

1N-03

11926-7

P-44

Analysis of Delay Reducing and Fuel Saving Sequencing and Spacing Algorithms for Arrival Traffic

Frank Neuman and Heinz Erzberger

(NASA-TM-103880) ANALYSIS OF DELAY
REDUCING AND FUEL SAVING SEQUENCING
AND SPACING ALGORITHMS FOR ARRIVAL
TRAFFIC (NASA) 44 p

N92-33194

Unclas

G3/03 0117764

October 1991



National Aeronautics and
Space Administration

Analysis of Delay Reducing and Fuel Saving Sequencing and Spacing Algorithms for Arrival Traffic

Frank Neuman and Heinz Erzberger, Ames Research Center, Moffett Field, California

October 1991



National Aeronautics and
Space Administration

Ames Research Center
Moffett Field, California 94035 - 1000

TABLE OF CONTENTS

	Page
SUMMARY	1
ABBREVIATIONS AND DEFINITIONS	1
INTRODUCTION	2
SCHEDULING ALGORITHMS	4
Spacing Requirements	4
Modified First-Come First-Served Algorithm	4
Time-Advance	5
Heuristic time-advance algorithm	5
The fuel saving TA algorithm using a performance function	5
The idealized fuel saving TA algorithm	5
Pure TA	9
The realizable fuel saving TA algorithm	9
Constrained Position-Shift Algorithm	9
Optimal CPS	9
Heuristic CPS	10
Scheduling Algorithm for Two Dependent Runways	10
TRAFFIC MODEL	10
Toward an Accurate Traffic Model	10
Traffic Model for Studying Scheduler Effect	11
RESULTS	13
Time Diagrams of Traffic Samples and Associated Delays	13
Analysis of Traffic Including Both Modes of Optimization	21
Further Traffic Analysis FCFS Only	28
Delay as function of length of traffic samples	28
Effect of the distribution of arrival times	29
Effect of winds and changes in interarrival times on delays	29
Effect of increasing in-trail spacing	29
Effect of more precise guidance using CTAS	29
Results for Scheduling for Two Dependent Runways	31
CONCLUSIONS	34
APPENDIX A FCFS MINIMIZES STANDARD DEVIATION OF THE DELAYS	37
APPENDIX B AN EXACT SOLUTION OF THE CONSTRAINED POSITION-SHIFT PROBLEM FOR THE SINGLE-POSITION SHIFT	39
REFERENCES	41

Summary

The air traffic control subsystem that performs sequencing and spacing is discussed. The function of the sequencing and spacing algorithms is to automatically plan the most efficient landing order and to assign optimally spaced landing times to all arrivals. Several algorithms are described and their statistical performance is examined. Sequencing brings order to an arrival sequence for aircraft. First-come-first-served sequencing (FCFS) establishes a fair order, based on estimated times of arrival, and determines proper separations. Because of the randomness of the arriving traffic, gaps will remain in the sequence of aircraft. Delays are reduced by time-advancing the leading aircraft of each group while still preserving the FCFS order. Tightly spaced groups of aircraft remain with a mix of heavy and large aircraft. Spacing requirements differ for different types of aircraft trailing each other. Traffic is reordered slightly to take advantage of this spacing criterion, thus shortening the groups and reducing average delays. For heavy traffic, delays for different traffic samples vary widely, even when the same set of statistical parameters is used to produce each sample.

This report supersedes NASA TM-102795 on the same subject. It includes a new method of time-advance as well as an efficient method of sequencing and spacing for two dependent runways.

Abbreviations and Definitions

AAR	Airport Acceptance Rate
ARTCC	Air Route Traffic Control Center (also called Center)
ASP	Arrival Sequencing Program (new operational version of the Center traffic management system (1991))
ATC	Air Traffic Control
CPS	constrained position-shift optimization scheduling method
CTAS	Center TRACON Automation System
DA	descent advisor calculates ETAs
ERM	en route metering (previous operational version of the Center traffic management system (before 1991))
ETA	estimated time of arrival at the runway (no interference from other aircraft)
FAA	Federal Aviation Administration

FCFS	first-come-first-served scheduling method
N	traffic density = demand = number of aircraft/hr wanting to land
STA	sequenced time of arrival at the runway (includes required delays)
STA_TA	sequenced time of arrival by the TA algorithm
t_a	the amount of time-advance
TA	time-advance optimization scheduling method
TMA	Traffic Management Advisor
TRACON	Terminal Radar Approach Control Facility
STA	scheduled time of arrival (FCFS)
VAR	Vertex Acceptance Rate

Estimated time of arrival (ETA): The ETA plays a decisive role in the operation of the TMA and its scheduler. 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 Trajectory Synthesizer - 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 $ETA - \text{Current time}$. It does not include time delays imposed by the scheduler.

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

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

Scheduling horizon: This time interval determines when an aircraft is first added to the list of aircraft currently being scheduled, referred to as the scheduleable 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.

Freeze horizon: This traffic manager specified time interval determines when an aircraft STA becomes frozen and is transferred from the scheduleable 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 scheduleable 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 scheduleable list depends solely on its ETA and not on its STA generated by the scheduler.

Introduction

An automated system for air traffic control (ATC) may be divided into three principal subsystems whose functions involve sensing, planning, and controlling. The subject of this report is the planning subsystem that performs sequencing and spacing. In this report, when both sequencing and spacing are referred to as a combined concept, this will be called scheduling for short. The function of the scheduling algorithms is to plan automatically the most efficient landing order and to assign optimally spaced landing times to all arrivals, given the times the aircraft are actually arriving at the Air Route Traffic Control Center (ARTCC). First, the present technology is discussed. Then, several important scheduling algorithms are described, and the statistical performance of the algorithms is examined. *This concept of scheduling must be clearly distinguished from "scheduling" as used by the airlines. It is the prerogative of the airlines to schedule their aircraft, namely to determine their nominal times of departure and arrival.*

An operational sequencing system used at some Air Route Traffic Control Centers (ARTCC) is called En Route Metering (ERM), which has recently been replaced (1991) by the Arrival Sequencing Program (ASP). The main use of ERM was to provide a specified aircraft arrival flow rate from the Center to the TRACON by assigning equally

spaced arrival times based on the inverse of the Airport Acceptance Rate (AAR). The difference between ERM and ASP is that ASP can specify more than one vertex per airport. This means that incoming aircraft can be separately sequenced to more than one runway or groups of runways, where the spacing is based on the inverses of the individual vertex acceptance rates. ASP generates rough approximations for the estimated times of arrival (ETAs) and bases the first-come first-served (FCFS) ordered times of arrival (STAs) on the order of the ETAs. This FCFS ordering method smooths the traffic as much as possible by minimizing the standard deviation of the delays (see appendix I). If more than one vertex is used, the aircraft are sequenced separately to each vertex, based on the individual Vertex Acceptance Rates (VAR). ASP does not perform any other optimization of the sequences, nor does it provide advisories to achieve accurate arrival times. Also, the order assigned by the ASP sequencer is often not followed; rather, it is used by the sector controllers primarily to indicate landing-slot availability to which they may assign any aircraft. Current sequencing and metering does not use spacing dependent on the weight classes of aircraft. However, the AAR or VARs can be changed at any time, which will change the time-intervals for sequencing aircraft. Although the FAA specifies minimum separations for the final approach based on the sequence of aircraft weight classes, this information cannot be used in sequencing and spacing for aircraft delivery to the TRACON, unless the runway destinations are known as each aircraft enters the Center and unless aircraft heading for different vertexes are on separate flightpaths.

In the discussion of the scheduling methods proposed here, it is assumed that the runway destinations are known as the aircraft enter the Center, and that aircraft destined for different runways are on separate flightpaths so that the scheduling problems for different vertexes are independent of each other. In practice, this is true for most aircraft. If, occasionally, different runways are assigned for a few aircraft, this will somewhat increase the predicted delays. In the approach followed here, several methods of scheduling the traffic for one vertex are applied one after the other, each method improving upon the next.

As in the ASP operational system, in the system proposed here, scheduling is always begun with FCFS ordering. With this method, after aircraft enter the ARTCC, they are sequenced in the order in which they are predicted to land by their ETA sequence, while using a flight plan which specifies the nominal path. For greatest fuel efficiency, the Trajectory Synthesizer, which calculates the ETAs, also needs to know on which runway the aircraft is to land as soon as it enters the Center airspace. In contrast to ASP with its equal time spacing between aircraft, the CTAS

FCFS sequencer adds appropriate delay times to insure proper spacing, which depends on the weight classes of the aircraft. There are two additional distinct differences compared to the ASP:

1. The new ETA calculations are based on accurate modeling of aircraft characteristics (lift, drag and thrust models) and pilot procedures, which generate accurate predictions, provided pilots follow the system generated advisories.
2. For heavy traffic, the specified landing sequence automatically guarantees the maximum landing rate based on the types of aircraft arriving, provided the specified runway is usable to its full capacity.

An effective method of reducing the average delay time without changing the FCFS order of the aircraft is called Time-Advance (TA). This method recognizes the beneficial effect of occasionally speeding up the lead aircraft during periods of heavy traffic in order to reduce delays that naturally occur in FCFS sequences. It is called Time-Advance herein and in reference 1 and is called the negative delay effect in references 2 and 3. If TA is used, and the Scheduled Time of Arrival (STA) of the lead aircraft of a group is reduced by one minute, the delays of the following aircraft are reduced by the same amount. The lead aircraft uses extra fuel, while most of the remaining aircraft will save fuel. Time-advance, however, increases the possible landing rate by a negligible amount, since, although pure TA sequences each aircraft t_a minutes earlier than FCFS scheduling, the gaps between the aircraft remain exactly the same as without TA.

For the heuristic time-advance method, which was analyzed in a previous report (ref. 4) a constant t_a of 1 minute was used for all delays, and the leading aircraft of a group was time-advanced only if it would benefit at least one following aircraft. In the earlier system implementation of TA (ref. 1) there was a choice of between 0 and 4 aircraft that must follow the leader of the group without a gap in the sequence before it is time-advanced. This could occasionally be wasteful. Assuming the ETAs are approximately spaced as required for FCFS, TA would force all aircraft in the group unnecessarily to fly faster. Also, a small group may be closely followed by a large group, then, if the small group is not time-advanced, the large group may have an insufficient gap to time-advance it for minimum fuel costs. Two additional elements were missing in the heuristic TA implementation. First, it was not known what is the best amount of t_a to choose for the prevailing demand, and second, it was not known how many aircraft should be time-advanced. In this updated report these elements are addressed, and presented in a simple algorithm which has now been implemented in the real time system described in reference 1.

The data will still be presented for the heuristic TA algorithm, since the data trends for different traffic conditions that were reported earlier do not change for an improved TA algorithm.

The spacing requirements mentioned earlier offer the opportunity to optimize the landing sequence further, thereby improving on the FCFS and TA methods by minimizing the average delay per aircraft. A scheduling optimizing method called constrained position shift (CPS) was developed several years ago by Dear (ref. 5). The CPS method assumes that an initial landing order has been determined by FCFS and that all aircraft are tightly packed, that is, that they have minimum time-separations. By rearranging the landing order, while not shifting any aircraft from its original position in the sequence by more than a few places, the total time between the first aircraft and the last aircraft can often be reduced. Though CPS is conceptually straightforward, its implementation in a real-time algorithm is more complex because of grouping and because of gaps in the arrival sequence. The groups and gaps are due to the randomness of the arrival times of aircraft in the terminal area. CPS must, therefore, be applied to individual groups of aircraft as was done here, or the algorithm's performance index must be rewritten from that given in appendix II so that it minimizes the sum of the sequenced flight times instead. Appendix II, which gives an exact solution of the single position constrained position-shift problem, was written by Jeffrey C. Jackson (School of Computer Science, Carnegie-Mellon University, Pittsburgh, PA).

Another method of optimization, the branch-and-bound technique, was used in an ATC advisory system called COMPAS (ref. 6). Both optimization methods, CPS and branch-and-bound attempt to sequence incoming aircraft in such a manner as to minimize the total delay for all aircraft, and for both methods various restrictions apply in order to obtain feasible solutions.

Three of these methods of scheduling, FCFS, TA, and CPS have been implemented in a Traffic Management Advisor (TMA) Station, which is part of an automated system for the management of arrival traffic (ref. 1). The sequencer in this system permits the selective use of any combination of these scheduling schemes, and it contains other features that are important for the human interaction with the automated sequencer but which are not discussed in the present report.

Besides describing the scheduling methods, the purpose of this report is to statistically evaluate those that are implemented in the TMA. This is done using a large number of realistic traffic samples to determine their overall effect on aircraft delays. Additionally, the analysis is used to show the effects of other variables on delays

such as traffic distributions, lengths of traffic samples, and winds. Also, the initial analysis of the results for an optimal time-advance algorithm and an optimal single-position-shift CPS, both of which cannot be implemented in an operational ATC system, permitted the design of an improved TA algorithm and of a heuristic CPS that are being implemented in the CTAS system.

First, the three scheduling algorithms, FCFS, TA, and CPS will be discussed, wherein each successive algorithm improves on the preceding one by further reducing the average delay for all aircraft. Then, a model of incoming traffic to a hub airport for the purpose of evaluating the scheduling algorithms will be built. Finally, a sufficiently large number of randomly chosen traffic samples will be generated to obtain the statistical characteristics of the scheduling algorithm as a function of mix of aircraft and traffic density. The primary criterion of performance is average delay per aircraft. In addition, a few individual traffic samples will be examined to determine where scheduling algorithms may be simplified or improved.

In reference 4, CPS always followed TA, because CPS had a much smaller effect on delay reduction. There are reasons to believe that CPS should be applied first, however, followed by TA. CPS increases the lengths of the gaps compared to FCFS. These longer gaps may sometimes be taken advantage of in the fuel saving TA algorithm.

Scheduling Algorithms

In order to sequence aircraft for landing at an airport in an efficient manner one has to know the spacing requirements for different types of aircraft. Therefore these requirements will be discussed before going to the actual scheduling algorithms.

Spacing Requirements

Spacing requirements are an essential input for all types of scheduling algorithms. As stated earlier, two types of aircraft are being dealt with, heavy and large. For each type, the FAA specifies a spacing distance at landing that is dependent on the sequence heavy-heavy, heavy-large, large-large, large-heavy. The minimum spacing matrix is shown below (spacing distances are in nautical miles).

