INCREMENTAL, PROBABILISTIC DECISION MAKING FOR EN ROUTE TRAFFIC MANAGEMENT

Craig Wanke and Daniel Greenbaum, The MITRE Corporation, McLean, VA

Abstract

En route airspace congestion, often due to convective weather, causes system-wide delays and disruption in the U.S. National Airspace System (NAS). Today's methods for managing congestion are mostly manual, based on uncertain forecasts of weather and traffic demand, and often involve rerouting or delaying entire flows of aircraft. A new, incremental decision-making approach is proposed, in which prediction uncertainty is explicitly used to develop effective and efficient congestion resolution actions. Decisions are made based on a quantitative evaluation of the expected delay cost distribution, and resolution actions are targeted at specific flights, rather than flows. A massively-parallel simulation of the proposed method has been developed, and results for an operational-scale congestion problem are presented.

Introduction

En route airspace can become congested through either excessive demand or capacity reduction, the latter often due to convective weather. Traffic managers in today's U.S. National Airspace System (NAS) control congestion primarily through manual processes, relying on experience and limited traffic prediction data to develop congestion resolution strategies [1].

Figure 1 depicts a typical situation. Convective weather reduces sector capacity by

Figure 1. Weather-Related Airspace Congestion

reducing available airspace and by inducing flight deviations, which make the traffic patterns harder to manage. Weather forecasts are inherently uncertain, making it difficult to forecast capacity. Traffic demand forecasts are also uncertain, due to changing flight schedules and pilot/airline choices. Predicting congestion, where demand exceeds capacity, is therefore also uncertain. This complexity is compounded by the size of en route congestion problems, which often involve hundreds of flights.

Because it is difficult for traffic managers to resolve en route congestion, conservative strategies are used, which induce excess flight delay and schedule disruption. Also, resolution actions typically involve entire traffic flows, since no tools are available to produce more tightly-targeted solutions.

Background

Present-Day Congestion Management

In the NAS, the Enhanced Traffic Management System (ETMS) provides demand predictions for most sectors in 15-minute bins, for prediction look-ahead times (LAT) of several hours [2]. This information is used to provide the congestion alerts shown in Figure 1. Future aircraft trajectories are predicted based on filed flight plans or schedule data, wind forecasts, and for airborne flights, radar track reports. The peak predicted count for each 15 minute period is compared to an alerting threshold called the Monitor/Alert Parameter (MAP). When the peak count is predicted to exceed the MAP for a sector, the sector is alerted in 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.[1] This is a crude way to capture uncertainty, since departure time estimates are quite uncertain.

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
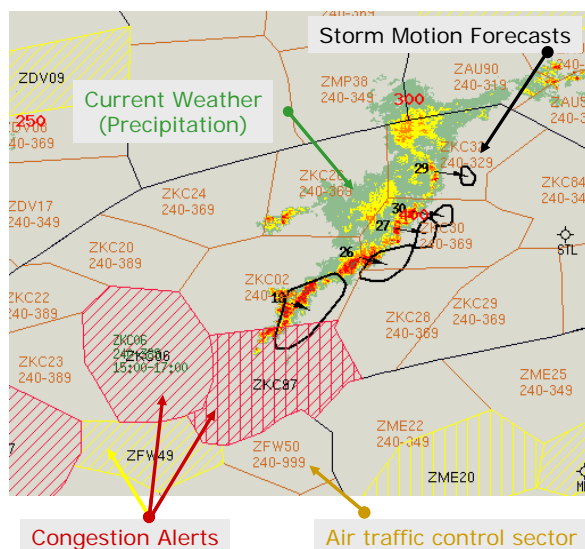
---

[1] The alerts shown in Figure 1 are a composite of all predictions for the next 2 hours. The 15-minute predictions can also be viewed as a matrix (Figure 8).

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 [3,4]. However, MAP is an easily-understood abstraction of workload for alerting purposes. Thus, simply defined, congestion occurs when demand exceeds the MAP.

In today's NAS, these predictions are of limited use. They are one data source for traffic managers who must develop ground delay or reroute initiatives to control congestion. There are no decision-support tools currently available to test traffic management initiatives, though extensive work has been done to develop such tools [5,6,7]. These initiatives typically affect flows of aircraft (e.g., rerouting all traffic between a pair of airports, or miles-in-trail spacing restrictions) rather than individual flights.

Prediction Uncertainty

A key limitation of present-day approaches is the uncertainty of traffic demand and sector capacity predictions. At the timeframes required for en route congestion management (30 minutes to several hours), these uncertainties are significant.

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. 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 extensively described [8], measured [9], and modeled in the context of sector load forecasting [10,11,12]. The models developed by Wanke *et al* [12] were used in this study.

Uncertainties also exist in predicting sector capacity. While ETMS generates alerts based on constant aircraft count thresholds, it is widely accepted that the real capacity of sectors depends on traffic complexity and weather, and should also be treated probabilistically. Research is underway to develop a probabilistic sector capacity prediction [13], but a practical form is not yet available.

