Design of Automated System for Management of Arrival Traffic

Heinz Erzberger, Ames Research Center, Moffett Field, California
William Nedell, San Jose State University, San Jose, California

June 1989

NASA

TABLE OF CONTENTS

ABBREVIATIONS

AAS	Advanced automation system
AI	Artificial intelligence
ARTCC	Air route traffic-control center, also referred to as the center
ATC	Air traffic control
C	Cruise speed mode
CLR	Cleared aircraft
D	Descent speed mode
DA	Descent advisor
DME	Distance measuring equipment
ERM	En route metering
ETA	Estimated time of arrival
FAA	Federal Aviation Administration
FAST	Final approach spacing tool
FCFS	First-come-first-served
HUD	Heads-up display
I	Inquiry mode
IAS	Indicated airspeed
ID	Identification
KIAS	Knots, indicated airspeed
n.mi.	Nautical miles
PVD	Plan-view display
RI	Route intercept mode of horizontal guidance

SP	Standard airline procedure descent profile
STA	Scheduled time of arrival
TMA	Traffic management advisor
TOD	Top of descent
TRACON	Terminal radar control (facility)
VORTAC	Type of navigation station providing range and bearing
WC	Waypoint capture mode of horizontal guidance

DESIGN OF AUTOMATED SYSTEM FOR MANAGEMENT OF ARRIVAL TRAFFIC

Heinz Erzberger and William Nedell[1]
Ames Research Center

SUMMARY

This paper describes the design of an automated air traffic control system based on a hierarchy of advisory tools for controllers. Compatibility of the tools with the human controller, a key objective of the design, is achieved by a judicious selection of tasks to be automated and careful attention to the design of the controller system interface. The design comprises three interconnected subsystems referred to as the Traffic Management Advisor, the Descent Advisor, and the Final Approach Spacing Tool. Each of these subsystems provides a collection of tools for specific controller positions and tasks. This paper focuses on the design of two of these tools, the Descent Advisor, which provides automation tools for managing descent traffic, and the Traffic Management Advisor, which generates optimum landing schedules. The algorithms, automation modes, and graphical interfaces incorporated in the design are described.

Information generated by the Descent Advisor tools is integrated into a plan view traffic display consisting of a high-resolution color monitor. Estimated arrival times of aircraft are presented graphically on a time line, which is also used interactively in combination with a mouse input device to select and schedule arrival times. Other graphical markers indicate the location of the fuel-optimum top-of-descent point and the separation distances between aircraft at designated points. Computer generated advisories provide speed and descent clearances which the controller can issue to aircraft to help them arrive at the scheduled times or with specified separation distances. Two types of horizontal guidance modes, selectable by the controller, provide advisories for managing the horizontal flightpaths of aircraft under various conditions.

The Traffic Management Advisor (TMA) comprises algorithms, a graphical interface and interactive tools for use by the Center traffic manager in controlling the flow of 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. Four scheduling algorithms have been implemented in the scheduler; they are run first-come-first-served with or without time advance and position shifting with or without time advance. 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. The Descent Advisor and the Traffic Management Advisor have been implemented on a network of workstations, which distributes the computational load, yet permits efficient exchange of information between the two types of automation tools.

---

[1]San Jose State University, San Jose, CA.

This report supersedes a recently published report which described the design of the Descent Advisor. In addition to updating the earlier report, this report contains a new chapter covering the design of the Traffic Management Advisor.

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 system described in this paper builds upon ATC automation concepts and algorithms described in reference 1. Its main purpose is to provide a variety of computer-aided tools that can assist controllers in achieving safe, orderly, and expeditious movement of traffic within the terminal area. The criteria for designing these tools revolve around the principle of human-centered automation. In the context of ATC this principle requires developing tools that complement the skills of controllers without restricting their freedom to manage traffic manually. Such capability is achieved in the design described herein by providing various modes to assist the controller in solving specific ATC problems and by letting the controller decide when and how to use these tools.

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 a discussion of design guidelines and an overview of the automation concept. This is followed by a detailed description of the Descent Advisor and the Traffic Management Advisor, which are the two key elements of the concept. Although the design of the Descent Advisor was recently published in reference 2, it is also included in this report both for the sake of completeness and to describe newly implemented enhancements and modifications.

Preparations are in progress to implement elements of this system at an en route center for evaluation on a noninterfering basis. In view of past experience, such operational testing is an essential step in

validating automation concepts for ATC. Then, the FAA can confidently decide what elements of this system warrant implementation in the AAS.

DESIGN GUIDELINES

Automation tools for ATC are defined here as systems with which controllers conduct informative as well as interesting dialogs that contribute to increased efficiency in performing their tasks. A successful dialogue consists of a rapid exchange of information using color graphics and possibly synthesized speech for system-to-human communications and a mouse, keyboard, touch-sensitive screen, and possibly a voice recognition system for human-to-system communications. This paper focuses on interactive color graphics and mouse input as the primary vehicle for system-human dialogue. As shall be explained, the graphics developed here are used to convey more complex information than is provided by aircraft tracks and alphanumeric lists displayed on current-generation ATC monitors.

In order to ensure desirable characteristics in the human-system interface and gain controller acceptance of automation tools, a list of guidelines has been established, which is presented in table 1.

These guidelines are based in part on experience in other aerospace fields with successful examples of automated system designs. The system design described in the following sections represents an attempt to be faithful to these guidelines.

AUTOMATION SYSTEM CONCEPT

An automated system for ATC may be divided into three principal subsystems whose function involves sensing, planning, and controlling. The sensing subsystem includes all of the components, including the ground radars, mode C transponders, and ground computers, that contribute to generating aircraft position tracks on ATC monitors. Since this subsystem is already highly refined in today's system, this paper concentrates on the design of the planning and controlling element of the automation system.

Figure 1 gives a diagrammatic representation of the proposed concept. 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.mi. 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

3

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.

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 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.

DESCENT ADVISOR

This section begins with a brief review of descent procedures and continues with a detailed description of the DA.

Center Controller Procedures for Managing Descent Traffic

An analysis of controller procedures for managing descent traffic provides important insight and motivation for designing the descent advisor system. Hence these procedures are briefly reviewed in preparation for describing the system in the next section.

Typically, arrivals enter the airspace of the ARTCC that is feeding traffic to a major destination airport at least 200 n.mi. from the airport. Initially, the arrival traffic continues along established jet routes at cruise speed and altitude. Center controllers direct the traffic to points in space called feeder gates, typically located about 30 n.mi. from the airport at 10,000 ft above ground level. Some airports utilize as many as four such gates or corner posts which approximately form a rectangle, with the airport at the center. At the gates, the Center controllers hand the traffic over to the Terminal Radar Control Facility (TRACON) for final sequencing to the runway.

Center controllers have the critical task of ensuring that the streams of randomly spaced aircraft are merged into a properly ordered and spaced single stream as the traffic arrives at the gates. They accomplish this by issuing a series of speed and altitude clearances and heading vectors.

Traffic over the gates is spaced either on the basis of distance or time. The choice of spacing distance depends on the level of traffic and other factors, but is typically 10 n.mi. Under the time spacing criterion, arrivals are separated either by a specified time interval such as 3 min or alternatively are delivered at a specified clock time. Clock time over the gate as a spacing strategy is used primarily when En Route Metering (ERM) is in effect. ERM is a program installed on Center host computers. Certain Centers such as the Denver Center rely on ERM during periods of heavy traffic to allocate delays and coordinate flow at the feeder gates.

Two types of strategies are commonly used to control spacing. In one type, controllers keep aircraft on their standard arrival routes as long as possible and control spacing with speed clearances before and during the descent. This strategy is appropriate for low and medium traffic conditions. In the second type, controllers take aircraft off their standard arrival routes while still at cruise altitude and then issue a series of heading vectors to control the spacing. This is the strategy of choice during heavy traffic.

Although accurately controlling the spacing at the gates is their primary objective, controllers also try to keep aircraft on fuel-efficient profile descents whenever possible by issuing appropriate descent clearances. Another factor that complicates the spacing control task is the necessity to compensate for strong variations in ground speed during the descent as a result of altitude-dependent winds and atmospheric effects. The management of descent traffic under these complex conditions is a difficult task that only experienced controllers with exceptional skill can perform efficiently. Thus, these tasks offer a timely opportunity for the development of automation tools.

Graphic Interface Design for Descent Advisor Automation Tools

In essence, management of arrival traffic in the Center airspace involves predicting and controlling the spatial and time relationships of aircraft at the feeder gates. With currently used manual procedures, controllers are able to visualize evolving traffic situations 5-10 min into the future, depending on skill

and circumstances. The DA increases this prediction time limit to at least 25 min in most situations, thus giving the controller more time to organize traffic flow efficiently.

The analytical foundation for the DA is a collection of numerical algorithms for accurately predicting and controlling aircraft trajectories in space and time, referred to as four-dimensional guidance. Although four-dimensional guidance was originally developed for on-board flight management systems, it also provides a powerful framework for solving analogous problems in ATC automation. The design of a four-dimensional descent advisor algorithm was reported in a series of papers (refs. 1 and 3). It was shown in these papers that with descent clearances generated by this algorithm, pilots could control their arrival time at a feeder gate to an accuracy of ±20 sec. Such accuracy is significantly higher than is achieved in current operations and provides the basis for increasing the efficiency of the ATC system. One major challenge of applying this algorithm in an ATC environment lies in the design of an efficient interface between the controller and the algorithm. An equally important challenge is designing sufficient flexibility into the DA so that controllers can adapt it effortlessly to handle the types of traffic management problems previously described.

The design of the interface builds upon the environment available in a modern engineering workstation. Thus, the interface uses a mouse or track ball as the primary device to enter information and special color graphics to output information. However, the mouse is the preferred device in this application because it is less cumbersome to use than a trackball for entering certain types of information.

Integrated controller display– The controller interface combines on a single high resolution color monitor both the traditional plan-view display of aircraft tracks as well as advisories generated by the automation tools. Furthermore, the monitor screen is used to display the complete menu of automation tools available for selection by the controller. The main problem in implementing this concept, referred to as an integrated controller display, is organizing the information on the screen so as to minimize confusion and avoid display clutter. An alternative to an integrated controller display is a separate monitor for displaying advisory information. However, this approach has the disadvantage of requiring the controller to shift his attention and viewing direction from one monitor to the other, possibly causing distraction and loss of concentration during crucial moments.

In a sense, an integrated controller display is somewhat analogous to a heads-up-display (HUD) installed in modern aircraft cockpits. A HUD superimposes flight director information on the out-the-window visual scene, a technique that is favored by most pilots.

