can therefore be managed by moving a few flights at a time.

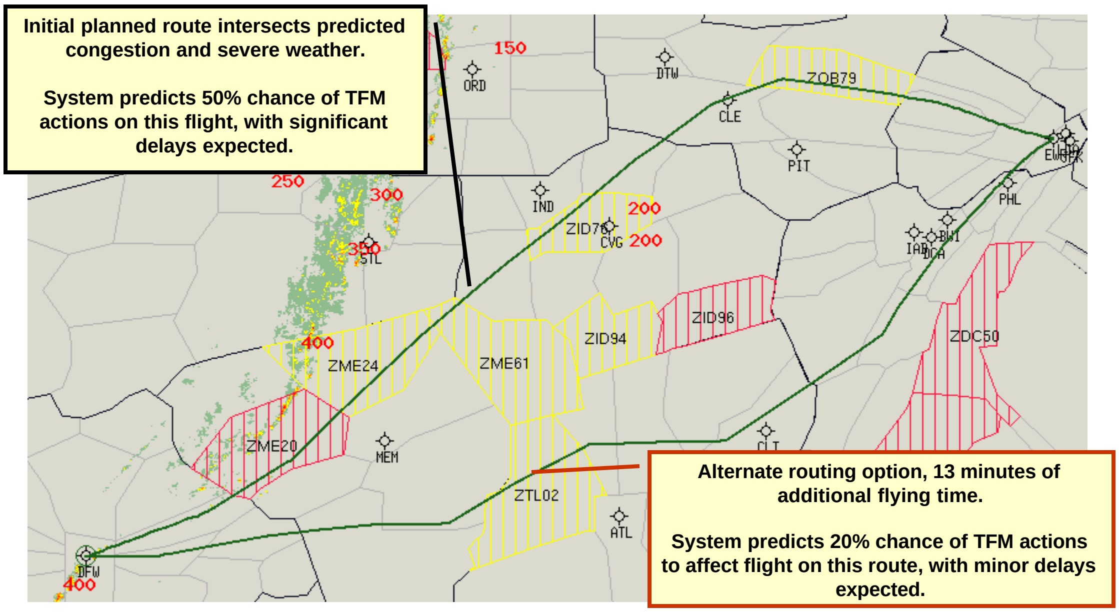

Figure 21. Flight Planning with Probabilistic Congestion Management

5.3 Automated Resolution Development

In the advanced concept, a more sophisticated approach is suggested for generating resolution maneuvers. First, we would like the incremental congestion management maneuvers to make operational sense. For example, it is inappropriate to repeatedly change the same flight from step to step. Also, we would like to execute the right number of incremental maneuvers at appropriate times to get the best long-term result. This is related to the "future flexibility" issue raised earlier. Second, the maneuvers should be chosen to reduce congestion risk while minimizing (or nearly so) a pre-defined and agreed-upon cost criterion. This criterion would set the relative value of airborne vs. ground delay, for example, and might be sensitive to fuel costs. It could also include measures of equity across air carriers or classes of flight operator, such that delays do not unduly impact a particular NAS user. Third, the resolution development process must work in real-time, so that resolutions can be quickly computed, reviewed, and executed.

The algorithm used in the current prototype provides feasible solutions in real time, but does not quite achieve the first and second of these goals. Research is underway to develop an algorithm that does, while maintaining computational feasibility in the probable deployment timeframe (projected to be 2010-2015).

6 Probabilistic Decision-Making for Humans

The concept of probabilistic congestion management involves a higher level of automation than any existing TFM decision support system. Probabilistic estimation and decision making are mathematically complex, and it is impractical for traffic managers to make detailed, probabilistic decisions in real time without automation support. In the concepts 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 have to be addressed through creative interface design and human-in-the-loop evaluations.

7 Conclusion

A concept for managing en route congestion in the presence of uncertainty has been presented. The proposed system can improve the quality of en route congestion management procedures in two primary ways. First, uncertainty is estimated and explicitly considered in problem identification and solution development. Second, flight-specific, rather than flow-oriented, congestion resolution maneuvers are generated. The expected effect of these improvements is to reduce the number of unnecessary flight restrictions while maintaining safe traffic levels more effectively and with less controller workload. An initial prototype has been developed, thus demonstrating the feasibility of the first phase of the concept, and is being used to do an initial benefits analysis.

Acknowledgments

The authors would like to thank Daniel Greenbaum, Sandeep Mulgund, Claude Jackson, Norma Taber, Laurel Rhodes, Dusty Rhodes, Gretchen Jacobs, Celesta Ball, Joseph Sherry, James Wetherly, and a host of others for contributions to the concepts developed here. We would also like to thank David Chaloux and James Woodside for developing the prototype software.

References

Eurocontrol, Pessimistic Sector Capacity Estimation, EUROCONTROL Experimental Centre Note 21/03, EUROCONTROL Experimental Centre, Brétigny-Sur-Orge, France, 2003.

FAA, *Severe Weather and Route Management 2005,* Federal Aviation Administration - Air Traffic Control System Command Center, http://www.fly.faa.gov/Operations/Strategic_Planning/svrwx_handbook.html, 2005.

Kopardekar, P., and Magyarits, S., "Dynamic Density: Measuring and Predicting Sector Complexity," *21st Digital Avionics Systems Conference*, Irvine, CA, 2002.

Masalonis, A., Callaham, M., and Wanke, C., "Dynamic Density and Complexity Metrics for Real-Time Traffic Flow Management," *5th USA/ Europe Air Traffic Management Research and Development Seminar,* Budapest, Hungary, 23-28 June 2003.

Masalonis, A., S. Mulgund, L. Song, C. Wanke, and S. Zobell, "Using Probabilistic Demand Predictions for Traffic Flow Management Decision Support," *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Providence, RI, 16-19 August, 2004.

Megenhardt, D., Mueller, C., Trier, S., Ahijevych, D., and Rehak, N., "NCWF-2 Probabilistic Forecasts," *11th Conference on Aviation, Range, and Aerospace Meteorology,* Hyannis, MA, 4-8 October 2004.

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

Volpe, *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.

Wanke, C., Callaham, M., Greenbaum, D., Masalonis, A., "Measuring Uncertainty in Airspace Demand Predictions for Traffic Flow Management Applications," *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Austin, TX, 11-14 August 2003.

Wanke, C., Rhodes, L., Connolly, K., DeArmon, J., and Topiwala, T., "Progressive Planning: A Decision-Support Framework for Collaborative TFM," *4th AIAA Aircraft Technology, Integration, and Operations Forum,* Chicago, IL, September 2004.

Wanke, C., Song, L, Zobell, S., Greenbaum, D., and Mulgund, S., "Probabilistic Congestion Management," 6th *USA/Europe Air Traffic Management Research and Development Seminar,* Baltimore, MD, 27-30 June 2005.

Wolfson, M. M., Forman, B. E., Calden, K. T., Dupree, W. J., Johnson Jr., R. J., Boldi, R. A., Wilson, C. A., Bieringer, P. E., Mann, E. B., and Morgan, J. P., "Tactical 0-2 Hour Convective Weather Forecasts For FAA," *11th Conference on Aviation, Range, and Aerospace Meteorology,* Hyannis, MA, 4-8 October 2004.