MODELING TIME AND SPACE METERING OF FLIGHTS IN THE NATIONAL AIRSPACE SYSTEM

Paul T. R. Wang
Craig R. Wanke

Center for Advanced Aviation System Development (CAASD)
The MITRE Corporation, Mclean, VA 22102

Frederick P. Wieland
AEgis Technologies, Alexandria, VA 22311

ABSTRACT

Metering flights at key points such as sector crossings is an important operational procedure in mitigating National Airspace System (NAS) traffic congestion due to high demand or changing weather conditions. The authors combine a mathematical model for minutes-in-trail or miles-in-trail (MIT) metering with discrete event simulation in a newly developed tool that can be used by analysts to examine or predict existing or developing bottlenecks within the NAS. We define a penalty function recursively in terms of MIT delays between leading and following flights. Such a recursive MIT penalty function has been implemented in a NAS-wide discrete event simulation at CAASD to provide predictive assessment of MIT-related delays locally and globally. It is possible to examine the anticipated MIT delays for all the flights scheduled to arrive at any crossing point. Impacts of flight cancellations, route changes, and additional en route delays as results of airport or sector congestion can all be evaluated during each simulation by updating the scheduled flight crossing times and the expected MIT delay penalties for all the trailing flights. A limited test case quantifying the impacts of increasing spatial arrival separation at a busy airport is provided to illustrate tradeoffs between en route queuing delay versus airport arrival queuing delay both locally and globally.

1 INTRODUCTION

The National Airspace System (NAS) handles over 50,000 daily flights. Scheduled flights may be connected through itineraries that show the flight legs traversed by a single airframe during the course of a day. If the originally intended itinerary is disrupted, e.g., by bad weather or excessive congestion, flights may be terminated, delayed, diverted, replaced, or rerouted to their departure airports. Flights may also be delayed, diverted, or metered while en route at key control points such as sector crossings, fixes, or waypoints. In addition, flights are also subject to handoff, sequencing, and metering for events such as takeoff, landing, and sector crossing. Air traffic flow management (TFM) (Ball, Connolly, and Wanke, 2003) procedures such as Ground Delay program (GDP), ground stop (GS), or miles-in-trail (MIT) metering are options available to the Air Traffic Management (ATM) authority to manage airway congestion and to respond to anticipated weather conditions (Wanke et al., 2003). The impacts of specific TFM actions on overall NAS performance can be measured with metrics such as flight delays and fuel use. Multiple simultaneous TFM actions may be highly interdependent, and the effects of a TFM action can ripple to other NAS resources and other flights during the day (Ostwald et al., 2003). The effects of such complex interactions can potentially be quantified with either discrete event simulation or mathematical models or both. In this analysis, the authors developed a recursive MIT penalty function to quantify the ripple effects of specific MIT programs over relevant sets of flights and flight restrictions within the NAS. In conjunction with discrete event simulation, it is possible to examine and quantify the total impacts of various TFM programs for alternatives analysis and provide a comparison across several alternative TFM programs available to air traffic flow management decision-makers. Combining the MIT penalty function with fast event-driven simulation, it is demonstrated that potential congestion “hot spots” in the NAS can be identified based on flight schedules. Potential or developing bottlenecks also can be simulated for anticipated or real weather conditions and the impacts of alternative GDP, GS, or MIT programs quantified either individually or collectively.

2 A RECURSIVE MIT PENALTY FUNCTION

Safety is the ATM authority’s ultimate reason for keeping aircraft separated. In some cases, aircraft separations may need to be maintained at distances significantly larger than minimum separation standards to accommodate flow control or other operational requirements. In principle, aircraft may be separated in time or space or both. When both aircraft type and instantaneous speed are known, spatial and time separation are theoretically interchangeable, although in actual operations, ATM procedures mandate minimum spatial separation standards for aircraft in en route and terminal airspace (Beaton et al., 2002). The enforcement of aircraft separation for a given pair of flights may ripple through the remainder of the flights' itineraries and affect other nearby flights as well. The net impact of such ripple effects can be modeled and quantified with a recursive penalty function that links all the relevant flights. In this paper, spatial separation is defined in terms of miles-in-trail (spatial MIT) restrictions while time separation is defined in terms of minutes-in-trail or time-based metering (TBM) i.e., temporal MIT restriction. In this sense, MIT restrictions encompass both time and space separations.