Probabilistic 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. A simple application is to replace the current point estimates of traffic demand with an estimate of known statistical properties, such as the median of the probability distribution. This would compensate for prediction bias without requiring the traffic manager to absorb any new information. For example, ETMS predictions at longer look-ahead times (LAT) are more frequently too low than too high, since they cannot include flights that have not yet filed plans. However, this gives little help to the traffic manager in determining when and how to start resolving a predicted congestion problem.

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. Numerous efforts are under way to incorporate uncertainty explicitly into air traffic management decision algorithms. Davidson *et al* [14] suggest a "probabilistic decision tree" approach for making large-scale ATM decisions, by estimating and continually adjusting the probability of the outcome of a set of alternate futures. Mukherjee and Hansen [15] demonstrate a decision tree approach coupled with an optimization method, to planning arrivals to a single, weather-impacted airport. Nilim *et al* [16] have developed a method in which convective weather is modeled as a dynamic, probabilistic process, and flight-specific solutions are found via dynamic programming. Ramamoorthy and Hunter [17] have built a framework to estimate probabilistic TFM decision-making benefits at the NAS level. Algorithms and an early prototype for solving en route congestion probabilistically in real-time, at a single point, have been presented by Wanke *et al* [18,19]. However, the general problem of deciding both *when* and *how* to solve a specific en route congestion problem probabilistically has not been solved.

Research Focus

We aim to provide real-time decision support for en route congestion problems, explicitly accounting for uncertainty, and adapting to changing conditions as problems evolve over time. The decision-making process for congestion resolution is shown in Figure 2. Traffic demand and sector capacity are predicted and compared to identify congestion problems. In the context of other operational constraints, the decision maker must determine whether action is required. If action is warranted, resolution strategies are constructed and executed to achieve the desired goal. This is an iterative process, done today by humans with limited information.
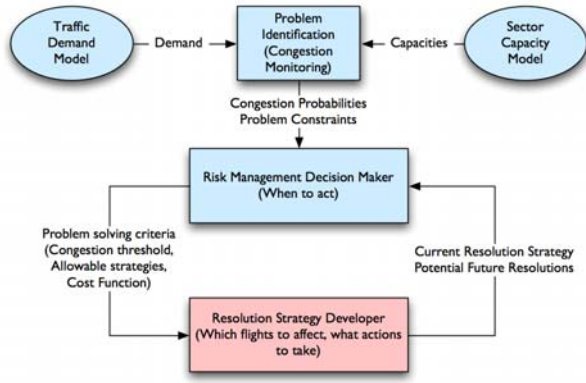
Figure 2. Operational Concept for Probabilistic Congestion Management

Given the uncertainties in the traffic demand and sector capacity predictions, the timing of resolution actions is as important as the type and magnitude of the actions. If action is taken early, unnecessary delays are likely to be incurred. If action is taken late, required maneuvers may be more costly (e.g., rerouting airborne flights). The challenge is to choose the correct set of actions at each point in time, to incrementally solve the congestion problem with the minimum overall impact on involved flights.

Figure 3 illustrates an incremental decision-making framework. There are a series of candidate decision points over time, starting at the time at which a significant congestion problem is detected. At each point, there is a range of discrete options, ranging from doing nothing to "fully" resolving the congestion (i.e., matching the predicted demand directly to the predicted sector capacities).

The method proposed here involves simulating the range of possible outcomes for each combination of current and future decision point options. This is a stochastic simulation, producing a distribution of resolution impact for each decision path. The best path is that which successfully resolves the congestion and incurs the most acceptable impact distribution (e.g., lowest expected value of total flight delays).
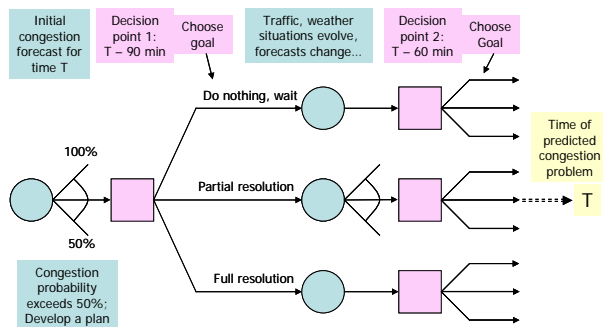

Figure 3. Abstract Incremental Decision Tree

The Incremental Decision-Making Simulation

This process is difficult to simulate. Although the target metrics of congestion and delay are aggregate values, the resolution actions require flight-specific maneuvers (ground delay, rerouting). Therefore, predictions and actual outcomes for individual flights must be simulated. This rules out traditional closed-form methods such as dynamic programming, which require modeling the system as a Markov process. Thus, a Monte Carlo simulation approach is used.

