

# The Traffic Management Advisor

William Nedell

San Jose State University, San Jose CA 95102

Heinz Erzberger, Frank Neuman

NASA Ames Research Center, Moffett Field CA 94035

## ABSTRACT

The Traffic Management Advisor (TMA) comprises algorithms, a graphical interface and interactive tools for controlling the flow of air traffic into the terminal area. The primary algorithm incorporated in it is a real-time scheduler which generates efficient landing sequences and landing times for arrivals within about 200 n.m. from touchdown. A unique feature of the TMA is its graphical interface that allows the traffic manager to modify the computer generated schedules for specific aircraft while allowing the automatic scheduler to continue generating schedules for all other aircraft. The graphical interface also provides convenient methods for monitoring the traffic flow and changing scheduling parameters during real-time operation.

## INTRODUCTION

Although automated decision systems for air traffic control (ATC) have been investigated for at least two decades, attempts to implement these systems in the current ATC environment have largely failed. Among the reasons for this failure are the use of obsolete ATC computers and displays, which are preventing the implementation of advanced concepts, and a tendency of developers to underestimate the complexity of automating even simple ATC functions.

Recently, the prospects for introducing higher levels of automation have improved because of two concurrent developments. First, a new generation of controller suites incorporating color graphics workstation technology together with new ATC host computers will remove many of the limitations impeding the implementation of automation concepts. The new controller suites, which will become operational in the mid-1990s, are the key element of the Federal Aviation Administration (FAA) Advanced Automation Systems (AAS). Second, recent research has provided new insights into the appropriate role of automation in ATC and has yielded promising methods for designing such systems.

The need for an effective controller-system interface imposes the most critical design constraint on ATC automation tools. To meet this constraint, the interface makes extensive use of on-screen switches and menus selectable by manipulating a mouse or trackball. Such

techniques improve the interface by minimizing the need for time-consuming and distracting keyboard entries. Also, computer-generated advisories are transformed, when possible, into a graphical format that enhances rapid perception of advisory information.

The report begins with an overview of the automation concept. This is followed by a detailed description of the Traffic Management Advisor. A description of the Descent Advisor can be found in reference 1, and a description of Final Approach Spacing Tool can be found in reference 2.

It is strongly recommended that elements of this system be implemented at an ATC facility for evaluation on a non-interfering basis. In view of past experience, such operational testing is an essential step in validating automation concepts for ATC. Past attempts at implementing automation tools have failed in part because they only worked well under the carefully controlled conditions of the laboratory. Unfortunately, such conditions rarely exist at ATC facilities. An automation tool which cannot handle the frequent departures from the quiescent design state will be quickly abandoned by controllers. Yet these are the times when automation assistance is most needed. Thus testing under realistic conditions can only be attained at operational ATC facilities. Then, the FAA can confidently decide what elements of this system warrant implementation in the AAS.

## AUTOMATION SYSTEM CONCEPT

Figure 1 gives a diagrammatic representation of the overall system. Its key ground-based elements are the Traffic Management Advisor (TMA), the Descent Advisor (DA), and the Final Approach Spacing Tool (FAST). The functions of each element and the relationships between elements are discussed below.

The primary function of the TMA is to plan the most efficient landing order and to assign optimally spaced landing times to all arrivals. These time schedules are generated while aircraft are 150 to 200 n.m. from the airport. The TMA algorithm plans these times such that traffic approaching from all directions will *merge on the final approach* without conflicts and with optimal spacing. The TMA also assists the Air Route Traffic

Control Center (ARTCC) Traffic Manager in rerouting traffic from an overloaded sector to a lightly loaded one, a process known as gate balancing. Another function of the TMA is to assist the Center Traffic Manager in efficiently rerouting and rescheduling traffic in response to a runway reconfiguration or a weather disturbance. In general, the functions of the TMA involve assisting the Center Traffic Manager in coordinating and controlling the traffic flow between Centers, between sectors within a Center, and between the Center and the Terminal Radar Approach (TRACON) Facility. Moreover, the TMA must permit the Center Traffic Manager to specify critical flow control parameters such as runway acceptance rate and to override computer generated decisions manually.

are then transformed by the TMA to gate arrival times by subtracting the time to fly from the gate to touchdown, and are sent to the DAs at the appropriate arrival areas.

