

A FLEXIBLE SIMULATION MODEL
OF
AIRPORT AIRSIDE OPERATIONS

by

JOHN PHILIP NORDIN

B.S., Michigan State University
(1976)
S.M., Massachusetts Institute of Technology
(1978)

SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS OF THE
DEGREE OF

DOCTOR OF PHILOSOPHY IN
CIVIL ENGINEERING

at the


MASSACHUSETTS INSTITUTE OF TECHNOLOGY

August 1980


© John Philip Nordin 1980

The author hereby grants to M.I.T. permission to reproduce
and to distribute copies of this thesis document
in whole or in part.

Signature of Author


Department of Civil Engineering
August 29, 1980

Certified by


Amedeo R. Odoni
Thesis Supervisor

Accepted by

ARCHIVES
MASSACHUSETTS INSTITUTE
OF TECHNOLOGY
Chairman, Departmental Graduate Committee

OCT 9 1980

LIBRARIES

NOTICE: THIS MATERIAL MAY BE
PROTECTED BY COPYRIGHT LAW
(TITLE 17 US CODE)

Abstract

A Flexible Simulation Model of Airport Airside Operations

by

John Philip Nordin

Submitted to the Department of Civil Engineering on
August 22, 1980 in partial fulfillment of the
requirements for the degree of Doctor of Philosophy

The expanding problem of airport congestion emphasizes the need for general purpose models of airport performance. Previous efforts to develop simulations that could analyze a wide range of airport geometries have not been entirely satisfactory. The resulting models could not treat significant problems in airport analysis. This work develops, from a new foundation, a flexible simulation model of airport operations. A number of contributions are made to modeling of airport performance. A new method for describing the demand for service is developed which avoids a significant bias of existing methods. A procedure for modeling landing operations permits analysis of the effect of exit location on overall airport congestion. Methods for separating aircraft are developed which can describe the actual behavior of the airport with greater accuracy than previous models and methods for modeling changes in operating policies in response to weather and congestion also are developed. Numerous additional optional provisions and the overall design of the model offer flexibility to the analyst and improve computational efficiency. These expanded capabilities permit the wide use of the model not only in analyzing existing conditions but also in studies of optimal design of airports. Tests of the model against several other airport models show this model to be satisfactory. This study of airport operations and modeling techniques has also resulted in the identification of a number of important topics for future research.

Thesis Supervisor: Amedeo R Odoni

Title: Professor in Aero & Astro

ACKNOWLEDGEMENTS

I wish to thank Professor Amedeo Odoni, Department of Aeronautics and Astronautics, M.I.T., for his advice and direction as my advisor over the last three years. His painstaking and prompt reading of each draft of this work was greatly appreciated. Professor Robert Simpson, Department of Aeronautics and Astronautics, M.I.T. gave valuable advice based on his extensive knowledge of airport operations. Professor Nigel Wilson, Department of Civil Engineering, M.I.T., asked many difficult questions, which hopefully have now been answered. Professors Odoni, Simpson and Wilson are also to be thanked for serving on my Ph.D. committee.

Professor William Dunsmuir, Department of Mathematics, M.I.T., was generous with his time and helped with a number of statistical questions. Professor L.E. Grosh, Department of Industrial Engineering, Kansas State University, provided guidance to the simulation literature and loaned me a number of items out of his personal library which were unavailable at M.I.T. Mr. William Hoffman, Flight Transportation Associates, Cambridge, Massachusetts, helped in defining the case study in Chapter IV.

Mr. Art Anger, Information Processing Center (IPC), M.I.T., provided uniquely accurate and unfailingly friendly advice on programming questions. Mr. Joseph Connors, Director of Purchasing, IPC, helped me obtain several important items. Mr. Joe Hevey and all the employees of User Accounts, IPC, have given me excellent and friendly service these past four years. Ms. Victoria Murphy, Graduate Program Administrator, Department of Civil Engineering was extremely indulgent with deadlines and provided good advice about the thesis requirements.

Mr. Harold Clark, Department of Materials Science, M.I.T. applied his amazing command of English grammar to greatly improve the readability of this work. Ms. Marjorie Shane, Department of Chemistry, M.I.T., assisted with figure preparation, insured continuity and generally cleaned up what the text editor could not handle. My wife, Barbara Grosh, overcame her natural antipathy to typing and performed excellently as "data entry operator", as well as acting as general editor and stylistic critic. I incurred a large debt to her which I expect to repay in several years when she finishes her Ph.D. thesis. A great deal of thanks goes to Dr. Linda Anthony, Bell Laboratories, Murray Hill, New Jersey, for all of her "thesis stories" and for proving that it was possible to complete a Ph.D. thesis.

And finally, I wish to acknowledge whoever it is that manufactures fortune cookies for causing me to receive the following fortune during the final weeks of my thesis: "Work hard and you will achieve your goals."

TABLE OF CONTENTS

	Page
Abstract.....	2
Acknowledgements.....	3
Table of Contents.....	4
List of Figures.....	7
List of Tables.....	10
Chapter I. Introduction.....	12
A. The Problem of Modeling Airport Operations.....	12
B. Approaches to Modeling Airport Operations.....	14
C. Contributions to Modeling Airport Operations.....	16
D. Overview of Chapters II Through V.....	20
Chapter II. Model Structure.....	23
A. Major Issues.....	23
B. Airfield Geometry and Ground Operations.....	27
1. Introduction.....	27
2. Description of airfield Geometry.....	28
3. Ground operations.....	32
C. Demand for Service.....	42
1. Introduction.....	42
2. Present schedule algorithms.....	44
3. Aircraft attributes.....	47
4. Modeling the level of demand.....	61
5. Comparison of the methods of demand generation.....	73
D. Runway Operations.....	83
1. Introduction.....	83
2. Aircraft performance on landing.....	85
3. Aircraft performance on takeoff.....	111
4. Separation of aircraft movements.....	113
E. Dependent Runway Operations.....	138
1. Introduction.....	138
2. Intersecting runways.....	140

3.	Parallel runways.....	149
4.	Projected intersecting runways.....	154
5.	Assignment of aircraft to runways.....	156
6.	Special problems.....	163
7.	Comparison of separation methodologies.....	170
F.	Control of the Airport.....	177
1.	Introduction.....	177
2.	Weather.....	178
3.	Congestion.....	186
4.	Other factors.....	190
G.	Statistical Issues in Modeling Airport Operations.....	191
1.	Introduction.....	191
2.	Survey of statistical issues relevant to simulation.....	191
3.	Selection of simulation outputs.....	209
Chapter III.	Model Validation.....	215
A.	The Problem of Validation.....	215
B.	Test Case: La Guardia.....	221
1.	Comparison to FAA capacity analysis.....	221
2.	Comparison with ASM and MITASIM.....	229
C.	Test Case: Bromma Airport.....	244
1.	Case description.....	244
2.	Discussion of results.....	251
D.	Conclusions.....	256
Chapter IV.	Example of Model Application.....	257
A.	Introduction.....	257
B.	The Airport.....	258
C.	Description of Operating Policies.....	262
D.	Characteristics of Demand.....	267
1.	Aircraft classes.....	267
2.	Aircraft performance.....	268
3.	Aircraft mix.....	268
E.	Air Traffic Control Rules.....	270
F.	Input Requirements.....	275
1.	Summary of input requirements.....	275
2.	Comparison of input requirements.....	276

G.	Capacity Analysis.....	280
1.	VFR results.....	280
2.	IFR results.....	290
H.	Delay Analysis.....	292
1.	Case description.....	292
2.	Discussion of results.....	297
I.	Model Properties.....	301
1.	Sensitivity Analysis.....	301
2.	Computational Efficiency.....	303
J.	Conclusions.....	305
Chapter V. Summary and Conclusions.....		307
A.	Review.....	307
B.	Further Work.....	311
1.	Needed research.....	311
2.	Potential uses for FLAPS.....	314
3.	Extensions to FLAPS.....	319
Bibliography.....		321
Biographical Note.....		325

LIST OF FIGURES

Figure No.		Page
II-1.	New York La Guardia: Airport Map.....	29
II-2.	New York La Guardia: ASM Modeling Geometry.....	30
II-3.	Selection of Aircraft Attributes.....	60
II-4.	Rate Function.....	62
II-5.	Rate Function and Integrated Rate Function.....	66
II-6.	Summary of Nomenclature.....	68
II-7.	Numerical Example.....	72
II-8.	Comparison of Schedule Generation Methods: Parameters.....	80
II-9.	Runway Operations.....	84
II-10.	Landing Aircraft Profile.....	90
II-11.	Interrelationships of Aircraft Floating Phase Parameters.....	95
II-12.	Effect of V_a on Floating Phase Parameters.....	97
II-13.	Arrival/Arrival Separations: Closing Case.....	121
II-14.	Arrival/Arrival Separations: Opening Case.....	123
II-15.	Separations for Mixed Operations on a Single Runway.....	130
II-16.	Dulles International Airport, Washington, D.C.....	141
II-17.	Intersecting Runway Departure Operations.....	144
II-18.	Mixed Operations on Intersecting Runways.....	145
II-19.	Airports with Projected Intersecting Runways.....	155

List of Figures, continued

Figure No.		Page
II-20.	Displaced Runway Thresholds.....	164
II-21.	Hold Short Arrivals.....	167
II-22.	Distribution of Average Delay.....	204
II-23.	Distribution of Number of Landings.....	205
III-1.	La Guardia: FLAPS Geometry Representation.....	222
III-2.	La Guardia Configurations.....	224
III-3.	Sensitivity Run 1A: Base Case.....	234
III-4.	Sensitivity Run 2A: Group 2 Separations.....	235
III-5.	Sensitivity Run 3A: Group 4 Separations.....	236
III-6.	Sensitivity Case 7A: $Q = 24$	237
III-7.	Sensitivity Case 8A: $Q = 1$	238
III-8.	Bromma: Demand Profile.....	248
III-9.	Bromma: Demand Functions.....	250
III-10.	Bromma Airport: 400 op./day.....	253
III-11.	Bromma Airport: 600 op./day.....	254
III-12.	Bromma Airport: 800 op./day.....	255
IV-1.	Boston Logan Airport.....	259
IV-2.	Logan Airport: FLAPS Geometry Representation.....	264
IV-3.	Logan Configurations.....	266
IV-4.	Logan VFR Capacity Input File.....	277
IV-5.	Interpolation-Extrapolation Method for "Percentage Arrival" Capacity Estimation.....	288

List of Figures, continued

Figure No.		Page
IV-6.	Logan Delay Configurations.....	294
IV-7.	Logan Delay Cases: Demand Function.....	295
IV-8.	Logan Delay Cases: Summary of Parameter Changes.....	296
IV-9.	Logan Delay Cases: Landing Delay.....	298
IV-10.	Logan Delay Cases: Takeoff Delay.....	299

LIST OF TABLES

	Page
Table II-1. Probabilistically Generated Schedules.....	76
Table II-2. Base Method Schedules.....	77
Table II-3. Comparison of Schedule Generation Method: Results.....	81
Table II-4. Float Phase Parameters.....	92
Table II-5. Landing Model Validation: Common Parameters.....	106
Table II-6. Landing Model Validation: Case Description.....	107
Table II-7. Landing Model Validation: Results.....	109
Table II-8. Terms used in Derivation of Separations.....	122
Table II-9. Arrival/Arrival Separations: Typical Values.....	126
Table II-10. Departure/Departure Separations: Typical Values.....	129
Table II-11. Terms used for Mixed Operation Separations.....	131
Table II-12. Mixed Operations: Summary of Rules.....	137
Table II-13. Dependent Runway Nomenclature.....	139
Table II-14. Normality Test: Summary Results.....	202
Table III-1. La Guardia Validation: Capacity Parameters.....	225
Table III-2. La Guardia Validation: Capacity Results.....	226
Table III-3. Sensitivity Runs: Parameters.....	232
Table III-4. Sensitivity Runs: Summary Statistics.....	239

List of Tables, continued

Table III-5.	Average Departure Queue Length Statistics.....	242
Table III-6.	Practical Separation Minima at Bromma Airport.....	245
Table III-7.	FLAPS Equivalent Separations.....	247
Table III-8.	Bromma: Summary Comparison of Results.....	251
Table IV-1.	Logan Airport: Geometry Elements.....	261
Table IV-2.	Logan Airport: Runway Modules.....	263
Table IV-3.	Assumed Aircraft Performance Parameters.....	269
Table IV-4.	Logan Airport: Aircraft Mix and Runway Assignment.....	271
Table IV-5.	Logan Airport: Separation Parameters.....	272
Table IV-6.	Air Traffic Control Dependencies.....	274
Table IV-7.	Logan VFR Results.....	282
Table IV-8.	IFR Capacity Estimates.....	291

I INTRODUCTION

A THE PROBLEM OF MODELING AIRPORT OPERATIONS

In the late 1960's severe congestion problems at several major air carrier airports attracted national attention. This problem eventually receded due, in part, to wide-body jets reducing the number of flights needed and to the economic slowdown of the mid-1970's. Following the demand surge of 1978, however, reports are being heard with increasing frequency and from a growing number of locations of a rising incidence of congestion and delay (LEF 80). For example, delays of 30 minutes or more increased 300% at Los Angeles from the first half of 1977 to the first half of 1978 (SWE 79 p.1-1). Delays of 30 minutes or more increased from 3 to 6 per thousand operations over the entire United States airport system from 1978 to 1979 (ATS 1980). The prevalence of delays is cited as one restraint on airline growth (SWE 79 p.1-1). The increase of airline fuel costs by 83% from December 1978 to December 1979 (ATW 80) raises significantly the cost of forcing aircraft to fly holding patterns while waiting to land.

The significance of the problem of airport congestion makes apparent the necessity for modeling airport performance as a function of demand for service and capabilities of airport facilities. Delay and congestion,

of course, are caused by many different factors, and it will require action on several fronts to control them. However, any significant change in the airport environment will have to be assessed as to its impact on airport performance before implementation. Because of the complex nature of airports and the amount of resources necessary to implement all but the most modest change, it would seem obvious that careful mathematical modeling of the consequences of each proposed change would be necessary.

The modeling of airport performance poses an interesting challenge to transportation analysts. Airports, like other complex transportation facilities, do not always behave in simple or intuitively obvious ways under fluctuating loads. The fact that each airport is different despite being made up of the same generic components and is subject to different loads, complicates the analysis. Flows of aircraft intersect which makes it difficult to accurately analyze portions of the airport independently. Increasing the capacity of one flow of aircraft may have the effect of restricting other flows that must be fitted between aircraft of the first flow.

Analysis can not be confined to predictions of steady-state results but must model the dynamic aspects of the airport. Demand for service is strongly time-varying. The characteristics of aircraft using the airport may also vary significantly over the course of a day.

Another level of analytical problems is introduced by the fact that the rules under which airports operate can change during a day. Priorities of various operations, runways in use, and minimum separations permitted between aircraft can vary in response to weather, congestion or equipment failures.

While the airport must be modeled in its entirety and with attention to its dynamic environment, information is needed not only about overall results but about performance during each time period and at each major facility of the airport. The need is for analytical tools that predict the performance of an airport (average delay, peak delay, average number of aircraft waiting to takeoff and land, number of operations) under a variety of loads (aircraft flow rates, types of aircraft) as well as estimate the change in performance due to altering the number or type of facilities or the rules for the use of facilities.

B APPROACHES TO MODELING AIRPORT OPERATIONS

A number of approaches have been taken to the problem of predicting the performance of the airside of airports. These approaches have included analytical capacity and delay formulas, special purpose simulations tailored to one particular airport and general purpose simulations. We term an airfield simulation, "general purpose", or "flexible" if

it can model a wide range of airport geometries, operating rules and service patterns.

Of the three approaches, a general purpose simulation has the widest potential usefulness. The tradeoffs between simulation and analytical models are well established (TEI 66 p.724, SHA 75 pp.10-13, FIS 73 p.4-5). One of the major advantages of a simulation over analytical models is the capability of a simulation to model the transient conditions and changes in operating rules that are an inherent part of the questions to be analyzed at airports. A simulation can also accurately model special situations that can only be approximated by analytical approaches. A general purpose simulation would also imply significant cost savings compared to a series of special purpose simulations that together might provide the same capability as the general purpose simulation. The widespread desire in airport analysis for answers to performance related questions insures that a general purpose model would have extensive application.

Efforts to develop a model flexible enough to analyze a wide range of airport geometries and operating policies have been only partially successful. Only one model developed by Peat, Marwick, Mitchell and Company for the FAA was intended to be truly general purpose (HOC 77, PMM 77, BAL 76). This model is commonly referred to as the airfield simulation model or ASM. Unfortunately, this model, an

evolutionary descendant of a model first developed more than ten years ago, was designed to examine only a limited range of operating policies. It requires much important analysis to be performed externally to the model. Because of its fundamental design, using the ASM model imposes large time and cost requirements on the user (NOR 78 pp.34-60). The specific features of the ASM model that support this conclusion are discussed in detail in chapter II below.