Figure 1 illustrates the relationship of a given flight to all the flights that will arrive at the point where a specific MIT restriction is applied during the time interval of interest in a simulation. In this Figure, time runs along the horizontal axis and different flights arrive at the point of the MIT restriction x at different times. Each black arrow represents a flight scheduled to arrive at a specific MIT restriction on the time axis associated with the restriction. Delayed flights or rescheduled flights are represented by arrows in red. Flights that do not arrive at the restriction on time are deleted and reinserted accordingly during the simulation as these events occur. Hence, an up-to-date sequencing of flights anticipated to arrive at a given restriction is always available and maintained during the simulation

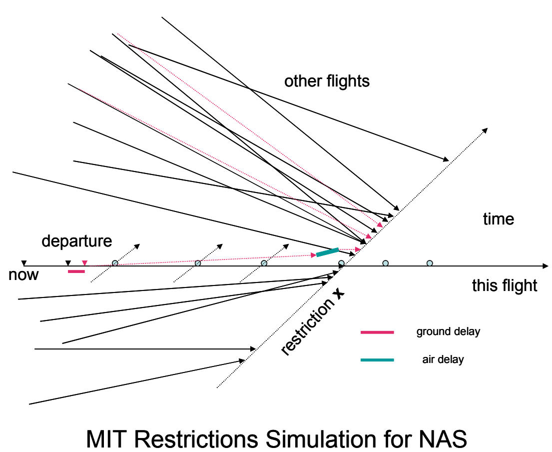

Figure 1. MIT Restrictions Simulation

From Figure 1, it is possible to anticipate the degree of congestion or potential MIT violations based upon initial scheduled flight plans during the simulated period. Such information is very important in identifying hot spots within the NAS that are results of scheduling conflicts. Ground delay/stop and/or air delay may then be designed to alleviate the anticipated hot spots or bottlenecks.

However, flights may be delayed, cancelled, diverted, or replaced by airlines due to various factors such as developing weather conditions, congestion at airports, procedural delays, or mechanical problems. From Figure 1, it is also clear that for any flight delay, there are ripple effects over the initial computed NAS MIT restrictions profile. First, the MIT delay will ripple down the remainder of the fight itinerary for each airframe. Second, for each restriction encountered, the MIT delay may also ripple over the group of flights anticipated to arrive at the point of restriction. To quantify these ripple effects, we take advantage of the recursive relationship between the expected MIT delays for each flight and the leading flight ahead of it.

Let t_i = estimated (scheduled or adjusted) arrival time for flight f_i at restriction X.

p_i = actual (recorded or simulated) arrival time for flight f_i at restriction X.

v_i = speed of flight f_i at restriction X.

S = miles-in-trail (spatial MIT) separation enforced at restriction X.

d_i = MIT delay (penalty) for flight f_i at restriction X.

From Figure 2, it is clear that we have the following recursive relation among flights subject to a given MIT restriction X:

$$d_i = \mathrm{MAX}\{0, d_{i-1} + t_{i-1} + S / v_{i-1} - t_i\} \qquad (1)$$

Equation (1) defines an MIT penalty function for the NAS. From the family $\{d_i\}_{i=1}^{|F|}$ where F is the family of all the flights scheduled to arrive at restriction X, the expected total MIT penalty at X can be quantified. Whenever t_i is adjusted or modified by a GS, GDP, or other TFM action, $\{d_i\}_{i=1}^{|F|}$ is updated according to the MIT penalty function (1) for all the relevant flights in F. Note that the ripple effects are embedded implicitly within the MIT penalty function. Note that in equation

(1) t_i stands for scheduled arrival time at the restriction for flight f_i; initially, t_i is simply the scheduled arrival time at restriction *X*, however, during a simulated period, flight f_i may be delayed as a result of airport or en route congestion, or GS, or GDP.

When such events occur, the delayed flights are removed from the ordered list of the flights and reinserted back to the ordered list such that correct and up-to-date t_i and d_i are maintained. The ripple effect of each change in every flight will automatically be reflected in the recomputed $\{d_i\}_{i=1}^{|F|}$.

