Deferability: A Concept for Incremental Air Traffic Management Decision Making

Stephen Zobell¹ and Craig Wanke² The MITRE Corporation, McLean, VA, 22102

Dealing with uncertainty poses a challenge for traffic flow management (TFM) decision making, for example, when rerouting flights to avoid forecast convective weather. These decisions impose costs and delays on flights, and weather does not always materialize where it was forecast or with the forecasted severity. Deciding on a good strategy to take now when faced with uncertain and probabilistic forecasts is difficult for human decision makers. We have developed the deferability concept as an approach for building TFM decision support systems that make decisions incrementally. Simulation studies show this concept can be an effective way to manage airspace congestion using uncertain traffic and weather forecasts.

I. Introduction

In incremental decision making, decisions are made knowing that further action will be taken at later times. This approach can improve decision making when forecasts are uncertain, by delaying some decisions to times when better forecasts could be available. If forecasts were perfect, then there would be no advantage to defer any part of the solution. But forecasts are imperfect, and accuracy tends to decrease with higher look-ahead times. Incremental decision-making automation will make decisions at periodic times throughout the day, such as every 15 minutes. These systems may only partially solve future congestion and weather impact at each decision time, knowing that the remaining problems will be managed at later decision times. The essential tradeoff in designing incremental decision-making strategies is to weigh the advantage of deferring decisions (better knowledge) against the disadvantages (e.g., fewer resolution options are available).

We have developed the deferability concept as an initial approach for building incremental decision-making systems that are fast enough to support real-time decisions. The deferability concept is aimed at solving tactical airspace congestion problems from 30 minutes to 2 hours into the future.

Deferability has been incorporated into a Monte Carlo simulation that is used to study incremental TFM decision support concepts. ¹ This simulation has been used to compare the effectiveness and costs of using the deferability concept with other incremental decision making approaches.

II. Capacity, Weather Impact, Demand and Congestion

Before automation attempts to solve future TFM problems it needs to quantify the problem. The approach we use is to define airspace resources, model the capacity of the resource, and predict the demand on the resource. A congestion problem exists when the predicted demand exceeds the predicted capacity.

We are using air traffic control sectors as the resources for determining capacity and demand. For now, sector capacity is defined using the existing Monitor/Alert Parameters (MAPs).² The MAP is a predefined number for each sector giving the number of simultaneous flights that can safely be in the sector in normal circumstances. The MAPs are not hard limits, and sectors can safely manage higher counts when traffic complexity is low. Ongoing research into sector workload and complexity could improve the sector capacity model.³⁻⁷ Also, the demand and capacity of other resources, such as route segments, may be included in future research.

The presence of weather in a sector tends to reduce the number of flights that can be managed. Predicting the impact of weather on sectors first requires a weather forecast. A 2-hour convective weather forecast from the Corridor Integrated Weather System (CIWS)⁸ was used in developing this concept. The future impact of weather on sector capacity is quantified as reductions in the sector MAP value at different look-ahead times (LATs). This reduction is calculated using a simple formula based on the forecast percent coverage of the sector by level 3 and

¹ Lead Software and Systems Engineer, M/S N450.

² Senior Principal Simulation and Modeling Engineer, M/S N450. Senior Member, AIAA.

above reflectivity. This approach is described in Ref 9, along with some more advanced techniques for predicting the capacity of sectors in the presence of forecast weather. Other researchers are also investigating advanced techniques for predicting the impact of weather on the capacity of sectors and jet routes.¹⁰⁻¹²

Sector demand is forecast using an Aggregate Demand Model (ADM).¹³ The input to this model is the predicted 4-dimensional trajectories of all flights in the National Airspace System (NAS). The ADM predicts the peak number of simultaneous flights in each sector for each 15-minute time period into the future. The ADM is a probabilistic model and produces the distribution of possible peak counts for each sector and each future time period. The

distribution accounts for demand prediction uncertainties such as errors in predicted departure times, cancellations, and pop-up flights.

Figure 1 shows the demand and capacity forecast for a single (ZID20). sector The sector capacity forecast is shown in blue. This sector has a MAP value of 13. Sector capacity is predicted to decrease over the next hour to a reduced MAP of 5 as the area of weather increases in the sector. For the second hour of the forecast, the area of weather in the sector is predicted to decrease. The MAP reduction forecast stops at 2 hours look-ahead, the limit of the CIWS weather forecast. The uncertainty of the forecast is shown by the light blue boxes, representing the 25% to 75% confidence interval for the reduced MAP prediction. The current system uses a rough estimate of forecast accuracy to generate this distribution. More



accurate Figure 1. Forecast capacity, demand, and congestion for one sector.

predictions of this distribution need to be developed based on the historical accuracy of the weather forecasts.

Figure 1 also shows the sector demand prediction produced by the ADM. These are the thin box plots in grey, violet, and magenta. Each box represents a 25% to 50% confidence range of the sector demand. The line through the middle of each demand box is the median of the distribution. Demand predictions are available out to 6 hours look-ahead.

Forecast congestion can be calculated by convolving the demand distribution and the capacity distribution for each 15 minute period. From this, the probability of congestion can be calculated. Figure 1 shows congestion probability as violet and magenta alerts when the probability of congestion exceeds 50% or 75% respectively. These alerted time periods are a concern because there is a danger the sector controllers may not be able to handle the demand if the congestion is not resolved.

III. Automated Heuristic Congestion Resolver

We have developed a heuristic algorithm for solving future congestion. It is a single-pass algorithm that processes flights using a priority ordering. This algorithm is fast enough to produce solutions in real-time, and also fast enough for use in Monte Carlo simulations.

The solutions are generally good but not optimal, and better solutions can be obtained using optimization techniques. However, these can be time consuming to compute. Capacity and demand forecasts can change while a solution is being generated, a disadvantage when solution generation takes a long time and the forecasts are changing rapidly. Long computation times can take away from the time available for evaluation and discussion of the solutions by decision makers.. Partial optimization approaches are being researched that can produce better solutions with only modest computational requirements.¹⁴

The congestion resolution algorithm needs a list of candidate flights that are eligible for maneuvers. Potentially this could be all flights in the NAS. To save time, we choose a smaller area using a polygon called the Congestion Resolution Area (CRA) that includes all the congested and weather impacted sectors that we want to solve. A candidate flight list is generated consisting of all flights that enter the CRA polygon, and only these flights are eligible for congestion resolution maneuvers.

