Sequential Congestion Management with Weather Forecast Uncertainty

Craig Wanke* and Daniel Greenbaum†
The MITRE Corporation, McLean, VA, 22102

En route airspace congestion, often due to convective weather, causes system-wide delays and disruption in the U.S. National Airspace System (NAS). Present-day 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 sequential decision-making approach is proposed, in which traffic and weather forecast prediction uncertainty is quantified and explicitly used to develop efficient congestion resolution actions. The method is based on Monte Carlo simulation of traffic and weather outcomes from a specific forecast. Candidate sequential decision strategies are evaluated against the range of outcomes to determine the best course of action. 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 weather-induced airspace congestion scenario is explored using the simulation, and three different levels of weather forecast uncertainty are postulated. Weather forecast uncertainty was shown to affect when and how aggressively to act to solve the congestion problem. Forecast uncertainty also increases the mean and variance of the congestion resolution cost. The simulation can be used both to learn about solving airspace congestion problems, and to do several types of cost benefit analyses. It is also a prototype of a future, real-time, probabilistic decision support aid for tactical traffic management.

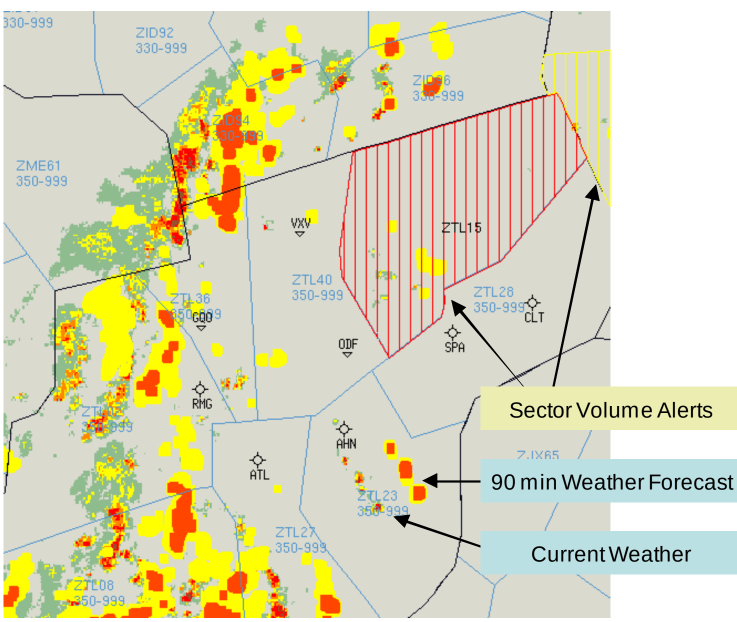

Figure 1. A Convective Weather Event. Weather moving west-to east is restricting use of airspace in the Atlanta Air Route Traffic Control Center.

I. Introduction

En route airspace can become congested either through excessive demand or capacity reduction, the latter often due to convective weather. Traffic managers in the present-day 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] Traffic managers must identify and resolve impending congestion based on uncertain predictions of weather and traffic demand. They must estimate the future loss of airspace capacity due to predicted weather, and implement strategies to restrict demand such that congestion is avoided. This process may involve hundreds of flights, and is done with little or no automated decision support.

At the strategic timeframe, more than 2 hours in advance of anticipated congestion, strategies are developed in a teleconference

---

* Senior Principal Simulation and Modeling Engineer, M/S N450. Senior Member, AIAA.
† Lead Software Systems Engineer, M/S N450.

of traffic managers, meteorologists, and airspace users (e.g., airlines). These strategies involve delaying or rerouting large groups of flights, to hedge against uncertainty. At the tactical timeframe (30 minutes to 2 hours in advance), traffic managers at en route and terminal air traffic control facilities apply smaller scale initiatives to resolve remaining congestion. But even at the tactical timescale, these actions are based on uncertain predictions and frequently affect more flights than necessary.

Figure 1 illustrates the tactical congestion management problem when convective weather is present. Current weather, in the form of rainfall severity, is displayed in a grid of 6-level Vertically Integrated Liquid (VIL) data. Areas of yellow or red (VIL level 3 or greater) represent potentially hazardous areas, thus reducing the effective capacity of the affected air traffic control (ATC) sectors. The data shown is from WSI Corporation's NOWRAD product. Weather forecasts in a similar form are now available from the Corridor Integrated Weather System (CIWS)[2], allowing prediction of capacity loss. Traffic demand predictions are provided by the Enhanced Traffic Management System (ETMS)[‡], and projected peak sector traffic counts are compared to a sector-specific threshold – the Monitor/Alert Parameter- to determine when to show an alert (sector ZTL15 in Figure 1). These systems are independent, however. The alerting thresholds are not reduced in the presence of weather, and indeed there is no accepted way to estimate a reduced threshold when weather is present. It is also difficult to execute precise resolution strategies at the tactical timeframe, because automation does not yet exist to distribute flight-specific resolution maneuvers quickly to the personnel who need to execute them (controllers, traffic managers, airline dispatchers, etc.)

As a result, current traffic management practice is to take aggressive action at strategic timescales, using large initiatives such as Airspace Flow Programs or National Playbook reroutes.[1] This often leads to unnecessary delay, since weather events often do not occur as forecasted. If tactical congestion resolution can be improved, then less needs to be done at the strategic time scale.

This project is developing a tactical congestion resolution decision support method which explicitly accounts for uncertainties in traffic and weather forecasting. This method employs sequential decision-making and flight-specific maneuvers (e.g., ground delays and reroutes) to develop tightly targeted congestion resolution strategies, rather than relying on large scale traffic management initiatives. In conjunction with the FAA's plans to deploy automated flow strategy execution tools, such a method could improve the ability of the NAS to adapt to severe weather disruption.

II. Previous Work

In a previous paper, we described an implementation of this method to account for traffic demand prediction uncertainty, but assumed the weather forecast was perfectly accurate.[3] This laid the foundation for the approach, but the weather forecast uncertainty and the corresponding sector capacity uncertainty are critically important. This must be addressed to obtain an operationally feasible approach.

The congestion resolution process aims to balance traffic demand with available sector capacity. Thus, it requires predictions of demand and capacity. Predicting capacity ideally requires knowledge of both how sectors behave under particular traffic conditions, a forecast of severe weather which will impact sector operations, and a model for determining the reduction in capacity due to the forecasted weather. Finally, if predictions and prediction uncertainties are known, there are several decision-making approaches which can be used to determine appropriate resolution actions. Current research in all of these areas is described below.

A. Traffic Demand Prediction

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[4], measured[5], and modeled in the context of sector load forecasting.[6-8] Gilbo and Smith[9] measured prediction errors for the ETMS itself, and proposed a

---

[‡] The ETMS is currently being replaced by the Traffic Flow Management System (TFMS), which will initially use similar prediction algorithms. It will have new software and hardware to facilitate upgrades, maintenance, and improved data distribution. At this writing the transition is incomplete, and though the name of the operational system has been changed to TFMS, the ETMS demand prediction software is still running. Thus, we still refer to the operational system as the ETMS.

regression model to obtain better estimates of actual traffic demand. Models developed by Wanke, et. al.[8] were used in this study.

B. Weather Forecasting for ATM

Some current tactical weather forecasts include uncertainty parameters as a gridded set of probabilities that significant convective weather will occur.[2,10] These are not yet easily applied to decision support applications, since they lack correlation; it is not enough to know the probability of severe weather at a particular point, but we must also know how the likelihood of weather at one point relates to likelihood at other points. For example, we might know that there is a 50% likelihood of departure sector A being impacted, and a 50% likelihood of departure sector B being impacted. But for planning purposes, we need to know the joint probabilities, i.e., are we likely to have one or the other sector available, or if one is blocked, is the other one likely to also be blocked? The answer to this question is critical to the traffic management strategy. Some researchers are pursuing ensemble forecasts, specifically designed for air traffic management, to provide information of this type.[11] For this project, we developed a simple and completely generic model of weather forecast uncertainty, based on uncertainties in weather propagation speed and growth or decay rate. This is sufficient to identify the relative effects of different levels of weather uncertainty on decision making, but we would like to replace it with a real weather forecasting product when a suitable one becomes available.