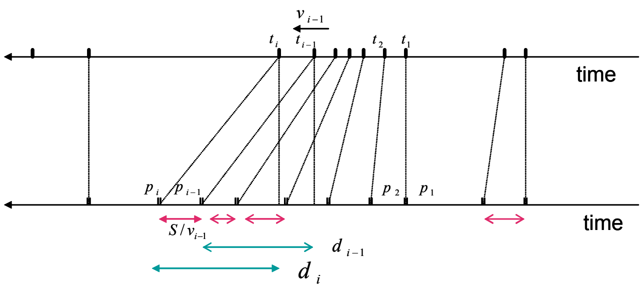

Figure 2. A Recursive MIT Penalty Relation

If equation (1) is implemented in a discrete event simulation package for the NAS, it will be possible for analysts to perform the following three tasks:

First, artificial MIT restrictions may be placed at key strategic points such as sector crossings, way points, or fixes to identify NAS bottlenecks resulting from scheduled demand. The hot spots and the extent of MIT space/time separation violation contributed by scheduling alone can be quantified in this way.

Second, once hot spots or bottlenecks within the NAS are identified, corrective actions or procedures may be designed to alleviate the problems. When alternative action plans are available, a comparison of their pros and cons can be assessed based on quantified performance metrics in simulated scenarios capturing the ripple effects of adjusted MIT restrictions over the entire NAS.

Third, such a tool will be able to evaluate or predict the overall NAS performance impact due to developing weather conditions, increased traffic demand, or reduced airport or airspace (sector) capacities. Furthermore, quantified NAS performance metrics can be captured in simulations to assess the relative merits of individual potential GS, GDP, MIT, or other TFM programs.

The main difference between the recommended recursive MIT penalty function vs. traditional summing of individual flight delays are the fact that the MIT penalty can be computed both locally and globally for all flights either as they are scheduled or during simulation. Such a recursive relationship offers realistic prediction prior to the actual occurrence of the events such that intelligent decision making for congestion relief policy can be better justified. The ripple effects of each MIT restriction across the NAS are also firmly embedded within the computed MIT penalty function.

3 TEST SCENARIOS

The implementation of the proposed MIT penalty function for all flights is made simpler with the use of Simulation Language with eXtensibility (SLX) (Brunner and Henriksen, 2003). We have already developed a NAS simulation in SLX, and the MIT capability was implemented in the context of the NAS simulation (Wieland, 2004). In our NAS simulation, a flight follows its scheduled itinerary from airport to airport. Various en route models with different degrees of detail for a wide range of aircraft types are available to analysts. In this study, flights follow a string of sectors. Takeoff, sector crossing, and landing are simulated with handoff request and acceptance. A handoff may be rejected or delayed if the target sector or airport has exceeded its capacity. We implemented two complementary data sets for MIT restrictions in the SLX simulation: the set of restrictions for each flight and the set of flights for each restriction. For any delay or change in scheduled flights, the MIT penalty is computed for all the relevant flights in these two sets.

Several scenarios were tested to provide quantified MIT-related delay profiles for baseline scenarios, with and without weather-related MIT restrictions. Figure 3 illustrates one such test scenario with 4 neighboring sectors, 8 airports, and 5 distinct routes with 2 MIT restrictions. In the baseline scenario, all sectors have identical capacity, identical Poisson flight arrival rates are maintained over all distinct routes. Weather conditions on S4 were simulated with reduced sector capacity and increased sector occupancy time. The impacts of imposing MIT restrictions at sector crossings, the interdependency between MIT penalties, reduced sector capacities, and ripple effects over neighboring sectors and flight streams were quantified with simulation for comparative analysis. When all relevant flights are tabulated for their end-to-end performance, it is possible to develop better understanding of and justification for specific GDP, GS, or other TFM options.

Figure 3 represents a simplified test in which the only variables in the model are arrival rate, airspace capacity, and the spatial MIT separation value. In reality, flights may be subject to additional constraints such as vectoring, slot assignments, and weather related amendments. The advantage of excluding these elements is to reduce the noise and unrelated factors in the model so that any shift in delay metrics is a direct consequence of the variation in model variables.