C CONTRIBUTIONS OF THIS WORK TO MODELING AIRPORT OPERATIONS

This work develops a flexible tool for general use in analyzing performance related issues at airports. It is directed towards meeting the needs of airport analysis and overcoming the serious limitations of existing models.

The model follows the movement of aircraft from the start of final approach through landing and ground movements then to the point of clearing the runway on takeoff. For convenience it is referred to in this document as FLAPS, for FLExible AirPort Simulation. The objective of this work was to strive for significant advances in the modeling of airport operations in order to achieve a comprehensive model that is both theoretically correct and economical enough to reduce the cost of such analysis to feasible levels.

The model that was developed, implemented and validated here extends the modeling of airport operations to include

for the first time in a general purpose simulation model the following aspects of the dynamics of airport operation:

1) The effect of changes of the runways used for operations over the course of a day. This includes being able to switch direction of operation on any given runway. Previous models were not able to represent an entire day's operation during the course of a single run because of limited capabilities in this area (see section II.F).

2) The impact of changes in exit location, exit type and additions of new exits on runway performance of landing aircraft. Previous models required that this analysis be performed externally to the model and the resulting performance to be an input for use in the general purpose model. (See section II.D.)

3) The impact of dynamically changing operating rules to minimize delay. Aircraft are often given runway assignments to balance queue lengths. The relative priorities of conflicting aircraft movements are often altered by air traffic controllers over the course of a day in order to eliminate long queues and take advantage of gaps in the flows of aircraft. Previous models require such rules to be held constant over the duration of a run.

4) The effect of wind speed and direction and rainfall on aircraft performance. Weather has a major impact on both the rules of operation of airports and the performance of aircraft but it has not been explicitly modeled. As in (2) above, the analysis had to be performed externally to the model and input.

5) The interaction of aircraft performance and separation logic. Previous models did not accurately handle cases involving aircraft on two different runways where the air traffic control separation is based on aircraft clearing an intersection or exiting from the runway. (See section II.E.)

6) Expanded techniques for generation of aircraft schedules. Included in FLAPS are both the standard technique of using a base schedule and a new probabilistic way of generating schedules. As explained in section II.C, the two techniques have different underlying assumptions. The base schedule method is suitable for short term analysis whereas the probabilistic method is more suited for long term prediction. The inclusion of the second method expands the type of problems that can be analysed. This new method facilitates analysis of the effect of changes in the type of aircraft using the airport over the course of a day. This capability permits the analysis of operating policies such

as restrictions on general aviation aircraft at peak periods.

7) Proposition of a simple method of modeling ground operations of aircraft. This method can approximate the actual taxiing delay without detailed modeling of such movements. Previous models used an extremely complex method of modeling taxiing operations.

8) Collection and computation of a full set of necessary outputs to support statistical testing of various hypotheses. This includes disaggregate as well as summary statistics and time series outputs.

The inclusion in a general purpose model of the dynamics indicated above results in an enhancement of the precision of the overall estimates of performance and in a significant expansion in the types of policies that can be analyzed compared to previous models. The aggregate effect of points (1) to (3) above is that a model is now available that can be used, not only for overall capacity and delay studies, but to assess various dynamic strategies for managing congestion at the airport.

The model has been implemented as an event based simulation. The increased capabilities of this model did not result in an increase in the resources required to use

the model compared to previous models. New techniques for specifying airfield geometry elements were developed (see section II.B) that result in a reduction of about 75% in resources required to prepare and use the simulation.

D OVERVIEW OF CHAPTERS II THROUGH V

Chapter II contains a detailed discussion of the various problems in the analysis of airport operations. Because the set of issues that arise in modeling airports is large, chapter II is organized on a topic-by-topic basis. Thus there is no separate review of previous work. Rather, this material, primarily the existing ASM model, is discussed in parallel with the presentation of the new model. After the problems associated with a particular topic are reviewed, the portion of FLAPS dealing with that topic is described.

A satisfactory degree of confidence in the validity of FLAPS was achieved by several approaches. The validity of the assumptions used in FLAPS is defended as they are introduced. The important landing performance sub-section was tested independently on two sets of data and found to give good results. Estimates of capacity and delay from FLAPS were compared to estimates from other models and to observed data. These tests are reported in chapter III. A variety of tests with data from two airports show that FLAPS can model airport performance acceptably.

To illustrate the scope of the capabilities of FLAPS, chapter IV presents an example of the use of FLAPS to analyze airport performance. A complete description of the selection and derivation of input parameters is given. Capacity and delay estimates are made for several runway configurations. Several runs are made to demonstrate various capabilities of the model.

In the course of conducting the study necessary for developing this model, a number of topics for further work were identified. These are discussed in chapter V. Studies of landing operations are needed to collect a comprehensive data set. Very few studies of any type on takeoff movements have been found. Collection of accurate delay statistics is still in its infancy. Work is hampered by the fact that much research on airport performance is not published in a conveniently available form but circulates as memorandums, unpublished papers and the like.

There are a number of areas of study in which FLAPS can be directly used. Capacity estimates are needed periodically by most airport authorities, as is detailed study of any contemplated change. Several research possibilities using FLAPS can be suggested. These include experiments on how changes in landing performance, exit location, and separation standards affect overall airport performance, and assessing the optimality of the various control strategies used by controllers to manage congestion

and delay at the airport. Several models of terminal airspace and of groundside operations have been developed. These could be operated in tandem with FLAPS to assess the interaction of operations in each area.

There are a number of logical extensions to FLAPS, also described in chapter V, that could enhance its usefulness. The extension with the highest priority is probably the implementation of the method of modeling ground operations, suggested in section II.B.

II MODEL STRUCTURE

This chapter describes in detail the major conceptual components of aircraft operations at airports. Current methods of modeling each component are discussed and evaluated. Where adequate methods exist, they are described and incorporated into FLAPS. Where no adequate methods currently exist, a new methodology is developed. Validation of the individual features of FLAPS is described. We begin by describing what the major components of aircraft operations at airports are and how they relate to each other and to the analysis problem.

A MAJOR ISSUES

Modeling airport operations requires analysis of many subjects. This section presents a brief overview of the scope of modeling issues addressed in depth in the following sections of chapter II.

First of all, if the model is to be able to analyze a variety of airports, a method of specifying the positions of airport facilities is needed. Present methods accomplish this, but at significant cost. A new method is proposed that can accomplish the same task at much lower cost. Aircraft to be modeled can flow over a large number of paths between runways and gates. Modeling this process is desired

to assess taxiway constraints for delay analysis or for taxiway design problems. Present manual enumeration techniques prohibit path switching to avoid congestion and are difficult to use. The difficulties associated with automatic generation of paths are discussed and a compromise procedure is proposed (but not implemented) that limits the effort required to model taxiing and provides an approximate means to analyze taxiway system problems. These topics are discussed in section II.B.

The characteristics of the various aircraft using the airport greatly influence the capacity and the delay experienced during operation. Any estimate of present conditions and especially any forecast of future demands for service involve uncertainties as to number, type and timing of the aircraft that will use the airport. Section II.C discusses how present methods of translating these estimates into a schedule of operations in the simulation systematically underestimate the uncertainty, with the consequence of underestimating the expected delay in the given situation. A new demand generation technique is developed to avoid this problem.

It is recognized that the operation of runways is the major determinant of the capacity of the airport. Sections II.D and II.E address this topic. Modeling of runways involves the analysis of both the factors influencing the performance of aircraft and the rules governing the separation of aircraft movements.

Significant study has been done to identify the factors that determine exit selection and runway occupancy time for landing aircraft. This study has been partially motivated by a desire to facilitate early exit and to reduce delay. However, thus far, models intended for direct estimation of delay have not modeled these factors. Instead, the predicted performance was an input to these models. A method of modeling landing performance is developed which allows a direct assessment of the impact on delay of changes in a number of factors such as the location and type of exits.

The rate at which aircraft can use the runway is also limited by the separations imposed between aircraft. Modeling separation rules is complex, due to the variety of ways in which aircraft can interact and because the rules vary in response to weather. Significant work has been done on analysis of single runway separations (both aircraft on the same runway). Analysis of dependent runway separations (aircraft on different runways) has not progressed as far as the single runway case. Current methods of specifying these separations allow only approximate solution to several interesting cases. Further, it will be seen that current methods rely on the analyst to assess any interaction between separations and aircraft performance. This makes certain cases difficult to analyze. We develop separation logic to model this interaction directly.

Airports operate in a dynamic environment. Section II.F will discuss the interactions between the airport and its environment. Weather and congestion can produce a complex set of changes in operating policies, demand for service, and aircraft performance. Because airports seldom operate with the same runway configuration and air traffic control rules for an entire day, it is important to model the impact of such changes. Existing models, however, have almost no ability to do this. We will develop an extensive capability for FLAPS to model changes in operating policies.

Many statistical issues underlie a simulation model. Unfortunately, they are not often sufficiently appreciated by applications oriented analysts. The proper construction of random number generators and the analysis of simulation outputs will be discussed. We will discuss the selection and design of the various statistics computed by FLAPS.

B AIRFIELD GEOMETRY AND GROUND OPERATIONS

1 Introduction

This section deals with the closely related issues of how the geometry of the airport and the routes aircraft take in moving over the airport are represented in the model. We need to model the location and attributes of airport facilities to provide a framework for modeling aircraft movements between runways and aprons.

In this section we examine and evaluate the methods currently used to model both airport geometry and aircraft taxiing operations. We will see that these methods are unsatisfactory both in terms of the restrictions they place on analysis and in the time and cost requirements they place on the user.

We then review a method of modeling airfield geometry proposed in an earlier work by this author (NOR 78). This is a satisfactory means of representing airfield geometry and has been implemented in FLAPS. It reduces the difficulty of specifying taxiing routes but, for reasons described below, does not represent a complete solution to the problem.

Finally, we propose a modification of this method that provides greater flexibility. Due to constraints on both time and cost the method was not implemented in FLAPS.

We first discuss geometry representation, then describe taxiway modeling.

2 Description of Airfield Geometry

The ASM model represents the airfield as a set of modules where a module consists of any one aircraft capacity element. Typical elements are runways, intersections, taxiway segments and runway crossing links.* Each taxiway segment has a length and a taxiing speed associated with it. Figure II-1 shows a map of La Guardia and figure II-2 presents the airport geometry as described by ASM.

This representation of airfield geometry is accurate but demanding. Taxiway links are so short that any airport large enough to be of interest in this study is very difficult to describe and prepare as an input to the model. For the model of O'Hare (BAL 76 p. 94-103), perhaps the most complex site, 545 links were defined, each associated with a length and taxiing speed. The model of La Guardia (PMM 77 p. 90-3) required 230 taxiway links.

Describing the airport geometry through this methodology also leads to very high costs in running the model. The model maintains an events list of up to 200 aircraft which must be scanned after every event to find the next event. Links range from 200 to 400 feet in length with an average of 300 feet and taxiing speeds range from 10

*taxiway segments that cross a runway.




FIGURE II-1. New York La Guardia: Airport Map.




FIGURE II-2. New York La Guardia: ASM Modeling Geometry.

to 30 mph with 20 being the average. Short links, not surprisingly, tend to be associated with slower taxiing speeds so that travel time across a link runs around 10 seconds. A taxiing aircraft thus requires updating and an events list scan of approximately 6 times per simulated minute.

A procedure making considerably less demands on resources was developed in NOR 73. The procedure was based on the realization that the airfield is a network with nodes and links. The links are the airport facilities and the nodes are the places where facilities join or intersect. A minimum set of modules to represent the links (runway, apron, intersection and taxiway) and a set of keypoints (with X and Y coordinates) to represent the nodes was defined. The keypoints provide the information necessary to locate the modules and each module provides the attributes of the facility to which it corresponds.

This procedure is implemented in FLAPS through the preparation of two files: 1) The keypoint file and 2) the module file.

The keypoint file consists of a list of keypoint numbers and the X and Y coordinates of each keypoint relative to a user defined origin. A keypoint is located at every point on the airfield where two or more streams of aircraft cross, for example: runway endpoints and exits, taxiway intersections and runway intersections. See the

examples for New York La Guardia (figure III-1) and Boston Logan (figure IV-1 and IV-2) for maps of airport and model geometry. Table IV-1 shows the keypoint file for Boston Logan.

The module file consists of a list of runways to be used in the module. Associated with each runway is a set of parameters specifying what keypoints define its limits (runway end points), what keypoints are located at exits and parameters defining exit velocities and the length of the final approach for the runway (see section II.D for descriptions of these parameters).

The model processes the airport geometry information to derive additional information about each type of module. For example, the compass orientation of the runway and the distance of each exit and intersection from the ends of the runway are derived from the runway information rather than being separate inputs.

3 Ground Operations

3.a Existing approaches

Aircraft ground movements between runways and aprons are provided for in the ASM model by a set of user-specified aircraft paths. One path is required between every combination of runway exits and each gate and between every combination of runway endpoints (where takeoffs begin) and

each gate. A path consists of an ordered list of the links, intersections, and runway crossing elements that an aircraft will traverse while taxiing. Only one path between each origin-destination pair is permitted and can not be modified during a run. This procedure was evaluated in an earlier work (NOR 78) which identified two major weaknesses.

First the method, while permitting modeling of many situations, imposes a significant burden in time and cost on the user. Again using the O'Hare example, 1293 aircraft paths with an average length of 20 links (BAL 76 p. 154) are required to permit landings on 6 exits of 2 runways and permit takeoffs at one end of a third runway. A full set of paths allowing landings at 37 exits and takeoffs from both ends of 7 runways would consist of about 5500 paths. La Guardia Airport requires 372 paths with an average length of 14 links (PMM 77 p. 125). Once an airport has been set up for simulation it can, of course, be used as often as one wishes with no additional preparation. However alterations in airport geometry could pose significant problems. Even adding a single taxiway would alter several hundred paths. Testing alternative taxiing routes would be a very time consuming task.

The second and more fundamental criticism of the ASM model is that it does not permit path switching to avoid congestion. A controller, when routing aircraft through the airport's taxiway system during a period of heavy

congestion, would be likely to divert aircraft to new routes that are less congested. Confining aircraft to a single pre-specified path will likely concentrate aircraft movements and exaggerate estimates of taxiway congestion. Having made this criticism, it must be pointed out that devising a flexible, simple alternative turns out to be difficult.

3.b Discussion of possible alternatives

In this section we discuss the possibility of devising automated methods of aircraft path generation with the objective of reducing the input requirements and enhancing flexibility by permitting multiple routes. We conclude that no perfectly satisfactory (both economical and efficient) solution can be achieved and that a range of options and levels of modeling detail might be the best choice in this case.

If the ASM model, with one user-defined, unalterable aircraft path between runway exit and gate, represents one extreme, then perhaps the other extreme would be to have the simulation select the best path for each aircraft to follow every time an aircraft reaches an intersection. To use this method a criterion for deciding which prospective path is best and a function or procedure for calculating the value of this criterion for each link would need to be established. The usual situation is to have minimum travel

time be the criterion and to define a supply function giving link travel time as a function of link volume and link characteristics. The well-known shortest path algorithm of Dijkstra (DIJ 59) could be used to determine which path between the aircraft's current location and its destination traverses links with the lowest sum of values for the criterion chosen. If our criterion, for example, were travel time, the Dijkstra algorithm would find the minimum travel time path.

There are, unfortunately, a number of serious shortcomings to such a scheme. The airport network differs from conventional traffic networks in that most taxiway delay to aircraft is likely to occur in trying to cross busy intersections or in waiting for an aircraft moving in the opposite direction to clear a link rather than from interference with other aircraft on the same link. To deal with this problem, the expected delay offered by crossing an intersection will have to be explicitly modeled, perhaps as a function of the utilization rate of the intersection.

An additional problem is that taxiways are one-way-at-a-time links rather than either two-way routes or one-way routes in any single direction. This means that, if aircraft are to be assigned to conflicting routes, some method must be developed to prevent or resolve such conflicts. This problem was not adequately dealt with in NOR 78. The suggested method did prevent certain types of

conflicts, for example, side-on collisions at taxiway intersections. However, the method was not able to resolve or prevent other conflicts, such as those occurring when taxiing aircraft have paths proceeding in opposite directions for several links.

A second problem is that the data needed to determine the best path may include more than the present state of the network. In highway traffic, hundreds of vehicles are on each link of a network at any one time and the number stays relatively constant over the length of time one vehicle takes to cross the network, then link characteristics and the choice of best path will remain stable. Neither condition applies to an airfield network. Flow rates are very low. The number of aircraft on a path being considered for additional traffic may change significantly during the time one aircraft will require to move over the network. This means that using a recent average of travel time as a best path criterion may be misleading.

