


AD-A205 651

Report No. 88001

Report No. 88001



ROYAL SIGNALS AND RADAR ESTABLISHMENT,
MALVERN

DTIC
ELECTE
FEB 07 1989
S D D

COMPUTER-BASED TOOLS FOR
ASSISTING AIR TRAFFIC CONTROLLERS
WITH ARRIVALS FLOW MANAGEMENT

Authors: S A N Magill, A C F Tyler,
E T Wilkinson, W Finch

DISTRIBUTION STATEMENT A
Approved for public release
Distribution Unlimited

PROCUREMENT EXECUTIVE, MINISTRY OF DEFENCE
RSRE
Malvern, Worcestershire.

September 1988

RSRE REPORT No. 88001

COMPUTER-BASED TOOLS FOR ASSISTING
AIR TRAFFIC CONTROLLERS WITH ARRIVALS
FLOW MANAGEMENT

S A N Magill, A C F Tyler, E T Wilkinson, W Finch


September, 1988

SUMMARY

This document reports a research project on the use of computer-based tools for arrivals flow management at major airports. It describes in detail two experimental prototype tools which were constructed as part of the research project, and exercised in a real-time simulation environment. These were a Landing Order Calculator and a Speed Control Adviser. Both give advice to air traffic controllers for aircraft which are in the vicinity of their top-of-descent points. The underlying concepts and the methods used in the tools are fully described. The real-time simulation environment and the experience gained from it are discussed.



Copyright
©
Controller HMSO London
1988



Contents

EXECUTIVE SUMMARY	v
1 INTRODUCTION	1-1
1.1 Context of the Research	1-1
1.2 Structure of the Report	1-3
2 CONCEPTS	2-1
2.1 The Arrivals Problem	2-1
2.2 Time Adjustment Mechanisms	2-2
2.3 Possibilities For Computer Assistance	2-3
2.3.1 Landing Order	2-3
2.3.2 Terminal Approach Time	2-4
2.3.3 Speed Control	2-5
2.4 The Time Horizon Principle	2-6
2.5 Trajectory Prediction	2-7
2.6 The Controller/Computer Relationship	2-8
3 LANDING ORDER CALCULATOR	3-1
3.1 Introduction	3-1
3.2 Landing Times Allocation	3-1
3.3 Functional Description of the Allocation Process	3-2
3.3.1 Overview	3-2
3.3.2 Algorithmic Details	3-2
3.4 Use of Trajectory Prediction	3-3
3.5 Data Structures	3-4
3.6 Outputs and Required Inputs	3-4
4 SPEED CONTROL ADVISER	4-1
4.1 Introduction	4-1
4.2 Functional Description	4-2
4.3 SCA Outputs and Required Inputs	4-4
5 Optimising the Landing Order	5-1
5.1 Introduction	5-1
5.2 Optimised Landing Times Allocation	5-1
5.3 Functional Description of the Optimised Allocation Process	5-1
6 TRAJECTORY PREDICTION	6-1
6.1 Introduction	6-1
6.2 Overview	6-1
6.3 Forward and Backward Prediction	6-1
6.4 Aircraft Performance Models	6-2
6.4.1 Flight Idle Descent Performance	6-2
6.4.2 Deceleration Characteristics	6-3
6.4.3 Acceleration Characteristics	6-3
6.4.4 Operational Envelope Limit Parameters	6-3
6.4.5 Approach Speed and Flap Deployment Data	6-4
6.5 Height Change Modes	6-4

6.6	Airspeed Constraints	6-4
6.7	Turns	6-5
6.8	Integration	6-5
6.9	The Approach Phase	6-5
7	EXPERIMENTAL ENVIRONMENT	7-1
7.1	Introduction	7-1
7.2	The Skeleton Control Centre Software (SCCS)	7-2
7.3	The Air Traffic Simulator	7-3
7.4	The Traffic Sample Generator	7-4
7.5	The Automatic Controller	7-5
7.6	Software Implementation	7-5
8	CONTROLLER INPUT AND DISPLAY INTERFACES	8-1
8.1	Overview	8-1
8.2	PVDs	8-1
8.3	EDDs	8-2
8.4	FPUDs	8-3
9	EXPERIENCES FROM USING THE PROTOTYPE TOOLS	9-1
9.1	Introduction	9-1
9.2	Overview	9-1
9.3	Conduct of the Experiments	9-2
9.3.1	Landing Order Calculator Experiments	9-2
9.3.2	Speed Control Adviser Experiments	9-2
9.4	Effects of Controller Intervention	9-3
9.5	Overtaking Aircraft	9-4
9.6	Landing Sequence Numbers	9-5
9.7	Factors Influencing the Choice of Time Horizon Value	9-6
9.8	Alternative Routings	9-7
9.9	Adjusting the Planned Flow	9-8
9.10	Creating Spare Landing Slots	9-9
9.11	Optimising the Landing Sequence	9-10
9.12	Display Issues	9-10
9.13	Impact on the normal task of the controller	9-10
10	CONCLUSIONS	10-1
11	ACKNOWLEDGEMENT	11-1
	REFERENCES	R-1
	APPENDIX A:	A-1
	SIMULATION STUDY OF "NEGATIVE DELAY"	
	A.1 The Simulation	A-1
	A.2 Results	A-1
	APPENDIX B:	B-1
	SIMULATION STUDY OF THE TIME HORIZON PRINCIPLE	
	B.1 The Simulation	B-1
	B.2 Results	B-1

APPENDIX C:	C-1
TRAJECTORY PREDICTION - EXAMPLE	
C.1 Introduction	C-1
C.2 The Problem	C-1
C.3 The Solution	C-1
C.3.1 From Cruise Down to 6000 Ft Altitude	C-1
C.3.2 From Ten Miles before LAM to Touchdown	C-1
C.4 Implementation	C-2
APPENDIX D:	D-1
EXPERIMENTAL ENVIRONMENT - CONTROLLER INTERFACES	
D.1 Flight Progress EDD Formats	D-1
D.2 FPUD Formats	D-1
GLOSSARY	G-1

List of Figures

3.1	LOC Functions	3-6
3.2	Predicted Flight Profile	3-7
3.3	Aircraft State Table Structure	3-7
3.4	Landing List Structure	3-8
5.1	OptO Method Functions	5-4
7.1	Experimental Environment	7-7
7.2	Skeleton Control Centre Software	7-8
A.1	Effect of Speedup on Mean delay	A-3
A.2	Effect of Speedup on Cumulative Delay Distribution	A-3
B.1	Effect of Displacement from Time Horizon	B-2
B.2	Effect of Allocation Time when not using Time Horizon	B-2
C.1	Inbound Route Plan View	C-6
C.2	Trajectory Specification	C-7
C.3	Prediction Events from Cruise down to 6000ft Altitude	C-8
C.4	Prediction Events from Ten Miles before LAM to Touchdown	C-9
D.1	Stack Inbound Sector Flight Progress EDD	D-4
D.2	Enroute Sector Flight Progress EDD	D-5
D.3	Rest Picture showing Callsigns and Accept Marks	D-6
D.4	Picture after BA069 Touched	D-7
D.5	Picture after Sequence Number Exchange Touch Sequence	D-8
D.6	Picture after Touching Height Field	D-9
D.7	Picture after Touching TAT Field	D-10
D.8	Picture after Touching Speed Field	D-11
D.9	Picture after Touching Heading Field	D-12
D.10	Picture after LRATE Touched	D-13
D.11	Picture after PARAMS Touched	D-14
D.12	Top Level State Transitions	D-15
D.13	Flight Strip State Transitions	D-15
D.14	Sequence State Transitions	D-16
D.15	Height State Transitions	D-16
D.16	Speed State Transitions	D-17
D.17	Heading State Transitions	D-17
D.18	TAT State Transitions	D-18
D.19	Landing Rate State Transitions	D-18
D.20	Parameters State Transitions	D-19

EXECUTIVE SUMMARY

This report describes research work which has led to the development of two experimental prototype computer-based tools for assisting air traffic controllers with the flow management of traffic arriving at busy major airports. The tools are:-

1. *A Landing Order Calculator (LOC).*

This performs automatic calculation of landing order and Terminal Approach Times (also known as Stack Departure Times), for aircraft which are still at their cruising levels.

The research indicates that the benefits to be derived from this tool are:-

- Reduction in controller workload because:-
 - controllers do not have to determine landing sequence numbers or communicate them to each other,
 - controllers receive an early indication of when holding is required,
 - early determination of optimised landing order (optimised to minimise the effect of wake turbulence separation rules on landing rate) allows the process of traffic re-ordering to begin outside the approach sequencing area.
- A smoother and more orderly flow of traffic into the approach sequencing area because:-
 - early display of landing sequence numbers enables TMA controllers to leave gaps in the stream for later merging of traffic from other directions,
 - early communication of Terminal Approach Times to aircraft enables pilots and avionic systems to plan accurate stack exit times.

2. *A Speed Control Adviser (SCA).*

This includes all the functions of the LOC, and in addition calculates descent speeds needed to achieve the planned Terminal Approach Times.

The research indicates that the benefits to be derived from this tool (in addition to those listed for the LOC above) are:-

- Reduced delays or increased capacity because:-
 - increasing the speeds of a few aircraft allows runway time to be used which would otherwise be wasted,
 - early determination of landing order together with speed control allows a greater degree of order optimisation (for aircraft which are not holding) than is possible today.
- Reduced holding because of use of speed variation as a delay absorbing mechanism.
- A smoother flow of traffic into the approach sequencing area (in addition to that gained from the LOC) through the use of speed control.
- Improved aircraft fuel economy through reduced holding.

The LOC and SCA were exercised extensively by Air Traffic Control Officers in a real-time simulation environment. This led to a very successful direct involvement of controllers in the development process. The experiments have of necessity been conducted on a small scale (two manned sectors), but have used the routes and procedures proposed for the CCF development of the London Terminal Area. As a result of this research, the CAA plans to integrate these tools into its CCF development programme, and now sees them as part of its operational requirement for the CCF.

A full list of conclusions is given in section 10 of the report.

1 INTRODUCTION

1.1 Context of the Research

This report is the primary record of a research activity sponsored by the UK Civil Aviation Authority (CAA), and carried out at RSRE by a joint RSRE/CAA research team known as the Terminal Control Systems Development Group (TCS DG). Although the TCS DG has been in existence since 1979, it adopted its current direction in the field of arrivals management early in 1984, and the work described here has been done since then.

The aim of the research activity was to investigate how computer technology could be used to assist air traffic control (ATC) in dealing with the special problems of traffic arriving at major airports such as London Heathrow. It was considered that if successful, the methods developed might eventually be incorporated into the operational system used at the London Air Traffic Control Centre (LATCC) during the early to mid 1990s. At that time traffic will be controlled by Air Traffic Control Officers (ATCOs) using radar displays and radio telephony much as today. The emphasis of the research was very much on "computer assistance" to the human controllers, rather than on "automation" of their functions. In practice this meant providing information and advice for controllers in a way which supported existing methods of performing the control task, and in fact trying to interfere as little as possible with existing methods, rather than trying to find new methods.

The general objective of all ATC systems is to produce a safe, orderly and expeditious flow of traffic. The specific objective of the computer-based tools described here was to maintain present levels of safety, while increasing the number of aircraft movements handled per unit time and/or reducing the costs of aircraft operation.

The process of introducing substantial changes into an operational ATC system is a long and careful one. In the UK three major stages can be identified: a research stage, a large-scale evaluation stage, and an operational implementation stage. Only the first stage can be carried out at a research establishment such as RSRE. The main output from the research stage is a statement that it is, or is not, worth proceeding to the much-more-expensive second stage and starting to plan the third stage. In the case of the work reported here, the CAA has now decided to proceed with the second stage at the Air Traffic Control Evaluation Unit (ATCEU), and is making plans for the third stage. A secondary output from the research stage is a description of the details of the methods developed, which can be used by later stages. That is provided by this report.

During the course of the research activity experimental prototype versions of two computer-based tools were developed, namely:-

- *A Landing Order Calculator (LOC).*

This tool plans a landing sequence, and recommends to the controller an appropriate position in the sequence for each inbound aircraft. The recommendation is made when the aircraft is in the neighbourhood of its top-of-descent point, typically 100 - 150 miles

out from the runway. The LOC also provides advice about Terminal Approach Time (TAT), and about whether or not the aircraft will be required to hold.

- *A Speed Control Adviser (SCA).*

This tool includes the functions of the LOC, but in addition recommends a descent speed for each inbound aircraft which will cause the aircraft to arrive in the approach sequencing area close to the planned time.

The construction of these experimental prototype tools, together with a real-time environment in which controllers could make use of the tools while controlling simulated traffic, played a central part in the research activity. The construction of such an environment and the use of man-in-the-loop simulation are both costly activities. The reasons for adopting this approach were:-

- *To stimulate the development of new ideas.* As Knuth observed [1]: "We often fail to realize how little we know about a thing until we attempt to simulate it on a computer". The ideas thus generated can be incorporated into the prototype tools and tried out. This leads to an iterative style of development.
- *To involve qualified practicing controllers in the research activity.* By using the prototype tools in the performance of a control task, the controllers can relate to them in a very direct way, and can perceive their advantages and disadvantages in a way which would not otherwise be possible.
- *To be in a position to demonstrate to other interested parties not directly involved in the research the appearance of an ATC system which includes the proposed computer-based tools.*

It is important that the reasons for using real-time simulation and software prototyping should not be misunderstood. The aim is NOT to attempt to quantify from the real-time simulations improvements in parameters of interest such as movement rates and delays. These parameters exhibit great variability compared with the small improvements in mean values being sought, and to estimate with useful statistical significance changes in their mean values would require very many hours of simulation and qualified controllers' time, which would be prohibitively expensive. Instead it is better at the research stage to attempt to quantify such improvements by non-real-time simulation.

While the TCSDG research activity was in progress, the CAA formulated a plan for development of LATCC in the early 1990s. The planned development is known as the Central Control Function (CCF). Its main aim is to increase substantially the traffic capacity of the London Terminal Area. Its main features are the relocation of Approach Control functions from the major London airports to LATCC, and a redesign of the airspace and ATC procedures in the London area. There has been close liaison between the CCF development team and the TCSDG. The ATC routes and procedures used in the TCSDG simulation environment were those developed for the CCF, and ATCO members of the CCF team have frequently taken part in TCSDG simulation exercises. The CAA now

plans to incorporate the tools described here into the CCF development during the forth-coming large-scale evaluations at ATCEU.

The TCSDG is at present also engaged in research into computer assistance for Approach Control, and this work could be considered to fall within the scope of this report. However the work on computer assistance for Approach Control is still in progress, and will be reported separately at a later date. The tools described here are concerned with traffic between cruising level and the edge of the approach sequencing area, up to about 30 miles out from the runway.

1.2 Structure of the Report

Section 2 discusses the basic concepts which lie behind the experimental work on arrivals management and introduces some terms used later on. The arrivals problem is introduced by reference to queuing theory and the arrivals task summarised as the merging of traffic into a single stream for landing, smoothing the flow and modifying the landing order to maximise the landing rate. The various mechanisms for adjusting aircraft arrival times are described and the notion of using speed control to advance as well as retard a flight is introduced. Some possibilities for computer assistance are postulated and these were incorporated in the experimental tools described later. These include computer generation of the landing order, the calculation of Terminal Approach Times (TATs) and the computation of an airspeed for descent. The principle of using a "time horizon" to determine when, for each aircraft, landing times are allocated, is explained and justified in terms of fairness to all aircraft. Brief mention of the requirements for predicting the trajectories of aircraft from the ground acts as preparation for a fuller treatment in section 6. Finally, some assertions are made about the nature of the relationship between controller and computer. The two extremes of complete computer-control and the machine as a passive advisor are defined. In addition, the "tightly-planned" and "loosely-planned" methods for planning the arrivals stream are contrasted and the reasons why the latter was selected are explained.

Section 3 describes how the LOC produces landing sequence number, TAT and predicted Stack Arrival Time when aircraft are required to hold. The two states "busy" and "non-busy" are defined; in the busy state aircraft might be asked to meet earlier than preferred landing times. The rules for absorbing delay and for establishing a stack are also described. The significant events which drive the landing times allocation process are defined together with the actions which are taken accordingly. The major functions are given in flowchart form, together with a description of the main data structures. The inputs and outputs to the tool are detailed.

Section 4 introduces the Speed Control Adviser by describing how the rules for absorbing delay used for the LOC can be modified to use speed-control as a delay-absorbing mechanism; the revised rules for computing TATs are also detailed. In the functional description a full account is given of the way in which the Calibrated Airspeed (CAS) for descent is computed using an interpolation method. Finally, some additional inputs to the tool are defined. Discussion of these tools and the experiences gained within the experimental system is deferred until section 9.

In section 5 a modified form of the landing times allocation algorithm used by both the LOC and the SCA is described which, instead of using a first-come-first-served order, attempts to achieve the best possible order taking account of the minimum separations required according to the turbulent wake vortex rules.

Section 6 describes the methods for predicting the trajectories used by the experimental tools, and the way in which aircraft performance is modelled. The notion of predicting in a forward or backward manner is introduced. A trajectory is divided into small time increments and numerical integration is used to produce times taken for various manoeuvres such as descents at constant CAS or constant Mach numbers. A number of different aircraft types are modelled using sets of polynomial coefficients.

The experimental environment is described in sections 7 and 8 with section 8 concentrating on the details of the controller interface. In section 7 the terms "Skeleton Control Centre Software", "Air Traffic Simulator", "Traffic Sample Generator" and "Automatic Controller" are fully described. It is shown how these elements combine with display and input hardware to provide an experimental environment which is sufficient to demonstrate the prototype versions of the LOC and SCA. There is a short description of the software implementation at the end of the section. Section 8 describes the three types of experimental display, viz. the radar Plan View Display, the flight progress data EDD and the flight progress updating device. Details of the display formats are contained in Appendix D.

As mentioned above section 9 discusses the experience gained with the tools in the experimental environment. It represents, possibly, the most significant section of the report and covers a number of aspects which arose directly as a result of running the tools with controllers. It was seen that the operation of the tools could be affected by the "tactical actions" of controllers and facilities for the controller to override manually the allocated landing sequence numbers and to request the computer to produce a revised speed for an aircraft were added as a result. Problems of planned overtakes on the same route, use of alternative routings, how to adjust the planned flow and how to deal with occurrences such as missed approaches, and also aircraft which depart from points within the time horizon are fully covered. A short assessment of the merits of using the OptO planning algorithm is given and, finally, some comments on the impact for the controller in using the proposed tools is made.

A list of conclusions is given in section 10 and the report ends with a set of four appendices which give fuller details of aspects mentioned in the main body of the document. In particular, the results of two fast-time simulation studies are given in Appendices A and B.

2 CONCEPTS

This section introduces the basic concepts which underlie the TCSDG research into computer assistance for arrivals management.