A sector list is needed to determine which sectors will be managed by the congestion resolution algorithm. The algorithm will attempt to manage the existing congestion in these sectors, but also will prevent the creation of new congestion in any managed sector. A solution that just moved the congestion to sectors outside the CRA would be undesirable. To prevent this, the sectors in the sector list are chosen to cover a larger area, such as all sectors within 200 nautical miles of the CRA. This larger area is defined as the Congestion Management Area (CMA) The CMA may contain flights that do not enter the CRA polygon. The system knows those flights contribute to sector demand, but will not maneuver them.

Figure 2 is an overview of the heuristic congestion resolution algorithm. The algorithm takes the flight list and sorts the list into priority order. The sort order is as follows, listed from highest priority (earliest in the processing priority) to lowest priority:

- 1) Post-departure and previously maneuvered,
- 2) Post-departure and not previously maneuvered,
- 3) Pre-departure and previously maneuvered,
- 4) Pre-departure and not previously maneuvered.



Figure 2. Overview of the heuristic congestion resolution algorithm.

Within each category flights are ordered by their estimated time of arrival to the CRA. Post-departure (active) flights are given priority because they have fewer congestion resolution options (i.e. they cannot be ground-delayed). Flights that have been previously maneuvered are given priority in order to reduce the possibility that a flight that has been previously maneuvered will be issued a new maneuver, but this is still a possibility when no other options are available.

The algorithm removes all the flights in the flight list from the demand in the managed sectors. Each flight is then replaced one at a time in the priority order. If the addition of the flight to any managed sector it crosses does not create a congestion alert, then the flight is given the original route and departure time. If the addition of the flight causes congestion, the algorithm examines reroute and ground delay options for the flight and determines if any combination of ground delay and reroute is able to cross the managed sectors without creating congestion. If more than one reroute and ground delay option is viable, then the one with the earliest arrival time is chosen. These flights with viable reroute or ground delay options become the maneuvered flights.

Reroutes for pre-departure flights are found by searching for reroute options using historically-filed routes, Coded Departure Routes (CDRs), and National Playbook plays that use the same origin and destination airports. Generating reroute options for airborne flights is more complex. These flights will have to leave their current route at an upstream fix and join a route that is going to their destination. More advanced reroute option generation capabilities are also being researched, such as ad-hoc route generation algorithms to build custom reroutes to avoid weather and congestion.¹⁵

If no viable reroute or ground delay option is available for a flight that contributes to a congestion problem, the flight is given the original route and schedule. These are cases where the algorithm could not resolve the congestion.

Flights later in the priority list are more likely to be maneuvered, because earlier flights in the list will tend to fill up the capacity of the sectors, particularly sectors with reduced capacity. The capacity of a sector can go all the way to zero. This happens when the sector is predicted to have more than 50% of the sector area covered by level 3 and above weather. In these cases, the algorithm attempts to keep all flights out of the sector for the duration of zero capacity.

The results of the algorithm are a set of flight maneuvers to manage the congestion. Figure 3 shows the results of congestion resolution for the congested sector from Figure 1. For this sector the algorithm was able to manage all the congestion to below 50% probability by rerouting and ground delaying several flights. The algorithm managed several sectors in addition to this one.

A. Delay Recovery

A delay recovery algorithm is used to determine if flights that have been previously rerouted or delayed can be given resolutions with less delay. Delay recovery is possible when newer congestion forecasts are less severe than the previous forecasts, either due to less severe weather forecasts, or customer cancellations and reroutes that open up capacity.

mirrors the resolution algorithm in



The delay recovery algorithm Figure 3. Heuristic congestion resolution to 50% probability.

that it makes a list of candidate flights, prioritizes them, develops a set of recovery maneuvers for each flight, and tests each maneuver against the probability of congestion to determine acceptability. The mechanics for each of these steps are somewhat different, however. For now we limit delay recovery to pre-departure flights, even if the airborne rerouting option was available for resolution. This avoids the difficulty of deciding whether it makes sense to return an airborne flight to its previous route, and if so, designing a path that achieves this. Thus, we structured the algorithm as follows: A flight is a candidate for recovery if it (1) has a currently-assigned reroute or ground delay, (2) has not yet departed and will not depart for at least A minutes, and (3) is within B minutes of its original departure time. Parameter A is a constant in this simulation, but in practice could be replaced by a value provided by the operator of that flight, such as the "earliest possible departure time" that airspace users provide in the context of ground delay programs. Parameter B defines how early we are willing to start delay recovery; if B is too large, there is a risk that recovered flights might have to be delayed again. For simulation results in this paper, A is 15 minutes and B is 60 minutes.

Candidate flights are then prioritized for recovery based on the arrival time delay resulting from their currentlyassigned maneuver. The flight with greatest delay is evaluated for recovery first, then on through flights with decreasing values of delay. This is only one possible ordering. One alternative would be to have airspace users identify which flights would be most important to recover.

The best recovery option is to return the flight to its original route and departure time (full recovery). Partial recovery is also possible, by shortening the assigned ground delay (moving up the departure time), returning to the original route (if rerouted), or a combination of the two. A range of possible departure times is established by taking 5-minute intervals from the earliest to the latest possible departure time. The latest possible time is the currently-assigned departure time. The earliest is the latest of the original departure time or A minutes from the current time. A list of possible recovery actions generated, with the full recovery option (if available) at the top, and the rest of the route/departure time combinations following in order of predicted arrival time, from earliest to latest.

This list is then evaluated in order to determine if the recovery action would create congestion. Since the recovery action list is sorted in order of desirability, the first action which does not create congestion is selected and executed.

IV. Incremental Decision Making

In incremental decision making, the decision support systems will reevaluate the problem periodically to determine if any new action is needed now to solve future congestion. This reevaluation period would be based on factors like how quickly solutions can be generated and how much human interaction is needed to evaluate and implement solutions. In a fully automated system, reevaluation could be very frequent, such as every 5 minutes (the update rate of the CIWS weather forecasts). If humans have to initiate or review each solution, then reevaluation will need to be less frequent, perhaps every 15 or 30 minutes. A previous study¹ indicated that 15 minutes is an effective decision interval, and we use it here.

The heuristic congestion resolution solution described above is not tailored for incremental use, since the solution assumes nothing more will be done. Previous work¹ indicates that this can produce situations where flights are maneuvered multiple times before a solution is reached, which is not desirable. The relative costs and effectiveness of different incremental decision making approaches are presented later.

The incremental decision making process is illustrated in Figure 4. Flight-specific congestion management maneuvers are calculated and executed at each decision time, every 15 minutes in this case. The process then waits until the next decision time, and generates another set of congestion management maneuvers using updated weather and traffic forecasts.



