Fast-time Simulation of an Automated Conflict Detection and Resolution Concept

Robert Windhorst[*] and Heinz Erzberger[†]

NASA Ames Research Center, Moffett Field, CA, 94035

This paper investigates the effect on the National Airspace System of reducing air traffic controller workload by automating conflict detection and resolution. The Airspace Concept Evaluation System is used to perform simulations of the Cleveland Center with conventional and with automated conflict detection and resolution concepts. Results show that the automated conflict detection and resolution concept significantly decreases growth of delay as traffic demand is increased in en-route airspace.

I. Introduction

In the current National Airspace System (NAS), air traffic controllers manually control the flights operating within their sector. The maximum number of flights that they can safely control is called the Monitor Alert Parameter (MAP). MAPs are strongly affected by controller workload, which is composed of several tasks one of which is Conflict Detection and Resolution (CD&R). At times, traffic demand grows to a level, which if unmanaged, will cause the number of flights in a sector to exceed its MAP. When this is predicted to occur, a Traffic Management Coordinator (TMC) initiates a Traffic Management Initiative (TMI), such as a Ground Delay Program (GDP) or a Miles-In-Trail (MIT) restriction, to reduce traffic flow into the affected sector and maintain its flight counts below its MAP. These initiatives produce large delays within the NAS across a broad range of flights.

Reference 1 proposes a future concept for automating CD&R called the Advanced Airspace Concept (AAC). It shows that the AAC could increase controller productivity by over 100 percent and increase the airspace capacity by 100–200 percent, thereby allowing MAPs to be increased and reducing the need for TMIs. This paper compares the delays produced by the AAC with those produced by the conventional concept described above as demand is increased. The approach is to use the Airspace Concept Evaluation System (ACES) [2-5] to perform simulations of both concepts using 1-, 1.5-, and 2-times current levels of traffic demand. Delays are calculated for each simulation and compared.

First, the approach used to simulate the concepts is presented. Then, the two algorithms that are used in the simulations (a Traffic Flow Management algorithm and an automated CD&R algorithm) are summarized. Finally, the results of the simulations in terms of delays are given.

II. Approach

The approach was to use the ACES to perform simulations of Cleveland Center (ZOB). ACES is a model of the NAS, consisting of software agents that model flights, Air Traffic Control (ATC), Traffic Flow Management (TFM), the Air Traffic Control System Command Center (ATCSCC), and the Airline Operational Centers (AOCs). Individual ATC and TFM agents are devoted to each airport, Terminal Radar Approach Control (TRACON), and Air Route Traffic Control Center (ARTCC) in the simulation. References 2-5 provide a more detailed description of ACES.

ACES can be configured to simulate only the parts of the NAS that are salient to the experiment. In this case, only the flights entering ZOB airspace were simulated. Other flights were filtered out of the simulation. Each of the simulated flight's departing and arriving airports and corresponding TRACONs were also simulated. Similarly, the ARTCCs that the flights entered when not in ZOB were also simulated.

One of the inputs to ACES is the flight demand set. This set consists of all of the flights in the simulation and their departure times and flight plans. The 1-times current level of traffic set used in this experiment was created from data recorded from the Enhanced Traffic Management System (ETMS) for a 24-hour period on 21 April 2005.

---

[*] Chief, Aerospace Operations Modeling Branch, Mail Stop 210-10, Senior Member AIAA.
[†] Adjunct Professor, University of California at Santa Cruz, Mail Stop 210-10, Fellow AIAA.

It was grown to create the 1.5- and 2-times sets by duplicating flights that were selected according to city-pair growth ratios and perturbing their departure times[6].

Another input to ACES is the MAP for each sector. The TFM agents constrain the sector flight counts during simulation so that they are less than or equal to the MAP. The TFM Algorithm section of the paper describes how the agents work.

During pre-processing, ACES computes unconstrained boundary crossing and airport arrival times for each flight in the simulation. At the end of the simulation, delays are calculated by comparing the simulation times to the unconstrained times. Delays are presented and compared for all the simulations.