2.1 The Arrivals Problem

If there were no coordination between airlines using busy major airports, then at peak periods the rate at which aircraft would arrive in the vicinity of an airport would be greater than the maximum possible landing rate. Long delays and chaos would result. However at most major airports the necessary coordination machinery does exist. It takes the form of a Scheduling Committee [5] which operates under the auspices of the International Air Transport Association (IATA). In most cases the scheduling committee holds a scheduling conference twice a year at which the proposed timetables of all interested airlines are coordinated to form a combined schedule such that mean arrival rate is closely matched to runway capacity.

If scheduled arrival rate was arranged never to exceed maximum possible landing rate, and if all aircraft flew exactly to schedule, there would be no arrivals problem. However aircraft operators are prevented from keeping closely to schedule by many factors beyond their control. The most significant factor (especially for long haul flights) is variation in wind conditions; for most modern jet airliners the speed margins available at cruising altitudes are much smaller than the day-to-day variations in along-track wind components. The next most important factor is random displacement of actual take-off times from scheduled take-off times at airports from which the aircraft originated. There are many other smaller factors which contribute to deviation from scheduled arrival times, such as ATC vectoring to avoid potential conflicts with other aircraft, variations in how different crews fly the aircraft, and so on. All these factors combine to cause each aircraft's actual arrival time to be randomly displaced from its scheduled arrival time. The magnitude of the displacement is many times greater than the mean inter-arrival time. One study [2] found that the difference between scheduled and actual arrival time had a mean close to zero, and a standard deviation of 27 minutes as compared with a mean inter-arrival time of about 1.7 minutes.

These random displacements from scheduled arrival times lead to the formation of a queue, as pointed out by Attwooll [3]. Why should a queue form when the mean flow is carefully matched to the runway capacity? The reason is that the random displacements of individual aircraft lead to local fluctuations in the total flow, and during troughs in the total flow some runway time is lost; aircraft which were scheduled to land in the lost time will then compete with other aircraft for the remaining time. Fortunately each busy period has only a limited length, so the queue dies away at the end of the busy period. Scheduling committees recognise the existence of the queue, and schedule on the basis that a mean delay of 5 minutes is acceptable during busy periods¹. However, if the mean delay is 5 minutes, the delay experienced by an individual aircraft will sometimes be several times this value.

Each landing aircraft needs to make exclusive use of the runway for a short time,

¹5 minutes is the UK figure; other figures are used in other parts of the world

and on the final approach path the separations between successive aircraft must not be less than the minima prescribed by ATC separation rules. Generally, landing rate is limited by separation minima rather than by runway occupancy times. In order to maximise landing rate, separations between successive aircraft should exceed the allowed minima by as little as possible. Because of wake turbulence effects, the minimum separation allowed between two aircraft is a function of the aircraft types of the leading and following aircraft. Therefore the maximum possible landing rate on any occasion will depend on the traffic mix on that occasion, and on the order in which the traffic lands.

The arrivals management task can now be stated as follows:-

- Aircraft from several different directions must be safely merged into a single stream before reaching the runway. At the same time they must be descended from cruising level to ground level.
- The random fluctuations in total inbound flow must be smoothed out so as to give a precisely-spaced even flow at the runway. It is this smoothing which is the essence of arrivals flow management.
- The natural first-come-first-served arrival order may be modified slightly to maximise landing rate.

2.2 Time Adjustment Mechanisms

To smooth the inbound traffic flow and achieve the required precise spacing at the runway, it is necessary to adjust the position of each individual aircraft relative to the rest of the landing stream. Each aircraft may be advanced or delayed. The possibilities for advancement are very limited, and once a queue exists it is delay which is required, so the time adjustment mechanisms are sometimes referred to as delay absorption mechanisms. There are three possible mechanisms:-

- *Path-length Variation.*

An aircraft can be delayed or advanced by causing it to fly a path longer or shorter than the standard path to the runway. It can be advanced by shortening the path only if some excess distance has been built into the standard path. Path-length variation is widely used in approach control for making fine adjustments to the positions of individual aircraft relative to the rest of the stream. The amount of advance or delay which can be obtained is limited by the amount of airspace available. In general the method cannot be used in en-route airspace in Europe because of insufficient airspace, though it is used in some specific instances.

- *Holding in a Stack.*

A Holding Stack is an arrangement in which many aircraft can fly closed loops (sometimes referred to as orbits) simultaneously, one stacked above the other, according to procedures [4] which ensure that they are safely separated from each other and from other traffic in the neighbourhood. Holding is a special case of path-length variation which requires a fixed amount of airspace, and which can provide only delay, not advance. The delay time available is not continuously variable, but is a multiple of the time for a single orbit, and this can be varied only within certain limits. The minimum time for an orbit is determined mainly by aircraft bank angle considerations and prevailing wind conditions. The maximum time is determined by the volume of protected airspace reserved for the stack [4].

- *Speed Variation.*

An aircraft's speed can be varied within certain limits which depend on the aircraft's type and weight and the altitude at which it is flying. For many modern jet airliners, the speed flexibility which the operators are prepared to use at cruising altitudes is about + or - 5%, whereas at lower altitudes (below flight level 270) it is about + or - 25%[6]. The amount of time adjustment obtainable from this mechanism depends on the amount of speed flexibility available, and the length of time over which it is applied. Aircraft operators prefer to operate at cruising altitudes as much as possible consistent with efficient descent profiles for fuel economy reasons. The speed flexibility available at cruising altitude is of little practical value for arrivals smoothing purposes, so the amount of time adjustment available is limited by the time spent in descent.

Today's ATC practice must of course be able to cope with busy arrivals peaks at major airports. The general method used is to hold aircraft in a number of Holding Stacks (typically 2, 3, or 4 stacks) some 30 to 60 miles out from the runway. The stacks can be thought of as buffers between the airways and the airport. Aircraft are released from the stacks at a steady rate. The block of airspace between stacks and runway is known as the approach sequencing area. It is used by approach controllers for path-length variation, in order to make fine adjustments to the position of each aircraft relative to the rest of the stream. Speed variation is sometimes used by approach controllers, and is occasionally used by en-route controllers before the stacks, but full exploitation of the possibilities of speed variation throughout the descent phase is not possible without computer assistance.

2.3 Possibilities For Computer Assistance

2.3.1 Landing Order

Possibly the simplest form of computer assistance for arrivals is advice about landing order. The computer could calculate a position in the landing sequence

for each aircraft while the aircraft was still 100 miles or more out from the runway, (the precise point at which such calculation should be done will be discussed in section 2.4). The calculated landing sequence number could then be displayed as part of the aircraft's data block on the radar display. This would help the controller by making it clear a long way out from the runway which aircraft were already in the correct relative positions, and should maintain those relative positions, and which ones should be separated to allow aircraft from other directions to be inserted between them later.

The calculated order need not be the first-come-first-served order. In current ATC practice approach controllers do some re-ordering of the traffic. They do this to reduce the effect of wake turbulence separation rules, and to take several successive aircraft from the same direction, in order to increase landing rate. Likewise, a computer calculated order could deviate from the first-come-first-served principle to increase landing rate. The computer calculated order would have the advantage that the new order would be available much earlier, and the re-ordering process could begin a long way before the approach sequencing area. However it might not be practicable for controllers routinely to re-order traffic in en-route airspace without the use of computer assisted speed control.

In order to determine positions in the landing sequence, it is necessary to establish for each aircraft a "Preferred Landing Time". This can be done by predicting the time the aircraft will take to fly from the point where the calculation is being done to the runway, taking account of the normal route to the runway (including ATC altitude constraints), the normal methods of operating the particular aircraft type, and wind conditions. However several aircraft might have the same Preferred Landing Time, so a procedure for establishing Allocated Landing Times is needed.

The difference between the Preferred Landing Time and the Allocated Landing Time gives the amount of advance or delay which must be applied to an aircraft. If the amount of delay is greater than that which can be absorbed by a combination of speed reduction and path-stretching, then it will be absorbed by holding in a stack. Thus calculation of the delay value can give the controller an early indication of which aircraft will have to hold.

2.3.2 Terminal Approach Time

It is convenient to have a term to describe points where aircraft enter the approach sequencing area. We refer to such a point as a Terminal Approach Fix (TAF). For an aircraft which is required to hold, the holding fix is the TAF, but an aircraft which is not required to hold might not overfly a holding fix, and so a term other than "holding fix" is needed for its point of entry into the sequencing area. The Terminal Approach Time (TAT) is the aircraft's planned time of arrival at the TAF. The term Stack Departure Time (SDT) is sometimes used instead of TAT, but we prefer the latter term because it is more meaningful in cases where aircraft are not holding.

Prediction of TAT is of most use for aircraft which are going to hold. Many modern aircraft are equipped with avionics which is capable of arranging the flight in such a way as to leave a hold close to a specified time, if given the

specified time enough in advance. The TAT prediction can be passed to the pilot for this purpose. When not holding, knowledge of the TAT may still be useful to the pilot, and this will be especially so in the future for aircraft equipped with Time Navigation Systems.

Once the Allocated Landing Time has been calculated, determination of the TAT is straightforward. The time planned for traversal of the approach sequencing area is subtracted from the Allocated Landing Time. The planned traversal time is the time predicted for the aircraft to fly a standard path through the sequencing area, plus any delay to be absorbed in the sequencing area.

Landing time allocation, delay determination, and TAT prediction have all been incorporated into a single program, the Landing Order Calculator (LOC).

2.3.3 Speed Control

So far the assumption has been made that controllers can get traffic into the order advised by the computer, and to meet the advised TATs, without computer assistance. As the TATs are known some time in advance, speed variation is a possible mechanism for use in meeting them. However the question of how much speed change will produce how much delay at the TAF is not one which can easily be judged by eye. This is especially so because speed instructions are given in terms of Calibrated Air Speed (CAS), and a CAS value translates into different True Air Speed (TAS) values (and thus different ground speed values) at different altitudes. The calculation of descent speeds which will cause aircraft to arrive at the TAF at the planned times is an obvious candidate for further computer assistance.

Computer-assisted speed control brings the following advantages:-

- Some delay (typically as much as 2.5 minutes) can be absorbed by speed reduction between top of descent and the TAF. This saves fuel and keeps traffic flowing well by reducing the number of occasions on which a holding stack is required. It is sometimes claimed that use of speed control can eliminate altogether the need for holding under normal weather conditions, but this is clearly not the case while a five minute mean delay is used as the runway scheduling criterion.
- Aircraft can be advanced by speed increase as well as being delayed by speed decrease. If queues form because runway time is lost during gaps in the incoming stream, as argued in section 2.1 above, these gaps can be predicted and at least partially filled by advancing some aircraft. The process of advancing aircraft in this way is referred to as giving them "negative delay". The process will have a significant effect on the statistics of the runway queueing process. The "negative delay" effect is probably the most significant benefit from computer-assisted speed control. It has been explored in a discrete event simulation study [7]. The results of the study are summarised in Appendix A.
- Use of speed control to meet planned landing times has the effect of separating traffic as it nears the approach sequencing area. This tends to reduce the number of potential conflicts which controllers must deal with, as

demonstrated by Attwooll [8] and tends to simplify the approach sequencing problem.

- As mentioned above, the allocated landing order may differ from the first-come-first-served order so as to increase landing rate, and with computer calculation of landing order the re-ordering process can begin much sooner. With computer-assisted speed control controllers are given a tool with which to achieve the re-ordering before the traffic reaches the approach sequencing area.

The problem of accurately calculating descent speeds for speed control purposes is not a trivial one. Given a route from a start point A to an end point B consisting of a number of straight legs joined by curved portions, and including altitude and speed changes, all through a three-dimensional wind field, and given the fact that aircraft normally operate at constant CAS, there is no direct way of calculating the CAS value which will cause the aircraft to take a specified time to fly from A to B. The best that can be done is to calculate the times which result from a set of different CAS values and use some form of interpolation procedure to find the CAS which will give the required time. A study has been made of such interpolation procedures [9].

2.4 The Time Horizon Principle

ATC authorities have an obligation to be equally fair to all classes of air traffic, and not to impose on any one class delays which on average are greater than those imposed on other classes of traffic. If a computer program is unfair, then it will be systematically unfair, and that will compound its failing. It is surprisingly easy to construct a landing time allocation procedure which is unfair.

The fairness of a landing time allocation procedure is very sensitive to the precise point in each flight where the allocation is done. To illustrate this point, consider a situation where there are two sets of aircraft. For one set assume that each aircraft has a landing time allocated when it is 30 minutes before its Preferred Landing Time, and for the other set assume that each aircraft has a landing time allocated when it is 20 minutes before its Preferred Landing Time. When allocating landing times for members of the 20-minute set, it will usually be found that most of the desirable landing times close to the Preferred Landing Time are already allocated to members of the 30-minute set, whereas when allocating times to members of the 30-minute set, most of the desirable landing times will be available. Therefore on average, members of the 20-minute set will suffer more delay than members of the 30-minute set.

The only way to be completely fair is to allocate a landing time to each aircraft when it is a fixed time interval before its Preferred Landing Time, and to use the same fixed time interval for all aircraft. We refer to this method of allocation as a "Time Horizon" method, and refer to an aircraft "crossing the Time Horizon" when it passes the point in its flight which is this fixed time before its Preferred Landing Time. Schemes which allocate landing times at some fixed distance out from the runway are unfair because the fixed distance corresponds to different times for different aircraft types. Schemes where aircraft arriving from different directions have landing times allocated at different distances out (because of

international boundaries or ATC sector configurations) are also unfair. A discrete event simulation study has been undertaken to quantify the unfairness in terms of mean delay caused by various sorts of deviation from the time horizon principle [10]. The results of this study are summarised in Appendix B.

The time horizon method was devised to achieve fairness, but it brings another important benefit. When using the method, each landing time allocated is always later than all times previously allocated. With other methods this is not necessarily the case. When not using the time horizon method a situation can arise where a gap of unallocated time is left between two Allocated Landing Times, and filled later. Because the allocation algorithm does not know what aircraft types the gap will eventually be filled with, it does not know the best sized gap to leave. Thus it might be necessary to alter the allocated landing time at the end of the gap later when the gap is filled. Such revision would significantly increase the complexity of the whole process, and could sometimes cause new instructions to be sent to aircraft, which would increase the workload for both pilot and controller. The time horizon method avoids these difficulties.

The time horizon method as described above only works in a first-come-first-served allocation procedure, and it must be modified for use in a situation where the landing sequence is re-ordered to increase landing rate. The obvious modification is to have two time horizons, an inner one and an outer one. Then aircraft become candidates for re-ordering when they cross the outer time horizon, and they have their landing times firmly allocated when they cross the inner one.

2.5 Trajectory Prediction

In order to do the calculations necessary for landing order allocation and speed control, it is necessary to predict the time which an aircraft will take to fly from some point on its route to the runway. Trajectory Prediction is a vital part of many applications of computers in ATC. The prediction must take account of the route to the runway in three dimensions (including ATC altitude constraints), wind conditions at all points along the route, and the characteristics of the particular aircraft type. The prediction is complicated by the fact that aircraft normally fly at constant CAS or constant Mach number, and both constant CAS and constant Mach number give rise to a TAS value (and thus a ground speed value) which varies with altitude. The prediction is further complicated by the fact that aircraft do not behave in a simple linear fashion: the rate of change of height in descent, and the rate of change of speed in acceleration and deceleration must be modelled separately for each aircraft type, or at least for each group of similar types.

The time prediction process is basically a summation or integration of small time increments along the route. At each point the CAS or Mach number is known, as is the altitude, so the TAS can be calculated. At each point values are known for wind speed and direction (though in reality these quantities might not be known very accurately), so ground speed can be calculated. The distance travelled during the time increment can thus be determined. The change in altitude and/or speed can be found from an aircraft performance model. The conditions are then known for the beginning of the next time increment. It can be seen that the

process of predicting a trajectory is in fact the process of solving numerically a set of differential equations with initial values specified.

2.6 The Controller/Computer Relationship

Computer-based tools for arrivals flow management cannot be implemented in isolation from the rest of the air traffic control system. In the vicinity of an airport there are departing and overflying aircraft as well as arriving aircraft. Even for those controllers who are primarily concerned with arrivals, maintaining safe separation is a higher priority task than smoothing the inbound stream. Because controllers are involved in things other than arrivals flow management, it is necessary to consider their relationship with their computer-based tools in some detail.

A whole spectrum of possible controller/computer relationships exists between the following two extremes:-

- *Controller-driven Relationship.*

Computer-based tools are used by controllers if and when they consider it appropriate to use them. For the rest of the time the tools do not alter the controllers' tasks in any way whatsoever.

- *Computer-driven Relationship.*

The computer is in control of everything, and uses the controller as a means of getting instructions to aircraft.

At the present time the only viable kind of relationship is the controller-driven one. There are many reasons for this. The computer-driven relationship requires that the computer program should address the total ATC task, not just the arrivals flow management component of it. Software technology has not yet developed to the point where large programs can be written with enough coverage of all possible circumstances, or enough correctness for such a task. There is a difficulty with relationships which are near but not actually at the computer-driven extreme. In order that controllers can make control decisions on the few occasions when the computer cannot make them, they must have a full and clear perception of the current traffic situation. It is difficult to see how they can have such a perception without active involvement in the control process. Even when the intention is to have a controller-driven relationship, care is needed to avoid accidentally reducing the controllers' involvement in the control process by having them relay instructions from computer to aircraft at a specified time.

The effect of the controller-driven relationship is most evident when considering the time-keeping accuracies being aimed for. When an aircraft trajectory is predicted forward from the time horizon to the runway there are many uncertain factors - wind conditions aloft, exactly how the crew will fly the aircraft, the time taken to react to control instructions, the possible need for conflict-avoiding manoeuvres - and so errors in the calculation are inevitable. Thus most aircraft will not arrive at the TAF at exactly the planned times. There are two ways of dealing with this fact:-

- *The Tightly-Planned Method.*

The aim is to deliver each aircraft to the TAF at a time which is as close as possible to the planned time. This method has the potential to derive the maximum benefit from computer assistance in terms of increasing movement rates and reducing delays and fuel consumption. The method also has the greatest potential to interfere with the rest of the controller's activity.

- *The Loosely-Planned Method.*

The aim is to minimise interference with the rest of the controller's activity, even though this may mean sacrificing time accuracy at the TAF. Although this method does not derive the maximum possible benefit from computer assistance, it can derive a substantial proportion of the maximum benefit by smoothing out the worst of the peaks and troughs in the inbound flow.

Aircraft can be delivered to the TAF closer to the planned times by monitoring their progress and revising their speeds as necessary. The more frequent the revision, the greater the obtainable time accuracy. However the more frequent the revision, the greater the controller workload. Increased workload tends to reduce the number of aircraft which can be handled, and thus loses the benefit of the computer assistance. Also more frequent revision requires prompting of the controller by the computer, which tends towards the computer-driven relationship.