		Second to land	
		Large	Heavy
First to land	Large	3	3
	Heavy	5	4

The time separations are based on FAA specified spacing distances on final approach, and on the speed profile of each aircraft weight class (ref. 7). The time separations for no wind are shown below (times in seconds). These time separations include pads for navigation and piloting errors as well as other uncertainties.

		Second to land	
		Large	Heavy
First to land	Large	78	78
	Heavy	125	125

Under headwind conditions, when the trailing aircraft flies at constant airspeed independent of the wind, its ground speed is reduced by the speed of the wind. Thus for a specific spacing in miles, under headwind conditions, the time-separation matrix will have larger required separations. When the headwind is 20 knots, the following approximate values are obtained. (The exact values depend on the assumed indicated airspeed profile). Again, the time separations are given in seconds.

		Second to land	
		Large	Heavy
First to land	Large	91	91
	Heavy	145	122

The above time separations are those used in the present report for analytical purposes. In CTAS, situational dependent time separations to meet the minimum separation standards are computed for individual aircraft pairs via fast time integration.

Modified First-Come First-Served Algorithm

The simple FCFS algorithm determines the aircraft landing sequence based on the order of the sequence of estimated times of arrival (ETAs) at the runway, computed by the Center at the time the aircraft cross the Center's boundary.

The modified FCFS algorithm recognizes two scheduling horizons: an initial scheduling horizon and a final scheduling horizon. The initial scheduling horizon is a spatial horizon, which is the position at which each aircraft enters the Center's space. The final scheduling horizon, called the freeze horizon is defined by a specific time-to-landing (ETA), given no interference from other aircraft. Once an aircraft has penetrated the freeze horizon, its STA remains unchanged, independent of ETAs of

other aircraft subsequently entering into the scheduling interval.

The scheduling algorithm receives the data for each new aircraft as it passes the initial spatial scheduling horizon: present time at which the aircraft crosses the initial sequencing horizon, ETA, aircraft weight class, and aircraft identification. If the (temporal) freeze horizon is a shorter time interval than the shortest estimated flight time from the Center boundary to landing, the scheduling algorithm establishes the landing sequence in order of the computed ETA, which is called the FCFS order, and computes the associated sequenced time of arrival (STA) at the runway. If the freeze horizon is a larger time interval than the shortest estimated flight time from the Center boundary to landing, the situation is more complicated, as will be discussed next.

For aircraft that enter the scheduling horizon, the STAs are computed as follows. If no other previously sequenced aircraft's ETA is later than that of the newcomer's ETA, then the STAs of the earlier sequenced aircraft are not disturbed, and the newcomer is assigned a time equal to its ETA or the time that ensures the minimum time-separation required for the types of aircraft that are following each other, whichever is larger. If a new arrival's ETA falls ahead of the time slots reserved for previously sequenced aircraft, and if none of the already sequenced aircraft had its sequence frozen, then the new arrival is inserted ahead of these aircraft in the order of the ETA and at the proper spacing from the next earlier aircraft. All aircraft following the new arrival are respaced with the proper time separation. If frozen sequenced aircraft have STAs later than the new arrival's ETA, it is first checked if a sufficiently large gap exists such that the new aircraft can be sequenced ahead of the frozen aircraft without changing any other aircraft's position. If proper spacing cannot be maintained, the new aircraft is sequenced in front of the first aircraft whose schedule has not yet been frozen, and the non-frozen aircraft behind the newcomer are rescheduled. (If frozen aircraft are present, this is not strictly an FCFS scheduling, even though it is called that in this report.)

Aircraft arriving at the boundary of the scheduling horizon appear unevenly spaced. Therefore the FCFS algorithm creates groups of tightly sequenced aircraft with larger gaps between individual groups. With the FCFS algorithm, the first aircraft in a group requires no delay whereas succeeding aircraft, on the average, require increasingly larger delays.

Time-Advance

Heuristic time-advance algorithm— The TA method, called the negative delay effect in reference 2, operates on the first aircraft of each group, and does not change the existing order (e.g., FCFS). The first aircraft in a group is speeded up to arrive sooner than its nominal ETA, and all aircraft in the group following it will have their delays decreased by the same amount of time. This also reduces or closes the initial intergroup gap. Since speedup is costly, the first aircraft is speeded up only when at least the immediately following aircraft requires a delay, which is shortened because of the speedup of the first aircraft. In this statistical evaluation, there exists no program that calculates maximum, minimum, and nominal ETAs from aircraft, navigation, and weather data. In the absence of actual minimum ETA data, a maximum time-advance of 1 min is chosen for all aircraft. In the implemented scheduling system of CTAS, the time-advance for each leading aircraft in a group is based on a fraction of the calculated values of the available time-advance.

When the (temporal) freeze horizon has a smaller value than the time of the shortest flightpath, FCFS and TA applied to the incoming traffic result in the same overall aircraft order.

The fuel saving TA algorithm using a performance function— Although, the following algorithm is called fuel saving TA, fuel saving is only one of the objectives of the algorithm. Whenever possible, it also attempts to preserve the desired ETAs, which the DA specifies, and it attempts to reduce mean delays compared to FCFS scheduling while keeping the standard deviation to a minimum. Time-advance is performed in two steps. First, all aircraft are time-advanced from their FCFS time by an equal amount regardless of future benefit. Second, since time-advance is costly, all time-advances that are not required to improve the following traffic are removed or reduced.

The idealized fuel saving TA algorithm: Before discussing the realizable fuel saving TA algorithm, one pretends to have the complete traffic sample available. Later the realizable case is examined, where traffic data are available only in the scheduling window. This permits the study of performance loss due to a finite scheduling window. To answer the question how much time-advance to use as a function of the demand, one must model the cost. First, fuel use was considered as cost only, without adding savings due to flight time reduction, but this did not prove practical. Table 1 shows some examples of fuel expenditures and times for nominal and fast incoming flights from 250 n. mi. out. The TA capability is larger for aircraft that enter the Center airspace at slow speed, since they can speed up briefly before descent. Also, for higher entry speed the fuel cost per minute time-advance

increases. The fast rise in fuel expenditure if t_a is further increased is not seen on this chart, since the present Descent Advisor (DA) sets conservative maximum speed limits. Table 1 shows data for a specific large aircraft. Heavy aircraft use about twice the amount of fuel, even if that fact was taken into consideration, heavy aircraft would not be sequenced preferentially, since for TA the FCFS order remains preserved. Only the magnitude of t_a may be changed, which then would favor heavy aircraft since the total fuel tends to be minimized. Also, Table 1 shows that the fuel use of an aircraft depends on many other factors such as initial speed and altitude. Even if one could take these factors into account computationally, certain aircraft would get preferential treatment. To prevent favoring fuel savings for any aircraft, in this evaluation heavy and large aircraft were given the same performance function.

There is another factor that suggests not to use overall fuel minimization for the fleet directly. The desired ETA calculated by the DA is not the point of minimum fuel use for a given flight path. In an example (not shown), fuel use was plotted versus ETA for a 250 n. mi. approach from 33,000 ft altitude to a landing at Denver while staying on the nominal horizontal flight path. Minimum fuel use was almost at the the slowest approach. In fact, compared to that for the nominal ETA, the fuel use was 80 lbs or 2.4% less at a cost of an additional flight time of 2 minutes. The nominal ETA has been chosen for several good reasons. The chosen ETA considers the cost of both fuel and time, and it also tends to center the achievable

range of time control without having to go off the specified path. Therefore, a simplified cost model was used, which treats all aircraft equally, preserves the fact that TA is costly for the lead aircraft, and emphasizes the fact that it is desirable to achieve the nominal ETA whenever possible.

For the estimate of cost for various amounts of t_a , the following simplified model for the costs and savings was adopted where j denotes the j th aircraft in a sequence:

- 1) calculate the STA_j 's without TA (first-come-first-served)
- 2) subtract desired t_a from each STA to obtain the time-advance STAs.

$$STA_TA_j = STA_j - t_a$$

This means that the sequence of aircraft and their spacings remain exactly as before, except for a time shift of t_a minutes.

- 3) calculate the delay after TA for each aircraft :

$$\text{delay}_j = STA_TA_j - ETA_j$$

- 4) when the delay is positive, the aircraft is still delayed in spite of the t_a , but one assumes that the cost decreases proportional to the decrease of delay (t_a), and the incremental cost is:

Table 1. Cost of time advance and gain for reduced delay as calculated by the descent advisor 250 miles from touchdown. Maximize TA in cruise and descent

Initial alt (ft)	Initial speed (mach)	w_f (lb/min)	t_n (min)	f_n (lb)	t_f (min)	f_f (lb)	$t_n - t_f$ (min)	$f_n - f_f$ (lb)	$\frac{f_n - f_f}{t_n - t_f}$ (lb/min)
27000	0.65	114.8	39.6	3793	33.8	4071	5.75	278	48
27000	0.68	120.7	38.0	3759	34.4	4063	3.63	304	84
27000	0.72	129.0	36.6	3793	34.3	4034	2.30	241	104
32000	0.72	116.8	36.5	3358	34.2	3606	2.26	248	109
32000	0.76	124.9	35.8	3400	34.1	3575	1.71	176	102
32000	0.80	135.3	34.6	3469	34.1	3542	1.61	74	121
37000	0.75	119.7	36.6	3202	35.2	3244	1.31	42	32
37000	0.79	126.3	35.3	3185	35.1	3210	0.28	25	89

Fuel flow rate at initial altitude and speed = w_f .
 Time for nominal speed profile = t_n .
 Total fuel for nominal speed profile = f_n .
 Time for fast profile = t_f .
 Total fuel for fast profile = f_f .

$$\Delta\text{cost} = -t_a$$

This is reasonable, since one assumes that, at least for the larger FCFS delays, the reduced delay saves fuel by reduced path stretching in cruise, and, of course, it saves time proportional to t_a .

5) when the delay after TA is negative,

$$\text{delay} < 0$$

the aircraft is asked to arrive before its nominal ETA. Unless there was no delay in the FCFS sequence the t_a will first bring the original delay to zero with a savings in cost for the j th aircraft

$$\Delta\text{cost1} = -(STA_j - ETA_j)$$

Then the remaining time-advance called "actual t_a ," incurs a positive cost, called cost2. When the actual time-advance is equal or less than one minute, one assumes that the additional fuel cost is of similar magnitude as the fuel saved for reducing the cruise time by the same amount. But, above one minute actual t_a , fuel consumption rises sharply. Because of the absence in this simulation of a DA to compute a maximum available t_a , for this analysis a maximum amount of allowable actual t_a of 2 minutes was assumed to account for the maximum speed capability of the aircraft and to provide a reserve for control. Then

$$\Delta\text{cost2} = -\text{delay}; \quad \text{where } 0 < -\text{delay} < 1 \text{ minute};$$

$$\Delta\text{cost2} = \text{delay} * \text{delay}; \quad \text{where } 1 < -\text{delay} < 2 \text{ minutes};$$

In other words 2 minutes of negative delay costs four times as much as 1 minute of negative delay. The total change of cost compared to FCFS for time advancing one aircraft then is

$$\Delta\text{cost} = \Delta\text{cost1} + \Delta\text{cost2}$$

6) The total cost change from FCFS for an air traffic sample is the sum of the incremental costs for all aircraft

$$\Delta\text{total_cost} = \sum \Delta\text{cost}_j;$$

This is the incremental cost to minimize. From table 1, in the above calculations a unit of cost is about 100 lbs of fuel per minute, but can vary widely dependent on the initial conditions when entering the Center airspace.

After initially time advancing each aircraft in a traffic sample, one removes or reduces all t_a s not required for maintaining the TA sequence of the next aircraft exactly as it was. This is illustrated in figure 1. When initially all aircraft are time-advanced, it may happen that some aircraft have negative delays (time-advance) with their inherent fuel costs, when the immediately following aircraft does not benefit. Therefore, as shown in figure 1, for each previously sequenced aircraft using FCFS plus pure TA, one checks if it,

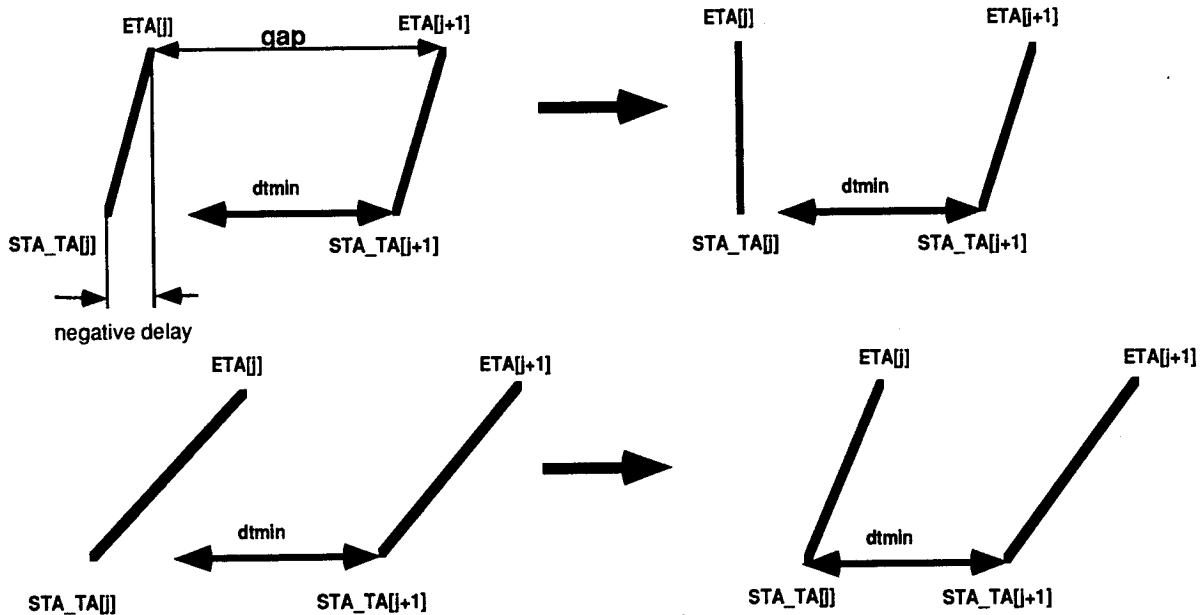


Figure 1. Part 2 of the fuel optimal TA algorithm. Removal or reduction of time advance if next aircraft does not require re-scheduling.