C. Sector Capacity Prediction

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. The impact of traffic complexity on sector capacity has been frequently studied[12-14], but quantifying the impact of weather on capacity is a relatively new topic. Song, et. al.[15] propose an approach based on traffic patterns and predicted weather coverage. This builds on work done by DeLaura and Evans[16,17] relating weather parameters to observed weather avoidance behavior by pilots. Those observations were extended to predict the impact of weather on air route utilization, which is another way to capture airspace capacity.[18] None of these concepts are quite ready for application to congestion management, however, and a simpler method was employed for this study. It relates observed weather coverage (VIL level 3 and above) to observed sector throughput, and applies a simple rule to reduce the effective MAP of a sector based on the predicted weather coverage.[19,20]

D. Probabilistic Decision Making

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. Hoffman et. al.[21] have developed a decision tree approach for making strategic traffic management decisions, by estimating the probabilities of a set of possible weather outcomes, and structuring traffic flows to achieve an optimal solution. Mukherjee and Hansen[22] demonstrated a decision tree approach coupled with an optimization method for planning arrivals to a single, weather-impacted airport. Nilim et. al.[23] 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[24] 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.[25,26] Finally, as noted above, we have developed a probabilistic, sequential decision making approach for tactical congestion management subject to traffic demand prediction uncertainty.[3]

While progress has been made, the general problem of deciding both *when* and *how* to solve a specific en route congestion problem probabilistically, taking both traffic and weather (capacity) uncertainty into account, has not been fully and practically solved. It is likely that a strategic decision support concept such as decision trees can be used with a tactical decision support method such as that described here to provide a more complete traffic management toolset.

III. A Sequential Congestion Management Concept

We aim to provide real-time decision support for tactical 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

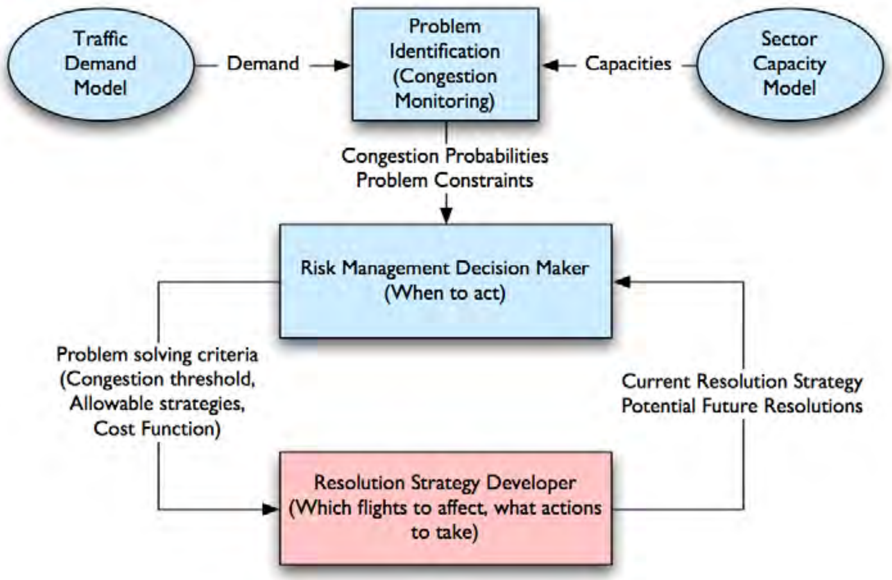

Figure 2. Operational Concept for Probabilistic Congestion Management

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.

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.

In Figure 2, the problem is quantified in terms of "congestion probabilities", and problem solving criteria include "congestion thresholds". This is one convenient way of including probabilities in the calculation of resolution strategies. If the traffic demand and sector capacity predictions can be generated such that they are probability mass functions (PMFs)[§] rather than point estimates, then these predictions can be combined through convolution to determine the probability that the demand will exceed the sector capacity at each prediction time. This is shown in Figure 3 for a single sector, via box plots of predicted distributions of demand and capacity for a single sector over a series of 15-minute intervals.

The heavy blue lines represent the 50th percentile prediction of airspace capacity, surrounded by a blue box representing the 25th to 75th percentile prediction range (a standard box plot). The capacity is well-known at short prediction look-ahead times (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. Similarly, the demand data (gray, purple, and magenta boxes) increases in uncertainty as LAT increases. The demand boxes are color-coded by the probability that the demand will exceed the capacity (i.e., there is congestion) during each interval. There are two periods (1515 – 1545, 1630 to 1730) during which the probability of congestion is greater than 75%.

[§] In this work, capacity and demand are expressed in terms of aircraft count, an integer quantity, thus the distribution is expressed as a PMF rather than a probability density function (PDF).

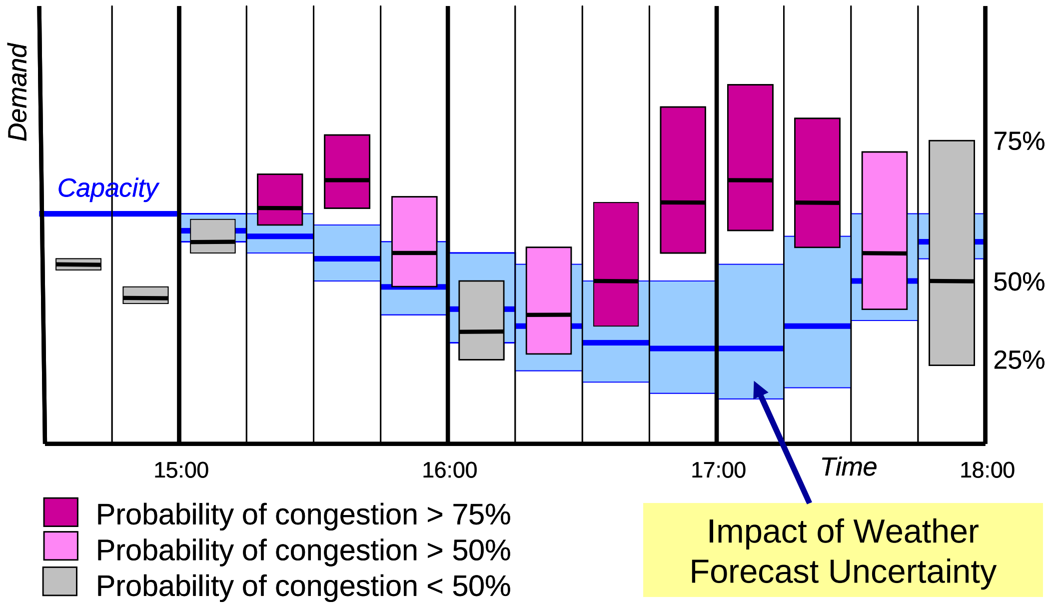

Figure 3. Predicting probability of congestion using uncertainty distributions for traffic demand and sector capacity.

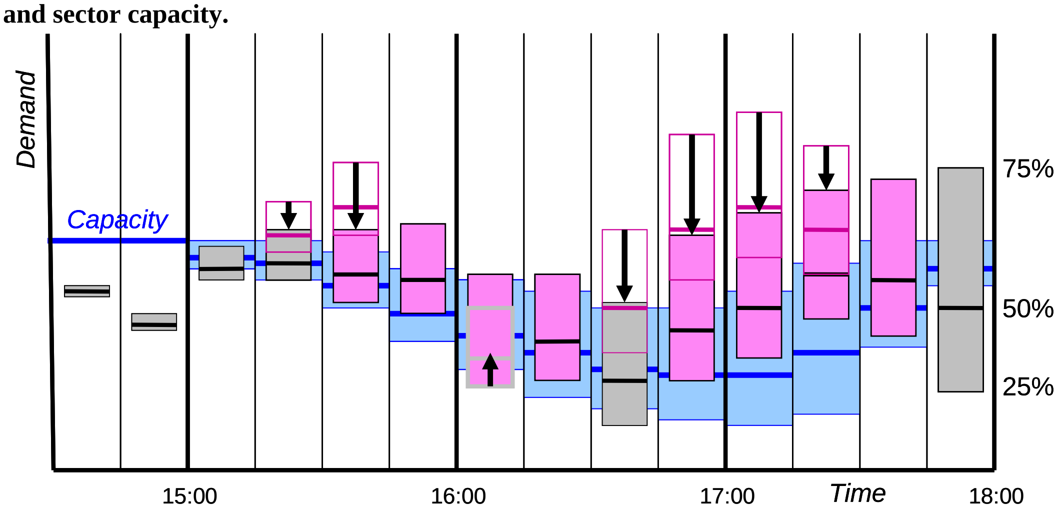

Figure 4. Reducing the probability of congestion to a maximum of 75%.

Once such a prediction is available, flight-specific maneuvers such as reroutes or ground delays can be selected to modify the traffic demand and reduce congestion risk. Figure 4 shows the effect of reducing the congestion risk to no more than 75% (again, for a single sector). For the work presented here, a relatively simple method is used to select maneuvers in a computationally efficient way. It can develop solutions to reach a target congestion risk, or even a time-varying target risk profile. The challenge, however, is to determine the appropriate target congestion risk at a particular time during a continually evolving traffic and weather situation. Thus, we formulate the problem as a sequential decision making process, in which we attempt to determine the appropriate target level of congestion risk for the current time, considering both prediction uncertainty and our opportunity to revise the strategy at a later time.