Simulations of the conventional concept and the AAC were performed. The conventional concept was modeled by turning off the automated CD&R algorithm and using the TFM algorithm to constrain sector flight counts to their present day MAPs. The simulation did not include a model of a controller performing CD&R, because it was assumed that delays produced by CD&R are much smaller than those produced by TFM. This assumption is supported by the results of this study. The AAC was modeled by turning on the automated CD&R algorithm and turning off the TFM algorithm. Turning off the TFM algorithm is the same as eliminating the MAP constraints on sector counts. The reason for turning the TFM off was that Ref. 1 showed that the automated CD&R algorithm could safely accommodate up to 2-times current levels of traffic and the largest flight demand set used in this study was 2-times. However, the results of Ref. 1 were based on the number of conflicts encountered and the volume of the airspace the automated CD&R algorithm had available for resolving them. En-route weather events, equipment failures, mixed equipage flights, and controller workload not associated with CD&R were not considered. Further study of these factors is required to determine how much the MAPs could be safely increased in the real NAS.

Because the study was focused on concepts for performing CD&R in en-route airspace, airport departure and arrival rates were not constrained during any of the simulations. Another study named the Preliminary System-wide Concept Assessment[7] analyzed the interactions between airspace and airport constraints.

Table 1 illustrates which algorithms were used to model the concepts. The following sections summarize the TFM and automated CD&R algorithms used in the simulations.

Table 1. Study matrix showing which algorithms were used to model the concepts

Concept	Traffic Flow Management	Automated Conflict Detection and Resolution
Conventional	On (current MAP values)	Off
Advanced Airspace	Off	On

III. Traffic Flow Management Algorithm

The ACES TFM algorithm does not directly model a particular TMI used in a specific region of the NAS today. For example, it does not model a MIT restriction imposed between a pair of ZOB sectors. However, it does implicitly model the gross effects of all the TMIs acting in the NAS at once by constraining all NAS sector flight counts and all NAS airport arrival and departure rates.

TFM relies on an accurate and updated flight schedule to predict where and when congestion will occur within the system. The flight schedule contains departure times, sector-crossing times, and arrival times for each flight scheduled to use the NAS. The flight schedule can be sorted and filtered to build departure and arrival schedules at airports and schedules of flights passing through sectors.

In the real NAS, the flight schedule is stored and updated by the ETMS. Weeks in advance of departure, AOCs insert schedules and flight plans for each of their flights into the system. The departure times and flight plans are updated as needed by the AOC or Air Traffic Service Provider until one-half hour before departure. TMCs visualize ETMS data via the Traffic Situation Display (TSD). Using this tool, they can identify where and when congestion is predicted occur and select a TMI to mitigate it. When flights are maneuvered or re-routed, their schedule, which is stored in ETMS, is updated to reflect the change of plan.

As discussed in Section II, a pre-processing function named "configuration" generates the initial, unconstrained flight schedule in ACES. In this function, flights are selected individually from the flight demand set, and, using their submitted departure times and flight plans, their trajectories are calculated from departure to arrival. The flight schedule is built from the submitted departure times and the computed sector crossing and arrival times, which were extracted from the predicted trajectory.

The flight schedule is updated as the simulation progresses by a process called "flight data distribution." This function updates the flight schedule every time an event occurs within the simulation that changes the schedule. Common events that trigger a recalculation of the schedule in ACES are delayed departure times, re-routes, and

maneuvers. Each time one of these events occurs, "flight data distribution" re-generates, based on the new event, the trajectory of the affected flight from the event time to gate arrival and updates the schedule with the new departure, sector crossing, and/or arrival times.

The TFM algorithm in ACES is distributed across airport, TRACON, and ARTCC TFM agents and the ATCSCC agent. The algorithm always begins with one of three types of assessments. These assessments are performed periodically on a 15-minute interval. The three types are airport departure rate, airport arrival rate, and sector count. Obviously, the airport arrival and departure rate assessments are performed in the airport TFM agent. The sector count assessment is performed in the ATCSCC agent. Once the assessment is complete, its results are distributed to the other TFM agents in the system via messages. The results are boundary-crossing restrictions that are enforced by the ATC agents imposing delays on particular flights.