- 1) has negative delay, and if,
- 2) the next aircraft follows after a gap greater than the minimum time spacing. If both conditions are true, one then checks if the aircraft's ETA is earlier or later than the STA_TA minus the minimum time spacing to the following aircraft (dtmin). If the ETA is earlier, one reduces the time-advance for that aircraft to zero (top of fig. 1)

$$STA_TA_j = ETA_j;$$

If the ETA is later, one moves STA_TA_j as close to that of the next following aircraft as possible (bottom of fig. 1)

$$STA_TA_j = STA_TA_{j+1} - dtmin$$

where dtmin is a function of the aircraft types. In both cases actual time-advance is reduced, with accompanying reduction in cost. No other aircraft's delay is affected. With this method one often reduces the time-advance of the aircraft to such degree that a new gap opens for the next earlier aircraft. Therefore, *the second part of the algorithm must be employed successively from the last to the first aircraft. The last aircraft in the traffic sample is treated as a separate case, since it never needs to be time-advanced.*

The average minimum costs versus t_a for various demands were determined by running one thousand 1.5 hour length

samples for a range of t_a, 0 ≤ t_a ≤ 2 minutes in steps of 0.1 minute using both parts of the algorithm and the average cost vs. t_a was plotted in figure 2. This gives the result: low t_a is best for low demand; high t_a for high demand. The curves also show the relatively low sensitivity to choosing a less than minimum cost t_a. This low sensitivity is mostly due to the second part of the fuel saving TA algorithm.

So far the minimum cost average t_a for the full capacity runway was determined. When the acceptance rate of the runway is reduced, delays build up more quickly with time, larger groups of sequenced aircraft will occur, and a larger t_a will be effective in reducing delays and saving fuel. In other words t_a is a function of both demand and acceptance rate. To find the proper t_a when the acceptance rate is reduced is simply a matter of scaling; e. g. if the runway acceptance rate falls by a factor of 2 it is equivalent to the doubling of the demand. We can therefore write the following equation

$$\text{demand (for looking up } t_a) = \text{actual demand} / (\text{per unit reduction in acceptance rate})$$




Figure 2. Effect of time advance on extra fuel used or saved per aircraft.

In the determination of the average t_a s for minimum cost each complete sample was analyzed. The minimum cost for the best t_a for a sample is, of course, smaller than that which can be achieved when the average minimum cost t_a for a specific traffic density is used as determined for all samples with the same demand. Unfortunately, in a real traffic situation, one has only the traffic in the scheduling window to operate on and cannot examine the complete traffic sample to minimize cost. Sequences for earlier aircraft are frozen, and data for later aircraft are not yet available. However, as will be seen in the results section, the penalty for this limitation is not very large.

Pure TA: A simpler version of the fuel saving TA algorithm is pure TA; that is, all FCFS STAs are reduced by the same amount without removing unnecessary time advances. When one finds the optimal t_a s versus demand for pure TA they are in general lower than the t_a s for the earlier case. Comparing the optimum curves for pure TA (not shown) with figure 2, one finds that they do not differ much for large demands. The reason for this is that for large demands only rarely can one remove unnecessary time-advances. Because of the limited size of the scheduling window, and the fact that frozen aircraft cannot be resequenced automatically, it may be best to use pure TA for heavy traffic. The small increase in the average cost may compensate for the large increase in the cost for a specific traffic sample when a time-advance has been removed and is later found to be needed after all.

The realizable fuel saving TA algorithm: When an aircraft's ETA enters the scheduling window, the aircraft is initially FCFS sequenced and time-advanced by the average minimum cost t_a determined in the last section. Then, depending on the demand, one applies the second part of the algorithm for possible t_a reduction to all aircraft but the newly appeared one starting with the next to last aircraft and ending with the earliest aircraft in the scheduling window. If no new aircraft appears in the window for a time interval greater than the largest possible minimum spacing between two types of aircraft "dtmin," the time-advance for the last aircraft in the window is removed, and the second part of the algorithm is applied to the remaining aircraft in the window as before. The possible performance loss compared to the idealized TA algorithm stems from the fact that earlier frozen aircraft may have also benefited from the t_a reduction.

In an operational system, a traffic sample does not have a specific duration. Since, in the simulation, traffic samples of finite duration have been used, one must transfer this concept to the operational system. One could say that a traffic sample begins when the traffic is light enough that delays are not required and it ends when the traffic becomes so light that new traffic requires no delays.

If it were possible to examine the complete traffic sample, a t_a resulting in the lowest cost for that sample could be chosen. This is not possible. Instead a t_a based on an estimate of the demand must first be chosen. The algorithm must therefore look into the future, by obtaining data for aircraft in adjacent centers. Such data are available in the form of calculated ETAs to the coordination fixes just inside the Center. These ETAs are available about 40 minutes before an aircraft reaches the coordination fix. This is sufficient to estimate the future demand as a function of time, which then determines the present t_a .

It is important to switch to the appropriate magnitude of the t_a early based on both demand and acceptance rate. Once one has begun scheduling a tightly spaced group of aircraft, one cannot increase t_a until another gap occurs. Even then, the gap might be too small to apply the full desired t_a . From the earlier discussion, the TA algorithm usually saves fuel even if it is applied with an incorrect t_a . It is therefore recommended to apply a larger t_a quite liberally even when present traffic is light, but when flight-plan data predicts heavy traffic ahead. The fuel saving TA algorithm will still remove unnecessary time-advances, even though the t_a chosen is not the minimum cost one for the sample.

From the foregoing discussion it is proposed not to use TA for light traffic, to use the two step fuel saving TA for medium dense traffic, and to use pure TA for heavy traffic. From the data for the two types of TA algorithm one approximates the t_a s vs demand as linear functions, where $t_a = 1$ means maximum allowable t_a . Then for demands in number of aircraft per hour

demand < 20	$t_a = 0$
20 < demand < 30	$t_a = 0.35 + 0.0325 (\text{demand} - 20)$, remove unnecessary t_a .
30 < demand < 40	$t_a = 0.5 + 0.04 (\text{demand} - 30)$, use pure t_a .
demand > 40	$t_a = 1.0$

Constrained Position-Shift Algorithm

Optimal CPS- As previously mentioned, the CPS method reorders the existing FCFS order by taking advantage of different spacing requirements for different aircraft classes. Reordering makes sense only within a group. It is theoretically most effective when the groups are long (heavy traffic). Two aircraft are considered for reordering by a single position, provided that they arrive at the airport from different directions. This prevents possible overtakes within a sector. An optimal single-position-shift algorithm was developed by Jackson (unpublished) and is described in appendix II. A necessary restriction is that

none of the aircraft in the group can be given a time-advance of more than 1 min. As the algorithm is written, this restriction can only be tested after the algorithm has proposed position switches. Various non-optimal methods of removing the violations to the restrictions make the overall algorithm non-optimal. Because of computational constraints on the scheduling algorithms operating in real-time, and because of the physical constraints of time-advancing aircraft, the use of larger position shifts is not practical. Similar to the TA method, in a more complete simulation, one could use a time-advance equal to a percentage of the available time-advance for the given aircraft.

Heuristic CPS— The optimal CPS is based on the dynamic programming principle, and the solution of all position shifts is found only at the end of each group of aircraft. In an actual system, a scheduling window exists with a time-interval often much shorter than that of a group of aircraft, especially for heavy traffic, when CPS would be useful. As discussed in the results section, optimal CPS switches aircraft positions in such a way as to group heavy aircraft when possible. The following groupings are searched for and position shifted by the heuristic CPS, where “H” means heavy and “L” means large aircraft. Groups of size 5 where one HL or LH switch suffices: LHLHH → LLHHH, HHLHL → HHHLL, LHLHL → LHLL. The sequence HHLHH is not switched. In addition, two longer patterns are considered if a sufficient number of aircraft are in the scheduling window: a pattern of size 6, LHLHL → LLHHLL, which requires one HL and one LH switch, and a pattern of size 7, which also requires one LH and one HL switch: LHLHLHL → LLHHHLL. Two additional patterns, those that require a two-position shift of a heavy aircraft, are also checked for: LHLLHH → LLLHHH and HHLHL → HHHLL. This makes it possible that under certain conditions the heuristic CPS may work better than the single-position-shift optimal CPS.

For all these patterns, the earliest aircraft can be below the freeze horizon, since it is never involved in a switch; nevertheless, it determines which pairs of aircraft are switched. One additional pattern is searched for at the end of each group of aircraft LHL → LLH. This switch is independent whether L or H follow the gap, and it increases the length of the gap. Before switching, the old order of aircraft types and STAs is preserved. After the new STAs have been calculated, using the minimum time-separation matrix shown previously, one checks to verify that all delays, $STA_j - ETA_j$, are less than -1 min (time-advance of less than 1 min). If this is not the case, the old unswitched values are restored. This automatically takes care of larger gaps in the original aircraft sequence. The longer patterns are checked by using the original sequence

of aircraft types, and the possible changes from matching a shorter pattern are again changed to match the longer pattern, provided all aircraft meet the delay criteria. As desired, attempted resequencing over a large gap will cause the delay criteria to fail. No performance criterion needs to be calculated, since each switch guarantees some reduction in average aircraft delay of the traffic sample.

Scheduling Algorithm for Two Dependent Runways

When two runways are available for landing, a limited further chance for optimization occurs. Here only the most restrictive case is treated: Aircraft arriving from the left land on the left runway, and aircraft arriving from the right land on the right runway and no switches between runways will be permitted. As an example Kennedy airport has two runways that are separated by one mile. Therefore, the traffic from the left and the right can be landed only semi independently. The following stagger rule applies: *Aircraft landing on different runways must be separated by at least 2 n. mi. longitudinally on final independent of aircraft type.* This means that, independent of the algorithm used for scheduling, the spacing of aircraft on final to one runway must be the larger distance of the two “minimum spacing required between two aircraft on final, (see section on “spacing requirements”)” and “two miles longitudinal spacing between aircraft on different runways.”

Each time a new aircraft appears in the scheduling window for either runway, based on the two separate FCFS orders, the algorithm calculates for each runway separately the scheduled landing time for the latest aircraft in the scheduling window, while using the dependent runway rules. Then the algorithm selects the aircraft scheduled with the earlier landing time while discarding the other. One may call this last step a one step optimization.

Traffic Model

Toward an Accurate Traffic Model

To evaluate the scheduling algorithms for a particular situation, one needs an accurate traffic model. Such a model might be based on scheduled airline arrivals for specific days of the week, including planned routes and aircraft types. Such lists are available at ARTCCs and are presently used by the traffic managers to predict peak traffic times. However, these lists are not accurate enough to predict aircraft arrivals at the Center boundary, because, as will be explained next, scheduled traffic must be considerably modified before it crosses the Center boundary.

All traffic samples discussed herein are based on arrivals at Denver Center. Arrival traffic scheduled by the airlines for Denver for a particular date and time-interval is illustrated in figures 3 and 4. It can be seen from these figures that the incoming traffic was heavily concentrated in the 30-min period from 7:45 A.M. to 8:15 A.M. (local time), and that almost all traffic was from the NE and SE. In fact, there were 56 aircraft scheduled in a 33-min time interval. Such peaks are somewhat flattened out by the natural statistical blurring owing to random delays in takeoffs, errors caused by winds, and flight technical factors. The flattening process is further aided by deliberate changes, such as ground holding and increases in in-trail spacing. To get a more precise model of the traffic, one would have to collect data for many days on aircraft crossing the Center boundary, along with aircraft type and planned route. Such data are difficult to obtain. Therefore, a somewhat less detailed model is used, which is based on gross traffic statistics.

Traffic Model for Studying Scheduler Effects

The purpose of this work is to describe a statistically accurate traffic model typical of peak hours at Denver, which was used to investigate different scheduling algorithms. When many traffic samples are analyzed, the

model provides a good insight into traffic problems resulting from the random nature of traffic arriving at the Center boundaries, even though the traffic scheduled by the airlines may be almost identical for many days. The aircraft arrival rates, in-trail distances, and their statistical variations are realistic for each jet route. These arrival rates may be changed, depending on what time of day is simulated. Also, traffic from one direction can be made heavier than that from the other direction. Moreover, the model assumes that the incoming traffic on different jet routes is not coordinated for conflict avoidance at the various route junctions or at touchdown. Coordination and conflict resolution have to be accomplished in the Center sectors (with the help of the sequencer) and finally in the TRACON area.

Jet traffic arrives in Denver Center's northwest arrival sectors along four routes, and in the northeast arrival sectors along three routes (fig. 5). The northwest traffic is handed to the TRACON through the Drako feeder gate, and the northeast traffic through the Keann feeder gate. Incoming traffic from the lower half-plane is not simulated, since it is landing on a separate runway during VFR operation.

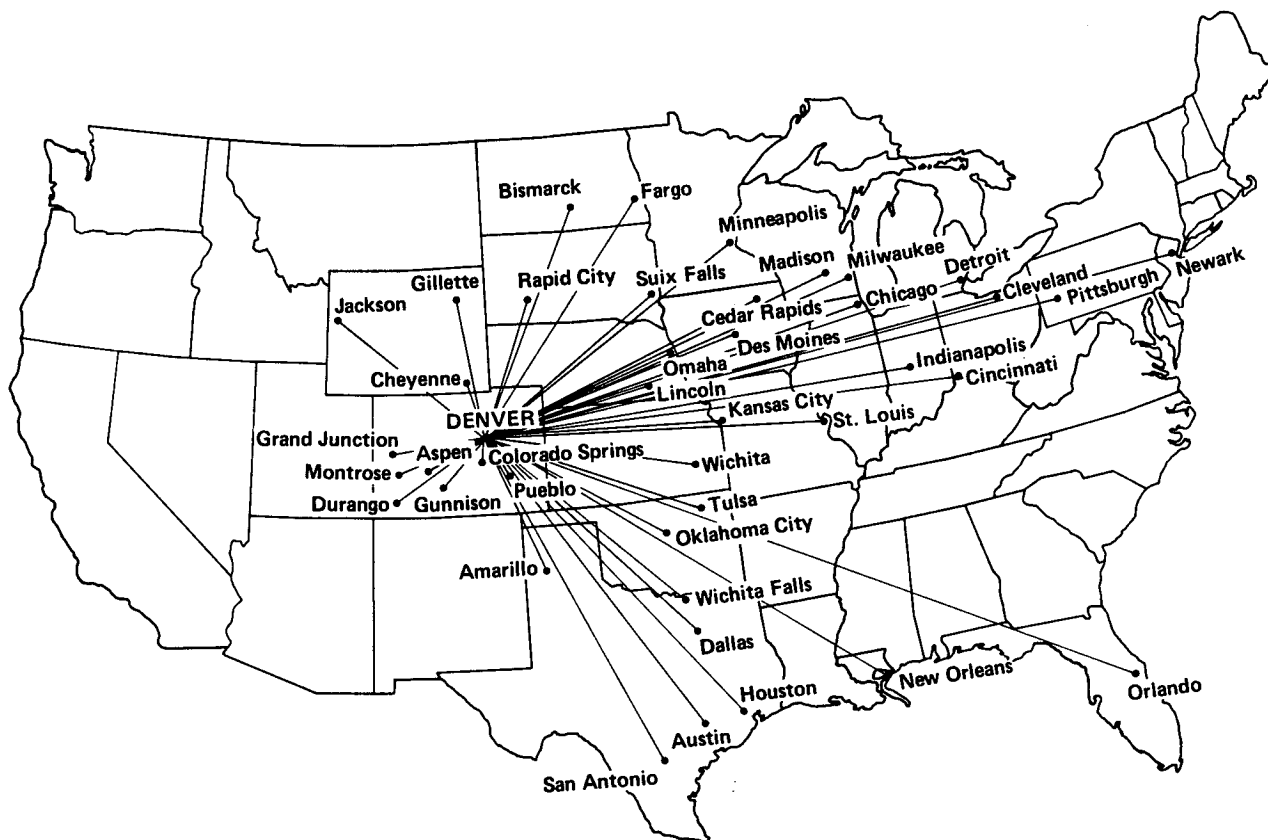


Figure 3. Denver traffic: 1 March 1987, 7:00-8:35 a.m.