Figure 5 illustrates this process. 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

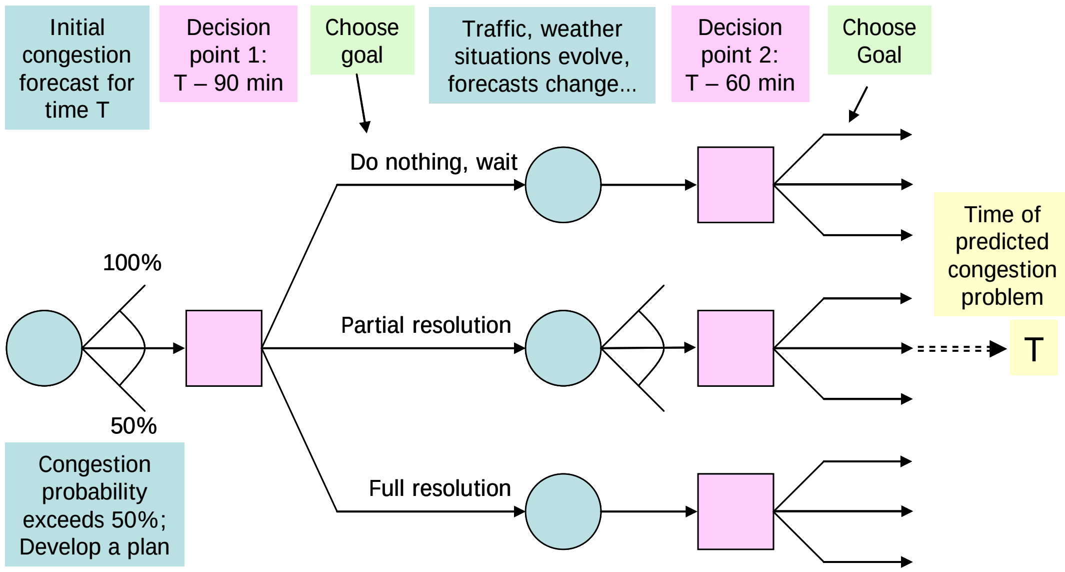

Figure 5. An Abstract Sequential Decision Tree for Congestion Management.

each decision path. The best path is that which strikes the desired balance between successfully resolving the congestion and incurring cost. Defining this balance is a challenge, since it is a mathematical representation of a traffic management objective, which is typically defined qualitatively.

IV. A Monte Carlo Simulation Method for Sequential Congestion Management

Tactical congestion management is a complex process, and complex to simulate. We are attempting to control aggregate quantities, specifically, balancing traffic loads against available capacities while minimizing delay and schedule disruption. However, we would like to control traffic loads efficiently, which requires flight-specific maneuvers. (ground delay, rerouting). Therefore, we must simulate predictions and actual outcomes for individual flights. This requirement 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 essential tradeoff in sequential decision making under uncertainty involves flexibility vs. knowledge. If we wait to make a decision, we will learn more about the possible outcomes and thus be able to make better decisions. However, some actions that were available earlier will no longer be available. For example, it is easier and less costly to ground delay a flight than to reroute it after it departs. So, a simulation of this concept must capture both the range of possible outcomes and the increase in knowledge as time progresses.

E. Overview

Our method begins by defining the congestion problem as (1) a set of ETMS-like trajectory predictions for traffic demand, (2) a nominal set of clear-weather sector capacities, and (3) an initial convective weather forecast. Next, we define the decision tree: a set of decision times and a set of congestion resolution goals at each of those times. The congestion resolution goals also include constraints on what kind of maneuvers are permitted, for example, whether or not we will allow rerouting of airborne flights to solve congestion.

The simulation is started by running a Monte Carlo simulation for both traffic and weather outcomes. The weather outcomes are converted to sector capacities using an empirical model. This provides the key advantage of this approach over the study of historical severe weather events, which can only have a single outcome. With the Monte Carlo simulation, we can explore the whole reasonable range of outcomes for a single scenario, and therefore determine what the best congestion resolution over the entire range of outcomes would be. Since historical weather events are as different as snowflakes, it is nearly impossible to gain statistical data on the goodness of a particular strategy by studying historical events.

Each possible path through the decision tree (decision path) is then evaluated for each of the Monte Carlo outcomes. Resolution maneuvers are computed to achieve the desired congestion risk (as in Figure 4) via a heuristic assignment method, based on closed-form statistical models of the uncertainty in demand and capacity predictions. These models are different from the Monte Carlo outcome models, since they must be computable in real-time. Rather than enumerating a series of outcomes, they compute the statistics of the aggregate prediction error in aircraft demand count and in sector weather coverage fraction. The weather coverage fraction is then converted to capacity using an empirical model. The resolution maneuvers are computed to achieve the target level of risk for all sectors in the desired region simultaneously, since each flight can affect several sectors within the congestion region.

As the simulation moves forward in time along each decision path for each Monte Carlo outcome, the predictions and traffic demand outcomes co-evolve. When resolution actions are taken, the traffic demand outcomes are modified to reflect those actions. Also, as time passes, the traffic and weather predictions improve based on the knowledge gained as events occur.

All decision paths converge at the last decision point, at which the remaining congestion is acted upon. Metrics are computed for each path to describe the impact and success of the resolution strategy.

F. The Decision Tree

Figure 5 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 6. 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 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 options and decision points can be increased at will, with only the cost of longer computation time.

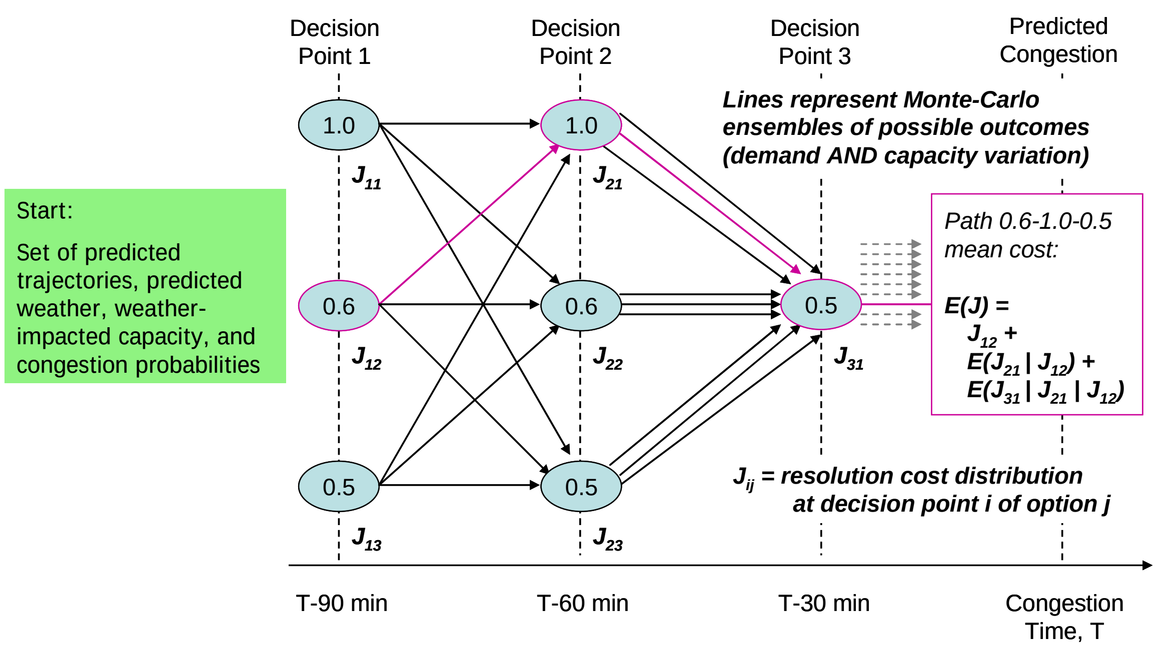

Figure 6. Three-Stage Decision Tree.

G. Modeling Traffic Outcomes

The initial traffic prediction is comprised of a set of predicted flight trajectories and departure times. Because we wanted to develop flight-specific resolution actions at each decision point, we developed a Monte Carlo model to simulate the possible "actual" flight trajectories that would be flown given such a traffic prediction.[25] 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"). Thus, for a single traffic prediction, the model will create N different outcomes. Each of these outcomes may contain a different number of flights (due to cancellations and pop-ups), and the flights will differ from the prediction due to the other prediction error distributions. These distributions were developed empirically, based on current prediction methods and procedures. If future ATM concepts are to be modeled, and these concepts would change the predictability of trajectories, the model distributions could be adapted to reflect that.