It is possible to remove the uncertainty in time prediction caused by lack of knowledge of how long the crew and aircraft will take to react to control instructions. The uncertainty can be removed by modifying standard ATC instructions to include the point (e.g. beacon and DME distance) where the instruction is to be implemented. This method was tried by the TCSDG in some work which pre-dates that reported here. The method had the effect of moving too far toward the computer-driven relationship. It required the controller to relay computer-generated instructions to the aircraft. It required the computer to prompt the controller to ensure that instructions were sent in time. Because some information had been added to standard instructions, and because controllers were relaying computer-generated instructions rather than creating instructions for themselves, it was more difficult for them to visualise the consequences of the instructions.

The method adopted in the tools described here was the Loosely-Planned method. If a large majority of aircraft could be delivered to the TAF within 30 seconds of the planned time, this was considered good enough to smooth out the worst of the peaks and troughs. Most aircraft could achieve this with a single speed instruction, but a few needed a second one. Using this method, there was no need for the computer to drive controllers in any way. Controllers could treat the computer output as advice which improved traffic flow when taken, but which did not get in their way when not taken.

3 LANDING ORDER CALCULATOR

3.1 Introduction

This section describes the program for calculating landing order and terminal approach times (TATs). The program is known as the Landing Order Calculator (LOC). The description in this section is for a first-come-first-served landing order; the additional functions needed for an optimised landing order are discussed in section 5.

The LOC represents a very simple form of computer assistance which constrains the controller to a minimal extent. Controller inputs are required only very occasionally (e.g. for setting the landing rate to be used by the LOC). The program generally performs landing order and TAT calculations for aircraft when they are between 20 - 25 minutes from touchdown. To do this, it makes use of sophisticated techniques for trajectory prediction which are described in section 6. Landing sequence numbers are displayed both in track labels on the radar displays, and in "strips" on an electronic flight progress display. TATs are displayed in the latter display only.

3.2 Landing Times Allocation

The LOC allocates landing times to arriving aircraft some distance before the airport when most inbound aircraft are at cruise altitude. Prediction of the trajectory of each aircraft is performed based on current aircraft position, height, speed, route and type.

The basic algorithm for the allocation process uses a first-come-first-served (FCFS) order allocation method. This is based on the arrival sequence at a time horizon (a specified time prior to the preferred landing time, typically 25 minutes).

As aircraft become known to the tool, the time horizon crossing time is predicted. When an aircraft crosses the time horizon a position in the landing sequence is allocated. The aim of the allocation process is to allocate TATs to the traffic in such a way as to achieve a known delivery rate into the approach sequencing area. This rate is determined by a notional landing rate value set by the duly authorised controller and a fixed interval between successive aircraft is implied.

For the purpose of allocation one of two states can be specified by the controller as follows:-

- Busy state - When traffic intensity is high the earliest possible landing times are allocated to aircraft in order to maximise runway capacity (see section 2.3.3). This might mean aircraft flying faster than preferred.
- Non-busy state - In a less intense traffic situation the earliest landing time allocated would be the preferred landing time.

In all cases a terminal approach time (TAT - also referred to as stack departure time) is computed. But before describing how this is computed it is necessary to define a control parameter. The variable parameter (referred to as MAXIMUM

ABSORB TIME) defines the amount of delay which it is feasible to absorb within the approach sequencing area; this is a function of the geometry of the available airspace. The value is potentially different for each direction of approach and affects the point in time at which stacking becomes necessary. The value chosen reflects the desired situation under normal operating conditions. Aircraft will be expected to hold if the required delay is greater than MAXIMUM ABSORB TIME and, in any case, if the preceding aircraft was holding and will not have departed the stack by some suitable margin ahead. The method of calculating the TAT is now described.

For an aircraft which is not required to fly a holding pattern, the TAT is the estimated time of arrival at the terminal approach fix derived by subtracting from the preferred landing time the normal time taken to traverse the approach sequencing area (assuming a defined standard routing) plus the value of the amount of delay to be absorbed in the sequencing area. If the aircraft is required to hold (because the required delay exceeds MAXIMUM ABSORB TIME) then the TAT is the time to depart the stack. This is computed as the allocated landing time minus the approach sequencing area normal traversal time.

Delay is either absorbed in the sequencing area or in the stack. If the value of MAXIMUM ABSORB TIME is less than the time taken to fly a minimum stack orbit or less than the difference between the time for a maximum stack orbit and the time for two minimum orbits then there will be some delay values which cannot be exactly achieved. This will lead to a small loss of runway capacity.

3.3 Functional Description of the Allocation Process

3.3.1 Overview

Figure 3.1 shows the names of the main functions of the allocation process together with an indication of the triggering events and associated outputs. The main outputs are the terminal approach times and landing sequence numbers.

3.3.2 Algorithmic Details

The allocation process is driven by a number of external events as follows:-

- New aircraft - an aircraft has "entered the system". This can be expected to occur at or just before the aircraft enters a designated region of airspace, typically 150 - 200 miles before the airport.
- Clock Tick - a regular event occurring typically at 10 second intervals.

- Landing rate change - an interaction by the controller to adjust the planned delivery rate into the approach sequencing area.
- Delete aircraft - radar tracking indicates that an aircraft has reached its final state of approach and is about to land or the aircraft has diverted to another airport.

The main functions which are invoked in response to the above events are now described.

New Aircraft Event

Compute the expected time horizon crossing time for the aircraft by performing a trajectory prediction from the aircraft's current position to the runway assuming a preferred speed profile (see section 3.4).

Clock Tick Event

Check if any aircraft have crossed the time horizon and allocate landing times accordingly. The earliest feasible free time is allocated, assuming a fixed inter-aircraft separation defined by the current landing rate. A "feasible" time is one which can be achieved within the performance limits of the particular type of aircraft. Those aircraft allocated landing times are added to the landing list. Three trajectory predictions are performed at this point for each of these aircraft assuming minimum, maximum and preferred descent speed profiles. This produces terminal approach fix ETAs, top of descent positions and times and Preferred Landing Times for flight at these speeds.

Landing Rate Change Event

A new set of landing times is computed for all aircraft already in the landing list to reflect a new delivery rate into the approach sequencing area and new Terminal Approach Times are issued.

Delete Aircraft Event

Remove aircraft from landing list.

3.4 Use of Trajectory Prediction

An overview of trajectory prediction was given in section 2.5 and a full treatment is given in section 6. What follows indicates how the allocation process uses the prediction functions.

Figure 3.2 illustrates the predicted flight profile for an arriving aircraft from cruise altitude to the runway. In predicting the profile a continuous, idle thrust descent is assumed to be flown using a preferred CAS/Mach number speed schedule specified in the model for a particular aircraft type. All descents are planned to

start as late as possible consistent with performing an idle-thrust descent to be level by the Terminal Approach Fix (or possibly some short distance before). At cruise level aircraft normally fly at constant Mach number which has an equivalent CAS depending on altitude. In order to descend at the demanded CAS a speed change is normally required. This is assumed to happen in one of two ways. If the demanded CAS is greater than the equivalent cruise CAS the descent is assumed to start at constant cruise Mach number with transition to constant CAS when the CAS has risen to the demanded value. If, however, the demanded value is less than the equivalent cruise CAS it is assumed that the aircraft descends until the necessary speed flexibility is available to decelerate to the demanded CAS (possibly by levelling out).

Aircraft may be required to conform to one or more speed limit points on the route, for example at the Terminal Approach Fix. A top of descent point is computed based on the speed of descent to the TAF. This must be computed starting from the TAF and working backwards to enable the aircraft to remain at high altitude as long as possible (for fuel economy reasons). To predict when an aircraft will cross the time horizon a profile based on a preferred descent speed (specified for each aircraft type) is computed. The profile and route between Terminal Approach Fix and runway is assumed to be standard. It is only necessary to perform a fairly crude prediction of this phase of the flight since it is accepted that there will be some variability of control in the approach sequencing area.

3.5 Data Structures

The data structures which the algorithms use are :-

- a. A representation of aircraft routes
- b. A set of aircraft performance models
- c. A table of aircraft state data records
- d. The airport landing lists.

Items a) and b) are assumed to exist in a form suitable for performing prediction of aircraft trajectories and are not described here (see [11] and section 6 of this report).

The essential features of c) and d) are shown in figures 3.3 and 3.4 respectively. Note that the landing list is ordered according to runway time/landing sequence number. Landing sequence numbers run from 1 to 99 and then revert to 1 again. Once allocated to an aircraft the sequence number is fixed (unless changed by the controller).

3.6 Outputs and Required Inputs

This section aims to give an indication of the requirements of the tool and the facilities available. A full description of the experimental environment is given in section 7 with a description of the displays being given in Appendix D.

The main outputs from the tool are referred to as plans. A plan is generated for each arriving aircraft and can be revised when the landing rate is changed. The plan comprises:-

- a. Landing sequence number for the appropriate airport
- b. Terminal Approach Time (TAT)
- c. Predicted Stack Arrival Time (same as TAT if the aircraft is not required to hold)

The TATs are displayed on a flight progress electronic data display (EDD) which contains data lines for a number of inbound aircraft on a sector. For each aircraft the planned landing sequence number is shown on the EDD and also on the radar plot label. Both the TATs and landing sequence numbers are displayed as soon as allocated, i.e. when the aircraft is about 25 minutes from the airport and can be shown on all interested sectors.

The LOC advises the controller of the requirement to hold an aircraft by marking the holding aircraft's data line on the EDD and also on the radar plot label (using a single character) and by displaying TATs on the EDD. (Note that it would be possible to display TATs on the radar picture, but the best way of doing this has not been investigated).

By communicating the TAT to holding aircraft early enough, e.g. 10 - 15 minutes before ETA at the stack it can be expected that most of them will be able to arrange their flight to depart the stack within 30 seconds of the planned TAT [6]. If desired the controller may update the display to indicate that the TAT has been communicated to the aircraft, but this is optional.

There are three other possible inputs to the LOC as follows:-

- a. The landing rate may be altered to reduce or increase the planned rate of delivery into the approach sequencing area. This would probably be adjusted by the approach control team and leads to the display of revised TATs for those aircraft which have not yet departed the stack.
- b. The allocated sequence numbers may be overridden by the controller swapping any two aircraft in the order.
- c. The parameter MAXIMUM ABSORB TIME could be adjusted if prevailing wind conditions or visibility changed significantly.

Note that in the experiments it was not possible for the controller to change item c) interactively.




Figure 3.1 LOC Functions




Figure 3.2 Predicted Flight Profile

Aircraft identity, route, position, current clearances, performance model type	
Preferred Landing Time	Predicted runway time assuming no delay
Allocated Landing Time	Predicted time taking account of other aircraft and landing rate
Stack ETA	Predicted ETA at stack
Terminal Approach Time	
Top of descent time/position	Predicted
Flight delay	Difference between Preferred and Allocated Landing Times
Planned hold	True or false
Landing sequence number	In range 1 to 99
Landing list entry number	Index to landing list

Figure 3.3 Aircraft State Table Structure




Figure 3.4 Landing List Structure

4 SPEED CONTROL ADVISER

4.1 Introduction

The Speed Control Advisor (SCA) includes the functions of the landing order calculator but goes further by computing a speed for each arriving aircraft. The speed is calculated so that the aircraft meets its allocated landing time by absorbing as much delay as possible enroute and increasing speed if necessary. The speeds are displayed to the controller to assist in the production of an orderly stream of traffic inbound to the Terminal Approach Fixes and to reduce the need for holding patterns to be flown. As a result of the allocation process a top of descent time is produced and this is also displayed to the controller.

Section 3 described the landing times allocation process and how Terminal Approach Times are computed. For the SCA the rules for absorbing delay are slightly different from the landing order calculator (LOC) and there are three possible methods:-

- a. absorb total delay by speed control
- b. absorb some delay by speed control and some in the approach sequencing area
- c. use stack orbits in conjunction with speed control

The rules for computing TATs are also different according to the chosen delay absorption combination. The three cases are:-

- a. If total delay can be taken between top of descent and the TAF then

$$TAT = \text{allocated landing time minus standard approach sequencing area traversal time}$$

- b. If a) is not possible and total delay does not exceed that which can be taken by flying at minimum descent speed and by absorbing an amount up to the value given by MAXIMUM ABSORB TIME in the sequencing area then

$$TAT = \text{ETA at TAF assuming minimum descent speed}$$

- c. If neither a) or b) is possible then at least one minimum stack orbit must be flown and

$$TAT = \text{allocated landing time minus standard approach sequencing area traversal time.}$$

There are also some extra actions for the SCA which take place during the landings times allocation process :-

- a. When an aircraft crosses the time horizon, in addition to a position in the landing sequence being allocated, an appropriate speed is computed to enable the aircraft to meet it as far as possible without flying holding patterns.

- b. The use of a busy/non-busy state comes into its own with the SCA. The tool can advise a higher than preferred speed to fill an empty landing slot ahead.
- c. In order for the aircraft to fly the optimal idle-thrust descent using the computed speed it must begin its descent near to the correct point. If descent is started early then time will be spent unnecessarily cruising at low level, if started late then a steeper descent trajectory will be necessary to be level by the Terminal Approach Fix. Both these manoeuvres will result in a performance below the optimum for fuel economy. The time for top of descent is therefore computed and advised to the controller.

The prototype Speed Control Adviser system displays landing sequence numbers both on the radar picture and a colour flight progress EDD at the air traffic control position. The Terminal Approach Times and advised speeds for each aircraft are shown on the EDD. Section 7, together with Appendix D gives a description of this environment.

4.2 Functional Description

The functions described in section 3.3 apply also to the SCA process but there are two additional functions and one new triggering event. The new functions are summarised below:-

- Compute Speed - at the same time as allocating a landing time an appropriate descent speed is chosen to meet the time.
- Update Landing List - in response to a landing rate change request the landing times for each entry in the landing list are adjusted to reflect a new landing rate. In addition to new TATs, new speeds are computed for those aircraft which still have scope for absorbing delay enroute (not implemented - see section 9 for further discussion).

The new triggering event is generated by a request from the controller for a new plan (a new speed). The function to compute a new speed just described is invoked. Since computing the speed is a major part of the SCA the details of how this is done are now presented.

In computing the appropriate calibrated airspeed (CAS) for descent to meet the Allocated Landing Time consideration is given to the required delay and the appropriate combination of delay absorbing mechanisms as described in section 4.1. In the cases where no stacking is planned the speed is chosen to enable the aircraft to arrive at the TAF at the TAT; in the case of combination b) this means minimum speed. When stacking is required the slowest possible speed for descent is computed which is consistent with arriving at the holding fix in time to fly at least a minimum stack orbit. The reason for requesting aircraft to reduce speed enroute as well as flying one or more holding patterns is because of ATC problems - see section 9 for further comment.

To compute the required CAS (target CAS) to meet the Allocated Landing Time (target time) the arrival times flying at minimum, maximum and preferred speeds appropriate to the performance of the type of aircraft are computed using trajectory prediction. Only the constant CAS is chosen, the constant Mach number portion of the descent being flown until transition to this CAS at the transition altitude (i.e. when the CAS has risen to the desired value). To describe the technique used to compute the speed the following notation is introduced:-

$$\begin{array}{ll} V_f = \text{maximum CAS (fast)} & T_f = \text{time at } V_f \\ V_s = \text{minimum CAS (slow)} & T_s = \text{time at } V_s \\ V_p = \text{preferred CAS} & T_p = \text{time at } V_p \\ V_c = \text{chosen CAS} & T_c = \text{time at } V_c \\ V_t = \text{target CAS} & T_t = \text{time at } V_t \text{ (target time)} \end{array}$$

The values (V_s, T_s) , (V_p, T_p) , (V_f, T_f) are used as a three point approximation to the function of speed against time to enable the target CAS to be derived using interpolation. If $V_p = V_f$ or $V_p = V_s$ then a trajectory is predicted at a speed midway between V_s and V_f to give a value for V_p . [9] gives a comparison of interpolation algorithms. The approximation used is given below:-

$$V + a = \frac{k}{(t + b)} \quad (1)$$

Speed (V) is assumed to be approximately inversely proportional to time (t) and a, b and k are constants.

When substituted into (1) the three points (V_s, T_s) , (V_p, T_p) and (V_f, T_f) give the equations:-

$$V_s + a = \frac{k}{(T_s + b)} ; \quad V_p + a = \frac{k}{(T_p + b)} ; \quad V_f + a = \frac{k}{(T_f + b)}$$

These equations can be solved to give the following values for b, a and k .

$$\begin{aligned} b &= \frac{V_f T_f (T_s - T_p) + V_p T_p (T_f - T_s) + V_s T_s (T_p - T_f)}{V_f (T_p - T_s) + V_p (T_s - T_f) + V_s (T_f - T_p)} \\ a &= \frac{V_p (b + T_p) - V_f (b + T_f)}{(T_f - T_s)} \\ k &= (V_f + a)(T_f + b) \end{aligned}$$

The values b, a and k are then substituted into (1) to give the approximation of speed against time. The three points (V_f, T_f) , (V_p, T_p) and (V_s, T_s) must be distinct in order to provide three solvable simultaneous equations.

Normally, the approximation gives the target CAS directly, assuming target time is required to within five seconds. In some cases, for example when target speed lies away from V_s, V_p and V_f , two or three iterations may be necessary. We define (V_L, T_L) , (V_M, T_M) and (V_U, T_U) as the lower, middle and upper speed value pairs which are the three points for the approximation at any iteration.

Initially,

$$\begin{aligned}(V_L, T_L) &= (V_s, T_s) \\ (V_M, T_M) &= (V_p, T_p) \\ (V_U, T_U) &= (V_f, T_f)\end{aligned}$$

After computing an approximation the points are set as follows:-

a. For $V_c > V_M$

$$\begin{aligned}(V_L, T_L) &= (V_M, T_M) \\ (V_M, T_M) &= (V_c, T_c) \\ (V_U, T_U) &\text{ unchanged}\end{aligned}$$

b. For $V_c \leq V_M$

$$\begin{aligned}(V_L, T_L) &\text{ unchanged} \\ (V_M, T_M) &= (V_c, T_c) \\ (V_U, T_U) &= (V_M, T_M)\end{aligned}$$

A new approximation is computed until target time is achieved to the required tolerance.

It should be noted that trajectory prediction is potentially expensive in terms of processing power requirements and the number of times the trajectory is computed needs to be kept to a minimum. This means it is important to minimise the number of iterations which the interpolation requires.

It is assumed that, generally, only a single speed control instruction will be issued to the aircraft and that this will be implemented by the pilot as part of the descent speed schedule. Note that the aircraft may not be able to implement the demanded CAS immediately at the beginning of descent due to Mach limit effects (see section 3.4).