The Decision Tree

Figure 3 is an abstract representation of the decision options. A concrete expression is needed for computation. To this end, we have defined a resolution action in terms of a target *maximum congestion probability*. It is assumed that at each decision point there exists a prediction of traffic demand, a prediction of sector capacity, and an estimate of the probability that the traffic demand exceeds the capacity. The resolution strategy developer (Figure 2) is invoked to find flight-specific actions (reroutes, delays), if needed, to reduce the probability of congestion over the airspace and time of interest to the target goal.

An example of this tree is given in Figure 4. There are three options at the first two decision points. A goal of 1.0 indicates that no resolution actions will be done, and progressively lower values will require more aggressive resolution actions to be taken. A goal of 0.5 is roughly equivalent to matching the most likely value of demand to the most likely capacity, which is analogous to today's deterministic demand/capacity management techniques. Thus, a goal of 0.5 means "resolve all the anticipated congestion". The third option, 0.6, indicates a partial resolution.

The final decision point contains a single goal: to resolve the remaining congestion, by reducing the exceedance probability to 0.5 as the congestion time is reached. At this short time before the problem, the prediction uncertainty is small, so any actual MAP exceedance would likely be small.

Each path through this tree is described by a sequence of maximum probability goals, and represents a single congestion management strategy. The overall mean cost along a decision path is the sum of mean cost at each decision point along the path. This sequence of options at discrete decision points is obviously an approximation of what is in reality a continuous process with many options. The most general form of this problem is likely impractical to solve. This approach provides a solvable approximation, and the number of
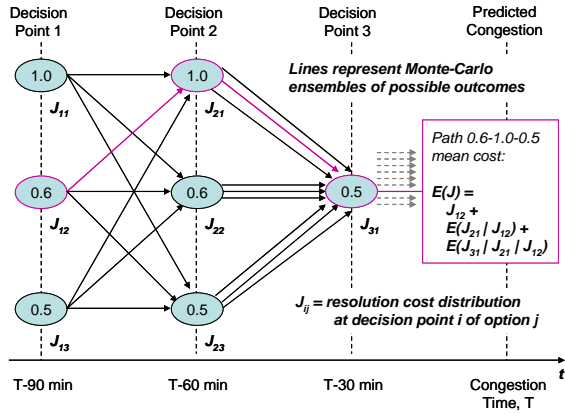
Figure 4. Three-Stage Decision Tree

options and decision points can be increased at will, with only the cost of longer computation time.

Traffic Modeling

There are two traffic models required for the simulation. The first is a traffic demand prediction uncertainty model, used to determine the probability of congestion given a standard deterministic prediction of future aircraft trajectories. It only needs to predict the magnitude of the peak traffic counts, not specific flight uncertainty. To meet this need, a closed-form, statistical uncertainty model for demand predictions was developed [18]. This model, the Aggregate Demand Model (ADM), is based on a comprehensive set of statistics on sector peak count prediction uncertainty compiled in prior research [9]. The model forecasts peak traffic demand distributions based on four variables: the look-ahead time, the deterministic predicted peak count, the number of airborne flights in the peak count prediction, and the primary sector traffic type (departure, en route, arrival, mixed). This model is very fast to compute and can be used in either simulation or real-time applications.

The second model is needed to simulate the range of possible traffic outcomes in a situation, given an initial prediction. It must be flight-specific, to allow resolution actions to be computed

at each decision point. A Monte Carlo model was developed for this purpose [19]. This model begins with a set of predicted flight trajectories, and based on empirically-derived distributions, determines a set of possible actual outcomes for that flight. It models the following, for predicted flights:

- Cancellations
- Departure time estimation errors
- Changes in route and cruise altitude
- Flight progress estimation (speed) errors

Also, the model will create and add a set of flights that have not filed at the time of the prediction but will appear before the time for which the prediction was made ("pop-ups").

Airspace Capacity Modeling

As noted earlier, there are not yet accepted methods for sector capacity prediction, especially in the presence of weather. For this study, the ETMS MAP was used as a baseline sector capacity, and the impact of weather on a sector was simulated by reducing the MAP over a time period. Once a model for capacity prediction uncertainty is available, it will be incorporated into the simulation.

Congestion Resolution

For each option at each decision point, a resolution strategy must be developed to meet the desired maximum congestion probability. The simulation uses a heuristic algorithm that can be rapidly computed [18], and has been shown to provide effective, though not optimal, flight-specific solutions [20].