H. Sector Capacity Modeling

Before developing the Monte Carlo model for weather outcomes, we needed to understand how weather parameters can be related to what we actually need for the congestion calculation, namely, future sector capacity outcomes. For this study, empirical data was used to determine a simple relation between the fraction of a sector which is covered by VIL Level 3 or greater and the effective capacity of the sector.[19,20] Capacity is measured in maximum desired peak aircraft count, in the same way that MAP is defined.

Observations indicate that sector throughput declines linearly with coverage, and that sector capacity becomes effectively zero as the coverage fraction increases beyond 50%. Figure 7 illustrates this. This includes data for many sectors over the course of a summer, and the highest throughput values for each coverage level represent an "optimistic" capacity estimate for that level of coverage. While the slope of the reduction is different for different sectors, we chose to take the simplest possible approximation. Sectors are assumed to have a peak aircraft capacity equal to the MAP when the weather is clear. This capacity value is assumed to fall linearly with increasing coverage fraction until the coverage fraction reaches 50%, and at 50% and above, the capacity is assumed to be zero.

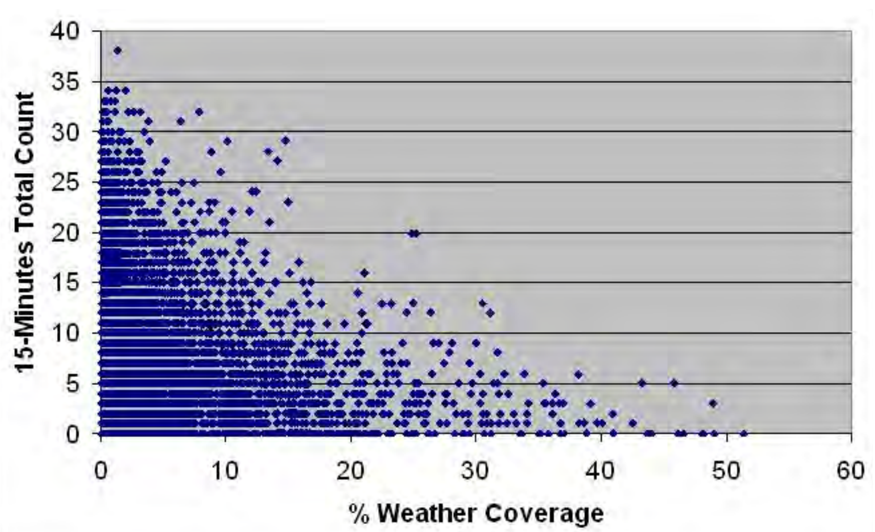

Figure 7. Sector throughput as a function of weather (VIL Level 3 or greater) coverage. From Davis et. al.[20]

This is a large simplification, since the position and shape of a weather event are important in determining where capacity will be lost. Song et. al.[15] and Martin[18] have proposed airspace capacity models that consider flow patterns and route organization, respectively. We may use such models in follow on work.

I. Modeling Weather Outcomes

Given this capacity model, we chose to represent weather as sector coverage fraction. Weather predictions are assumed to take the form of predicted sector coverage fractions as a function of look-ahead time. These time series are created for each sector of interest in the simulation, and referred to as “coverage traces.” A sector coverage trace can be converted to a sector capacity trace, via the algorithm described in the previous section. So, the Monte Carlo model needs to start with a nominal coverage trace and generate a set of “actual” coverage traces which represent a range of weather outcomes.

As noted earlier, some current weather forecast products contain measures of uncertainty. However, these are not easily adapted to computing the uncertainty in predicting sector coverage fractions. Thus, a simplified method of quantifying coverage uncertainty was adopted. Prediction uncertainties in the tactical TFM timeframe are assumed to occur in two independent forms: errors in predicting weather evolution speed, and errors in predicting the growth and decay of weather intensity.

Nominal traces for each sector are created either from an observed weather event (as in Figure 1) or synthetically, to represent a desired capacity reduction scenario. These traces are taken as the initial weather prediction for the scenario. Then, a Monte Carlo set of “actual” weather outcomes are generated by first computing speed and intensity variations from closed-form statistical distributions, and then dilating and scaling the nominal coverage traces to represent the “actual” coverage traces. Note that speed and intensity errors are assumed to be consistent over the airspace of interest. If weather worsens earlier than expected in sector A due to a speed forecast error, it also worsens earlier in adjacent sector B. This process is illustrated in Figure 8. The blue trace represents the nominal weather coverage, the red indicates weather which is both faster to evolve and more intense than predicted, and the green trace indicates weather that is less intense and slightly slower to evolve than expected.

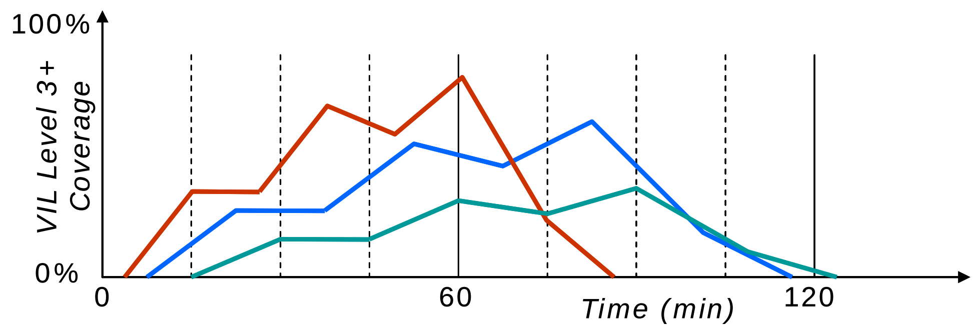

Figure 8. Speed and intensity variations on a weather coverage trace.

In the absence of empirical data on this process, which we plan to obtain and use in future work, we employ independent, zero-mean, symmetric triangular distributions for speed and intensity prediction error. The half-width of the distributions are varied to represent different levels of uncertainty. In reality, these and other parameters would be functions of the weather prediction model, storm type, etc., if indeed this is a reasonable way to represent prediction error in the tactical timeframe. For this study, we used three uncertainty levels: no variation (perfect forecast), moderate uncertainty, and high uncertainty. Moderate uncertainty parameters were set such that the maximum evolution speed error was ±25% and the maximum growth rate error was ±12.5%. The high uncertainty parameters were set at ±50% and ±25%, respectively.

J. Modeling Prediction Uncertainty for Congestion Resolution

As noted above, closed form, aggregate models for traffic demand and weather prediction uncertainty are needed to calculate the congestion resolution maneuvers. For traffic demand, we used a previously developed model.[26] 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.[5] 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). It treats demand in all sectors as independent, which is a simplification. However, this model is very fast to compute and can be used in either simulation or real-time applications.

For weather, we developed a similar model for this research. It computes a distribution of predicted weather coverage error based on the nominal prediction and look-ahead time. This was developed in a different way, however, since the Monte Carlo weather model described above does not correlate specifically with a real weather forecast of known accuracy. So, we developed a model based on the overall statistics generated by the Monte Carlo simulation, using the weather pattern from the example problem in Section V. It is based on simple, closed-form distributions of coverage prediction error, conditioned on the nominal predicted coverage for each outcome, the LAT, and the weather forecast uncertainty level (moderate or high). Both of these models predict aggregate quantities, but to be effective, they must approximate the statistics for traffic count and weather coverage that are produced by the Monte Carlo simulations. In the weather case, this is true by design. Assuming the simulations are good, they should also approximate real world statistics, and thus we are using them in parallel research into real-time decision support tools. We will see in the results that sometimes the aggregate models do not match the Monte Carlo models, degrading congestion resolution performance.

K. Congestion Resolution Algorithm

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[26], and has been shown to provide effective, though not optimal, flight-specific solutions.[27]

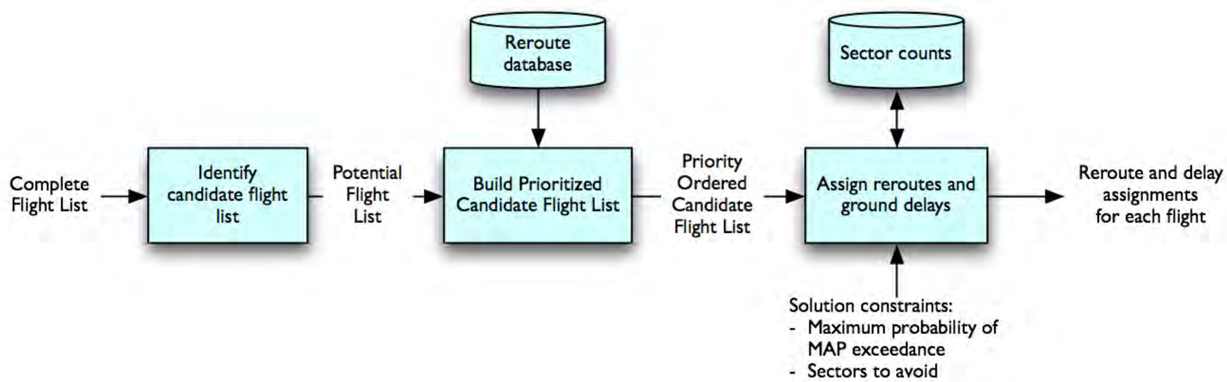

Figure 9. Overview of Heuristic Congestion Management Algorithm

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

Figure 9 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. A final modification for the sequential decision making was to take flights that had been maneuvered at a previous decision point and place them at the top of the priority list, to make it very unlikely that they would be maneuvered again; multiple maneuvers of a single flight causes schedule problems for operators.