The conclusions drawn from the difficulties described above is that any automated dynamic routing procedure would be very difficult to implement. If ground delay were believed to be the major component of delay in most situations, this would provide a motivation to seek such an accurate model of ground operations. It is generally agreed, however, that ground delay is usually not an important problem at all. Ground delay may sometimes be

significant in particular locations of an airport but only very rarely is ground delay a major problem over an entire taxiing system. The operation of airports is so idiosyncratic and future aircraft flows and routings so easily shifted that it is not clear what significance an "exact" estimate of ground delay would have. The analyst is much more likely to ask questions such as "Can the taxiway handle a demand increase of 30%?" and look for answers such as "Yes - with no problems", "Yes - with some trouble spots that may require redesigning" or "No- it causes airport-wide problems." So while ground delay is not often a significant part of total delay, it is very important that the analyst be able to prove this for any given airport configuration. Therefore, we cannot simply abandon the modeling of ground operations. We do, however, need a more flexible modeling technique. The following technique is proposed as a solution. We will explain the method by showing how it would be used in the analysis of an airport.

3.c Proposed solution

The following discussion assumes the geometry input described in part 3a. The method is basically iterative. In the first pass through the process the analyst prepares a set of aircraft paths, one for each origin-destination pair over which aircraft will travel. Each path consists of the list of keypoints, in order, on the route taken by the

aircraft. When the model is run, the length of each path is calculated by summing the distance between each pair of keypoints on the path. As each aircraft moves onto a path a taxiing speed is determined (from user-supplied inputs) and the taxiing time over the path calculated. Each aircraft is then made to move without interference from any other aircraft, even aircraft moving in opposite directions. This procedure will result in no ground delay being recorded by the simulation.

An additional provision in the model provides an indication of how much taxiway delay would have been recorded with a complete modeling of intersections. Each time an aircraft crosses a keypoint the time and direction of the crossing is noted. At the end of the run the average number of times that two aircraft cross a keypoint within a specified interval of time is displayed for each keypoint. A separate (presumably larger) interval is used for aircraft moving in opposite directions across the keypoint to pick up head-on encounters. This output enables the analyst to assess the magnitude and location of the congestion problem. For example, if the results indicated that two aircraft came within 20 seconds of each other at keypoint X an average of three times during the course of the day, then it is clear that even if the most sophisticated intersection module were to be used to model this intersection there would be no discernible impact on overall airport performance. If this

is the situation at all keypoints, then the analyst will be satisfied that ground operations do not encounter serious delays and turn to other aspects of the analysis without proceeding any further. Conversely, if the average number of close encounters is very large at a significant number of keypoints, then the analyst knows that serious problems exist and that some action will be necessary in order to expedite the aircraft flow on the ground.

Intermediate cases will require a revised approach. If a number of intersections are showing severe congestion, then the analyst may either try to redesign the aircraft paths to reduce the problem at these locations or may elect to place intersection modules at the most congested intersections and run the model again. An intersection module would be designed to permit passage of only one aircraft at a time across an intersection. An intersection module at a keypoint would hold aircraft if necessary to prevent two aircraft from "colliding" at the intersection. If the outputs of this case reveal no remaining intersections with serious congestion problems then the results may be taken as accurate. The same can be said if some congestion exists but overall delay figures are largely unchanged from the first iteration. If, however, the original case had many intersections with severe delay or if the addition of a few intersection modules has caused the taxiway capacity to be reduced so that the congestion has

propagated to other intersections, then the analyst has several options. A new set of paths may be designed or additional intersection modules may be installed to see if significant additional delay will result.

As more and more intersection modules are added, two problems can occur. First the simple "low level" control of conflicting movements via individual intersection modules will not prevent or resolve head-on conflicts between aircraft. Secondly the danger arises that ground delay may be increased over the real world situation because the model has only one path between each origin and destination. If this limits the analysis, then the user may wish to take advantage of the model's controller module to design and implement a path switching mechanism. This would involve developing a "high level" control strategy that could oversee the entire state of the airport's ground operations. Any path switching mechanism could be as simple or complex as desired and may be tailored to the specific airport or be a general procedure. No such mechanism is proposed here. The point is that, while it is very difficult to anticipate all of the options that may be needed to model operations at any airport, a procedure can be provided for allowing airport specific changes within the overall model.

An additional benefit of the time of crossing record is that it enables checking of the very important runway logic. The intersection of two runways is a keypoint. Thus the

average number of near-collisions can be recorded for each runway intersection as well. This would be valuable both as a model development tool and as proof in production runs that the runway and separation logic is performing as desired.

C DEMAND FOR SERVICE

1 Introduction

The previous section of this chapter described the characteristics of the airport, the supply side of the model. This section considers the aspects of the model that are related to the generation of demand for service. This involves modeling how many aircraft use the airport at what times of the day, i.e., the level of demand. It also involves the specification of the attributes of the aircraft. By attributes we mean any aspect of an aircraft (such as class) that could affect the manner in which it uses the airport. This section of the model produces a schedule of the operations that will take place at the airport during a given replication of the model. Contained in the schedule are those attributes of aircraft that are exogenous to the airport.

Three modeling issues are significant in the construction of aircraft schedules. 1) What aircraft characteristics or attributes are in fact external to the airport, and what attributes are determined dynamically by conditions at the airport? 2) Once the correct set of attributes has been chosen for inclusion in the schedule, what relationships (correlations or logical dependencies) exist among the attributes? 3) What method should be used

to generate the number and times of entry of aircraft that are to be modeled?

We will first describe and briefly evaluate the schedule algorithm used in the ASM model. This procedure reads a schedule prepared externally to the model and shuffles the aircraft in the schedule for each replication. We term this a "base schedule" procedure. Evaluation of the assumptions of this procedure (in section 5 below) shows that while it permits analysis of a number of situations, significant improvements are needed with respect to all three of the modeling issues in order to correctly simulate all the situations of interest.

We will then discuss in detail each of the three modeling issues raised above and propose a new procedure for generating aircraft schedules. This new procedure is termed a "probabilistic schedule." In this method, interarrival and interdeparture times are drawn from a probability distribution, usually an exponential distribution with a time varying rate. Other aircraft attributes are drawn from other appropriate probability distributions. Each replication thus uses a schedule independent of other replications. This method expands the capability of the present schedule generation technique to permit the use of FLAPS for forecasting future airport performance. FLAPS provides for the generation of schedules using either a base schedule procedure or a probabilistic schedule procedure.

2 Present Schedule Algorithms

2.a Description

The schedule generation method used in the ASM model consists simply of preparing a schedule as an input to the main simulation. The schedule contains one record for each aircraft to be modeled, including

- 1) Airline name (UA, EA, AA etc.)
- 2) Flight number
- 3) Gate
- 4) Flight type code (originating, terminating, through, turnaround, touch & go*)
- 5) Aircraft class (category of aircraft, i.e., heavy, medium, light)
- 6) Arrival time at threshold of landing runway
- 7) Departure time from the gate
- 8) Approach fix
- 9) Landing runway
- 10) Takeoff runway
- 11) Departure fix

In each replication of the model the schedule is used to control the insertion of aircraft into the simulation. In each replication the arrival and departure times are perturbed by the addition of a random "lateness distribution" (which may include negative values) to

* Intended for modeling GA training maneuvers. It is not intended to model missed approaches induced by separation violations

simulate the day-to-day variations in time of the arrival and departure of the aircraft. The number of aircraft and the set of attributes for each aircraft do not change from replication to replication. The attributes are used as they are specified in the schedule because there is no internal mechanism in the model logic for altering attributes during the simulation. For example, runway assignments cannot be varied to minimize delay or to avoid weather problems. Because the original input schedule is used as a base for future schedules, this method will be referred to as a "base schedule" method.

2.b Evaluation

In this section we will briefly present some observations about the ASM schedule procedure. In sections 3 and 4 below we will examine in more detail the modeling implications of this method.

The most obvious point to make about the ASM base schedule procedure is the amount of information that the analyst must know in order to prepare a schedule. For example, not only the distribution of the classes of aircraft must be known (e.g. 5% of class 1, 20% of class 2, etc.) but the exact sequence is also required: class 1 followed by class 3 followed by class 2. The relationship among various attributes is needed (the class 2 aircraft at 3:15 used runway 3). Collation of OAG entries will provide

the analyst with much of this information but only for existing airports and only for past or present situations. For analysis of new airports or forecasting studies of existing airports the analyst will have to create this information outside the ASM model. In any case, arrival fix, departure fix, and runway assignment will have to be supplied by the analyst.

The second point that should be made is that regardless of the source of schedule information, the base schedule technique imposes certain restrictions on the situations that can be correctly modeled. The use of the same set of aircraft in each replication eliminates the possibility of accurate analysis of situations where only an approximate estimate of some schedule parameter is available. Thus, this approach implies, in effect, that information about aircraft is deterministic, not stochastic. Using the same aircraft in each replication in analyzing operations at some future date implies more exact knowledge about the future than it seems reasonable to assume. Once we have developed the probabilistic schedule method this point will be discussed in depth in section 5 below.

A third difficulty with the ASM model's schedule generation procedure is in the choice of aircraft attributes that are included in the schedule. As mentioned, for example, runway assignments are among the attributes included in the schedule and cannot be altered. If we wish

to model the dynamics of airport operation, this capability will have to be created.

In summary, the basic problem with the ASM method is that the analysis of how schedules affect performance is, in effect, forced outside the main model. One cannot easily study the effect of uncertainty in the prediction of aircraft flows, nor the effect of varying the amount of uncertainty in the schedule. Since it is reasonable to expect that this kind of simulation will be used primarily for forecasting congestion problems, this schedule procedure is a serious deficiency of the ASM model. In the following two sections we develop an alternative schedule generation methodology that is more responsive to the types of questions an analyst will be likely to ask. The new method is explicitly oriented toward modeling the uncertainty inherent in the forecasting process.

3 Aircraft Attributes

In this section we take up the first two modeling issues posed in the introduction.

3.a Implications for analysis

Briefly stated the schedule should contain that information about the aircraft which is known before its arrival at the airport. It should not include information

which is in fact determined during the simulation, such as runway assignment or specific path over the airport surface. This distinction is an important one. If a given aircraft attribute, which in the real airport is determined during the course of operations and can be altered in reaction to developments (i.e. the landing runway used by an aircraft), were to be specified in the model as an input, fixed before the simulation run, then the potential arises for misleading results. It may very well be that for a base or test case, data about the given attribute can be obtained and used as input to the model. This information will enable the model to produce accurate results for the base case since the input data were collected for that particular case. However, when the model is used to analyze a different situation, a case for which no data exist on the given attribute, one of two situations may occur. The user may continue to use the base case data on the new case. This will create misleading results as the real airport would not exhibit the same pattern of the given attribute in the new situation as it did in the base case. Alternatively the user may estimate the attribute from the user's expectation of what it should be for the new case. Such estimated data may turn out to be very accurate or may be quite different from what the airport would actually be like.

The point is that in the example just described, the attribute is not really an independent variable but is

dependent on what is occurring on the airport, and thus should be modeled as such. In fact, if one expands the scope of airport analysis sufficiently, only the distribution of aircraft classes will remain fixed - under severe conditions even the apron area used, direction of approach to the runways, path taken after takeoff and the number of movements can change. The course taken in this work is dictated by the primary purpose of the model developed here, which is to analyze questions of capacity and delay on the airport and how the number of movements accommodated can vary by changes in the operating rules of the airport. Thus the items in the schedule for FLAPS should reflect this orientation in that they include those attributes of aircraft that are not altered except in extreme situations.

3.b Selection of attributes

A wide variety of attributes could potentially be in the schedule, including:

1) Entry attributes

arrival time, departure time

type of operation - arrival, departure, through

2) Routing attributes

arrival fix, departure fix,

arrival runway, takeoff runway

apron area or gate

3) Performance attributes

Class, % loaded or landing weight, noise, fuel consumption rate

4) Identifying attributes

airline, flight number, index number (1, 2, 3, etc., i. e. a set of numbers identifying aircraft which would be without significance outside the model.)

We consider each of these attribute categories in turn.

Entry attributes

Entry attributes describe when, and by implication, where aircraft enter the simulation. There are basically three ways an aircraft may use an airport: 1) as an arrival - that is the aircraft enters the terminal area, lands, taxies to a gate and then either stays there for the remainder of the day or some time later is removed to a hangar. 2) As a departure - an aircraft that is on the airfield at the beginning of the day and at some point leaves a gate, taxies to the end of a runway, takes off and leaves the airport. 3) As a through aircraft, one that both arrives, lands, taxies to a gate, and later taxies out and takes off. Touch-and-go operations, simulated by the ASM simulation are not modeled in FLAPS as they constitute a negligible fraction of operations at major air-carrier airports.

The point of entry is different for arrivals and through aircraft on one hand and departures on the other. It would reduce the possibilities of confusion if these point of entry times could be defined to correspond both to a field-observable event and be measured at some convenient point.

The choice of how to define the point of entry for arrivals and for through aircraft turns out to be very difficult. The ASM simulation uses time over the landing runway threshold for arrivals. A problem arises because this is not a field-observable event. The observed threshold arrival time, obtained from watching aircraft land, already includes any delay the aircraft has experienced due to runway and airspace congestion. This delay is what the airport model should be trying to predict. If observed threshold times were used for point of entry, the model might give excellent results when simulating the existing conditions that produced the observed threshold times but would be very misleading when run for different conditions. The actual runway delay would be different for the new case but the input data would implicitly assume the old delays. This would cause the misspecification discussed in section 3.a above.

If the base schedule is obtained from the OAG listings an additional error may exist. In this case, a natural course of action would be to simulate scheduled threshold

time by subtracting some constant time from each of the scheduled gate arrival times in the OAG. The time subtracted would be intended to represent the time an aircraft required to land and taxi to the gate. This procedure, however, would also be incorrect since it is widely known that airlines add expected delay to trip times when computing scheduled arrival times. Thus, uncritical use of the OAG information will, to some extent, smooth out the arrival pattern.

However, there is unlikely to be any field-observable time that corresponds to an undelayed entry. This is because under extreme delays aircraft will be delayed at great distances from the destination airport. In fact, aircraft may be held at the departure gate of their originating airport rather than circle above the destination airport. Therefore the point of measurement of arrival time should be chosen at the point where the model actually begins to follow the movement of aircraft. The use of threshold time in the ASM model is in slight violation of this idea as the ASM model does follow the final approach phase of operations to the extent of imposing minimum arrival/arrival separations on the landing aircraft. The arrival time for FLAPS is chosen to be at the top of the common approach path since FLAPS does completely model the final approach phase of landing (see section D.2).

For departing aircraft there is an obvious event to use: scheduled time to leave the gate. This value is both observable and the point where we start modeling departures.

Through aircraft use the same entry time as arrivals. Departure time for these aircraft cannot be generated directly as the time an aircraft leaves the gate is a function of the time the aircraft enters the gate and a certain minimum on-gate time. An aircraft has a scheduled departure time and will not leave before this time. There is also some minimum time that the aircraft must be at a gate in order to complete the boarding of passengers, refueling, etc. If an aircraft arrives at its gate late, it may still leave at the scheduled departure time but it must stay at the gate for the minimum length of time. Therefore two sets of mean values and standard deviations are required per aircraft class. The first is used to set the scheduled stay at the gate and the second to draw the minimum time the aircraft must stay at the gate. Once the arrival time is determined, the sum of the arrival time and a time drawn from the first distribution is used as scheduled departure time. Note that the first distribution should include an estimate of time the aircraft requires for landing and taxiing to its gate. When the aircraft actually arrives at the gate the maximum of 1) the scheduled departure time and 2) the sum of the arrival time at the gate and the minimum on-gate time (drawn from the second distribution) is the

actual departure time. The ASM model uses a similar procedure save that the scheduled departure time is supplied to the model by the external schedule.

Given the inclusion of an arrival time and a departure time it is redundant to input a separate code for the type of operation (arrival, departure or through aircraft). This information may be easily coded in the arrival and departure time. A negative arrival time may signify a departing aircraft. A departure time of infinity may signify an arrival. In summary, entry attributes included in the FLAPS schedule are 1) an arrival time at the head of the common approach path and 2) a departure time from the gate.

Routing attributes

Some information must be provided as to the route aircraft take over the airport. As the ASM model provides no mechanisms for dynamically assigning aircraft to routes a full set of information must be provided to that model. As will be discussed in detail in section II.F below, it is important to be able to simulate the dynamic assignment of aircraft to runways and routes. Therefore, we must try to pre-specify in the FLAPS schedule only those routing attributes that do not change as conditions change on the airport. Of the five attributes used in the ASM simulation (numbers 3, 8, 9, 10 and 11 in the table in section 2.a above) the two that define runway use should be removed. As has been discussed above, runway assignment is not

predetermined. The runways in use change over the course of the day and runway assignments to aircraft are made dynamically based on the class of the aircraft, runways in use, aircraft currently waiting to takeoff at each runway and, perhaps, destination of the aircraft.

The choice of apron area should, however, remain a decision independent of airport conditions. The usual case is for an airline to occupy a fixed area and have all of its flights use the same set of gates. If this inflexibility turned out to be a problem in a particular situation, a special controller module could be designed to alter apron assignment.