The TFM departure and arrival assessments were not used here. However, they are described, because they share functionality with the sector count assessment, which was part of this study.

A. Airport Departures

The airport departure assessments are performed every 15 minutes over a 6-hour time horizon for each airport in the ACES simulation. They are carried out by the airport TFM agents and compare the requested departure schedules with the Airport Departure Rate (ADR) constraints applied at each airport. The airport TFM agents set the ADRs, which are a function of the requested departure schedule, the predicted arrival schedule, and the weather conditions. The requested departure and predicted arrival schedules for each airport are obtained from the flight schedule. In addition, they are updated as needed by "flight data distribution." If an assessment determines that a requested departure schedule will cause the departure rate to exceed the ADR, departures for that airport are re-scheduled. The re-scheduling algorithm spaces the departure times such that the departure rate remains at or below the ADR. The spacing is achieved by delaying individual flight departure times. Departure times are never moved ahead, and once a delay is imposed on a flight, that flight can never recover the lost time, even if the circumstances requiring the delay have changed and it has not yet departed.

Imposing delays on flights that have not yet departed is a simple task, because no NAS domains other than the departure airport are involved. The airport TFM agent simply requests its respective ATC agent to delay the flight at the gate.

B. Airport Arrivals

Similar to the departure assessments, the airport arrival assessments are performed every 15 minutes over a 6-hour time horizon for each airport in the ACES simulation. They are carried out by the airport TFM agents and compare the predicted arrival schedules with the Airport Arrival Rate (AAR) constraints applied at each airport. The airport TFM agents set the AARs, which are a function of the requested departure schedule, the predicted arrival schedule, and the weather conditions. The requested departure and predicted arrival schedules are obtained from the flight schedule that was built during configuration. In addition, they are updated as needed by flight data distribution. If an assessment determines that the predicted arrival schedule will cause the arrival rate to exceed the AAR, arrivals for that airport are re-scheduled. The re-scheduling algorithm spaces the arrival times such that the arrival rate remains at or below the AAR. The spacing is achieved by delaying arrival times. Arrival times are never moved ahead.

Up to the re-scheduling point, the arrival assessment is similar to the departure assessment. However, when flight arrivals are pushed back, the delays need to be imposed on flights while they are in NAS domains upstream of the arrival airport. Delaying flights upstream is accomplished by sending arrival restrictions via messages to the upstream TFM agents. The upstream TFM agents are the TFM agents of the NAS domains (TRACON, ARTCC, or departure airport) that the flights being delayed pass through upstream of the arrival airport. Arrival restrictions are generated and sent out on a per-flight basis.

The first NAS domain upstream of the arrival airport is always the arrival TRACON. A message containing an airport arrival restriction for that flight is sent to the TRACON TFM agent. Upon receiving that message for a particular flight, the TRACON TFM agent identifies when the flight will enter the TRACON by searching its flight-crossing schedule. Then it calculates how much delay is required to be imposed on that flight as it passes through the TRACON to meet its airport arrival restriction. If that amount of delay is under the maximum allowable delay restriction for flights in a TRACON, the delay is imposed and the process ends. If that amount of delay is above the maximum allowable delay restriction for flights in a TRACON, the maximum allowable delay is imposed and a TRACON arrival restriction message, specifying the amount of residual delay still needing to be absorbed, is passed to the upstream ARTCC TFM agent.

3

American Institute of Aeronautics and Astronautics

Upon receiving a TRACON arrival restriction for a particular flight, the ARTCC TFM agent uses the same process as that of the TRACON TFM agent. Restrictions are passed upstream until all of the delay is allocated or the restrictions are passed all the way back to the departure TRACON and airport. If the restrictions reach the departure airport, all of the residual delay is absorbed by a gate departure delay.