Figure 4. Incremental airspace congestion management.

We have implemented two key incremental decision-making features for this study: congestion tolerance profiles, and deferability. Each of these approaches can be incorporated into the heuristic congestion resolver, and are fast enough to be used in real time with today's computers.

A. Congestion Tolerance Profiles

The congestion resolution solution shown in Figure 3 attempts to limit the probability of congestion to 50% and below. But the algorithm can also resolve to any probability of congestion tolerance, and can have different tolerance for each 15 minute LAT. For example, the following profile (Table 1) tells the resolver to manage congestion to 50% probability for up to 60 minutes into the future, 60% probability for the period from 60 - 90 minutes into the future, and 70% probability for 90 - 120 minutes into the future.

Tabl	e 1	l. A	4n	examp	le	conge	estion	to	lera	nce	pro	fil	e.
------	------------	------	----	-------	----	-------	--------	----	------	-----	-----	-----	----

0-15 min	15-30 min	30-45 min	45-60 min	60-75 min	75-90 min	90-105 min	105-120 min
0.5	0.5	0.5	0.5	0.6	0.6	0.7	0.7

Generally congestion tolerance profiles should have 0.5 at the shortest LATs, as long as the final goal is to eliminate congested sectors, and the definition of congested is probabilities greater than 50%.

Resolving congestion using this profile will tolerate some predicted congestion at higher LATs. Later incremental decisions will then attempt to solve this remaining congestion when the LAT to the congestion drops below 60 minutes and the congestion tolerance drops to 50%. Profiles that resolve to higher probability of congestion at longer LATs tend to reduce overall costs. ¹ Doing less congestion management using the less reliable longer look-ahead forecasts can reduce the number of unnecessary maneuvers.

The downside is that this can also increase the risk that the incremental process will not be able to eventually manage all congestion to below 50% probability. Consider a sector with 60% probability of congestion at 90 minutes LAT, where all the flights involved will depart within the next 30 minutes. This profile will wait 30 minutes before managing the sector to 50% probability of congestion. By then all the flights in the congestion will be active and will have fewer congestion resolution options. We have lost the ability to ground delay any flights, and the reroute options for each flight are likely to fewer than were available before departure. In some cases there is enough flexibility now to manage congestion to 50%, but later decisions will not have enough flexibility.

For delay recovery, the congestion tolerances are not allowed to exceed 50%. Allowing delay recovery at higher percentages such as 60% means that the delay recovery algorithm could move flights into sectors from which they will need to be removed later when the congestion tolerance goes back to 50%.

Congestion tolerance profiles are completely dependent on the prediction of the uncertainty and range of possible outcomes of forecasts. If forecasts are treated as perfect or deterministic, then managing to 70% will produce the same results as managing to 50%. Currently we do not have good techniques for predicting this uncertainty of weather forecasts. This is an area where more research is needed.

B. The Deferability Concept

The deferability incremental decision making concept was developed to make a more intelligent determination of which congestion resolution maneuvers should be made now, and which can be deferred to later decision times.

At each decision time, flights are divided into two categories: deferrable and non-deferrable. This division is based on the departure time for each flight and a constant called deferability time. For this discussion we use 60 minutes for the deferability time. Deferability times of 45, 60, and 75 minutes were evaluated in simulations to study how changes to this parameter impact the solutions.

Figure 5 illustrates the deferability time line for a single flight. Any flight where the departure time is more than 60 minutes in the future is deferrable, after that time the flight is non-deferrable.

While a flight is deferrable, the congestion resolution algorithm will defer decisions and not maneuver this flight. While deferrable, a flight could be subject to other TFM initiatives, such as Ground Delay Programs (GDPs) and Airspace Flow Programs (AFPs). At any time customers can also make proactive decisions, such as changing routes to avoid weather. To aid in making proactive decisions, customers should be warned when their flight is at risk for delays and maneuvers that are currently being deferred.



Figure 5. Deferability time line for one flight.

6 American Institute of Aeronautics and Astronautics

Deferrable flights are highly flexible, capable of delaying on the ground or implementing a wide range of reroutes. Deferring maneuvers for these flights helps preserve this flexibility for future decisions. This flexibility may be needed if the new forecasts unexpectedly predict higher congestion. Also, if the current forecasts have overestimated congestion, maneuvers that appear needed now for deferrable flights may become unnecessary.

When the departure time is less than 60 minutes in the future a flight becomes non-deferrable. In each incremental decision time after that point the flight is eligible for maneuvers and delays in order to manage congestion. For the flight's first non-deferrable incremental decision, the full range of potential reroutes and ground delays are possible. As time progresses this flexibility decreases.

The deferability concept has been incorporated into the heuristic congestion resolution algorithm through two changes. First, the priority sort order is modified to include the deferrable classification for each flight. The sort order is as follows, listed from highest priority (earliest in the processing priority) to lowest priority:

- 1) Non-deferrable and post-departure and previously maneuvered,
- 2) Non-deferrable and post-departure and not previously maneuvered,
- 3) Non-deferrable and pre-departure and previously maneuvered,
- 4) Non-deferrable and pre-departure and not previously maneuvered,
- 5) Deferrable and previously maneuvered,
- 6) Deferrable and not previously maneuvered.

Within each category flights are ordered by arrival time to the CRA.

Second, the congestion resolution maneuvers are generated based on this priority order, but only the maneuvers for non-deferrable flights are included in the maneuver flight list.

These changes mean that deferrable flights are not added to the predicted demand when managing congestion using non-deferrable flights. Non-deferrable flights only compete with other non-deferrable flights for capacity, and only the maneuvers for non-deferrable flights are used.

Figure 6 shows congestion resolution using deferability for the congested sector from Figure 1. The demand for all flights (deferrable and non-deferrable) is shown. Congestion is resolved until 17:15, but very high congestion from 17:30 to 18:00 is being allowed and deferred for later decision times.

The situation in Figure 6 appears very risky. But the deferability solution has retained sufficient flexibility to manage this congestion later. Figure 7 shows the sector congestion if the congestion resolution maneuvers for deferrable flights are added to the solution in Figure 6. This shows that if the forecasts do not change radically, maneuvers to deferrable flights at later decision times should be able to manage the congestion that is being tolerated now. These deferrable flight maneuvers are calculated but are not issued now, because changes to the forecasts could invalidate those maneuvers. Knowledge of these deferrable flight maneuvers could be useful, for example, to warn customers of flights that are at risk for future maneuvers.