From the controller's viewpoint it is possible to issue the speed instruction at any point between it being displayed and the aircraft being in a position to implement it. A recognised radio telephone phrasing would need to be agreed such as "Your speed for descent is 260 knots" if the instruction is to be given ahead of implementation.

4.3 SCA Outputs and Required Inputs

In addition to the outputs of the LOC the SCA includes:-

- a. CAS for descent
- b. Predicted time for top of descent

The additional inputs which the SCA enables are:-

- a. When a landing rate change is requested in addition to producing revised TATs the SCA recomputes speeds for all aircraft that have not yet passed the TAF.

- b. The busy-state of the system may be set to "busy" or "non-busy" (see section 3.2). In the busy state the controller is given assistance in achieving earlier than preferred landing times via the speed control advice.
- c. A request for computation of a new speed may be made for any aircraft. A trajectory prediction will be performed from the aircraft's current position to deduce this.
- d. A check can be requested for an aircraft to ascertain whether it can still achieve its Allocated Landing Time. This input could be combined with c).

5 Optimising the Landing Order

5.1 Introduction

The descriptions in sections 3 and 4 of the LOC and SCA assumed a first-come-first-served order across the time horizon in the process of allocating landing times to aircraft. This section describes another method, applicable to both tools, which attempts to improve runway utilisation by arranging aircraft in an order which minimises the inter-arrival spacing for a stream of aircraft taking account of the requirements for separation due to turbulent wake.

In current practice, approach controllers already reorder aircraft to optimise, as far as they are able, final runway spacings. However, by using the computerised tools this process can begin much earlier and thus reduce the workload for the approach controllers.

5.2 Optimised Landing Times Allocation

The Optimised Order landing times allocation method (OptO) attempts to produce a landing order which is the optimum based on the separation rules for turbulent wake vortices and uses two time horizons, an inner and an outer one typically at 25 and 30 minutes before the Preferred Landing Time, respectively. Any aircraft between these time horizons is eligible for inclusion in a reordering process which may result in a deviation from the first-come-first-served (FCFS) order. The sequence may also depend on the desirability of changing the order of two aircraft sharing a common route.

In this method the position in the landing sequence is allocated when an aircraft crosses the inner time horizon. The aim of the planning process is still to deliver the traffic into the approach sequencing area at a known rate, but this time the notional landing rate value set by the controller implies an inter-arrival interval which is an average figure that does not cause the declared delivery rate for any hour to be exceeded.

5.3 Functional Description of the Optimised Allocation Process

The OptO method is a two-stage process. As an aircraft crosses the outer time horizon it becomes eligible for inclusion in the optimisation process and ceases to be eligible when it has been allocated a sequence number. When an aircraft crosses the inner horizon the optimisation process is activated which fixes the position of the aircraft in an optimised sequence. At this point the only landing times which are firmly allocated are those for aircraft which have crossed the inner time horizon and any of the eligible aircraft (those between the time horizons) which will be allocated earlier sequence numbers than the aircraft which triggered the allocation process. (It is necessary to fix these earlier aircraft to avoid the problem of having to later fill gaps of arbitrary size as described in section 2.4 The turbulent wake vortex separation rules define the minimum spacing between landing times, but the separations can be increased to achieve the currently set landing rate.

Section 3.3.2 described the triggering events for the basic allocation process. The OptO method affects the actions taken following the *new aircraft* and *clock tick* events with the latter representing the most significant changes. The main functions are shown in figure 5.1 and described below.

New Aircraft Event

Compute the expected times at both the inner and outer time horizons using trajectory prediction.

Clock Tick Event

Add any aircraft which have crossed the outer time horizon to the list of aircraft that are eligible for inclusion in the landing sequence order optimisation process (these are referred to as the "eligible" aircraft). If any one of the "eligible" aircraft has crossed the inner time horizon compute a proposed set of landing times for *all* the "eligible" ones. Allocate firm times to those which have crossed the inner horizon and to any which, although they have not yet reached this horizon, will be allocated earlier times than those firmly allocated in the proposed set. Those aircraft allocated firm landing times are added to the landing list. Three trajectory predictions are performed at this point for each aircraft as for the FCFS order method.

To produce the optimal set of proposed landing times two attempts are made. Firstly, a permutation which avoids any aircraft stacking and with fewest deviations from the FCFS order is sought. If there is no such feasible ordering then a set which gives minimum total delay is proposed. In order to produce a proposed ordering the following is required. For each possible permutation of aircraft in the "eligible" list (excluding those with more than two order position shifts) :-

- i. *Check that the permutation is valid.* There are two options here. Firstly, if re-ordering of aircraft which are flying common routings is not allowed (due to overtaking control problems - see section 9) then permutations which interchange two such aircraft are eliminated. Secondly, if the traffic state is specified as "non-busy" then any permutations which would require aircraft to increase speed are also eliminated.
- ii. *Try the permutation for feasibility based on the performance limits of each aircraft.* Compute the runway time for the last aircraft in the order. Compute the order displacement index which is a value derived from the total disturbance from the FCFS order (the higher the index the greater the disturbance).

The landing separations used to allocate runway times for a permutation of aircraft are based on aircraft types and wake vortex separation rules. The time separations are derived by transforming the distances specified in the ATC operations manual [18] into times by assuming a notional average final approach

speed of 140 knots. A time separation of 78 seconds is used where 3nm separation is assumed. These separations are adjusted on a proportional basis to ensure a flow of traffic into the approach sequencing area which does not exceed the currently set airport landing rate. Note that these figures represent a simplification and in practice the times may need to be a function of aircraft type to achieve the required distance separation for all aircraft.

Assuming there is a feasible permutation with no aircraft holding, select the permutation which gives the earliest runway time for the last aircraft in the permutation. If two or more permutations give similar times then select the one with the lowest order displacement index.

*Triggering
Event*




Figure 5.1 OptO Method Functions

6 TRAJECTORY PREDICTION

6.1 Introduction

The tasks of identifying a landing order and allocating landing times need to know how long each aircraft will take to reach the runway, and what arrival time flexibility is possible. Making the necessary predictions involves detailed calculations based on simulating the remainder of the aircraft's trajectory to touchdown. A range of simulated flight paths may need to be examined, to explore the full potential of the aircraft's speed envelope. This section describes the software that supports the prediction functions. The software is implemented as a library of procedures and state variables which are called by the LOC and the SCA. An illustration of how the functions are used by the tools is given in Appendix C.

6.2 Overview

Prediction is based on the assumption that the trajectory to touchdown will be confined laterally to follow a known route, specified as a series of waypoints. However, the start of the predicted trajectory can be off route if appropriate. (An aircraft might be off route following radar vectors at the time that trajectory prediction is called for). The prediction algorithm maintains a set of state variables which identify the current position and flight conditions of a notional aircraft flying along the route. Before prediction begins the state variables must be initialised to reflect the required start conditions. Information such as aircraft type, starting position, altitude and airspeed as well as intended route must be supplied. Progress along the route is determined by "flight manoeuvre" directives issued to the algorithm. A directive can specify the prediction of descending flight, or a speed change, or a defined period of level flight.

Predicting flight progress along a route involves representing both straight line flight (when tracking steadily towards a waypoint), and a curved flight path when turning from one waypoint to intercept the trackline towards the next. A single flight manoeuvre may be completely contained within one straight section, or within a curved section, but may also span several sections. The detailed mapping of a manoeuvre on to the route is not known to the calling program. After each manoeuvre is completed the state variables can be examined to establish where on the route the notional aircraft has reached. Other summary data such as total lapsed time and distance covered since the start of the prediction is also available. In practice the state variables are divided into two sections, one section describes a forward prediction state, the other section describes the status of a backward prediction state.

6.3 Forward and Backward Prediction

The purpose of forward prediction is to calculate the final position, the distance travelled, and the time taken by an aircraft in completing any specified phase of flight. Calls for speed change, descent, or simply level flight are examples of manoeuvres which may occur within such a phase. In all cases the forward

prediction process navigates the notional aircraft along the route towards the end of the route (assuming a route to have a defined start and end). However, when planning a trajectory, it is frequently necessary to arrange to complete a manoeuvre at a pre-determined position or time. A good example is the problem of calculating a top of descent point so as to arrive at bottom of descent, say ten miles before a TAF. A direct solution is in general not possible when using forward prediction and an iterative approach would be required.

Backward prediction starts with the position of an aircraft on completing a flight manoeuvre, such as a descent, level flight or a speed change, and predicts back to the start of the manoeuvre identifying the lapsed time, the distance travelled, and the position immediately before the start of the change.

Thus, the state variables describing the forward processing state show the position and flight conditions of a notional aircraft *after* completing forward flight manoeuvres. The backward processing state defines the position and flight conditions of a notional aircraft *before* beginning the commanded flight manoeuvres designed to be completed on reaching the previously specified end point. Since the forward and backward prediction state variables define two independent states a trajectory design strategy can be employed that uses both forward and backward prediction methods.

6.4 Aircraft Performance Models

Aircraft performance related data is available to represent a range of aircraft types. Not all types are uniquely represented. Because of limited data availability some performance details are applied to several aircraft types. Five performance aspects are modelled three of which, idle thrust descent, acceleration and deceleration are based on the PARZOC method developed by Benoit [14].

6.4.1 Flight Idle Descent Performance

Instantaneous flight idle descent rate is modelled as a function of descent CAS and current altitude:-

$$\text{Descent rate}(v, h) = \frac{1}{J + Mv + Pv^2 + 2h(K + Nv + Qv^2)}$$

where:-

$$\begin{aligned} v &= \text{Descent CAS} \\ h &= \text{Cruise Altitude} \end{aligned}$$

and J, M, P, K, N and Q are type dependent coefficients.

(Benoit models the time taken for an aircraft to descend from cruise altitude to an altitude of 5000 feet as a second degree polynomial in descent speed and cruise altitude. This relationship however, is inconvenient when calculating the effect of wind on ground speed, and hence on the distance travelled throughout the descent. The more useful descent rate characteristic given above is derived from the Benoit model).

Note that the coefficients are optimised for the altitude range between cruising level and 5000 feet and for a speed range not requiring flaps to be deployed. The

model is applied below 5000 feet altitude and gives acceptable results for speeds not requiring flap deployment. Gross Errors are avoided in the flap deployment speed range by augmenting the descent model with an approach model (see section 6.4.5).

A model for descent rate when descending at constant Mach number is derived from the constant CAS descent rate at the same airspeed. The relevant expression is:-

$$\frac{\text{Descent rate}(Mach = const)}{\text{Descent rate}(CAS = const)} = 0.9941 + 0.0475Mach + 0.612Mach^2$$

6.4.2 Deceleration Characteristics

Deceleration is modelled as a function of current speed and altitude. This too is derived from the PARZOC work by Benoit. The relationship is:-

$$\text{Deceleration}(v, h) = \frac{1}{0.6(L + Mh + Nh^2 + 2v(O + Ph + Qh^2))}$$

where:-

$$\begin{aligned} v &= \text{Descent CAS} \\ h &= \text{Cruise Altitude} \end{aligned}$$

L, M, N, O, P and Q are type dependent coefficients.

6.4.3 Acceleration Characteristics

The form of the acceleration model is the same as for the deceleration model requiring a further six coefficients per aircraft modelled.

6.4.4 Operational Envelope Limit Parameters

This part of the model includes the following values:-

- Minimum CAS with flaps retracted
- Maximum operating CAS
- Preferred descent CAS
- Normal turn bank angle
- Expedited turn bank angle
- Mach limit
- Expedited descent rate (relative to flight idle descent rate)

6.4.5 Approach Speed and Flap Deployment Data

The effect of flap and undercarriage deployment on flight idle descent and deceleration performance is represented by a series of coefficient sets. One set of four coefficients is needed to define each flap configuration change. The four parameters are:-

- a. altitude at which deployment occurs
- b. aircraft speed after deployment
- c. descent rate resulting from deployment (relative to descent rate given by model in section 6.4.1).
- d. level flight deceleration resulting from deployment (relative to deceleration given by model in section 6.4.2).

6.5 Height Change Modes

Both forward and backward descent prediction can be flexibly specified. The descent can be achieved under any of three specified speed regimes. These are:-

- a. constant CAS throughout
- b. constant Mach throughout
- c. constant Mach followed by constant CAS

The way in which the height loss will be achieved can be specified independently of the selected speed mode. Four options are available:-

- a. flight idle descent rate
- b. expedited descent rate
- c. specified constant slope (e.g ILS glideslope)
- d. specified constant rate

When a constant slope or a constant descent rate is specified the implementation is modified where necessary so as not to exceed the expedited descent rate.

6.6 Airspeed Constraints

The prediction process takes note of operational limit CAS and Mach values and will modify its activity to avoid exceeding limits. For example, when predicting a descent trajectory from flight level 270 down to flight level 80 at Mach 0.70 the CAS would increase as the descent progressed reaching the typical max operating CAS of say 340 kts at an altitude of around 17000 feet. Below that altitude the descent prediction would be continued, not at the demanded Mach number but, at max operating CAS.

6.7 Turns

The prediction of turning flight assumes that a constant angle of bank is applied throughout the heading change. Unless specified to the contrary both forward and backward prediction processes will not overfly route waypoints. Calculations are made in advance to establish the distance prior to the waypoint at which an anticipated turn should start. Starting a turn towards the next waypoint at this point results in the correct interception of the next trackline. The pre-waypoint distance is influenced by windspeed, true airspeed and bank angle. In circumstances when these parameters may be continually changing the turn calculations are refined as the pre-waypoint distance is approached in order to maintain accuracy.

6.8 Integration

Numerical methods are used to integrate instantaneous vertical and horizontal velocity vectors in order to predict future position. Since the vector values rarely remain constant throughout a trajectory it is necessary to decompose the integration problem into a succession of integration calculations, each calculation predicting ahead through a short time interval from the position determined by the previous calculation. The total predicted displacement and the time taken is calculated from the summation of the individual time intervals. In order to minimise the number of time steps, and hence the computation effort required, the time step size is made as large as possible consistent with the need to maintain accuracy. When the velocity vectors are not changing or are changing only slowly a large time step can be used. Non-linear rates of change demand a much smaller time step size. The size of time step is chosen so that the following simplifying assumptions can be made without introducing significant errors:-

- a. Acceleration and deceleration (ie rates of change of CAS) are assumed constant throughout the time interval.
- b. When executing an idle thrust descent the rate of descent is assumed constant throughout the time interval.
- c. The distance travelled during the time interval is calculated from the average of initial and final groundspeed.

6.9 The Approach Phase

Both forward and backward prediction processes take account of the stepped speed reductions and drag increases associated with flap and landing gear deployment when descending on approach to the runway. However, it is acknowledged that for one set of flight conditions the modelling technique used is inadequate. This concerns the situation where speed reduction occurs when an aircraft is in descent and where the aircraft is constrained to continue the descent throughout the deceleration. (When following the ILS glideslope, for example). The deceleration model together with the appropriate part of the flap deployment model apply only for level flight deceleration. Under descent conditions a much

reduced deceleration would be observed in practice. In the work described in this report this modelling shortcoming is not apparent, however, since the predictive software and the air traffic simulator share the same aircraft model.

7 EXPERIMENTAL ENVIRONMENT

7.1 Introduction

In order to conduct air traffic control experiments with real live air traffic controllers it is necessary to simulate a large proportion of the controllers' normal working environment, that is, the environment found in an air traffic control centre. There is a conflict of interests between achieving more and more realism on the one hand, and keeping the total experimental system as small as possible on the other. Achieving insufficient realism lays the experimenter open to the charge that the tasks being performed by the controllers in the simulated environment are so different from those performed in a real operational environment that no useful conclusions can be drawn from the experiments. Achieving more realism increases the cost of producing and manning the system, and makes it more difficult to change the system when the need arises. Ease of change is a most important property of experimental systems.

The TCSDG experiments were based on the route structure and ATC procedures which formed part of the CCF (see section 1). Normally, two ATC sectors were manned and represented in detail. Traffic was simulated in other sectors, and a very basic level of air traffic control was provided in these by a program known as "The Automatic Controller". Traffic arriving at and departing from all the major London airports was simulated, as was overflying traffic. The complete set of components which provided the working environment for the controllers and experimental prototype tools is as follows:-

a. *Skeleton Control Centre Software.*

This provides the software environment for the experimental prototype arrivals management tools. It also provides a database from which radar displays and electronic flight progress displays can be generated.

b. *Display and input hardware.*

This provides the means of displaying radar data, flight progress data, and computer-generated plans to the controller. It also provides the controller with a means of making inputs to the system.

c. *Air Traffic Simulator.*

This program simulates the flights of many aircraft, and generates a stream of radar plots which represents their changing positions. It accepts control inputs for the simulated aircraft.

d. *Traffic Sample Generator.*

This program generates realistic but pseudo-random samples of traffic for use in the experiments. Its output is used both for driving the Air Traffic Simulator and for providing "flight plan" input to the Skeleton Control Centre Software.

e. *Automatic Controller.*

This program provides very rudimentary ATC functions for unmanned sectors. It has the ability to accept traffic handed over to it, and to initiate hand-over of traffic to a manned sector at appropriate points.

Item b. of the above list is discussed in section 8. All other items are discussed in the remainder of this section. The relationship between these components is shown in figure 7.1.

7.2 The Skeleton Control Centre Software (SCCS)

The SCCS consists of a database of information needed by the controllers or by the experimental prototype tools, and a set of processes which maintain and interact with the database. The main subdivisions of data in the database are:-

- *Aircraft State Data.*

An Aircraft State Record is maintained for each aircraft known to the SCCS. Each record contains aircraft type and callsign, destination and route, latest radar position and altitude, current controlling ATC sector, and current ATC clearance. Information generated by the experimental tools, such as TAT and recommended speed, is also included in the Aircraft State Record. Aircraft State Data is the most important and central part of the SCCS.

- *Route Data.*

This contains a description of all ATC routes known to the SCCS. Each route is essentially a list of waypoints, and a waypoint is defined by a range and bearing from a radio navigation aid at a known position. As well as geographical information, the Route Data contains information about certain ATC procedures. For example it contains locations of inter-sector boundaries, details of holding stacks, and speed and altitude constraints at various points. There can be several variants on each route to allow for the difference between easterly and westerly runway operations, and the different requirements of high-performance and low-performance aircraft types.

- *Aircraft Performance Data.*

This contains for each aircraft type known to the SCCS details of performance parameters, such as maximum and minimum speeds, normal descent speeds, and maximum cruising altitudes. It also contains numerical coefficients for the formulae which the trajectory prediction functions use to model aircraft performance.

- *Experimental Environment Data.*

This is a miscellaneous collection of data on topics such as wind conditions aloft, runway operating directions, radio frequencies to be used by each manned sector, display configurations and formats, and so on.