TFM agents do not actually impose delay on flights; ATC agents do. The TFM agents pass sector exit restriction messages to the ATC agents, which then delay the flights. ATC agents delay flights differently depending on their location in the NAS, because they are propagated through different domains by varying level-of-fidelity models. In the ARTCC domains, a four degree-of-freedom dynamic model propagates flights, whereas in the TRACON domains, a transit-time model propagates flights. ARTCC ATC agents impose delay by path-stretching the flight's trajectory, whereas TRACON ATC agents impose delay by simply adding it to the transit-time of the flight.

C. Sector Flight Counts

Similar to departure and arrival assessments, the sector count assessments are performed every 15 minutes over a 6-hour time horizon in ACES. The ATCSCC agent performs this assessment. It compares predicted sector counts to their MAPs. Predicted sector counts are generated using the sector boundary crossing times in the flight schedule that was generated during configuration. The assessments identify sectors and corresponding 15-minute time intervals within the 6-hour horizon where the flight count is predicted to exceed the MAP.

Once the assessment is complete, congestion alert messages, containing sector identifiers and corresponding time-intervals, are sent to the ARTCC TFM agents. These agents analyze each sector during its congested time-interval to identify the last flight that entered the sector and caused it to exceed its MAP. A delay is imposed on the flight upstream of where it enters the sector, so that it does not enter until after another flight has exited. In this way, the sector count is not permitted to exceed the MAP. If all of the necessary delay cannot be absorbed by the flight within the ARTCC, the ARTCC TFM agent passes an arrival restriction, specifying the residual delay back to the next upstream TFM agent, and the delay is propagated exactly the same way that it is propagated to relax the AAR constraint.

Figure 1 exemplifies the results of the process. It contains a plot illustrating maximum flight count during 15-minute time-intervals vs. time for the Cleveland Center sector 66 (ZOB66) for the 2-times demand simulation. Three curves are shown. The horizontal "capacity" line is the MAP value for ZOB66. The "unconstrained" curve shows the flight counts when the simulation was executed without the TFM algorithm. The "constrained" curve shows the flight counts when the simulation was executed with the TFM algorithm. Figure 1 shows that, at times, the flight counts exceed the MAP by several counts. These small constraint violations are typical of real NAS operations. In fact, MAP values are selected such that small overruns are manageable by controllers. Plots of sector counts for the other sectors demonstrate similar effects.
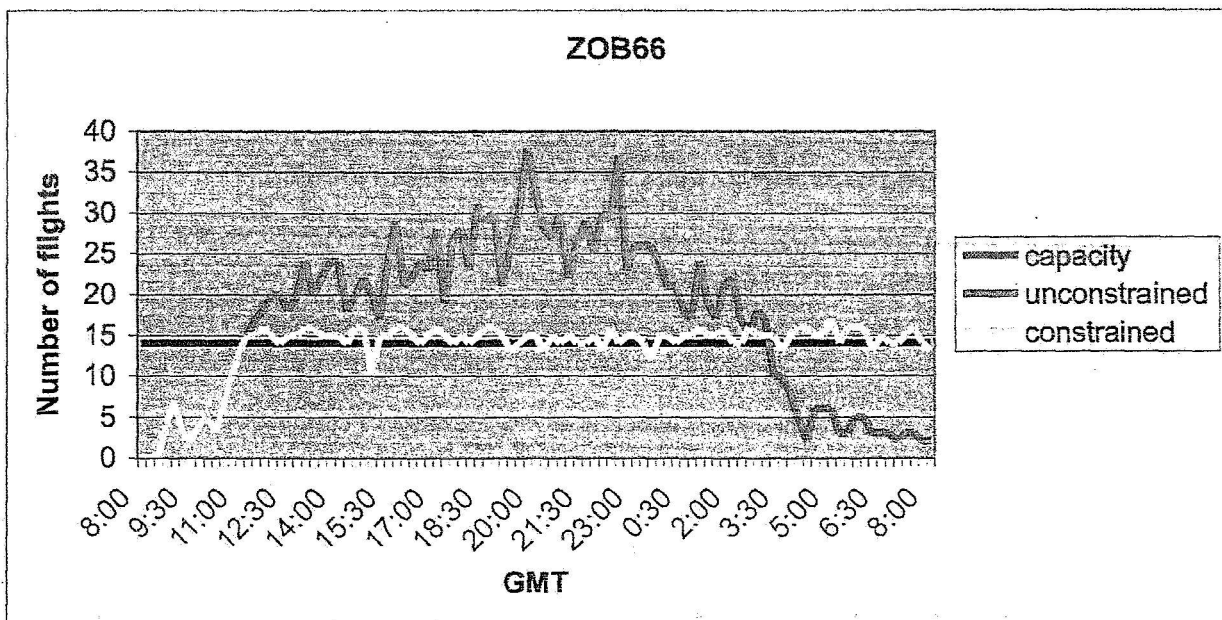