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 are 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.

Several parameters can be adjusted to represent different solution constraints. These include whether or not to allow rerouting of airborne flights, the number of reroute options to explore, the maximum ground delay allowable, and the minimum time before planned departure that a flight can be delayed or rerouted.

L. Prediction Evolution

In order to capture the interesting features of probabilistic decision-making, we must simulate how the state of knowledge (i.e., the updated prediction) changes as simulation time passes. A single traffic and weather prediction exist at the start of the simulation. Many Monte Carlo outcomes are 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 traffic 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. Thus, we created models for how traffic and weather predictions "learn" from observed events.

The traffic prediction evolution model is simple, but realistic enough to generate interesting results. 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 may be replaced in future with more realistic, statistically-modeled behavior based on empirical studies.

Weather predictions are evolved by "inverting" the Monte Carlo weather variation model. Recall that the sector coverage prediction error is captured in two parameters (speed and growth). These are generated from the triangular error distributions for each outcome. As time advances along the outcome, the speed and intensity errors become apparent, and the prediction is "reset" at each decision point. The coverage at that time becomes the new baseline, and new speed and intensity values are computed from the error distributions to represent the new prediction as variations from the new observation. The weather prediction thus learns from the experienced weather, though not perfectly.

M. Evaluating the Decision Tree

Figure 10 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 and predicted weather coverage trace, which are used as a basis by the Monte Carlo traffic simulation to generate a set of N possible "actual" outcomes for the flights and sector capacities. These characterize the variety of ways that the situation can play out.

Next, decision point 1 (DP1) is evaluated (see the large brown boxes marked "*A*" in Figure 10). For each option, the initial prediction is used to compute a resolution action that meets the desired maximum congestion probability. The resolution impact is saved in several forms: number of affected flights, amounts of ground and airborne delay, and estimated additional operating cost. A new predicted flight 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 to adjust flight outcomes as a result of this step. If the resolution maneuver for a 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.

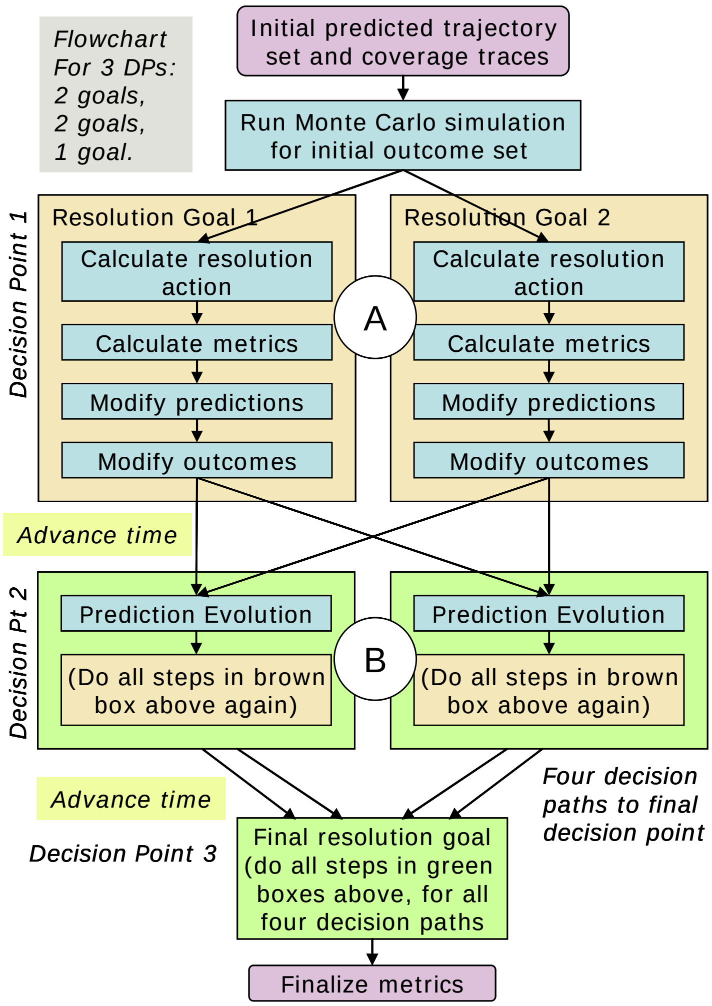

Figure 10. Simulation flowchart.

Once all options are calculated, time is advanced to DP2 (boxes marked "**B**"). 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.

N. 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 8 hours to complete on 8 dual-processor/dual-core systems (N = 1000).

V. Simulation Results for a Weather-Induced Congestion Problem

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. We ran the simulation for three levels of weather forecast accuracy and compared the effectiveness of the various decision paths.

O. Congestion Scenario

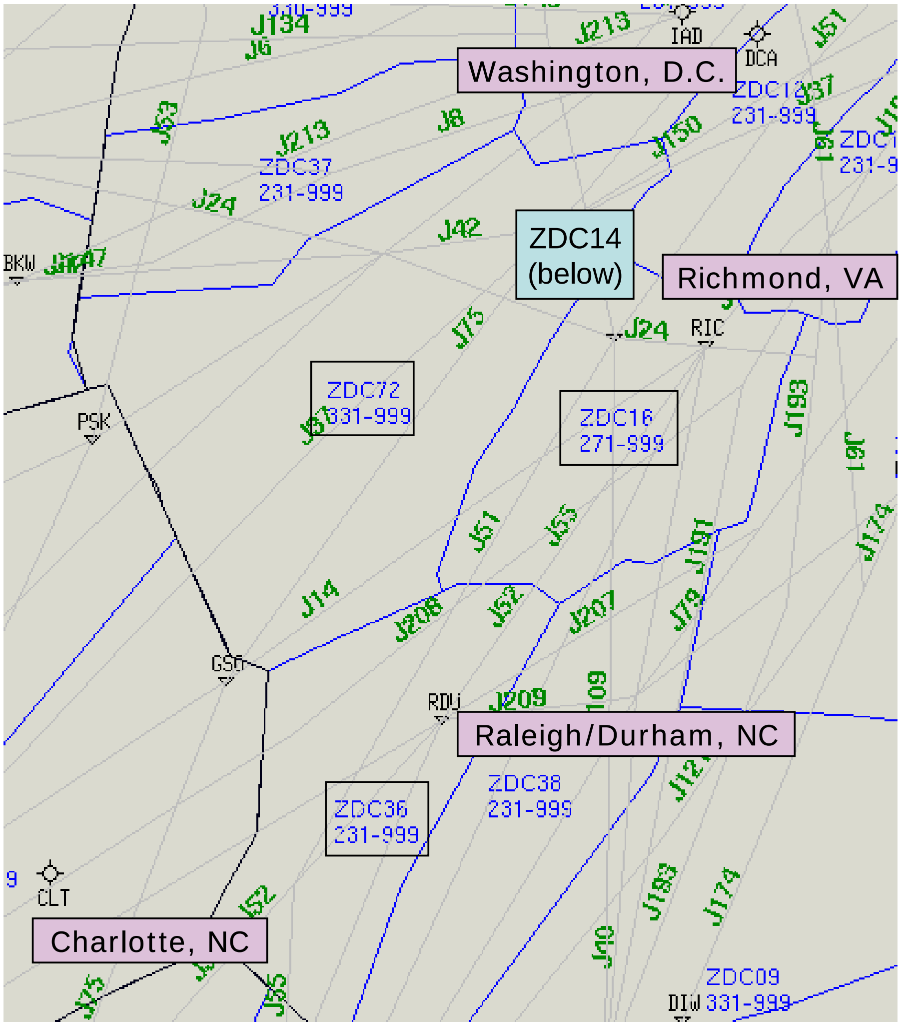

Figure 11. Congestion Scenario Airspace

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 (ZDC sectors 72, 16, and 36) are visible in Figure 11; the fourth, ZDC14, is a low altitude sector below 16 and 72. It is assumed that the predicted weather coverage trace for these sectors induces a capacity reduction of 5 below their nominal MAP values for the period 1800 to 2000 UTC, and zero reduction before and after those times. This is clearly simplistic, but adequate to simulate the effects of weather forecast uncertainty. 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, based on the nominal weather and traffic predictions, there is congestion predicted in the CRA sectors at 1700 UTC. Figure 12 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. As in Figure 3, dark purple boxes indicate a greater than 0.75 probability that the actual demand exceeds the sector capacity. Pink boxes indicate a greater than 0.50 probability, and boxes with a gray background indicate a less than 0.50 probability. The majority of the anticipated congestion lies in the time period 1815 to 1915 UTC, and this 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 6, with two exceptions. First, resolution goals of {1.0, 0.7, 0.6, 0.5} were available at decision points 1 and 2, producing 16 total decision paths.