Figure 4. Scheduled Denver arrival traffic in 1-min intervals.

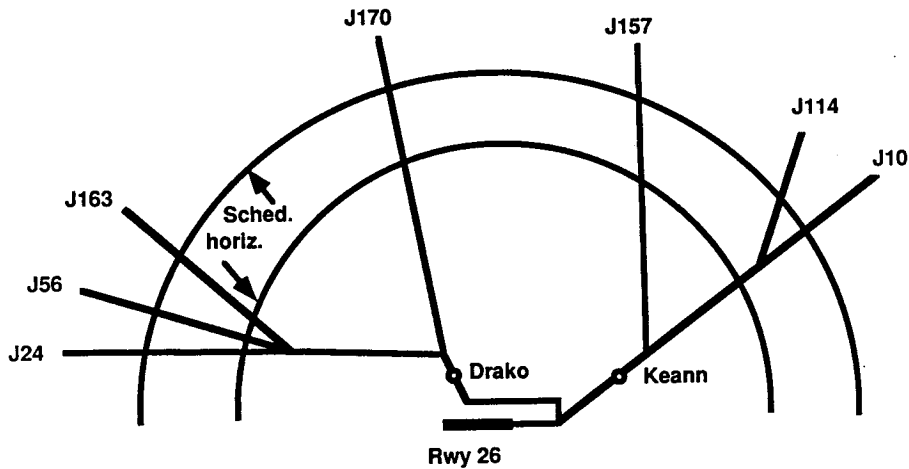


Figure 5. NW and NE jet routes into Denver.

One of the traffic directions usually carries high-density traffic and the other direction usually carries low-density traffic. In the simulation, the high-density direction carries about 70% of the traffic. Any other ratio of high-density-to-low-density traffic can be chosen. For this simulation, on the average, each of the four routes in the Drako area carries 25% of the NW traffic, and each of the three routes in the Keann area carries one third of the NE

traffic. Also, on the average, of all aircraft arriving, unless otherwise specified, 30% of all traffic is heavy jets, 70% is large jets. Presently only two types of aircraft are dealt with, heavy and large. At Denver, small aircraft usually land on a different runway. These assumptions will sometimes be varied to observe the effects on the delay statistics.

Actual route-traversal times within the Center boundary are shown in table 2. These times vary considerably and make it difficult to develop a sequence that remains fixed in time, one that a controller can use. Hence, the times for each route were approximately equalized (the "total" columns in the table), which is equivalent in a real system to extending the shorter routes, J170, J10, and J157 into the adjacent Centers. The total route-traversal times are not arbitrarily made equal, since in a real system route-traversal times vary as a function of aircraft types and winds, and the sequencer must be able to handle routes of various lengths. This simulation is limited to constant route-traversal times for each route, independent of the type of aircraft, thus avoiding the study of possible conflicts on the same route, when a faster aircraft may pass a slower one.

Table 2. Route traversal times

Jet route no.	Within Center boundaries, min	Total, min	
1	J163	42.30	} Through Drako
2	J156	45.45	
6	J170	28.50	
7	J24	47.78	
8	J114	41.43	} Through Keann
9	J10	25.10	
10	J157	34.11	

Given the chosen statistical traffic parameters, such as the landing rate and the sample time-interval, the start times and routes for exactly M aircraft are generated uniformly for the time-interval specified, where $M = (\text{landing rate} * \text{time interval})$, and is rounded to an integer. The aircraft arrival rate for each route is chosen based on the traffic load at Denver. Since time en route varies between 42.3 and 47.78 min for different routes, the nominal landing times have been rearranged so that the rectangular distribution of ETAs for each route centers on one-half of the time-interval specified. This results in an overall nonuniform distribution for start times [ETA - Route_traversal_time], where only the first and last few minutes are affected. The uniform distributions of start times for individual routes sometimes violate the minimum spacing standard of 3 min on each route. (For the newer results with the fuel saving TA algorithm, the minimum spacing is reduced to 2 min, which is thought to be more realistic). Hence, the times at which the aircraft cross the original Center boundary for each route are modified iteratively, starting from the earliest aircraft, by shifting each aircraft that violates the specified minimum in-trail spacing time-interval to a later time until all air-

craft meet the specified in-trail spacing. This often generates several equally spaced aircraft, especially in heavy traffic, which duplicates real traffic situations. Also, this modification sometimes makes the traffic sample longer than the specified interval, especially when large in-trail spacing time-intervals are specified. Again, this is thought to be realistic, since a scheduling time-interval for a fixed number of aircraft must get stretched out, as the example in figures 3 and 4 showed.

To study the sensitivity of calculated delays as a function of the distribution of aircraft in the specified time-interval, the choice of a triangular distribution of aircraft for each jet route is also provided.

Results

First, individual traffic samples will be discussed to give a clear picture of the generation of traffic samples and their statistical character, as well as to demonstrate the effect that scheduling has on delays. Second, statistical results will be discussed in terms of cumulative probability distributions.

Time Diagrams of Traffic Samples and Associated Delays

In this section a variety of traffic samples are presented. Since traffic samples in tabular form are hard to grasp, a graphical presentation has been developed. The graphical presentation of the sample affords a quick way of understanding the interrelations of the various time-ordered lists and of grasping causes of delays, as well as suggesting some remedies. Using an example, first the time diagram will be described in detail and then various traffic samples will be presented, limiting the discussion to major points. Figures 6(a) and 6(b) show a theoretical traffic sample at Denver (arrival rate of 25 aircraft/hr from the NE and NW). Given a set of parameters, such as arrival rate (demand), sample time-interval, percent of total traffic on each route, minimum in-trail spacing, and freeze horizon, each random-number generator seed defines one traffic sample. Knowing this seed, one can examine in detail unusual traffic sequences detected during the statistical runs.

In all the traffic samples shown, two thirds of the traffic is through Drako and one third through Keann. The mix of large-to-heavy aircraft is 70% to 30%. The traffic sample time-intervals are 1.5 hr, with no traffic before or after. The minimum in-trail spacing in the Center is 3 min, which often results in several aircraft on a route exactly 3 min apart.

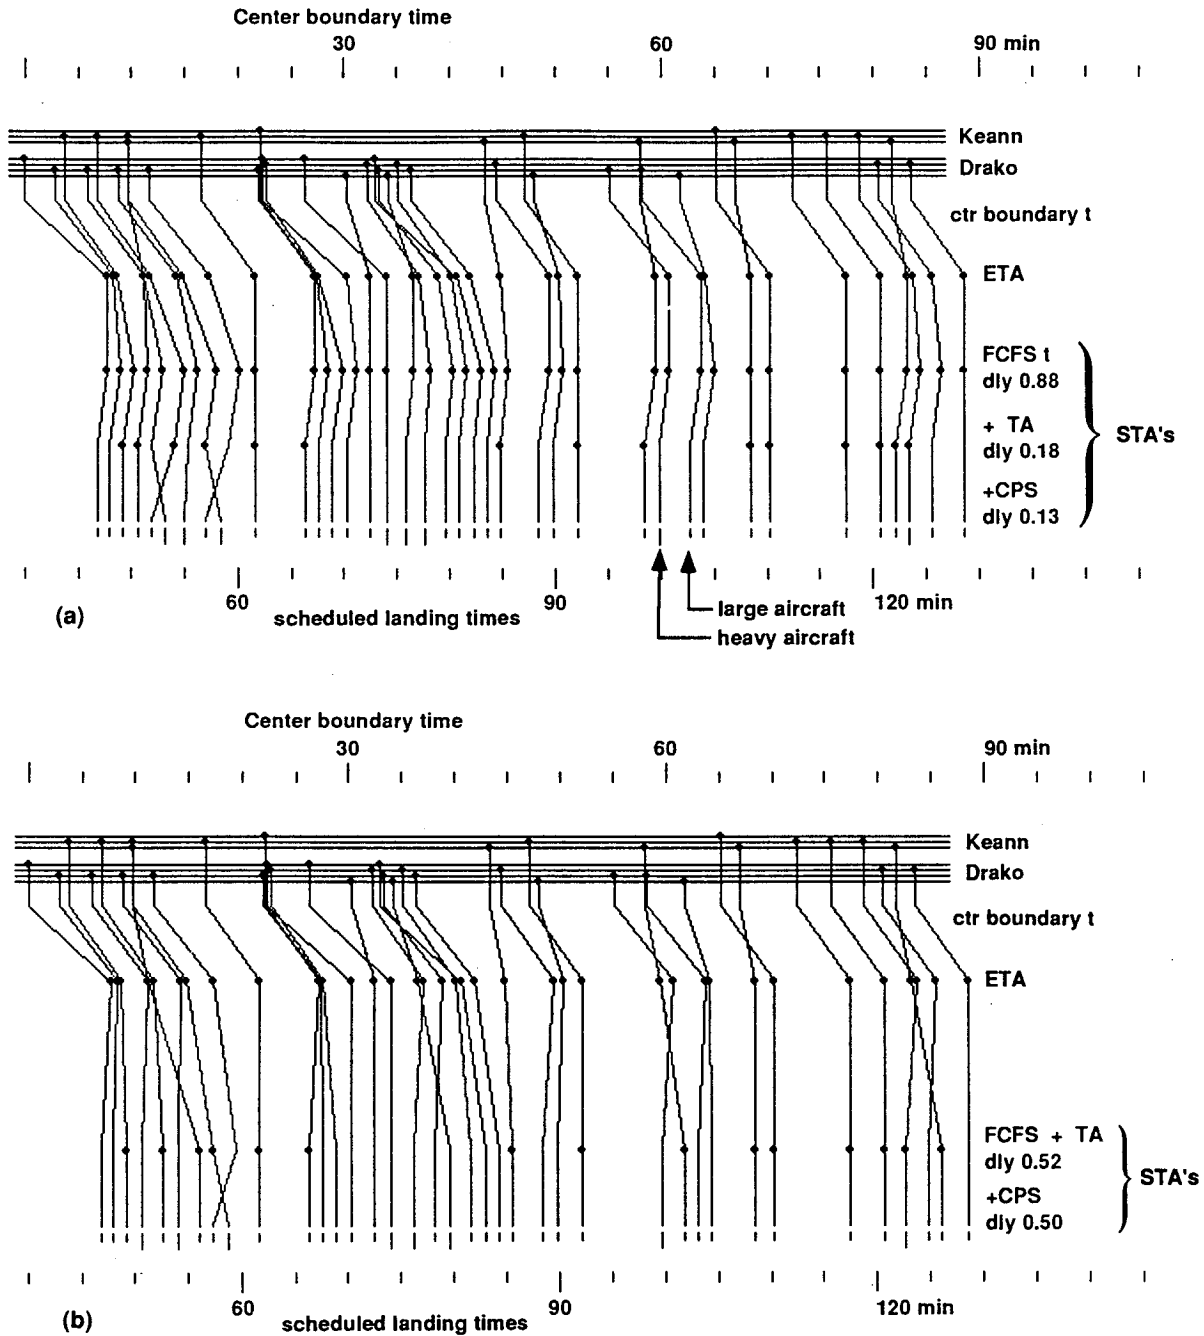


Figure 6. Representation of scheduled traffic sample. (a) Low freeze horizon. (b) Large freeze horizon

The closely spaced top horizontal lines in figures 6(a) and 6(b) are time lines for each jet route. They are from top to bottom, J157, J10, J114, J24, J170, J56, and J163 (see fig. 5). The dots on each horizontal time line show when an aircraft is crossing the Center boundary on a given jet route. The time-scale for these time-lines is given above the lines. The time-scale for the ETAs and STAs has been

shifted by a constant amount (40 min) to make the figure more compact. This scale is shown below the graph.

Each downward slanting line is called a scheduling line for one aircraft. The vertical top portion of each scheduling line begins at the appropriate jet route time-line and ends at an imaginary horizontal line, the Center boundary arrival-time-line. A slanted straight line

connects the vertical line's lower end to the ETA. This time represents the time the aircraft would arrive at the runway, if there was no interference from any other aircraft or from unknown navigation errors and environmental conditions. The sequence of all ETAs determines the FCFS order to be preserved (at least approximately) for fair scheduling.

Any two lines that cross between the Center boundary arrival time and the ETA belong to two aircraft on different routes, where the aircraft on the shorter route is arriving later at the boundary, but whose ETA is earlier than that of the other aircraft.

The horizontal component of the line between ETA and FCFS STA in figure 6(a) or the ETA and FCFS + TA STA in figure 6(b) represent the sequenced delay to meet spacing requirements. If the scheduling line is vertical, no delay is required for the particular aircraft. The more delay is required the slant of the line. If none of the lines intersect, as in figure 6(a), the FCFS order has been preserved, which is the case when the scheduling horizon is selected below the time for the shortest route. The average delay per aircraft in minutes is shown for each scheduling method; for example, under FCFS the average delay is 0.88 min, and further sequence optimization reduces the average delay.