The integrated-controller-display concept is illustrated in a series of photographs of the monitor screen (figs. 2-4, 6, and 8). These photographs provide the basis for describing the graphic interfaces in this paper. They illustrate the detailed implementation of this concept for an arrival area at the Denver Center. Specifically, the pictures show a plan-view map of the airspace through which arrival traffic flows toward the Drako feeder gate, one of four such gates feeding traffic to Denver Stapleton International Airport. The Drako point can be found near the right lower corner of the display. The names of several waypoints referred to in later discussions are also shown. Blue lines indicate standard arrival routes leading through the gate into the TRACON airspace. Markers depicting air routes, sector boundaries and other map features are shown in white. The white broken circular arcs are 25-n.mi. range circles centered at Drako. Except for the color coding, this type of information is identical to that found on the existing plan view display (PVD) at the Denver Center. A complete set of display management tools,

including zooming and panning, has also been implemented. It allows the display to be reconfigured for any other sector.

Aircraft positions, identification (ID) tags and associated data blocks for Denver arrivals are drawn in green. The data blocks are organized in standard format, with the second line showing mode c altitude in hundreds of feet, and with the third line showing the computer identification code and the ground speed in knots. Aircraft data blocks drawn in white identify overflights. Other display features shown for some aircraft in the figures are 5-n.m. range circles which controllers use to check for separation and 3-min trend vectors. The end of the trend vector predicts the position of the aircraft in 3 min by extrapolating the current position using the current estimate of ground speed and course. These features can be turned on and off for each aircraft individually using a combination of mouse pointer and pop-up menu.

The next sections describe in detail the graphical artifacts and automation modes comprising the DA tools.

Time line– The graduated white scale on the left side of the screen is the so-called time line, whose purpose is to show graphically the future time relationships of aircraft at a designated time control waypoint. The scale covers a time range of about 30 min, with future time increasing toward the top in 1-min increments. The corresponding hour on the scale can be inferred from the current time shown just below the time line. During operation the controller observes the time-line scale sliding steadily toward the bottom of the display. At the point where the downward sliding scale terminates just above the current time, future time becomes current time.

The time line concept is a natural by-product of four-dimensional guidance theory, which underlies the accurate estimation of arrival time. It was used in a slightly different format as part of an artificial intelligence (AI) based scheduling system for ATC flow management (ref. 4). It also plays a key role in the ATC planning system, COMPASS, developed in Germany (ref. 5).

In this report the feeder gate/time control point has been changed from the traditional one at Drako for this area to Jasin (shown in blue in the figures, except fig. 4). With the increased precision of control provided by the DA it is feasible as well as more optimal to push the effective gate location closer to touchdown. This change increases both fuel efficiency and arrival time accuracy by allowing arrivals to continue their fuel efficient profile descents without interruption to 11,000 ft (approx. 6000 ft above the runway) instead of terminating them at 17,500 ft at Drako. The operational implications of this change will be determined in the future.

When aircraft destined for Denver first enter the airspace of the sector that is feeding traffic to the Drako arrival area, the DA algorithm computes an ETA to the time control waypoint, Jasin. In this computation the algorithm assumes the aircraft will continue to fly at cruise speed and altitude along the planned arrival route to the point of descent and then follow a standard profile descent at idle to the altitude specified at Jasin. This type of descent trajectory will also be referred to as a standard procedure profile (SP). The ETAs with their corresponding ID tags are then posted in green on the right-hand side of the time line. The arrival times are automatically updated approximately every 30 sec by the DA algorithm using the current position, velocity, altitude, and route information.

After learning to interpret the information on the time line, the controller can use it to quickly assess the traffic situation at the designated time control point at various future times. Furthermore, he or she can gauge the effect of speed clearances and heading vectors by monitoring the ETAs posted on the time line.

Descent clearance advisors– In addition to providing ETAs on the time line, the DA algorithms can also provide more extensive assistance in solving the traffic management problems previously reviewed. To invoke the DA algorithms, the controller must first choose one or more aircraft of interest and then designate his choice to the computer by an action referred to as "selecting" an aircraft. An aircraft is selected by moving the mouse pointer to the position marker of a chosen aircraft, a small green diamond, and depressing (clicking) the left mouse button. If the pointer is within the acceptance aperture—a small region containing the diamond when the button is clicked—the color of the aircraft and all information associated with it on the screen turns from green to yellow. Similarly, a selected aircraft is deselected by clicking on its position marker, thereby changing its color back to green. For operational and computational reasons, the number of aircraft in the selected state should be kept to fewer than about five.

Selecting an aircraft causes two panels (also known as windows) to pop up on the screen (fig. 3). The Mode Select Panel near the bottom of the screen provides a menu of labeled "buttons" for selecting advisory modes and display options. The functions of these modes will be explained in the next section. At the top of the screen, the DA Panel displays profile descent information generated by the DA algorithm.

Each line in the panel provides various profile parameters and route data for one selected aircraft. These parameters are computed specifically for the type and weight of the arriving aircraft as well as the current estimate of the wind field. They are updated whenever the profile is refreshed, about every 10 sec. As on the time line, the ETAs of the selected aircraft determine the ordering of the lines in the panel, with the last-to-arrive aircraft occupying the top line. The parameters that define the profile descent trajectory of a selected aircraft consist of the desired cruise speed, the DME range of the top of descent point (TOD) from a nearby station (Denver VORTAC in this case) and the descent speed profile consisting of constant Mach and constant indicated airspeed (IAS) segments. Table 2 gives three example lines in the DA panel as well as the definition of each item in the line. A full explanation of some of the items will be found in later sections.

For convenience, those parameters in the DA Panel that comprise the profile clearance, which the controller issues to the aircraft, can also be displayed on an extra line in the aircraft data block. Since the extra line may cause excessive crowding of the display under heavy traffic conditions, it can be turned off by a mouse actuated on-screen switch. Some of the data blocks for selected aircraft in figures 3 and 4 have this profile clearance line switched on.

A significant attribute of the system design is its rapid refresh feature. Every radar sweep or approximately every 10 sec the DA algorithm automatically recomputes the descent profiles and the ETAs of all selected aircraft. The profiles of the other aircraft continue to be updated at a slower rate. Thus, the information on the time line and in the DA panel adapts continuously (at a 10-sec rate for selected aircraft) to changes in aircraft position, altitude, heading, and airspeed. In addition to increasing the accuracy and timeliness of the information, this adaptive capability can be exploited by a controller skilled in its use to solve various traffic management problems.

8

Time range bar– When an aircraft is selected, a yellow vertical bar with brackets at the ends is displayed next to the time line at the ETA of the aircraft. The time range enclosed by the brackets indicates the maximum variation of arrival time achievable through speed profile management along a fixed arrival route. In particular, the ends of the brackets point at the maximum and minimum arrival times obtained by flying the aircraft along its minimum and maximum speed envelopes, respectively. It therefore follows that the ETAs for the SP profile must fall between or on the end points of the brackets. This is seen to be the case in figures 3 and 4. Overlapping brackets, such as those in figure 3, indicate that the time ranges of some aircraft overlap. To avoid ambiguities in this case, the overlapping brackets are slightly offset horizontally.

Top of descent marker– From aircraft performance studies it is well known that the location of the TOD point for a fuel efficient descent trajectory depends significantly on the choice of the descent speed profile as well as on many other factors. Therefore, showing its location on the plan view display (PVD) can give the controller useful cues for the efficient management of descent traffic. Its location on the PVD is indicated with a small purple marker which is labeled with the appropriate aircraft ID (see fig. 3).

The location of the TOD marker should normally be at least several miles in front of the aircraft position at the time the descent clearance is issued. However, if the TOD and position markers appear to coincide, it indicates that the aircraft is either at or has already passed the optimum point of descent. In this situation the DA algorithm uses a mid-descent logic for generating speed profiles, as explained in a later section. As long as this logic can still find a feasible descent profile during the profile refresh cycles, the TOD marker will be seen to track the position marker. The colocation of these markers alerts the controller that the opportunity to initiate a descent along the currently predicted horizontal path is about to disappear, unless the aircraft is already flying a descent profile. If this loss of descent opportunity occurs, it is indicated by the simultaneous disappearance of the TOD marker and the posting of the message "Path Stretching Required" in the DA Clearance Panel for the affected aircraft. The controller can remedy this situation by vectoring the aircraft to stretch its horizontal path until the lost information reappears. The primary use of this marker is as a visual indicator to help the controller in judging the distance remaining to the TOD point, and in properly timing the issuance of the descent clearance. Also, the markers can assist the controller in monitoring the descents in mountainous terrain and in detecting potential conflicts with descending traffic or with overflights.

Distance spacing markers– Distance spacing markers give a graphical representation of the future spatial relationship of traffic converging on a designated waypoint, called the spacing checkpoint. They provide a graphical method for visualizing potential conflicts and for checking the spacing between aircraft at any waypoint on the arrival routes. The location of the markers on the screen are based on the DA-generated trajectory parameters currently displayed in the DA Clearance Panel, and therefore are refreshed at the same update rate. The time control point or any waypoint upstream of it can be designated as the spacing check point by choosing a mouse-actuated pop-up menu item. The name and position marker of the spacing check point are drawn in dark blue. (Drako in fig. 3).

The markers are displayed on the PVD by clicking on the spacing button in the Mode Select Panel. They consist of small blue dots which have leader lines to an aircraft identification tag and a number representing the distance in nautical miles to the spacing check point (Drako) (fig. 3). The aircraft ID associated with the marker at the spacing check point identifies the first of the selected aircraft, referred

to as the leading aircraft, to arrive at the spacing check point. Markers for all other selected aircraft are placed at locations corresponding to the instant that the leading aircraft arrives at the spacing check point. For convenience, the separation distances from the leading aircraft are also displayed in the DA clearance panel. The markers play an analogous role in the spatial domain as do the ETA markers on the time line. The spacing markers can be used to probe quickly for conflicts at various points on the arrival routes. For example, the controller can check for conflicts at critical route junction points such as Estus, as well as at Drako or at the time control point, Jasin. Since the markers are computed from the DA trajectories, they provide the ability to check for future conflicts along complex descent trajectories.

The display of both ETAs and spacing markers on the PVD provides controllers with complementary tools for visualizing the future relationship of aircraft at a point in space. These graphic tools are more versatile and efficient for controllers to use than are the lists and tables found in the current ATC system.

Arrival time selection and advisories– An on-screen, interactive procedure is used for selecting arrival times and generating arrival time advisories. It is intended to be used primarily as a stand-alone manual mode. A more automated mode used in conjunction with the TMA is described in the next section.

The first step in the procedure is to "select" an aircraft of interest and to watch for the time range bar to appear on the time line. The appearance of the bar indicates that the system is ready to accept a time command, and its length gives the time range available for selection. The aircraft ID tag located to the left of the time range bar serves both as a cursor for selecting a time and as a command input device for the DA algorithm. When an aircraft first enters the arrival area, the location of the cursor ID tag is initialized to the predicted time for the SP profile. Thus, both tags are initially located at the same point on the time line.

The cursor tag can be relocated by moving the mouse pointer on top of the tag, holding down the left mouse button, and dragging it with the mouse to any point on the time line. A capture of the cursor tag by the mouse pointer is indicated by a change in color of the tag from yellow to white when the left mouse button is held down.

The decisive moment for acceptance of a time command occurs when the mouse button is released. At that moment, the cursor position becomes the commanded time that is fed into the DA algorithm. If the cursor position is outside the time range brackets, the time of the nearest bracket is instead chosen as the commanded time.