The resolution process begins by defining two airspaces. The first, the Congestion Resolution Area (CRA), contains those sectors identified as being congested. Flights that penetrate the CRA during the congestion period are candidates for resolution maneuvers. The second, the Congestion Management Area (CMA), is a larger group of sectors surrounding the CRA. These sectors are monitored during the resolution development
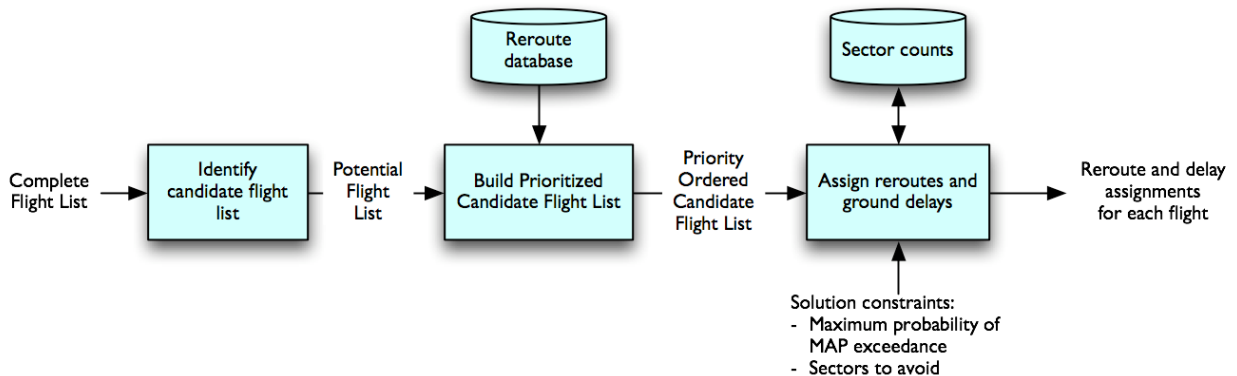

Figure 5. Overview of Heuristic Congestion Management Algorithm

process so that resolution maneuvers do not create additional congestion in the CMA.

Figure 5 illustrates the process. Candidate flights are subtracted from the CMA traffic count predictions. Then, the flights are placed in priority order. Airborne flights are first, sub-ordered by arrival time to the CRA. Pre-departure flights are then added to the bottom of the list, also sub-ordered by arrival time. Next, a series of alternate route options are generated for each flight. These are selected from a database of predefined and historically-flown routes, keyed by origin-destination pair. Pre-departure flights also have the option of taking ground delays up to a set maximum value.

Resolution maneuvers are assigned in a single pass through the ordered candidate list. First, the current flight trajectory is added to the CMA sector counts, and the ADM is used to evaluate the resulting congestion probabilities. If the maximum congestion probability is not exceeded, then the flight is not rerouted or delayed. If the maximum probability is exceeded, then predicted trajectories for all combinations of alternate routes and (if the flight is pre-departure) ground delays are constructed. Of the trajectories that do not violate the congestion constraint, the one with the earliest arrival time at destination is chosen as the best option. If no trajectories work, the flight is not modified, and the congestion probability goal will not be achieved.

Flights that are early in the prioritized list tend to be easier to solve. As the processing reaches the end of the order, it is harder to find options that do not exceed the congestion threshold, so later flights may experience more severe reroutes and delays. This processing order is a key factor in determining the optimality and equity of a proposed solution, and it remains an area for experimentation to try other sorting approaches. The current implementation divides the overall list into those that can be ground delayed and those that cannot. Flights that cannot be ground delayed are processed first, since they have less flexibility in terms of actions that can be taken. Also, rerouting airborne flights is generally more difficult and expensive than rerouting flights that have not yet departed.

Prediction Evolution

In order to capture the interesting features of probabilistic decision-making, it is not sufficient to simulate how actual traffic ensues from a given prediction. We must also simulate how the state of knowledge (i.e., the updated prediction) changes as simulation time passes. There is a single traffic prediction at the start of the simulation. There are N Monte Carlo outcomes modeled from that prediction. When simulation time is advanced to the next decision point, each of those outcomes will also have an updated prediction, and that prediction will reflect what has become known since the last decision point. For example, if flight ABC123 is contained in the initial prediction, but in a particular outcome ABC123 is cancelled, then there is some time at which this becomes known. If the flight is cancelled between the first and second decision points, then the prediction at the second decision point should not contain ABC123, and ensuing resolution actions will not attempt to delay or reroute that flight.
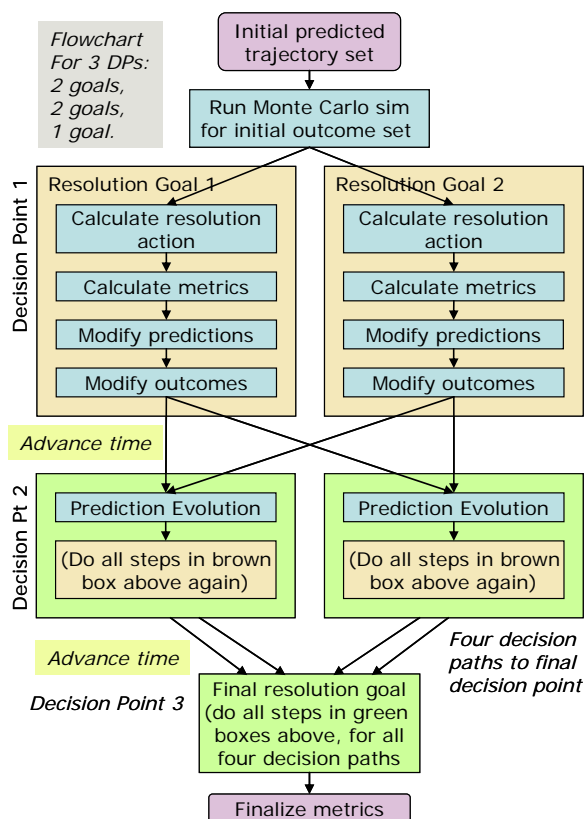
For the initial simulation runs, a very simple model of prediction evolution was used. Flights which cancel do so 15 minutes before their planned departure time. Pop-up flights file 30 minutes before their planned departure time. Flights which leave later than predicted are discovered to be late when their initial departure time passes. Flights which are rerouted receive the new route at takeoff. These rules are simplistic, and will be replaced with more realistic, statistically-modeled behavior based on empirical studies. But they are realistic enough to generate interesting results.