The role of arrival and departure fix information in the ASM model is rather obscure. Documentation implies that aircraft actually arrive at this point and are "vectored or put in holding patterns" (PMM 77 p.4). However, examination of the actual program code does not show any such use of the fix information save as a way of partitioning total arrival delay among several categories of aircraft.

Fix information has a much more important potential use. Aircraft that arrive from different directions may be put on different runways to segregate aircraft streams. For departures the information is more crucial in that aircraft that depart on the same route must be separated by larger distances than successive departures that follow diverging courses (FAA 78 parag. 340). Departure fix can

provide information to decide what separation is required. Departure fixes may also be used as a proxy for destination if it is desired to assign aircraft to runways on this basis.

Performance attributes

There is wide variation in the performance of aircraft in a number of respects. These include: speed on landing, exit selection, runway occupancy time, separation rules used, noise levels, number of passengers or quantity of cargo, wake vortex generation, taxiing speeds, fuel consumption rates, time on gate. If one were to use all of these it would result in a very complex and unwieldy model. Fortunately, the physical size of the aircraft is well correlated with most other performance characteristics. Thus we include only aircraft class as an indication of the performance of each aircraft generated. The class of an aircraft is then used internally in the model as an input parameter to various functions to specify the necessary additional performance attributes.

Usually five classes are distinguished:

- 1) "Heavy" jets (B747, DC10, L1011)
- 2) "Large" jets (B707, DC8)
- 3) "Medium" jets (B727, B737, DC9, BAC111)
- 4) Large propeller aircraft (DC6, Convair 580)
- 5) General aviation

Class 5 aircraft may be divided into a class 5 and class 6 for high performance general aviation and small general aviation aircraft. It should be emphasized that the association of class categories to size of aircraft, while an obvious and useful procedure, is not required. If desired, the class parameter may be used to analyze any problem where it is useful to segment the input stream according to some criterion. For example, each class might represent aircraft from a different airline if this was felt to be the major performance-related difference among aircraft.

Identifying attributes

The ASM simulation model requires that each aircraft be identified by airline name and flight number. There are two difficulties with using this procedure for the probabilistic schedule. First, when the number of aircraft varies across replications then there is no unambiguous way to identify the same aircraft from replication to replication. Secondly, this level of identification is probably meaningless for any use of the model in medium or long term forecasting given the rapid rate at which flights are revised by airlines. Note that the airline to which the aircraft belongs can be taken into account through the way apron areas are assigned. Airlines usually occupy one (or sometimes two) apron areas. The different aircraft fleets of a large airline flying mostly medium and large jets and a

small airline flying mostly propeller aircraft with a few medium jets can be modeled by assigning different percentages of the various classes to each apron.

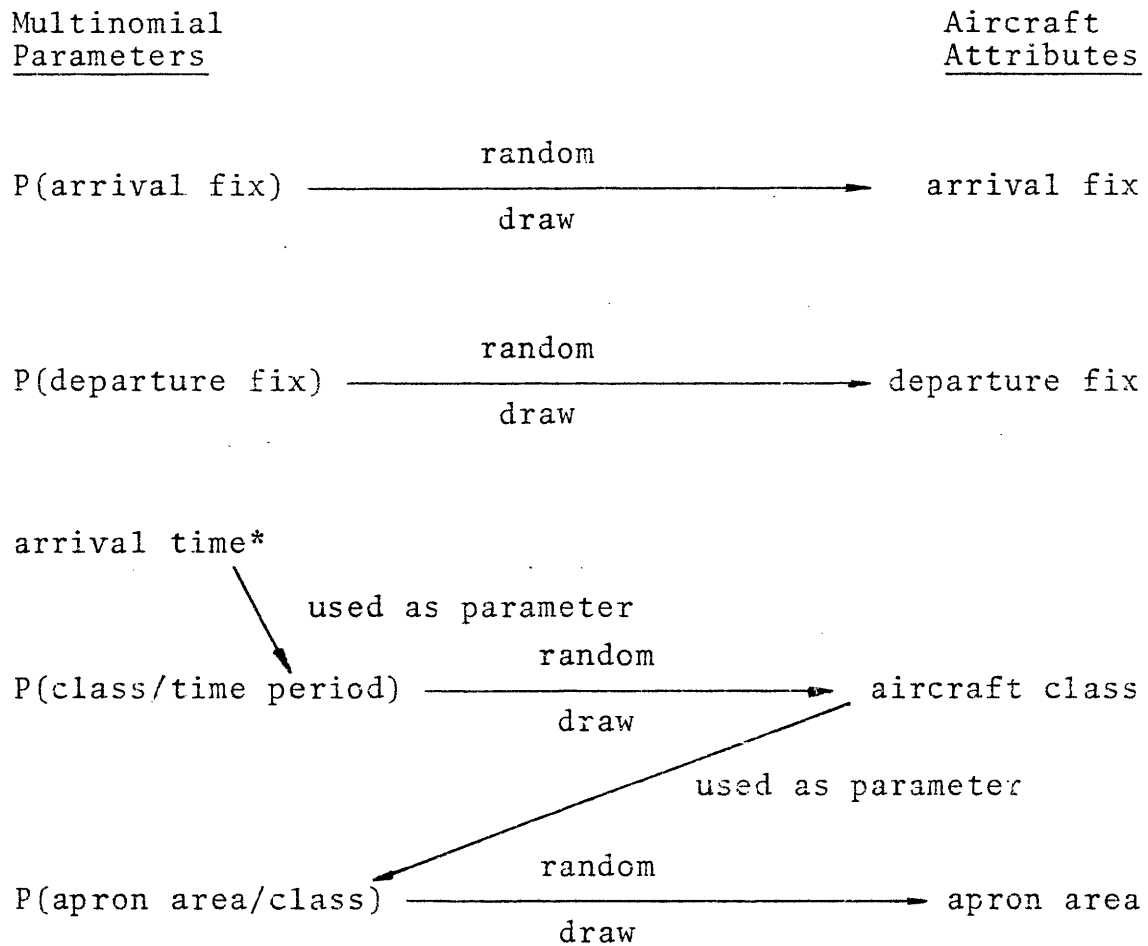
There is a second dimension to identifying aircraft and that is whether the identity of each aircraft should be retained during one replication. In other words, when an aircraft completes a takeoff, is it important to know the history of this aircraft, i. e. what apron it came from or how much delay it experienced on landing? This does turn out to be necessary for FLAPS as we want to prepare aggregate statistics about the aircraft. As this is the only function of aircraft identification in FLAPS, this number is assigned internally by the model and need not be a concern of the analyst.

3.c Relationship among attributes

The question of the relationship among aircraft attributes does not arise when a base schedule method is used. This is because any such relationship is determined by whatever procedure is used to create the external schedule. In the schedule procedure being developed here the schedule is created internally by the model. Thus it is necessary to consider the relationships among the aircraft attributes.

Four aircraft attributes must be set: class, arrival fix, apron area and departure fix. The obvious procedure

would be to define four multinomial distributions and select each attribute independently. This procedure, however, would restrict the range of situations that could be modeled. It may be the case that we would want to model situations where certain types of aircraft (e. g. general aviation) were restricted from operating at certain periods or situations where certain apron areas cannot accommodate certain aircraft types (e. g. all jumbo jets must use one of the aprons or there is a GA hangar that only GA aircraft use). In order to handle these and similar situations two first order interactions are modeled among the parameters: 1) aircraft class as a function of time period; and 2) apron area as a function of aircraft class. These are set by two matrices of parameters. The first is a set of multinomial parameters, one set for each time period, which give $P(\text{aircraft class} | \text{time period})$ - probability of a particular aircraft class given a time period. The second is a set of multinomial parameters, one set for each aircraft class, which give $P(\text{apron area} | \text{aircraft class})$ - probability of a given apron area given aircraft class. Used sequentially the two generate aircraft class and apron area. Arrival and departure fixes are selected from simple multinomial distributions. See figure II-3 for a schematic of how aircraft attributes are selected.



* From arrival rate function

FIGURE II-3. Selection of Aircraft Attributes.

4 Modeling the Level of Demand

This section discusses how the number and entry times of each aircraft are determined, the third modeling issue raised in the introduction.

4.a Probabilistic schedule parameters

A number of parameters need to be specified to completely define the probabilistic type of schedule. The number of aircraft to be simulated is controlled by specifying two time-varying rate functions, one for aircraft that arrive and for through aircraft, and one for aircraft that only depart. The rate functions could be specified in one of a number of increasingly complicated ways: stepwise constant, piecewise linear, or second or higher order curve. The data that is typically available on airport demand is hour by hour flow rates. This would seem to argue for using a stepwise constant function. But flow rates are very unlikely to jump suddenly from one value to the next and then remain constant for an hour. Using a piecewise linear function will ensure that flow rates in the simulation build up and decline gradually. We thus assume that a piecewise linear function provides an adequate approximation to any given demand profile. (See figure II-4 for details.) As

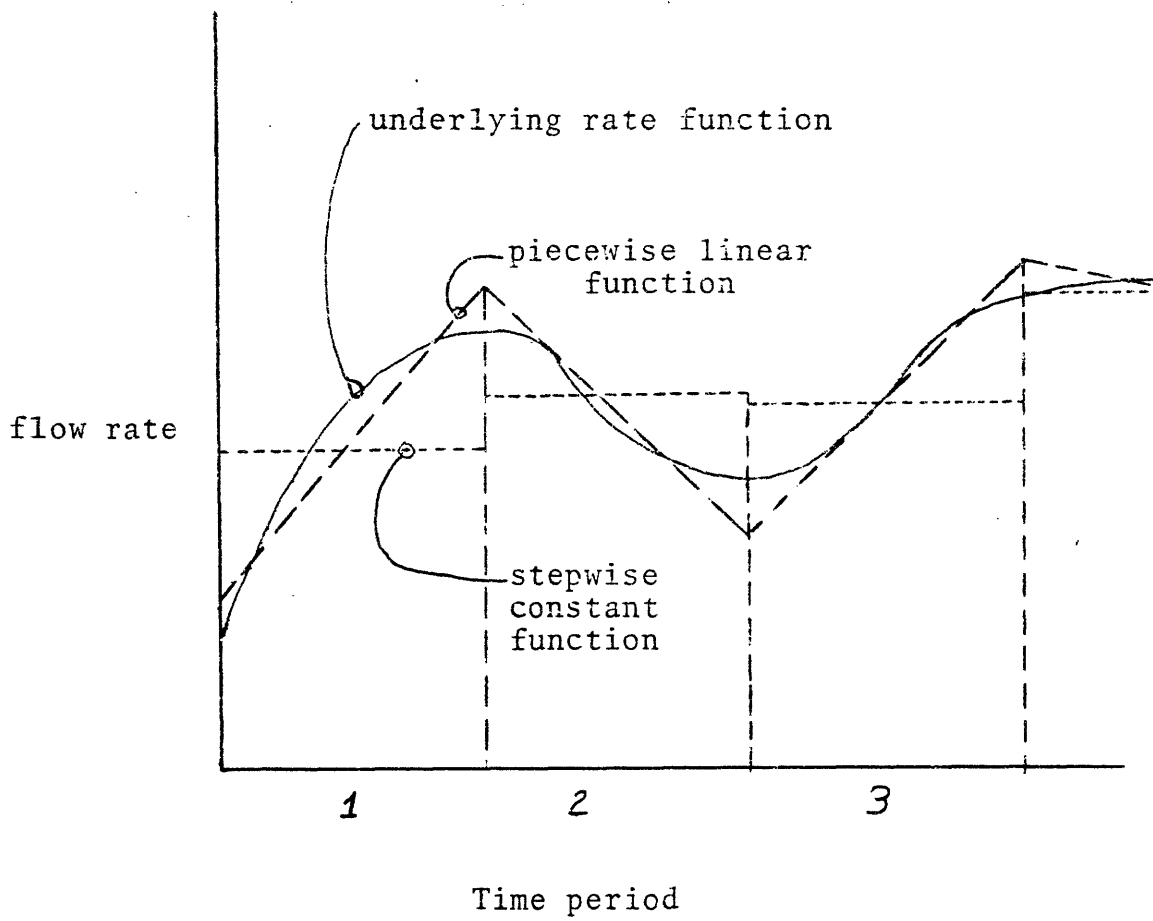


FIGURE II-4. Rate function.

will be explained below, use of the piecewise linear function results in a very simple procedure for probabilistic schedule generation.

With each of the rate functions a type of probability density function must be specified that will control the distribution of inter-aircraft times. An obvious assumption would be to use the exponential function. Data from Logan Airport tends to support the validity of this assumption.

One additional binomial parameter is needed to set the percentage of arrival aircraft that are through aircraft. Whether each particular arrival aircraft departs is thus determined randomly using this parameter to set the odds.

4.b Method for generating a time-varying stochastic process

In this section* we discuss how inter-aircraft times are generated. To simplify matters we will confine our discussion to generating interarrival times for aircraft that arrive and for through aircraft. The same procedure is used independently to generate departure times for aircraft that only depart.

The basic procedure used is to start at the beginning of the time to be simulated, draw an interarrival time $t(1)$ from a distribution, place the first arrival at time $t(1)$,

* The author wishes to acknowledge the assistance of Dr. William Dunsmuir, Dept. of Mathematics, MIT with this section.

draw a second interarrival time $t(2)$, place the second arrival at time $t(1) + t(2)$, and proceed in like manner until the sum of $t(i)$ is greater than the end time to be modeled. A problem arises, however, because we wish the rate of arrivals to vary over time. We will first consider the common case where interarrival times are assumed to be Poisson distributed and then consider extensions to the non-Poisson case.

Two basic methods exist for generating samples from nonhomogeneous Poisson processes (LEW 78): The rejection method and the time-scale transformation method. The rejection method involves generating samples from a homogeneous process with rate equal to the maximum rate of the time-varying process and then accepting or rejecting each point of the homogeneous process based on a random draw with the probability of acceptance equal to the ratio of the rate of the time-varying process at that point over the rate of the uniform process. For many of the cases of interest to airport analysts, with significant variation between the occasional maximum rate of 30-35 aircraft per hour and the more common off-peak rate of 10 to 20 aircraft per hour, this method would entail generation of a large number of points which would be rejected. This method is certainly feasible, however, we prefer the time-scale transformation method for this application as it enables the straightforward generation of the sample.

The time-scale transformation method is "a direct analogue of the inverse probability integral transformation method for generating (continuous) nonuniform random numbers." (LEW 73 p.2) The method is thus analogous to the way uniform random numbers in the range of 0 to 1 are mapped into negative exponentially distributed random numbers. The method is implemented in the following way (figure II-5 should be consulted in parallel with the following discussion). Samples from a negative exponentially distributed random variable of mean 1 are generated starting at zero and continuing until the sum of the sample values that are generated has a value greater than the total number of expected arrivals in the course of the run. This process is represented in the figure as $x(1)$ to $x(I)$. Note that the number of points, I , in the process $x(i)$ will vary from replication to replication and will be distributed according to the Poisson probability distribution with a mean value equal to the total number of expected arrivals for the case being considered. The lower graph in Figure II-5 indicates the arrival rate function in terms of aircraft per time period. The upper graph is the integrated rate function which is merely the integral of the lower function. The value of the integrated rate function at any time is the number of aircraft expected to have arrived by that time in the simulation. This function is used to map the process $x(i)$ onto the process $t(i)$. The process $x(i)$ which extends


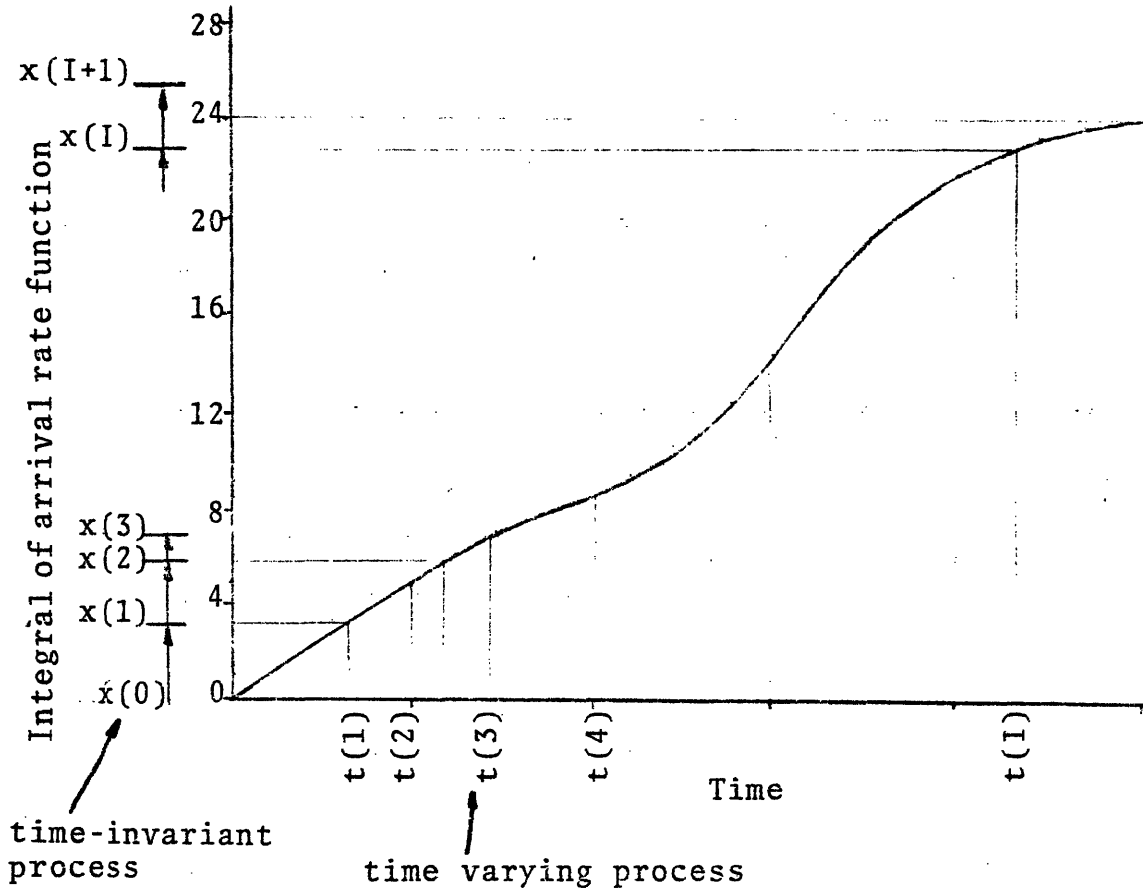


FIGURE II-5. Rate Function and Integrated Rate Function.

from zero to total number of expected arrivals in the dimension of cumulative expected arrivals is thus mapped onto the dimension of time, producing the process $t(i)$ extending from zero to the time length of the run being considered. This set of values $t(i)$ constitutes the times aircraft will enter the simulation. It is hoped that the intuitiveness of this method will be enhanced by consideration of the following derivation of the equations and numerical results for the type of cases in which we are interested.

