Proceedings of the 2004 Winter Simulation Conference
R. G. Ingalls, M. D. Rossetti, J. S. Smith, and B. A. Peters, eds.

ESTIMATING EFFICACY OF PROGRESSIVE PLANNING FOR AIR TRAFFIC FLOW MANAGEMENT

Lynne Fellman
James S. DeArmon
Kelly A. Connolly

The MITRE Corporation
Center for Advanced Aviation System Development
7515 Colshire Drive
McLean, VA 22102-7508, U.S.A.

ABSTRACT

Air traffic flow management (TFM) is a set of processes and procedures which seek to balance the demand for airspace resources with the capacity of these resources. Examples of resources are airports, sectors (airspace volumes managed by air traffic controllers), and fixes (imaginary points in space used for navigation). The Federal Aviation Administration (FAA) is continually looking for ways to provide new tools and techniques for TFM personnel. As the TFM function improves, flight efficiency improves, and the experience of the flying public is likewise improved.

In this paper, we describe a simulation modeling exercise to assess the benefit, if any, of a proposed new feature of TFM called Progressive Planning (P2). P2 allows the flow manager to model the impact of multiple concurrent flow management actions. It is envisioned that the improved modeling leads to better decision-making, which leads to greater flight efficiency.

1 BACKGROUND

Traffic flow management is used by the FAA to balance air traffic demand with airspace and airport capacity constraints during times of high demand or resource limitation. FAA TFM personnel at national and regional facilities are tasked with developing and implementing plans to solve demand/ capacity imbalances. This is accomplished through TFM Initiatives (TMIs), which typically include modifying flight plans to reroute aircraft around severe weather or congested regions of airspace, and restricting rerouted flows by spacing aircraft via enforcement of miles-in-trail (MIT) restrictions. Alternately, a TMI may hold aircraft on the ground prior to departure. TMIs are created prior to the expected event, based on assumed aircraft flight paths and weather predictions, and are modified as needed to respond to the actual weather and demand as these develop.

Currently, there are limited automation tools available to the traffic managers to assist in this process. MITRE and the FAA are developing a TFM Integrated Impact Assessment Capability (IIAC) in which a traffic manager can specify any combination of rerouting and MIT spacing restrictions (Ball, Fellman, Taber, Wanke and Yee 2003). The capability predicts the airspace demand impact and the imposed delays, and allows the traffic manager to adjust route and MIT restriction parameters to improve the proposed solution. With this tool, the traffic manager can investigate the effectiveness of possible resolutions prior to implementing the TMI. In the current version of IIAC, each TMI is modeled independently, in isolation from other TMIs.

An enhancement to the TFM IIAC, called Progressive Planning (P2), is currently under investigation. P2 allows traffic managers to evaluate multiple TMIs in the composite (Rhodes, Connolly, Fellman and Wanke 2003). P2 is expected to improve airspace demand predictions, leading to more accurate traffic management strategies. Improved predictions should result in fewer missed or false alerts, more accurate planning for the geographic extent and time duration of TMIs, and should minimize the number of aircraft affected by TMIs.

To test the question of the utility of the P2 enhancement, several research questions are posed (Ostwald, DeArmon and Wanke 2003): What is the efficacy of P2 for congestion management planning? Were the right TMIs established? Did they impact the correct flights? Were the number of flights and the duration of the TMI minimized? The approach is to perform simulation runs, capture results, and present these to human subjects who have operational experience. Because of the difficulty (or even impossibility) of real-world trials, simulation modeling was selected to address the question of the efficacy of P2 for congestion management.

In the parlance of simulation modeling (see Law and Kelton 2000), we have a *dynamic, deterministic* simulation model. It is dynamic in that the situation (state of flights, sector demand, etc.) change over time. Our model is deterministic in that there are no random variables—a set of TMI inputs create the same outputs each time the simulation model is run.