Evaluating the Decision Tree

Figure 6 illustrates the simulation flow. This assumes that the congestion resolution and management areas has been identified (CRA and CMA), sector capacities are defined, and supporting data has been assembled (wind forecasts, Monte Carlo distribution parameters, etc.) The process begins with a predicted trajectory set, which is used as a basis by the Monte Carlo traffic simulation to generate a set of possible "actual" outcomes for the predicted flights. These characterize the variety of ways that the situation can play out, based on the statistical model of flight-specific variations.

Next, decision point 1 (DP1) is evaluated (see the large brown boxes in Figure 6). For each option, the initial prediction is used to compute a resolution action that meets the desired maximum congestion probability. The incurred cost is saved. A new prediction set is developed by substituting in the resolution maneuvers. Finally, for flights modified by the resolution, the corresponding flights across the full set of actual outcomes are altered to capture the effect of the resolution maneuvers. The Monte Carlo model is used again for this step. If the resolution maneuver for a specific flight involves a ground delay, then new "actual" departure times are generated for that flight for each of the N traffic outcomes. If the resolution maneuver includes a reroute, then the flight progress and route/altitude variability models are re-applied to that flight for all outcomes. This maintains consistency. Flights that are not been

maneuvered at a decision point retain the same set of trajectory variations at the next decision point, and the combined unmodified and modified set represents the altered range of traffic outcomes resulting from the executed resolution maneuver.

Once all options are calculated, time is advanced to DP2 (the large green boxes). Each DP2 option must be calculated for each DP1 result, to fully explore each possible decision path. The first step is to apply the prediction evolution rules described earlier to reflect the passage of time and increase in knowledge about the outcomes. Then, new resolution maneuvers are calculated much the same as for DP1, with two important differences. First, there is now a different prediction associated with each actual outcome, so there will be N different resolution maneuvers calculated. This produces a distribution of resolution cost, rather than a single value. Second, flights that were maneuvered at a previous DP are exempted from further maneuvers, reducing flexibility. This is to avoid making repeated schedule changes to the same flight, which is considered highly undesirable by airlines and pilots.

This process continues through the DPs until all possible paths converge at the final DP, and the final resolution goal is computed.

Statistical metrics are captured at each decision point along each decision path. If a resolution strategy is applied, the number of flights affected, type and number of maneuvers generated, and the delays produced are saved. Predicted post-maneuver congestion probabilities are calculated to determine if the resolution strategy succeeded. "Actual" congestion probabilities, based on the modified Monte Carlo outcomes, are also saved.

Implementation

The simulation has been developed in Java, and is highly parallelized. Because of the computational structure, groups of Monte Carlo outcomes can be independently carried through the decision paths. Intermediate results are saved, and recombined for analysis after all parallel runs are complete. The run described below for the sample scenario took approximately 30 minutes to complete on 8 dual-processor/dual-core systems.

Example Congestion Scenario

A realistic congestion scenario was developed by selecting traffic from a busy period of a clear-weather day, and creating congestion by postulating sector capacity reductions due to weather. This was done to avoid a situation where significant traffic management actions were actually taken in response to congestion, which would make it difficult to assess the performance of the proposed congestion management technique.

The area of interest for this scenario comprises four laterally or vertically adjacent sectors in the Washington Air Route Traffic Control Center (ARTCC), denoted ZDC. Three of the sectors

Figure 6. Flowchart for Three Decision Points

Figure 7. Congestion Scenario Airspace

	1700	1715	1730	1745	1800	1815	1830	1845	1900	1915
ZDC14 [14]	8	11	8	6	6	9	11	13	12	11
ZDC16 [15]	6	6	6	9	8	10	11	11	5	5
ZDC36 [16]	5	9	7	13	8	12	14	9	6	7
ZDC72 [18]	15	17	11	7	14	14	17	18	13	13

Figure 8. Predicted Congestion

(ZDC sectors 72, 16, and 36) are visible in Figure 7; the fourth, ZDC14, is a low altitude sector below 16 and 72. It is assumed that these sectors have a a capacity reduction of 5 below their nominal MAP values for the period 1800 to 2000 UTC. This area was designated as the CRA. The CMA, composed of all sectors adjoining the CRA either laterally or vertically, includes 38 sectors.