After the DA algorithm has finished its profile calculation (in about a second) the new descent parameters are read into the DA Clearance Panel. At that time, the profile mode designator is changed from SP to I, where I stands for "Inquiry." Also, the ETA and its associated ID tag are relocated to the new position on the time line. This process was accomplished for TA 321 in figure 3. The "inquiry" designation indicates that the profile shown is the DA's response to the controller's "what if" question. The clearance corresponding to the inquiry profile has not been issued to the aircraft, which continues on its previous profile.

The process of selecting new arrival times and generating the corresponding descent clearances can be repeated any number of times. Each time a new clearance is generated, the TOD and spacing markers

are also updated. Thus, the controller can quickly assess the effect of different time commands on the trajectory and, by making use of the spacing markers, determine whether it creates a conflict situation with other traffic at various waypoints along the route.

As before, the DA algorithm automatically refreshes the profile and corresponding descent clearance approximately every 10 sec. At each refresh cycle, the algorithm reads the location of the cursor tag as well as the current aircraft state. It is also possible to return to the original SP profile by selecting a "Show SP" item in a pop-up menu.

When the controller has found a profile with a conflict-free arrival time, he can lock in the profile and clearance by clicking on the associated ETA ID tag located to the right of the time range bar. This action simultaneously freezes the current profile and clearance, stops the cyclic refreshing of the profile and changes the profile type designation from I to CLR, which stands for cleared. It also places a * symbol next to the ETA tag. This locked profile mode is illustrated by PA001 in figure 3. In order to prevent the locked profile from being changed inadvertently, the cursor tag of the aircraft with the locked profile is rendered inactive. However, the locked profile state can be cancelled if necessary and the aircraft returned to the I-state by clicking again on the starred tag. When an aircraft with a locked profile is deselected, the * remains visible, thereby providing a marker to remind the controller that the aircraft had previously been issued a profile descent clearance. Furthermore, the locked clearance can be redisplayed by reselecting the appropriate aircraft.

Meeting gate times sent by traffic management advisor– In the traffic management concept described earlier, optimum landing times are generated by an automatic scheduler, which resides in the TMA. As explained in a following section, the scheduler assigns landing times that merge the traffic flow from all four feeder gates while ensuring that the aircraft spacing requirements on final approach are observed. Then, the TMA transforms the landing times to gate arrival times and sends them to the DA located at the controller position feeding traffic to that gate. The DA tools and procedures described in the preceding section are applied here to assist the controller in generating advisories for meeting these scheduled times.

The controller accesses the automatic scheduler, referred to as the TMA mode, by a keyboard command. In this mode, the green cursor, which was used in the preceding section to input a time command manually, is removed from the time line. In its place, a blue ID tag appears at the gate time determined by the automatic scheduler. The time line in figure 8 illustrates the TMA mode of operation. If necessary, however, the controller can still assign his own gate times by returning the DA to the manual mode described in the preceding section. To obtain advisory clearances for meeting the gate time, the controller first selects the appropriate aircraft as before, and then clicks on the blue ID tag. That causes the DA algorithm to enter the inquiry profile mode, and begin generating profiles and advisories for the selected aircraft. If the blue ID tag position falls within the time range bar, the speed and descent advisories shown in the DA panel correspond to permissible aircraft trajectories that will cause the aircraft to arrive at the scheduled gate time. As previously explained, conflicts along the route can be checked by using the spacing markers. The controller can lock any currently displayed clearance he finds acceptable by clicking on the ETA tag. He can also redisplay and revise the clearance by applying previously described procedures.

Spacing advisories– Two advisory modes are available for generating descent clearances that meet specified spacing distances at the time control point. The two modes are invoked by clicking on one or the other of the buttons labeled SEP_ONE and SEP_TWO in the Mode Select Panel.

Before using either mode, the controller must choose a desired spacing distance in nautical miles and enter it via the keyboard. The current value of 10 n.mi. shown in figures 3 and 4 is displayed in the Mode Select Panel to the right of the SEP_TWO button. Another requirement for activating these modes is that exactly two aircraft must be in the selected state. As a rule, the controller should pick two consecutively arriving aircraft that are 15 to 35 min of flight time from the feeder gate. The arrival time sequence displayed on the time line and the spacing markers provide information to help the controller choose an appropriate pair.

After selecting a pair of aircraft and observing that the spacing markers predict a spacing sufficiently different from the specified value at the time control waypoint, the controller can choose one or the other of the two spacing modes to assist in achieving the desired spacing.

The SEP_ONE mode should be chosen if only the speed profile of the trailing aircraft can be modified to change the spacing, while the speed profile of the leading aircraft must remain unchanged. Clicking on SEP_ONE causes the DA algorithm to search for a speed profile for the trailing aircraft such that the resulting spacing distance will match the specified value as closely as possible. After the DA algorithm has completed the calculation, the new clearance is posted in the DA Clearance Panel and the spacing markers are updated. Also, the profile mode indicators are changed from SP to I. As for all SP and I profile types, the DA algorithm refreshes the profile and its corresponding clearance in a 10-sec cycle, each time using updated aircraft state information. The clearance for TA 321 and the corresponding spacing distance between TA 321 and AA 404 obtained by invoking SEP_ONE is shown in figure 4. In this example, both the time control point and the spacing check point are located at Drako. The spacing distance is also displayed in the DA Clearance Panel.

If the spacing distance achieved with SEP_ONE is inadequate, the controller can invoke the SEP_TWO mode. This mode gives the DA algorithm the freedom to change the speed profiles of both trailing and leading aircraft. As in the SEP_ONE mode, the DA algorithm first attempts to achieve the specified spacing by changing only the trailing-aircraft speed profile. If this attempt fails, the DA algorithm then changes the speed profile of the leading aircraft in an effort to achieve the desired spacing.

The controller must decide which of the two modes to use based on an assessment of the prevailing traffic situation in his sector. For example, he would probably decide not to use the SEP_TWO mode if he previously had issued a descent clearance to the leading aircraft or if the leading aircraft is itself in trail and at a minimum spacing distance behind another aircraft.

If the controller judges the spacing distances to be acceptable, he would then proceed to issue the appropriate clearance(s) to the aircraft. At that point, he should lock in the profiles and clearances by clicking on the "ok" button in the Mode Select Panel. This action terminates the cyclic refreshing of the profiles and freezes the currently displayed clearance. As was the case for the Time Control Advisories, locking the profiles changes the profile type indicator from I to CLR and places a * symbol after the ETA tag. The profiles can be unlocked by previously described procedures.

The last transaction in using these two modes is to deselect the pair of aircraft. Then the controller can concentrate on identifying and picking another pair of aircraft whose spacing needs to be controlled.

Finally, it should be noted that the SEP1 and SEP2 modes control the spacing only at the time control point. However, the spacing can be checked upstream of the time control point by turning on the spacing markers at various spacing checkpoints.

Horizontal Guidance Modes

In the preceding explanation of descent profiles and automation tools, the crucial problem of constructing efficient horizontal paths for the aircraft was not addressed. In the previously discussed examples, all of the aircraft were flying along standard jet routes. Here the methods used to construct the horizontal profiles are defined for both on-route and off-route situations.

Construction of a horizontal path always begins at the current position and heading of the aircraft and terminates at the time control point. Its calculation is a prerequisite for the subsequent synthesis of the descent profile. The controller has a choice of two horizontal guidance modes, referred to as Route Intercept (RI) and Waypoint Capture (WC). They are selected by clicking on the appropriately labeled button in the Mode Select Panel. The controller also can assign a horizontal guidance mode to an individual aircraft without changing the modes of the others. This is done by displaying a pop-up menu at a specific aircraft and selecting the desired mode with the mouse cursor.

Route intercept (RI) mode– This mode operates in conjunction with the set of standard jet arrival routes converging on the time control point. The routes recognized by the DA algorithm in the Drako feeder gate approach sectors are those drawn in blue on the PVD. They have a corridor width of ±4 n.mi. relative to its center line. Other feeder gates have their own set of standard arrival routes.

In calculating a profile, the DA algorithm first determines if the aircraft position falls within a corridor of one of the standard routes and if its course is within 120° of the route direction at its current location. If the aircraft passes these two tests, the DA algorithm then uses its on-route logic to construct the horizontal path. This is done by concatenating all standard arrival route segments traversed when moving from the current aircraft position to the time control point. It includes circular arc segments to transition smoothly from one segment to another where a change in course occurs at junction points.

If the aircraft fails at least one of these tests, the DA algorithm declares it off-course and invokes the RI logic. This logic, illustrated in figure 5, seeks to create an RI segment connecting the current aircraft position to a point on a standard route segment. First, the algorithm generates a directed line, called the aircraft course vector, which emanates from the current position and points in the direction of the aircraft course over the ground. Then it searches for points of interception of the course vector with standard route segments. As a rule, the first point of interception found establishes where the aircraft joins the standard route. Next, the logic creates a new route by concatenating the RI segment with the segments of standard route traversed between the RI point and the time control point. From here on, the DA algorithm generates speed and altitude profiles in exactly the same manner as if the new route were a standard route. One exception to the rule of choosing the first point of interception arises when that point falls 25 n.mi. or less in front of a junction point and a second point of interception occurs on the downstream side of the junction point. In this case, the RI logic chooses a path to the second interception

point. There are other exceptional cases incorporated in the RI logic, but a description of these is deferred to a future publication. Here it suffices to say that the exact conditions of capture are highly dependent both on the geometry of the route structure and the established controller procedures in an arrival sector.

Route intercept points are indicated on the PVD by white markers with attached aircraft identification tags. In addition, the section of circular flight path that merges the aircraft onto the standard route is also drawn on the PVD. The RI profiles shown for the two selected aircraft in figure 6 illustrate both a regular RI point for CO102 and an exceptional case for TWA61. The first RI point for TW61 on its current heading was found to be closer than the minimum allowed distance (25 n.mi.) from the junction of routes J56 and J170 at Estus. Hence, the second RI point located between the junction point and Drako was the one selected by the algorithm.

The RI logic combined with the automatic refresh of profiles provides considerable flexibility for solving arrival time and spacing problems that require off-route vectoring. This flexibility derives from the fact that changes in heading between refresh cycles reflect in migration of the RI point and therefore in changes of the predicted arrival time. The controller can use this characteristic to expand the arrival time range beyond what is available by speed control alone. For example, in figure 6, the time spacing at Jasin between TWA61 and CO102 can be adjusted by a heading change of CO102. A change to the left will move the RI point upstream on the intercepted segment, thus delaying its arrival at Jasin. At the same time it also causes the downward motion of the time range bar to slow down or even migrate up the time line. One way for a controller to take advantage of this capability is to first make gross changes in arrival time with vectoring and then to make fine adjustments with speed clearances.