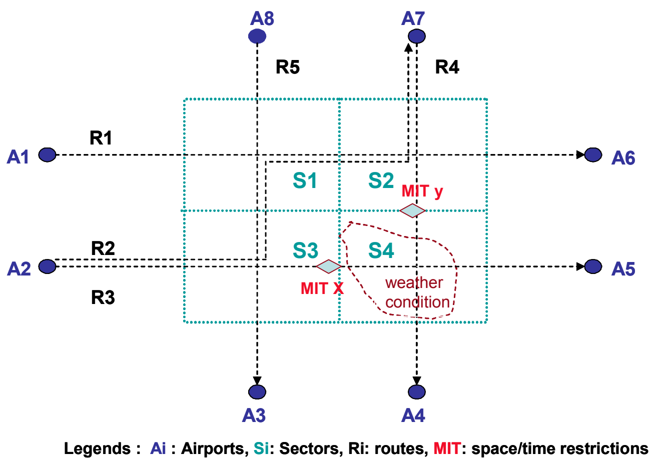

Figure 3. MIT Restrictions Simulation with SLX

4 SIMULATION RESULTS

In this section, we discuss the computed spatial MIT penalty profile for scheduled flights, the impacts of reduced sector capacity at S4, and the sensitivity to MIT space separation settings for MIT restrictions at x and y (see Figure 3). We also demonstrate cases when the penalty from MIT restrictions clearly exceeds its benefits. However, we also identify cases when a significant reduction in total trip time can be achieved with properly fine-tuned MIT restrictions.

Key parameters that determine the distribution of the MIT penalty $\{d_i\}_{i=1}^{|F|}$ given in equation (1) are:

C_j : the sector capacity at S_j

t_j : the sector occupancy time at S_j

R_x and R_y : the value of MIT separation at *x* and *y*.

λ_k : the intensity of traffic rate for each route R_k .

The key measurements of the MIT penalty are:

n: the total number of flights with $0 < d_i$

m: the instantaneous MIT metering queue size

D: total MIT metering penalty (= $\sum_{i=1}^{|F|} d_i$)

For a given scenario, m, n, and D can be determined by simulation. Figures 4, 5, 6 illustrate such a test case with various intensities of flights scheduled as Poisson arrivals for each route from 8 to 20 flights per route per hour. Note that the artifice of scheduling flights with exponential inter-departure times is only a reference point. In the real world flights can be scheduled more or less regularly than simulated here.

The sensitivity of expected total MIT metering penalty over time to the value of MIT spatial separation restriction (or metering) is illustrated in Figure 7. As the simulation clock moves ahead, flights are metered at the restriction x, and the expected total MIT delay penalty for the remaining flights scheduled to arrive at x is gradually reduced. Delayed flights will result in a reordering of the arrival sequence with ripple effects on the computed D for the remaining flights. In Figure 7, all model parameters for sector capacity, scheduled flights, sector occupancy time, and MIT restriction at x are fixed; while the value of spatial separation at MIT restriction x varies from 1 mile to 20 miles. Note that the more flights are separated at x, the larger the total MIT metering penalty D becomes and the later the last flight will be arriving at its destination. This illustrates how unnecessary MIT restrictions can produce large cumulative penalty across multiple flights.

