Probabilistic TFM: Preliminary Benefits Analysis of an Incremental Solution Approach

James DeArmon, Craig Wanke, Daniel Greenbaum
MITRE/CAASD
7525 Colshire Dr
McLean VA 22102

Introduction

The National Airspace System (NAS) is the collection of airports, airspace, and other resources which enable air transportation in the U.S. When considered in light of the interactions with airspace users (airlines and non-commercial flight operations) and environmental conditions (e.g., winds and weather), this larger gestalt is quite complex. One of its notable attributes is uncertainty – there is uncertainty in nearly every aspect: flight take-off times, airport and airspace capacities, effects of decisions made by airspace users and airspace managers, etc. Traffic Flow Management (TFM) is the function, performed by the Federal Aviation Administration (FAA), which seeks to balance demand and capacity for airspace/airport resources. Also, more informally called "resource management", TFM typically must manipulate air traffic demand, since resource capacity (e.g., weather impacting air routes, or runway repair impacting airport operations rate) is typically non-negotiable.

In the last years, a new idea in the area of TFM planning and analysis has evolved, known as "Probabilistic TFM" (PTFM) [[GNC 2004][ATM 2005][Hunter and Ramamoorthy, 2006]. PTFM pursues an explicit recognition of uncertainty – it attempts to incorporate probability theory and thereby to improve the decision-making. PTFM can be thought of as applying broadly to two areas: 1) visualization and event-alerting technology, and 2) decision support analysis. Regarding visualization, consider notional figures 1 and 2.

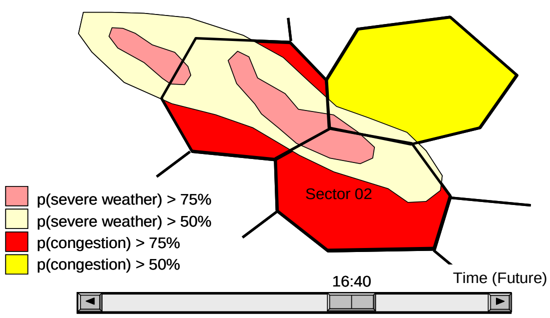

Figure 1. A probabilistic airspace congestion

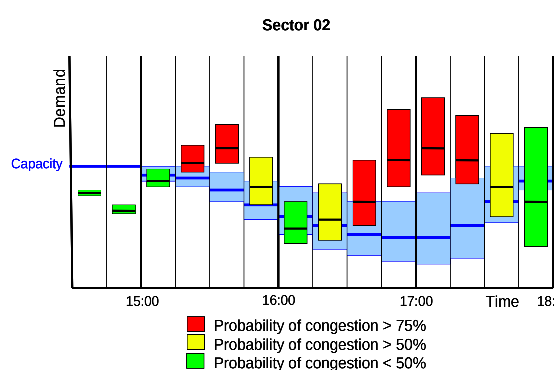

Figure 2. The congestion forecast over time shows transient and longer term conditions

In Fig. 1, using weather forecasts, contours have been constructed to delineate regions of likelihood of severe en route weather, at levels >50% and >75% likelihood. Considering these impacts on airspace sector capacity and staffing, as well as expected air traffic demand, a probabilistic score can be associated with affected sectors, here >50% and >75% probability of congestion. Figure 2 shows the demand vs. capacity situation for Sector 2 over time. The blue boxes show capacity, expressed as an expected or mean value in the middle of the box, plus 50% error bounds at the top and bottom of the box. Green, yellow, and red boxes, corresponding to increasing probability of congestion, show expected demand, using the same format of mean and confidence interval bounds. (Boxes have mean lines not necessarily evenly splitting the vertical extent of the box, since the probability distributions are typically non-normal and even non-symmetric). A traffic flow manager could manage resources using an automation tool with a display as in Fig. 2, by reducing demand (either in time via delay or in space via alternate routing) until the red boxes become lower on the display (the color would thence become yellow), matching better the blue, underlying capacity distributions.

Note that this example is notional – concepts and prototypes are currently being pursued, but are not in the field. Current TFM systems do not consider uncertainty. Alerts, for quarter-hour intervals some hours into the future, are either green ("do nothing"), yellow ("monitor the situation"), or red ("investigate to see if intervention is necessary"). A congestion risk of 50%, as this phrase is used in this paper, approximately corresponds to a "red" alert in the field currently, even though today's alerts are not explicitly probabilistic.