Figure 6. Congestion resolution using deferability.

The results of congestion resolution using deferability exhibit features that are similar to congestion tolerance using profiles that allow higher congestion at longer LATs. But the deferability resolution is more intelligent, only tolerating congestion that should be resolvable at later decision times.

C. Deferability and Congestion Tolerance Profiles

Deferability and congestion tolerance profiles can be used simultaneously. The ability to vary the congestion tolerance profile provides a way to change the balance between the overall costs of the incremental solutions (the costs associated with delays and reroutes) and the risk of not being able to eventually manage all the congestion. Profiles that percentage contain higher tolerances at longer LATs tend to have lower overall cost, but higher risk of unmanaged congestion.



Figure 7. Congestion resolution using deferability plus deferrable flight maneuvers.

Congestion tolerance profiles that are below 50% at higher LATs may also be useful with deferrable solutions. Deferrable solutions already act similar to solutions using profiles with high percentages at longer LATs. Adding a profile with lower than 50% tolerance at longer LATs could offset this, increasing the proactive congestion management at longer LATs. This risk management approach can decrease the possibility of unmanaged congestion, but at the expense of higher costs. Table 2 shows an example of a decreasing tolerance profile:

0-15 min	15-30 min	30-45 min	45-60 min	60-75 min	75-90 min	90-105 min	105-120 min
0.5	0.5	0.5	0.5	0.49	0.48	0.47	0.46

 Table 2. A decreasing risk tolerance profile for use with deferred maneuvering.

Some of the added cost of using decreasing risk profiles for one incremental decision can be recovered at later decisions. These profiles can cause more maneuvers for non-deferrable flights at the current decision time, decreasing the demand on congested sectors. This decrease could be needed later if the forecasts have underestimated the eventual weather impact. If the forecasts are accurate or overestimate the weather impact, then later decisions may be able to use that available capacity for flights that are deferrable now. This can mean fewer maneuvers at the later decision times. Decreasing risk profiles are thus more proactive, and could also shift some of the congestion management burden from short-haul flights to long-haul flights.

Having the ability to use different profiles could allow traffic managers to tune congestion resolution in realtime. If they feel they are facing weather and congestion that is more perilous than usual they can request more aggressive risk management solutions. The system could then use a congestion tolerance profile with lower risk targets, increasing the costs but decreasing the risk of unmanaged congestion.

The challenge in evaluating the effectiveness of profiles is quantitatively evaluating congestion risk. It is insufficient to study past weather events to do this, because an actual event played out to a single conclusion, and there are not enough truly similar weather events to compare the potential effectiveness of a new strategy in a statistically-valid way. Thus, we used Monte Carlo simulation to test and refine the proposed strategies.

V. Monte Carlo Simulation

A Monte Carlo simulation has been developed for researching incremental decision making concepts. This simulation is described in more detail in Ref 1. This simulation can be used to evaluate and compare different incremental decision making approaches. The key benefit of Monte Carlo simulation is that many outcomes of a predicted congestion situation can be generated, hopefully spanning the complete range of possible outcomes, and statistical measures of success can be computed. We used it here to evaluate the deferability concept and a range of congestion tolerance profiles. Other parameters were also varied, such as the deferability time and whether airborne rerouting is allowed in solutions.

A. Simulated Incremental Decisions

For the simulation results presented here, the Monte Carlo simulation attempts to resolve congestion over a 2 hour time period. The simulation makes congestion management decisions for 7 decision times, covering the first 105 minutes.

Figure 8 illustrates the simulation flow. An initial set of flight trajectories and weather coverage forecast are specified before the simulation. At the first incremental decision time the simulation uses the initial flight trajectories and weather forecast to calculate the demand and capacity of heuristic congestion sectors. The resolution algorithm is then used to congestion generate resolution maneuvers. These solutions use the specified incremental strategy, including the congestion tolerance profile settings and whether deferability is used.

The simulation then applies randomly-generated modifications to the traffic and to the weather, generating an ensemble of possible "actual" traffic and weather outcomes. Details of these modifications are given later. These variations simulate the fact that the traffic and weather do not play out over time in exactly the way the forecasts predict at the current decision time.

The simulation then advances to the next decision time, modifies the predictions based on what was learned about the traffic and weather over the interval. and calculates another incremental congestion management solution using the modified weather and traffic predictions. Delay recovery is also applied to determine if any previous maneuvers can be undone or reduced. This loop is repeated until all decision times are processed. This is done along each of the "actual" traffic and weather outcomes, and taken together, represents



Figure 8. Monte Carlo simulation flow chart.

the range of possible outcomes achievable by a given incremental decision-making strategy.

Information about the flight maneuvers from each decision time is saved for later analysis of the costs of using this incremental congestion management approach. Information about any sectors that have unresolved congestion

9 American Institute of Aeronautics and Astronautics after the last decision time at which it could be resolved is saved for analysis of the effectiveness of the approach. Metrics are assembled for all of the simulated outcomes (1000 in this study) for statistical analysis.

B. Modifying Traffic

The Monte Carlo simulation creates random modifications of demand for each incremental decision. These represent the types of unexpected changes that can happen to our knowledge of flights and flight plans as time progresses from one incremental decision to the next. The following types of modifications can be made to the flight trajectories:

- Flight cancellations,
- Flights departing earlier or later than predicted,
- Changes to the routes of flights,
- Changes to cruise altitudes,
- Active flights progressing faster or slower than expected,
- Unexpected new flight plans ("pop-ups").

The number and magnitude of these variations are scaled to simulate the types of variations that have been observed in the NAS. This is based on an analysis of several months of archived flight plan and track data.¹³

The Monte Carlo simulation applies these random flight modifications after each incremental decision. New 4dimensional trajectories are calculated for each modified flight, and this is used to generate new sector demand forecasts using the ADM. New track positions for active flights are also generated from the 4-dimensional trajectories, including flights that depart during the simulation.

C. Modifying Weather

The Monte Carlo simulation models the uncertainty of weather forecasts by varying the forecast between each incremental decision. This simulates the fact that weather does not always evolve exactly as forecast.

The approach used is to vary the graph of the forecast coverage of Vertically-Integrated Liquid (VIL) level 3 and

above weather for each sector. Figure 9 shows an example of the graph of the forecast weather coverage over 2 hours for one sector. The blue trace is the coverage of the sector using the initial weather forecast at the start of the simulation. Variations in the weather are generated by applying scaling factors to the time and coverage axis to produce new coverage traces.