Waypoint capture (WC) mode– This mode provides advisories for predicting and controlling the arrival time of aircraft during off-route vectoring. It uses only the aircraft initial position and course and the capture waypoint position in generating the horizontal path. Therefore, it differs from the RI mode in that it does not depend on knowledge of standard arrival routes for its calculations. Any recognized waypoint on the arrival routes including the time control waypoint can be selected for capture. Selection is made by the combination of mouse pointer and pop-up menu.

The horizontal path synthesized in this mode consists of an initial circular arc turn starting at the current position and course followed by a straight flight segment leading directly to the time control point. The path is computed such that the end of the circular arc turn is tangent to the straight flight segment. The geometry of this construction is illustrated in figure 7. The algorithm determines the radius of the turn from the airspeed, wind speed, and maximum allowable bank angle. Furthermore, the direction of the turn toward the time control point is chosen so that the total length of the path is minimized. In order to compensate for computational delays and to allow for controller response time, the algorithm also moves the start of the turn at each computational cycle a distance equivalent to 15 sec of flight time ahead of the current aircraft position. Once this logic has determined the parameters of the path, the DA algorithm synthesizes the speed and altitude profiles in exactly the same way as in the RI mode.

As in other modes, the DA algorithm refreshes the waypoint capture profile in approximately a 10-sec cycle using updated aircraft state information.

14

As illustrated in figure 8, the initial circular turn of the capture profile is drawn on the PVD and is refreshed at the same time that the DA refreshes the trajectory. Thus, the arc tracks the position of the aircraft and maintains its tangency to the current course of the aircraft, even as the aircraft performs complex maneuvers such as holding or path stretching.

A white marker indicates the end of the circular arc and the beginning of the straight line segment on the PVD. A number next to the marker gives the magnetic heading in degrees of the straight line segment leading to the capture waypoint, Estus, whose name and position markers are drawn in yellow (fig. 8). The spacing checkpoint is located at Drako and the time control point at Jasin in this figure.

The most sophisticated application of this mode is in advising the controller when an off-route aircraft should be vectored toward the capture waypoint in order to capture a time slot at the time control waypoint. This includes solving the related problem of when to break out of a holding pattern to capture a time slot.

In such applications, the general procedure is to vector an aircraft initially on a heading 30° to 150° away from a direct course to the capture waypoint. In figure 8, CO102 is heading in the proper direction for using this mode. As the aircraft continues on this course with the DA algorithm generating a sequence of waypoint capture profiles, the motion of the time range bar will slow down relative to that of the time line scale, indicating increasing delays. The greater the angle between the aircraft heading and the direct course to the time control point, the faster the delay will increase. When a specified arrival time on the time scale passes near the middle of the time range bar, the optimum time has been reached for the controller to issue the WC heading advisory shown on the PVD. This places the aircraft on a path to arrive approximately at the specified time. If the residual time error is excessive after the aircraft has executed the heading vector, the controller can zero out that error using the time control advisories described in an earlier section. It is recognized that this procedure (as well as others described above) will be made more automated in future refinements of this tool. However, by initially providing the tool in this form, the controller retains maximum flexibility, and thus can tailor its use more easily to a variety of specific traffic situations.

Speed Profile Modes

Three modes representing different constraints on generating speed profiles for time control have been implemented on the DA algorithm. They can be selected by clicking on a speed-profile-mode button located in the Mode Select Panel. Repeated clicks of the button cause each of the modes to be selected in succession. The names of the modes, in the order in which they are selected, are Descent (DA), Cruise-plus-Descent (C+D), and Cruise (C). The characteristics of each mode are described below.

Descent speed mode– In this mode the DA algorithm achieves a specified arrival time by iterating on the descent speed profile only. Thus, while the aircraft is at cruise altitude, the cruise Mach number/indicated airspeed (IAS) is not controlled by the algorithm. However, at each cyclic refresh of the profile, the cruise speed used as an initial condition by the algorithm is updated with information received from the National Airspace System Host Computer.

15

For each aircraft type, the speed range available for time control is specified by the maximum descent Mach number, the maximum descent IAS, and the minimum descent IAS. For a 727 these limits have been chosen as Mach 0.84, 350 KIAS, and 230 KIAS, respectively.

In addition to providing time control, the DA algorithm also ensures that the descent profile is optimally fuel efficient. It achieves this first by choosing the TOD point so as to minimize powered flight at low altitude and second by generating altitude-speed profiles that permit idle thrust flight throughout the descent. The types of profiles are similar to those generated by advanced flight management systems.

The idle thrust condition produces an interdependence of the speed and altitude profiles. That is, for each speed profile there is a unique altitude profile that can be flown at idle thrust with the aircraft in the clean configuration (no flaps and speed brakes). This interdependency changes the TOD point whenever the speed profile changes. Furthermore, the altitude profile and TOD point also depend on aircraft performance parameters and winds. For example, the location of the TOD point can move 25 n.mi. or more when the speed profile is changed from the fastest to the slowest descent speed.

The best time to invoke this mode is while an aircraft is in cruise and still some distance from the TOD point. However, the algorithm also provides speed advisories after the aircraft has passed the TOD point or while the aircraft is descending. The part of the DA algorithm handling these situations is called mid-descent logic. If an aircraft has moved passed the optimum TOD point or is already in descent, this logic first computes the constant descent angle required to reach the desired altitude of the time control point from its current position. Then it determines the speed range within which the aircraft can fly at this descent angle and still be able to decelerate to a specified final speed. In this step, the use of speed brakes is permitted in any part of the descent. The speed range obtained will grow smaller as the angle of descent increases. Eventually, as the angle of descent reaches a critical value, the speed range shrinks to zero and the DA algorithm fails to generate a feasible profile.

Cruise speed mode– In this mode the DA algorithm iterates only on the cruise speed segment in order to achieve a specified spacing or arrival time. For the descent portion of the flight, the algorithm always assumes the airline SP profile. The limits on the cruise speed envelope used in this mode consist of a maximum cruise Mach number, which depends on weight, altitude, and temperature, and a minimum cruise indicated airspeed. Both limits depend on aircraft type and engine model. Figure 9 shows plots of Mach numbers, IAS, and altitude as a function of range for the maximum, minimum, and nominal cruise speeds.

The cruise speed advisory (CS) appears as the fifth item in the DA Clearance Panels of table 2 and in figure 6. It consists of a desired cruise Mach number or, if the aircraft is cruising below 28,000 ft, a desired indicated airspeed in knots.

In order to compensate for controller and pilot delays in executing a cruise speed advisory, the DA algorithm also includes a 15-sec delay in the start time of cruise speed changes, measured from the aircraft position at the time of calculation. This delay compensation combined with the periodic refresh of the profiles and clearances every 10 sec helps to reduce time errors in executing the cruise speed advisory.

This mode should be used early in arrival sequencing when an aircraft is at least 50 n.mi. from the TOD point. At a smaller distance, the available time range is too small to be useful in arrival time control. By using cruise speed control only, the controller can complete most of the repositioning of an aircraft necessary to achieve a specified arrival time before the aircraft begins its descent. Since the descent profiles in this mode are all based on standard procedure types, the speed differential between in-trail descending aircraft is likely to be small. This has the advantage of minimizing the occurrence of overtakes and thereby helps reduce the controller's effort in monitoring the descent traffic.

Cruise-plus-descent speed mode– This mode provides the greatest possible time range by allowing the DA algorithm to control both the cruise speed and the descent speed profiles. The speed limits that determine the extreme values of the time range are generally identical to those used in the other two speed-control modes. An exception to this rule is the maximum Mach number in descent which in this case is set equal to the maximum cruise Mach number. This assumption sacrifices a small amount of time range but has the advantage of simplifying both the algorithm and, as will be shown below, the specification of the speed clearances.

In generating a cruise-plus-descent speed profile, the algorithm first attempts to meet the specified arrival time by cruise speed control only. If the attempt succeeds, the profiles generated are similar to those in the cruise speed mode. If the attempt fails, cruise speed will have reached one of its limits without achieving the specified arrival time. Descent speed control is then initiated with the cruise speed set from the limit value. Another condition imposed on the descent speed profile is that the Mach number of the constant Mach descent segment, if one is necessary, be the same as the Mach number in cruise. This condition limits the number of the speed clearance parameters that the controller has to issue to two, which is the same as for the other two modes.

When an aircraft is less than 25 n.mi. from the TOD point of the SP profile, the algorithm reverts to descent speed control only. Plots of Mach number, IAS, and altitude as a function of range for maximum, minimum, and nominal speed profiles in cruise and descent are given in figure 10.

Computation of Profile Tracking Errors

As explained in a preceding section, the controller normally "locks" a profile immediately after issuing to an aircraft the clearances displayed in the DA clearance panel. Locking the profile sends a signal to the DA algorithm to cease generating new profiles and initiate tracking of the locked profile. The profile tracking process consists of computation of errors in time, lateral position, and altitude between the current position of the aircraft and the locked profile.

As the first step in computing these errors, the error tracking algorithm determines if the current aircraft position lies within a three-dimensional corridor centered around the locked profile. In the horizontal plane, the corridor extends 4 n.mi. on each side of the center of the locked profile, while in the vertical plane it extends 1000 ft above and below the altitude of the locked profile. These corridor dimensions can be changed for different conditions. Where the horizontal path consists of a circular turn segment, such as in a transition between two straight line segments, the corridor takes the shape of an annulus (fig. 11). If the current aircraft position is found to be outside this corridor, the aircraft is considered to be off track, and the rest of the error computations are skipped.

17

If the aircraft is within the route corridor, the error computation algorithm first projects the actual aircraft position onto the path of the locked profile in order to locate a projected position (see fig. 11). Then it determines the increment or decrement in current time needed to bring the locked profile position into coincidence with the projected position. An increment in time corresponds to early arrival (negative time error) and a decrement to late arrival (positive time error). The time error, in seconds, is displayed in the DA panel to the right of the locked profile arrival time. In figure 3, PA 001 has accumulated a small time error relative to its locked (DA) profile. The time error is also incorporated in the position of the predicted arrival time along the time line via algebraic addition of the time error to the locked profile arrival time. Thus, even when the locked clearance is no longer displayed in the DA clearance panel, the controller can still monitor the time error by comparing the blue scheduled time (or the green cursor position, as the case may be) with the predicted arrival time.

In the current implementation, the last computed time error in the DA panel is appended with a question mark, if the current aircraft position lies outside the route corridor. Similarly, a question mark is also placed after the corresponding predicted arrival time ID tag on the time line. The error computations are repeated approximately every 20 sec.

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 estimated times of arrival (ETAs) 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.

Definitions

The variables and parameters that play crucial roles in the operation of the scheduler are defined and explained below. These definitions apply equally to schedulers designed for the Center and the TRACON.

Estimated time of arrival (ETA): The ETA plays a decisive role in the operation of the TMA and its scheduler. The TMA does not generate this quantity itself but instead obtains it by request from the DA designed for the arrival area where the aircraft is currently located. Since ETAs can change during the approach, the TMA obtains updated ETAs about every four radar sweeps or approximately twice per minute. The method of obtaining ETAs from arrival area DAs rather than from a master DA implemented in the TMA workstation has two advantages. First, it distributes the computational load over several workstations, thus allowing the TMA workstation to concentrate its computational resources exclusively on the generation of schedules. Second, it produces the most accurate values for the ETAs because each arrival area DA can more easily account for the effect of recently issued controller clearances as well as changes in aircraft descent procedures.