Upon receiving these STAs the DA algorithm generates cruise and descent clearances which controllers can use to keep aircraft on schedule. For aircraft that drift off their planned time schedules, the controller can request revised clearances that correct such time errors to the extent possible. If this concept is implemented in today's environment, the controller would have to issue the clearances by voice, but in the near future it will be more efficient to issue them via the proposed ground-to-air data link.

The TRACON controllers take over control of traffic at the feeder gates. They merge the traffic converging on the final approach path while making sure that aircraft are properly spaced. If the Center controllers have delivered aircraft at the gates on time using the DA tools, the TRACON controllers ordinarily will need to make only small corrections in the relative positions of aircraft to achieve the desired spacing. The FAST assists the controller in making these corrections with high accuracy and a minimum number of heading vectors and speed clearances. Achieving precise spacing between aircraft on final approach ensures that landing rates will always be close to the theoretical capacity of the runway.

Another type of tool designed for the TRACON controller is the Tactical Advisor. This tool helps the controller to replan traffic quickly in response to several special situations, such as missed approaches, runway changes, and unexpected conflicts.

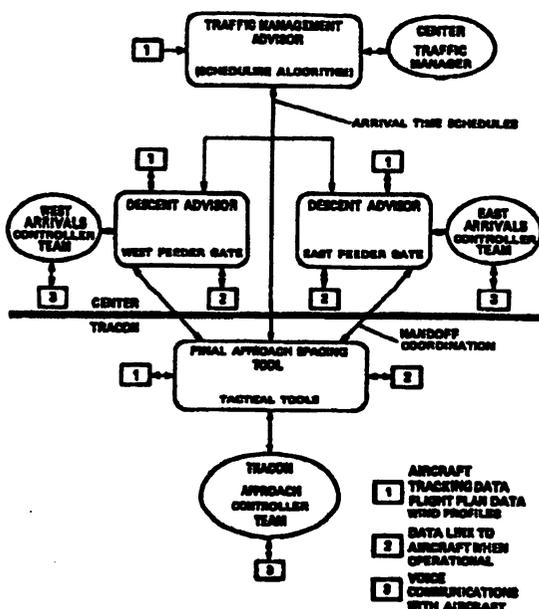


Figure 1.-Automation concept.

At a Center, the controller positions requiring the highest skills and mental workload are those handling descent traffic. These positions are responsible for producing an orderly flow of traffic into the TRACON. The Descent Advisor (DA) is intended to provide controllers in these positions with flexible tools to implement the traffic plan generated by the TMA.

For all aircraft entering an arrival sector, the DA implemented at that sector computes estimated times of arrival (ETAs) at its respective arrival gate. These ETA computations take into account the airspace structure and ATC procedures of each arrival sector. For simplicity, only two DAs are shown in figure 1, but in general there can be four or more, at least one for each arrival gate feeding traffic into the TRACON. The ETAs from all arrival sectors are sent as input to the TMA which uses them to calculate efficient, conflict free-landing schedules. These scheduled times of arrivals (STAs) at the runway

### TRAFFIC MANAGEMENT ADVISOR (TMA)

This section describes elements of the TMA, with emphasis on the design of the scheduler and graphical interface.

#### Overview

The Traffic Management Advisor (TMA) comprises algorithms, a graphical interface, and interactive tools for use by the Center traffic manager or TRACON controllers in managing the flow of traffic within the terminal area. The primary algorithm incorporated in it is a real-time scheduler which generates efficient landing sequences and landing times for arrivals within about 200 n.m. from touchdown. Its graphical interface and interactive tools are designed to assist the traffic manager in monitoring the automatically generated landing schedules, to override the automatic scheduler with manual inputs and to change scheduling parameters in real time. It has been implemented on a separate workstation that is interfaced with the workstations running the DAs at the various arrival areas.

In essence, the scheduler is a real-time algorithm that transforms sequences of arrivals into reordered sequences of scheduled times of arrival (STAs) using one of several scheduling protocols selected by the traffic manager. Operation in real time implies that the algorithm generates the STAs in a small fraction of the time it takes each aircraft to fly from its initial position to touchdown. This condition places important computational constraints on the algorithm.

Since the scheduler is the main computational unit of the TMA, its functions and operations are described first. Next is a description of the graphical interface and the set of tools the traffic manager uses to monitor and to interact with the scheduler in real time.