An Aircraft State Record is initialised by a Flight Plan Process some minutes before a new flight comes into the system. Thereafter the Aircraft State Record is updated by a simple Radar Tracking Process (which obtains radar plots from the Air Traffic Simulator), by a process which obtains input from dialogues with controllers, and by the experimental tools. A graphics display process obtains information from the Aircraft State Data for generating synthetic radar pictures. A tabular display process uses information from the Aircraft State Data for output to electronic flight progress displays. These processes are shown in figure 7.2.

The remaining three types of data listed above - Route Data, Aircraft Performance Data and Experimental Environment Data - are loaded into the SCCS at the beginning of each run, and are constant throughout the run. These data contain much information about assumed ATC procedures and experimental conditions, and it is important that they can be read easily by experimenters and controllers, and can be changed easily from one run to the next. To achieve this end, they are held in readable text form on three disk files. They can be edited and printed in the same way as any other text files. The Route Data can be thought of as being represented in a simple Routes Description Language [11], and similarly for the other types of data. Each type of data is converted from readable text form to an internal form suitable for computation when loaded into the SCCS at the beginning of a run.

7.3 The Air Traffic Simulator

The Air Traffic Simulator [12] simulates the flights of many aircraft simultaneously. It obtains details of each flight to be simulated - start time, aircraft type and callsign, route and flight level - from the traffic sample file. In the absence of any control instructions, each simulated aircraft flies along its route at its initial speed and flight level, and ceases to exist at the end of the route. Control instructions can be given to change speed or flight level, to leave the route and fly a specified heading, to fly a holding pattern, to change radio frequency, and so on. Simulated aircraft can spontaneously report various conditions of interest to ATC, e.g. reporting "On Frequency".

There are two mechanisms for sending instructions to simulated aircraft and receiving reports from them. For an aircraft under the control of a manned ATC sector, the controller communicates by voice with a person who plays the role of a pilot. Such a person is known variously as a "Pseudo-Pilot", an "Aircraft Control Operator", or a "Blip Driver". When the controller gives an instruction to the pseudo-pilot, the latter enters it via an Aircraft Control Terminal. When an aircraft reports a condition to ATC, the details are displayed on the Aircraft Control Terminal to the pseudo-pilot, who in turn communicates them by voice to the controller. Normal ATC radio telephony procedures are used in the voice

communication between controller and pseudo-pilot. In a simulation exercise there is one pseudo-pilot for each manned ATC sector. A pseudo-pilot might have to "fly" as many as fifteen aircraft at the same time, and so it is important that the interface with the Aircraft Control Terminal should facilitate fast easy interaction. For an aircraft which is not being controlled by a manned sector, control instructions are given, and reports are received via a data channel within the computer. This mechanism allows the Automatic Controller program to generate control instructions and receive reports. However, apart from the difference in representation of instructions and messages, the two mechanisms are equivalent.

The Air Traffic Simulator models the flights of the simulated aircraft in some detail. All translational motion takes account of wind conditions. Climb performance is modelled by the EROCOA method devised by Benoit [13]. Performance in descent and speed change is modelled by a method based on the PARZOC method developed by Benoit [14], but modified substantially to take full account of wind conditions, to model descent at constant Mach number, and to represent the approach phase of flight [15]. Wind compensated holding patterns are simulated properly. Some attempt is made to introduce a degree of variability into each flight so that the traffic does not behave in a way which is too predictable [12].

The Air Traffic Simulator needs to use the same Route Data and Aircraft Performance Data as the SCCS uses. The same software modules are used for loading this data and converting it from readable text form to internal form, but the simulator has its own distinct copy of the data. This is to avoid any accidental sharing of data between the simulator and the SCCS which could give the SCCS access to data which it would not have in the real world.

7.4 The Traffic Sample Generator

A central ingredient in any ATC simulation exercise is the Traffic Sample. The Traffic Sample is simply the set of flights to be simulated, complete with times, aircraft types and callsigns, route information and the like. Each hour of exercise may require a sample of more than a hundred flights. The sample used in a particular exercise determines the quantity and type of work which the controllers will have to do, and thus has a significant effect on the value of the simulation. A sample cannot be used too many times with the same controllers because they soon get to know it. Controllers can be quite sensitive to what is, or is not, a realistic sample.

Traffic Samples can be generated by hand, but this is a very lengthy and tedious task. As well as generating the details of the flights to be simulated, it is necessary to control the scheduled arrival rates at airports and the loadings on various routes. Samples can be generated by observing the traffic which actually flies on a particular day, and later using the details in simulation exercises. A range of samples can be generated by selectively combining traffic from several days' observations, or by perturbing some details of the observations. However, to generate samples with a specified set of properties can still be a very labour-intensive task.

For the TCSDG experimental work a computer program was written for traffic

sample generation [16]. This program reads as input a sample specification file and generates pseudo-random samples which comply with the properties specified in the input file. Scheduled arrival and departure rates can be specified for airports, as can the proportion of traffic on each route. For each route the mix of aircraft types and airline operators (as reflected in callsign letters) can be specified, as can the set of flight levels to be used. Using this method, an almost unlimited number of different samples can be generated from one specification. The samples generated were considered sufficiently realistic by the controllers who took part in the experiments.

7.5 The Automatic Controller

The Automatic Controller provides a rather ad hoc collection of control functions for aircraft which are not being controlled by one of the manned sectors. These functions are needed both to support the experiments being conducted, and to improve the realism of the control task in the manned sectors. The Automatic Controller can accept control of flights either when they first come into the system, or when handed over to it by a manned sector. For flights under its control it performs the following functions:-

- For aircraft about to enter a manned sector, it initiates the handover.
- It issues instructions to comply with altitude and speed constraints recorded in the Routes Data part of the SCCS database.
- For arriving aircraft, it issues a descent instruction at the top-of-descent point. For an aircraft whose route will later pass through a manned sector, the descent clearance is to a level appropriate for entry to the manned sector.
- For arriving aircraft, it issues holding instructions and speed control instructions in response to advice from the LOC or SCA.
- For aircraft in a holding stack, it controls "laddering down" and stack exit.
- For aircraft in the approach phase, it controls descent and ILS clearance.
- For departing aircraft, it controls take-off times and climb to cruising altitudes.

In all these cases the Automatic Controller updates the Aircraft State Data in the SCCS database to record the clearances issued, and generates messages to cause Flight Progress EDDs to be updated where necessary. Note that the program does not attempt to maintain safe separation for aircraft under its control.

7.6 Software Implementation

The Traffic Sample Generator is a conventional program with a primary input file and a primary output file, but the SCCS, the Air Traffic Simulator, and the Automatic Controller are real-time multi-tasking programs. All software was written in CORAL 66 [21], and built and run on a Digital VAX-11/780 computer with the VMS operating system. The Air Traffic Simulator comprises about

20,000 lines of source code, the SCCS about 40,000 and the Automatic Controller about 5,000 lines. Each real-time program consists of a set of cooperating asynchronous pseudo-parallel processes. Communication between processes is via message queues and global data areas. The approach used follows the spirit (but not the letter) of the MASCOT software construction method [20].

Processes operate in an "event driven" manner. In most cases the driving event is the arrival of a message from another process, but in some cases the operating system signals the event in response to input from a hardware device. Using the event-driven principle means that all processes have the same basic structure, and this aids understanding and ease of modification. Inter-process messages are of variable length, and each includes message-type and message-length information. A process has a single input message queue on which it receives all types of messages of interest. When a message arrives the receiving process decodes its type to determine the action required, executes the action, and returns to await the next message.

Any process can send messages, and these are automatically copied to the message queues of all interested processes. As part of its initialisation each process declares the types of messages it is interested in receiving. Because sending processes have no knowledge of which processes will receive their messages, processes can be added to or removed from the real-time programs without changing the code of other processes which send messages to them. This aids ease of modification.




Figure 7.1 Experimental Environment




Figure 7.2 Skeleton Control Centre Software

8 CONTROLLER INPUT AND DISPLAY INTERFACES

8.1 Overview

This section describes the design and operation of the displays which were used by the controllers in the LOC and SCA experiments. Both experiments used substantially the same design of displays. The only difference was the display of the CAS in the SCA. Only the SCA formats are described. The interfaces which were developed were considered sufficient for the experimental environment and obviated the need to generate paper flight progress strips. It was never intended that the displays should be developed into operational use. However, the design of the interfaces did evolve in response to controller reactions and in that sense reflect some aspects of the operational requirement. The required facilities were:-

- A simulated radar picture with the capability to display aircraft labels which could include landing sequence numbers
- A tabular display to show sequence numbers, Terminal Approach Times and planned speeds for each aircraft
- A display containing flight progress data for each aircraft
- A means for updating the flight progress data
- A means for interacting with the planning process.

These requirements were met by using three displays on each sector position. A plan view radar display (PVD) was mounted vertically in front of the controller with a vertical electronic data display (EDD) of flight progress data adjacent to it. A data entry device referred to as the Flight Progress Updating Device (FPUD) was mounted below the PVD at an angle of approximately 30° to the horizontal. The device used was a conventional colour VDU with touch-sensitive overlay and it provided the facilities for flight progress data update and interaction with the planning process.

Aircraft appeared on the EDDs some minutes before sector handover and on the FPUD at handover. The controller could not update the EDD until the aircraft had appeared on the FPUD and been accepted (by the controller interacting with the FPUD). At acceptance the position of the aircraft data line on the EDD was changed to indicate that the aircraft was now under that sector's control.

Clearances were shown on the EDD as they were entered by a controller on any sector. Computer advice(sequence numbers, Terminal Approach Times and speeds) was shown as part of the aircraft data line on the flight progress EDD.

8.2 PVDs

The PVDs showed the simulated traffic in the airspace selected. Two versions were used during the project. The first was a monochrome, 16" diameter, cursively scanned device (Plessey Mark 8), of similar characteristics to those used in current operational centres. The second device was a colour raster scan device (Ferranti VARS) on a 20" diagonal tube with a resolution of 1280 × 1024 pixels.

On both displays aircraft positions were marked with a cross and an attached label which included callsign (top line), secondary radar derived height and destination code (second line). For arriving traffic the landing sequence number was shown on the third line when allocated. Route waypoints and airports were shown, some with names displayed. The centre and scale of the picture was independently adjustable for each sector position by a controller input via the FPUD, allowing any part of the simulated airspace and associated traffic to be viewed.

The raster scan device in addition enabled separate colour coding of arrivals and departures, colour-coded airport designators and "attention getters" on the landing sequence numbers and height indicators. The attention getters were used to indicate a newly allocated or revised sequence number and to show when an aircraft had reached its planned top of descent point. The former would be cancelled when the cleared speed was entered (for new allocations) or after a specified timeout (for revisions). The latter would be cancelled when the initial descent clearance was given.

Some commonality between colours used on the PVD and the EDDs was attempted since this seemed to be recognised as a good idea, but choice of colours was purely in response to controller reactions and had no greater significance.

8.3 EDDs

The EDDs were vertically mounted, 24 line $\times$ 80 column colour VDUs and displayed flight details of a list of inbound aircraft appropriate to each sector position. The formats followed a basic style tailored to suit the requirements of a particular sector (e.g. Terminal Approach Times were only shown on the stack inbound sector display). The formats for the two types of sector used in the SCA experiments (stack inbound and enroute) are shown in Appendix D. The screen was partitioned into a pending and an active area. Initially, aircraft appeared in the pending area on adjacent lines. When the aircraft had been accepted from the adjacent sector the controller informed the computer of this by touching the accept mark for that aircraft on the FPUD. At that point the data line for the aircraft was deleted from the pending area and redrawn in the active area. Data lines in the active area were separated by one blank line to simplify the visual search task. The data lines on the display could be ordered according to a criterion set at run-time which included landing sequence number, stack arrival time, top of descent time, predicted sector entry and sector exit time.

If the screen became full subsequent aircraft data lines were held until there was room to display them and a blue marker appeared on the top left of the screen. The data shown on the line for an aircraft was divided into three categories as follows:-

- Fixed (In white) - callsign, aircraft type and stack beacon
- Planned (In yellow) - landing sequence number, stack ETA/TAT or top of descent time (according to sector) and planned CAS
- Cleared (In green) - cleared height, speed and heading

The stack ETA/TAT field was shown only on the stack inbound sector. For non-holding aircraft only the TAT was shown; for those required to hold the predicted arrival time at the stack assuming the aircraft flew the planned CAS (or the preferred CAS in the case of the LOC) was also displayed. The planned data only appeared when the aircraft crossed the time horizon (or inner time horizon). The planned CAS appeared initially in red as an attention getting mechanism. Red was also used to highlight the top of descent time when within 20 seconds of the time.

In addition to the above fields a single letter was shown to designate the preceding sector, a QSY mark in background yellow adjacent to the callsign to show that the aircraft had been transferred to the next sector controller and a foreground blue tick mark was used to show aircraft whose TATs had been communicated to the aircraft. The data line for an aircraft remained on the EDD for a specified time after transfer - usually one minute.

8.4 FPUDs

The flight progress updating device was a 24 line $\times$ 80 column colour VDU with a pressure-sensitive touch overlay. The device provided a means of updating the information on the EDD. The formats for the various menus together with the syntax tree for touch inputs are given in Appendix D. The device was operated by touching labelled areas on the screen. When an area had been activated by touch this was indicated by changing the background to red.

From the rest picture an aircraft callsign could be selected (once it had been accepted by touching "ACC"). This caused a copy of the aircraft data line (in a format similar to that shown on the EDD) to be displayed on the FPUD. Fields of the data line could then be touched in order to activate updating menus appropriate to the field. For example, to enter a cleared height the height field of the data line was touched causing a menu of flight levels to be selected. A flight level was selected by a single touch followed by touching ENTER. Multiple field updates could be made for an aircraft by successively touching fields and terminating the sequence by touching ENTER. Reversion to the rest picture could be achieved at most points by touching ENTER to complete an update or REJECT cancelling all entered updates.

9 EXPERIENCES FROM USING THE PROTOTYPE TOOLS

9.1 Introduction

Earlier sections have described the experimental prototype tools for computer assistance. The aim of this section is to describe the way the experiments have been conducted, the issues which were raised and the impact this had on the development of the tools.

Previous TCSDG arrivals management work aimed at investigating the "metering" of arriving traffic in what could be described as a "tightly planned" manner. A plan was generated for each aircraft which comprised a series of recommended air traffic control instructions (height and speed changes), some of which were given in advance of their implementation (e.g. "Descend to flight level 190 at 10 miles before Longsands"). The plan was complicated because of the airspace design used, involving the sharing of high-level stacks between airports. A rigorous monitoring of conformance to the plan was also performed. The controller's actions were therefore quite tightly constrained by the need to follow the planned actions.

The work described in this report was initiated in an attempt to produce a simpler form of assistance which would interfere with the normal tasks of the controller (such as conflict resolution) as little as possible. Indeed it was hoped that a stand-alone system could be introduced operationally which could be optionally used by controllers, but which would not hinder their task should they wish to ignore the advice.

Three main stages can be identified in the experimental work. Firstly, a Landing Order Calculator was built, partly as the simplest type of tool possible and partly as a system to be used as a reference against which to compare more sophisticated tools. The second stage was the development of the Speed Control Adviser (SCA) and, finally, an order optimisation version of the SCA was produced.

The experiments covered three years and involved controllers from NATS. The experiments have relied heavily on inputs from the CCF airspace planners in order to represent the proposed new route structures and procedures for the London Terminal Area due for implementation in the early 1990s. Many of the issues raised from the experimental work have been related to particular "real-life" airspace problems (e.g. the positioning of the major holding fixes).

9.2 Overview

The experimental work described here is based on real-time simulation of two air traffic control sectors as described below. It must be recognised that the results which come out of the work are concerned more with assessing the viability from a man/machine viewpoint and with stimulating the design thought processes than producing statistics showing by how much capacity can be increased and workload reduced. These results can be demonstrated by other forms of simulation which do not involve the interaction of the controller (e.g. discrete event simulation). In order to derive statistically significant results from real-time simulation a large number of human resources and many hours, if not days, of operation are required.

9.3 Conduct of the Experiments

Generally, experiments were run in one-day sessions approximately once or twice per month during the project. The tools were usually run for several periods of about one hour followed by a debriefing session with the controllers at which system performance was analysed.

9.3.1 Landing Order Calculator Experiments

For the LOC experiments two sectors were configured and two combinations tried:-

- a. *A TAF inbound sector and an approach control sector.* The TAF inbound sector controlled traffic from typically flight level 190 to the TAF altitude (6000 - 11000 feet). The controller managed the stack when necessary including communicating the TATs to the aircraft. The approach control sector performed the number one radar director function.
- b. *A TAF inbound sector and an enroute sector.* The TAF inbound sector was as for configuration a). The enroute sector controlled traffic from cruise altitude to an agreed handover level to the TAF inbound sector (typically at flight level 190).

For the two configurations the landing sequence numbers were displayed at both sector positions on the EDDs and the radar picture. For configuration a) the TATs were also displayed on each sector EDD. For configuration b) the TATs were only shown on the TAF inbound sector. For configuration a) the automatic controller program handled the initial descent phase for aircraft in the enroute sector with automatic handover to the TAF inbound sector and for configuration b) it handled the approach control function.

Most of the experiments were concerned with handling traffic into the Heathrow TAFs, with a selection of TAF inbound sectors and associated enroute sectors over a number of runs. Planning information, however, was produced for all four airports meaning that sequence numbers for all arriving traffic were shown on the radar picture. Departing traffic was also present in the system to increase the realism of the control task.

Usually both sectors were manned, although it was possible to run either automatically for the purpose of concentrating on a particular problem in one sector.

9.3.2 Speed Control Adviser Experiments

These were run in a similar manner to the LOC experiments, but because the interest was in the issue of initial descent clearances and speed control instructions which for most aircraft occurred in the enroute sector, only configuration b) was used.

In all the experiments the main aim was to assess the initial planning concept and, although many observations were made which resulted in changes to the displays,

this was a secondary aspect of the experiments. This is not to understate the importance of a well developed man/machine interface, but merely to emphasise that further work would be necessary in the light of operational constraints such as the integration of flight progress data with the planning information.

The remaining sub-sections represent a list of observations from the experiments which give an indication of the techniques which are most likely to lead to the successful development of an operational system for arrivals flow management. Many of the observations apply to both the LOC and the SCA, but where this is not the case this will be indicated.

9.4 Effects of Controller Intervention

When the LOC or SCA allocates a landing sequence number and TAT, it assumes that a certain trajectory will be flown from the time horizon to the TAF. It is inevitable that on some occasions controllers will have to cause an aircraft to deviate from its planned trajectory in order to resolve potential conflicts. Such controller intervention will cause aircraft to arrive at the TAF earlier or later than planned.

Consider a typical situation where three inbound routes converge towards a common TAF. The usual (or preferred) merge point is about 40 miles before the TAF. Simultaneous arrival of traffic on each of the routes implies the need for conflict resolution action on the part of the controller. The worst situation is if the aircraft are descending to a common altitude, e.g. an agreed sector handover level.