We will derive results for the case where the arrival rate function is assumed to be piecewise linear. Figure II-6 gives the nomenclature used in this derivation. Our objective is to derive an expression for $L(t)$, the integral of the arrival rate function. We can simplify this integral by defining:

$$C(n) = L(nP) = \int_0^{nP} 1(s) ds \quad (1)$$

It is apparent that:

$$C(0) = 0$$

$$C(1) = [R(1) + R(0)] * P/2$$

$$\begin{aligned} C(2) &= [R(2)+R(1)]*P/2 + [R(1)+R(0)]*P/2 \\ &= [R(2)+2R(1)+R(0)]*P/2 \end{aligned}$$

and that

$$C(n) = P[R(n)/2 + R(0)/2 + \sum_{j=1}^{n-1} R(j)] \quad 1 \leq n \leq N \quad (2)$$

Let:

$l(t)$ - arrival rate function (aircraft/unit time)

t - time

N - number of periods of process


P - length of period

s - time from start of period n

n - current period

$R_n \equiv l(nP)$ rate at junction of periods n and $n+1$

We assume $l(t) \geq 0$ at all values of t



Further let:

$L(t)$ - integrated rate function

$$L(t) \equiv \int_0^t l(s) ds$$

FIGURE II-6. Summary of Nomenclature.

The above may also be written as:

$$C(n) = C(n-1) + [R(n)+R(N-1)]P/2 \quad 1 \leq n \leq N \quad (3)$$

Thus:

$$L(t) = C(n-1) + [R(n-1)+l(t)]s/2$$

but

$$l(t) = (1-s/p)R(n-1) + R(n)s/P \quad (n-1)P \leq t \leq nP$$

So

$$\begin{aligned} L(t) &= C(n-1) + [R(n-1)+(1-s/p)R(n-1)+R(n)s/P]s/2 \\ &= C(n-1) + [2*R(n-1)+(R(n)-R(n-1))s/P]s/2 \\ L(t) &= C(n-1) + R(n-1)*s + (R(n)-R(n-1))*s^2/(2*P) \quad (4) \end{aligned}$$

We generate points $x(1)$ to $x(I)$ along the dimensions of $L(t)$ and we want to map them onto the dimension of time. $x(i)$ gives the value of $L(t)$ but we want to know the value of t (or equivalency, n and s) for which $L(t) = x(i)$. Therefore (4) must be solved for s and n . The value of $L(t)$ must lie between $C(n-1)$ and $C(n)$ when t is in period n , so for a value of $L(t)$, n can be determined from examination of the set of $C(n)$. With n fixed, (4) may be written and solved as a simple quadratic.

$$0 = [R(n)-R(n-1)]*s^2/(2P) + R(n-1)s + C(n-1) - L(t)$$

Applying the standard formula we obtain:

$$\begin{aligned} s &= -PR(n-1)/w \\ &\pm (P/w)*SQRT[R(n-1)^2 + 2w(L(t)-C(n-1))/P] \quad (5) \end{aligned}$$

where $w = R(n) - R(n-1)$

This is a well-behaved equation. The positive root is always the desired one. The negative root represents a time beyond one end of the period in question when the projection of $l(t)$ based on $R(n-1)$ and $R(n)$ alone has a negative value such that the total area of the projected $l(t)$ from the beginning of the period up to the value of the root is equal to $l(t)$. See the numerical example in the following figure for an illustration of this.

The discriminant can be shown to be always greater than or equal to zero. We will prove that the following inequality is true:

$$R(n-1)^2 + 2[R(n)-R(n-1)] (L(t)-C(n-1))/P \leq 0 \quad (6)$$

Proof

As $0 \leq C(n-1) \leq L(t) \leq C(n)$, $R(n) \geq 0$ and $P \geq 0$ only the term $R(n) - R(n-1)$ can be less than 0. The left side of (6) would be smallest when $R(n) - R(n-1) < 0$ and $L(t) - C(n-1)$ is largest, that is when $L(t) = C(n)$. But from (3) it is known that:

$$C(n) - C(n-1) = [R(n)+R(n-1)]P/2$$

so (6) may be written as:

$$R(n-1)^2 + (2/P)[R(n)-R(n-1)](R(n)+R(n-1))(P/2) \geq 0$$

Simplifying:

$$R(n-1) + (R(n)^2 - R(n-1)^2) \geq 0$$

$$R(n)^2 \geq 0$$

As the above must be true we have established that (6) is true.

One other case must be considered, namely, $R(n) = R(n-1)$. In this case (4) becomes

$$L(t) = C(n-1) + R(n-1)t \quad (7)$$

Thus

$$t = [L(t) - C(n-1)] / R(n-1) \quad (n-1)P \leq t \leq nP \quad (8)$$

Note that should $R(n) = R(n-1) = 0$ then

$$C(n) = C(n-1) + [R(n) + R(n-1)]P/2$$

$$C(n) = C(n-1)$$

and no points of the process $x(i)$ (and thus no values of $L\{t\}$) will appear on this portion of the interval so neither (5) or (8) will ever be used in this situation. Figure II-7 presents a specific numerical example for this case.

Extension to non-Poisson processes

Although the above method was presented in terms of Poisson processes the derivation is not dependent on this assumption. Tests of the method with other distribution functions (normal, k order Erlang) show correct results using the algorithm in 4.c below. The only caution is that some probability density functions permit negative inter-aircraft times and this must be accounted for in the implementation of the process.

4.c Algorithm for generated schedules

Step 1) Input $P, R(n), N$

Step 2) Use (3) to find the set of $C(n)$

Let $R_0 = 5$ aircraft/period
 $R_1 = 10$ aircraft/period
 $P = 1$ hour

From (3) $C_0 = 0$ $C_1 = (5 + 10)\frac{1}{2} = 7.5$

As $R_0 \neq R_1$ (5) applies $w = 10 - 5 = 5$

$$s = \frac{-(1)(5)}{5} \pm \frac{1}{5} \sqrt{5^2 + \frac{(2)(5)}{1} (L_9t) - 0} \quad n = 1$$

$$s = -1 \pm \frac{1}{5} \sqrt{25 + 10L(t)}$$

For various values of $X_i = L(t)$ this equation gives

<u>i</u>	<u>L(t)</u>	<u>s</u>	<u>Positive Root tp(i)</u>	<u>Negative Root tn(i)</u>
1	0	-1 ± 1	.0	-2
2	2.5	$-1 \pm \sqrt{2}$.414	-2.414
3	3.125	-1 ± 1.5	.5	-2.5
4	7.5	-1 ± 2	1	-3




FIGURE II-7. Numerical Example.

- Step 3) Set $i = 1$, $x(0) = 0$
- Step 4) Generate X from the chosen distribution function with mean value of 1.
- Step 5) Set $x(i) = x(i-1) + X$
- Step 6) If $x(i) > C(N)$ then stop (process complete)
- Step 7) Locate n such that $C(n-1) < x(i) < C(n)$
- Step 8) Use (5) or (8) to find s using $x(i)$ for $L(t)$
- Step 9) $t(i) = s + (n-1)P$
- Step 10) $i = i + 1$
- If $i > \text{size of storage provided for schedule}$
then issue warning and stop
- Step 11) Go to Step 4.

5 Comparison of the Methods of Demand Generation

The two methods of demand generation discussed here, the ASM base schedule and the probabilistic schedule, are not equivalent even if the parameters of the two methods are matched. The probabilistic schedule will exhibit considerably greater variability than the base schedule method. For example, in the base schedule the number of aircraft will be held constant in all replications whereas that number will vary in the probabilistic schedule. The distribution of aircraft attributes will remain constant

over all replications in the base schedule whereas the distribution will vary in the probabilistic schedules. For example, if the base schedule contains 7% heavy jets, then all replications using this schedule will have exactly 7% heavy jets. The set of probabilistic schedules will average 7% heavy jets but the percentage will vary across replications. The probabilistic schedule has additional variability by comparison to the base schedule in that the distribution of aircraft attributes can vary over time within a replication. To continue the previous example, two replications of a probabilistic schedule might both contain 7% heavy jets but differ greatly with regard to the time these heavy jets appeared. In the base schedule aircraft times can be altered only via the lateness distribution which would mean for most situations that reordering of aircraft would be limited to moving aircraft one or two positions up or down in the schedule. Increasing the variance of the lateness distribution would increase the size of the sort. In fact, two replications of the base schedule method might contain periods in which identical sequences of aircraft types appear. This difference with the probabilistic schedule is of greater consequence where the demand pattern has significant peaks, since performance at peak demand is affected more by such shifts in the order of aircraft arrival than is performance in the off-peak period.

DEPARTMENT OF CIVIL ENGINEERING
Massachusetts Institute of Technology

MEMORANDUM

TO: Institute Archives

DATE: March 20, 1981

FROM: V. H. Murphy - Room 1-281, Ext. 3-7106

SUBJECT: Missing Theses Pages

Page 75 of the PhD thesis of John Nordin (Course I, 9/80) cannot be located.

A	B	C	D	E	F	G	A	B	C	D	E	F	G
1	3	2	360.8	394.3	4	2	1	4	1	361.5	384.3	5	2
2	3	3	361.0	390.3	3	2	2	1	1	362.9	414.2	2	1
3	2	2	364.2	400.9	2	1	3	3	3	364.1	405.7	2	2
4	2	2	365.5	383.5	5	1	4	4	1	369.7	409.6	4	2
5	2	1	367.8	392.6	4	2	5	2	2	382.4	402.0	2	1
6	3	2	367.9	417.7	1	1	6	2	1	384.2	419.6	4	1
7	2	3	372.4	419.8	3	2	7	3	2	386.2	436.3	3	1
8	3	2	375.5	424.1	2	2	8	4	2	387.9	411.5	3	1
9	3	1	376.5	451.3	2	2	9	3	3	390.9	425.6	2	2
10	4	1	378.1	432.8	1	2	10	2	2	392.0	390.4	2	2
11	4	1	388.2	433.7	4	1	11	3	3	397.4	450.3	1	2
12	3	3	389.1	432.3	2	2	12	4	3	398.6	431.0	5	2
13	1	3	392.6	452.7	4	1	13	2	1	400.6	455.0	2	1
14	2	3	393.5	424.3	4	2	14	4	1	402.5	439.3	5	1
15	3	3	396.1	424.7	2	1	15	3	1	404.8	414.1	4	2
16	4	3	396.8	421.5	5	2	16	3	3	410.1	413.6	4	1
17	3	2	408.1	453.6	4	1	17	2	2	412.0	459.4	4	2
18	2	1	408.9	444.1	2	1	18	3	1	414.2	464.5	3	2
19	1	1	409.1	428.9	2	1	19	2	1	419.4	467.9	3	1
20	2	1	412.8	445.5	2	2	20	4	3	419.4	464.1	3	1
21	2	2	413.3	471.5	2	2							
22	3	1	414.5	439.6	2	1							
23	4	3	417.8	454.2	2	1							

Replication 1

Replication 2

- A - Index number
- B - Aircraft class (1 = heavy, 2 = large, etc.)
- C - Arrival fix number
- D - Arrival time
- E - Departure time
- F - Apron number
- G - Departure fix number

Table II-1: Probabilistically Generated Schedules

A	B	C	D	E	F	G	A	B	C	D	E	F	G
1	3	2	361.1	394.3	4	2	1	3	2	361.6	394.4	4	2
2	3	3	363.3	390.3	3	2	2	3	3	361.8	390.3	3	2
3	2	2	374.9	400.9	2	1	3	2	2	363.9	400.9	2	2
4	4	2	363.0	388.5	5	1	4	4	2	376.2	388.5	5	1
5	2	1	366.9	392.6	4	2	5	2	1	370.2	392.6	4	2
6	3	2	374.6	417.7	1	1	6	3	2	375.3	417.7	1	1
7	2	3	380.5	419.8	3	2	7	2	3	370.5	419.8	3	2
8	3	2	384.0	424.1	2	2	8	3	2	385.3	424.1	2	2
9	3	1	382.1	451.3	2	2	9	3	1	381.1	451.3	2	2
10	4	1	378.7	432.8	1	2	10	4	1	375.4	432.8	1	2
11	4	1	383.8	433.7	4	1	11	4	1	392.2	433.7	4	1
12	3	3	390.5	432.3	2	2	12	3	3	400.2	432.3	2	2
13	1	3	393.3	452.7	4	1	13	1	3	399.4	452.7	4	1
14	2	3	398.2	424.3	4	2	14	2	3	399.0	424.3	4	2
15	3	3	403.8	424.7	2	1	15	3	3	399.7	424.7	2	1
16	4	3	401.9	421.5	5	2	16	4	3	406.9	421.5	5	2
17	3	2	410.1	453.6	4	1	17	3	2	405.0	453.6	4	1
18	2	1	411.4	444.1	2	1	18	2	1	420.6	444.1	2	1
19	1	1	414.3	428.9	2	1	19	1	1	415.7	428.9	2	1
20	2	1	412.9	445.5	2	2	20	2	1	412.1	445.5	2	2
21	2	2	414.9	471.5	2	2	21	2	2	413.3	471.5	2	2
22	3	1	415.3	439.6	2	1	22	3	1	415.7	439.6	2	1
23	4	3	424.2	454.2	2	1	23	4	3	421.3	454.2	2	1

Replication 1

Replication 2

- A - Index number
- B - Aircraft class
- C - Arrival fix number
- D - Arrival time
- E - Departure time
- F - Apron number
- G - Departure fix number

Table II-2: Base Method Schedules

aircraft but the order in which the aircraft enter the simulation has been shuffled.

The likely consequences of the additional variability of the probabilistic schedule include both a larger variance in performance estimates produced by the model and a higher level of delay. The increased variance of model outputs is an intuitively obvious consequence of increased variation of the inputs. The higher level of delay comes from the general observation that systems operating in very stochastic environments do not perform as well as similar systems in less stochastic environments.


Consider, for example, the family of single-server queues that are $E_k/M/1$ (interarrival times are Erlang of order k). This family includes $M/M/1$ ($k = 1$) and $D/M/1$ ($k = \infty$) queues as special cases. The variance of interarrival times declines as k increases. It has been shown (FIC 74) that the total expected waiting time decreases monotonically as k increases. For systems with $u = 1$, for example, waiting time decreases 28% moving from $k = 1$ to $k = 2$ at $r^* = .9$, 39% at $r = .5$, and 73% at $r = .1$ (FIC 74 pp. 901-2).

Further confirmation of this point is provided by a comparison between the two methods on a test case. A 6 hour simulation of a one runway airport was performed using FLAPS. All aircraft were through aircraft. Fifty


* r - utilization rate
 u - arrival rate

replications were performed. The arrival rate function in figure II-8 was used. The only difference in the two runs was the schedule generation method. Each method scheduled the same average number of aircraft into the airport in each period. The results in table II-3 show a halving of landing delay when the base schedule method is used. Variances are 60% to 90% lower in the base schedule method.