	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 12. Predicted congestion at start of scenario.

Second, the congestion actually starts 75 minutes, rather than 90, after the first prediction (though the first red alert is 90 minutes later). The system is therefore tasked to resolve potential congestion from 1800 to 1930 UTC, which includes one period before and one after the severe congestion time period. Approximately 1500 flights pass through the CMA during or near this period, 191 of which also penetrate the CRA.

The resolution algorithm was asked to solve each congestion goal using ground delays of up to 30 minutes and/or by rerouting pre-departure flights on one of three alternate routes. No reroutes or ground delays were given to flights within 30 minutes of their planned departure time. These routes were chosen from an adapted route set which includes FAA-defined routes (e.g., Coded Departure Routes) and a selection of historical routes that were found to be often used. At least one of these routes was guaranteed to avoid the CRA (but not the CMA).

We determined from the statistical features of the output distributions that 1000 Monte Carlo outcomes were required to obtain a 95% confidence that the estimate of the mean number of aircraft affected by each resolution action was within one aircraft of the actual mean.[**]

P. Results for Perfect Weather Forecasting

Before we explore the impact of weather forecast errors, it is useful to compare the decision paths in the context of traffic demand uncertainty alone. These results are similar to those found in the previous study[3], but not exactly the same because of changes in assumptions and modeling. We have added constraints to implementation of resolution actions, as described in the previous section, and also provided more decision paths to choose from. Table 1 shows metrics for each of the 16 possible decision paths.

Table 1. Congestion Resolution Metrics for Perfect Weather Forecast Simulation

Decision Path	Number of aircraft affected		Minutes of Positive Delay		Direct Op Cost (dollars)		Cells over 50% congestion risk	
	Mean	StdDev	Mean	StdDev	Mean	StdDev	Actual	Predicted
0.5 - 0.5 - 0.5	90.95	9.90	604.08	88.09	17820.63	4164.32	2	0
0.5 - 0.6 - 0.5	86.27	9.01	582.52	83.42	17120.21	4019.76	2	0
0.5 - 0.7 - 0.5	85.30	8.29	579.03	81.57	17017.66	3962.54	2	0
0.5 - 1.0 - 0.5	85.00	8.01	577.57	80.75	16957.43	3922.90	2	0
0.6 - 0.5 - 0.5	58.79	10.57	276.10	86.98	9795.52	4211.34	3	0
0.6 - 0.6 - 0.5	51.41	9.47	232.77	80.77	8645.48	4069.04	3	1
0.6 - 0.7 - 0.5	49.37	7.77	222.71	75.34	8454.01	3861.36	3	1
0.6 - 1.0 - 0.5	48.97	7.16	221.34	73.81	8431.20	3790.22	3	1
0.7 - 0.5 - 0.5	53.00	10.94	224.82	87.33	8770.29	4294.59	2	0
0.7 - 0.6 - 0.5	44.07	10.73	178.35	84.33	7398.91	4222.60	3	0
0.7 - 0.7 - 0.5	39.65	8.08	161.87	73.97	6983.23	3767.42	4	1
0.7 - 1.0 - 0.5	38.26	6.06	157.24	70.64	6849.52	3623.53	4	1
1.0 - 0.5 - 0.5	53.19	11.87	264.11	94.64	9921.57	4308.37	2	0
1.0 - 0.6 - 0.5	41.91	11.28	196.00	84.49	8058.78	4178.33	2	1
1.0 - 0.7 - 0.5	33.94	9.02	162.94	69.04	7331.13	3479.90	3	1
1.0 - 1.0 - 0.5	28.91	5.20	147.23	60.57	7016.39	3133.03	4	3

Each row of the table contains the results of a single decision path, evaluated across the range of Monte Carlo outcomes. They are delineated by the three resolution goals encountered along the path, for example, 0.5 – 0.5 – 0.5 means that the congestion risk threshold was 0.5 at each of the three decision points. They are ordered from most aggressive solution (earliest resolution) to least aggressive (wait until the last moment).

There are three metrics related to impact on the airspace users: number of aircraft affected by the solution, minutes of positive delay, and approximate direct operations cost (DOC) in dollars. Positive delay is used so that credit is not given for rerouting an aircraft on a shorter path with an earlier arrival time. That situation is treated as "zero delay", since it is still a disruption, and users would generally like to fly their filed flight plan if possible. DOC is calculated from the assumption that airborne delays cost \$50 per minute and ground delays cost \$25 per minute. These values are subject to much variation (fuel prices, etc.) but were a reasonable approximation of costs in January 2008. In 2004, DOC was estimated by GRA, Inc.[28] for large aircraft as \$43/minute for airborne delay and \$23/minute for ground delay. Means and standard deviations are given for all of these metrics. The standard deviation is important, as it represents the range of cost outcomes and thus the predictability of the strategy.

[**] In the previous study[3], which used the same scenario but with perfect weather forecasting, only 250 evaluations were required. It is necessary to evaluate the required number of evaluations for each simulation situation, in order to properly balance the precision of the results against the computation time.

The last two columns represent the number of sector/time cells (as in the matrix in Figure 12) during the congested time period which still have a greater than 50% probability of congestion after the last decision point has been computed. For the "actual" column, this is taken from the Monte Carlo outcomes, and the probability is calculated simply by dividing the number of outcomes where a sector's capacity is exceeded during each 15 minute time period by the total number of outcomes. For the "predicted" column, this is taken from the final *predictions* after the last decision point has been computed, and thus represents whether or not the algorithm worked as it was designed to. The disagreement between the columns occurs when the predicted uncertainty distributions used in the resolution algorithm do not match the actual probability distributions resulting from the Monte Carlo simulation. For this scenario, there were two cases in which the ADM distribution did not agree very well with the Monte Carlo traffic simulation statistics, and so even when the predicted number of cells < 50% risk is zero, there are often two actual cells which were slightly greater than 50%.

The first four decision paths, which resolve to a congestion risk of 0.5 at the first DP, generate the most delay. In fact, they produce almost twice as much delay as any other strategy, though less than twice as much cost, since more of the delay is taken on the ground. They also solve all the predicted congestion (last column). The last path (in red), which attempts to solve all congestion at the last DP, leaves 3 cells unsolved on average. There are not enough viable maneuvers left at the last DP to resolve the congestion in most of the outcomes, so this is not an acceptable solution. The other paths involve partial solutions at one or both of the first two DPs, producing tradeoffs between cost and congestion resolution effectiveness. Details of the solution can be explored further to choose the best one (more on this later), but if we require all the predicted congestion and as many actual congestion cells to be resolved as possible (0 and 2, respectively) then 0.7 – 0.5 – 0.5 and 1.0 – 0.5 – 0.5 are good choices. The first of these (in green) incurs the lowest cost of all congestion solutions with 2 actual and 0 predicted unsolved cells; if we allow 3 actual cells to go unsolved, then 0.7 – 0.6 – 0.5 has an even lower cost. In general, if the situation is such that a certain amount of unresolved congestion is acceptable, lower cost solutions are available.

Q. Effect of Weather Forecast Uncertainty

We now explore how the decision strategies perform in the presence of weather (and capacity) forecast uncertainty. Recall that we chose two sets of parameters, one representing "moderate" uncertainty and one representing "high" uncertainty. When moderate uncertainty is simulated (Table 2), the results change in several ways. First, the cost increases for all decision paths where congestion is reduced to a similar level. The variability, as indicated by the standard deviations, also increases. For example, the 0.7 – 0.5 – 0.5 path is still a successful congestion solution, but at 6% higher mean cost and a 24% higher cost standard deviation than with perfect forecasting. Also, the last path is even less viable than before, since there are many outcomes where the weather is worse than anticipated and the remaining maneuvers at DP2 are inadequate to solve the congestion. If we define all paths with at least 2 unsolved predicted congestion cells as non-viable, there are now three non-viable paths. Paths which begin with a 1.0 goal are less effective and have much higher variability than before, suggesting that earlier action is advisable in the presence of weather forecast uncertainty.