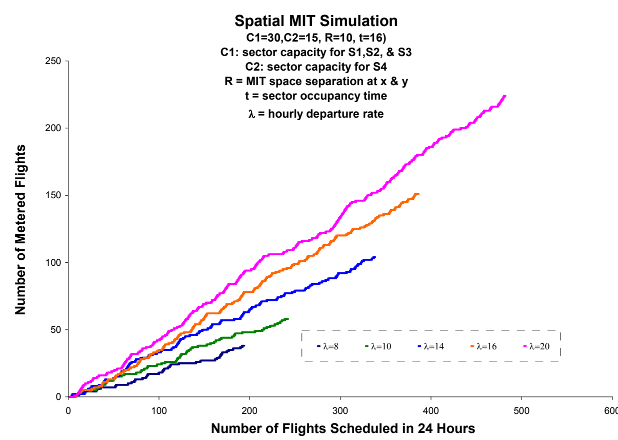

Figure 4. Intensity of MIT Metering Simulation

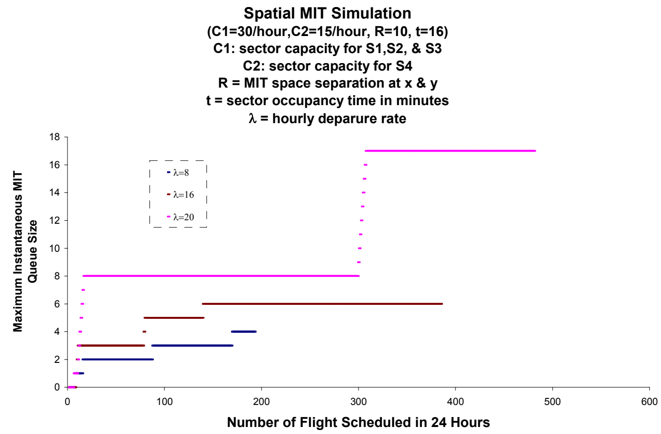

Figure 5. Instantaneous MIT Metering Queue Size

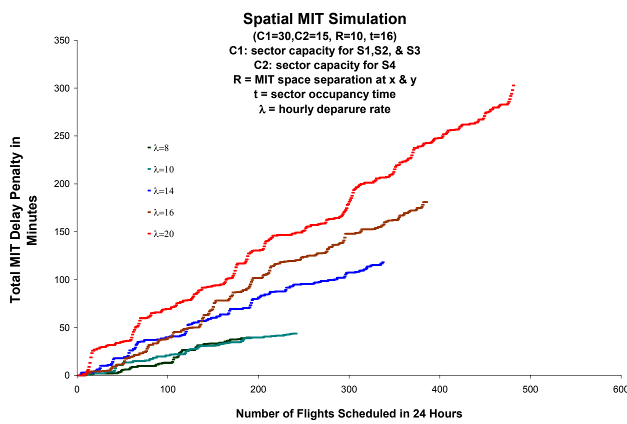

Figure 6. Total MIT Metering Penalty Simulation

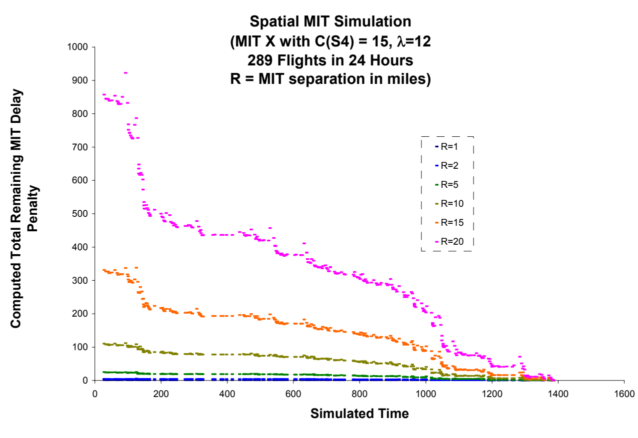

Figure 7. Sensitivity of D with Respect to R

One can also examine the relationship between sector capacity at S4 and the total MIT penalty across all relevant flights for a given MIT restriction x. Figure 8 is the result of such a simulation for MIT x in Figure 3, given the specific model parameter settings shown in Figure 8.

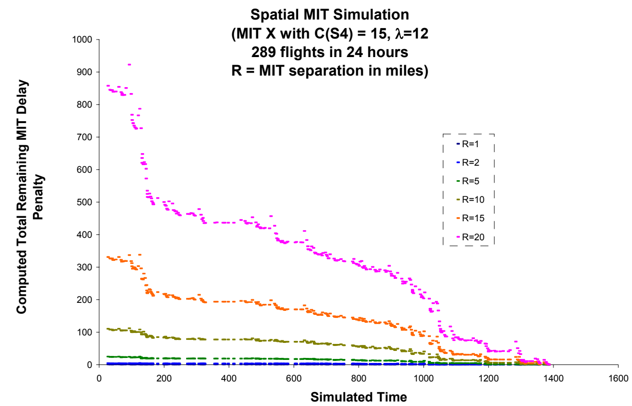

Figure 8. Sector Queuing Delay vs. MIT Penalty

Figure 8 suggests that MIT penalty can be beneficial only if the reduction of target sector queuing delay and sector occupancy time are larger than the expected or computed total MIT metering penalty D. Otherwise, the MIT metering penalty will become a real penalty as an added penalty to the total trip time for all flights. In Figure 8, we changed the capacity of the target sector, S4, from 2 to 30 flights per hour as C2 to C30. However, we did not increase the sector occupancy time of 16 minutes per flight. Clearly, the sum of MIT delays for all flights increases noticeably as the target sector's capacity is reduced from C5 to C2. To demonstrate the feasibility of such MIT metering benefits, we increased the sector occupancy for S4 from 16 minutes to 36 minutes and reduced the sector capacity of S4 from 15 to 5 flights per hour. We also provided sufficient sector capacity at the neighboring sectors S2 and S3 so that the MIT restrictions at x and y will not penalize flights with queuing delays at S2 and S3. Such a scenario mimics the situation of a severe thunder storm at S4. With increasing spatial separation at MIT restriction x, we observe that the reduction of queuing delay at S4 is greater than the total MIT metering penalty incurred at x. Figures 9 to 11 illustrate the results of such a scenario for R_x=0 to 40. Figure 9 tabulates flight arrival times. Figure 10 compares trip delay for all flights. Figure 11 quantifies the net benefit as reduction in total trip time for the entire population of 768 flights among all different routes. Figures 9 to 11 illustrate the feasibility of quantifying global benefits from implementing two spatial MIT restrictions around an area of poor weather conditions.