2 APPROACH

The approach used in this research is to employ recorded real-world traffic collected for a day when TMIs were in effect and apply P2 to the process of investigating and resolving the predicted traffic flow problems. Two scenarios are modeled. The first scenario models each of the actual TMIs put into operation on that day, in isolation, using IIAC *without* P2. The second scenario models the same day but with the TMIs modeled in the composite using IIAC *with* P2. The airspace demand predictions for each of these two scenarios are captured for various times during the day, including the approximate time when the TMIs are developed prior to their start time, and at various times when the TMIs are active. The results of these two scenarios are then presented to personnel with operational experience to assess whether a different action might have been taken if P2 airspace loading information were available. Based on the assessment, an estimate of the amount of TFM-imposed delays, assuming different actions are taken, can be modeled and benefits derived.

3 IIAC SIMULATION

The IIAC uses recorded traffic of all flights throughout the continental U.S. to establish current and future airspace demand. Depending on the time being modeled, either the actual aircraft position reports for airborne flights or the filed flight plans for flights scheduled for departure at a future time are used in this estimate. The airspace demand is calculated for each of the approximately 700 sectors, a sector being the volume of airspace for which a single air traffic controller or team is responsible. The IIAC calculates the number of aircraft expected to occupy each sector, for each 15 minute period, starting at the current time through the ensuing 6 hours. These counts are displayed graphically in two forms: the NAS (National Airspace System) Monitor shows the entire U.S., and the Sector Monitor shows counts for a single en route Air Traffic Control center.

Associated with each sector is what is called the "Monitor Alert Program" (MAP) value—it specifies the (approximate) maximum number of aircraft that can be safely handled at any given time and is specific to each individual sector. A sector is said to be in alert status if the predicted demand exceeds this threshold capacity.

The IIAC also includes a capability to assess "what-if" questions, i.e., to model the impact of TMIs on sector demand. The Traffic Manger can use the IIAC as a means for testing the aptness of a TMI on resolving a problem, prior to implementing the TMI. The TMIs can be adjusted and retested as needed to resolve the problem, and potentially be adjusted to minimize the number of aircraft impacted. Currently, the TMIs that can be modeled by the IIAC include calling-up a stored reroute, or creating an ad hoc reroute, for all flights meeting filter criteria, and restricting flows of aircraft crossing a geographic location by spacing aircraft a certain distance in-trail.

The results of the predicted demand and airspace capacity within the NAS are presented to the Traffic Manger on two displays. The graphical display of the predicted sector demand, summarized for each of the 20 air traffic control centers, is shown on what is called the NAS Monitor display and the breakdown of each of the sectors within each air traffic control center is shown on the Sector Monitor display. These displays are described later in the paper.

Presently, the IIAC models each TMI in isolation. Often, on extreme weather days, multiple TMIs are active concurrently to circumvent bad weather and to resolve the resulting airspace congestion problems. The drawbacks of modeling each TMI in isolation is that it is possible for reroutes to overlap geographically and temporally and the combined effect of multiple TMIs would not be reflected in the sector counts. As a result, the IIAC was enhanced via P2, which models the impact of concurrent TMIs on the prediction of airspace demand.

4 EXPERIMENT

To answer the questions on tool utility presented earlier, scenarios using recorded real-world data for a specific bad weather day were developed to compare decisions which may be made given the availability of the P2 capability. The day selected was one in which TMIs were in effect and also had a likelihood of demonstrating benefits, if any, from the P2 capability. The IIAC was used to model the scenarios. The first scenario modeled each of the TMIs that were actually used for the bad-weather day in isolation using the IIAC without P2. The second scenario modeled the same TMIs in the composite using IIAC with P2. Charts showing the airspace demand predictions for each of the scenarios were captured at various times during the simulation day at the approximate time when the TMIs were developed and at various times while the TMIs were active.

The day selected for the initial experiment was 9 May 2003, a day that had severe weather in the mid-Atlantic region, and which had multiple overlapping TMIs in effect (see Figure 1). The time periods investigated included both before the TMIs were issued (when the Traffic Manager would be planning the TMI) and while several TMIs were active.