Regarding PTFM used for decision support analysis, that is the topic of this paper. We will show an application of probability theory – incremental problem solving addressing the TFM challenge in stages: given a problem such as severe en route weather, multiple decision points are defined, and some decision activity takes place at each decision point. By contrast, today's TFM approach uses first a strategic approach, and then a tactical one. A first pass is undertaken from a strategic perspective, by the national TFM facility. It is understood that this is not a complete problem solution, and will have to be fine-tuned – the fine tuning is performed, as needed as events unfold, by the local TFM facilities, using a tactical perspective. In terms of benefits analysis, a contrast can be created: today's approach of strategic, then tactical treatments vs. an incremental approach of successive partial solutions, backed by quantitative analysis.

It can be argued that, on the face, the incremental approach should best a single, early, one-time solution, but is that intuition borne out in quantitative experimentation? This paper presents experimental results to answer that question.

Background: A Challenging Problem involving Severe En Route Weather

Consider Figure 3, a good example of challenging TFM in light of en route capacity shortfalls due to en route weather. The figure has key visual features: actual weather is shown on the geographic display as green, yellow, and orange/red regions. Predictions for the movements of high-intensity weather are delineated as the black-outlined polygons with accompanying movement vectors. Other visual features: airspace sectors with red cross-hatching are where demand is predicted to exceed capacity; sectors with yellow cross-hatching are where demand may exceed capacity.

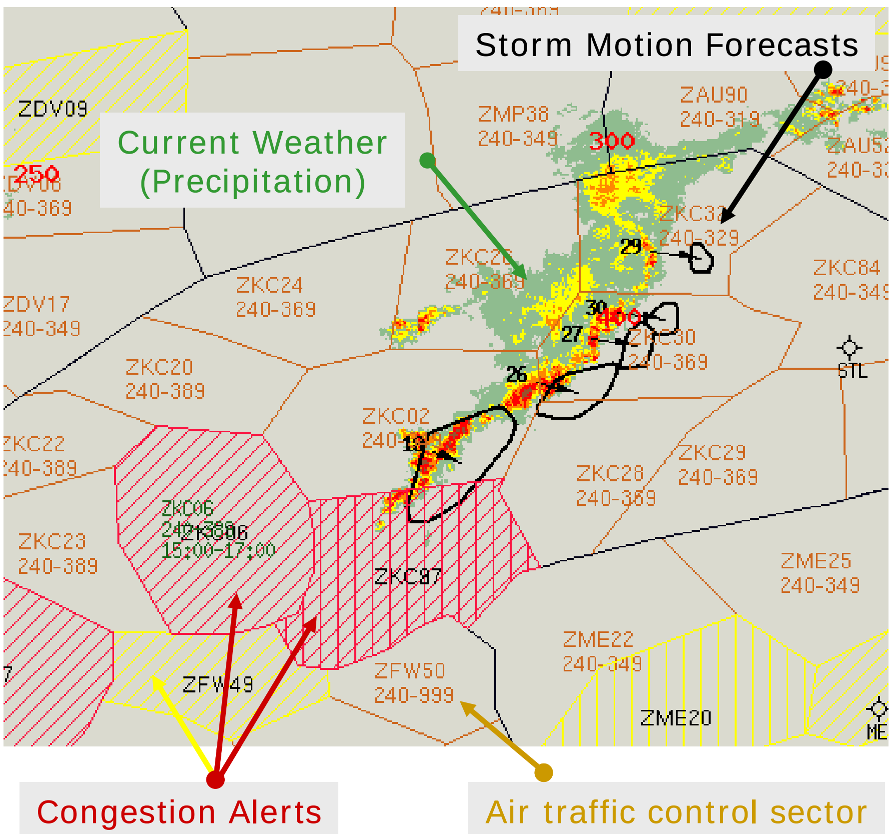

Figure 3: Example of Severe En Route Weather, with Predicted Movement of Weather and Capacity-Impacted Airspace Sectors

This figure is a typical challenging severe en route weather TFM scenario. Uncertainty overshadows all of the decision-making, and some questions arise:

1) When should air traffic be restricted?
2) Which flights should be affected?
3) How can NAS operators participate?

The challenge presented in this scenario is addressed in the subsequent sections of this paper.

Solving the Problem Incrementally

We propose a process for systematic, incremental intervention to solve the problem of en route demand exceeding available sector capacity, which we call the Probabilistic Incremental Congestion Alleviator (PICA). Figure 4 gives a diagram of the process. (For a fully-detailed description, see [ATM 2007]).