Figure 9. Speed and intensity variations on a weather coverage trace.

Scaling time is analogous to variations in the speed of storm movement and development. Scaling the coverage is analogous to variations in storm size and intensity. The red trace represents a variation where the storm moves faster and develops more intensity than the initial forecast predicted. The green trace represents a variation where the storm moves slower and develops less intensity than the initial forecast predicted.

Random coverage trace variations are generated for each sector. These new coverage forecasts are used to calculate new sector capacity forecasts, based on a simple rule which assumes that capacity decreases linearly from 100% of MAP with no severe weather coverage to zero capacity with 50% coverage. The capacity forecasts are convolved with the demand forecasts to produce probability of congestion forecasts, which are then used for generating congestion resolution and delay recovery maneuvers.

Two models for weather variability are used for the studies in this paper. The moderate uncertainty model limits time variation to $\pm 25\%$ and intensity variation to $\pm 12.5\%$. The high uncertainty model limits time and intensity variation to $\pm 50\%$ and $\pm 25\%$, respectively.

Another approach for generating weather forecast variations is to use ensemble or scenario forecasts where multiple weather forecasts are generated each representing a different possible outcome. Weather researchers are studying these techniques, ^{16,17} and this could be considered for use in future research as these techniques mature.

VI. Congestion Scenario

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 10; the fourth, ZDC14, is a low altitude sector below 16 and 72. This area was designated as the CRA. The managed sector list contains 38 sectors, including of all sectors adjoining the CRA either laterally or vertically.



Figure 10. Congestion scenario airspace.

It is assumed that weather moves from west to east through the CRA such that coverage in each sector can be approximated as a half-sine wave (Figure 11). The solid curves indicate the weather coverage percentage as a function of time for the four CRA sectors. The corresponding sector capacities are shown by the dashed lines (referenced to right Y axis). Each sector capacity drops from the nominal value to approximately half that value, and recovers about 3 hours later. It is assumed that the other managed sectors are not impacted by weather.

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 1, magenta boxes indicate a greater than 0.75 probability that the actual demand exceeds the sector capacity. Violet 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 1914 UTC, and this represents a serious congestion situation that needs to be resolved. During this period, approximately 1500 flights pass through the managed sectors, of which about 200 pass through the CRA and are eligible for maneuvers.



Figure 11. Predicted weather coverage and resulting sector capacities at start of scenario.

	1700	1715	1730	1745	1800	1815	1830	1845	1900	1915
ZDC14 [14]	9	12	8	6	6	9	10	12	11	10
ZDC16 [15]	7	7	7	10	9	10	11	11	5	5
ZDC36 [16]	6	10	8	14	9	13	15	9	6	7
ZDC72 [18]	16	19	12	8	15	15	18	19	13	13

Figure 12. Predicted congestion at start of scenario.

D. Risk Management Strategies

The risk management profiles evaluated in this scenario are summarized in Table 3. The first four profiles are used with the basic resolution strategy, and increase in aggressiveness from "BasicIncHigh" (basic strategy, increasing to a high risk tolerance) to "BasicLevel" (basic strategy, resolve all LATs to 0.5 probability of congestion). Three profiles are used with the deferred resolution strategy, and include an increasing risk profile ("DeferInc"), a level profile ("DeferLevel"), and a decreasing risk profile ("DeferDec"). The deferability time is another important parameter for the deferred resolution strategies, and three values were used: 45, 60, and 75 minutes. Thus, there are four basic strategies and nine deferred resolution strategies for a total of 13 strategies.

Profile ID	0-15	15-30	30-45	45-60	60-75	75-90	90-105	105-120
	min	min	min	min	min	min	min	min
BasicIncHigh	0.5	0.6	0.65	0.7	0.75	0.8	1.0	1.0
BasicIncMed	0.5	0.5	0.6	0.6	0.7	0.7	0.8	0.8
BasicIncLow	0.5	0.5	0.5	0.5	0.6	0.6	0.7	0.7
BasicLevel	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
DeferInc	0.5	0.5	0.5	0.5	0.52	0.54	0.56	0.58
DeferLevel	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
DeferDec	0.5	0.5	0.5	0.5	0.49	0.48	0.47	0.46

Table 3.	Selected	risk	management	profiles.

Each strategy was run for a series of decision points starting at 1700 and ending at 1830, with a LAT of two hours, using a 15-minute decision point intervals for a total of 7 decision points. Pre-departure flights could be rerouted and/or assigned up to 30 minutes of ground delay, and maneuvers could be issued up to 15 minutes before their departure time. Airborne routes could be assigned after takeoff; scenarios were run with and without airborne rerouting as an available strategy. For either ground or airborne reroutes, up to four reroute options were available.

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 are frequently used. At least one of these routes was guaranteed to avoid the CRA, but can pass through other managed sectors. Delay recovery was also possible for all strategies, starting 60 minutes before the planned departure time.

We tested the strategies under two different weather forecast uncertainty levels (moderate, high) and with and without airborne rerouting available, yielding 52 total cases. In today's NAS, it is much easier to assign ground delays and reroutes than it is to assign airborne reroutes, so it is interesting to observe the relative benefit of allowing airborne rerouting as a strategy.

The statistical features of the output distributions indicated 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.

E. Results: Ground-Based Resolutions

Table 4 shows the results of the simulation runs without airborne rerouting in terms of mean metric values; being a stochastic simulation, all the metrics are distributions, and the nature of those distributions is important and will be discussed later. For both congestion resolution and delay recovery actions, there are three metrics related to impact on the airspace users: number of flights affected by the solution, minutes of delay, and approximate direct operations cost (DOC) in dollars. Note that, for resolution actions, only positive delays are included 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 are a reasonable approximation of costs in January 2009. In 2004, DOC was estimated by GRA, Inc.³¹ for large aircraft as \$43/minute for airborne delay and \$23/minute for ground delay.