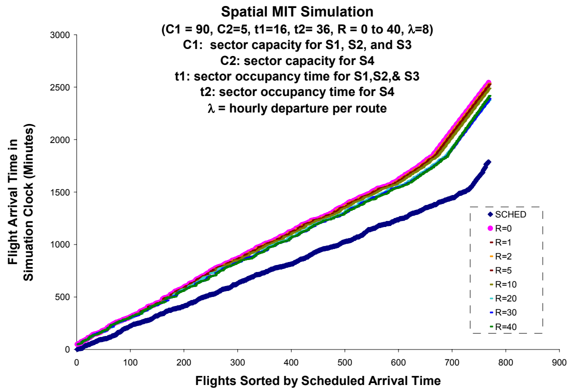

Figure 9. Flight Arrival Time with MIT Restrictions

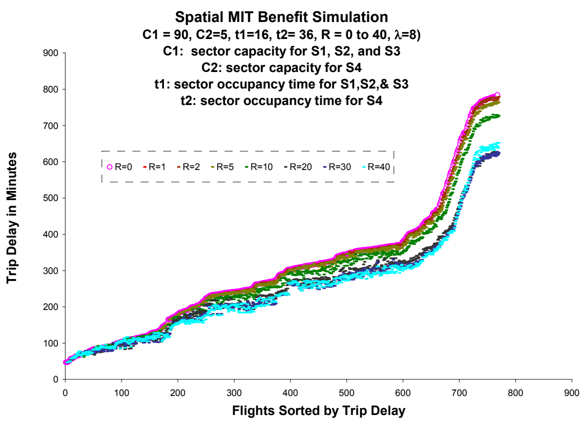

Figure 10. Flight Trip Delay with MIT Restrictions

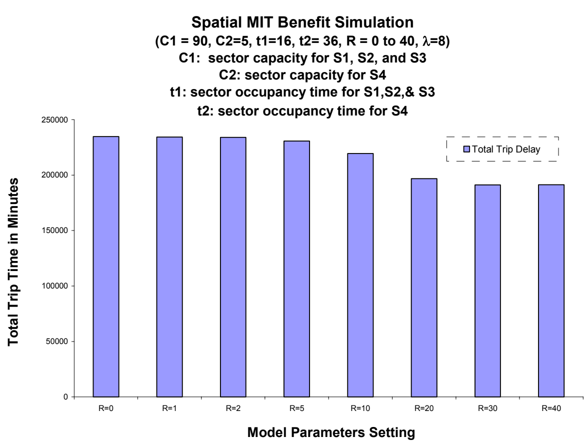

Figure 11. A Case of Beneficial Spatial MIT Restrictions

5 TEST CASE OVER THE NAS

The proposed MIT penalty function has been implemented in CAASD's latest NAS-wide air traffic control (ATC) simulation in SLX. The NAS is a very complex system with a large number of control parameters and procedures. Calibration of a NAS-wide ATC simulation is beyond the scope of this paper. However, it is possible to simulate a typical NAS scenario with a well-calibrated set of itineraries, airport configurations, airport and airspace (sectors) capacities, and nominal handoff procedures. In this example, the authors select a busy airport X among the 35 major airports and place spatial MIT restrictions for all flights entering the arrival terminal space. There are 12 en route sectors adjacent to the terminal airspace where airport X is located. Only 6 of the 12 sectors are currently used for flight arrivals. On a typical busy day, there are 63,000 flights across the NAS. Among the 63,000 flights, there are 1,500 flights arriving at airport X as they are reported in Enhanced Traffic Management System (ETMS). In our NAS-wide simulation, we calibrated model parameters to reflect a typical day for the NAS. Only the value of the spatial MIT separation for arrivals at airport X was changed from 0 to 12 miles. We also examined the impacts of spatial MIT separation for arrival flights at airport X for bad weather day with reduced airport capacity (at 70% of the normal airport arrival capacity).