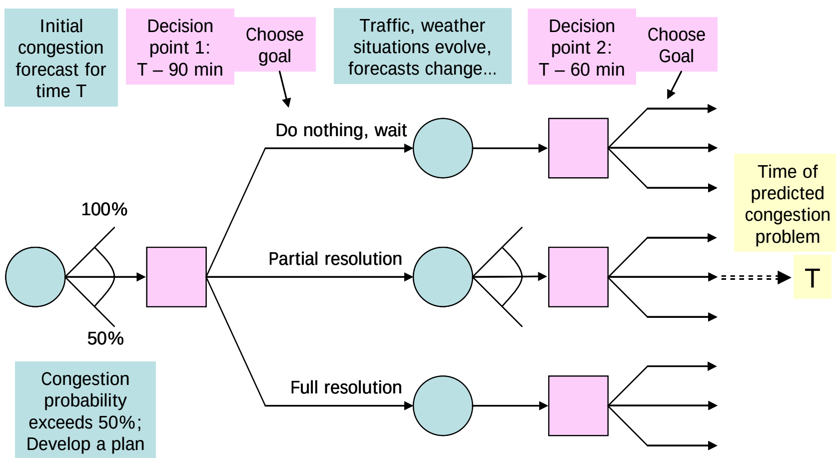

Figure 4: Congestion Resolution Decision Tree

Starting on the far left of the figure, a set of flight trajectories are subjected to Monte Carlo modeling, creating a population of realizations for the predicted flights. Congestion evaluation is performed, and sector congestion is forecast for future time T. Since the probability of congestion in one or more sectors exceeds 50%, it is necessary to intervene. Moving right in Figure 4, Decision Point 1 (DP1) takes place 90 minutes prior to the predicted congestion problem, at time T-90. Here we show that three options are considered and modeled: do nothing, partial resolution, and full resolution. The partial and full resolution options apply ground delays and/or reroutes to address the expected sector congestion. Trajectories with these maneuvers are substituted in each set of flight realizations. The costs of the delays and reroutes are saved, as they will be used in a later step.

Now, half-way through Figure 4, at the three vertical circles, simulated time elapses during which certain events occur (in the context of this modeling system): flight cancellations, new flight plan filings ("pop-ups"), and changes in flights': departure times, route, cruise altitude, and speed. These events have the effect of altering the "virtual history" of the simulation system, as in the real world – as time passes, some uncertainty evaporates. Predicted events unfold in one way or another: sector capacity is impacted, or not; scheduled flights show up for service, or they don't. DP2 at T-60 minutes is the next column of pictographs moving right in Figure 4, the three vertical squares. Given the current population of realizations, spawned by the selected option from Decision Point 1, three intervention options are again considered. Again congestion is solved to some level per one of three options; delays and reroutes are substituted in the populations of realizations.

Note the multiplicity of solution paths brought forward. Coming from Decision Point 1, 3 populations of realizations, corresponding to the 3 options, move forward in the process. Each of these three populations is considered for the 3 options at Decision Point 2, yielding 3 x 3 = 9 populations of realizations to be brought forward. As the number of decision points and options at each point expand, the problem grows ever larger. However, the final decision point must have only a single option, 50% probability of congestion, since any residual congestion (any remaining after the earlier interventions) must be removed. This is shown in the next section.

A Simple Example

An example problem shown in Figure 5 will make some of the above explanations clearer. Corresponding roughly to the region depicted in Figure 3, we assume that the en route airspace sectors known as ZKC31 and ZKC84 (ZKC is the identifier for the Kansas City Air Route Traffic Control Center) are predicted to experience congestion, and the decision points are 90 and 30 minutes prior to that time. As part of the problem solution, sectors adjacent to the subject impacted sectors are also monitored – these sectors likewise will be managed to avoid congestion – this could happen if rerouted traffic avoided ZKC31 and ZKC84, only to “bunch” in the adjacent sectors.