For this scenario, there is congestion predicted in the CRA sectors at 1700 UTC. Figure 8 shows the median peak traffic counts and congestion alerts. Each row of the matrix is a time-series prediction for one sector, at 15 minute intervals. The normal peak count threshold (MAP value) for each sector is next to the sector name. The number in each cell indicates the median peak traffic count value from the ADM. Red alerts indicate a greater than 0.75 probability that the actual demand exceeds the sector capacity. Yellow alerts indicate a greater than 0.50 probability. Thus, the period outlined in blue represents a serious congestion situation that needs to be resolved.

The decision tree used in this scenario is nearly identical to that shown in Figure 4, with the exception that congestion actually starts 75 minutes, rather than 90, after the first prediction (though the first red alert is 90 minutes later). The

system is tasked to resolve potential congestion from 1815 to 1915 UTC. Approximately 1500 flights pass through the CMA during or near this period, 191 of which also penetrate the CRA.

By studying the statistical features of the output distributions, it was determined that 250 Monte Carlo outcomes were required to obtain a 95% confidence that on the estimate of the mean number of aircraft affected by each resolution action was within one aircraft of the actual mean.

Results

The overall impact results from the nine decision paths are summarized in Table 1. The metrics used are the mean number of flight affected (across all Monte Carlo outcomes), mean total positive delay (TPD), and the mean TPD per flight. Only positive delays were tabulated so as not to be cancelled out by negative delays (early arrival times). While negative delays are desirable in some cases, they may also represent a disruption in other cases. Therefore negative delays were treated as neutral in our analysis of decision path goodness.

The first decision path 1.0-1.0-0.5 (in red) failed to produce a successful congestion solution in many of the outcomes, as observed from the final
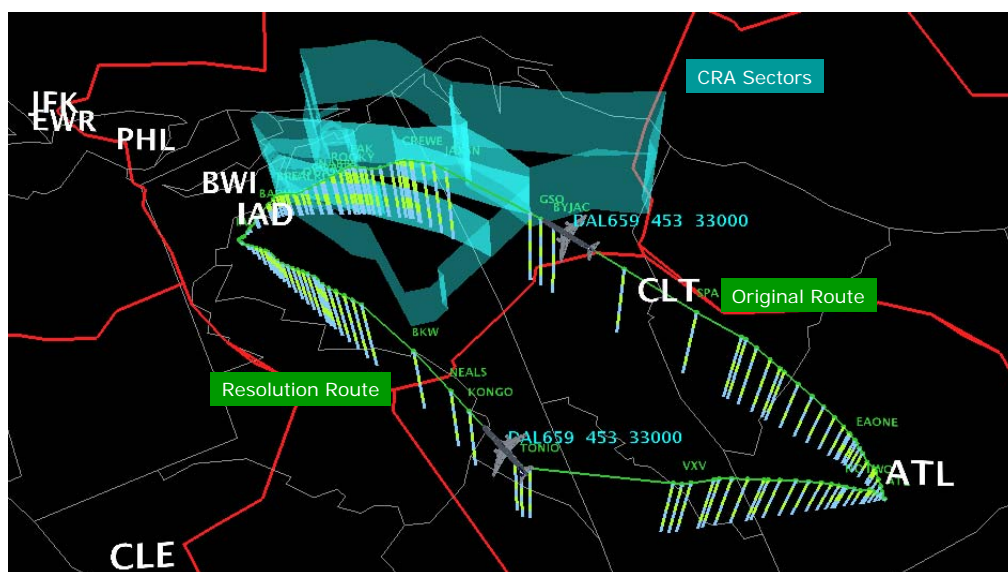

Figure 9. Sample Resolution Reroute Maneuver at Decision Point 1

Table 1. Decision Path Impact Summary

Decision Path	Mean No. of Flights Affected	Mean Total Positive Delay (min)	Mean TPD per Flight
1.0-1.0-0.5	29.5	271	9.2
1.0-0.6-0.5	53.6	400	7.5
1.0-0.5-0.5	71.1	557	7.9
0.6-1.0-0.5	48	318	6.6
0.6-0.6-0.5	58.6	392	6.7
0.6-0.5-0.5	71.1	489	6.9
0.5-1.0-0.5	83.6	643	7.7
0.5-0.6-0.5	87.5	678	7.7
0.5-0.5-0.5	93.3	752	8.1