Two methods of resolving this conflict have been observed during the experiments and were used according to controller preference. The first method is to separate the aircraft vertically and gradually clear the aircraft for descent as separation rules permit. Since this technique implies constant monitoring by the controller of aircraft heights and deviation from the ideal of a continuous descent some controllers prefer to use a second method. In this method the aircraft are given radar vectors to fly so that they turn on to parallel courses with appropriate lateral separation, thus allowing continuous descent.

Both methods can upset the planned flow as is now described.

Vertical Separation Method

The severity of the effect from using this technique is a function of the altitudes of the converging aircraft relative to the sequence numbers which the computer has allocated. Indeed the order may not be achievable at all using vertical separation alone. However, the most likely consequence is that an aircraft's descent will be far from continuous. Since the predicted trajectory assumes continuous descent, this may mean either that the aircraft will arrive late at the TAF due to being descended early (assuming constant CAS during descent and any level sections during descent) or early due to remaining longer at high altitude. In the latter case there is a limit to how long an aircraft's descent can be delayed consistent with achieving the required height by the TAF since a faster rate of descent will be required.

Using this technique the significant factor is the difference between the time taken to fly a given distance at high and low levels. For example, if descent starts two minutes early for a descent from flight level 280 to 8000 feet at a typical CAS of 260 knots approximately an additional 40 seconds will be required to cover the equivalent high-level distance, leading to late arrival². Descending late would lead to a similar early arrival error.

Lateral Separation Method

The effects from this method are due to either increased or reduced track distance being flown by the aircraft. The amount of error introduced in terms of time clearly depends on the true airspeed of the aircraft. For an aircraft between flight level 250 and 350 a five nautical mile track increase or reduction is typical and the controller often takes advantage of natural turns on the route. This leads to a time error of between 40 - 60 seconds.

From the above discussion it can be seen that there will be cases when the landing sequence numbers and speeds which the computer advises will not be achievable. It was in the light of observing such situations experimentally that consideration was given to providing the controller with a facility to modify the computer-generated plan. Two such facilities were implemented.

- a. The landing sequence numbers for two aircraft could be exchanged by the controller issuing a command to the computer. On its own this simply serves to communicate a sector controller's plans to other interested controllers. In the opinion of the controllers it was undesirable to allow an aircraft to be swapped with any other aircraft except one under the same sector control, because to swap aircraft on different sectors implies inter-sector coordination.
- b. A new speed could be computed for any aircraft and advised to the controller for implementation. It is essential with this facility to give an indication of whether the Allocated Landing Time is still achievable. Experimental observations indicated that this is only infrequently likely to be the case, but if it is the landing sequence would need to be modified using facility a).

The philosophy behind these facilities was to provide them as options for the controller. The decision as to when it was necessary to use them relied on the judgement of the controller.

9.5 Overtaking Aircraft

A typical mix of arriving aircraft will exhibit a spread of speeds and cruising altitudes. This implies true airspeed differentials which, assuming similar wind fields will cause the faster aircraft to overtake the slower ones. Where the speed differential is large or the height difference significant, overtaking is usually acceptable from an ATC viewpoint. However, where aircraft are at similar altitudes and speeds it may take a significant track distance before two aircraft pass each other. Since the aircraft would most likely be put on to radar vectors

²similar wind fields are assumed at both altitudes

for separation during the overtake a large amount of airspace may be needed to accomplish the overtake and for a significant length of time. The controller would, generally, prefer to avoid this situation and indeed may be forced to do so by airspace restrictions.

In the light of the above discussion and the observations made of overtaking situations during the experiments it became clear that a computer-generated plan which would require aircraft to overtake in order to achieve the planned landing order would not always be possible to implement. Because the main problem encountered was when aircraft were overtaking near to (or in some cases even beyond) the TAF in order to achieve the planned order, it was felt that the requirement was for the ability to mark a point on each inbound route beyond which no overtaking must occur. The planning algorithm would be able to identify this point and avoid generating an overtaking order, based on a prediction of the arrival times there. (This was not implemented at RSRE.)

In section 4.2 it was indicated that stacking aircraft must fly as slowly as possible enroute. The reason for making this rule was because of overtaking problems observed during the experiments when other strategies were adopted. For example, it might seem fairer to operators to allow the preferred descent speeds to be flown, in view of the required fuel penalty of low-level holding. However, fuel optimal speeds for modern jet aircraft generally lie above minimum speed. During the onset of stacking the allocation of the preferred speed to the first aircraft on a route to be required to hold can mean this aircraft overtaking one or more aircraft which have just avoided stacking (by reducing speed) and are flying minimum (clean) speeds.

Clearly, the existence of an overtaking constraint may theoretically reduce capacity, but if two aircraft are close together as they near the TAF, the effect is, in practice, unlikely to be significant.

9.6 Landing Sequence Numbers

In early experiments all landing sequence numbers were dynamically updated as an aircraft landed so that the next aircraft to land was always shown as number one and so on. This meant that allocated sequence numbers for all aircraft were continually changing. This was very confusing to the controllers and it was decided that sequence numbers should be fixed once allocated, thus giving only a relative indication of the position of other aircraft in the sequence.

The question of when to allocate and display sequence numbers to the controller is also interesting. The time of allocation is determined by the position of the time horizon (or inner time horizon for the OptO allocation method) and normally this would be the time to display the sequence number. However, one type of situation was observed in which some aircraft on a route had sequence numbers shown and a slow aircraft ahead of these did not because its time horizon crossing point was nearer to the airport. If there was a significant delay before this allocation was made the controllers would be uncertain whereabouts in the sequence the slow aircraft was planned to fit.

One possible solution to the problem would be to look ahead to the predicted time horizon crossing time for the slow aircraft and perform a preliminary landing

times allocation for all aircraft in the system which had not yet been allocated times. This would require considerable processing and it could not be guaranteed that the preliminary allocation for the slow aircraft would be the eventual allocation (due perhaps to controller intervention on another sector). Since displaying misleading information is probably worse than showing no information this is not a recommended approach. In fact, use of the time horizon principle guarantees that the aircraft would be allocated a number later than any already allocated and this is probably sufficient indication.

9.7 Factors Influencing the Choice of Time Horizon Value

Section 2.4 discussed the advantages of the time horizon method of landing times allocation. During the experiments the effects of choosing different time horizon values was observed in the context of the CCF airspace structure. It emerged that the choice of value can be critical in the context of problems on a particular ATC sector.

Section 9.4 described a scenario with three converging inbound routes and some approaches to controlling traffic in this sort of airspace. The merge point (near to Longsands on the Clacton sector) was close to the point of handover to the next inbound sector. As far as possible the task of the enroute sector controller was to pass the aircraft to the next sector in an order which, bearing in mind the computer-allocated landing order, presented the next sector controller with as few control problems as possible in terms of, for example, overtaking aircraft.

Choosing the time horizon so that sequence numbers were allocated to some aircraft well before the merge point sometimes interacted with the technique for separating the aircraft. It was found that moving the time horizon in towards the airport by two minutes (from 25 to 23 minutes) meant that most aircraft were close to or had passed the merge point when crossing the time horizon. Indeed, for the FCFS order allocation method the controller was able to exert some small influence over the allocated landing order by, for example, shortening an aircraft's route and causing it to cross the time horizon earlier.

This technique worked well, but by choosing the time horizon value based on a specific problem in one sector a problem on another sector was introduced. In the CCF structure the locations of the holding fixes meant differences between the track lengths from each stack to the runway. In particular the stack at Milton Keynes gave about a 45 mile stack to runway track compared with 20 miles from the fixes at Lambourne and Oxshott. Therefore moving the time horizon value in closer for the Clacton sector as described above meant bringing the time horizon for most aircraft passing through Milton Keynes to typically about 15 miles before the stack. The effects of this were twofold:-

- a. The TATs were not available until aircraft were close to the holding fix
- b. The amount of time available for speed control was reduced, assuming no speed control after the TAF.

The effects from a) are not serious until it is necessary to stack, when TATs are most useful. For most aircraft there would still be time to issue hold instructions

and communicate TATs, but the advantage of having this information early would be lost. For slow aircraft the situation could be much more serious since they might actually have passed the holding fix before crossing the time horizon. In the RSRE experiments this case hardly ever arose and the solution was to force allocation at a point just prior to the holding fix for any aircraft which would not cross the time horizon before that point.

The effect from b) was reduced in the case of the Milton Keynes stack by allowing speed control beyond the stack. The term "allocation fix" was introduced to cater for this situation. No speed control was planned to take place after this point³. For most routes the allocation fix coincided with the TAF. For the Milton Keynes case it was about 15 miles beyond the stack and there was a maximum speed constraint between the TAF and the allocation fix.

It can be appreciated from the above that there may be occasions due to airspace design in which slight deviation is made from the time horizon concept. It must be emphasised, however, that this will introduce a bias in favour of some approach routes in terms of landing times allocation and this must be considered in the light of other unfairnesses in the system such as, for example, re-ordering aircraft according to the turbulent wake rules.

9.8 Alternative Routings

This section discusses the issue of alternative routes to touchdown and proposes some modifications to the algorithms described in sections 3 and 4. These were not implemented at RSRE but are considered to be straightforward to incorporate and necessary for realistic operation.

As aircraft near the approach sequencing area there are possible variations in routing, depending on whether an aircraft is to enter the hold or not and whether path length variation in the sequencing area is necessary. The allocation process is parameterised (by MAXIMUM ABSORB TIME) to allow path length variation to be used for delays of less than the minimum stack orbit time.

The allocation process should take account of the alternative routings as follows :-

- a. Define the shortest route to touchdown as the "standard" route
- b. Define the route via the stack as the "stacking" route
- c. Define two TAFs, one for the standard and one for the stacking route
- d. Compute stack ETA for holding aircraft in usual way
- e. Compute TAT
- f. Compute landing times (both Preferred and Allocated) assuming aircraft follows the standard routing
- g. In deciding whether an aircraft must hold or not and computing the terminal approach time the parameter MAXIMUM ABSORB TIME should be modified to represent the amount of time which can be assumed to be

³The approach controllers may of course require to adjust speeds, but this is not part of the SCA

absorbed between the TAF and the runway. The time that can be absorbed will continue to be a function of airport approach geometry.

With this design the interpretation of the TAT would depend on whether an aircraft is to enter the hold or not and would relate to the appropriate TAF.

9.9 Adjusting the Planned Flow

The means of controlling the rate of delivery of aircraft into the approach sequencing area in the RSRE experiments was by the adjustment of the planned landing rate parameter available interactively to the controller. Only a small amount of experience was gained with this facility. The need for and the implications of such flow control mechanisms in both the LOC and SCA are now discussed.

The flow needs to be adjusted from time to time for the following reasons:-

- a. Runway capacity is temporarily lost due to some runway incident (e.g. runway inspection, aircraft slow in clearing, burst tyres etc.)
- b. Low visibility conditions
- c. Approach control workload limits.

Four mechanisms for controlling the flow have been identified:-

- i. *Reduce the planned landing rate and modify plans for all aircraft.* By this mechanism the separations between aircraft are adjusted to give a delivery rate which does not exceed the set landing rate.
- ii. *Reduce the planned landing rate but only alter the separations for future allocations.*
- iii. *Delay the complete arrivals stream by a fixed amount of time.* By this mechanism the separations between aircraft remain the same but all the landing times are delayed by the specified amount.
- iv. *Delay only future allocations by the specified amount.*

In cases i. and iii. new TATs would be computed and displayed immediately to the controller, together with revised indications of which aircraft must stack (if any). The SCA would also compute new speeds for those aircraft with revised runway times (not all aircraft will necessarily need new times if there are natural gaps in the stream).

When a flow adjustment takes place the aircraft which have firm landing times allocated will be typically within 25 minutes from the runway and less than 12 minutes from the TAF. Consider, firstly, the recomputation of TATs and stacking states (which aircraft are to stack). If the new information indicates the onset of stacking then aircraft which are subject to sufficient delay will be required to stack or, in the case of the SCA, to fly a slower speed (if possible). For those aircraft which have already passed the holding fix no delaying action is possible

(unless speed control is permitted beyond the TAF, e.g. at Milton Keynes in the CCF route structure). The tools give no assistance for these aircraft. Aircraft will also have difficulty in absorbing extra delay enroute unless they are near to the time horizon due to the limited scope for delay absorption.

There are further implications for revising TATs for those aircraft in the stack. If the extra delay required for an aircraft is less than the time required to fly a minimum stack orbit then the aircraft might not be able to depart exactly as requested. In this case the decision as to whether an aircraft departs immediately (and early) or flies an orbit and departs (late) should probably be based on which gives the least deviation from the planned time⁴.

When new TATs and/or speeds are displayed for several aircraft as a result of a flow adjustment there is a potential sudden increase in the controller's workload (due to the need to communicate the new information to the aircraft). Care must therefore be taken when considering how and when the new data is displayed.

For aircraft which have no landing time allocation yet, the planning information will reflect the new delivery rate when it is generated.

9.10 Creating Spare Landing Slots

There are at least three occasions on which it would be necessary in an operational system to plan for one or more spare slots in the arrivals stream. The first is when an aircraft performs a missed approach (overshoot) procedure and the second when it is required to allocate slots to aircraft whose point of departure lies within the time horizon, e.g. an aircraft flying from Birmingham to Heathrow.

In the first case it would be expected that the overshooting aircraft could be resequenced within five or six slots after the missed one. For a short time there would be more aircraft in the approach sequencing area than desirable, but one of the flow adjustment mechanisms described in the previous sections would be used to slow down the arrivals stream.

In the second case the spare slot would need to be reserved before the aircraft was airborne. It is assumed that a preliminary assessment of slot availability would be made before the aircraft was cleared to depart. It would also be undesirable for the aircraft to depart and then be required to fly a holding pattern. To avoid this the planning algorithms would need to be able to take account of the fact that the required delay (or some part only to give a margin for errors in flight) was to be absorbed on the ground. Once airborne the normal facilities for changing sequence numbers and recomputing speeds would of course be available.

Thirdly, the facility to create spare landing slots would be an essential component in applying the tools to shared runway operations, where slots would be reserved for departing aircraft.

⁴There could be a case for consistently bringing aircraft off early because it may not be possible to make up time in the approach sequencing area

9.11 Optimising the Landing Sequence

The final stage of the experiments was to implement a version of the SCA which attempts to optimise the landing sequence to take account of the requirements for separating aircraft according to turbulent wake rules. This meant a departure from the first come, first served principle. The main conclusion was that reordering is possible but that if two aircraft on the same route are swapped then extra problems for the controllers are sometimes introduced (in the form of overtaking). A few other points are worth discussing.

The value for the inner time horizon is governed by the considerations discussed in section 9.7. The outer time horizon is set n minutes earlier. Because the aircraft between the horizons are those eligible for inclusion in the optimisation process, which involves computing all the possible permutations the value of n must not be so large as to cause a combinatorial explosion. A maximum of six eligible aircraft was considered realistic (this represents 720 permutations). The value n is also governed by aircraft performance limits and the assumption that it is reasonable to move an aircraft by a maximum of two positions from the first-come-first-served order. This greatly reduces the number of possible permutations. The value of n used in the experiments was five minutes.

Another effect of reordering aircraft in this way is that aircraft with the same wake vortex categories tend to be grouped together in the sequence. This can have a detrimental effect on minority aircraft types which are the most likely to suffer the maximum position displacement.

9.12 Display Issues

As already noted the development of the various displays in the experiments was not the primary objective. It was necessary to display the planning information and to keep a record of flight progress data without using paper flight strips. No final assessment was made of the best place for the planning information, but it was felt that in a system using paper flight progress strips the data could be incorporated into the radar picture. For a system using electronic flight strips the data should be displayed as part of the strip. Although the displays did not contain the amount of detail required for an operational system, they were generally well accepted by all the controllers who participated in the experiments.

9.13 Impact on the normal task of the controller

As mentioned earlier the experimental tools are advisory and need not affect the normal control task. However, when used, there are some implications for the controller which are now summarised.

The tendency, by the use of speed control, to smooth out the flow of traffic should reduce the number of times controller action is required in order to resolve conflicts. Less time should be spent in the management of the stacks, due to en-route (linear) delay absorption.

There is extra use of the radio telephone to communicate the speed control instructions. These must be issued in time to be implemented by the aircraft

during the descent phase, but may be communicated earlier (see section 4.2, page 4-4). In the experiments the issuance of the TAT implied clearance to depart the holding pattern at the specified time with no further instruction expected. This method of operation was accepted by the controllers involved.

In the Speed Control Adviser TATs need only be communicated when a holding pattern is required to be flown and in that case should be communicated to the aircraft early enough to enable the most accurate possible time of departure to be achieved, i.e. well before the aircraft arrives at the holding fix.

In the Landing Order Calculator it is possible that, if TATs were communicated as soon as available, suitably equipped aircraft could use the information to compute the best top of descent point corresponding to the preferred profile. In this way it might be possible for some aircraft to perform their own speed control.

There is, clearly, extra workload for the controller in requesting revised speeds following tactical intervention and also in relaying changes of the landing sequence numbers to the display system.

10 CONCLUSIONS

This section offers some conclusions both about the possibilities for computer assistance in arrivals flow management, and about the methods used by the research project.

The following conclusions can be drawn about computer assistance in arrivals management:-

1. An improved understanding of the possibilities and practicalities of computer assistance for arrivals flow management has been gained.
2. An experimental prototype Landing Order Calculator (LOC) was developed and exercised in a real-time simulation environment by a number of Air Traffic Control Officers. The LOC provided landing sequence numbers as part of the labels on the radar display. It also provided a good estimate of Terminal Approach Time, and an early indication of whether or not holding would be required, on the flight progress display.

The controllers who took part in the simulation experiments found the LOC to be of benefit to them, and made a substantial contribution to its refinement.

3. The research indicates that the benefits to be derived from the LOC are:-

- Reduction in controller workload because:-
 - controllers do not have to determine landing sequence numbers or communicate them to each other,
 - controllers receive an early indication of when holding is required,
 - early determination of optimised landing order (optimised to minimise the effect of wake turbulence separation rules on landing rate) allows the process of traffic re-ordering to begin outside the approach sequencing area.
- A smoother and more orderly flow of traffic into the approach sequencing area because:-
 - early display of landing sequence numbers enables TMA controllers to leave gaps in the stream for later merging of traffic from other directions,
 - early communication of Terminal Approach Times to aircraft enables pilots and avionic systems to plan accurate stack exit times.

4. An experimental prototype Speed Control Adviser (SCA) was developed and exercised in a real-time simulation environment by a number of Air Traffic Control Officers. The SCA included the functions of the LOC. In addition the SCA gave advice about the descent speeds which would cause aircraft to arrive at the Terminal Approach Fixes close to the planned times. This advice was shown on the flight progress display.