The ETA obtained when an arrival is first tracked by the Center's radar upon penetrating the Center airspace is referred to as the original estimated time of arrival (OETA). It is based on the planned arrival route and standard procedure descent profile. This value is used by the TMA as a reference value in computing the accumulated delays of an arrival during its transition through the Center's airspace.

Minimum time to landing: This quantity is defined as the earliest time an aircraft can arrive at the runway from its current location and altitude. It is based on the DA-computed time range and used by the scheduler to determine the earliest feasible time an aircraft can be scheduled to land.

Time to landing: This is defined as ETA – Current time. It does not include time delays imposed by the scheduler.

Scheduled time of arrival (STA): This time is generated by the scheduler.

Scheduling horizon: This is a time interval specified by the traffic manager. It determines when an aircraft is first added to the list of aircraft currently being scheduled, referred to as the schedulable list. An aircraft is added to this list when its time to landing based on its current ETA first penetrates (becomes less than the time of) the scheduling horizon. This parameter is individually adjustable for each of the four arrival areas. Typical values will be discussed shortly.

Freeze horizon: This traffic manager specified time interval determines when an aircraft STA becomes frozen and is transferred from the schedulable list to the frozen STA aircraft list.

Scheduling window: The time interval between the scheduling horizon and the freeze horizon is the scheduling window. Aircraft whose time to landing (based on ETA) fall in this window make up the list of schedulable aircraft. The scheduler generates new STAs for this list when an aircraft in it receives an updated ETA, when a new aircraft is added to the list or when the traffic manager makes parameter changes. Thus, the STA of an aircraft in the scheduling window is subject to revision until it drops below the freeze horizon. It should be noted that the placement of an aircraft in the schedulable list depends solely on its ETA and not on its STA generated by the scheduler.

Blocked time interval: This is an interval of time in which aircraft are not allowed to be scheduled. Such intervals can be entered manually into the TMA data base by the traffic manager. It provides a means to specify a temporary closure of the runway, or to reserve time intervals for takeoffs and emergencies.

Blocked time slot: This designates a time slot for an aircraft not being tracked by the Center radar but expected to land at the specified time. Such aircraft may originate at nearby feeder airports or fly at altitudes too low to be tracked by the Center radar. The scheduler treats a blocked time slot as a specific type of aircraft with a specified landing time.

Although the scheduling and freeze horizons were defined above as 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 12 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. This control process takes place between the freeze horizon where scheduled times are broadcast to the arrival controllers, and the feeder gates where traffic leaves the Center airspace and enters the TRACON. 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 time range of 35-45 min to touchdown as a reasonable time range for the freeze horizon.

In order for the scheduler to have the greatest freedom in optimizing arrival schedules and sequences, the scheduling horizon should be chosen as far out as possible thereby maximizing the width of the scheduling window. However, its maximum width is constrained by the location of the Center boundary unless detailed aircraft data can be obtained from adjacent Centers. In general, the greater the width of the scheduling window, the more nearly optimal will be the arrival sequences generated by the scheduler.

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 12 as the TRACON rescheduling region. In this transition region, the scheduling process described above can be repeated to a limited degree by the TRACON scheduler which is embedded in the FAST. Since the scheduling window is narrow and close to touchdown, rescheduling in FAST 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.

Design of Scheduler

Optimization of aircraft arrival schedules has been the subject of numerous studies in recent years (refs. 5-9). However, a comprehensive analysis of the benefits of optimization, such as its potential for reducing delays, is still lacking. 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. A survey of recent studies leads to the conclusion that the average delay reductions obtainable from schedule optimization is not likely to exceed 10-15% at representative airports. However, additional analyses should be done to provide a more precise estimate of delay reductions and other benefits compared to today's operations.

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. 5) 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.

Four scheduling techniques, selectable by the user, have been implemented in the TMA. They are referred to as first-come-first-served (FCFS) without time advance, first-come-first served with time advance, position shift without time advance, position shift with time advance. The characteristics of each method are briefly described below.

The FCFS method, either without or with time advance, assigns the landing order in the same sequence as the ETA order. If no time advance is permitted, the ETA of an aircraft determines its earliest permissible STA, whereas if time advance is permitted, the minimum time to landing determines its earliest permissible STA.

The first step in all four scheduling methods is to arrange aircraft in the schedulable list in order of their ETAs. Thus, the aircraft with the earliest ETA is placed at the bottom of the list and the aircraft with the latest ETA is placed at the top of the list.

The next step performed by the FCFS scheduler is to check 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. Here, heavy aircraft are primarily wide body, large aircraft are narrow body, and light aircraft are general aviation 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 6. 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.

The operation of the FCFS scheduling algorithm on the schedulable list can be illustrated graphically with the help of time lines drawn side by side as in figure 13(a). The time line on the right shows the ETAs of several large and heavy aircraft within a scheduling window. As in the DA time line, 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 FCFS scheduler is illustrated in figure 13(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. 7) and subsequently by others (ref. 8). The scheduler implemented here optimizes the STA list with respect to one position shift. This means that the landing order of an aircraft may not be moved more than one aircraft ahead of or behind the FCFS order. The schedule produced by the position shift scheduler with time advance is illustrated in figure 13(c) for the same schedulable list as before. 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 13(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.

Description of Graphical Interface and Tools

Similar to the interface for the DA previously described, the interface for the TMA is also based upon exploiting the interaction between workstation screen and the mouse.

The workstation screen is divided into one large and three small nonoverlapping windows. The large window on the left in figure 14 displays aircraft arrival schedules on several reconfigurable time lines. The window on the upper right gives an overview of traffic in the Center in a miniature PVD. The middle window gives the status of all DAs providing ETA data to and receiving STA data from the TMA. The window on the lower right acts as the traffic manager's control panel where various scheduling parameters such as the airport acceptance rate and the configuration of the time lines are selected by mouse actuated switches and sliders as well as keyboard entries.

An additional window is available for displaying information about the schedule. This window, which is displayed by a keyboard command, overlays the other windows and can be positioned anywhere on the screen. Data displayed in this window includes the currently selected acceptance rate, average and peak delays, and various data on individual aircraft which the controller can select by picking the aircraft time line tag.

The time line window contains two pairs of time lines on which three types of time schedules can be selectively displayed as follows (from right to left): (1) ETAs of aircraft that have not yet entered the Center airspace; such ETAs are contained in flight plans sent to the Center ahead of time; (2) ETAs of aircraft tracked by the Center radars and sent to the TMA by the various DAs; and (3) STAs of all aircraft which will be or have been sent to the various DAs. These time schedules can be selectively displayed on both the left and right side of each time line. Furthermore, the display of these time schedules can also be segregated by arrival area through use of toggle switches in the control panel window. Thus, the traffic manager has available numerous options to optimize the display configuration for

monitoring and managing traffic flows. For example, one possible display configuration, which is illustrated in figure 14, uses the left track of the leftmost time line to display STAs for the two northern gates, Drako and Keann and the right track of that time line to display the STAs for the two southern gates, Byson and Kiowa. The same rule is used to segregate the ETAs on the second time line, and the flight plan ETAs on the third time line. This leaves the fourth time line as a spare (shown as turned off in fig. 14). Since traffic from the northern and southern gates generally lands on the northern and southern parallel runways, respectively, at Denver, this display configuration also has the advantage of segregating the flows by landing runway.

As previously described for the DA time line, aircraft time schedules move toward the bottom of their respective time lines as time increases. Moreover, under normal circumstances, aircraft schedules will be observed to migrate from the right to the left time lines. An aircraft first appears on the flight-plan time line at the time the Center receives its flight plan and planned ETA. When the aircraft becomes active in the Center airspace, its ETA is updated and it is simultaneously removed from the flight plan time line and displayed on the ETA time line. Finally, when its ETA penetrates the scheduling horizon, its STA is computed and then displayed on the STA time line.

Color coding of aircraft IDs and graphical markers are used on the time lines to convey the aircraft scheduling status and critical scheduling parameters. The scheduling windows appear as yellow vertical bars on each side of the ETA time line. Aircraft IDs appear in green on the flight plan and ETA time line. On the STA time line, scheduled aircraft with yellow IDs designate aircraft whose ETA is in the scheduling window. At the time the ETA of an aircraft falls below the freeze horizon, its color on the STA time line changes from yellow to blue. At this time the aircraft's scheduled time is sent to an appropriate PVD for display on the PVD time line.

The traffic manager also can select a TMA mode that broadcasts the STAs to the PVDs immediately after they are first calculated; that is, while aircraft are still above the freeze horizon. This mode is desirable when the queue of aircraft in the STA list includes significant delays (more than 5 min). By broadcasting the STAs early under these conditions the arrival controllers are given additional time to plan optimum strategies for absorbing large delays. Since the STAs that are broadcast above the freeze horizon can be revised by the scheduler, they are identified on the PVD time line by orange colored tags, which turn blue when the aircraft cross the freeze horizon.

Interaction with time lines– Of particular importance to the implementation of the TMA was the requirement that the traffic manager be able to interact with the automatic scheduler. The TMA interface has been structured to allow traffic managers considerable flexibility in modifying the computer generated schedule for specific aircraft, while allowing the automatic scheduler to continue generating schedules for the other eligible aircraft. The operational philosophy of the TMA is analogous to the concept of text flow in desktop publishing applications. An area of the page is designated for specific uses, (usually insertion of graphics) and the text flows around the designated area. In a similar fashion the TMA uses the controller inputs to determine time slots and intervals in which it cannot schedule aircraft and works around these areas by scheduling aircraft in the next available slot.

Also of importance is the immediate feedback available to the traffic manager when he or she does modify the computer plan. The computer immediately modifies the scheduled aircraft display to reflect the traffic manager's input, making it easy to see the effect of his/her actions.

Some of the ways in which the traffic manager can interact with the automatic scheduler will now be covered. In the following explanations, frequent reference will be made to figure 15. This figure shows two time lines. The time line on the right displays ETAs while the time line on the left displays STAs for the same set of aircraft.

Manual scheduling– The traffic manager can alter the scheduled time of any aircraft currently in the system (including those which have not yet been scheduled, and those whose times are already frozen). This is done by placing the mouse cursor over the aircraft time line tag, depressing the middle button of the three-button mouse used in Sun workstations and dragging the tag to a new location (time) on the time line. As soon as the middle button is released, the computer will generate and display an updated schedule. Aircraft scheduled in this fashion are displayed in purple to highlight the fact that they have been manually scheduled by the controller. In the figure both PA001 and SP404 have been manually scheduled.