Three alternate routing plans ("reroutes"), called-up from a collection of stored reroute patterns called a "playbook," were in effect during the period of time investigated on 9 May 2003: a play called "No J6 1" in which flights

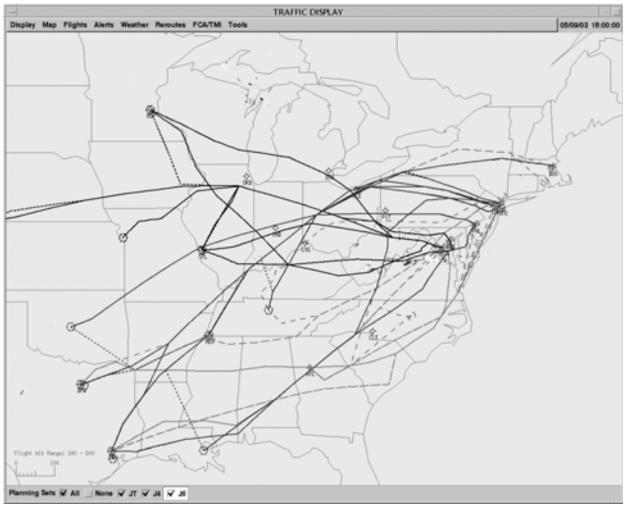

Figure 1: May 9, 2003 Reroutes

from the northeast bound for Atlanta, Dallas, Houston, Cincinnati, and Memphis were rerouted off of airway J6 typically used for those flights; "No J42 2" in which flights to the major airports in New York, Boston and Philadelphia were rerouted off J42, and "JOT 1" which rerouted traffic for flights bound for the DC metropolitan airports, crossing a geographical location called Joliet, designated JOT. Associated with each of the reroute plans were spacing restrictions including 30 miles-in-trail per airport for flights to the airports listed above. The predicted sector loads and miles-in-trail delays were captured at the approximate times when the reroutes were developed and at times when all three reroute plans were active for both the scenario showing the reroutes individually and for the scenario where the reroutes were developed in the composite. The NAS Monitor and Sector Monitors for both scenarios were presented to personnel with operational experience to determine if a different action might be taken, given the improved accuracy of the P2 predictions.

5 INITIAL RESULTS

Prior to running the experiment it was thought that, using P2 sector loading information, the experimental subjects would suggest modifications to the original TMIs and develop a more efficient plan to circumnavigate the predicted weather while also avoiding creating secondary congestion problems. Initial outcomes of this experiment did not produce this result.

For the 9 May 2003 experiment the displays showing center and sector load for each individual reroute and for the combined P2 view did not show potential problems. The NAS Monitor showed a slightly different view of sector load using P2 (Figure 2) than was shown in the single reroute with the most impact (Figure 3) and the subjects did not change the plan as originally developed. Color coding and notations on the display are as follows. If the predicted number of flights for a sector exceeds the MAP value, the entry is displayed in red for that 15 minute time period, *except:* if the predicted number exceeds the MAP values but not all of the aircraft have departed, then the sector is displayed in yellow for the 15 minute time period. Otherwise the cell is colored green. The cells on the NAS Monitor present the number of sectors predicted to exceed their MAP parameter. The results of TMI modeling are also coded on the NAS Monitor and Sector Monitor. Any changes resulting from the implementation of a TMI are indicated on the two displays by outlining the cell in either dark blue to indicate the change caused a sector to go from non-alerted to alerted status, or in light blue to indicate the change caused a sector to go the other way.