### Design Issues

The scheduling and freeze horizons are key parameters in determining the performance of the scheduler. The scheduling horizon is a time interval specifying when an aircraft becomes eligible for scheduling. The freeze horizon is a time interval specifying when an aircraft's STA will no longer be changed without manual intervention. The TMA gets periodic updates of the estimated times of arrival (ETAs) for all aircraft. These updates come from the DA at the sector currently controlling the aircraft. The TMA then uses the ETAs to determine when aircraft are eligible for scheduling or having their scheduled times frozen.

Although the scheduling and freeze horizons are time dependent quantities, they can also be approximated in the spatial domain by concentric circles with the arrival airport at the center. The circles representing the horizons are superimposed in figure 2 on the arrival airspace structure of the Denver Center. The location of the freeze horizon (in space and time) must balance two conflicting objectives. On the one hand, it must be chosen sufficiently early in the approach in order to give the arrival controllers adequate time and airspace to meet the scheduler-generated STAs. At the very latest, the freeze horizon must be chosen before arrivals reach the area where descent clearances are issued. This area is about 30 min to touchdown. On the other hand, the location of the freeze horizon should not be chosen so early that arrivals from nearby airports often appear later than the freeze horizon, thus missing the scheduling window altogether. Also, an early freeze horizon increases the probability of schedule-disturbing events occurring between the freeze horizon and the TRACON boundary, such as weather disturbances. The rescheduling of frozen aircraft necessitated by such occurrences causes an undesirable increase in controller workload, and thus should be minimized. These considerations, as well as the results of simulation tests and experience with the current metering system at the Denver Center suggest a freeze horizon 35 minutes to touchdown, with a scheduling window 10 minutes long. Aircraft whose ETAs are in the scheduling window are

operated upon by the scheduling algorithm to generate the optimum sequence.

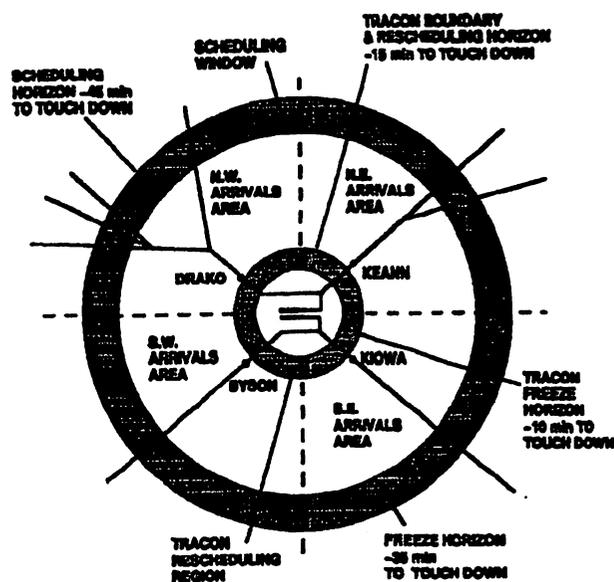


Figure 2.-Scheduling Regions.

When it becomes necessary to reschedule certain aircraft, the last and best opportunity to do so occurs in the region where the arrivals transition from the Center into the TRACON at the feeder gates. This region is identified in figure 2 as the TRACON rescheduling region. In this transition region, the scheduling process described above can be repeated. Since the scheduling window is narrow and close to touchdown, rescheduling in the TRACON region consists primarily of fine tuning the Center-determined arrival sequence. Extensive changes in the schedule for arrivals this close to the runway are neither necessary nor feasible. Frequent reordering of arrival sequences or large changes in STAs at this point would disrupt the orderliness of the arrival flow and produce complex trajectories in the TRACON airspace, thereby increasing controller workload. The primary reason for rescheduling aircraft in the TRACON airspace arises from the need to handle missed approaches, emergency aircraft, and changes in runway.

One of the most critical aspects of designing a scheduler is defining the procedures for establishing the arrival sequence. In today's ATC system controllers generally attempt to maintain a first-come-first-served (FCFS) sequence when vectoring aircraft for arrival in a terminal area. Generally first-come refers to projected arrival time at the runway but it is possible to define first-come in other ways. For instance, first-come could mean the order in which the aircraft cross a point in space such as the Center boundary. Time based orderings can be static or dynamic. The initial estimate of arrival time for each aircraft establishes a certain arrival order but subsequent ETA updates may change the arrival order.