The controllers found the SCA to be of benefit to them, and made a substantial contribution to its development through their comments and suggestions.

5. The research indicates that the benefits to be derived from the SCA (in addition to those listed for the LOC above) are:-
- Reduced delays or increased capacity because:-
 - increasing the speeds of a few aircraft allows runway time to be used which would otherwise be wasted,
 - early determination of landing order together with speed control allows a greater degree of order optimisation (for aircraft which are not holding) than is possible today.
 - Reduced holding because of use of speed variation as a delay absorbing mechanism.
 - A smoother flow of traffic into the approach sequencing area (in addition to that gained from the LOC) through the use of speed control.
 - Improved aircraft fuel economy through reduced holding.
6. For both LOC and SCA, the great importance of having the right controller/computer relationship must be recognised. The experimental LOC and SCA largely achieved the aim of providing advice rather than "driving" the controllers. The advice did not generally get in the way when not taken. However it proved to be possible for the tools to influence the ATC processes in subtle ways. For example, modifications had to be made to the algorithms to prevent the generation of unnecessary overtaking situations.
7. When the SCA performs its speed calculations, the precise trajectory to the runway is not known. The SCA assumes a "standard" trajectory, but aircraft deviate from this for many reasons including: conflict-avoiding manoeuvres, variable time taken to react to ATC clearances, and uncertainty about wind conditions aloft. These uncertainties mean that aircraft are not delivered to the Terminal Approach Fix at exactly the planned times. Most aircraft deviate from the planned times by less than 30 seconds, but a few have deviations large enough to require controller action.
- It was found necessary to provide controllers with two mechanisms for dealing with this situation: a means of over-riding the computer-planned landing order, and a means of requesting the re-calculation of descent speed.
8. The potential benefits to be gained from the SCA because of the "Negative Delay" effect were quantified by a discrete event simulation study. The benefits were expressed in terms of delay/traffic intensity curves for the arrivals queueing process. By using the SCA, delay can be reduced while maintaining the same traffic intensity, or traffic intensity can be increased while maintaining the same mean delay.
9. In current ATC practice the landing sequence is optimised by approach control to minimise the effect of wake vortex separation rules, and thus maximise landing rate. It was demonstrated in the real-time simulation environment that the SCA enables the re-ordering to begin between top-of-descent and the Terminal Approach Fix. However it was found to be preferable to avoid exchanging the positions in the landing sequence of two aircraft on the same route.

10. The importance of the time horizon method for completely fair allocation of landing times was demonstrated by a discrete event simulation study. The effect of where landing sequence numbers were allocated relative to merging points in the CCF airspace design was studied in the real-time simulation environment. It was found that the time horizon principle might sometimes have to be compromised slightly, to reduce ATC problems.
11. A trajectory prediction module sufficient for the needs of the SCA and LOC was developed. This was based on the aircraft performance modelling work (PARZOC) of Benoît, but the PARZOC method was considerably extended to allow incorporation of wind information, to represent the approach phase of flight, and to model descent at constant Mach number. However, no great difficulty is envisaged in upgrading the module to take account of the new aircraft performance modelling work currently being undertaken by Eurocontrol [17].
12. While it was not the purpose of the TCSDG project to research electronic display of flight progress data, it was necessary to provide some such form of display in the simulation environment, and an electronic method was chosen. The way in which flight progress data was displayed, the way in which controllers interacted with it, and the way data from the LOC or SCA was embedded in it, all developed considerably during the experiments as a result of controller suggestions and comments.

Within the limitations of the types of sector being simulated, controllers who took part in the simulation exercises had no great difficulty working with the electronic flight progress data. Some expressed considerable enthusiasm for it.

The LOC and SCA could not be developed much further within the small-scale simulation environment (two manned ATC sectors) available at RSRE. It is recommended that the LOC and SCA be integrated into the much larger simulation environment available at ATCEU, and that they be exercised within the more realistic ATC environment which can be provided there. From the TCSDG research work, there is enough confidence in the eventual success of the LOC and SCA to begin planning their eventual operational implementation and this recommendation has now been accepted by the CAA.

The following conclusions can be drawn about the methods adopted by the TCSDG research project:-

13. The method adopted - use of a small easy-to-change real-time simulation environment, with fast incorporation of ideas and suggestions from all concerned with the project - was generally successful. A more realistic ATC environment would have cost more to produce and man, and would have slowed down the process of translating ideas into computer programs. A less realistic environment would have been too far from the real world to be useful. The balance was about right.
14. A great deal of experience and understanding was gained of ATC processes, trajectory prediction by computer, and real-time simulation of ATC environments. This will be of great value in future RSRE projects on computer assistance for other aspects of ATC.

It is recommended that a similar method be considered for future *RSRE* research projects. It is recommended also that care be taken to harness as much as possible of the experience gained from the *TCSDG* project to the needs of other research projects.

11 ACKNOWLEDGEMENT

The authors wish to acknowledge the Civil Aviation Authority sponsorship of this research project.

The authors would like to thank the many Air Traffic Control Officers who took part in the TCSDG real-time simulation exercises, and who contributed many valuable comments, criticisms and suggestions. Without their participation this research project would not have been possible.

The authors would also like to thank all those people – scientists and engineers, controllers and managers – who although they did not actually take part in the simulation exercises, saw demonstrations of the experimental system and made useful comments as a result.

References

- [1] D E Knuth. *The art of computer programming*, volume 1. Addison-Wesley, 1969.
- [2] R P Maggs. Private communication. Civil Aviation Authority, London, 1987.
- [3] V W Attwooll. *Some mathematical aspects of air traffic systems*. Journal of Navigation 30,3 Sept 1977.
- [4] *Procedures for Air Navigation Services - Aircraft Operations, Vol 2: Construction of visual and instrument flight procedures*. International Civil Aviation Organization Document 8168-OPS/611, volume 2, Montreal 1982.
- [5] L G Hewson. *The Scheduling Committee Function*. UK Guild of Air Traffic Control Officers conference CONVEX 87, London, 1987.
- [6] Private communications with several UK airlines.
- [7] S A N Magill and A J Budd. "Negative delay" in the airport arrivals process. Proc. UKSC Conference on Computer Simulation, Bangor, Sept 1987.
- [8] V W Attwooll and D Rickman. *A study of conflicts within a Zone Of Convergence*. Paper presented at a meeting of the Eurocontrol Specialist Panel on Automatic Conflict Detection And Resolution, Brussels, 1984.
- [9] A J Budd. *Comparison of interpolation algorithms for speed control in air traffic management*. RSRE Memo 4131, 1988.
- [10] A J Budd. *A simulation study of landing time allocation procedures for use in computer-assisted air traffic management systems*. RSRE Memo 4132, 1988.
- [11] W Finch. *Route data representation for air traffic control simulation experiments*. RSRE Memo 4133, 1988.
- [12] S A N Magill, W Finch, A C F Tyler and E T Wilkinson. *A real-time simulator which provides an environment for air traffic control experiments*. Proc. SCS European Simulation Conference, Antwerp, Sept 1986.
- [13] A Benoît and P Brown. *A presentation of EROCOA climb and descent performance data with illustrations of the approximation quality*. Eurocontrol Document FPP E1/T/2211/5, Brussels, 1973.
- [14] A Benoît and S Swierstra. *Simulation of air traffic control operation in a zone of convergence: 2. Aircraft (PARZOC) performance data*. Eurocontrol doc 812031-2, Brussels, April 1982.
- [15] E T Wilkinson. *Aircraft performance modelling in real-time simulation experiments at RSRE*, presented at a meeting on *Aircraft performance models for present and future ATC needs* at the Eurocontrol Experimental Centre, Nov 1986. Published as Eurocontrol document 872002, Jan 1987.
- [16] S A N Magill, G Ord and K A L Burek. *A traffic sample generator for ATC simulation experiments*. RSRE Memo 3818 (to be published).

- [17] J-L Renteux. *Aircraft modelling standards for future ATC systems*. Eurocontrol Doc. 872003, Brussels, 1987.
- [18] *Manual of Air Traffic Services - Part 1*. NATS, CAA UK.
- [19] J.O. Cook. *Calculation and Display of Stack Departure Times for aircraft inbound to Heathrow Airport*. RSRE Memo 3750, August 1984.
- [20] *Official Handbook of MASCOT (MASCOT II)*. Joint IECCA and MUF Committee on MASCOT (JIMCOM), March 1983.
- [21] *Official Definition of CORAL66*. H.M.S.O., 1970.

APPENDIX A

SIMULATION STUDY OF "NEGATIVE DELAY"

As pointed out in section 2.3.3, speed control in the main descent phase has the possibility of advancing aircraft by speed increase as well as delaying them by speed decrease. Thus it is possible to advance some aircraft to make use of portions of runway time which would otherwise be wasted. The term "negative delay" has been coined to describe what happens when aircraft are advanced in this way. A discrete event simulation study was undertaken to demonstrate and quantify the effect of negative delay on the statistics of the runway queueing process. This appendix gives a brief summary of the simulation and the results. A fuller description can be found in [7].

A.1 The Simulation

The simulation modelled an aircraft arrival process, an approach and runway service process, and two alternative landing time allocation processes (both performed at a time horizon). One landing time allocation process assumed that aircraft could be delayed but not advanced, while the other assumed that aircraft could be advanced within certain limits as well as being delayed. From queueing theory, traffic intensity and delay are inextricably bound up with each other, so each cannot be discussed separately from the other. Graphs of mean delay plotted against traffic intensity were produced for the two landing time allocation processes.

Airport arrival peaks do not last long enough for the queue to reach a steady state, so it was necessary to simulate a large number of short arrival peaks. Within each peak, scheduled arrival times were assumed to be equally spaced, (wh. this does not happen in practice, the assumption was good enough for the purpose of this simulation). Actual arrival times were generated by adding normally distributed random perturbations to the scheduled arrival times. This substantially re-orders the scheduled arrival sequence.

The service time was modelled as a constant. There are two reasons for doing this. Firstly, a computer program which is allocating landing times some 20 - 25 minutes before touch-down cannot do anything about the probabilistic variation of inter-landing time, and must work with a mean value. Secondly, the speed-controlled phase of flight planned by the SCA will be followed by a normal approach sequencing phase which will essentially buffer it from the runway.

A.2 Results

Figure A.1 is the main result of the simulation. Mean runway service time and maximum possible negative delay both depend heavily on airport and airspace geography, and other local conditions. Because of this, mean runway service time was taken as the unit of time, and delay is expressed in terms of this unit in figure A.1. The figure shows plots of mean delay against traffic intensity without negative delay, and with three different values of maximum negative delay. The effect of negative delay can clearly be seen. Putting in numbers typical for a large

airport in Western Europe, it is found that mean delay can be reduced from 5 minutes to 3 minutes while keeping traffic intensity constant, or traffic intensity can be increased by 5% while keeping mean delay constant.

How figure A.1 changes with variation of the magnitude of schedule perturbation, and with the length of the arrivals peak, is explored in [7].

In order to investigate the effect of negative delay on the percentage of aircraft required to hold, cumulative delay distributions were plotted for the two cases, without and with negative delay, see figure A.2. A cumulative delay distribution curve gives for each delay value d the percentage of aircraft with delay $\leq d$. If d is now chosen to be the threshold value at which holding begins, figure A.2 shows that a greater percentage of aircraft do not have to hold when speedup is used. Putting in some typical numbers again, the following percentages result:-

No speed control at all	63%hold
With speed reduction	39%hold
With speed increase and decrease	22%hold.




FIG A.1 EFFECT OF SPEEDUP ON MEAN DELAY




FIG A.2 EFFECT OF SPEEDUP ON CUMULATIVE DELAY DISTN

APPENDIX B

SIMULATION STUDY OF THE TIME HORIZON PRINCIPLE

As pointed out in section 2.4, the only completely fair way to allocate landing times is to use the Time Horizon principle. A discrete event simulation study was undertaken to demonstrate this fact, and to quantify in terms of mean delay the unfairness caused by various forms of deviation from the principle. This appendix gives a brief summary of the simulation and the results. A fuller description can be found in [10].

B.1 The Simulation

Two types of deviation from the Time Horizon principle were investigated:-

- *Situations where most of the traffic has landing times allocated according to the time horizon principle, but a fraction has times allocated a fixed time away from the time horizon.* The aim was to show that if the fraction had its landing times allocated before the time horizon, it suffered less mean delay than the rest of the traffic, and if it had landing times allocated after the time horizon, it suffered more mean delay than the rest of the traffic.
- *Situations where the time horizon method is not used.* These situations were modelled by allocating landing times at some random time interval (uniformly distributed between two limiting values) before Preferred Landing Time. The aim was to demonstrate the correlation between the time before Preferred Landing Time at which the allocation was done, and mean delay.

The arrival process and service process were modelled as described in Appendix A.

B.2 Results

Figure B.1 shows a family of delay curves for the case where 90% of aircraft have landing times allocated according to the time horizon principle, and the remaining 10% have landing times allocated a fixed time interval away from the time horizon. Mean delay is plotted against traffic intensity for several different values of the fixed time interval, (note that delay is expressed in minutes, not the time units used in Appendix A). It can clearly be seen that traffic which has landing times allocated before the time horizon suffers less mean delay than the rest, and traffic which has its landing times allocated after the majority suffers more mean delay. At 95% traffic intensity, the traffic with landing times allocated 5 minutes early has mean delay reduced by 2.7 minutes, and traffic with landing times allocated 5 minutes late suffers 4.4 minutes extra mean delay.

Figure B.2 is for the case where the time horizon method is not used. Landing times are allocated at points which are uniformly randomly distributed between +10 and -10 minutes from a fixed time before Preferred Landing Time. Several curves are shown for different traffic intensity values. It can clearly be seen that there is an almost linear relationship between the time at which the landing time is allocated (relative to Preferred Landing Time) and the mean delay. At 95%

traffic intensity, a displacement of 10 minutes in the time at which the allocation is done causes a change in mean delay of about 8 minutes.




FIG B.1 EFFECT OF DISPLACEMENT FROM TIME HORIZON




FIG B.2 EFFECT OF ALLOCATION TIME WHEN NOT USING TIME HORIZON

APPENDIX C

TRAJECTORY PREDICTION - EXAMPLE

C.1 Introduction

Set out below is an illustration of the way that trajectory prediction is used to estimate the timing and position of occurrence of the significant events in an aircraft's flight path. In the example the steps involved in computing the trajectory for a specified speed and descent regime are shown for the last twenty-four minutes flying time to touchdown. This covers the period from before the aircraft leaves its cruising altitude. The illustration is based on a Boeing 737 aircraft flying the hypothetical London (EGLL) inbound route shown in figure C.1.

C.2 The Problem

Figure C.2 shows a general outline of the required vertical and speed profile, starting at a cruising level of flight level 290 and cruising speed of Mach 0.74. Referring to figure C.2, the initial descent from flight level 290 to 6000 ft altitude is to begin at cruising Mach 0.74, (equivalent to 284 kts CAS at flight level 290). Descent at constant Mach number results in a continuing increase in CAS, and once the airspeed has reached 290 kts CAS the descent is to be completed at a constant CAS of 290 Kts. The descent to 6000 ft is required to be complete some ten miles before the holding fix LAM. Deceleration to 250 kts CAS starts when level 10 miles before LAM. Glideslope interception is required to occur at 2500 ft altitude and the descent from 6000 ft to 2500 ft must be completed two miles before glideslope intercept occurs.

C.3 The Solution

The flight path is considered in two parts. One part comprises the route from the starting point at flight level 290 down to the bottom of the intermediate descent at 6000 ft altitude. The remainder of the flight path specification, from ten miles before the holding fix down to the runway is considered separately.

C.3.1 From Cruise Down to 6000 Ft Altitude

Since the end of descent point is defined (10 miles before LAM), backward prediction is used to predict back to the top of descent point. Having identified this point a forward prediction can then be made from the given start point to this calculated top of descent point.

C.3.2 From Ten Miles before LAM to Touchdown

A combination of forward and backward prediction is needed to economically complete the calculations. Forward prediction is used to calculate the time and distance involved in completing the deceleration to 250 kts CAS, starting ten

miles before LAM. It is convenient to predict back from the runway to find the point where the descent from 6000 ft begins. The time taken to traverse the remaining route mileage between the end of the deceleration to 250 kts CAS and the start of the descent to 2500 ft can be calculated using either forward or backward methods.

C.4 Implementation

The sequence of steps required to complete the trajectory design are shown in diagrammatic form in figures C.3 and C.4. A square label represents a demand for predictive activity and the attached arrow points to the route position described by the relevant prediction algorithm state variables when the prediction calculations start. When the demanded prediction calculations are complete the state variables point to a position identified by the arrow with an appended round label having the same numerical ident as the square label. The following description covers the sequence of events in greater detail and for illustration gives appropriate database values.

1F Initialise forward prediction state variables to position shown.

1F Resulting forward prediction state variables.

Range from next waypoint	45.00 n.mi
Next waypoint	LSD
Altitude	29000 ft
Airspeed	Mach 0.74, equivalent to 284 kts CAS
Cumulative flight time	0.0 seconds
Cumulative distance	0.00 n.mi

1B Initialise backward prediction state variables to position ten miles before LAM.

1B Resulting backward prediction state variables.

Range from next waypoint	10.00 n.mi
Next waypoint	LAM
Altitude	6000 ft
Airspeed	290 kts CAS
Cumulative flight time	0.0 seconds
Cumulative distance	0.00 n.mi

2 Predict backwards descent.

Command specification	
Terminating altitude	29000 ft

Speed regime	<i>Constant CAS at low altitude, changing to constant Mach 0.74 at high altitude</i>
Descent regime	<i>Idle thrust throughout</i>

② Resulting backward prediction state variables.

Range from next waypoint	18.80 n.mi
Next waypoint	<i>LSD</i>
Altitude	29000 ft
Airspeed	<i>Mach 0.74</i>
Cumulative flight time	569.1 seconds
Cumulative distance	58.98 n.mi

③ Predict forward in level flight.

Command specification	
Terminating condition	<i>position on route</i>
Position definition	<i>Current backward prediction route position</i>

③ Resulting forward prediction state variables.

Range from next waypoint	18.80 n.mi
Next waypoint	<i>LSD</i>
Altitude	29000 ft
Airspeed	<i>Mach 0.74</i>
Cumulative flight time	215.2 seconds
Cumulative distance	26.22 n.mi

④F Initialise forward prediction state variables to position ten miles before LAM.

④F Resulting forward prediction state variables.

Range from next waypoint	10.00 n.mi
Next waypoint	<i>LAM</i>
Altitude	6000 ft
Airspeed	290 kts CAS
Cumulative flight time	0.0 seconds
Cumulative distance	0.00 n.mi

④B Initialise backward prediction state variables to flight conditions at touchdown and position at runway threshold.

④B Resulting backward prediction state variables.