Figure 1. ZOB66 Flight count vs. simulation time

IV. Automated Conflict Detection and Resolution Algorithm

The automated CD&R algorithm used in this study is described in detail in Ref. 8. It consists of two main parts: the conflict detector (CD) and the conflict resolver (CR). The CD predicts where and when conflicts will occur and sends them to the CR. The CR generates resolutions that resolve the conflicts. The resolutions conform to current flight plan standards of operation and could be parsed into a sequence of conventional clearances.

The CD initiates the CD&R algorithm. It activates periodically every 10 sec. to 2 min., depending on user preference, and searches for conflicts that would occur within the next 20-minute time horizon. The positions of the flights during that time horizon are predicted by the Trajectory Synthesizer (TS), which propagates the position of a flight forward in time using a four degree-of-freedom flight dynamics model and aircraft-specific performance models. The TS assumes that the flights are conforming to their flight plans. The position of each flight is stored at times with 10-second intervals. At each time the position of each flight relative to the positions all other flights in the airspace is checked against the separation criteria, which are user adjustable. Typically, they are set to 1,000 ft vertically and 5nm horizontally plus some buffers to account for sensor errors. A conflict is predicted when any flight pair violates the criterion during the 20-minute time horizon. A list of predicted conflicts is sent to the CR, which sorts it by time to loss of separation (LOS). The CR algorithm resolves the predicted conflict with the earliest predicted LOS and then moves down the list.

Given a predicted conflict, the CR generates a resolution for one of the flights in the pair. The resolution resolves the predicted conflict and merges the flight back onto its original flight plan. Then, the TS generates a trajectory conforming to the resolution, and the CD checks that trajectory against the trajectories of the other flights in the airspace for the selected time horizon (up to 20 min.) to insure that it does not create any secondary conflicts. If it does create secondary conflicts, it is rejected, and the CR generates a new resolution. The process starts with one flight of the pair and moves on to the other. It is repeated until a resolution is found that does not generate any secondary conflicts and resolves the primary conflict or until the number of iterations allowed is exceeded, in which case it is allowed to go unresolved. Most unresolved conflicts are resolved the next time the algorithm is triggered, which is within two minutes. Conflicts that repeatedly go unresolved, and have a predicted time to LOS of less than a specified parameter (set to 2 minutes here), would be resolved by an independent safety critical function not modeled in this analysis, but described in Ref. 1.

V. Results

Table 2 lists the delays produced in the simulations of the conventional concept for the 1-, 1.5-, and 2-times demand sets. The average delay per flight sharply increases with increasing demand. At 1.5-times demand, it already exceeds an operationally acceptable level. The delays at 2-times demand are shown here to further illustrate the complete impracticality of the conventional concept when traffic density dramatically exceeds the MAPs required for manual CD&R.

Table 2. Delays produced in the simulations of the Conventional Concept.

	1x	1.5x	2x
number of flights	7,602	11,983	15,495
average delay in seconds per flight	38	1990	6014

Table 3 lists the delays produced in the simulations of the AAC. These delays were published in Ref. 7. Note that the delays here are in seconds of delay per conflict as opposed to per flight, which was used above. These simulations were performed using a slightly different demand set. When demand is increased 2-times, the number of conflicts dramatically increases. However, the average delay per conflict does not proportionately increase. This is because, even at this 2-times level of demand, relatively few secondary conflicts are being produced by the resolutions, indicating that there is a large amount of airspace still available with which to perform resolutions.