	Defer	Masthau	Mean delay metrics			Mean re	covery m	etrics	Means of other metrics				
Strategy	Time (min)	forecast unc.	No. of flights	Delay (min)	DOC (\$)	No. of flights	Delay (min)	DOC (\$)	MTBD (min)	Delays/ affected flight	Recovery percent.	Unsolved congestion	
BasicIncHigh	N/A	Mod	96	466	19590	5	-26	-189	53	1.21	4	3	
		High	101	558	21588	5	-28	-262	53	1.25	3	3	
BasicIncMed		Mod	118	614	23477	8	-34	-257	57	1.24	4	2	
		High	132	794	27876	9	-43	-534	56	1.32	4	2	
BasicIncLow		Mod	146	807	28094	13	-67	-857	58	1.27	5	2	
		High	171	1186	37543	15	-81	-1281	59	1.40	5	2	
BasicLevel		Mod	219	1478	45602	94	-685	-13735	60	1.41	19	0	
		High	260	2107	61705	101	-880	-18669	60	1.59	18	0	
DeferInc	45	Mod	85	543	18330	2	-23	-519	33	1.14	1	2	
		High	97	772	24233	3	-29	-666	34	1.21	1	2	
	60	Mod	116	811	26076	6	-58	-1312	42	1.19	1	2	
		High	136	1158	34996	7	-67	-1521	43	1.29	2	2	
	75	Mod	154	1062	33406	14	-109	-2147	51	1.26	4	2	
		High	179	1493	44482	15	-128	-2661	52	1.36	3	2	
DeferLevel	45	Mod	88	593	19683	2	-28	-647	33	1.14	1	2	
		High	103	865	26711	3	-34	-789	34	1.23	1	2	
	60	Mod	125	904	28529	8	-78	-1764	43	1.20	2	2	
		High	146	1302	38737	9	-85	-1937	43	1.31	2	2	
	75	Mod	166	1232	37733	22	-186	-3744	52	1.26	6	2	
		High	195	1722	50177	23	-210	-4485	52	1.39	5	2	
DeferDec	45	Mod	91	618	20396	3	-32	-748	33	1.14	1	2	
		High	106	911	27925	3	-42	-973	33	1.24	1	2	
	60	Mod	131	955	29822	10	-95	-2168	42	1.20	2	2	
		High	153	1387	40840	11	-106	-2427	42	1.31	2	2	
	75	Mod	175	1320	39868	29	-243	-5036	51	1.27	6	2	
		High	207	1886	54321	30	-288	-6253	51	1.41	6	2	

Table 4. Strategy evaluation metrics: no airborne rerouting allowed.

Three other metrics measure the characteristics of the resolution and recovery activity. Mean time before departure (MTBD) gives the average time before the original departure time that a flight received its first delay maneuver. Longer times are better for predictability, except when the flight is re-maneuvered later for resolution or delay recovery. Delays-per-affected-flight indicates how often a flight that was maneuvered once was acted upon again. Since we allowed a maximum of 30 minutes of ground delay in a single action, a non-trivial number of flights needed to be affected more than once. The recovery percentage indicates what percentage of the affected flights received a full or partial delay recovery action.

Finally, the unsolved congestion column indicates how many CRA sector intervals (i.e., cells in the matrix in Figure 12) had a congestion risk greater than 0.5 when the last decision point which could solve them was completed. This means that for each decision point at 1730 or later, the ensuing 15 minute periods (e.g., 1730-1759 for the 1730 decision point) for the CRA sectors were examined, and any periods with congestion risk greater than 0.5 were added to the unsolved list. This is a simplistic way to evaluate the congestion effectiveness, because the value of the final risk is equally if not more important than the fact that it exceeds 0.5. A more detailed metric to capture this feature is discussed later.

Some clear trends are evident in the results. First, let us examine the basic strategies. More aggressive strategies delay more flights and incur more cost, but do a better job resolving congestion. The mean cost of resolving the problem, for moderate weather forecast uncertainty, ranges from about \$15,000 to \$35,000 (if we sum the DOC means for delay and recovery actions). Also, more aggressive resolution action leads to more delay recovery actions, though this is not a strong effect until the most aggressive strategy is employed (BasicLevel). More aggressive strategies are also more likely to maneuver a flight more than once (delays per affected flight), which NAS customers find undesirable as it adversely affects schedule integrity.

The unsolved congestion periods metric works well enough to stratify the results into three groups. Two of the sector periods were extremely difficult to solve, given the option space and prediction errors, and could only be solved with very aggressive action. Thus, only the BasicLevel strategy succeeded in doing so, at a large cost. BasicIncMed and BasicIncLow solved everything but these two periods, and BasicIncHigh left an additional period unsolved. If we choose "no more than 2 unsolved periods" as our congestion effectiveness target, then BasicIncMed is the best basic strategy for both moderate and high levels of weather forecast uncertainty. This is too simplistic, however, and will be revisited later.

All of the deferred resolution strategies work equivalently well by the unsolved periods metric, so more detailed study is required to pick the best. Note that the lowest-cost option, DeferInc with deferability time (DT) of 45 minutes, is better than the BasicIncMed strategy in almost all ways. It incurs lower net cost, affects fewer flights, causes fewer recovery actions, and has a lower re-maneuver rate (14% vs. 24% for BasicIncMed). In general, the deferred strategies by their very nature will minimize recovery maneuvers, because maneuvers are not issued until they are deemed necessary. This produces a difference in the time a flight is first affected, shown by the MTBD metric. For the basic strategies, it appears that flights will be first maneuvered on average between 50 and 60 minutes before their planned departure time. This value decreases in the deferred strategies, as expected, and is governed by the DT parameter. For DT=45 minutes, the average MTBD is only 33 minutes. Having more decisions that occur closer to departure time appears to be an undesirable effect. However this disadvantage is offset by the fact that deferred solutions tend to produce fewer maneuvers overall and very few recovery maneuvers, and require a smaller fraction of flights to be re-maneuvered.

Increasing the weather forecast uncertainty drives up the mean cost of solving the problem and slightly increases the need for delay recovery maneuvers, but does not strongly affect congestion resolution effectiveness because the congestion resolution algorithms are explicitly accounting for the level of uncertainty in the calculations. The weather forecast uncertainty has a larger effect on the variability of the measures, as shown in Figure 13 by the error bars, which show plus and minus one standard deviation. The ability to measure this effect has useful implications for testing the utility of new, probabilistic weather forecast products.



Strategy (no airborne rerouting)

Figure 13. Variability of net mean cost across strategy and weather forecast uncertainty level.

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 possible metrics for this kind of information, but we first need to understand how to define the "best" strategy looks like. This is an operational question, and requires input from operational traffic managers to determine the precise criteria, but it is likely that they will be concerned with the likelihood of bad outcomes. This kind of goal can be expressed in natural terms, for example: "Choose the lowest cost strategy which reduces the probability of exceeding capacity by 3 or more flights to 10% or less." Monte Carlo simulation is well-adapted to answering this type of question, as the likelihood of such events can be directly computed.



Figure 14. Risk of severe congestion vs. net mean cost for strategies without airborne rerouting.