The different underlying structure of the probabilistic and base schedules and their different effects suggest different applications for each method. The base schedule is an appropriate model for analysis of near term changes in the airport environment such as the addition of several new flights or the opening of several new gates in one apron area, or for any case where current, exactly known conditions are to be altered in some precisely known way. The base schedule method also permits simulation of any very unusual demand situations that cannot be described statistically. The probabilistic schedule method is appropriate for longer term situations where the parameters are not as well known or when the effects of some change should be studied over a variety of demand patterns. Situations in this category include all forecasts of more than a year or two into the future or analysis of situations involving major changes in the airport environment (new runways, new aircraft types, large demand shifts) where the demand rate and pattern will be substantially different from the present.



Lateness Distribution (Base Schedule only)



Replications: 50
 Time periods: 6
 Run length: 6 hours

FIGURE II-8. Comparison of Schedule Generation Methods; Parameters.

Result	Probabilistic Schedule	Base Schedule	Units
Landings - mean	153.1	152.8	aircraft
- s. d.	9.3	0.9	
- c. i.	2.8	0.3	
- min.	136	150	
- max.	173	154	
Landing Delay:			
- mean	10.2	4.5	minutes
- s. d.	5.1	0.8	
- c. i.	1.5	0.2	
- min.	2.7	2.8	
- max.	25.9	6.2	
Landing queue at end of period 3:			
- mean	11.3	5.6	aircraft
- s. d.	6.1	1.7	
- c. i.	1.7	0.5	
- min.	2	2	
- max	25	10	
Takeoffs - mean	127.3	134.7	aircraft
- s. d.	7.9	2.5	
- c. i.	2.2	0.7	
- min.	111	125	
- max.	145	138	
Takeoff delay:			
- mean	19.7	17.8	minutes
- s. d.	9.2	5.3	
- c. i.	2.6	1.5	
- min.	1.8	9.7	
- max.	41.9	33.8	

s. d. - Standard deviation
c. i. - 95% confidence interval of the mean
min., max. - Lowest and highest values observed in
any one of the replications.

Table II-3: Comparison of Schedule Generation
Method: Results

Using one type of schedule in a situation intended for the other constitutes a misspecification of the model because of the different assumptions implicit in each schedule type. Using the base schedule implies that the user knows exactly the aircraft mix, number of aircraft and time-of-day distribution of aircraft parameters, an inappropriate assumption for many situations.

Because of the difference between the two types of schedules the model allows use of both the base schedule method (with the actual base schedule obtained externally to the model) and the probabilistic schedule method. A third option is permitted: the base schedule method may be used with the base schedule being created by using the algorithm normally used to create probabilistic schedules. This method was used to do the schedule comparison given above.

D RUNWAY OPERATIONS

1 Introduction

The purpose of this section is to analyze the movement of aircraft on runways and formulate a model of these activities. The scope of runway operations is different for landings and for takeoffs, as shown in figure II-9. Landing operations are followed from the start of the common approach path, typically 5 or 6 miles out, through touchdown on the runway, braking and exiting from the runway. Once a landing aircraft clears the runway becomes a ground operation and its movements are discussed in section F below. Aircraft takeoffs are considered runway operations from the moment they taxi into position on the runway for their takeoff roll. Their movement is followed until they pass over the far end of the runway.

There are two major aspects of modeling runway movements: 1) the performance of aircraft, i.e. how long an aircraft occupies the runway and what exit an aircraft uses. As landings and takeoffs have very different performance characteristics, they are discussed separately below. 2) The separations between successive aircraft that use the runway. This topic is discussed in part 4 below for all types of aircraft movements.

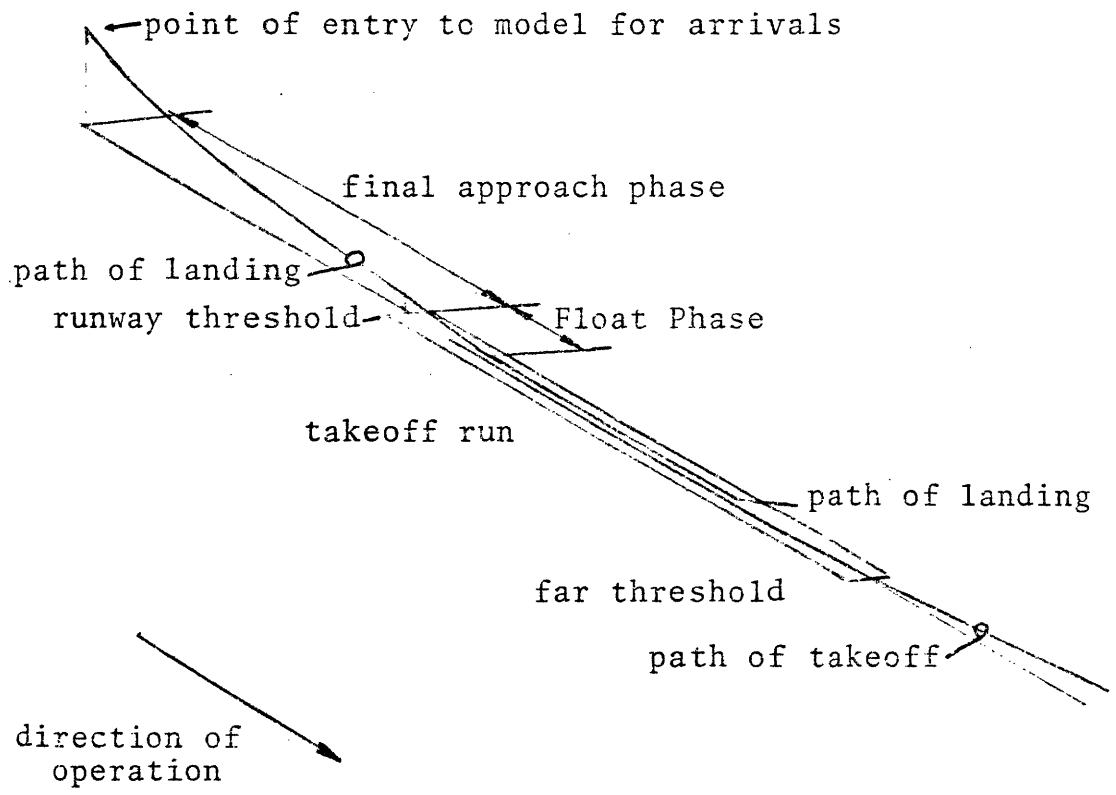


FIGURE II-9. Runway Operations.

2 Aircraft Performance on Landing

We will first review two methods that have been used to model the movement of landing aircraft and the advantages and disadvantages of each method. In response to the problems of existing methods, we will discuss an alternative approach to modeling landing performance. This new method was tested with two sets of data and the results of the validation are discussed.

2.a Methods of representing aircraft performance

We are interested in modeling two key aspects of the performance of landing aircraft: runway occupancy time and choice of exit. If only one runway is to be used for landings then, the obvious way to proceed is merely to input the appropriate data directly. For each exit the probability of exit, average time to exit and standard deviation of time to exit would be provided, probably as a function of aircraft class. This procedure was used in the MITASIM airport simulation (NOR 79). This procedure has several advantages. It can provide the exact behavior desired and it is easy and fast to program and use. If more than one landing runway were to be used, then additional sets of information for each runway would need to be prepared and input. Note that if one runway is used in both directions, two sets of information must be provided as

travel time to a given exit from one end of the runway will not be equal to travel time to the same exit from the opposite end of the runway.

The problem with this direct input method is that it is not behavioral. The user is required to specify as an input an aspect of performance (exit selection) which is really a result of both aircraft and runway characteristics. The method implicitly assumes that exit use, occupancy time, exit location and exit type are all independent of each other. If the approach speed of a landing aircraft changes or the exit location changes, then, in reality, runway occupancy time and exit choice are also affected but unless the user specifically alters them they will not change in the model.

A second method which avoids some of these problems is to input two functions for each aircraft class: 1) runway occupancy time as a function of exit distance, and 2) probability of exit as a function of exit distance. This method, using piecewise linear functions, was employed in the ASM airport simulation (PMM 77). In this procedure, the average occupancy time and probability of exit are both determined from the input functions depending on the distance of the exit from the runway threshold. The runway occupancy time is used directly in the model. The probability of exit obtained in this way is used as the parameter in a binomial draw to determine if the aircraft actually exits at this point.

The procedure will result in different aircraft performance if exits are moved, since the functions use distance of exit to derive occupancy time and exiting probabilities. However, changing approach speeds does not change aircraft performance since the functions do not use approach speed to derive runway performance. Additionally, the procedure as implemented does not allow performance to be affected by exit type (90 degree, angle*, high speed) but could easily do so by providing one function for each type of exit. The input requirements of this method are greater than the first method discussed, though the absolute amount is not a major burden.

The more significant problem with this method is that it is still the "inputting of conclusions". If the effect of exit location on runway performance is one of the questions to be answered by using the model, then the results cannot be an input. Having runway performance as an input requires that a separate, external model of aircraft performance be used to generate this data for use in the airport model. The method proposed in the next section is an attempt to solve this problem.

2.b Landing performance logic

In this section we start with the physical description of the basic, well-known procedure used by landing aircraft and derive the underlying equations and parameters of the

* Exits with an angle to the runway of 45 to 75 degrees.

process.

The relevant literature has been reviewed for data on the values of the parameters. It is unfortunate that many aspects of this procedure have not been researched, at least in ways that are helpful for our objective of comprehensive model building.


Data regarding runway occupancy times and related analyses are relatively scarce. Almost all work to date has focused on relationships of optimal performance characteristics of specific aircraft types and/or the placement and design of individual exits. This emphasis tends to disregard the operational need for consistent, sufficiently low runway occupancy times. . . . While few question the fact there exists a disparity between optimal and presently observed runway occupancy times, little attention has been directed towards reasons for these differences as they relate to airline, exit, aircraft, runway and airport. (KOE 78 p.1-1).

The published literature tends to supply primarily ad-hoc estimates of mean values of critical parameters. Some information is available on variances and still less on distributions of parameters. Little work has been reported on the relationship among parameters. This scarcity of information forces us into developing a model of landing performance that cannot be validated in a precise statistical sense. Nor are we always able to test alternative specifications of the model. What we must do instead is to derive a structure of the model from what seems most likely to be the case and limit the parameters in use to reasonable values. Despite this disappointing situation, it is possible to derive a fully behavioral model of landing performance that produces acceptable results.

Aircraft landing movements are modeled in five segments (see figure II-10): 1) final approach, 2) float, including flare and transition, 3) deceleration, 4) coasting to exit, 5) exiting. Each of the five phases will be considered in turn. For each phase we will consider information in the literature relevant to this aspect of landing and derive the logic for that phase.

Final Approach

The final approach phase refers to the section of landing from the beginning of the common approach path until the aircraft is over the runway threshold. The standard method of modeling the final approach phase is to assume that each aircraft flies at a constant velocity over a final approach of about 6 miles in length. This assumption is apparently adequate, as little discussion of this phase has been found in the literature. An important observation is that care must be used in specifying approach speeds. Minimum values, obtained from aircraft performance data are probably lower than speeds typically flown. The average value of approach velocity is clearly a function of aircraft type, varying from 140kts. for aircraft such as the B747 to 115kts. for a DC9. Thus a set of average approach speeds, $V_a(i)$, is specified for each aircraft class, i , to use the model. The standard deviation of approach speed, $V_{as}(i)$, among aircraft in any one class is approximately 4kts. A triangular probability density function is often used to



$v(i)$ - Approach Velocity
 V_f - Velocity drop during float phase
 V_c - Coasting velocity
 $V_e(n)$ - Exit velocity for exit n

FIGURE II-10. Landing Aircraft Profile.

draw $v(i)$, the actual approach velocity for a particular aircraft i , from $V_a(i)$ and $V_{as}(i)$. Once $v(i)$ has been determined for a given aircraft, then the time the aircraft requires to fly the approach path is $C/v(i)$, where C is the length of the common approach path.

In summary the parameters needed to model this phase are:

- 1) $V_a(i)$ - mean approach speed, for each aircraft class (knots)
- 2) $V_{as}(i)$ - standard deviation of approach speed (knots)
- 3) type of probability density function - used to draw samples from $V_a(i)$ and $V_{as}(i)$
- 4) C - length of common approach path (miles)

Float

The float phase refers to aircraft movement from the runway threshold to the point on the runway when all landing gear have made firm contact with the runway and braking may begin. This phase includes the flare and transition maneuvers. Transition is the period from the point of first touchdown of the main gear until the nose wheel has contacted the runway.

The typical transport aircraft is intending to come over the runway threshold at an altitude of 50 feet, flare and touch down on the main gear some distance down the runway. The aircraft then takes several seconds to

transition, i.e., to have the nose gear touchdown (HOR 75 p. 226, KOE 78 p.4-1).

Review of the literature shows that the mechanics of this process are very subtle. Table II-4 lists the

Parameters:

- Df - Actual distance used by a given aircraft from runway threshold to the end of the float phase (feet).
- A - Deceleration rate in air during the float phase (feet/s/s).
- Vf - Total velocity drop during float phase (knots).
- v(i) - Approach speed of aircraft i (knots).
- Va(i) - Average approach speed of aircraft class i (knots).
- Vas(i) - Standard deviation of approach speed of aircraft of class i (knots).
- Xf(i) - Mean float distance of all class i aircraft (feet).
- Xfs(i) - Standard deviation of float distance of all class i aircraft (feet).

Table II-4: Float Phase Parameters

parameters for this phase and defines additional quantities used in the derivations in this section.

Despite the critical nature of this process, little comprehensive analysis of it seems to have been done. Various authors provide data on one or more parameters and occasionally on the ranges of the parameters. Horonjeff indicates (HOR 75 p.228) that the floating distance is 1500

ft. for transport aircraft and 1000 ft. for twin-engine general aviation aircraft. Swedish measured this distance at the Boston and Atlanta airports (SWE 72) and found values of 1360 feet (standard deviation 469 ft.) for DC9 aircraft and 1509 feet (standard deviation 453 ft.) for B727 aircraft. Large aircraft (DC8, B747) had averages of 1650 feet. Convair 580's averaged 930 ft. Swedish also provides selected graphs of the distribution of touchdown distances which indicate that this parameter is highly variable. It seems important therefore, to model this variation.

Horonjeff indicates that the in-air deceleration rate is about 2.5 f/s/s. He also indicates total velocity drop is about 5 to 8kts (HOR 75 p.229) as does Boeing (BOE 69 p.3-164). Time through this phase is indicated by Boeing to be 7 to 11 seconds. Approach speeds have been discussed above.

In addition to the parameters, we also need to consider the equations that govern this process. In the absence of more detailed information we assume that the standard equations for movement of a body under constant acceleration are sufficiently accurate.

$$V = A*t + V_0 \quad (9)$$

$$X = A*t^2/2 + V_0*t \quad (10)$$

Equations (9) and (10) state the time-distance-deceleration relationship for a general acceleration, A, distance, X, initial velocity, V_0 , and time, t. These two equations will

be applied at several points in this section to various runway performance calculations.

Applying these equations using the float phase parameters given above reveals the unfortunate fact that the various given values are not exactly compatible. See figure II-11 for the graph of the parameter space for an aircraft with a V_a of 120kts, a typical value for aircraft like the B727. Figure II-12 shows how the parameter space moves with differing values of V_a . Matching the given values of float distance will result in either a very low deceleration or a very high V_f . Reducing the distance brings the time in this phase below the values in the literature. Some of the problem may be due to the omission of the transition phase in the touchdown data. Swedish's data apparently does not include transition distance. Koenig indicates that some deceleration occurs in the transition (KOE 78 p.4-3). It would seem unlikely that this deceleration is very significant, since transition takes only two or three seconds. If we assume that V_f is somewhat higher (say 8kts.), then deceleration rates will be within reason but still low. Logically float distance should be increased (to around 1800ft) to include the transition phase but this turns out to cause problems. In the course of testing the model it was determined that longer float distances necessitated higher braking rates to achieve the correct exit selection probabilities. This rapid deceleration





FIGURE II-11. Interrelationship of Aircraft Floating Phase Parameters.

resulted in unacceptably low runway occupancy times. It would be possible to reduce the significance of this problem by increasing V_f still further to reduce the amount of braking needed. Unfortunately a V_f of 8kts is already at the high end of the reported range. The author has discussed this problem with individual pilots and has read a number of operations oriented articles for private pilots. These sources indicate that aircraft fly the final approach at 1.3 times the stalling speed and touch down at or near stalling speed. For typical air transport aircraft this would involve a velocity drop of 30 knots. If all of this velocity drop were taken after the threshold was crossed it would result in unacceptably large deceleration rates. It may be then that some deceleration occurs prior to crossing the threshold. We chose to resolve this problem in the direction of leaving float distance at 1500ft. and V_f at 8kts. Rigorous resolution of this problem must await the availability of better data. Studies are needed which measure all relevant parameters for each aircraft in order to permit study of the interaction of parameters.