Controllers frequently deviate from a strict FCFS order in order to accomplish specific objectives. When decisions need to be made in sequencing of a group of aircraft factors such as wind direction, aircraft type, and the route topology are all taken into account by the controller in deciding on a sequence. This takes considerable skill and judgment, and different controllers may handle similar conditions quite differently. The approach taken in the TMA to address these concerns is to provide various choices to the controller which are selectable in real time.

### Design of Scheduler

Optimization of aircraft arrival schedules has been the subject of numerous studies in recent years (refs. 3-8). However, in the studies cited, optimization benefits have been difficult to quantify because they are sensitive to many factors that are difficult to measure or estimate. Such factors include the choice of representative arrival sequences, the distribution of aircraft weight classes and the selection of base line conditions against which schedule optimization benefits can be accurately gauged. In the most recent study of this problem (ref. 8) scheduling efficiency is computed by a Monte Carlo simulation for the three types of scheduling methods implemented in the TMA. Results of this study will be summarized after describing the real time scheduler.

A theory for the design of real-time schedulers capable of handling the diverse conditions arising in ATC has not been treated comprehensively in the research literature. In the U.S.A., the best known implementation of a real-time scheduler is the En Route Metering (ERM) system, which has been in operation at various Centers, including the Denver Center, for a number of years. ERM has evolved, with fair success, as a tool for controlling the flow of traffic into the TRACON under capacity limited conditions. However, it is not designed to produce conflict-free, optimum arrival schedules at the runway for a mix of aircraft weight classes, as is the objective in this design. In West Germany, the COMPAS system (ref. 3) undergoing tests at the Frankfurt Airport also incorporates a real-time scheduler. The design described herein expands on features in ERM and COMPAS and also incorporates new graphical and interactive concepts that capitalize on the capabilities of high performance workstations.

The TMA can be configured in a variety of ways. The user can select between time based and distance based sequencing and also whether the sequence is static or is updated dynamically. The user can also enable or disable time advance and optimization. All of the options can be changed at any time and will take immediate effect on the current schedule.

The first step in the scheduling process is to arrange the aircraft in an ordered list based on the currently selected sequencing method. For example, if the sequence method is set to dynamic time the aircraft with the earliest ETA

is placed at the beginning of the list and the aircraft with the latest ETA is placed at the end of the list.

Once the type of sequence to use has been established the scheduler next checks the interaircraft time spacings on final approach. For those with less than the minimum allowed, the scheduler adds just enough time to meet the minimum distance separation standards required by FAA regulations. It should be noted that the minimum separation distances depend on the aircraft weight classes (heavy, large, and light) of the leading and trailing aircraft. Since the scheduler works on the basis of time and not distance, it is first necessary to transform the distances into equivalent time separations using procedures described in reference 4. The results of these transformations are a set of time intervals which specify the minimum time spacings on final approach for all nine possible landing sequences of aircraft with three weight classes. A complicating factor is that the transformations depend implicitly on the ground speed of aircraft on final approach. Since ground speed depends on both final approach air speed and wind speed, it becomes necessary to update the time spacings in real time. This cumbersome and complicated procedure should be eliminated by developing new criteria specifically for time-based minimum separation standards.

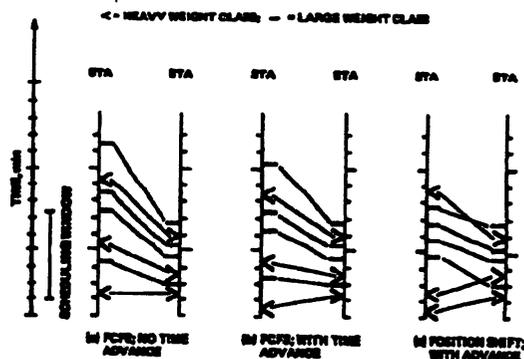


Figure 3.-Effect of scheduling methods on delays.

The operation of the basic scheduling algorithm without time advance or optimization can be illustrated graphically with the help of time lines drawn side by side as in figure 3(a). The time line on the right shows the ETAs of several large and heavy aircraft within a scheduling window. The earliest ETA is at the bottom of the list and increasing future time is toward the top. For illustrative purposes, only two minimum separation times are used, 2 min. for a heavy followed by a large aircraft and 1 min. for all other sequences. Since the time separations between ETAs in this list are generally smaller than the minimums, the scheduler has to delay aircraft to conform with the minimums. The result of this operation is shown on the STA time line. Here, horizontal lines connect STAs and ETAs of the same aircraft. The original ETA order has been maintained as

indicated by the fact that none of the connecting lines cross each other.