Table 2. Congestion Resolution Metrics for Moderate Weather Forecast Uncertainty Simulation

Decision Path	Number of aircraft affected		Minutes of Positive Delay		Direct Op Cost (dollars)		Cells over 50% congestion risk	
	Mean	StdDev	Mean	StdDev	Mean	StdDev	Actual	Predicted
0.5 - 0.5 - 0.5	95.08	15.12	628.41	108.88	18609.11	4645.78	2	0
0.5 - 0.6 - 0.5	90.01	12.75	602.82	98.25	17798.38	4449.03	2	0
0.5 - 0.7 - 0.5	87.7	10.83	592.74	91.70	17488.56	4224.16	2	0
0.5 - 1.0 - 0.5	86.81	10.12	588.74	89.08	17349.98	4137.86	2	0
0.6 - 0.5 - 0.5	60.92	16.98	309.39	128.15	10798.49	5280.05	2	1
0.6 - 0.6 - 0.5	53.98	14.28	266.29	110.36	9682.21	4859.58	3	1
0.6 - 0.7 - 0.5	50.49	10.83	247.41	88.62	9253.27	4234.33	2	1
0.6 - 1.0 - 0.5	48.98	8.83	239.26	80.26	9055.98	3981.43	3	1
0.7 - 0.5 - 0.5	54.21	16.85	242.36	125.65	9297.88	5303.53	2	0
0.7 - 0.6 - 0.5	46.22	15.51	194.24	107.93	7884.58	4853.19	3	1
0.7 - 0.7 - 0.5	41.34	12.19	172.46	90.27	7308.44	4281.86	4	1
0.7 - 1.0 - 0.5	38.47	8.32	160.59	77.42	6974.02	3840.85	4	2
1.0 - 0.5 - 0.5	53.72	17.85	279.69	143.22	10443.21	5625.93	2	0
1.0 - 0.6 - 0.5	43.53	17.35	210.30	116.74	8540.96	5110.33	2	1
1.0 - 0.7 - 0.5	35.76	14.13	172.70	89.48	7624.35	4134.64	3	2
1.0 - 1.0 - 0.5	28.96	7.11	148.29	65.75	7075.59	3321.79	4	6

Table 3. Congestion Resolution Metrics for Large Weather Forecast Uncertainty Simulation

Decision Path	Number of aircraft affected		Minutes of Positive Delay		Direct Op Cost (dollars)		Cells over 50% congestion risk	
	Mean	StdDev	Mean	StdDev	Mean	StdDev	Actual	Predicted
0.5 - 0.5 - 0.5	101.97	20.77	705.47	161.11	20836.36	6177.4	1	0
0.5 - 0.6 - 0.5	95.91	17.83	667.66	127.60	19642.42	5193.02	1	0
0.5 - 0.7 - 0.5	92.97	15.24	652.21	109.92	19152.19	4698.97	1	0
0.5 - 1.0 - 0.5	90.55	11.77	640.22	94.88	18776.19	4273.21	1	0
0.6 - 0.5 - 0.5	66.8	21.33	363.60	177.40	12359.59	6499.98	2	1
0.6 - 0.6 - 0.5	58.96	20	306.39	151.17	10809.51	5827.82	2	1
0.6 - 0.7 - 0.5	54.33	16.73	275.06	125.79	10020.68	5188.22	2	1
0.6 - 1.0 - 0.5	50.47	10.16	250.13	85.50	9358.82	4170.18	3	1
0.7 - 0.5 - 0.5	60.51	22.23	317.12	177.38	11274.56	6519.24	1	0
0.7 - 0.6 - 0.5	51.18	21.58	252.22	149.70	9542.02	5838.52	2	2
0.7 - 0.7 - 0.5	45.01	18.03	215.55	124.22	8566.2	5024.09	3	3
0.7 - 1.0 - 0.5	39.23	9.79	184.72	81.17	7704.7	3956.19	4	3
1.0 - 0.5 - 0.5	56.95	22.99	323.51	187.86	11653.89	6670.74	1	2
1.0 - 0.6 - 0.5	46.04	23.56	239.98	159.23	9353.53	6097.64	3	2
1.0 - 0.7 - 0.5	37.57	19.96	190.05	126.44	8026.97	5113.36	3	4
1.0 - 1.0 - 0.5	28.43	9.08	148.67	70.64	6995.67	3513.62	4	6

The high variability case (Table 3) emphasizes this last point. In this case, the last 7 paths are non-viable. The 0.7 – 0.5 – 0.5 path is still the best, but with 29% higher mean cost and 72% higher cost standard deviation than for the perfect weather forecast case. The 0.6 – 1.0 – 0.5 path, although it does not completely resolve the congestion, does not see such a large increase in mean cost and variability (11% and 10% respectively). This further supports early action as desirable when weather uncertainty is large, especially if some residual congestion can be tolerated.

R. Defining Criteria for the "Best" Strategy

Summary statistics do not tell the whole story, since they obscure two important features of the solutions. First, they do not measure how severely capacity will be exceeded. Second, they do not show the variability in the results well, namely, how frequently across the outcomes sector capacities are exceeded, and what the range of possible cost values looks like. There are many ways to present this information, but we first need to decide what the best strategy looks like. This is an operational question, and we will seek input from operational traffic managers to determine the criteria. It might make sense to choose based on something like: "choose the lowest cost strategy which reduces the probability of exceeding capacity by 3 aircraft to 5% or less." This type of goal probably requires exploring a larger decision tree, including risk targets below 0.5, to assure success under high uncertainty situations. Figure 13 provides distributions of congestion outcome for sector ZDC14 at 1845 (at left) and resolution cost (at right) for two strategies under the moderate high weather forecast uncertainty conditions.

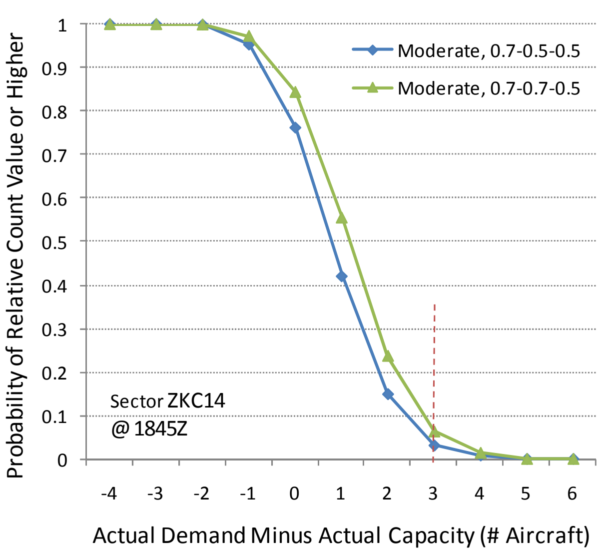

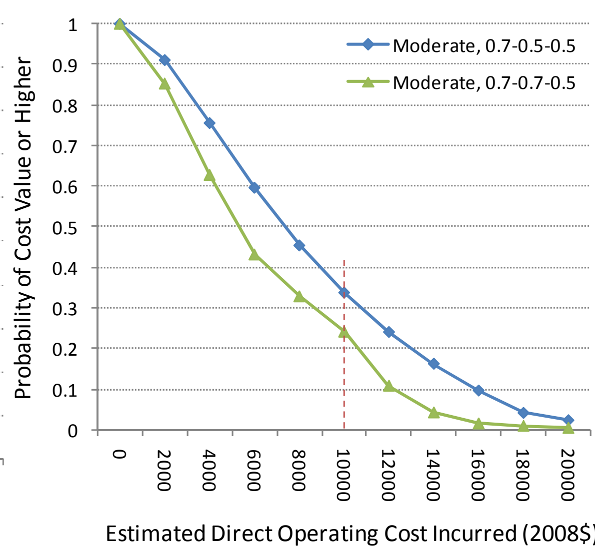

Figure 13. Congestion Resolution Cost and Effectiveness Distributions for Two Strategies.

The congestion outcome chart shows the probability distribution for the value of "actual" demand minus the "actual" capacity across the Monte Carlo simulation runs. The probability that the capacity will be exceeded by any amount is the probability that this difference will be one or more, which is 0.42 for the more aggressive 0.7 – 0.5 – 0.5 strategy and 0.55 for the less aggressive 0.7 – 0.7 – 0.5 strategy. In this case, the probability of exceeding capacity by 3 or greater is relatively low; 0.03 for the 0.7 – 0.5 – 0.5 strategy and 0.06 for the 0.7 – 0.7 – 0.5. This can be weighed against the cost profiles at right, which are plotted in the same way. The more aggressive strategy incurs cost of $10000 or greater 34% of the time, while the less aggressive strategy does this 23% of the time. Delay metrics or number of aircraft affected could also be considered. This illustrates the tradeoff between cost and congestion risk, and choosing the "best" strategy may depend on other factors as well.