Intuitively, without spatial MIT separation, airport X would be overloaded with arrivals and the penalty for airport arrival queuing delays could reach a high level. Excessive airport arrival queuing delays will also result in increased en route delays as handoffs for arrivals into terminal air space may be rejected. As the spatial MIT separation for arrivals increases, airport arrival queuing delay decreases while en route queuing delays for arrival flights increases. Hence, the tradeoffs between airport arrival queuing delays versus en route queuing delays for individual flights or all flights at airport X and their impacts both locally and globally for all flights across NAS can be quantified with simulation.

Figure 12 illustrates the simulation of spatial MIT separations for all flights arriving at airport X from each of the 12 en route sectors adjacent to the terminal airspace of airport X.

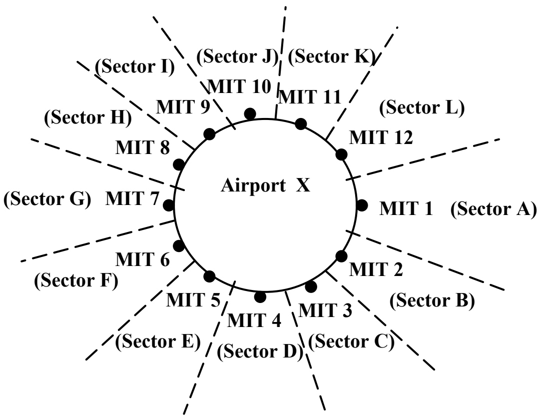

Figure 12. MIT Arrival Spatial Separation Simulation At Airport X

Figures 13 and 14 plotted the simulated results for average flight arrival queuing delays versus en route queuing delays for all the 63,000 flights at 35 major airports and the 1,500 flights arrived at airport X.

In Figure 13, average flight arrival at-gate delays are negative since flights are typically scheduled to arrive early. When airport X is operated under Instrument Meteorological Condition (IMC, which corresponds to bad weather) runway capacity, average airport arrival queuing delay at airport X and across the NAS, as a whole, are noticeably higher than that for Visual Meteorological Condition (VMC, or "normal"). In both cases, the advantage of spatial MIT separation may be lost due to rapid increase in average en route queuing delays as spatial separation increases. Figures 13 and 14 suggest that proper setting of spatial MIT restrictions for busy airports can be quite beneficial. Additional model calibration and validation are also very desirable.

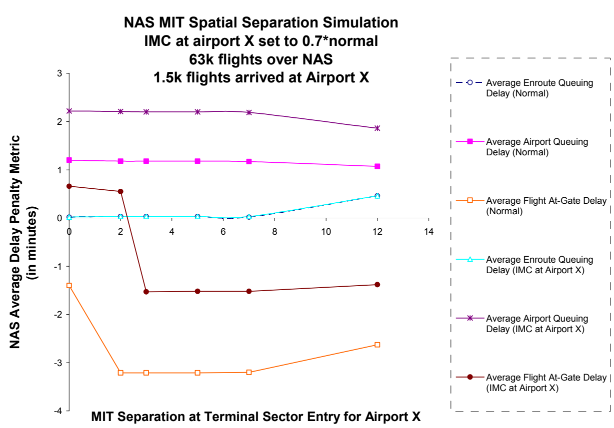

Figure 13. Global Impacts of Spatial MIT restrictions at Airport X over NAS

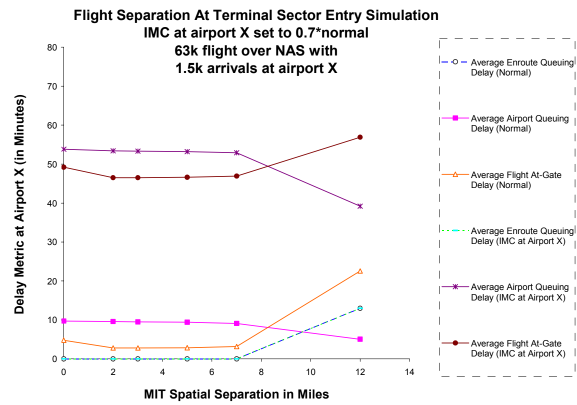

Figure 14. Local Impacts of Spatial MIT Restrictions At Airport X

6 CONCLUSIONS

The authors have provided an explicit spatial MIT metering equation that relates delay penalty between two consecutive flights crossing a sector boundary or fix. The concept of recursive flight delays computation is an extension of the flight delay tracking discussed earlier by Wang, Wieland, and Wojcik, (2001). It is demonstrated that such a recursive MIT penalty function can be implemented in a discrete event simulation for the entire NAS and be used as a tool to quantify the tradeoffs of various MIT programs against different model parameter settings such as flight schedules, sector occupancy, capacity, weather conditions, and imposed ground or air delays. The model can also be used for evaluating different TFM programs and provide quantitative performance metrics to help understand the impacts of different potential actions.