Table 3. Delays produced in the simulations of the Advanced Airspace Concept

	1x	2x
number of flights	7,602	15,495
number of conflicts	830	3,574
average delay in seconds per conflict	22	25

Table 4 compares the total delays in seconds produced in each simulation. The conventional concept produces more delay, and the delay increases proportionately faster as demand is increased. The automated CD&R algorithm imposes delay more judiciously than the TFM algorithm. Whereas the TFM algorithm imposes delays on broad groups of flights, the automated CD&R algorithm imposes delay on only the flights that are predicted to be in conflict. Since the CD&R algorithm can accommodate the traffic densities produced by the 2-times level of current demand, the traffic densities allowed within a Next Generation Air Transportation System incorporating the AAC could be increased significantly. This increase would substantially decrease delays due to TMIs, because such initiatives would seldom be triggered.

Table 4. Comparison of total delay in seconds produced by the Conventional Concept and the Advanced Airspace Concept

	1x	2x
Conventional Concept	288,876	93,186,930
Advanced Airspace Concept	18,260	89,350

VI. Conclusions

The delays produced by a conventional concept for performing conflict detection and resolution were compared to the delays produced by a concept that automated conflict detection and resolution. Results showed that the delays produced by the automated concept were smaller and grew slower as demand was increased than those produced by the conventional concept. The automated concept accommodated 2-times level of current demand traffic without massive increases in average delay per conflict, indicating that there was sufficient airspace to resolve conflicts without creating secondary conflicts. This concept, if incorporated within the Next Generation Air Transportation system, could enable a large increase in airspace capacity without introducing large delay.

Acknowledgments

References

[1]Erzberger, H., "Transforming the NAS: The Next Generation Air Traffic Control System," NASA/TP-2004-212828, October 2004.

[2]Sweet, D. S., Manikonda, V., Aronson, J. S., Roth, K., and Blake, M., "Fast-Time Simulation System for Analysis of Advanced Air Transportation Concepts," *AIAA Modeling and Simulation Technologies Conference and Exhibit*, AIAA 2002-4593, Monterey, CA, Aug. 2002.

[3]Couluris, G., Hunter, G., Blake, M., Roth, K., Sweet, D., Stassart P., Phillips, J., and Huang, A., "National Airspace System Simulation Capturing the Interactions of Air Traffic Management and Flight Trajectories," *AIAA Modeling and Simulation Technology Conference and Exhibit*, AIAA 2003-5597, Austin, TX, Aug. 2003.

[4]Meyn, L., Windhorst, R., Roth, K., Manikonda, V., Van Drei, D., and Hunter, G., "Build 4 of the Airspace Concept Evaluation System," *AIAA Modeling and Simulation Technologies Conference and Exhibit*, AIAA 2006-6110, Keystone, CO, Aug. 2006.

[5]Windhorst, R., Meyn, L., Manikonda, V., Carlos, P., and Capozzi, B., "The Airspace Concept Evaluation System Architecture," *AIAA Guidance, Control, and Navigation Conference and Exhibit*, AIAA 2006-6231, Keystone, CO, Aug. 2006.

[6]Huang, A., Scheicher, D., and Hunter, G., "Future Flight Demand Generation Tool," *AIAA 4th Aviation Technology, Integration, and Operations (ATIO) Forum*, AIAA 2004-6400, Chicago, IL, Sept., 2004.

[7]"Virtual Airspace Modeling and Simulation System-Wide Concept Report," Project report, Virtual Airspace Modeling and Simulation Project, NASA Ames Research Center, June 23, 2006.

[8]Erzberger, H., "Automated Conflict Resolution for Air Traffic Control," Paper No. 083, ICAS, Sept. 2006, Hamburg, Germany.