S. Adaptive Behavior

We assumed in these simulations that when a flight is assigned a maneuver, it cannot be altered later. This was done because it is very disruptive to schedules to modify the same flight more than once. However, it also prevents adapting to situations where the weather is not as severe as predicted, and the operator is willing to "un-delay" a flight. There are several issues involved in using unexpectedly recovered capacity. The resolution developer sometimes generates maneuvers for some flights well in advance (perhaps an hour before departure). If the weather changes such that the maneuver is not necessary, these could be undone, with the cooperation of the flight operator. Alternately, maneuvers could be generated but not disseminated until the necessity for them is clear. A similar feature is used in today's Ground Delay Program (GDP) software, in which operators can provide an "earliest acceptable departure time" to allow some or all of the delay to be removed automatically. Reroutes could be similarly undone. It is possible to exploit this effect by re-prioritizing the flight list used in the heuristic solver (section K) such that flights with later departure times are pushed down the list and thus become more likely to be maneuvered. By building in deferrable maneuvers, the resolver is building in later flexibility, at a certain cost in schedule predictability. This change, or any significant change to the resolver algorithm, could produce a different "best" strategy. We will explore this in future work.

T. Continual Congestion Resolution

The decision tree and example problem explored here were aimed at deciding a goal sequence for congestion resolution in a static future time window. Operationally, congestion resolution is a continual process. Congestion may be present anywhere in the tactical flow management time interval (say, up to 2 hours in the future) and predictions of congestion will be constantly changing. To be applied operationally, the congestion resolution goal should be a profile of congestion risk targets as a function of LAT, rather than a single value. The decision to be made at each step is between profiles to use, i.e. between more aggressive and less aggressive profiles. This decision could be made either by running a variant of this simulation in real time, or by running a large number of cases in offline simulations and developing heuristics which map common situations to effective strategies.

VI. 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 real probabilistic weather forecasts and a better weather-impacted sector capacity model, 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. Also, the resolution strategy developer we used can be improved. As noted in section S, we need to consider flexibility when developing strategies, such that we can recover when weather does not materialize. Also, the algorithm is far from optimal, and we have done work[29] on genetic algorithms and a generalized random adaptive search procedure (GRASP) [30] to gain improved solutions.

The second application is to develop heuristics for near-term congestion resolution tools and procedures, as described in Section T. 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 sequential, probabilistic decision-making, as compared to today's approaches. We are also exploring, as presented here, the effect of weather forecast uncertainty on decision making and cost. Thus, the simulation is a platform for evaluating the potential tactical congestion management benefits from proposed probabilistic weather forecasting products, provided an operationally-acceptable model for weather-impacted sector capacity can be developed.

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.

VII. Conclusion

A Monte Carlo simulation technique for evaluating sequential, probabilistic decision-making in en route congestion management has been developed. A sample congestion problem, caused by severe weather, was explored using the simulation. Weather forecast uncertainty was shown to have a significant effect on when and how aggressively to act to solve the congestion problem. 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 tactical traffic management.

Acknowledgments

The authors would like to thank all who helped with this work, especially Stephen Zobell, Sandeep Mulgund, James DeArmon, Lixia Song, Neera Sood, and Claude Jackson for their help in both concept development and software implementation.

Notice

References

[1]Federal Aviation Administration (FAA), *Severe Weather and Route Management Handbook*, http://www.fly.faa.gov/Operations/Strategic_Planning/svrwx_handbook.html, 2006.

[2]Wolfson, M., Forman, B., Calden, K., Boldi, R., Dupree, W., Johnson, R., Wilson, C., Bieringer, P., Mann, E., and Morgan, J, "Tactical 0-2 Hour Convective Weather Forecasts for FAA," *Proceedings of the 11th Conference on Aviation, Range and Aerospace Meteorology*, Hyannis, MA, October 2004.

[3]Wanke, C., and Greenbaum, D., "Incremental, Probabilistic Decision Making for En Route Traffic Management", *7th USA/Europe Air Traffic Management R&D Seminar*, Barcelona, Spain, July 2007.

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

[5]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, August 2003.

[6]Meyn, L. A.. "Probabilistic Methods for Air Traffic Demand Forecasting," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Monterey, CA, August 2002.

[7]Mueller, K. T., Sorenson, J. A., and Couluris, G. J., "Strategic Aircraft Trajectory Prediction Uncertainty and Statistical Sector Traffic Load Modeling," *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Monterey, CA, August 2002.

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

[9]Gilbo, E, and Smith, S., "A New Model to Improve Aggregate Air Traffic Demand Predictions," *AIAA Guidance, Navigation, and Control Conference and Exhibit,* Hilton Head, SC, August 2007.

[10]Megenhardt, D., Mueller, C., Trier, S., Ahijevych, D., and Rehak, N., "NCWF-2 Probabilistic Forecasts," *Proceedings of the 11th Conference on Aviation, Range and Aerospace Meteorology*, Hyannis, MA, October 2004.

[11]Steiner, M., Bateman, R., Benjamin, S., Brown, B., Carmichael, B., Davidson, G., Krozel, J., Mahoney, J., Mueller, C., Pinto, J., Rasmussen, R., and Weygandt, S., "Integration of Probabilistic Weather Information with Air Traffic Management Decision Support Tools – A Conceptual Vision for the Future," *33rd Conf. on Radar Meteorology*, American Meteorological Society, Cairns, Australia, August 2007.

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

[13]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, June 2003.

[14]Song, L., Wanke, C., and Greenbaum, D., "Predicting Sector Capacity for TFM," *7th USA/Europe Air Traffic Management R&D Seminar*, Barcelona, Spain, July 2007.

[15]Song, L., Wanke, C., Greenbaum, D., and Callner, D., "Predicting Sector Capacity under Severe Weather Impact for Traffic Flow Management," *7th AIAA Aviation, Integration, and Operations Conference*, Belfast, Northern Ireland, 18 - 20 September 2007.

[16]DeLaura, R. and Evans, J., "An Exploratory Study of Modeling Enroute Pilot Convective Storm Flight Deviation Behavior," *12th Conference on Aviation, Range and Aerospace Meteorology*, Atlanta, GA, 28 Jan. – 2 Feb. 2006.

[17]DeLaura, R., and Evans, J., *An Exploratory Study of Modeling Enroute Pilot Convective Storm Flight Deviation Behavior*, MIT Lincoln Laboratory Project Report NASA/A-6, 26 June 2006.

[18]Martin, B., "Model Estimates of Traffic Reduction in Storm Impacted En Route Airspace," *7th AIAA Aviation, Integration, and Operations Conference*, Belfast, Northern Ireland, 18 - 20 September 2007.

[19]Zobell, S., Song, L., and Wanke, C., "Translating Weather Forecasts into Sector Capacity for En Route Traffic Management," Annual Meeting of the American Meteorological Society, 23rd Conference on IIPS, San Antonio, TX, January 2007.

[20]Davis, E., Hargroves, B., and Holden, T., "The Effect of Convective Weather on Sector Capacity," MITRE Product MP 05W0000200, The MITRE Corporation, McLean, VA, September 2005.

[21]Hoffman, R., Krozel, J., Davidson, G., and Kierstead, D., "Probabilistic Scenario-Based Event Planning for Traffic Flow Management," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Hilton Head, SC, August 2007.

[22]Mukherjee, M., and Hansen, M., "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, June 2005.

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

[24]Ramamoorthy, K., and Hunter, G., "Modeling and Performance of NAS in Inclement Weather," *6th AIAA Aviation, Integration, and Operations Conference*, Wichita, KS, September 2006.

[25]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 2005.

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

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

[28]GRA, Inc., "Economic Values for FAA Investments and Regulatory Decisions, A Guide," prepared for the FAA Office of Aviation Policy and Plans, Dec. 31, 2004. Web address: http://www.faa.gov/regulations_policies/policy_guidance/benefit_cost/media/050404%20Critical%20Values%20Dec%2031%20Report%20/07Jan05.pdf

[29]Sood, N., Mulgund, S., Wanke, C., and Greenbaum, D., "A Multi-Objective Genetic Algorithm for Solving Airspace Congestion Problems," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Hilton Head, SC, August 2007.

[30]Taylor, C. and Wanke, C., "A Generalized Random Adaptive Search Procedure for Solving Airspace Congestion Problems," *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Honolulu, HI, August 2008.