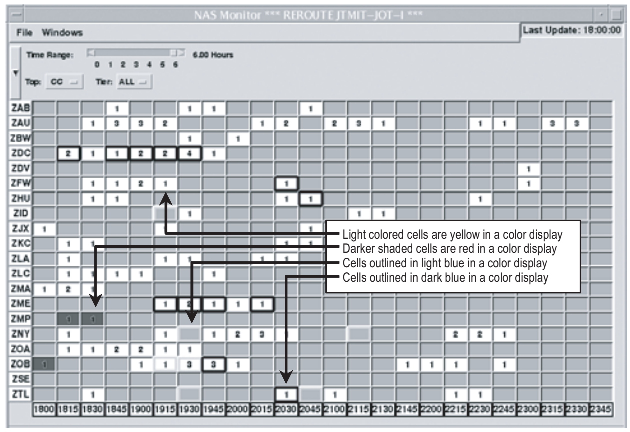

Figure 2: Sector Loading, Composite of Three Reroute Initiatives

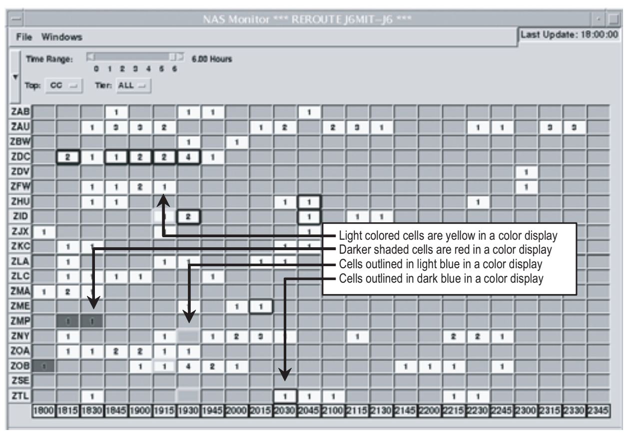

Figure 3: Sector Loading, "No J6 1" Reroute Initiative

The response from the experimental subjects was that the reroutes as specified were acceptable. However, it was noted that the miles-in-trail restrictions were unnecessary and resulted in unnecessary flight delay and workload for air traffic controllers (who are charged with implementing the restriction). It is believed that, given more accurate predictions of airspace demand provided by IIAC, Traffic

Managers will have better information needed in order to avoid routinely issuing miles-in-trail restrictions. But per initial results of experimentation, the composite TMI modeling feature, which we've referred to in this paper as P2 (or Progressive Planning), seems to have little or no utility, since subjects couldn't articulate how they might use the information to improve flow efficiency. A fortiori, for the single scenario examined, there was virtually no difference in the visual representation, between the highest impact TMI and the P2 display. Further analysis is required to either confirm or refute this initial result.

6 COST OF MILES-IN-TRAIL

The number of flights and the delay time imposed by the MITs used on 9 May 2003 as computed by the IIAC was used to estimate the savings that may be attributed to removal of MITs. This is an assessment of the benefit of IIAC, and *not* P2, since it was IIAC *without* the P2 enhancement that led subjects to assert that MIT restrictions were likely unnecessary. Although it was not the intent of this study to discern benefits of IIAC, the information is useful in planning future functions for implementation in the TFM system. Flight and delay statistics are shown in Table 1.

Table 1: MIT Delays

Reroute	Number Flights Delayed	Total Delay (minutes)
JOT 1	14	61.5
No J6 1	19	35.4
No J42 2	28	110.9

The estimated operating cost of airborne flights was \$57.15 per minute. (This amount includes both Airline Direct Operating Costs and passenger value of time. See (FAA 2001).) Additional MITs associated with moving flights on the "JOT 1" reroute imposed another 126 minutes of delay making the total cost of delay for MITs on all three reroutes of \$18,974.

7 PREDICTIVE ACCURACY OF P2

The value of P2 modeling was evaluated in an alternate way: examination of the "predictive accuracy" using statistical analysis. Predictive accuracy is defined here as the correctness of estimates of future sector demand counts. As with the above experiment, the IIAC was used with recorded data. The experiment set up was as follows. Four different treatments of the 9 May 2003 scenario were considered as competing "prediction engines," and each was compared to a set of "truth" values (best estimates of actual sector counts). The four prediction engines were:

- No TMIs
- TMI #1
- TMI #2
- Composite of TMI#1 and #2 (P2).

For various look-ahead times (every half-hour, starting four hours ahead), predicted counts were compared to the truth values. As an overall measure of accuracy, the average of the absolute difference between predicted and truth values was computed. Figure 4 shows, for the four prediction engines, this measure at various look-ahead times. As the look-ahead time diminishes, left to right in the figure, the prediction time gets closer to the evaluated time, and all four prediction engines have progressively improved accuracy, as expected. Note, though, how close together the four measures are. As a further step in this approach, individual sector errors (in the form of absolute difference between prediction and truth) were considered in an Analysis of Variance context. The ANOVA was constructed as a single factor, four level analysis, with the null hypothesis:

Ho: Prediction values are not different
across prediction engines

Using just two look-ahead times, 60 minutes and 120 minutes, the null hypothesis was not (nearly) rejected, with p-values for the two look-ahead times exceeding 0.97. This analysis gave us further evidence that P2 predictions are not very different from the higher impact initiative of TMI #1 and TMI #2.