Table 2. Decision Path Impact Details

Decision Path	Metrics	DP1	DP2	DP3
Best Strategy : 0.6-1.0-0.5	No. of AC	31.00	0.00	16.98
	Rerouted	15.00	0.00	9.08
	Ground Delayed	21.00	0.00	9.84
	Reroute Delay(min)	-33.88	0.00	101.76
	Ground Delay(min)	115.00	0.00	77.46
	Positive RR Delay	7.92	0.00	117.59
Worst Strategy : 0.5-0.5-0.5	No. of AC	73.00	12.94	7.35
	Rerouted	34.00	6.77	4.72
	Ground Delayed	60.00	9.32	2.74
	Reroute Delay(min)	-41.05	26.75	100.39
	Ground Delay(min)	430.00	104.28	23.56
	Positive RR Delay	44.35	44.37	105.08

predicted congestion probabilities. Thus, the low number of flights and total delay values are misleading. This path is the "wait until the last minute" solution, and it is not surprising that it can fail. Also, the maneuvers generated were much longer per flight than for other paths. The rest of the paths produced successful congestion resolutions for all runs. In other words, after the final step, the remaining congestion probabilities for the CRA sectors were all 0.5 or less.

Path 0.6-1.0-0.5, representing a moderate early maneuver followed by full resolution at the end, had the best overall performance. 48 flights were affected, with an average of 6.6 minutes of arrival delay per flight.

The last three paths represent strategies which solve all the congestion at the first DP, and the next two DPs are used to ensure that the congestion stays solved. These prove to be the most expensive. The last path, where all the congestion is solved at all three points, has twice the impact of the best successful path.

It is useful to examine some details of the solutions generated. Table 2 breaks down the impact statistics by decision point for the best and worst strategies (paths 4 and 9, respectively). Values for DP2 and DP3 are means across the 250 outcomes. In both strategies, the mean reroute delay at DP1 is negative. This is because reroutes chosen well in advance of the congestion can frequently be selected to arrive at the same time or earlier than the original route.

An example reroute is shown in Figure 9. This flight from Atlanta to Washington was originally filed to penetrate several sectors of the CRA (highlighted in blue). The resolver selected a predefined alternate route which, given the wind forecast on this day, was computed to have approximately the same arrival time as the original. It avoids the CRA and does not produce any new congestion in the CMA.

Table 2 also shows that ground delays are a more prevalent part of the strategy at DP1 and DP2 than at DP3. This naturally follows from the composition of involved flights at the different DPs. At DP1, most flights involved in the predicted congestion are still on the ground, and the fraction of airborne flights in the congestion area increases as time advances. In summary, resolving the congestion later requires longer and a higher proportion of reroutes, but because the prediction uncertainty has decreased, fewer flights need to be moved. This is the essential tradeoff.

Table 3 shows some statistics of the resolution distribution for the best strategy at DP3. It is apparent from the standard deviations that there is a wide variation across outcomes. The uncertainty in the traffic predictions at DP1 is quite high, and thus for this scenario it seems sensible to resolve part of the predicted congestion at DP1, and then wait as long as possible to identify and resolve the rest. More scenarios are needed to study this effect.

These results are subject to a few caveats. First, it is a single problem. The ADM and the Monte Carlo traffic models do not always agree, because the ADM is an aggregated model and the

Table 3. Resolution Statistics Example

Metric	Mean	Standard Deviation
No. of AC	16.98	6.77
Rerouted	9.08	6.18
Ground Delayed	9.84	4.40
Total Positive Delay (minutes)	179.2	76.6

Monte Carlo results are specific to the particular airspace and traffic situation. Thus, more problems need to be simulated to draw good conclusions. Also, the statistics used to generate the Monte Carlo model are from a few months of data, and should be expanded and updated to get more useful results.

Applications

This simulation has several useful applications. First, assuming computational power continues to increase, it represents a prototype of a real-time congestion resolution decision-support system. Many issues would need to be addressed, including: automatic updating of probability models, cognitive engineering of the human-computer interface, incorporation of probabilistic weather/capacity forecasts, and how to best allow airspace users to participate in resolution maneuver generation. The last could be handled by allowing users to submit preferred resolution options for their flights (already being discussed in government/industry working groups), or perhaps by automated negotiation between the resolution generator and airline flight planning software.

The second application is to develop heuristics for near-term congestion resolution tools and procedures. We plan to run a matrix of interesting congestion problems, and analyze the results to derive rules for effective congestion resolution actions and timing.

Thirdly, the simulation is useful for cost-benefit analysis. In the current form, the simulation is being used to evaluate the benefits of incremental, probabilistic decision-making, as compared to today's approaches. The baseline case includes probabilistic traffic forecasts, but is being extended to capture probabilistic capacity forecasts. This will provide a platform for evaluating the potential benefits from proposed probabilistic weather forecasting products. If those weather products can be used to provide a probabilistic forecast of airspace capacity, then the utility of those forecasts for congestion resolution can be directly simulated and evaluated.