No information has been found on how V_a is related to float distance. It would seem logical that aircraft with fast approach speeds (for their class) would tend to float a longer distance down the runway than aircraft of the same type with slower approach speeds. On the other hand, the



$V_f = 8$ kts

FIGURE II-12. Effect of V_a on Floating Phase Parameters.

correlation between the two is most unlikely to be perfect. Therefore the following compromise procedure is used, in lieu of data describing the correlation. A mean float distance ($X_f(i)$) and standard deviation of float distance ($X_{fs}(i)$) are specified for each aircraft class. The actual distance (D_f) that any given aircraft takes to float and to transition are drawn according to the following equation:

$$D_f = X + B(v - V_a)/V_{as} \quad (11)$$

where V_{as} is the standard deviation of V_a and v is the actual approach speed of the aircraft. We have omitted the subscript i on $v(i)$ and $V_a(i)$ in the equation and following discussion for clarity. The intent of the chosen equation is to have actual float distance be an equal function of actual approach speed and a random component. The first term, X , is intended to represent this random component and the second term, $V(v - V_a)/V_{as}$, is intended to represent the contribution of approach speed to determining actual float distance. X is assumed to be drawn from a normal distribution with:

$$E(X) = A \quad \text{VAR}(X) = C$$

where $E(X)$ is the expected value of X and $\text{VAR}(X)$ is the variance of X . We wish to select A, B and C so that $E(D_f) = X_f$, $V(D_f) = X_{fs}^2$ and the contributions to the variance of the two terms are equal. Taking expected values:

$$\begin{aligned} E(D_f) &= E(X) + E(B(v - V_a))/V_{as} \\ &= A + (B/V_{as}) * (E(v) - E(V_a)) \end{aligned}$$

but as the expected value of v is V_a

$$E(Df) = A$$

Thus A should equal the expected value of Df which is X_f .

Taking variances:

$$\begin{aligned} \text{VAR}(Df) &= \text{VAR}(X) + (B/V_a)^2 * \text{VAR}(v - V_a) \\ &= \text{VAR}(X) + (B/V_a)^2 * \text{VAR}(v) + (B/V_a)^2 * \text{VAR}(V_a) \end{aligned}$$

But $\text{VAR}(V_a)$ is zero and

$$\text{VAR}(v) = V_a^2$$

so

$$\text{VAR}(Df) = \text{VAR}(X) + B^2$$

If we wish the two parts to have equal weight then

$$\text{VAR}(X) = \text{VAR}(Df)/2$$

or

$$\text{VAR}(X) = X_f^2/2$$

and

$$\text{VAR}(Df)/2 = B^2$$

or

$$B = X_f/\text{SQRT}(2)$$

where $\text{SQRT}(2)$ is the square root of 2.

To summarize:

X is normal with mean X_f and standard deviation $X_f/\text{SQRT}(2)$

and

$$Df = X_f * (v - V_a) / (V_a * \text{SQRT}(2)) \quad (12)$$

The application of this formula can be illustrated by the following example. Suppose that an aircraft class had an average approach speed of 120kts., a standard deviation

of approach speed of 4kts., a mean float distance of 1500ft. and a standard deviation of float distance of 500ft. Then all the aircraft of this class that were modeled as having an approach speed of 124kts. (1 standard deviation over the mean) would have a float distance (Df) described by

$$Df = C + 500(124-120)/(4*\text{SQRT}(2))$$

Simplifying

$$Df = X + 355$$

Fast aircraft of a given class are therefore modeled as taking a longer distance to flare and to transition than slower aircraft of the same class. When implementing the model v is first chosen randomly from V_a and V_{as} . Then X is chosen randomly from $X_f(i)$ and $X_{fs}(i)$. X and v are used to determine Df . Then, assuming $V_f = 8\text{kts.}$, we solve for t using equations (9) and (10) with v in place of V_o and $v - 8$ in place of V . The resulting value of t is time in the float phase.

Deceleration and coasting

Deceleration and coasting is the phase of landing from the beginning of braking to the moment of exit from the runway.

It appears that little attention has been given to measuring actual deceleration rates for transport aircraft. Boeing does present (BOE 69 p.3.60) a very complete equation for acceleration and deceleration of an aircraft on the runway as a function of thrust, drag, lift, velocity,

braking and runway slope but does not give values for the various constants in the equation. Horonjeff (HOR 75 p.228) indicates that deceleration averages 5 f/s/s. This figure may originate from studies done prior to the introduction of aircraft types like the B737 and DC9 that are designed to operate from short fields and thus might have higher deceleration rates. Horonjeff does not indicate how this value might vary from aircraft to aircraft. Koenig states (KOE 78 p.4-1) that deceleration varies from 5.5f/s/s to 10f/s/s depending on whether moderate or hard braking is used and indicates that from his data, moderate braking is more typical. No information is provided in either source on how this would vary over aircraft class. It would also seem logical that deceleration would be lower on wet compared to dry runways, but no information as to the magnitude of the change has been found.

The information that is usually collected is (not surprisingly since it is readily observable) runway occupancy time. In the next section we use this information to calibrate the landing performance model, and we find that using mean values in the range of 5.25 to 5.75 feet per second squared with a standard deviation of 0.75 f/s/s gives acceptable results across all aircraft classes and runways.

Little information was found about how aircraft deceleration rates vary during the time when an aircraft is on the runway. It seems logical that aircraft maintain full

deceleration down to some safe speed and then coast (at a low deceleration rate) until they are near the exit to be taken in order to minimize time on the runway. It would also seem logical that this would occur at a higher speed than aircraft taxi which is 10 to 20 kts. In the validation it was found that a coasting speed of 25 knots gave good results. For simplicity, we assume that the coasting phase is conducted at constant speed.

Exiting from the runway

There are various types of exits from runways. These vary from high-speed exits to sharp 90 degree turns (and sometimes even greater than 90 degree exits). The speed at which aircraft may clear the runway will vary depending on the type of exit. To use the model, an exit speed, V_e , is specified for each exit. Suggested values for V_e are zero kts. for 90 degree exits, 10 kts. for angle exits and $V_e = 30$ for high speed exits.*

An important question is whether there is any interaction between deceleration rate and exit location. One possibility is that pilots would try to leave the runway closer to their intended terminal gates. Koenig divided his data between "motivated carriers" who had an incentive to exit early and other carriers, who did not. On a majority of runways examined motivated carriers had runway occupancy times of 4 to 6 seconds less than the remaining carriers (KOE 78 p. 3-3). On several other runways there was no

* Conversation with Professor Odoni, MIT.

significant difference. This behavior was not modeled. The data available does not permit the fitting of a model of this process, since Koenig does not report which airlines were motivated to exit early for each runway. However, the basic framework does exist in FLAPS for implementing such a model, since each particular aircraft is tracked throughout the simulation. Therefore, the ground destination (apron area) of each aircraft is known at the time the aircraft begins to land. An extension of FLAPS could model this interaction.

Because of the data limitations, a second, more limited interaction between deceleration rates and exit location, dependent only on runway conditions, was developed. It was found to be useful in enhancing the performance of the model across several runways. It is based on the idea that pilots would like to avoid long periods of coasting on the runway caused by narrowly missing an exit. If an aircraft using normal braking will miss an exit by less than some critical amount (so that the miss is perceived as "close") and there is no close subsequent exit thus necessitating a long rollout, it is assumed the pilots will employ harder braking to make the earlier exit. This is conditional on the harder braking being within the performance capabilities of the aircraft. Conversely, if a long rollout is inevitable given the placement of exits, it is assumed that a pilot will elect lighter braking to reduce runway occupancy time.

Experimentation with this model showed that using 600ft. as the definition of "near miss" and as the lower bound on "long roll" narrowed the divergence between predicted and actual performance. Aircraft deceleration rates were adjusted only if they were on the "wrong" side of average braking rates, that is, harder braking was employed on those aircraft that just missed exits only if their original braking rate was below average for that aircraft class. Aircraft with long rollouts were decelerated more slowly only if their original braking rate was above average.

Because the deceleration, coasting and exit phases are closely related, we summarize them together.

Once all of the landing gear have made solid contact with the runway, the deceleration phase begins. Aircraft are assumed to decelerate at a constant rate, A_r , that is drawn from a class specific mean and standard deviation. For each exit, n , on the runway an exit velocity, $V_e(n)$, is specified. $V_e(n)$ is the maximum speed at which aircraft can negotiate exit n . Note that $V_e(n)$ is direction specific. A high speed exit in one direction is a very sharp exit in the other direction. It is assumed that aircraft will take the first exit for which the aircraft can decelerate to a speed $V_e(n)$ or less for that exit. The deceleration rate can be adjusted either up or down to minimize near misses of exits and the runway occupancy time for aircraft on long rollouts.

This determines where aircraft clear the landing runway. The aircraft is assumed to actually brake from the speed it had when entering this phase (V_a - 3kts) to a coasting speed V_c , which is a specified value constant across all aircraft types. The aircraft coasts until it is near the place where it will clear the runway when it resumes decelerating at the previous rate to arrive at the chosen exit with speed equal to $V_e(n)$.

2.c Validation

Both Swedish and Koenig report data on runway occupancy times for a number of airports, runways and aircraft types. The landing performance model described above was checked against this data for four airports. The model was run for two Boston runways (data from SWE 72) and for one runway from New York LaGuardia, Los Angeles and Buffalo (data from KOE 78). For the Boston airport three different cases were run for each runway, one for each of three aircraft models. The Koenig data is for "class 3" aircraft which include BAC111, DC9 and B727 types. Note that the only change of parameters to model these cases was to change approach speed and the airport specific parameters of exit distance and exit speed. An adjustment between the Koenig and Swedish data sets had to be made as the two sets of data are incompatible. Koenig reports consistently earlier exits than Swedish. This discrepancy was resolved by using the

specific float distances reported by Swedish to fit the Swedish cases and the specific deceleration rates reported by Koenig when fitting the Koenig data. Swedish does not report deceleration rates nor does Koenig report float distance. A longer float distance and lower deceleration rate was used on the Koenig data. Table II-5 describes the

Parameter	Value used for		Units
	Swedish data	Koenig data	
Deceleration rate:			
mean	6.00	5.25	f/s/s
standard deviation	.75	.75	
p. d. f.	normal	normal	
Floating distance:			
mean (DC9,B727,B707)	1500	(not used)	feet
mean (class 3)	(not used)	2000	
standard deviation	450	450	
Approach speed:			
mean DC9	115		knots
mean B727	125		
mean B707	140		
mean class 3		120	
standard deviation	4	4	
Coasting speed:		25	knots
Velocity drop in float phase:		8	knots
Near miss cutoff*:		600	feet
Critical coasting distance*:		600	feet
Increment on deceleration rate*:		1.0	s. d.

p. d. f. - probability density function

s. d. - standard deviation

*see text for details

Table II-5: Landing Model Validation:
Common Parameters

common assumptions used for each case. The specific runways

Runway	Exit	Distance	Type	Ve
Boston	E	4500 feet	High Speed	30 knots
Logan	S	7000	90 degree	0
27				
Boston	F	4150 feet	High Speed	30 knots
Logan	J	5200	Angle	10
33L	T	6800	Angle	10
	North	7650	Wide 90 deg.	10
New York	6	3678 feet	90 degree	0 knots
La Guardia	7	4189	90 degree	0
22	8	4955	High Speed	30
	9	5160	90 degree	0
	10	6233	90 degree	0
Los	31	4136 feet	90 degree	0 knots
Angeles	32	4666	45 degree	15
25R	33	5515	45 degree	15
	34	6787	45 degree	15
	35	7424	45 degree	15
Buffalo	6	4768 feet	45 degree	15 knots
	7	5178	90 degree	0
	8	6203	45 degree	15
	10	7997	90 degree	0

Table II-6: Landing Model Validation:
Case Description

used for the validation are described in table II-6.

The procedure used to validate the model was to try to find a set of parameters (A, Vf, Vc, D, Ar, Ve) that would produce acceptable results across several runways and aircraft types without 1) straying beyond reasonable values for the parameters, and 2) without extensive alteration of

values over the cases being considered. The presentation of the model in the section above discussed a number of these tradeoffs. Table II-7 presents the results for the final model specification.

The model matches the two Logan runways quite well. The major discrepancy is for B707 aircraft on runway 33L. Two factors should be considered. First, note the small sample size of the existing data - only 2 or 3 aircraft per exit. So exit distributions could in truth be quite different from those reported. Secondly, Swedish reports exit by exit runway occupancy times (SWE 72 p.89). For exit T, 6800 feet down, B707's exit in 51 seconds. For exit "North", 7650 feet down B707's exit in 83.7 seconds. This 32 second additional runway occupancy time when using the further exit compares to 4 and 12 second reductions in occupancy time reported for DC9 and B727 using the further exit on the same runway. So it may very well be that were more data taken, B707 average runway occupancy times would decline from 55 seconds.

The model's performance on the Koenig data is close but not generally as good as on the Logan data. The model tends to spread aircraft out over the exits more than the data indicate should happen. Model runway occupancy times tend to be low. Sample size is not a problem here.

A number of changes could be explored in the model parameters to bring the predicted performance into closer

Aircraft type	Exit number					Runway occupancy time		
	1	2	3	4	5	mean (st. dev.)	N	
Boston Logan runway 27:								
DC9	Actual:	.80	.20			50.1	30	
	Model:	.83	.17			47.5 (15.0)	2000	
B727		.64	.36			48.3	21	
		.55	.44			53.2 (18.5)	2000	
B707		.50	.50			53.2	8	
		.15	.85			60.2 (15.2)	2000	
Boston Logan runway 33L:								
DC9	Actual:	.74	.16	.06	.04	42.4	52	
	Model:	.66	.29	.04	.00	40.9 (7.9)	2000	
B727		.41	.20	.32	.07	49.1	41	
		.33	.46	.16	.00	43.4 (9.8)	2000	
B707		.2	.3	.2	.3	55.4	9	
		.06	.33	.54	.06	50.2 (9.0)	2000	
New York La Guardia runway 22:								
Class 3								
	Actual:	.01	.07	.59	.06	.26	43.3 (9.5)	314
	Model:	.02	.07	.40	.00	.51	47.8 (7.8)	2000
	Adjusted model:	.01	.09	.53	.00	.37	44.9 (7.9)	2000
Los Angeles runway 25R:								
Class 3								
	Actual:	.02	.14	.55	.25	.05	52.6 (14.1)	138
	Model:	.06	.19	.41	.25	.08	48.5 (6.9)	2000
Buffalo runway 23:								
Class 3								
	Actual:	.15	.10	.72	.02		55.9 (8.7)	124
	Model:	.26	.20	.43	.11		50.7 (9.0)	2000

Table II-7: Landing Model Validation: Results

agreement with the data. The obvious solution, of reducing the variance in the random variables, is hard to evaluate as so few data exist on the variances of a number of the parameters. Variance in approach speed and float distance account for most of the variance in performance, but these values are also the best documented and thus the most difficult to justify altering. One possible change in parameter values is presented in table II-7 for La Guardia under the title of "Adjusted Model". This is intended to show that the degree of discrepancy in the model is in line with the uncertainty in the parameters. The adjusted model for La Guardia differs from the original model in that 1) the standard deviation of float distance is reduced from 450 to 300 feet, and 2) average deceleration rate is increased from 5.25 to 5.75 f/s/s. These changes bring the results into line with observed data. The same changes on the other test airports, however, do not produce similar improvements. Extensive model fitting and sensitivity analysis is beyond the scope of this work and was not done.

It is worth emphasizing again that no attempt was made to fine tune the model for runway specific behavior. Obviously such tuning of deceleration rates, coasting speeds, floating distances and exit velocities could cause an exact match of predicted to actual behavior as the model has many more parameters than there are independent data points. Doing this, however, would negate the value of the

exercise as a validation. This is because the objective of the text is to answer the question: can this form of the model accurately represent the major modes of behavior of landing aircraft or is a more detailed model necessary? It appears from the results given above that the present model is quite satisfactory.

3 Aircraft Performance on Takeoff

The principal issue in modeling departure movements on runways is to determine how long each aircraft will occupy the runway. Much less analysis of this issue has been done for takeoffs than for landings. There are no studies of takeoffs equivalent to the landing studies reviewed above. The only information available is runway occupancy time data. This is somewhat surprising since the determination of the minimum required length of runways is set primarily by the requirements for takeoffs (HOR 75 pp.66-78).

There are a number of reasons that may account for the absence of analysis in this area. First the question of proper location of exits does not arise for takeoffs. Whereas landing runway occupancy time can be minimized by accurate placement and location of exits, there is little the airport designer can do to affect takeoff performance other than changing the grade and orientation of the runway. Secondly, as will be seen in the next section, the minimum

separation between takeoffs is both simpler and shorter than between landings. Finally, the dynamics of takeoffs are simpler and less variable than for landings. The combination of these factors may account for the comparative neglect of analysis of takeoffs.