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 in red; intervals and slots. An interval is represented by double red bars. The scheduler will not place any aircraft in the area delimited by the blocked interval. This area is created by placing the mouse cursor anywhere within the two vertical lines on either side of the minute numbers and dragging with the middle mouse button depressed. 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 placing the mouse cursor in the same area as before but using the right mouse button to display a menu of slot options. Heavy, large, and small slots can be added by selecting the appropriate menu item. 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. This is because not enough time exists behind UA134 based on its current ETA to provide the minimum separation required between it and the heavy slot. Since the slot is fixed, the scheduler has to move UA134.

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.

The rescheduling feature is needed when the controller makes changes that affect aircraft which have passed the freeze horizon. This could occur during runway changes, missed approaches, pop-ups, and so on. Once aircraft are frozen they are not touched by the automatic scheduler unless the controller specifically requests it. The menu can be used to reschedule all frozen aircraft, or only those behind the aircraft whose tag was used to select the menu. For example, if CO711 is selected, only CO564 will be rescheduled.

The rebroadcasting feature is needed because once aircraft have passed the freeze horizon, their scheduled time will not be automatically rebroadcast unless the controller explicitly requests it. This is to

prevent confusion and any excessive workload of arrival sector controllers caused by too-frequent schedule changes. If an aircraft scheduled time is different from that previously broadcast to a PVD the aircraft tag will be displayed with an exclamation mark appended to it. In the figure, PA001 has an exclamation mark, warning the traffic manager that its currently scheduled time (manually generated in this case) is different from that previously broadcast to the PVD.

The menu command for returning an aircraft to automatic scheduling is generally used when the traffic manager has been scheduling an aircraft manually but has decided to return the aircraft to automatic scheduling, prior to the aircraft having crossed the freeze horizon. If this command is chosen, the aircraft will be again displayed in yellow and scheduled in with the rest of the currently schedulable aircraft.

CONCLUDING REMARKS

The human-centered automation concept described in this paper deliberately places the automation tools 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 two types of tools described in this paper are designed to assist Center controllers in managing arrival traffic. Thus, the Traffic Management Advisor generates efficient landing schedules and assists the Center traffic manager in flow control. Then, the Descent Advisor tools are used by controllers at each arrival gate to implement these schedules accurately. Although both tools are indispensable elements of the concept, the design of Descent Advisor presents the greatest technical challenge.

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.; and Tobias, L.: A Time-Based Concept for Terminal-Area Traffic Management. Proceedings of the 1986 AGARD Conference No. 410, Efficient Conduct of Individual Flights and Air Traffic, pp. 52-1 - 52-14.

2. Erzberger, H.; Nedell, W.: Design of Automation Tools for Management of Descent Traffic, NASA TM-101078, Dec. 1988.

3. Green, S.; and Davis, T.: A Piloted Simulator Evaluation of a Ground-Based 4D Descent Advisor Algorithm. AIAA-87-2522-CP, Guidance, Navigation and Control Conference, Monterey, CA, Aug. 17-19, 1987.

4. Tobias, L.; and Scoggins, J.: Time-Based Air Traffic Management Using Expert Systems. Proceedings of the 1986 American Control Conference, Seattle, WA, June 1986, pp. 693-700; also NASA TM-88234, 1986.

5. 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.

6. 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.

7. Dear, R.: Increasing Capacity with Computer-Assisted Decision Making., Proc. International Air Transportation Conf., Vol. 2, May 1979.

8. Luenburger, Robert: A Traveling-Salesman-Based Approach to Aircraft Scheduling in the Terminal Area. NASA TM-100062, March 1988.

9. 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.

TABLE I.– GUIDELINES FOR DESIGN OF AUTOMATED ATC SYSTEMS

DO NOT: Automate unique skills or enjoyable tasks of controllers

    Automate complex or poorly understood tasks

    Automate in ways that reduce situational awareness

    Automate such that a system failure leaves controller with an impossible problem to solve

DO: Automate to enrich controller's work environment

    Automate to increase situational awareness

    Automate to complement controller's skills

    Involve controllers from the start in selection and design of automation tasks

TABLE 2.– PROFILE PARAMETERS DISPLAYED IN THE DA CLEARANCE PANEL

1	2	3	4/5	6	7/8	9	10	11	12	13	14
PA001	CLR	RI	0.70/–	68	0.71/255	JASIN	FL110	J56	42:06	+5	38.5
TA321	SP	RI	0.71/–	62	0.75/260	JASIN	FL110	J56	38:03	–	7.7
AA484	I	RI	0.74/0.79	61	0.79/310	JASIN	FL110	J170	36:28	–	0.0

1. A/C ID
2. Profile type
3. Horizontal. guidance mode
4. Current Mach
5. Desired cruise Mach, if applicable
6. DME range to TOD
7. Descent Mach, if applicable
8. Descent KIAS
9. Time control waypoint
10. Crossing altitude
11. Jet route
12. ETA, min:sec
13. Arrival time error, sec
14. Spacing distance, n.mi.

Note: Item numbers 6-8 optionally displayed in A/C data block.
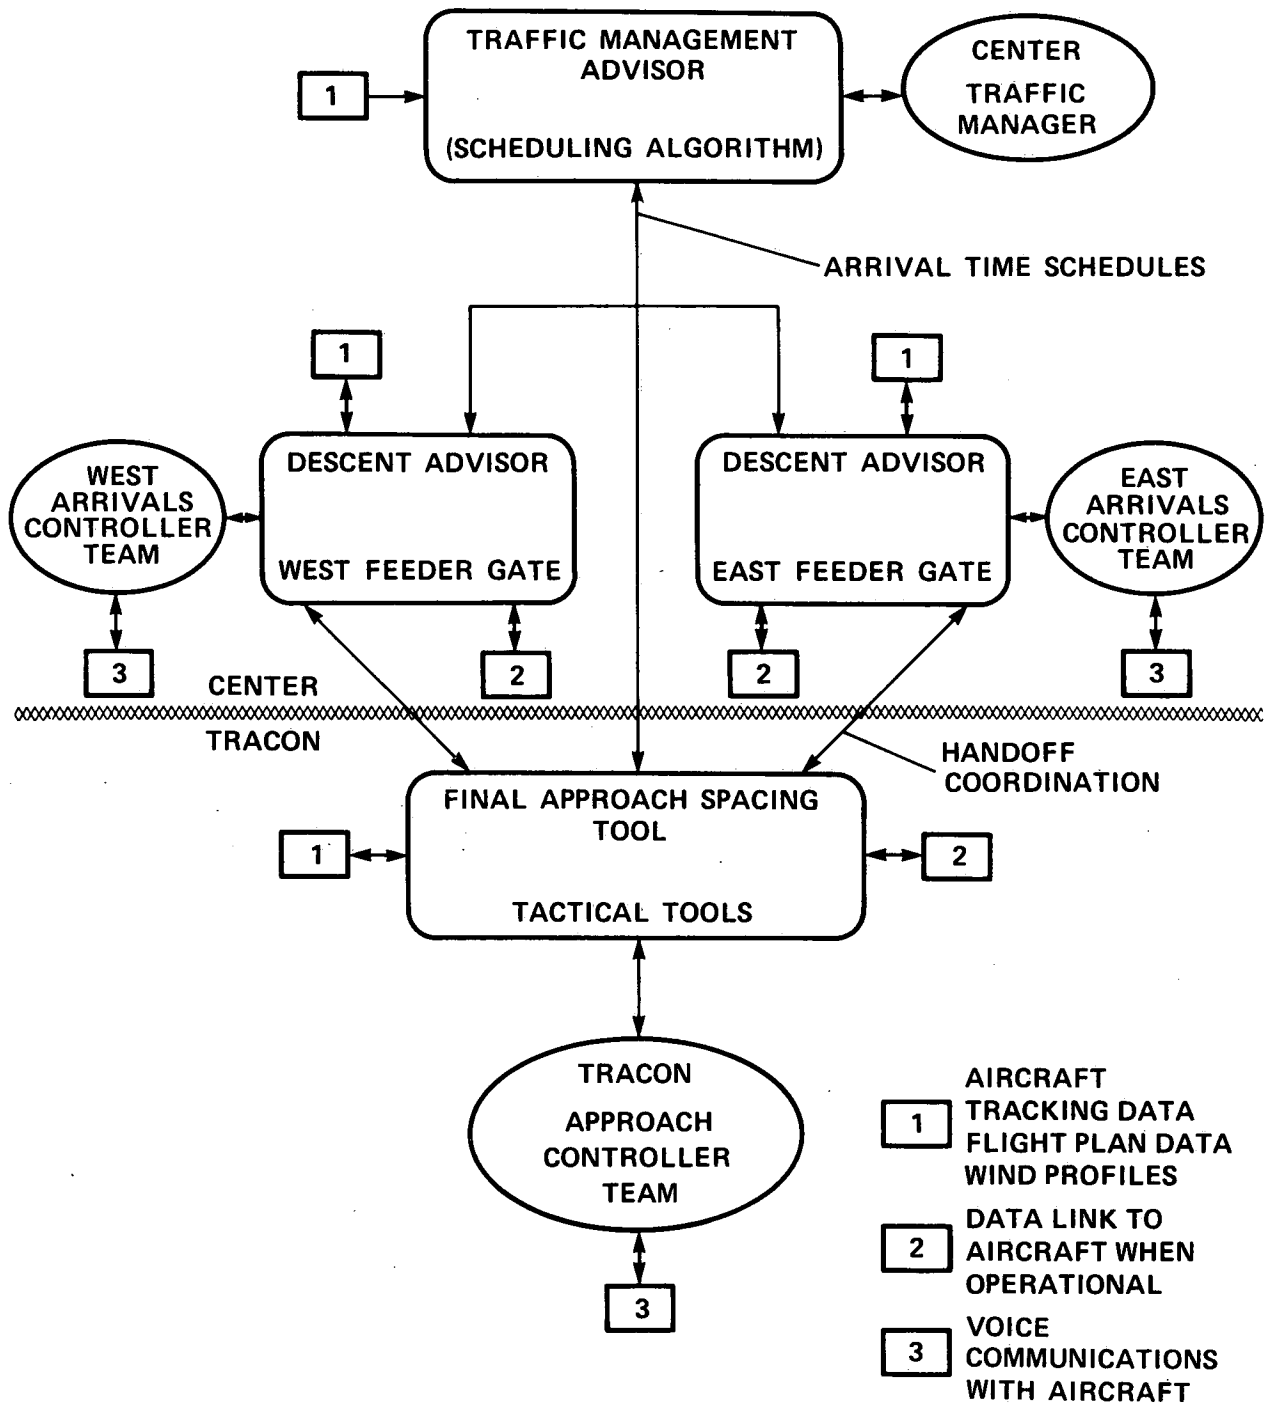
Figure 1.– Automation concept and hierarchy of automation tools.

Figure 2.— Integrated controller display with time line.

Figure 3.– Integrated controller display showing three selected aircraft.