8 HOW TYPICAL WAS THE SELECTED DAY?

The date 9 May 2003 chosen for this analysis was selected explicitly to test the utility of P2. On this date multiple overlapping reroutes were active concurrently. The question becomes: How often does this type of situation occur? Since we found little utility of P2 on this date, we need to look for days which are even more complicated with respect to multiple overlapping reroute initiatives. A separate analysis is being performed to examine a year's worth of data to see how often reroutes were issued which overlap in space and time.

The analysis is proceeding as follows. For each day in 2003, recorded data of all reroutes that were issued along with their start and stop times are examined. A table of each sector penetrated by the reroute is constructed. The table will show for each day, for each half-hour period, the number of reroutes penetrating the sector. Using that result, a daily total of all sectors with more than one reroute will be tallied. A histogram of the year can thence be developed. Comparing 9 May 2003, a day we've analyzed in detail, to the other days in the histogram will allow us to select days with potential P2 benefit, in light of multiple, concurrent reroutes. Stating the obvious, a sufficient sample size is necessary for drawing conclusions about the utility of P2.

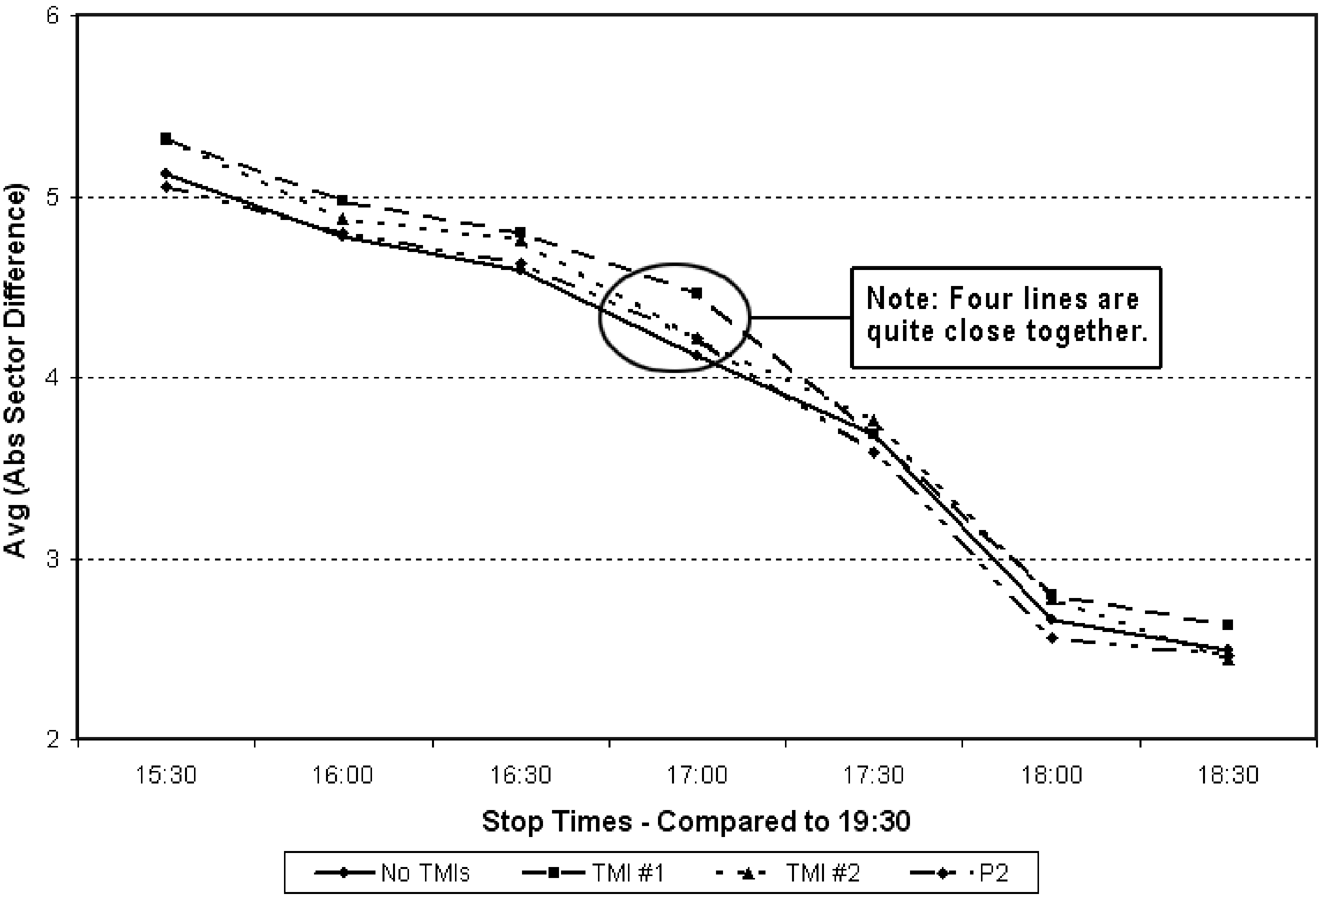

Figure 4: Predictive Accuracy of Four "Prediction Engines"

9 NEXT STEPS

The airspace loading in the P2 composite version of the experiment showed minor differences from the non-integrated scenario, and the traffic managers did not change the plan as originally developed. Further work is ongoing to test additional real-world scenarios with operationally experienced staff to see if this early conclusion is maintained. Simulation modeling was an invaluable tool for this work, since real-world trials would be prohibitively expensive or simply not possible.

ACKNOWLEDGMENTS

Debra Berry of The MITRE Corporation provided the information on the P2 predictive accuracy, and Phil Brown of The MITRE Corporation provided a description on the full-year analysis of overlapping reroutes.

REFERENCES

Ball, C. G., L. Fellman, N. J. Taber, C. R. Wanke, and M. Yee. 2003. *A Concept of Use for an Initial Integrated Impact Assessment Capability*, Report No. MP 01W0000139R1, The MITRE Corporation, McLean, VA.