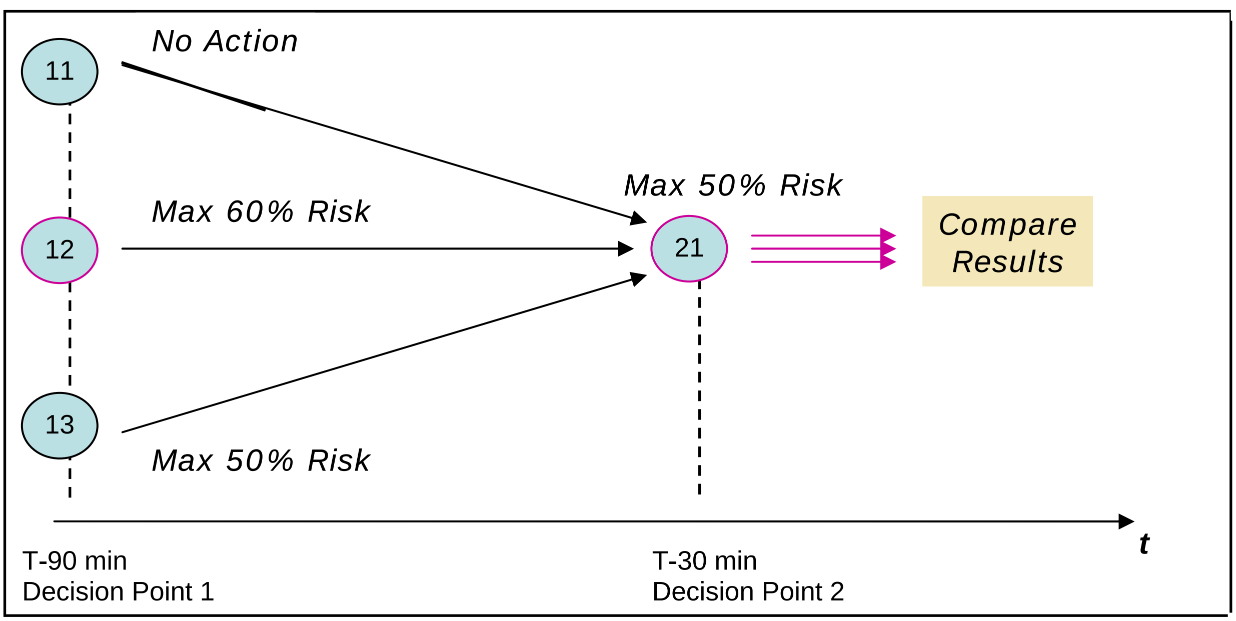

Figure 5: Example Problem: Decision Points and Options

At DP1, three options are considered – allow 100%, 60%, or 50% chance of congestion. At DP2, the final decision point, a single option must be taken – all remaining congestion must be handled, and so 50% is the single option to be undertaken. The right-most box says "Compare Results" – this is the step of evaluating which path through the decision tree, i.e., which sequence of options taken yielded the lowest cost. Cost is the only measure of interest, since "TFM efficacy" is insured by the final decision point wherein all remaining congestion is handled.

Translating these into English, they are as follows:

DP1=100%, DP2=50% – "Wait until the end, then solve everything"
DP1= 60%, DP2=50% – "Partial solution up front, revise later"
DP1= 50%, DP2=50% – "Solve everything up front"

Using total positive delay minutes as the cost measure (not shown), results show what might be guessed, a "middle path" solution approach is reasonable. By solving a portion of the problem early, and then later solving the remainder is better than either of the other two approaches. The (50%, 50%) "Solve everything up front" approach is overly conservative – more delay was assigned than was necessary, since there was so much uncertainty at DP1. Likewise (100%, 50%) "Wait till the end, then solve everything" could have done some of TFM intervention earlier. The best approach was (60%, 50%) "Partial solution early, then revise". It may be that 90 minutes prior to predicted congestion in an en route sector is simply too early to act – too many of the involved flights have yet to depart, and take-off time uncertainty is the largest single contributor to overall flight uncertainty.

Although these results may be intuitive, we now have a simulation facility to experiment with various problems and the parameter settings, yielding a new (as far as the authors know) area of knowledge for the TFM research community.

<u>Note regarding final version of the paper</u>

In the final version of this paper, we will present results for several larger problems, and explore the impact of parameters such as: number of decision points, number of choices, timing of decision points, sample sizes in the Monte Carlo modeling, etc. One problem of especial interest was presented at ATIO 2006 [ATIO 2006] – one of the worst days of 2004 with respect to severe weather, involving nearly 1000 flights, and scores of affected sectors. Benefits measures will include counts of flights, miles of path deviation, ground delay, plus a "monetarizing" step which yields a rough dollar estimate of the potential of this technique.

Notice

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

REFERENCES

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

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

[Hunter and Ramamoorthy, 2006] Hunter, G., and Ramamoorthy, ?, "A Trajectory-Based Probabilistic TFM Evaluation Tool and Experiment" ICNS Conference, Baltimore, MD, May 2006.

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

[ATM 2007] Wanke, C. and Greenbaum, D., "Incremental, Probabilistic Decision-Making for En Route Traffic Management", submitted for presentation to ATM 2007: 7th USA/Europe R & D Seminar, to be held in Barcelona, Spain, July 2007.