33
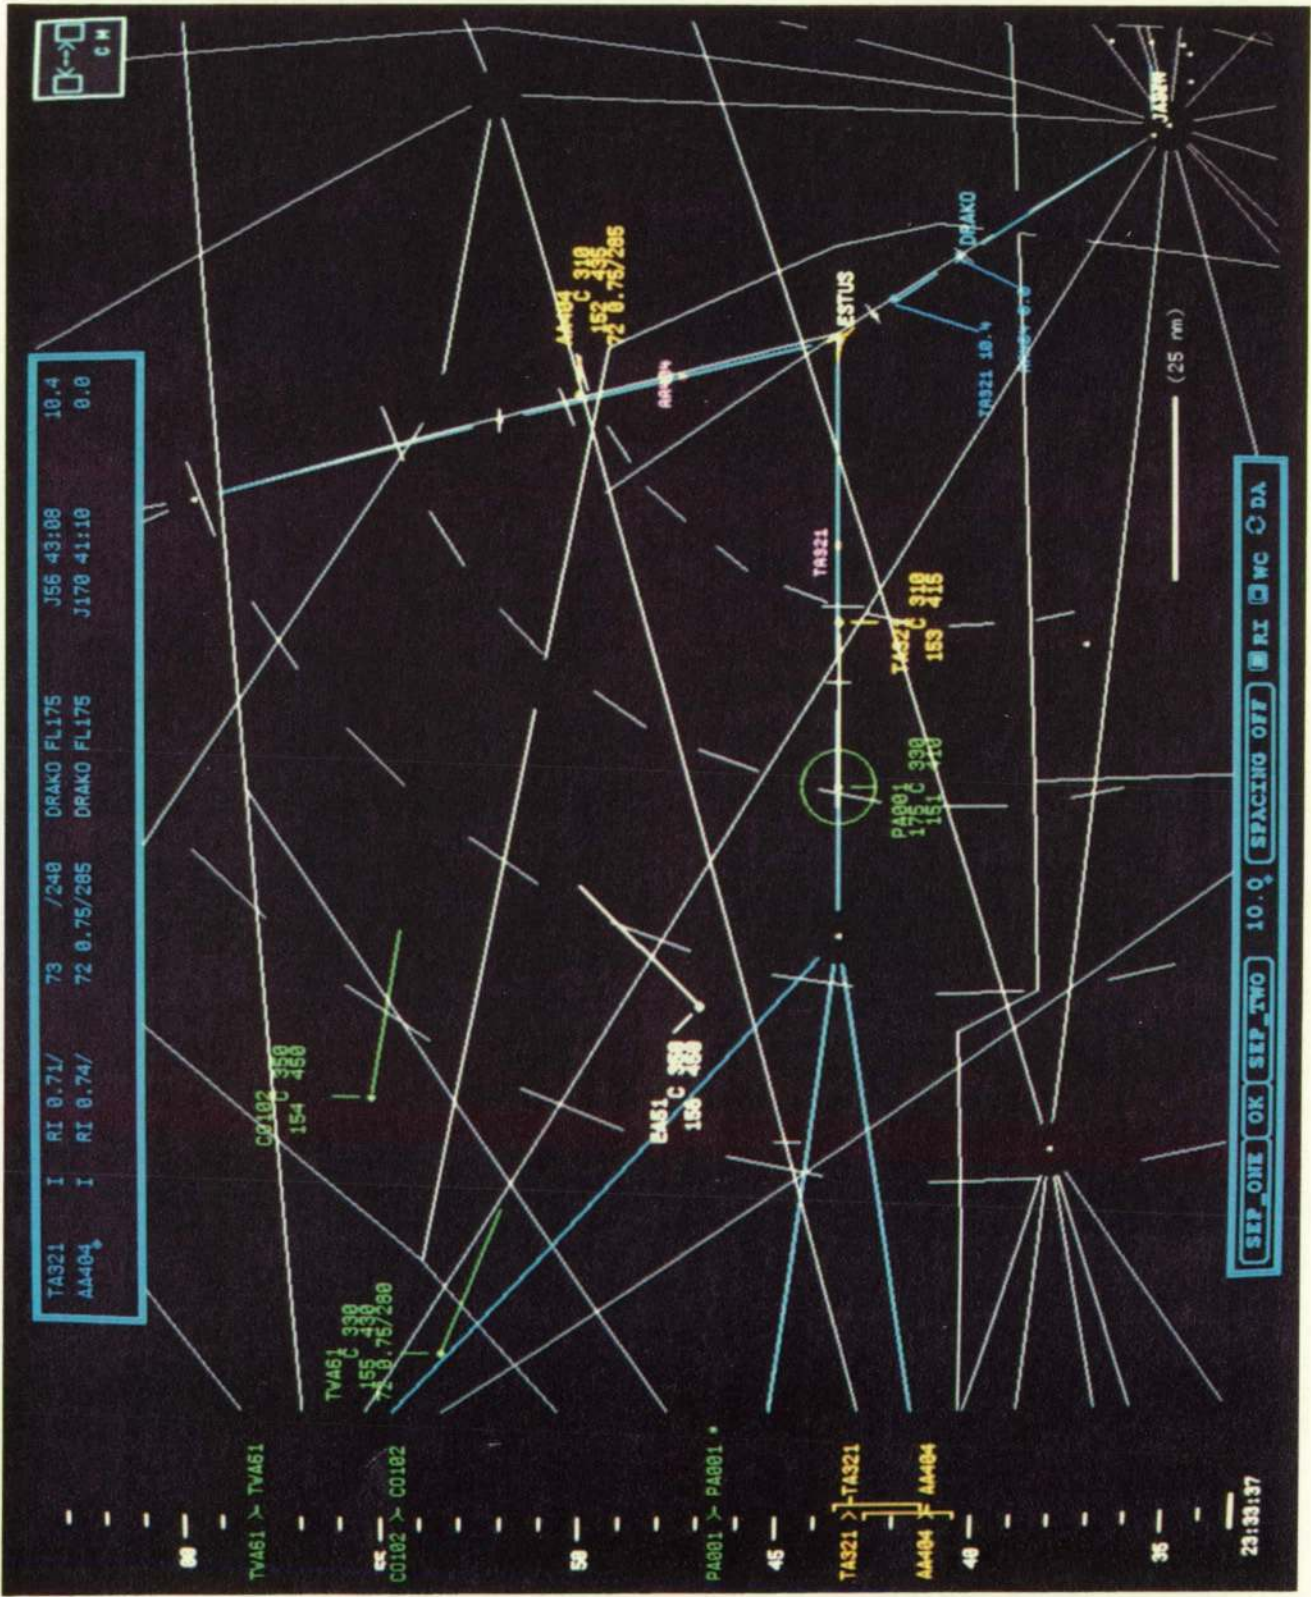
Figure 4.– Integrated controller display with spacing advisories.

ORIGINAL PAGE
COLOR PHOTOGRAPH

PRECEDING PAGE BLANK NOT FILMED

PAGE 34 INTENTIONALLY BLANK

Figure 5.– Sketch of route intercept guidance.

Figure 6.— Integrated controller display illustrating route intercept guidance and cruise speed advisories.

Figure 7.– Sketch of waypoint capture guidance.

41

Figure 8.– Integrated controller display illustrating waypoint capture guidance to Drako and STAs on the time line.

ORIGINAL PAGE
COLOR PHOTOGRAPH

43

Figure 9.– Cruise-only speed and altitude profiles. (a) Mach number; (b) KCAS; (c) altitude.

PRECEDING PAGE BLANK NOT FILMED

Figure 10.– Cruise-descent speed and altitude profiles. (a) Mach number; (b) KIAS; (c) altitude.

LOCKED PROFILE REFERENCE
POSITION AT TIME t

POSITION OF
AIRCRAFT
AT TIME t

PROJECTED
AIRCRAFT
POSITION

BOUNDARIES
OF CORRIDOR

DIRECTION
OF TRAVEL
ALONG ROUTE

4 n. mi.

4 n. mi.

CENTER
OF TURN

LONGITUDINAL
ERROR

REFERENCE ROUTE

Figure 11.– Definition of profile corridor and computation of time error.

TRACON BOUNDARY
& RESCHEDULING HORIZON
~15 min TO TOUCH DOWN

SCHEDULING
WINDOW

SCHEDULING
HORIZON ~45 min
TO TOUCH DOWN

N.W.
ARRIVALS
AREA

N.E.
ARRIVALS
AREA

DRAKO

KEANN

S.W.
ARRIVALS
AREA

KIOWA

BYSON

TRACON
FREEZE
HORIZON
~10 min TO
TOUCH DOWN

S.E.
ARRIVALS
AREA

TRACON
RESCHEDULING
REGION

FREEZE HORIZON
~35 min
TO TOUCH DOWN

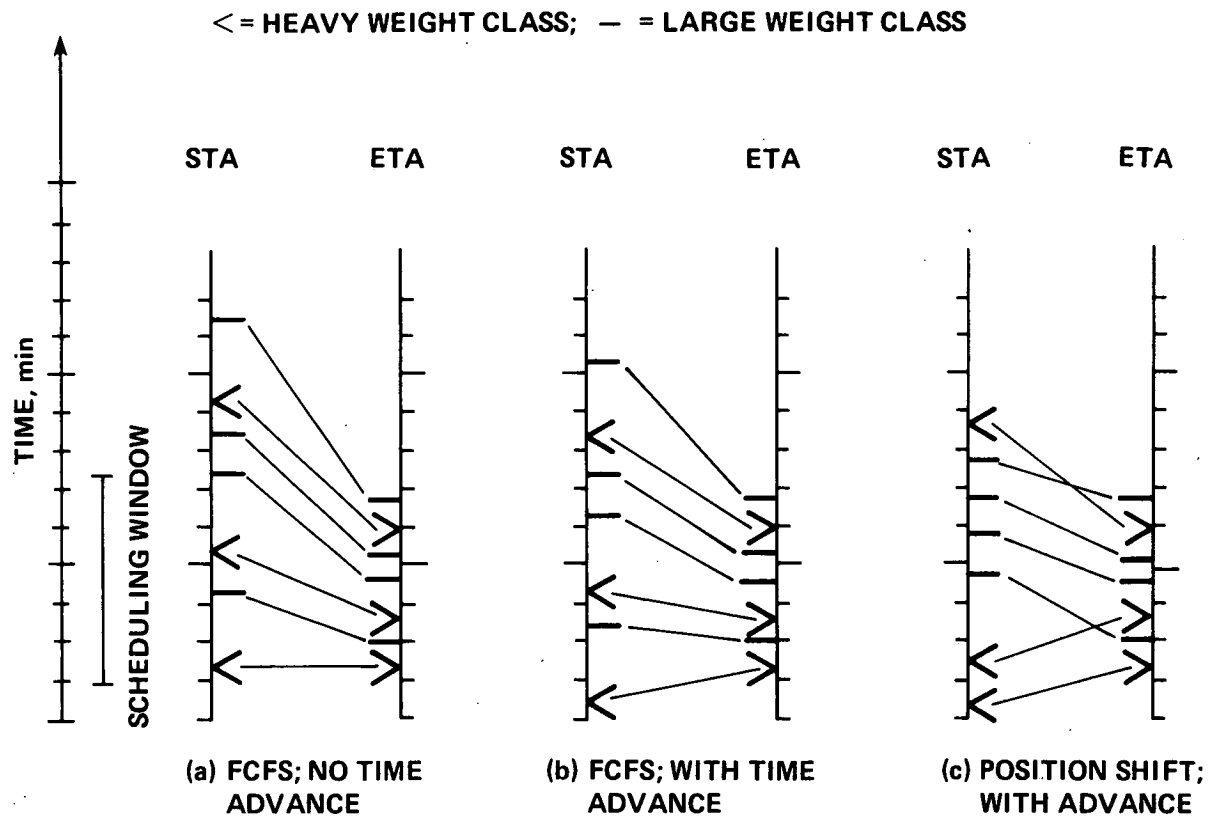Figure 12.– Arrival routes and scheduling regions at a Center.