FAA. 2001. *Investment Analysis Study Report: The Business Case for Free Flight Phase 2*, Federal Aviation Administration, Washington, D.C.

Law, A. and W. D. Kelton. 2000. *Simulation Modeling and Analysis*, McGraw-Hill, New York.

Ostwald, P. A., J. S. DeArmon, and C. R. Wanke. 2003. *Analysis Plan for Progressive Planning*, Report No. MP 03W0000258, The MITRE Corporation, McLean, VA.

Rhodes, L. R., K. A. Connolly, L. Fellman, and C. R. Wanke. 2003. *Updated Operational Concept for Progressive Planning*, Report No. MTR 03W0000036, The MITRE Corporation, McLean, VA.

AUTHOR BIOGRAPHIES

LYNNE FELLMAN is a Senior Staff at The MITRE Corporation's Center for Advanced Aviation System Development (MITRE/CAASD) and has worked for the past 18 years on air traffic control applications, most currently relating to performance and benefits analyses. She holds a B.S. degree in mathematics from the University of Maryland

and M.S. degree in computer science from Virginia Polytechnic Institute and State University.

JAMES S. DeARMON is a Principal Engineer with The MITRE Corporation's Center for Advanced Aviation System Development (MITRE/CAASD). He has worked for the last several years in research and modeling of air traffic flow dynamics. He holds a Master's degree in Operations Research, and has taught graduate courses at the University of Maryland. He has written numerous papers in aviation and applied math. He is a member of the Editorial Advisory Board of the journal *Computers and Operations Research*.

KELLY A. CONNOLLY is a Project Team Manager at The MITRE Corporation's Center for Advanced Aviation System Development (MITRE/CAASD) specializing in Air Traffic Management Decision Support Systems and Future Automation Research. She currently leads a diverse team of engineers focusing in the areas of Human Factors, Experimental Design and Execution, Operational Procedures, and Technology Transfer of large-scale systems. She earned an M.S. in Engineering Management from George Washington University and a B.S. in Information Systems Management from the University of Maryland.

DISCLAIMER

The contents of this material reflect the views of the author. Neither the Federal Aviation Administration nor the Department of Transportation makes any warranty, guarantee, or promise, either expressed or implied, concerning the content or accuracy of the views expressed herein.