Figure 14 illustrates the relationship between this type of goal, the incurred cost, and the weather forecasting uncertainty. The metric shown on the vertical axis represents the percentage of Monte Carlo outcomes in which one of the four CRA sectors has a peak count of 3 or more aircraft above its weather-impacted capacity during the 15

15 American Institute of Aeronautics and Astronautics minute period following the last decision point in the simulation. This represents the effectiveness of the strategy in reducing the risk that significant congestion occurs at this point. For reference, this value is 59% if no resolution action is taken at all. This metric is plotted against the net mean cost (the resolution cost minus the recovered cost due to delay reduction) on the horizontal axis, and all 13 combinations of strategy (no airborne rerouting), DT, and weather forecast uncertainty are plotted. The red circles are for those combinations under high weather forecast uncertainty, fitted by a red regression line ($R^2 = 0.92$). Similarly, the blue triangles and blue fit line ($R^2 = 0.85$) represent these combinations under moderate uncertainty. Solid symbols represent the basic strategies, and hollow ones represent the deferred resolution strategies.

Several interesting features are revealed by this analysis. First, the difference between the linear fits indicates that under increased weather forecast uncertainty, it costs quite a bit more to reduce the congestion risk metric to a similar level. Also, there is a relatively smooth linear relationship between the congestion risk and mean cost. The variations from this line indicate strategy-dependent differences, and points that fall below the line represent strategies that have a lower (i.e. better) risk/cost ratio.

If we are looking for the best strategy that meets the 10% goal, it is easy to see that the DeferInc strategy with DT=60 is the lowest cost solution for both levels of weather uncertainty. The general characteristics of the strategy types are also evident. The basic strategies (solid symbols) fall roughly along the fit lines, progressing from the least aggressive (BasicIncHigh) at the upper left of the plot to the most aggressive (BasicLevel) at the lower right. It is a design choice as to how aggressive a strategy to pursue. The BasicLevel strategies control risk well and fall below the fit lines, but also incur many multiple maneuvers and recovery actions (see Table 3). For moderate forecast uncertainty, the deferred strategies with DT=60 are all below the fit line, the best of these being the DeferInc strategy. For deferred strategies, the DT value controls the aggressiveness more so than the congestion profile, with longer defer times producing lower congestion risk and higher cost. Note that while the DT=75 runs produce higher delay than BasicLevel, they do so without incurring recovery actions and with lower re-maneuver rates.

The effect of weather forecast uncertainty can be roughly seen by comparing the cost of strategies that achieve the same risk level, or simply by comparing the fit lines themselves. In this case, it is approximately 25% more expensive to achieve equivalent risk at the high weather forecast uncertainty level than at the moderate level.

Declining profiles (DeferDec) increase costs but produce little improvement in congestion management. The exception is the DT=75 case for the high uncertainty weather. One question is: how can this approach manage congestion better for the high uncertainty weather forecast than for the moderate uncertainty forecast? This is likely because the high uncertainty forecasts admit up-front that their forecasts are not very good, and this is translated into broader sector capacity forecast distributions. The DeferDec profile manages to 46% probability of congestion at 120 minutes LAT. With high uncertainty forecasts, managing to 46% becomes harder, requiring more maneuvers than would be needed if the forecast were more certain. Using DT=75 also increases the number of flights available for proactively managing congestion. Greater forecast uncertainty magnifies the risk aversion of the DeferDec profile, causing more proactive maneuvers. Taking more risk management action when faced with greater uncertainty can be appropriate, and this is what the declining profiles are designed to do. However, the final costs are high compared to other approaches.

Figure 15 shows the severe congestion risk compared to the number of affected flights. There is less difference in this case between the high and moderate forecast uncertainty levels, but there is greater difference visible between the basic (solid symbols) and deferred resolution (hollow symbols) strategies at a given level of risk. This is a direct consequence of deferring resolutions, in that fewer unnecessary flight maneuvers are issued.





F. Results: With Airborne Rerouting

The problem becomes easier to solve when airborne rerouting maneuvers are allowed, especially for the two sector/time combinations that proved difficult. In fact, all tested strategies can achieve the 50% risk probability goal for all sector/time combinations in the simulation when airborne rerouting is allowed (Table 5). The cost, however, is higher than without airborne rerouting. This is for two reasons. First, the algorithm is trying to solve the congestion first, and thus given more options it is doing more maneuvering to achieve this goal. Second, airborne maneuvers are more expensive; ground delay is not an option, and most airborne reroutes will increase flying time.

	Defer	Weather	Mean delay metrics			Mean re	covery m	etrics	Means of other metrics			
Strategy	Time (min)	forecast unc.	No. of flights	Delay (min)	DOC (\$)	No. of flights	Delay (min)	DOC (\$)	MTBD (min)	Delays/ affected flight	Recovery percent.	Unsolved congestion
BasicIncHigh	N/A	Mod	119	613	26871	4	-11	220	20	1.09	3	0
		High	131	799	33693	4	-17	46	18	1.15	3	0
BasicIncMed		Mod	141	795	32552	7	-18	169	27	1.13	4	0
		High	163	1078	42190	7	-27	-98	25	1.23	3	0
BasicIncLow		Mod	169	1018	39195	12	-48	-280	32	1.18	5	0
		High	203	1499	54037	13	-59	-690	31	1.31	4	0
BasicLevel		Mod	238	1647	56597	81	-531	-10379	39	1.31	14	0
		High	289	2388	78209	90	-726	-15319	37	1.48	13	0
DeferInc	45	Mod	119	796	30395	2	-21	-437	9	1.17	1	0
		High	140	1171	42817	3	-28	-632	7	1.24	1	0
	60	Mod	150	1083	38481	7	-52	-1109	21	1.21	1	0
		High	179	1576	53646	7	-65	-1443	20	1.31	1	0
	75	Mod	183	1300	44248	14	-111	-2144	34	1.26	3	0
		High	218	1882	61895	15	-130	-2666	32	1.37	3	0
DeferLevel	45	Mod	124	868	32590	2	-25	-553	9	1.17	1	0
		High	146	1260	45281	3	-35	-797	8	1.25	1	0
	60	Mod	159	1181	41208	9	-74	-1636	23	1.21	2	0
		High	190	1719	57673	9	-86	-1924	21	1.32	2	0
	75	Mod	195	1461	48329	24	-186	-3702	35	1.27	5	0
		High	234	2108	67376	24	-214	-4521	34	1.40	4	0
DeferDec	45	Mod	127	902	33682	3	-30	-685	7	1.17	1	0
		High	150	1319	46886	3	-41	-952	6	1.25	1	0
	60	Mod	165	1250	43067	11	-89	-1983	20	1.22	2	0
		High	196	1802	59758	12	-105	-2366	19	1.33	2	0
	75	Mod	205	1551	50264	30	-236	-4903	32	1.28	5	0
		High	119	613	26871	4	-11	220	20	1.09	3	0

Table 5. Strategy evaluation metrics: with airborne rerouting.