The effect of adding the time advance option to the basic scheduler is illustrated in figure 3(b). A 1-min time advance relative to the ETA was allowed for each aircraft. The effect for many aircraft is a reduction of delay and fuel consumption. On the other hand, those aircraft whose time is advanced may experience increased fuel consumption because of higher-than-optimal cruise and descent speeds. Therefore, in assessing the overall benefit of time advance, it is necessary to balance time and fuel savings for those aircraft whose delays are reduced against an increase in fuel consumption for those whose time is advanced. Nevertheless, in most situations, time advance is likely to be advantageous.

In consideration of these trade-offs, the scheduler attempts to be intelligent in applying time advance by not advancing aircraft when the benefits to be gained are minimal. The scheduler does this in two ways. First, the amount of advance is controlled by specifying both a maximum advance and a fraction of the total advance available that is to be applied. The scheduler first determines the minimum time to landing for a given aircraft as previously defined. Only a fraction of the total advance available is used by the scheduler. This amount is compared with the maximum allowable advance and the smaller of the two quantities is used to arrive at the aircraft's scheduled time. The second technique used by the scheduler is to advance a given aircraft only when it is part of a closely spaced group of aircraft. A closely spaced group is defined as a set of consecutive aircraft which are spaced at or below the minimum allowable separation. The number of consecutive aircraft which defines a group is an adjustable parameter, typically set to four.

A position shifting scheduler with or without time advance removes the constraint of preserving the ETA order when generating the STA list. Position shifting for aircraft scheduling was studied by Dear (ref. 5) and subsequently by others (ref. 6). The scheduler implemented here optimizes the STA list with respect to a user specified maximum number of position shifts. This means that the landing order of an aircraft may not be moved more than the specified number of aircraft ahead of or behind the FCFS order. The schedule produced by the position shift scheduler with time advance is illustrated in figure 3(c) for a single position shift. It can be seen that position shifting has provided additional delay reduction. However, these reductions are highly dependent on the mix of aircraft in the list. There would be no advantage in position shifting if the minimum time separation between all aircraft were the same. Position shifting tends to bunch aircraft of the same weight class as in figure 3(c). Although position shifting can reduce delays, it is not always feasible to implement. For example, position shifting of two in-trail aircraft generally requires one aircraft to overtake the other. This procedure increases controller workload, making position

shifting undesirable. Therefore, the scheduler has been designed to allow position shifting only if it can be completed before the position-shifted pair has merged on a common route.

### Estimated Performance of Scheduler

A study has been made to determine the effect of the various scheduling methods used in the TMA on aircraft delay. A full discussion is available in reference 8. In this study a traffic model was developed to simulate peak arrival traffic at the Denver Center. Traffic samples with a varying mix of heavy and large aircraft types were created representing an hour and a half of data. Figure 4 illustrates the effect of using time advance and single position shift optimization.

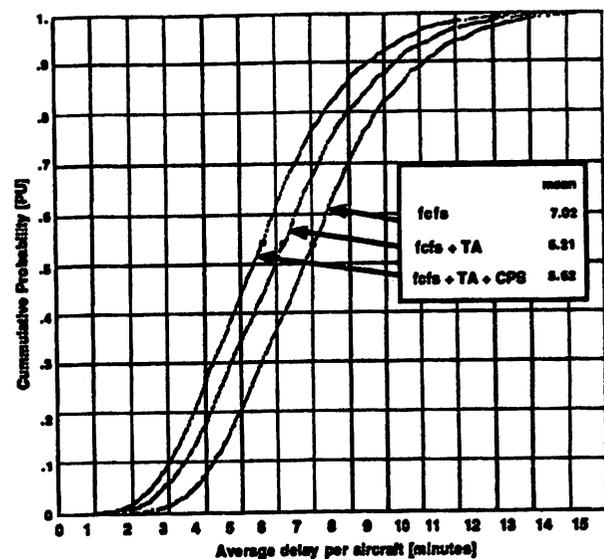


Figure 4.- Comparison of Scheduling Methods.