In figure 6(b), for the same arrival data, the scheduling freeze horizon was deliberately chosen larger than the time it takes to fly most routes (45 min), hence, lines between ETAs and (FCFS + TA) STA sometimes intersect, showing that the FCFS order has been altered. Since FCFS and TA are not separable in their effects (the FCFS order is not preserved), only the joint sequence is shown (i.e., FCFS + TA). Scheduling around frozen aircraft often has the effect of increasing the total delay for the traffic sample when compared with strict FCFS scheduling, as demonstrated by comparing the average delays in figures 6(a) and 6(b), where the FCFS + TA average delay increased from 0.18 to 0.52 min. (In a few samples of the several thousand analyzed, this trend was reversed in cases when changing the FCFS order mimics an intelligent CPS.)

The straight line between FCFS time and TA time in figure 6(a) shows the effect of time-advance. An aircraft that had zero FCFS delay is a candidate for time-advance, provided that it is the leader of a group of at least two aircraft (heuristic TA algorithm). The aircraft is speeded up by 1 min or until the gap to the preceding aircraft is reduced to the minimum allowable, whichever is the smaller time-advance. Commercial jet aircraft have only limited capability of speeding up in the descent phase, and a maximum of 1-2 min time-advance is thought to be typical, where the larger time advance has disproportion-

ally larger fuel costs. The leading aircraft incurs a fuel cost flying above its preferred speed. All other aircraft in the group that are not speeded up beyond their ETA will benefit by having their FCFS delay reduced by the amount of time-advance of the leading aircraft. For time-advance, none of the aircraft scheduling lines cross, and the previous order is preserved.

The final portion of the aircraft scheduling line shows the absence or presence of CPS. Since only a single position shift was allowed, only adjacent lines cross (see figs. 6(a) and 6(b)). Only the scheduling lines for aircraft going through Keann have a dot on the FCFS + TA line to indicate whether position switching is considered for two aircraft from the same direction NE (Keann) or NW (Drako), with a resulting overtake condition, which would add to controller workload. Notice that in this example, for each constrained position shift, one aircraft arrives from the NW, the other from NE. Thus, possible overtakes are prevented.

The short vertical lines underneath each aircraft time line indicate the type of aircraft, a longer line for heavy aircraft and a shorter line for large aircraft. Where the traffic is tightly grouped, it can be noted that the separations differ, depending on the successions of types of aircraft discussed earlier.

The total number of time-advance commands to aircraft goes up for smaller numbers of aircraft per hour, because the groups of aircraft for a 1.5-hr traffic sample become shorter and more numerous, and each leading aircraft of a group must be time-advanced. The sum of the time advances of the aircraft with negative delays is the cost of time advancing a sample of aircraft. This is illustrated for the heuristic TA for four traffic samples in figures 7(a) to 7(d). The traffic samples are chosen for light traffic (25 aircraft/hr) and heavy traffic (40 aircraft/hr), one sample each with relatively low average delay and the other sample with exceptionally large delay. Figure 7 shows what causes relatively small and large delays. Small average delays occur when the ETAs are uniformly spread over the time-interval considered and are without large gaps, and large average delays occur when the opposite is true. One can see that for low-density traffic or well-spread traffic, TA should not be used, since delay is small already, and the cost in time-advance for the modest delay reduction is high, 12.16 min in figure 7(a) and 14.49 min in figure 7(c). That is, many aircraft had to be speeded up to reduce the delays for the remaining aircraft. There are many short groups, and many aircraft would have to fly faster than their preferred speed profiles. On the other hand, the cost in time-advance for heavy or grouped traffic is relatively small, 2.55 min in figure 7(b) and 2.22 min in figure 7(d), since only three aircraft


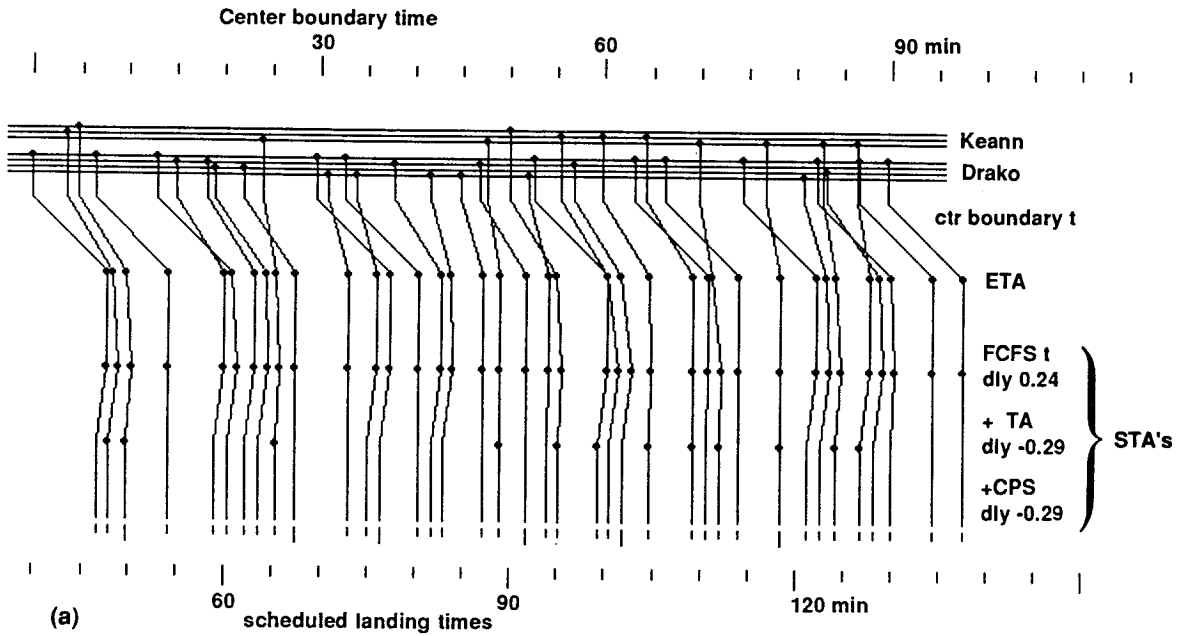


Figure 7. Traffic sample with heuristic time advance. (a) Light traffic with well-spread aircraft. (b) Light traffic with high delay.





Figure 7. Concluded. (c) Heavy traffic, well spread out. (d) Heavy traffic with large delay.




Figure 8. Comparison between scheduling methods.

needed to be speeded up in both cases. Figures 7(b) and 7(c) also show the modest improvement that can be achieved when CPS is added to TA. For figures 7(a) and 7(d), CPS found no position shift that gave reduced delays. It is difficult to determine a break-even point for heuristic TA versus no TA, since both time and fuel are involved either as savings or as cost for all aircraft whose sequences are affected.

The two time-advance methods, heuristic and fuel saving, are compared for one traffic sample by means of traffic diagrams. In the top part of figure 8 the fuel saving TA algorithm is compared with FCFS. In the bottom part of figure 8 the same traffic sample is sequenced with the heuristic TA algorithm, where at least one aircraft following must have its delay reduced before a lead aircraft is time-advanced. The resulting sequence is similar, but not quite as fuel efficient.

Figures 9(a) and 9(b) show a traffic sample in which CPS is applied with and without permitting overtakes. In this example, two additional heavy aircraft could be grouped together with overtakes permitted, resulting in a reduction of the average delay per aircraft from 2.87 min to 2.72 min.

It was shown in the scheduling algorithm section that a 20-knot headwind upon landing increases the required time-separations. A traffic sample illustrates this in figure 10, for FCFS only, for both no wind and for a 20-knot headwind. In this example, for an identical sequence of ETAs, the average delay for FCFS scheduling is increased from 2.31 to 4.05 min. Therefore, winds can play a major role in causing delays.

Figure 11 shows parts of the traffic-sample diagrams having to do with CPS only. CPS tries to reduce the length of a group of aircraft, which reduces the average delay of all aircraft. The cost of such delay reduction is the fuel cost for those aircraft that have to be time-advanced beyond their ETA. Therefore, CPS shows the most benefit in reduced average delay when the position switching is done early in a large group, thus reducing the time delay for all following aircraft in that group. Switching at the end of a group is of little benefit in reducing the average time delay (top example of fig. 11), but controllers prefer to place a heavy aircraft at the end of a group. The remainder of figure 11 shows how CPS groups the heavy aircraft together by either time-advancing or by delaying the heavy aircraft. In this





Figure 9. Sample of heavy traffic with CPS. (a) Overtakes not permitted. (b) Overtakes permitted.

manner, groups of two, three, or four heavies are formed. Figure 11 shows only the reduction in delay: the group becomes shorter. The cost of such position switching depends on data not shown here; namely, whether the aircraft that are switched toward an earlier arrival time simply have their delay reduced, or if they have to speed up to arrive earlier than their desired time of arrival.

The optimal CPS algorithm assumes that a tightly sequenced group of aircraft can be reordered such that it is again tightly sequenced with no more than the minimum required gaps. After calculating the new order of aircraft, one finds that sometimes the delay of some aircraft in the group decreases so that they have a negative delay of less than 1 min. In such cases, two alternative choices were




Figure 10. Scheduled traffic sample. (a) Without wind. (b) With 20-knot headwind.

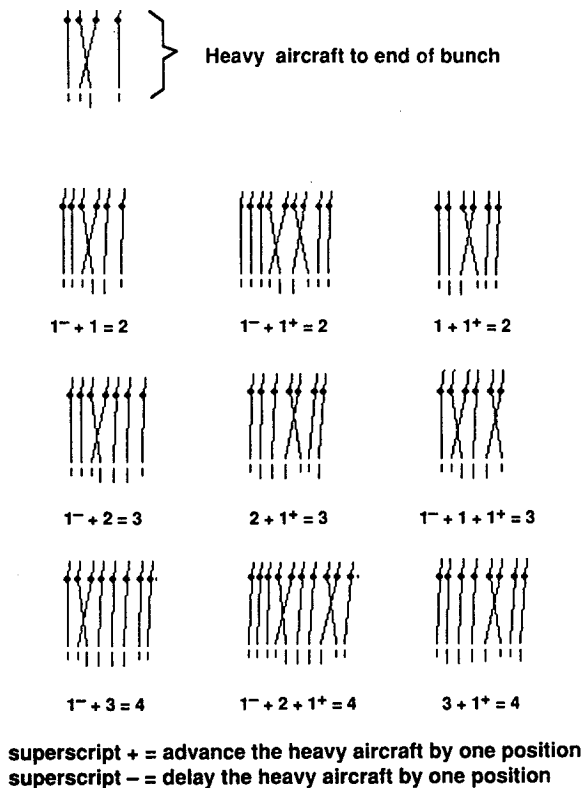


Figure 11. Constrained position-shift examples.

made to meet the restriction of a maximum negative delay of 1 minute (see the captions of figs. 12(a) and 12(b)). Either choice satisfies the restrictions at only a small loss of optimality when many samples are considered. Comparing the total delays for all aircraft in the sample of figures 12(a) and 12(b) with those of 12(d) and 12(e), one sees that there is no clear choice of method for meeting the maximum negative-delay restriction. In 12(a) and 12(b) the latter is better, and in 12(d) and 12(e) the earlier is better. To build this restriction into the algorithm directly would unnecessarily increase its complexity. This is not warranted, since the algorithm, as it stands, is not useful for an operational system which has a finite scheduling window. Another minor improvement to the optimal CPS algorithm was made by deleting position switches only after an unacceptable negative delay was detected in a group of aircraft, and by retaining the earlier switches. The optimal algorithm was mainly used to get an upper bound on the performance of a heuristic algorithm, which has been derived from the insights gained by observing the performance of the optimal algorithm. In figures 12(a)-12(f), various equivalent sections of traffic have been marked by double-headed arrows of equal lengths. The arrow in figure 12(a) shows that although different switches have been made by the




Figure 12. Comparison of two versions of post-processing of the optimal CPS and the heuristic CPS. (a) Optimal CPS, remove all position switches in a group if one negative delay <-1 min exists: SEED 109. (b) Optimal CPS, keep all position switches but reduce negative delay to <-1 min: SEED 109. (c) Heuristic CPS, locally proposed switches are not made if this would result in negative delay <-1 min exists: SEED 109.

heuristic and the optimal algorithms, the section containing the same aircraft (a1) is only slightly shorter for the optimal algorithm. The arrows in figure 12(b) show that the optimal CPS unnecessarily lengthened the sequence by one slot; figure 12(b) still has the overall shortest delay, owing to many earlier switches in the same group of aircraft. The arrows in figure 12(c) show that for the same algorithm, the two unnecessary switches in a group of aircraft increased the delay for six of the nine aircraft, but the overall delay for all aircraft is only 8.5 min longer.