17

American Institute of Aeronautics and Astronautics

Figure 16 shows the real benefit of adding the airborne rerouting option. It greatly decreases the risk of severe congestion with a modest cost increase. Almost all the strategies succeed in reducing the risk to 10% or below. If we were to choose a 5% goal, the DeferInc strategy with DT=45 is the best for both levels of weather forecast accuracy. Looking at the moderate weather prediction uncertainty case, this means that we have achieved half the level of severe congestion risk at a 29% increase in cost. Note that if we wanted a 10% goal, we could choose the BasicIncHigh strategy (leftmost solid triangle in Figure 16. But this is actually more expensive than the best strategy for 10% risk without airborne rerouting (DeferInc, DT=60), because the deferred strategies are generally more efficient than the basic ones.



Figure 16. Risk of severe congestion vs. net mean cost for strategies with airborne rerouting.

VII. Applications

The method proposed here is suitable for real-time applications. Once a suitable congestion management goal is chosen, and simulations have been done to choose the best strategy and parameters to achieve it, the real-time computational requirements are small. Most of the method has already been implemented in a real-time decision support prototype. However, many issues still need to be addressed before an operationally-acceptable version could be deployed. These include: automatic updating of traffic and weather prediction uncertainty 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.

The results shown here are for a single, synthetic-weather scenario. More scenarios need to be analyzed to refine the strategies. The deferred resolution strategy has some clear benefits when compared to a simpler priority scheme, but is sensitive to the choice of congestion tolerance profile and deferability time parameters.

A second use for the method and simulation is cost-benefit analysis. It is possible, for example, to quantify 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.

VIII. Conclusion

A heuristic congestion resolution algorithm (deferability) designed to promote flexibility has been incorporated into an incremental probabilistic en route congestion management method, and evaluated with Monte Carlo simulation. The method has small computational requirements, and is thus suitable for real-time applications. A sample congestion problem, caused by severe weather, was explored using the simulation.

The deferability method develops effective, flight-specific congestion resolution actions (reroutes and ground delays), and was found to do so better than previous fast heuristic approaches. Congestion risk was managed to a target level at lower cost, with fewer unnecessary maneuvers and fewer cases in which flights were maneuvered more than once over the duration of the congestion problem. In short, it was found that adding a simple rule which explicitly attempts to conserve future flexibility can help manage the risk created by uncertain predictions.

Weather forecast uncertainty was shown to have a significant effect on the cost of solving the congestion problem. Also, the value of airborne rerouting for congestion resolution was shown, which must be considered along with the difficulty of rerouting airborne flights in the present NAS.

Acknowledgments

The author would like to thank all who participated in this project, particularly Sandeep Mulgund, James DeArmon, Lixia Song, Daniel Greenbaum, and Claude Jackson for their help in concept development and software.

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 these views.

References

¹Wanke, C., "Continual, Probabilistic Airspace Congestion Management", AIAA Guidance, Navigation and Control Conference, Chicago, IL, 10 - 13 Aug 2009.

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.

³EUROCONTROL, "Pessimistic Sector Capacity Estimation", EUROCONTROL Experimental Centre Note No. 21/03, EUROCONTROL Experimental Centre, Brétigny-Sur-Orge, France, 2003.

⁴Masalonis, A. J., et al., "Indicators of Airspace Complexity for Traffic Flow Management Decision Support," *12th International Symposium on Aviation Psychology*, Dayton, OH, 14-17 April 2003.

⁵Masalonis, A. J., M. B. Callaham, and C. R. Wanke, "Dynamic Density and Complexity Metrics for Realtime Traffic Flow Management," *5th USA/Europe Air Traffic Management Research and Development Seminar*, Budapest, Hungary, 23-27 June 2003.

⁶Song, L., Wanke, C. and Greenbaum, D., "Predicting Sector Capacity for TFM," 7th USA/Europe Air Traffic Management Research and Development Seminar, Barcelona, Spain, 2-5 July 2007.

⁷Kopardekar, P., and S. Magyarits, "Measurement and Prediction of Dynamic Density," 5th USA/Europe Air Traffic Management Research and Development Seminar, Budapest, Hungary, 23-27 June 2003.

⁸Weber, M., Evans, J., Wolfson, M., DeLaura, R., Moser, B., and Martin, B., "Improving Air Traffic Management During Thunderstorms," *12th Conference on Aviation, Range, and Aerospace Meteorology*, Atlanta, GA, 29 January - 2 February 2006.

⁹Song, 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.

¹⁰Krozel J., Mitchell, J., Polishchuk, V. and Prete, J., "Capacity Estimation for Airspaces with Convective Weather Constraints", *AIAA Guidance, Navigation, and Control Conference*, Hilton Head, SC, August 2007.

¹¹Martin, 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.

¹²DeLaura, R., Crowe, B., Ferris, R., Love, J., and Chan, W., "Comparing Convective Weather Avoidance Models and Aircraft-Based Data," *Aviation, Range and Aerospace Meteorology Special Symposium on Weather-Air Traffic Management Integration*, AMS Annual Meeting, Phoenix, AZ, January 2009.

¹³Wanke, C., Song, L, Zobell, S., Greenbaum, D., and Mulgund, S., "Probabilistic Congestion Management," 6th USA/Europe Air Traffic Management Research and Development Seminar, Baltimore, MD, 27 - 30 June 2005.

¹⁴Taylor, 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.

¹⁵Taylor, C. and Wanke, C., "Dynamic Generation of Operationally Acceptable Reroutes," 9th AIAA Aviation Technology, Integration, and Operations Conference, Hilton Head, SC, September 2009.

¹⁶Steiner, 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.

¹⁷Steiner, M., Bateman, R., Megenhardt, D., and Pinto, J., "Evaluation of Ensemble-based Probabilistic Weather Information for Air Traffic Management," Aviation, Range and Aerospace Meteorology Special Symposium on Weather-Air Traffic Management Integration, AMS Annual Meeting, Phoenix, AZ, January 2009.