The curves shown are cumulative probability distributions for the average delay per aircraft. The distributions are based on 2500 traffic samples each, with a traffic density of 40 aircraft per hour. The model assumes a rectangular probability distribution for aircraft arrival times at the Center boundary. The resulting distributions shown in the figure are for a 50% heavy, 50% large traffic mix. It can be seen that both time advance and position shift provide approximately equal, incremental reductions in average delay per aircraft, compared to the FCFS scheduler. Note that because the distributions are not symmetrical the mean delay per aircraft does not fall at the 50% point. These curves show the best possible mix of traffic in terms of reducing delays. With a more typical traffic mix of aircraft types the reduction in delay per aircraft would be smaller. The average delay distributions are very sensitive to the arrival distribution and winds. Assuming a triangular instead of a rectangular distribution nearly doubles the average delay.



**Blocked time intervals and slots-** The controller can block out times in which he does not want aircraft to be scheduled as previously defined. Two kinds of blocked times are displayed in the figure; intervals and slots. The scheduler will not place any aircraft in the area delimited by the blocked times. The figure shows an interval which caused delays for CO409. Notice that the scheduler has placed aircraft right at the limits of the blocked interval. Blocked slots are slightly more complicated. They are created by using a menu option. The figure shows a heavy slot just past the 55-min mark. Unlike the procedure for intervals, the amount of airspace reserved by the slot depends on the weight classes of other aircraft being scheduled, just as though a slot were an actual aircraft. Notice in the figure that UA134 has been moved behind the heavy slot even though it was ahead of it on the time line.

**Time line tag pop-up menu-** Various other scheduling options are available on a menu brought up by depressing the right mouse button while over an aircraft time line tag. These include selectively rescheduling aircraft after they have passed the freeze horizon, rebroadcasting the current scheduled time to a PVD, and returning a manually scheduled aircraft to automatic scheduling status.

#### CONCLUDING REMARKS

The Traffic Management Advisor described in this paper deliberately places the automation tool in a subordinated position relative to that of the human controller, who will remain the cornerstone of the air traffic control process in the foreseeable future. The controller selects the automation levels and functions in response to specific traffic management problems. He or she can combine his/her own procedures and decisions with computer generated advisories by choosing tools that complement his own control techniques. At one end of the spectrum of computer assistance, the controller can use the tools in a passive mode to gain insight into the effect of the planned actions. At the other end of the spectrum, he or she can use the tools actively by issuing the computer generated clearances to the aircraft.

The interactive graphic interfaces adopted in the design are probably the most innovative as well as the least proven design feature. They build upon the user environment incorporated in modern high-performance engineering workstations. That this workstation technology can be so readily adapted to air traffic control automation is remarkable and fortunate for progress in this area.

Controller acceptance of these interfaces, more than any other issue, will determine the viability of this concept. Here, real time simulations are the main avenue for evaluating controller response, for refining the interface, and for developing baseline controller procedures. Ultimately, however, only tests with live traffic can establish their effectiveness with a high level of

confidence. Such tests, which are considered an essential step in the development of an advanced automation system, can begin as soon as access to aircraft tracking data is obtained at an en route center.

#### REFERENCES

1. Erzberger, H.; Nedell, W.: Design of Automated System for Management of Arrival Traffic, NASA TM-102201, June 1989.
2. Davis, T.; Erzberger, H.; Bergeron, H.: Design of a Final Approach Spacing Tool for TRACON Air Traffic Control, NASA TM-102229, Sep. 1989.
3. Volckers, U.: Computer Assisted Arrival Sequencing and Scheduling with the COMPAS System. Conference Proceedings, Efficient Conduct of Individual Flights and Air Traffic, AGARD No. 410.
4. Tobias, L.: Time Scheduling of a Mix of 4D Equipped and Unequipped Aircraft, Proc. of 22nd IEEE Conf. on Decisions and Control, pp. 483-488, Dec. 1983.
5. Dear, R.: Increasing Capacity with Computer-Assisted Decision Making., Proc. International Air Transportation Conf., Vol. 2, May 1979.
6. Luenburger, Robert: A Traveling-Salesman-Based Approach to Aircraft Scheduling in the Terminal Area. NASA TM-100062, March 1988.
7. Magill, S.A.N.; Tyler, A.C.F.; Wilkenson, E.T.; Finch, W.: Computer-Based Tools for Assisting Air Traffic Controller with Arrivals Flow Management; Royal Signals and Radar Establishment, Malvern, Worcestershire, U.K., Report No. 88001, Sept. 1988.
8. Neuman, F.; Erzberger, H.: Analysis of Sequencing and Scheduling Methods for Arrival Traffic. NASA TM- 102795, April 1990.