Analysis of Traffic Including Both Modes of Optimization

In interpreting the following data, one must remember that the model that is being used for traffic-sample generation assumes that there is a rectangular probability distribution

for arrival times at the Center boundary and that there is no traffic outside the interval under consideration, except where the 2 or 3-min minimum spacing requirements necessitated pushing some traffic beyond the maximum time. Almost certainly, the actual arrival-time distributions at the Center boundary are not completely rectangular, which would further modify the cumulative distributions. This means that the data given in this report are meant to show trends rather than precise values. The curves shown in figures 13-18 are approximations of the cumulative probability of the average time-delay per aircraft for a random traffic sample being equal to or less than the value given on the abscissa, with traffic density (demand in aircraft per hour) as parameter, where the average time delay per aircraft for a random traffic sample is defined as the sum of the individual aircraft delays divided by the number of aircraft in the sample. All

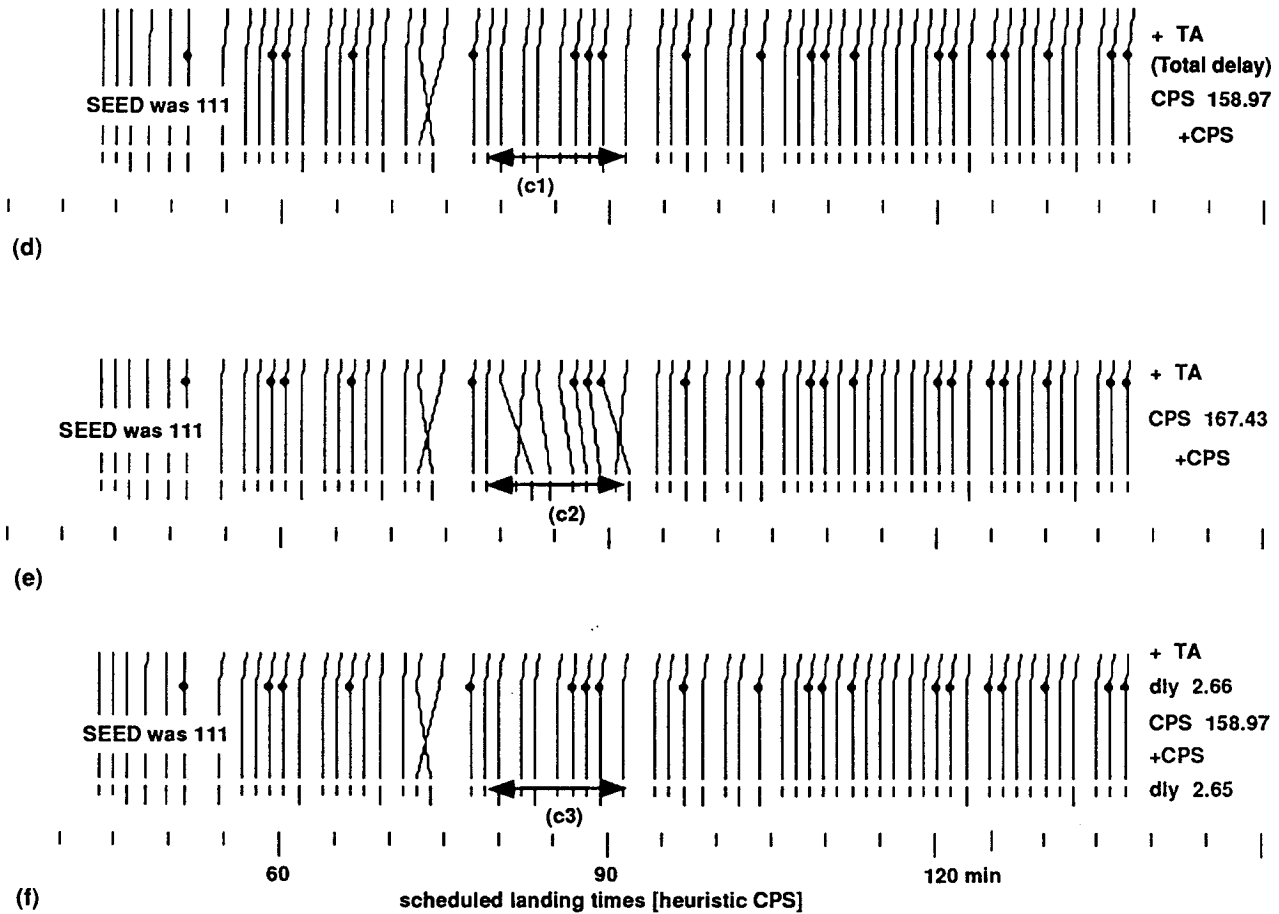


Figure 12. Concluded. (d) Optimal CPS, removal all position switches in a group if one negative delay < -1 min exists: SEED 111. (e) Optimal CPS, keep all position switches but reduce negative delay to < -1 min: SEED 111. (f) Heuristic CPS, locally proposed switches are not made if this would result in negative delay < -1 min exists: SEED 111.

cumulative distributions are based on 2500 traffic samples each, and data points are shown individually as dots to give an indication of the statistical noise in the data. The cumulative distributions are presented rather than parameters such as expected value and standard deviation, since the distributions are neither Gaussian nor any other common distribution. For figures 13 to 18 the heuristic TA method was used. The trends shown would be similar for the improved TA method.

Figure 13 shows the cumulative probability distributions for the average delay per aircraft in a given traffic sample, with the parameter N , the traffic density or demand in number of aircraft per hour. The traffic mix (traffic from NW and NE) and the aircraft mix (heavies vs large) have been chosen such that it should show the greatest benefit for CPS optimization, namely both 50%/50%. Figure 14 shows data similar to those in figure 13, but for the traffic and aircraft mix chosen for most of this simulation, which is described in the Traffic Model section. As an example,

if one studies the $N = 45$ curves in figure 14, the benefits of TA and TA + CPS can be readily seen. For FCFS scheduling, an average delay of 8 min or less is realized for 46% of the traffic samples. With the addition of TA, the same average delay per aircraft or less is realized for 58% of all traffic samples. With the further addition of CPS, this delay, or less, occurs 64% of the time. In the remaining cumulative distribution figures, the groups of curves representing FCFS, FCFS + TA, and FCFS + TA + CPS are not always labeled separately, since they are always in the same order.

By looking at the complete cumulative distribution curves, the TA curves are moved to the left of the FCFS curves by somewhat less than 1 min, as was expected, since that was the assumed maximum time-advance for each aircraft. In actual traffic, the allowable time-advance for a given aircraft depends on the type of aircraft, the aircraft state, and the proposed path. This may be somewhat more than 1 min on the average. In both figures 13




Figure 13. Cumulative probability distributions for traffic and aircraft type mix where CPS has best performance (traffic NE/NW and heavies/large both = 50%/50% for 1.5-hr traffic samples, using heuristic TA).




Figure 14. Cumulative probability distributions for nominal traffic and aircraft type mix NE/NW traffic 66.66%/33.33%, large/heavy = 70%/30% for 1.5-hr traffic samples.




Figure 15. Comparison of cumulative distributions for different aircraft type mixes, 1.5-hr samples, 40 aircraft/hr arrival rate.




Figure 16. Cumulative distribution for nominal traffic, 40 aircraft/hr including CPS with overtake. (a) Optimal CPS performance.




Figure 16. Concluded. (b) Showing slight decrease in performance for the heuristics CPS vs optimal CPS. Insert shows performance of the heuristic CPS as function of scheduling window size in minutes.




Figure 17. Effect of length of the traffic sample with otherwise same statistical parameters as shown in figure 16.

and 14, comparing the reduction of the average time-delay when CPS is added, one notices that CPS is more effective for greater traffic densities, which is fortunate. This is so, because longer groups occur in heavy traffic, and long groups can be optimized more effectively than short ones. However, compared with TA, the benefit of CPS is relatively small. Even in the best case, the delay reduction is less than 0.5 min per aircraft. In this simulation, CPS was calculated only once for each traffic sample by dividing it into groups of aircraft and applying CPS to each separate group. In an actual system, the STA calculations would have to be started for each aircraft as it arrives at the Center boundary and finished as it passes the freeze horizon. Since the present CPS algorithm is an example of the dynamic programming principle, the algorithm determines the final sequence only after the last aircraft of each group has passed the Center boundary. Making earlier decisions on position switching will cause some loss in performance.

In figure 15 data from figures 13 and 14 are combined to compare different aircraft mixes for the same arrival rates. The larger number of heavies in the 50% heavy/50% large aircraft mix curves require more spacing and therefore have more delay. However, CPS is more effective in this




Figure 18. Effect of changing freeze horizon.

case, since more switching opportunities exist. Since the slopes of the CPS curves are steeper than those of the TA curves, CPS is also statistically more effective for samples with a higher average delay for a given traffic density.

So far all CPS data have been shown for the case in which overtakes are not permitted. That is, position-shifting for two aircraft was not considered unless one aircraft was traveling through the Keann waypoint, and the other through the Drako waypoint (see fig. 5). As shown in figure 16(a), when this restriction is removed, the reduction in average delay CPS versus no CPS has almost doubled. The cost is a higher workload for the air traffic controller. Figure 16(b) shows similar data for the heuristic CPS as compared with the optimal. As can be seen, the heuristic CPS has only a minor loss in performance compared to the optimal.

The effectiveness of the heuristic CPS depends on the size of the scheduling window. As shown in table 3 and in the inset in figure 16(b), the larger the window, the closer the performance of the heuristic CPS approximates that of the optimal single-position-switch CPS. The mean values shown as dots on the inset of figure 16(b) are above the 0.5 cumulative probability point, since the tails of the

probability distributions are skewed toward large delays. For the large window sizes and a 0.5 traffic mix, the heuristic CPS even performs slightly better than the optimal single-position-shift CPS. This happens because it checks for two extra patterns, which shift one heavy aircraft either forward or backward by two spaces, and because those patterns are more frequent for the 50/50 traffic mix.

So far all cumulative probability curves shown were for 1.5-hr samples. Figure 17 gives the reduction of average delay when the length of the traffic sample is reduced. In figure 17, where the same parameters were used as in figure 13 for 40 aircraft/hr, one can see that the reduction in sample time interval by a factor of 3 reduced the average delay by a factor of more than 2. However, one notices that the benefit of CPS for short samples is much smaller. The effect of longer and shorter sample time-intervals on delays will be investigated later in more detail for FCFS only.

Figure 18 shows the effect of specifying a freeze horizon above the minimum flight time from the Center boundary to landing, in an effort to make a frozen sequence available early to the air traffic controllers. The FCFS

Table 3. Mean delays for 40 aircraft/hr demand for different scheduling algorithms and traffic mix

Traffic mix	Heavy/large = 0.3	Heavy/large = 0.5
Schedule algorithm	Delay, min	Delay, min
FCFS	5.03	7.02
FCFS + TA	4.25	6.20
+ CPS opt	3.88	5.75
+ CPS heur 15-min window	3.93	5.71
+ CPS heur 10-min window	3.94	5.71
+ CPS heur 5-min window	4.02	5.83
+ CPS heur 3-min window	4.15	6.04

curves have been omitted to prevent curves from overlapping. As can be seen, there is a relatively high cost involved in scheduling new arrivals around already frozen aircraft slots. For the high-density traffic, the cost is almost as high as the gain from TA, and it is somewhat smaller for the lower-rate traffic (demand of 30 aircraft/hr).

Now, results for the realizable fuel saving TA algorithm will be given followed by the heuristic CPS algorithm and visa versa. First, the cumulative probability plots of delays will be shown (see fig. 19). Here the average minimum cost t_a s in minutes that have been used for the different traffic demands are shown. Since the curves for the minimum cost t_a include the effects of reducing unnecessary time-advances, the minimum cost t_a curves are not moved to the left by t_a but show a reduction. This is especially noted for lower demands, where more unnecessary t_a s are reduced or removed.

As stated before, a better measure of performance is the cost discussed earlier. Cumulative probability curves for fuel cost are shown in figure 20. Due to the scaling chosen for the cost, the curves for FCFS delays that were shown in figure 19 and those for FCFS costs are identical. For demands of 35 and 45 aircraft per hour, figure 20 also shows the results for reversing the order of the optimization algorithms. Indeed statistically, CPS followed by TA is slightly better than TA followed by CPS, although the reverse is true for about one-third of the individual samples. Unless this change is accomplished easily in the system of reference 1, it may not be worth implementing.

Schedule optimization reduces fuel used in the terminal area and it reduces delays. It increases runway acceptance only indirectly. To show this, demands are chosen which

are above runway capacity. The average landing rates for 10,000 samples plotted against time are shown in figure 21. By postulating two extreme traffic samples one can determine the limits of the acceptance rates. For the case of 50% heavy and 50% large aircraft and only minimum gaps in the sequence, the worst case is a runway acceptance rate of 36 aircraft per hour when aircraft types are exactly alternating, and in the best case 40 aircraft per hour when they are separately grouped. This means that, when the demand exceeds the runway acceptance rate, delays must build up linearly after a transient. When the pool of delayed aircraft is large, it becomes less likely that sequence gaps exist. In that case CPS packs the aircraft more tightly than their FCFS order, and throughput increases. At the smaller demand, TA seems to contribute to the acceptance rate, at least for the early part of the sample. This is probably due to the fact that the TA algorithm spreads out the traffic more evenly when the second part of the algorithm removes the unproductive time-advances for certain aircraft. At large delays for high demands, this process is essentially inactive. But CPS now finds long groups which benefit from position switching. Therefore, for a demand of 45 aircraft per hour, there is an increase in landing rate of 0.75 aircraft per hour, which is an increase of 2%.

Finally, the heuristic TA results will be compared with the fuel saving TA algorithm. In addition the fuel saving TA algorithm, which can be implemented will be compared with the algorithm which minimizes the performance function for each particular sample but which cannot be implemented. In figure 22 data are given for the demand of 35 aircraft per hour, since this is close to saturation, where any improvement counts. The lowest improvement




Figure 19. Cumulative distribution for average delays in traffic samples for different demands.

over FCFS is for the old TA system with 1 minute t_a (curve 1). This is the system, in which an aircraft is time-advanced only if at least the next aircraft benefits. For the old system this performance is improved somewhat if a larger t_a is chosen, which, in this case, was the t_a appropriate for the improved TA algorithm (curve 2). If, in addition, the t_a is removed or reduced when the following aircraft do not benefit, as shown in figure 1, an additional improvement for the old system is obtained (curve 3). This case was tested to explore the possibility of an add-on to the present scheduling implementation. The improved TA algorithm outperforms the original algorithm by a significant margin (curve 4). Finally, the performance of the non-realizable algorithm is shown, which would require the availability of the complete data sample. Here the t_a to minimize the performance index for each particular 1 1/2 hour data sample was chosen (curve 5). The difference between curves 4 and 5 is the small cost one has to pay for a 15 minute windowing of the data for the realizable TA algorithm.

A comparison can be made between the proposed algorithm and the British work presented in references 2 and 3. They propose TA "on" for heavy traffic and "off" for light traffic. This is the proper thing to do; the problem

is to decide when the switch is to occur. The same problem occurred with the algorithm presently implemented in the Traffic Manager Station in the NASA Ames simulation, reference 1. The variable t_a with demand, and especially the second part of the fuel saving algorithm avoids this problem.

Further Traffic Analysis FCFS Only

In the preceding Results subsection it was shown what optimal scheduling can accomplish under various conditions by presenting complete cumulative distributions. Various other effects owing to change in the traffic model or environment will next be briefly treated by discussing the effect on the 50% frequency point of the cumulative distributions. That is, 50% of the samples have higher average delay. Because of the asymmetry of the distribution, the expected value is somewhat higher. This will be reported on FCFS with low horizon only, since the effects of optimization have been pretty well demonstrated in the last section.