Range from next waypoint	0.00 NM
Next waypoint	EGLL
Altitude	0 ft
Airspeed	140 kts CAS
Cumulative flight time	0.0 seconds
Cumulative distance	0.00 n.mi

⑤ Predict forward deceleration.

Command specification	
Terminating airspeed	250 kts CAS
Descent regime	Level flight

⑤ Resulting forward prediction state variables.

Range from next waypoint	7.44 n.mi
Next waypoint	LAM
Altitude	6000 ft
Airspeed	250 kts CAS
Cumulative flight time	31.4 seconds
Cumulative distance	2.56 n.mi

⑥ Predict backward descent.

Command specification	
Terminating altitude	2500 ft
Descent regime	Constant slope, 300 ft per n.mi
Speed regime	Standard approach speed profile

⑥ Resulting backward prediction state variables.

Range from next waypoint	8.33 n.mi
Next waypoint	EGLL
Altitude	2500 ft
Airspeed	190 kts CAS
Cumulative flight time	183.7 seconds
Cumulative distance	8.33 n.mi

⑦ Predict backward in level flight

Command specification	
Terminating condition	Distance: 2 n.mi

⑦ Resulting backward prediction state variables

Range from next waypoint	10.33 n.mi
--------------------------	------------

Next waypoint	<i>EGLL</i>
Altitude	<i>2500 ft</i>
Airspeed	<i>190 kts CAS</i>
Cumulative flight time	<i>20.4 seconds</i>
Cumulative distance	<i>10.33 n.mi</i>

8 Predict backward descent

Command specification	
Terminating altitude	<i>6000 ft</i>
Descent regime	<i>Idle thrust</i>
Speed regime	<i>Standard approach speed profile.</i>
	<i>250 kts CAS when above flap</i>
	<i>deployment altitude</i>

8 Resulting backward prediction state variables

Range from next waypoint	<i>13.14 n.mi</i>
Next waypoint	<i>AA01</i>
Altitude	<i>6000 ft</i>
Airspeed	<i>250 kts CAS</i>
Cumulative flight time	<i>430.7 seconds</i>
Cumulative distance	<i>24.29 n.mi</i>

9 Predict forward in level flight

Command specification	
Terminating condition	<i>Position on route</i>
Position definition	<i>Current backward prediction</i>
	<i>route position</i>

9 Resulting forward prediction state variables.

Range from next waypoint	<i>13.14 n.mi</i>
Next waypoint	<i>AA01</i>
Altitude	<i>6000 ft</i>
Airspeed	<i>250 kts CAS</i>
Cumulative flight time	<i>167.3 seconds</i>
Cumulative distance	<i>12.85 n.mi</i>




Figure C.1 Inbound Route Plan View




Figure C.2 Trajectory Specification




Figure C.3 Prediction events from cruise down to 6000 ft altitude




Figure C.4 Prediction events from ten miles before LAM to touchdown

APPENDIX D

EXPERIMENTAL ENVIRONMENT – CONTROLLER INTERFACES

D.1 Flight Progress EDD Formats

The formats for the two sector types are shown in figures D.1 and D.2. The pending area is the part of the screen above the dashed separation line and the active area the part below. Both parts of the screen scroll independently and can grow or contract by the movement of the separation line to allow additions to either part. The elapsed time is shown at the top right in hours, minutes and seconds. Other times are shown in minutes and seconds in the form "*minutes m seconds*". A lower case "h" indicates a hold.

D.2 FPUD Formats

Figures D.3 to D.11 show the formats after various screen touches. The accept mark is shown as "ACC", the QSY mark as "Q" – both alongside the callsigns. The "LRATE" box is for updating the landing rate parameter and the "PARAMS" box for adjusting PVD scale and centre. The "QSY" box is for issuing a request to the next sector to transfer control of the aircraft. When this box is touched (followed by ENTER) an accept mark is produced on the appropriate sector's FPUD and a QSY mark shown on this FPUD. When the next sector controller accepts the aircraft by touching the accept mark, the QSY mark and callsign are deleted from this FPUD. All sequences are terminated by ENTER to complete an update or REJECT to cancel a sequence. Most menus allow errors to be corrected simply by touching the correct box, e.g. if a callsign is selected in error a new callsign may be immediately touched (provided no intervening touches have taken place). Once selected several fields within the flight strip data may be updated in a single sequence (including QSY).

Figure D.3 shows the menu rest picture. Up to ten callsigns can be displayed simultaneously on the screen. Thereafter a marker appears on the screen to indicate that there are further callsigns to be displayed. Aircraft are deleted from the screen as they land or are transferred to another sector (and accepted). When aircraft are deleted any callsigns awaiting display are immediately drawn. A touch sequence is initiated by touching any of the following boxes:-

- | | | |
|-----------------|---|---|
| Callsign | - | Touching this field initiates a flight progress data update. |
| Callsign Accept | - | This is touched to accept an aircraft. Touching the callsign is invalid until the aircraft has been accepted in this way. |
| LRATE | - | This is touched to enter a new landing rate value |
| PARAMS | - | This is touched to change the radar picture scale and/or centre |

Whilst a sequence is in progress no additions or deletions to the callsign list are allowed; this would be very confusing to the operator and could lead to mistaken data entry. The callsign list scrolls⁵ as new aircraft are inserted or as aircraft are removed. The ordering of the callsign list is specified at set-up time (i.e. fixed for a given run). Time is shown at the top of the screen in hours, minutes and seconds and is continually updated.

Figure D.4 shows the format after a callsign touch has been made. A copy of the aircraft data line is displayed at the top of the screen. Most fields of the strip are touch-sensitive; these are marked in the figure by a dagger. Touching one of these fields leads to subsequent menus as described below.

Figure D.5 shows the picture after two aircraft sequence numbers have been exchanged (SK469 and SK886). To do this the touch sequence is: first callsign, sequence number field of flight strip, second callsign and then ENTER (which causes a reset to the rest picture). If the callsign list was ordered on sequence number then the two callsigns would also be redrawn in swapped positions on the screen.

Figure D.6 shows the picture after touching the cleared height field of the flight strip and selecting a height value. Having made a wrong height value selection another value may be immediately touched to correct.

Figure D.7 shows the picture after touching the TAT field of the data line. Touching the TICK box followed by ENTER causes the EDD to be marked with a blue cross to indicate that the SDT has been communicated to the aircraft.

Figure D.8 shows the picture after touching the cleared speed field of the flight strip. Touching the "C" box causes the planned speed (if set) to be shown on the speed digit pad. Other speeds may be entered by touching appropriate digits. Corrections are made by simply touch reselecting the digit(s) in error. The "R" box is touched to request a new speed for an aircraft – the new speed will be shown on the EDD.

Figure D.9 shows the picture after touching the cleared heading field of the flight strip and selecting heading "265" degrees. Touching the "X" box sets the heading field to blanks both on the FPUD and the EDD. Note that invalid headings (i.e. $> 360^\circ$) are automatically coerced to valid ones by changing the background highlighting accordingly. For example, with the value "265" set an attempt to change the first digit to a "3" would cause the third digit to be set to "0" instead of "5".

Figure D.10 shows the picture after touching the LRA¹⁰ box and making a selection to update the landing rate.

Figure D.11 shows the menu used to change the centre and scale for the PVD. Notice that a circular picture is assumed.

Figures D.12 to D.20 give the finite state transition diagrams for the FPUD and thus define the valid touch inputs for each state. The states are shown by circles containing the name of the state. The transition arcs are generally labelled with the name of the input which causes the transition. States are defined in a

⁵ Although this technique was used satisfactorily for the experiments, it was considered that callsigns should remain fixed in position. Occasionally the wrong callsign was touched if the list scrolled just as a touch was being made

hierarchy with Level 1 as the topmost level. Note that, for clarity, some transitions are not shown. In particular, the inputs ENTER and REJECT in any state cause reversion to the rest state. Also QSY state can be entered from any Level 3 or Level 4 state.

00:22:50

Landing sequence number				Terminal Approach time			Cleared CAS	
	Sector Designator	Stack Beacon	Cleared Flight level	TAT or Stack ETA			Planned CAS	
		Callsign	Aircraft Type					Cleared Heading
05	CCVY	SK469	DC9	280	37m30		240k	
08	CCVY	AF202	EA30	↓ 280	37m20/42m20h		270k	

03	CCVY	FW0TP	DA20	↓ 100	33m10		225k	
02	CCVY	BR908	B727	↓ 160	30m20		300k	
01	CCVY	BA069	B737	080	29m00		250k	265

Figure D.1 TAF Inbound Sector Flight Progress EDD

00:22:50

Landing sequence number	Sector Designator				Top of descent Time	Cleared CAS		
	Stack Beacon		Cleared Flight level			Planned CAS		
	Callsign							
			Aircraft Type				Cleared Heading	
	CCVY	BA835	B737	350	25m50			
	CCVY	BA032	EA30	240	29m50			

08	CCVY	AF202	EA30	↓ 280	21m20	270k	280	
07	CCVY	RJ114	B757	↓ 190	21m30	240k	280	
06	CCVY	SK886	DC9	240	25m40	230k		
05	CCVY	SK469	DC9	280	22m00	240k/230k		
03	CCVY	FW0TP	DA20	↓ 100	19m00	225k		
02	CCVY	BR908	B727	↓ 160	17m50	300k		

Figure D.2 Enroute Sector Flight Progress EDD

00:22:50					
ENTER					
FWOTP	ACC				
BR908					
	ACC	QSY	LRATE	PARAMS	REJECT
BA069					

Figure D.3 Rest picture showing callsigns and accept marks

01†	CVY	BA069	B737	↓80†	29m00	†	250k†	00:22:50 265†
								ENTER
FWOTP		ACC						
BR908								
		QSY		LRATE		PARAMS		REJECT
<u>BA069</u>								

Figure D.4 Picture after callsign BA069 touched

† Touch-sensitive Data-line Fields

06	CVY	SK469	B747	↓180	23m00	240k	00:22:50 ***
							ENTER
<u>SK886</u>							
<u>SK469</u>							
FWOTP	Q	QSY	LRATE	PARAMS	REJECT		
BR908	Q						

Figure D.5 Picture after sequence number exchange touch sequence

05	CVY	SK469	B747	<u>↓180</u>	23m00	240k	00:22:50 ***
					190	290	390
					<u>180</u>	280	380
					170	270	370
					160	260	360
					150	250	350
					140	240	340
					130	230	330
SK886		ACC			120	220	320
<u>SK469</u>					110	210	310
FWOTP		Q			100	200	300
BR908		Q	QSY				REJECT

Figure D.6 Picture after touching height field showing pad for entering heights

[illegible]

Figure D.7 Picture after touching TAT field

05	CVY	SK469	DC9	280	23m00	240k	00:22:50

						9 9	
						8 8	ENTER
				R		7 7	
						6 6	
				C		5 5	
						<u>4</u> 4	
SK886						3 3 3	
<u>SK469</u>						<u>2</u> 2 2	
FWOTP						1 1 1	
						0 0	
BR908		QSY					REJECT

Figure D.8 Picture after touching speed field with number pad for entering speeds. Note the 'C' box for entering computer-advised speed and the 'R' box for requesting a revised speed.

01	CVY	BA069	B737	↓80	29m00	250k	00:22:50 265
	9	9					
	8	8					ENTER
	7	7					
	6	6					
	5	5					
	4	<u>4</u>					
	<u>3</u>	3					
	2	2					
	1	1					
	0	0	QSY	<u>LRATE</u>	PARAMS	REJECT	

Figure D.10 Picture after LRATE touched showing number pad for entering landing rates.

00:22:50

		9	9	9	
<u>XCENTRE</u>	=	25	8	8	ENTER
			7	7	
			6	6	
<u>YCENTRE</u>	=	15	5	5	<u>5</u>
			4	4	
			3	3	
<u>RADIUS</u>	=	45	2	<u>2</u>	2
			1	1	1
<u>+</u>		-	<u>0</u>	0	0
					REJECT

Figure D.11 Picture after PARAMS touched showing menus for changing picture scale.




FIGURE D.12 TOP LEVEL STATE TRANSITIONS




FIGURE D.13 FLIGHT STRIP STATE TRANSITIONS




FIGURE D.14 SEQUENCE STATE TRANSITIONS




FIGURE D.15 HEIGHT STATE TRANSITIONS




FIGURE D.16 SPEED STATE TRANSITIONS




FIGURE D.17 HEADING STATE TRANSITIONS




FIGURE D.18 TAT STATE TRANSITIONS




FIGURE D.19 LANDING RATE STATE TRANSITIONS




FIGURE D.20 PARAMETERS STATE TRANSITIONS

GLOSSARY

Air Traffic Simulator The suite of software which provides a set of aircraft which fly according to specified performance models, along defined routes as directed by a traffic sample. Facilities are provided to control the aircraft via aircraft control terminals.

Allocated Landing Time The landing time allocated by the LOC or SCA.

Allocation This term is used when arriving aircraft are assigned landing times and associated sequence numbers.

Approach Sequencing Area The block of airspace between the holding stacks (or TAFs) and the runway.

ATCEU Air Traffic Control Evaluation Unit.

ATCO Air Traffic Control Officer.

Automatic Controller The program which simulates the necessary functions performed by controllers on sectors which can not be manned during the experiments. No conflict resolution is performed.

CAA Civil Aviation Authority.

CAS Calibrated air speed.

CCF Central Control Function. The new system of airspace organisation, ATC procedures and equipment planned for the the London Terminal Area in the early 1990s.

Delay The difference between the preferred and allocated landing times (note that this may be negative).

Discrete Event Simulation A simulation method where state variables change only at discrete points in time, and no attempt is made to characterise continuous change. The simulation clock is not synchronized with real time.

EDD Electronic data display. This refers specifically to the flight progress data device used by the landing order calculator and speed control adviser.

FCFS First-come-first-served

FCFS Order Method The method of allocating landing times based on the order of crossing a single time horizon.

FPUD Flight Progress Updating Device used to update the EDD.

Hold same as Stack.

IATA International Air Transport Association.

Inner Time Horizon The point in time at which landing time allocation occurs for arriving aircraft in the LOC and SCA in the OptO method. It is measured as a fixed time before the preferred landing time.

Landing List A list of aircraft in chronological order of allocated landing times.

Landing Rate The assumed arrival capacity for an airport used to plan landing times. The figure is based on the setting entered via the FPUD at a sector position and it is intended to be used only to adjust (increase or decrease) the flow into the approach sequencing area.

LATCC London Air Traffic Control Centre.

LOC Landing Order Calculator.

MAXIMUM ABSORB TIME A parameter used by the arrivals planning process to define how much delay it is feasible to absorb within the approach sequencing area for a particular TAF routing.

NATS National Air Traffic Services.

OptO Method The method of landing times allocation used by the speed control adviser that attempts to produce an optimised landing sequence based on reordering from the first-come-first-served arrival order taking account of turbulent wake categories.

Order Displacement Index A value used in the OptO allocation method to determine the total disturbance from the FCFS order for a proposed sequence of aircraft. The index is computed by counting the number of place shifts for each aircraft and summing these values for all aircraft in the proposed sequence. The larger the index the greater the total disturbance.

Outer Time Horizon The point in time at which aircraft become eligible for inclusion in the landing sequence optimisation process - applicable only to the OptO method. It is measured as a fixed time before the Preferred Landing Time.

Preferred Landing Time The predicted time of arrival at the runway assuming the aircraft flies its preferred profile with no intervention from the controller.

PVD Plan view display used to show radar picture.

RSRE Royal Signals and Radar Establishment

SAT Predicted stack arrival time (term used in LOC).

SCA Speed Control Adviser.

SCCS Skeleton Control Centre Software. This is the part of the real-time environment in which the LOC and SCA were exercised.

SDT see Stack Departure Time.

Stack A set of aircraft flying orbits at a designated fix according to ATC procedures which ensure safe separation.

Stack Departure Time The predicted departure time from the stack when one or more orbits is flown. Analogous to TAT in the holding case.

TAF see Terminal Approach Fix

TAS True air speed.

TAT see Terminal Approach Time

TCSDG Terminal Control Systems Development Group

Terminal Approach Fix A point on the edge of the approach sequencing area used to mark the beginning of the approach control function. The TATs apply to these points and there are generally several per airport.

Terminal Approach Time The predicted time of departure from the designated terminal approach fix.

Time Horizon The point in time at which landing time allocation occurs for arriving aircraft in the LOC and SCA. It is measured as a fixed time before the Preferred Landing Time.

Time Navigation System A flight management system which permits an aircraft to arrive at a given position at a specified altitude and time (also known as 4D navigation).

Traffic Intensity The ratio between mean arrival rate and mean landing rate where there is traffic queueing to land.

Traffic Sample Generator An offline program which uses a statistically-defined specification to produce a traffic sample with a pseudo-random distribution of aircraft types, cruising altitudes, routes and callsigns.

Trajectory Prediction The process of predicting an aircraft's flight path as a function of time taking account of manoeuvres in response to ATC instructions, routing and wind.

VDU Visual Display Unit.

DOCUMENT CONTROL SHEET

Overall security classification of sheet ... UNCLASSIFIED

(As far as possible this sheet should contain only unclassified information. If it is necessary to enter classified information, the box concerned must be marked to indicate the classification eg (R) (C) or (S))

1. DRIC Reference (if known)	2. Originator's Reference REPORT 88001	3. Agency Reference	4. Report Security Classification U/C	
5. Originator's Code (if known) 7784000	6. Originator (Corporate Author) Name and Location ROYAL SIGNALS & RADAR ESTABLISHMENT ST ANDREWS ROAD, GREAT MALVERN WORCESTERSHIRE, WR14 3PS			
5a. Sponsoring Agency's Code (if known)	6a. Sponsoring Agency (Contract Authority) Name and Location			
7. Title COMPUTER-BASED TOOLS FOR ASSISTING AIR TRAFFIC CONTROLLERS WITH ARRIVALS FLOW MANAGEMENT				
7a. Title in Foreign Language (in the case of translations)				
7b. Presented at (for conference papers) Title, place and date of conference				
8. Author 1 Surname, initials MAGILL, S.A.N.	9(a) Author 2 TYLER, A.C.F.	9(b) Authors 3,4... WILKINSON, E.T.	10. Date 1988.9	pp. ref. VP
11. Contract Number	12. Period	13. Project	14. Other Reference	
15. Distribution statement UNLIMITED				
Descriptors (or keywords)				
continue on separate piece of paper				
Abstract This document reports a research project on the use of computer-based tools for arrivals flow management at major airports. It describes in detail two experimental prototype tools which were constructed as part of the research project, and exercised in a real-time simulation environment. These were a Landing Order Calculator and a Speed Control Adviser. Both give advice to air traffic controllers for aircraft which are in the vicinity of their top-of-descent points. The underlying concepts and the methods used in the tools are fully described. The real-time simulation environment and the experience gained from it are discussed. <i>(Text Continues on page 2)</i>				