It is also shown that MIT restrictions can be harmful if misapplied. With the right model parameter settings, the results suggest that MIT restrictions can reduce the severity of performance degradation resulting from bad weather or high traffic loads.

Future work would include the integration of this newly developed tool into other simulation and modeling tools at MITRE for analysts and operators to help with analysis or, perhaps in the future, real time applications.

7 ACKNOWLEDGMENTS

Most of the work presented in this paper was performed as a MITRE Sponsored Research (MSR) project under Dr. Frederick P. Wieland. The authors are indebted to Dr. Leonard A. Wojcik for his review of the initial draft and editorial improvements. Dr. Robert S. Conker and Mr. Jerome E. Freedman provided test cases in NAS for partial verification of initial coding of the MIT penalty function used in the simulation. Figures 13 and 14 are a direct application of a test case provided by Freedman. Ms. Sheila M. McHale provided final editorial improvements to the paper. We also acknowledged helpful assistance from Mr. Thomas C. Holden and Ms. Shane L. Miller for their expertise on Collaborative Routing Coordination Tools (CRCT) and MIT-related software and documents. Finally, encouragement with very helpful suggestions and advice from the referees of the *Air Traffic Control Quarterly* are also greatly appreciated by the authors.

REFERENCES

Ball, C. G., Connolly, K. A., and Wanke, C.R., 2003, "TFM Integrated Impact Assessment Capability: Results of November 2002 Evaluation," Report No. F045-B03-045, McLean, VA.: The MITRE Corporation.

Beaton, E. K., et al., 2002, *CRCT Capabilities Detailed Functional Description*. Report No. MTR00W000302R1. McLean VA.: The MITRE Corporation.

Brunner, D. T. and Henriksen, J. O., 2003, *Simulation With SLX,* Wolverin Software Corporation.

Ostwald, P. A., Esters, S. L., Efllman, L., Miller, S. L., Topiwala, T., and Wanke, C. R., September, 2003, *Detailed Functional Requirements for an Integrtaed Miles-In-Trail and Rerouting Decision Support Capability*, MP03W0000171, McLean, VA: The MITRE Corporation.

Wang, P. T. R., Wieland, F. P., and Wojcik, L. A., 2001, *DPAT Flight Delay Modeling and Itinerary Tracking*, ATIO-3, 2001 ATIO Conference, American Institute of Aeronautics and Astronautics (AIAA).

Wanke, C. R., et al., 2003, *Human-in-the-loop evaluation of multi-strategy traffic management decision support capability*. Proceedings of the 5th Eurocontrol/FAA ATM R&D Seminar, Budapest, Hungary

Wieland F. P., 2004, "Next Generation Model of the National Airspace System," MITRE Technology Program briefing, McLean, VA.: The MITRE Corporation.

AUTHOR BIOGRAPHIES

PAUL T. R. WANG has been with the The MITRE Corporation's Center for Advanced Aviation System Development (CAASD) since 1984. He graduated from The Ohio State University with MS and PhD degrees. His interests include optimization, modeling, simulation, and queuing theory. Paul's work on simulation and optimization were published by WSC, ORSA, DASC, *Air Traffic Control Quarterly*, AIAA, and the *Journal of Computer Networks and ISDN Systems*. His e-mail address is <pwang@mitre.org>.

CRAIG R. WANKE is a senior principal engineer also with MITRE/CAASD. For the past several years, he has been working on advanced decision support systems for traffic flow management applications. He holds S B, MS, and PhD degrees in Aeronautics and Astronautics from the Massachusetts Institute of Technology. His e-mail address is <cwanke@mitre.org>.

FREDERICK P. WIELAND is with The AEgis Technologies. He was a Project Manager at MITRE/CAASD. He has been active in the simulation field for over 20 years, and currently is leading the development of system-wide models of the National Airspace System. He holds a PhD from George Mason University in applied probability theory and a BS degree from Caltech in Astronomy. His e-mail address is <fwieland@aegistg.com>.

DISCLAIMER