Delay as function of length of traffic samples— An individual traffic sample can be thought of as a segment of traffic in which traffic before and after the sample is very light. Figure 23 shows that for relatively brief




Figure 20. Cumulative probability distributions for fuel costs for aircraft with aircraft type mix 50% heavy 50% large.

segments of intense traffic, the average delay per aircraft remains small, even when the arrival rate is higher than runway saturation. Here, the delayed aircraft can be landed quickly after the initial rush is over. However, as the length of the rush period increases, the delays increase sharply, especially for large arrival rates.

Effect of the distribution of arrival times— To obtain the previous results, rectangular center boundary arrival time distributions were always used, which were modified by the requirement of 3-min in-trail spacing upon arrival at the Center borders. Figure 4 showed that the actual sequenced traffic is quite peaked. Although no actual arrival data have been studied as yet, it is likely that the distributions are not rectangular. Therefore results for rectangular distributions will be compared with the same total number of arrivals for triangular distributions over the same time-span. This means that in the center of the studied time-span the traffic is especially heavy with light initial and final traffic. Figure 24 shows that such moderate peaking of traffic about doubles the delays. One can conclude that delays are very sensitive to the distributions of ETAs.

Effect of winds and changes in interarrival times on delays— It has been shown that a 20-knot headwind upon landing increases the required time-separations. Figure 25 shows the statistical results, which are very similar to the results for triangular landing-time distributions. The delays approximately double.

Effect of increasing in-trail spacing— The last few changes that were studied increased the aircraft delays. One of the methods of decreasing the delays taken by the Center is to take delay outside the Center by increasing in-trail spacing. The inset in figure 26 shows schematically how this changes the distributions of incoming traffic. The number of aircraft is the same, but they are spread more evenly and the excess traffic is added as a tail over a longer period of time. The example is for 1.5-hr samples. It is clear that in-trail spacing is very effective in reducing the average delay at the Center. Of course, this assumes that no second traffic peak is expected in the near future.

Effect of more precise guidance using CTAS— As a result of more precise guidance, automation has the potential of reducing the errors in interarrival times, which means that the pad in the separation matrix, which prevents violation of the separation minimums can be reduced. In a detailed simulation of CTAS (ref. 8), it has

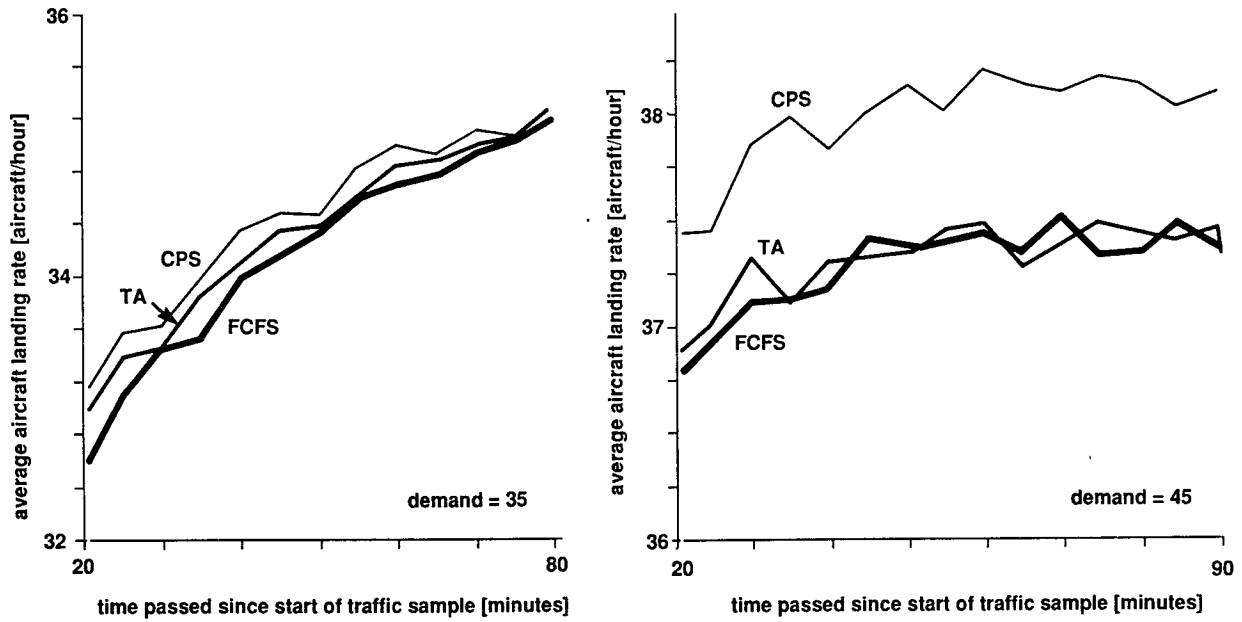


Figure 21. Increase in landing rate with CPS.




Figure 22. Cumulative probability distributions for a demand of 35 aircraft/hour.




Figure 23. Delay as function of traffic sample length.

been shown that one can reduce the interarrival times given in the Scheduling Algorithms section by 10 sec without violating the minimum separation requirements. Using 2500 traffic samples each with FCFS scheduling, this means that, for the nominal traffic mix and a demand of 40 aircraft/hr, the mean delay per aircraft in a 1 1/2 hour traffic sample is reduced from 4.8 to 2.4 minutes, and for a demand of 45 aircraft/hr, the mean delay is reduced from 8.3 to 4.3 min.

Results for Scheduling for Two Dependent Runways

The result of the one step optimization for staggered approaches on 2 dependent runways is exactly the same

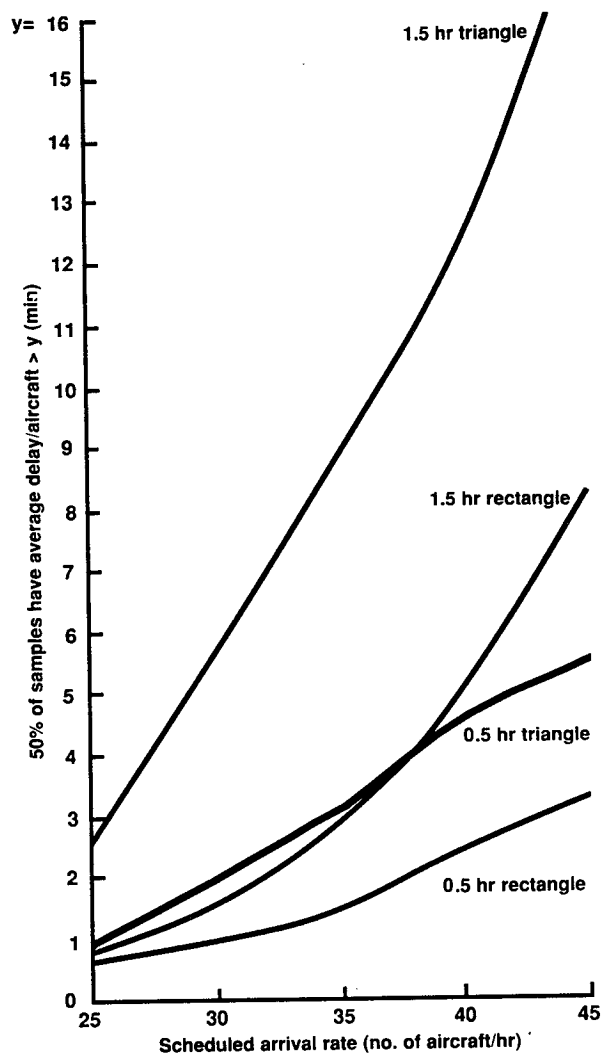


Figure 24 Increase of delay triangular vs. rectangular distributions.

whether one takes the two separately computed streams of ETAs or the STAs computed for independent runways without CPS. This is so, since the algorithm depends only on the order of these sequences, and the order remains exactly the same.

This method results in an ordering which involves largely alternating aircraft landing between runways, especially, when both have the same demand (fig. 27(a)). This method has a distinct advantage over the method where the ordering is based either on the order of the merged ETAs alone, or the order of the independent STAs (fig. 27(b)) alone, although the latter method is better than the first. (See fig. 28 for summary results).




Figure 25. Increased delays with headwind for rectangular traffic distributions.




Figure 26. Effect of variable in-trail spacing.





Figure 27. Scheduling of dependent closely spaced parallel runways with 2 n.mi. stagger. (a) Using 1 step optimization. (b) Same traffic sample as in figure 27(a) using strict FCFS order.

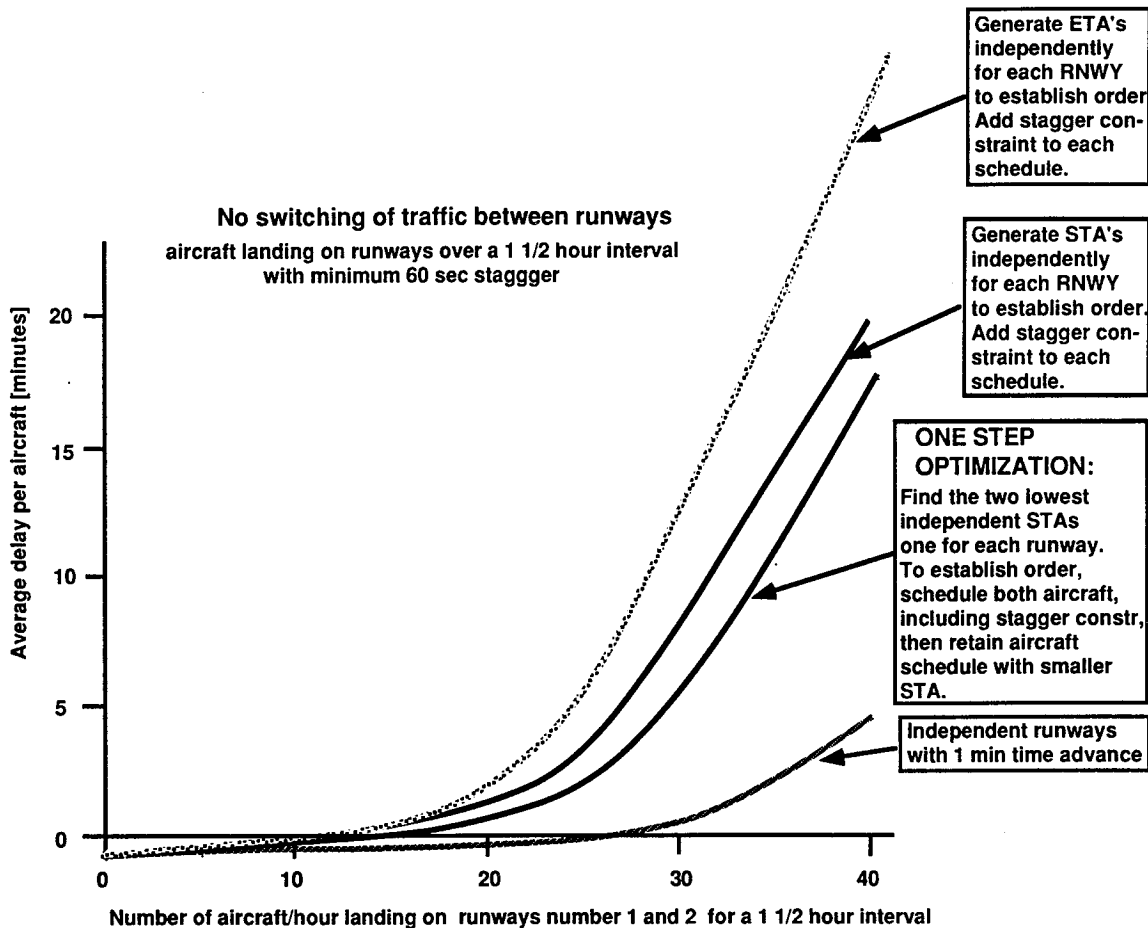


Figure 28. Closely spaced parallel runway operation.

Conclusions

Scheduling was performed in a three-step sequence: one initial ordering FCFS, and two optimization steps, TA and CPS, where CPS is computationally more complicated than TA. Unfortunately, the incremental reduction of the average delay time per aircraft is less for CPS than for TA, but CPS results in a small increase in runway capacity while TA does not.

The heuristic TA method can at best reduce the delay of each sequenced aircraft by the same amount that the first aircraft in each group has been time-advanced, which was 1 min in this report and can be somewhat more in practice. Although the left shift of the cumulative distribution curve owing to TA is almost independent of the traffic density, TA for light traffic is more costly for the airlines. This is because more leading aircraft of smaller groups must be time-advanced, which is unnecessary since delays are already small.

It has been shown that the two step fuel saving TA algorithm increases fuel savings over the heuristic TA algorithm. It has also been shown that proper choice of t_a as a function of demand together with the post processing of the resulting sequence with the second part of the new TA algorithm (fig. 1) results in fuel savings even for the heuristic TA system. However, the best result is obtained when all aircraft in the scheduling window are first time shifted by the proper amount (such a mode is available in the present system), and if then, for the intermediate demands, the algorithm of figure 1 is applied for all but the last aircraft in the window starting with the next to last one. Overall, this will result in a simplification of the scheduling system and it will reduce the traffic manager's workload, by removing his choice of how many aircraft must benefit for time-advance to occur. It will only add slight computational complexity, namely table lookup for the correct t_a , given the estimated demand and runway acceptance rate, and the implementation of the iterative algorithm of figure 1. In this report it was impossible to

treat the degradation of performance that will result when all aircraft do not have the same TA capability. For instance, if in a tightly sequenced group of aircraft the first aircraft has a 1 minute t_a capability, and all other aircraft have a 2 minute capability, then the lead aircraft may force all the remaining aircraft to use only 1 minute t_a . In these situations all aircraft individually should use the maximum t_a possible smaller than the minimum cost one, which does not cause a spacing violation.

CPS is most effective for heavy traffic with large groups of aircraft. For such traffic, CPS can reduce the average time-delay per sample by an additional 20 to 30 sec provided that there exists a relatively even mix between heavy and large aircraft and that traffic density is approximately equal from all directions. For a given traffic sample, this method reduces the average delay per aircraft by a reasonable amount only when position-shifting occurs at the early part of a group, since then all following aircraft in the group have a reduced time-delay. However, in the early part of a group, position-shifting may cause an unrealizable time-advance requirement, and thus cannot always be used.

The effects of increasing levels of scheduling improvements (TA and CPS) are reasonably independent of the actual ETA probability distributions. The basic FCFS delays, however, are very sensitive to these distributions and to the lengths of the traffic peaks. Hence, the data given are meant to show trends rather than to give hard values.