Also, if a new technology is proposed that reduces uncertainty in demand or capacity prediction (e.g., a surface management system, which would reduce departure prediction uncertainty), then the delay reduction benefits can be estimated via simulation.

Conclusion

A Monte Carlo simulation technique for evaluating incremental, probabilistic decision-making in en route congestion management has been developed. It has been demonstrated using moderately-sized traffic congestion scenario, and used to examine a variety of possible congestion resolution strategies. The simulation can be used to learn about when and how to solve a variety of airspace congestion problems, and to aid in cost benefit analyses of several types. In particular, it can be used to evaluate the potential benefits of advanced, probabilistic aviation weather forecasts. It also represents a prototype of a future congestion management decision support system, in which probabilistic information is used directly to do more efficient en route traffic management.

References

[1] *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

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

[4] 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.

[5] 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.

[6] Sridhar, B., G. Chatterji, S. Grabbe, and K. Sheth, "Integration of Traffic Flow Management Decisions," *AIAA Guidance, Navigation, and Control Conference,* Monterey, CA, August 2002.

[7] Wanke, C., N. Taber, S. Miller, C. Ball, and L. Fellman, "Human-in-the-Loop Evaluation of a Multi-Strategy Traffic Management Decision Support Capability," 5th *USA/Europe Air Traffic Management R&D Seminar*, Budapest, Hungary, June 2003.

[8] Krozel, J., D. Rosman, and S. R. Grabbe. "Analysis of En Route Sector Demand Error Sources," AIAA 2002-5016, AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, CA, USA, 5-8 August 2002.

[9] Wanke, C., M. Callaham, D. Greenbaum, A. Masalonis, 2003. "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.

[10] Meyn, L. A., 2002. "Probabilistic Methods for Air Traffic Demand Forecasting," AIAA 2000-4766, AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, CA, USA, 5-8 August 2002.

[11] Mueller, K. T., J. A. Sorenson, and G. J. Couluris. "Strategic Aircraft Trajectory Prediction Uncertainty and Statistical Sector Traffic Load Modeling," AIAA 2000-4765, AIAA Guidance, Navigation, and Control Conference and Exhibit, Monterey, CA, USA, 5-8 August 2002.

[12] Wanke, C., S. Mulgund, D. Greenbaum, and L. Song, 2004: "Modeling Traffic Prediction Uncertainty for Traffic Management Decision Support," AIAA Guidance, Navigation, and Control Conference and Exhibit, Providence, RI, 16-19 August 2004.

[13] Song, L., D. Greenbaum, & C. Wanke, "Predicting Sector Capacity for TFM Decision Support," 6th AIAA Aviation, Integration, and Operations Conference, Wichita, KS, 25-27 Sept. 2006.

[14] Davidson, G., J. Krozel, S. Green, and C. Mueller. "Strategic Traffic Flow Management Concept of Operations," 4th Aviation Technology, Integration, and Operations Conference, Chicago, IL, 20-22 September 2004.

[15] Mukherjee, M., and M. Hansen. "Dynamic Stochastic Optimization Model for Air Traffic Flow Management with En Route and Airport Capacity Constraints," 6th USA/Europe Air Traffic Management Research and Development Seminar, Baltimore, MD, 27-30 June 2005.

[16] Nilim, A., L. El Ghaoui, and V. Duong. "Multi-Aircraft Routing and Traffic Flow Management under Uncertainty," 5th USA/Europe Air Traffic Management Research and Development Seminar, Budapest, Hungary, 23 – 27 June 2003.

[17] Ramamoorthy, K., & G. Hunter, "Modeling and Performance of NAS in Inclement Weather," 6th AIAA Aviation, Integration, and Operations Conference, Wichita, KS, 25-27 Sept. 2006.

[18] Wanke, C, Zobell, S., and Song, L. "Probabilistic Airspace Congestion Management," 5th Aviation Technology, Integration, and Operations Conference, Arlington, VA, 26-28 September 2005.

[19] Wanke, C., Song, L., Zobell, S., Greenbaum, D., and Mulgund, S. "Probabilistic Congestion Management," 6th USA/Europe Seminar of Air Traffic Management R&D, Baltimore, MD, June 27-30 2005.

[20] DeArmon, J., C. Wanke, D. Greenbaum, L. Song, S. Mulgund, S. Zobell, and N. Sood, "Solving Probabilistic Airspace Congestion: Preliminary Benefits Analysis," 6th AIAA Aviation, Integration, and Operations Conference, Wichita, KS, 25-27 Sept. 2006.

Notice

The contents of this material reflect the views of the authors and The MITRE Corporation and do not necessarily reflect the views of the FAA or the DOT. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty or guarantee, or promise, expressed or implied, concerning the content or accuracy of the se views.

Keywords

Acknowledgements

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.

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.