47

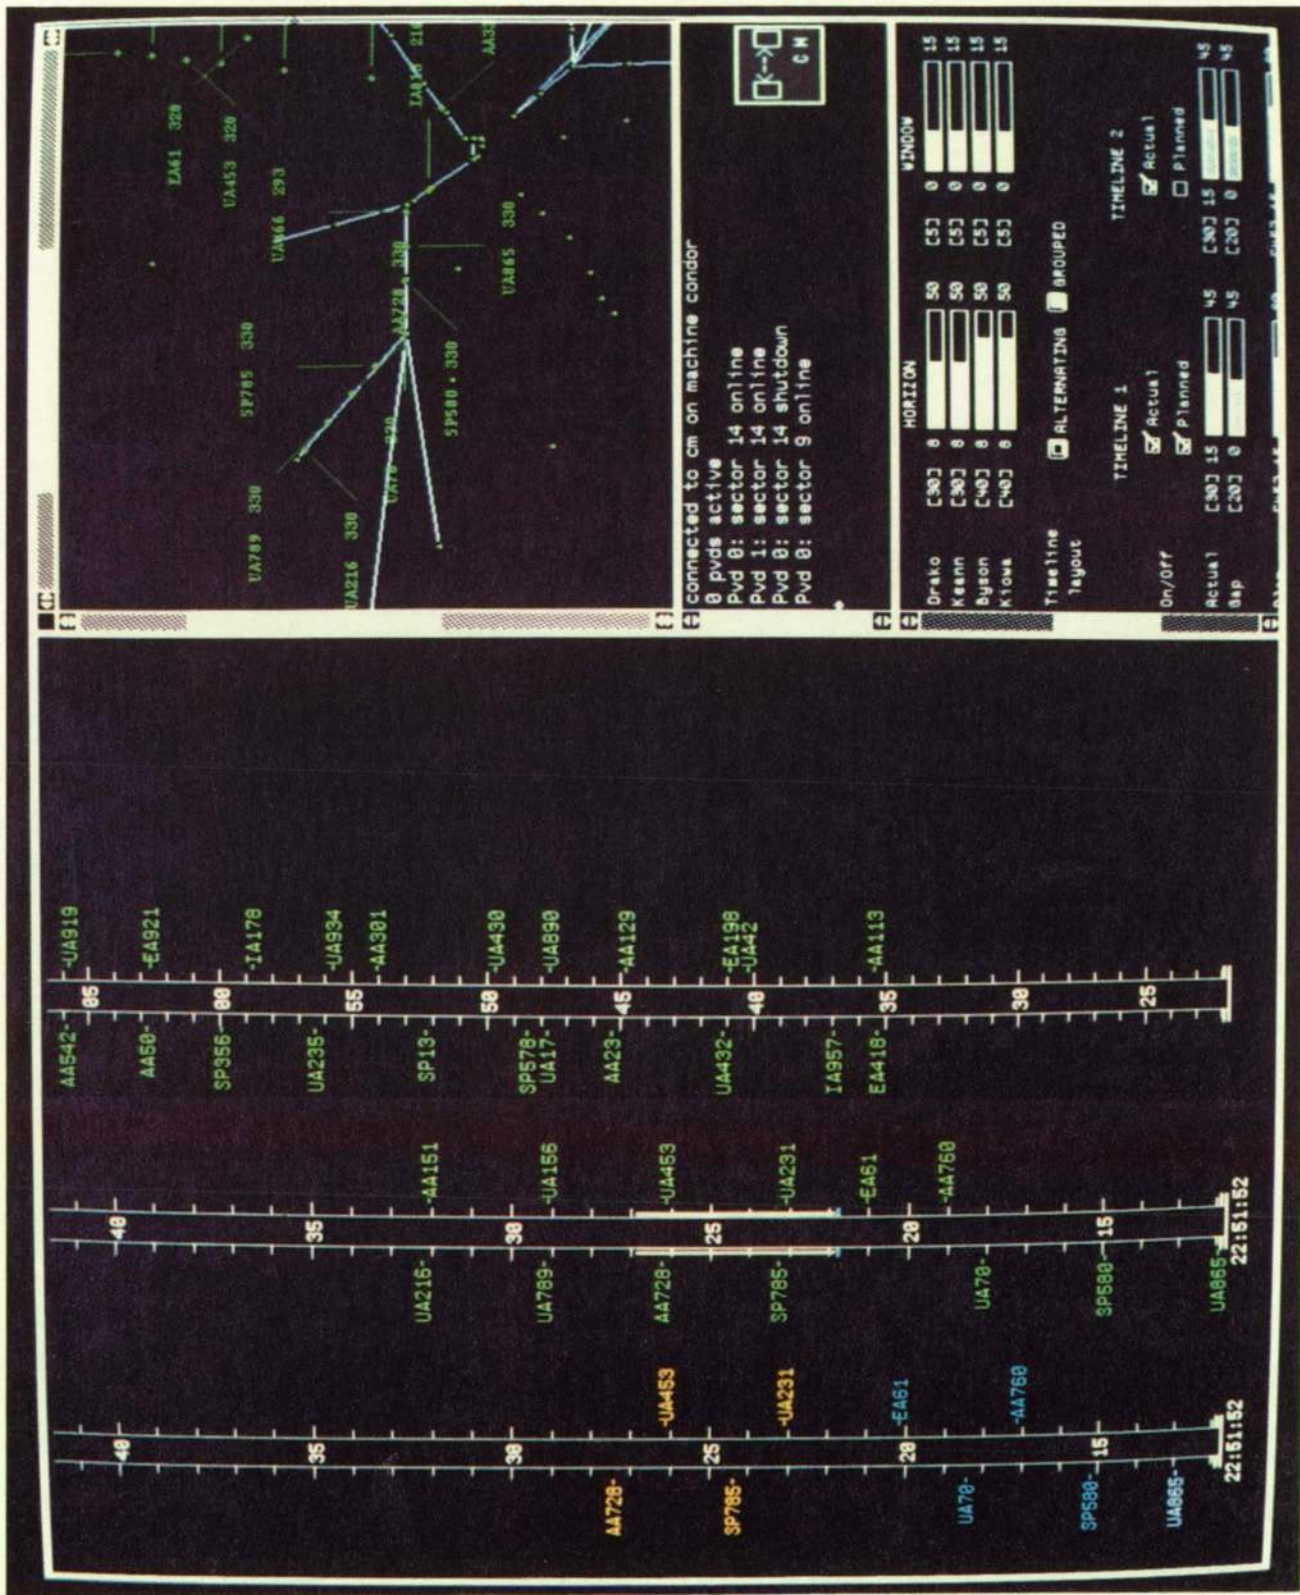Figure 13.– Effect of three types of scheduling methods on delays.

Figure 14.– Screen photograph of Traffic Management Advisor display.

ORIGINAL PAGE
COLOR PHOTOGRAPH

Figure 15.– Time lines for flow monitoring and manual scheduling.

NASA
National Aeronautics and
Space Administration

Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA TM-102201		

4. Title and Subtitle		5. Report Date
Design of Automated System for Management of Arrival Traffic		June 1989
		6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Heinz Erzberger and William Nedell (San Jose State University, San Jose, CA)	A-89116
	10. Work Unit No.
	505-67-21

9. Performing Organization Name and Address	11. Contract or Grant No.
Ames Research Center Moffett Field, CA 94035	
	13. Type of Report and Period Covered

12. Sponsoring Agency Name and Address	Technical Memorandum
National Aeronautics and Space Administration Washington, DC 20546-0001	14. Sponsoring Agency Code

16. Abstract

This paper describes the design of an automated air traffic control system based on a hierarchy of advisory tools for controllers. Compatibility of the tools with the human controller, a key objective of the design, is achieved by a judicious selection of tasks to be automated and careful attention to the design of the controller system interface. The design comprises three interconnected subsystems referred to as the Traffic Management Advisor, the Descent Advisor, and the Final Approach Spacing Tool. Each of these subsystems provides a collection of tools for specific controller positions and tasks. This paper focuses on the design of two of these tools, the Descent Advisor, which provides automation tools for managing descent traffic, and the Traffic Management Advisor, which generates optimum landing schedules. The algorithms, automation modes, and graphical interfaces incorporated in the design are described.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
Air traffic control Air traffic control automation Automated system design	Unclassified — Unlimited Subject category — 03

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of pages	22. Price
Unclassified	Unclassified	53	A04

Report Documentation Page

1. Report No.	2. Government Accession No.	3. Recipient's Catalog No.
NASA TM-102201		

4. Title and Subtitle		5. Report Date
Design of Automated System for Management of Arrival Traffic		June 1989
		6. Performing Organization Code

7. Author(s)	8. Performing Organization Report No.
Heinz Erzberger and William Nedell (San Jose State University, San Jose, CA)	A-89116
	10. Work Unit No.
	505-67-21

9. Performing Organization Name and Address	11. Contract or Grant No.
Ames Research Center Moffett Field, CA 94035	

12. Sponsoring Agency Name and Address	13. Type of Report and Period Covered
National Aeronautics and Space Administration Washington, DC 20546-0001	Technical Memorandum
	14. Sponsoring Agency Code

15. Supplementary Notes

Point of Contact: Heinz Erzberger, Ames Research Center, MS 210-9, Moffett Field, CA 94035
(415) 694-5425 or FTS 464-5425

16. Abstract

This paper describes the design of an automated air traffic control system based on a hierarchy of advisory tools for controllers. Compatibility of the tools with the human controller, a key objective of the design, is achieved by a judicious selection of tasks to be automated and careful attention to the design of the controller system interface. The design comprises three interconnected subsystems referred to as the Traffic Management Advisor, the Descent Advisor, and the Final Approach Spacing Tool. Each of these subsystems provides a collection of tools for specific controller positions and tasks. This paper focuses on the design of two of these tools, the Descent Advisor, which provides automation tools for managing descent traffic, and the Traffic Management Advisor, which generates optimum landing schedules. The algorithms, automation modes, and graphical interfaces incorporated in the design are described.

17. Key Words (Suggested by Author(s))	18. Distribution Statement
Air traffic control Air traffic control automation Automated system design	Unclassified — Unlimited Subject category — 03

19. Security Classif. (of this report)	20. Security Classif. (of this page)	21. No. of pages	22. Price
Unclassified	Unclassified	53	A04