For each landing rate and using the present model of traffic, large deviations from the mean delay occur as a function of the randomness of grouping of the traffic. Although the average delay in the Center airspace can be reduced by reducing the traffic density into the Center by

means of ground holding or in-trail spacing, samples with large delays will still occur occasionally, since traffic from different directions is not time-coordinated. Even global scheduling cannot wholly avoid this occurrence, since random atmospheric effects and other uncertainties will always be present.

When the scheduling freeze horizon is set so that aircraft on shorter routes are inserted into the frozen part of the sequence, the average delay per aircraft increases compared with scheduling with a low freeze horizon.

Parametric studies showed that the actual probability distribution of arrival times (triangular vs flat), presence of headwinds on landing, and an increase in the lengths of traffic samples cause large increases in average delays.

Limited results have been presented for dependent runways. This work can be much expanded if restrictions are relaxed.

In summary, scheduling brings order to an arriving sequence of aircraft. FCFS scheduling establishes a fair order, based on the ETAs and determines proper separations. Because of the randomness of the traffic, gaps will remain in the sequenced sequence of aircraft. The first gap is filled, or partially filled, by TA while preserving the FCFS order. Tightly sequenced groups of aircraft remain with a mix of heavy and large aircraft. Spacing requirements differ for different types of aircraft trailing each other. CPS takes advantage of this fact through mild reordering of the traffic, to shorten the groups, thus reducing the average delays. Actual delays for different samples with the same statistical parameters vary widely, especially for heavy traffic. Both methods of improving the schedule, TA and CPS, work best for heavy traffic.

Appendix A

FCFS Minimizes Standard Deviation of the Delays

We are looking for an ordering of scheduled aircraft, which **minimizes the standard deviation of the delays**. This is thought to simply the air traffic controller's job. To simplify the search for an answer we limit ourselves either to the Denver scheduling method, where, in heavy traffic, all aircraft are scheduled evenly spaced in time, or we limit ourselves to partial groups of aircraft scheduled without a gap and with only one type of aircraft present. This is reasonable, since we can first apply CPS which reorders the aircraft types to minimize mean delays, and then we look at subgroups of the same type of aircraft.

An example of a FCFS sequence is shown in figure A-1 where




Figure A-1

the e_j = ordered sequence of ETA's.
That is if $i > j$ $e_i > e_j$

and s_j = ordered sequence of STA's.
That is if $i > j$ $s_i > s_j$

The e_j 's are given from the incoming traffic flow and cannot be altered. The s_j are given to meet the separation standards and cannot be altered. However we can reassign the aircraft to different time slots within certain restrictions: e.g. the delays must be positive. $(s_j - e_j) > 0$. This statement also means that we cannot interchange positions across a gap. One such possible sequence is shown in figure A-2.




Figure A-2

Let us first look at the mean delay

$$\mu = 1/n \sum_{i=0}^n (s_j - e_j) = 1/n \left(\sum_{i=0}^n s_j - \sum_{j=0}^n e_j \right)$$

where the associated j 's in $(s_j - e_j)$ depend on the order, but each term appears only once, hence we can separate the sums. We can see that mean delay does not change with the assigned landing order of the aircraft.

Now let us look at the standard deviation of the errors. The general equation for the sample variance is

$$S^2 = 1/(n-1) (\Sigma y^2 - n \mu^2)$$

S is the value we want to minimize. Since in our case the mean delays or the number of samples do not change with the altered landing order we only need to minimize

$$\min = \Sigma y^2$$

$$= \sum_{i=0}^n (s_i - e_j)^2$$

where the j may have any order, subject to the realizability restrictions. Expand the last expression

$$\begin{aligned} \min &= \sum_{i=0}^n (s_i^2 - 2s_i e_j + e_j^2) \\ &= \sum_{i=0}^n s_i^2 - 2 \sum_{i=0}^n s_i e_j + \sum_{i=0}^n e_j^2 \end{aligned}$$

The 1st and last sums are independent of the order of the terms hence minimizing the above expression is equivalent to maximizing the 2nd term where the subscript pairs $i-j$ depend on the altered sequence being investigated

36

$$\max = \sum_{i=0}^n s_i e_j$$

Expressing the s's in terms of ΔT (see fig. A-2)

$$\begin{aligned} \max &= \sum_{i=0}^n (T + i \Delta T) e_j = \sum_{i=0}^n T e_j + \sum_{i=0}^n i \Delta T e_j \\ &= \sum_{i=0}^n T e_j + \Delta T \sum_{i=0}^n i e_j \end{aligned}$$

We can now remove all terms that are independent of the order of the e's and we need to maximize

$$\max = \sum_{i=0}^n i e_j$$

Writing out this series for the FCFS example $i = j$ for all i

$$0 e_0 + 1 e_1 + 2 e_2 + 3 e_3 + \dots + n e_n$$

But in FCFS

$$e_0 \leq e_1 \leq e_2 \leq \dots \leq e_n$$

Therefore, the larger the term e_j the larger is the factor that it is multiplied with. By inspection, any change in the order of the i 's from the FCFS order will reduce the value of the expression. Hence, the **FCFS schedule minimizes the standard deviation of the delays when all aircraft are to be spaced equidistant from each other**. Exploring sequences with a mix of aircraft is of little value, since this has already been done to reduce the mean delay.

Appendix B

An Exact Solution of the Constrained Position Shift Problem for the Single Position Shift

Introduction

The FAA mandates that various separations be maintained between landing aircraft based on their weights, and generally the lighter the aircraft the greater the spacing required. Clearly, then, the amount of time required to land a given set of aircraft can depend on the landing order.

One approach to finding an “optimal” landing order is the Constrained Position Shift (CPS) concept of Roger Dear. He posited that, given an initial arrival ordering, real-world constraints would preclude moving any of the aircraft more than some small number of positions from its original place in the arrival list. However, he did not present an exact solution to the CPS problem; his method was to examine a window of $2 \cdot \text{MPS} - 1$ positions, optimize it (exhaustively) for a single position shift, move the window down one position, and repeat the process. A later effort (Luenberger) improves on the performance of the Dear algorithm but still fails to consistently achieve optimal performance.

This paper presents an algorithm for finding an optimal solution to the CPS problem for a single position shift. This method has been developed by J. C. Jackson, who was at Ames Research Center in 1989.

The Algorithm

Finding the optimal ordering of a set of aircraft can be thought of as a search for the least “cost” path through a tree of possible aircraft orderings, where the cost is the sum of the time separations required between each pair of aircraft. For the CPS problem, an initial ordering of aircraft is given, along with a list of delays from the ETA and the maximum possible time-advance for each aircraft. In the final ordering each aircraft is constrained to lie within one position of its initial position, and no aircraft must have a time of arrival earlier than permitted by the maximum allowable time-advance. Figure B-1 illustrates the tree of possible orderings for the simplest case of $\text{MPS} = 1$. Note that the first aircraft (A) in the initial ordering is in our method constrained to be the first aircraft in the output ordering.




Figure B-1. Illustration of all possible re-orderings of aircraft, given initial ordering A-B-C-D-E-...

Thus, the only aircraft which can appear in position two of the final ordering are B and C, due to the MPS constraint. If the final ordering begins A-B, then C or D may be in the third position. However, if it begins A-C, B must appear next in the sequence since B can appear no later than the third position. Reasoning along these same lines produces the rest of the tree.

The algorithm for finding the least cost path through this tree is essentially an application of the dynamic programming principle: only extend the shortest path through a given set of nodes terminating at a particular node. For example, the paths A-B-C-D and A-C-B-D are both valid MPS = 1 paths terminating with aircraft D and containing the same aircraft. However, in general one of these paths will have less cost than the other, and only that path need be considered in further computations by the algorithm. This is because the optimal ordering of the remaining aircraft is independent of the order of B and C in the path to D. So if, for example, the path A-B-C-D is 15 units cheaper than the other path, the cheapest complete ordering beginning with A-B-C-D will be 15 units cheaper than the one beginning with A-C-B-D; that is, the optimal ordering of the remaining aircraft will be the same for both. This simple idea allows a great savings in the computation of the least cost path. For the MPS = 1 case, the algorithm begins by computing and storing the (time) cost of having B follow A and that of having C follow A (the paths A-B and A-C). It then computes and stores the costs of A-B-C, A-B-D, and A-C-B, discarding the two previously computed values. In the remainder of the processing, the dynamic programming principle is applied. For example, both A-B-C-D AND A-C-B-D are computed, but only the value of the lesser cost path is stored. Once all the values at each level of the tree have been computed, the previous level's values are discarded. It turns out that in this MPS = 1 case there is only one set of aircraft which can precede a given aircraft at a given level (e.g. A, B, and C in some order must precede D if D is going to be in the fourth position of the final ordering). Thus only six values (three for the current level of the tree and three for the previous) must be stored by the program to compute the value of the optimal path. This process of extending least cost paths eventually terminates when each path has N (the number of aircraft) aircraft along it. For the MPS = 1 case, only the last two aircraft in the initial list are candidates for being last in the optimal ordering. Thus the least cost paths leading to these aircraft at the lowest level of the tree are compared and the smaller cost path is chosen as the final optimal path.

For example, assume that the spacing times required for various pairs of five aircraft are as given in table B-1.

Table B-1. Example time spacing requirements for a set of five aircraft

Costs	A	B	C	D	E
A	-	2	2	4	3
B	6	-	3	2	4
C	3	2	-	3	2
D	4	5	6	-	4
E	3	2	2	2	-

Each value represents the time that the aircraft labeling the column must follow the aircraft labeling the row by. Notice that the values are not symmetric (e.g. it costs less for A to follow B than vice versa).

Tracing through the tree of figure B-1 (and ignoring the now undefined aircraft F) we find that there are two paths to aircraft D at level four, A-B-C-D and A-C-B-D, and that their respective costs are eight and six. Thus A-C-B-D is chosen as the preferred path to this node. Likewise, A-B-C-E is the low cost path (seven) to E at this level and A-B-D-C is the only path (cost 10) to C at this level. Extending these three paths to the fifth and final level, we find that A-C-B-D-E (cost 10) is better than A-B-D-C-E (cost 12), but A-B-C-E-D is preferable to both of these (cost nine). This final path is therefore chosen as the overall optimal path.

An additional detail of the algorithm which has so far been neglected is the maintenance of the list of best paths to each node of the search tree. This can be handled in a number of ways; a particularly simple way for the MPS = 1 case is to simply maintain three vectors which represent the best path thus far to the leftmost, middle, and rightmost nodes of the tree. For example, when A-C-B is chosen as the best path to D at level four in the example above, this path (the leftmost path at level three) can be copied to the vector for the middle path (position of D at level four) and D can be appended. Of course, care must be taken not to overwrite a vector representing a path at the previous level before that level has been completely processed, so two sets of three vectors (one for current level and one for previous) can be used.

References

1. Erzberger, H.; and Nedell, W.: Design of Automated System for Management of Arrival Traffic, NASA TM-102201, June 1989.
2. Magill, S. A. N.; and Budd, A. J.: Negative Delay in the Airport Arrival Process, Proc. UKSC Conference on Computer Simulation, Bangor, Sept. 1987.
3. Magill, S. A. N.; Tyler, A. C. F.; Wilkinson E. T.; and Finch, W: Computer-Based Tools for Assisting Air Traffic Controllers with Arrivals Flow Management, Royal Signals and Radar Establishment Report No. 88001, Malvern, Worcestershire, 1988.
4. Neuman, F.; and Erzberger, H.: Analysis of Sequencing and Scheduling Methods for Arrival Traffic. NASA TM-102795, April 1990.
5. Dear, R. G.: The Dynamic Scheduling of Aircraft in the Near Terminal Area: M. I. T. Flight Transportation Laboratory Report R76-9, September 1976.
6. Voelkers, U.: Computer Assisted Arrival Sequencing and Scheduling with the COMPAS-System, AIAA/ASCE - International Air Transportation Conference, Dallas/Texas, June 2-4,1986.
7. Alcabin, M.; Erzberger, H.; Tobias, L.; and O'Brien, P. J.: Simulation of Time-Control Procedures for Terminal Area Flow Management, 1985 American Control Conference, Boston MA June 19-21, 1985. TA1 - 11:00
8. Davis, T. J.; Erzberger, H.; and Green, S. M.: Simulator Evaluation of the Final Approach Spacing Tool, NASA TM-102807, May 1990.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE October 1991	3. REPORT TYPE AND DATES COVERED Technical Memorandum	
4. TITLE AND SUBTITLE Analysis of Delay Reducing and Fuel Saving Sequencing and Spacing Algorithms for Arrival Traffic			5. FUNDING NUMBERS 505-64-13	
6. AUTHOR(S) Frank Neuman and Heinz Erzberger				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Ames Research Center Moffett Field, CA 94035-1000			8. PERFORMING ORGANIZATION REPORT NUMBER A-91203	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) National Aeronautics and Space Administration Washington, DC 20546-0001			10. SPONSORING/MONITORING AGENCY REPORT NUMBER NASA TM-103880	
11. SUPPLEMENTARY NOTES Point of Contact: Frank Neuman, Ames Research Center, MS 210-9, Moffett Field, CA 94035-1000 (415) 604-5437 or FTS 464-5437				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Unclassified-Unlimited Subject Category - 03			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) The air traffic control subsystem that performs sequencing and spacing is discussed. The function of the sequencing and spacing algorithms is to automatically plan the most efficient landing order and to assign optimally spaced landing times to all arrivals. Several algorithms are described and their statistical performance is examined. Sequencing brings order to an arrival sequence for aircraft. First-come-first-served sequencing (FCFS) establishes a fair order, based on estimated times of arrival, and determines proper separations. Because of the randomness of the arriving traffic, gaps will remain in the sequence of aircraft. Delays are reduced by time-advancing the leading aircraft of each group while still preserving the FCFS order. Tightly spaced groups of aircraft remain with a mix of heavy and large aircraft. Spacing requirements differ for different types of aircraft trailing each other. Traffic is reordered slightly to take advantage of this spacing criterion, thus shortening the groups and reducing average delays. For heavy traffic, delays for different traffic samples vary widely, even when the same set of statistical parameters is used to produce each sample. This report supersedes NASA TM-102795 on the same subject. It includes a new method of time-advance as well as an efficient method of sequencing and spacing for two dependent runways.				
14. SUBJECT TERMS Air traffic control, Aircraft scheduling, Optimization			15. NUMBER OF PAGES 45	
			16. PRICE CODE A03	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT	