In view of this situation there is little that can be done to represent this operation other than to adopt a simple model using runway occupancy time as the primary variable. A mean and standard deviation of runway occupancy time for each class of aircraft are used as inputs to FLAPS. Each aircraft's actual runway occupancy time is drawn from this distribution.

The primary use in FLAPS of runway occupancy time is to determine the time each aircraft needs to cross taxiways and intersecting runways while occupying the runway. This is needed to assess proper separation for conflicting movements. To derive intermediate times we assume that aircraft accelerate uniformly. To adjust for different runway lengths, runway occupancy times are defined as time from a standing start at the end of the runway to a point 6000 feet down the runway. Then time to any intermediate point from the start can be found as follows:

If X is distance, t time and a acceleration, then from

(9) and (10):

$$X = (a/2)t^2$$

and

$$t = \text{SQRT}(2X/a)$$

Let T be time to 6000 feet

$$T = \text{SQRT}(2*6000/a)$$

or

$$2/a = T^2/6000$$

so

$$t = \text{SQRT}((T^2/6000)X)$$

or

$$t = T*\text{SQRT}(X/6000) \quad (13)$$

4 Separation of Aircraft Movements

4.a Introduction

The separations imposed between aircraft are both a major factor in determining the capacity of a runway and a major area of work for improving capacity (SIN 75 p. 3-1). This section discusses the procedures and rules by which aircraft are separated on a runway that is being used either for landings, or for takeoffs or both. A discussion of separations involving operations on parallel or intersecting runways is presented in section II.E. We will first introduce the framework necessary to discuss this topic. Each type of potential conflict will then be analyzed and the appropriate rules translated into consistent minimum separation formulas.

4.b Conceptual Framework

Separations are required to maintain the safety of operations which are potentially in conflict with each other. They are intended not only to prevent collisions but also to eliminate the hazard to aircraft of encountering the wake turbulence of large aircraft.

Obviously the concept of a separation involves a pair of aircraft. As there are two types of runway movements, landings and takeoffs, there are thus four categories of separations for a runway:

- 1) Arrival/Arrival (A/A) Arrival movement followed by another arrival.
- 2) Arrival/Departure (A/D) Arrival followed by a departure.
- 3) Departure/Arrival (D/A) Departure followed by an arrival.
- 4) Departure/Departure (D/D) Departure followed by another departure.

The first term of the pair refers to the type of movement that has already begun its operation. The second of the pair refers to the operation being scheduled.

Some separations are logically specified in terms of times between crossing some reference point, others as distances. FLAPS allows all separation types to be input either as a time or as a distance and converts each internally to a time for use during the run. Because of the wake turbulence issue and the different times aircraft require to cover the same distance, the majority of

separation types involve a matrix of values, one for each possible pair of aircraft classes.

Aircraft undertaking runway operations are at some point "cleared" by air traffic control to proceed through to completion of the desired movement. Beyond this clearance point the aircraft is committed and is not diverted or delayed, save in the case where an actual collision becomes a real danger. For takeoff movements this point of commitment clearly occurs when the controller gives permission to begin the takeoff roll. For landings the point is less clear as aircraft are funneled into the final approach pattern while spacing and sequencing has already been accomplished to a large extent upstream from the runway. The close cooperation among various controllers means that the point of transition from terminal airspace to the landing runway is blurred. However, at some point aircraft converge onto a common final approach path, usually 5 or 6 miles from the runway threshold. This, as will be discussed, is the portion of flight where the final interarrival separations are applied, so the point of convergence to the final approach path makes a convenient point at which to assume that landings are first committed. The point of commitment is significant to the issue of aircraft separations because the separations are checked and enforced when an aircraft is first eligible to be committed to a runway operation. The separations are applied against

aircraft that are already committed. If the minimum separation requirement is met, the aircraft proceeds to execute the runway operation in accordance with the performance rules discussed in the previous sections. If the minimum separations are not met, the aircraft is held until the first time when the separations are met. Where more than one separation type applies, all must be met before the aircraft can be committed. This means that the longest applicable separation controls the release of the aircraft.

The entire set of separations do not apply to each aircraft. On a runway used just for arrivals the arrival/arrival separation will be the only type that is used. On a departures only runway, only the departure/departure separation is relevant. To discuss the situation for runways using mixed operations, that is, runways used for both arrivals and departures at the same time, it is necessary to anticipate somewhat the discussion of airport control below in section II.F. Discussions with airport personnel and controllers indicate that runways used for mixed operations are operated in one of three ways (Also see WEI 80 p.37).

1) Arrival priority. In this scheme arrivals are given priority over departures. Arrivals thus need only be scheduled to meet adequate separation on final approach using the A/A separation without consideration of

departures. Departures are fitted in between arrivals whenever the space between arrivals permits. The arrival/departure separation is thus used (in addition to D/D separations) for departures, but the departure/arrival separation is not used for scheduling of arrivals. Departures can only go if adequate time exists before the next arrival. D/A separations are used to assess this, as explained in the operations section below.

Occasionally in the literature this mode of operation is incorrectly referred to as a "preemptive priority" system. It is actually a non-preemptive priority system, as once a departure operation (the lower priority) is committed to a takeoff roll, the appearance of a higher priority arrival does not cause the takeoff to abort its roll.

2) Alternating operations. Here, added spacing is used between arrivals, when necessary, to insure that at least one departure can be released between successive arrivals. Implementation of this procedure requires that all four separations be used.

3) Departures only. In this method, used only when many departures have queued up for takeoff, landings are "turned off" until the departure queue has been worked down to some acceptable level. This method differs from the case of a departures only runway in that here, arrivals continue to be assigned to the runway and form a queue waiting for the runway to open again for landings. This method will need

only the D/D separations whenever it is in effect. There will be a transition period during which arrivals already committed to land complete their operations. A/D separations will be needed for this case.

Throughout the remainder of this work we will refer to the three operating schemes as "modes of operation" for runways. Our attention will be primarily directed to the first two modes.

4.c Arrival/Arrival separations

Arrival/Arrival separations have probably been the most carefully analyzed in the literature. There is now a generally accepted procedure for studying this case. See HOR 75 p.140-5, WEI 80 pp.43-9 among others. We will briefly outline the technique, the details of which are in the cited references. We will then explain why this method must be revised slightly for use in FLAPS and then derive the specific formulas for the revised procedure.

The method begins with a matrix of minimum separations for each possible combination of leading and trailing aircraft classes. These separations are stated as the minimum allowable distance between the given pair of aircraft at any point after both are on final approach. These distance separations are then translated into time separations for use in FLAPS. It is assumed that each aircraft in a class maintains the same, constant speed,

Va(i) during final approach. As the minimum separation is applied at the point of closest approach between the two aircraft, two cases must be considered: 1) Lead aircraft faster than trail aircraft - opening case, and 2) Trail aircraft faster than lead - closing case. The situation in which lead and trail aircraft are of the same class and thus have the same speed may be included in either case. Once these time separations are determined a buffer is added to each time to allow for the imprecision in aircraft positioning at the approach gate. This resulting separation can then be combined with information on the mix of aircraft classes to derive the capacity of the runway.

One modification is necessary to use this system in FLAPS. In the standard presentation of the technique the point of reference is the runway threshold. For analytical work this poses no difficulties. Directly translated into the simulation structure it would imply that aircraft are scheduled only upon arrival at the threshold. Even though this is not what in fact happens, it will not cause inaccurate results in a simulation as long as the actual position of an aircraft on final approach is not needed in the simulation. However, since it is our desire here to be able to model dependent runway operations and apply separations that depend on the distance of aircraft from the runway threshold*, we must change the point of application of the A/A separations from the threshold to the beginning

* One example of this would arise when deciding to release

of the common approach path. This will ensure that any delay of landing operations due to congestion of the runway will be applied before the calculation of the time of arrival at the threshold is performed. Thus the actual position of the aircraft on final is available to be used in the simulation. Simulations, like the ASM model, that impose A/A separations at the threshold restrict the analyst to specifying time separations for these cases.


Closing case

Figure II-13 presents the time-distance diagram for the closing case. Table II-8 defines common terms used throughout the separation discussion.


As this is a closing case, $V_a(j) > V_a(i)$, the minimum separation applies when aircraft i is at the threshold. The resulting time separation at the approach gate ($S_{aa}(i,j)$) is the time i takes to fly the final approach ($C/V_a(i)$) less the time j requires to fly down to $R_{aa}(i,j)$ miles from the threshold.

In addition, a buffer is needed due to the inevitable errors in timing the arrival of aircraft at the approach gate. It is convenient to parameterize the buffer so as to control directly the fraction of aircraft that violate the separation. If an average buffer size (B) of $Z_{aa} * B_{saa}$ is chosen with $Z_{aa} = 1.65$, for example, then 5% of the aircraft will violate the separation. In summary we impose a

departures on a runway that intersects with an arrival runway.



$$Saa(i,j) = \frac{C}{Va(i)} - \frac{C - Raa(i,j)}{Va(j)}$$



Arrival error distribution

FIGURE II-13. Arrival/Arrival Separations: Closing Case.

C	- Length of the common approach path (miles).
Va(i)	- Average speed on final approach for class i (knots).
Vas(i)	- Standard deviation of speed on final approach for class i (knots).
v(i)	- Actual speed on final approach for aircraft i (knots).
Raa(i,j)	- Minimum distance required on final approach between class i in lead and class j following (miles).
Saa(i,j)	- Minimum time separation to be imposed on class j at the approach gate when class j follows class i on final approach in order to insure that Raa is met (minutes)
Bsaa	- Standard deviation of error in arrival position at approach gate (seconds).
Zaa	- Number of standard deviations of the arrival error distribution to be used in setting buffer length (dimensionless).
B	- Average buffer length (seconds).
b	- Actual buffer length imposed on a given aircraft (seconds).
Tag(i)	- Time that aircraft i crosses the approach gate (minutes).
Tt(i)	- Time that aircraft i crosses the threshold (minutes).
Tc(i)	- Time that aircraft i clears the runway (minutes).
Nt(a,b)	- Normal distribution, truncated at plus or minus 2 standard deviations with mean of a, standard deviation of b.

A number of these terms are shown graphically in figures II-13 and 14.

Table II-8: Terms used in derivation of separations




FIGURE II-14. Arrival/Arrival Separations: Opening Case.

separation of:

$$Saa(i,j) + b \quad (14)$$

where

$$Saa(i,j) = C/Va(i) - (C-Raa(i,j))/Va(j) \quad (15)$$

and b is drawn from a normal distribution with a mean of $Zaa*Bsaa$ and a standard deviation of $Bsaa$.

Opening case

Figure II-14 presents the case where $Va(i) > Va(j)$. The same terms are used here as in the closing case. Here $Raa(i,j)$ applies when trailing aircraft j reaches the approach gate. The time separation is merely the time it takes the first landing, j , to fly $Raa(i,j)$ miles. A buffer, b , is added as before. The separation is:

$$Saa(i,j) + b$$

where

$$Saa(i,j) = Raa(i,j)/Va(i) \quad (16)$$

Note that should $Va(i) = Va(j)$, the separation applies at all points when both aircraft are on final and (15) reduces to (16).

Once the appropriate $Saa(i,j) + b$ is drawn for a given pair of aircraft the second aircraft, j , is not allowed to begin its descent from the approach gate until the separation is met. Once an aircraft is permitted to proceed, an actual approach speed, $v(i)$, is drawn using $Va(i)$ and $Vas(i)$. This actual speed is used to determine the time needed to fly from approach gate to threshold.

Numerical values

$R_{aa}(i,j)$ - Under instrument flight rules (IFR) conditions R_{aa} is 3 miles except for situations in which wake turbulence must be considered. For visual flight rules (VFR) conditions no formal standards exist but a set of values has come to be recognized as typical of actual operating conditions. The two sets are given in table II-9.

Buffer - The most often used values for the buffer size are $Z_{aa} = 1.65$ and $B_{saa} = 18$ seconds (SIN 75 p. 4-1, 6-1).

If we assume $V_a(i) = 140, 120, 110$ kts. for $i = 1$ (heavy), 2 (large), 3 (small), respectively, then the set of time separations at the outer marker (S_{aa}) that result are those shown in the lower part of table II-9. Note that the upper right elements in each matrix are opening cases, the lower left elements are closing cases.

Analysis of the procedure

A question may arise as to the validity of the separation procedure described above, since the separation formulas for $S_{aa}(i,j)$ assume that each aircraft in a class has the same approach velocity, $V_a(i)$, whereas it is known that approach speeds vary among aircraft of a given class. We can calculate the actual separation that results when the aircraft separations are determined under the assumption of fixed $V_a(i)$. We let R indicate the actual closest separation between two arrivals that results when a separation of $S_{aa}(i,j) + b$ is imposed on aircraft with

Minimum Separations (miles):

		IFR Rules			VFR Rules		
Trailing Class - -		1	2	3	1	2	3
Lead Class -	1	4.0	5.0	6.0	2.7	3.6	4.5
	2	3.0	3.0	4.0	1.9	1.9	2.7
	3	3.0	3.0	3.0	1.9	1.9	1.9

Derived Time Separations, minutes ($S_{aa}(i,j) + B$):

		IFR Rules			VFR Rules		
Trailing Class - -		1	2	3	1	2	3
Lead Class -	1	1.99	2.36	2.73	1.51	1.84	2.13
	2	1.99	1.80	2.24	1.53	1.33	1.67
	3	2.10	1.92	1.86	1.69	1.44	1.36

Class 1 is heavy jets, class 2 large jets, class 3 small jets
 Minimum Separation data from HAI 78 p. 3-3,
 also see WEI 80 pp. 36-3.

Table II-9: Arrival/Arrival Separations: Typical Values

actual approach speeds $v(i)$ and $v(j)$. We examine only the more complicated closing case.

The desired minimum distance separation, $R_{aa}(i,j)$, is used to derive a time separation, $S_{aa}(i,j)$, at the outer marker. The same formula (13) can be used in "reverse": given an $S_{aa}(i,j)$, what separation R in fact results when the aircraft fly the final approach not at $V_a(i)$ and $V_a(j)$ but at $v(i)$ and $v(j)$? The actual separation is given by:

$$Saa(i,j) + b = C/v(i) - (C-R)/v(j)$$

Solving for R:

$$R = C + v(j)*(Saa(i,j) + b - C/v(i)) \quad (17)$$

Inserting Saa(i,j) from (15):

$$R = C + v(j)*[C/Va(i) - (C-Raa(i,j))/Va(j) + b - C/v(i)] \quad (18)$$

This describes the actual separation.

Taking expected values:

$$E(R) = C + C/Va(i)*E(v(j)) - (C-Raa(i,j))/Va(j)*E(v(j)) + E(b*v(j)) - c*E(v(j)/v(i))$$

Simplifying:

$$E(R) = Raa(i,j) + C*Va(j)/Va(i) + E(b*v(j)) - c*E(v(j)/v(i))$$

As b and v(j) are uncorrelated

$$E(b*v(j)) = E(b)*E(v(j))$$

$$E(b*v(j)) = B*Va(j) = Zaa*Bsaa*Va(j)$$

The remaining term $E(v(j)/v(i))$ cannot be evaluated exactly, unless the probability distributions of the v(i) for all classes of aircraft are known.

As v(i) and v(j) are uncorrelated it is known that:

$$E(v(j)/v(i)) = E(v(j))*E(1/v(i))$$

Therefore:

$$E(v(j)/v(i)) = Va(j)*E(1/v(i))$$

The expected value of $1/v(i)$ is not known. However, in the cases of interest v(i) is on the order of 120kts. with a standard deviation of 4kts. Because the mean of v(i) is large relative to the spread of v(i), we assume that

$E(1/v(i))$ can be adequately approximated by $1/Va(i)$, and therefore

$$R = . Raa(i,j) + Va(j)*Zaa*Bsaa$$

proving that the average error due to the assumption of fixed arrival speeds is insignificant.

Note that the second term above, $Va(j)*Zaa*Bsaa$, represents the contribution of the buffer to increasing the average separation. Using the typical values, as indicated above ($Va(j)=120$ kts., $Zaa=1.65$, $Bsaa=18$ seconds), this term has a value in practice of approximately 1 nautical mile.

4.d Departure/Departure separations

The procedure for D/D separations is much simpler than for the A/A case. Separations are directly specified by FAA regulations as minimum time between departures, so no conversion is necessary (FAA 79 para. 340, WEI 80 p.38). Additionally the trailing departure remains stationary at the end of the takeoff runway until clearance is given, so any error in beginning the roll is usually neglected. Thus, no buffer is needed. As with A/A separation there is a standard minimum, supplemented to avoid wake vortex problems behind heavy aircraft. Typical D/D separation values are given in table II-10.

The separation between departures is also a function of the route aircraft use. Shorter D/D separations are

Trailing Class - -	IFR Rules			VFR Rules			
	1	2	3	1	2	3	
Leading class -	1	90	120	120	90	120	120
	2	60	60	60	60	60	60
	3	60	60	60	50	45	35

Class 1 is heavy jets.

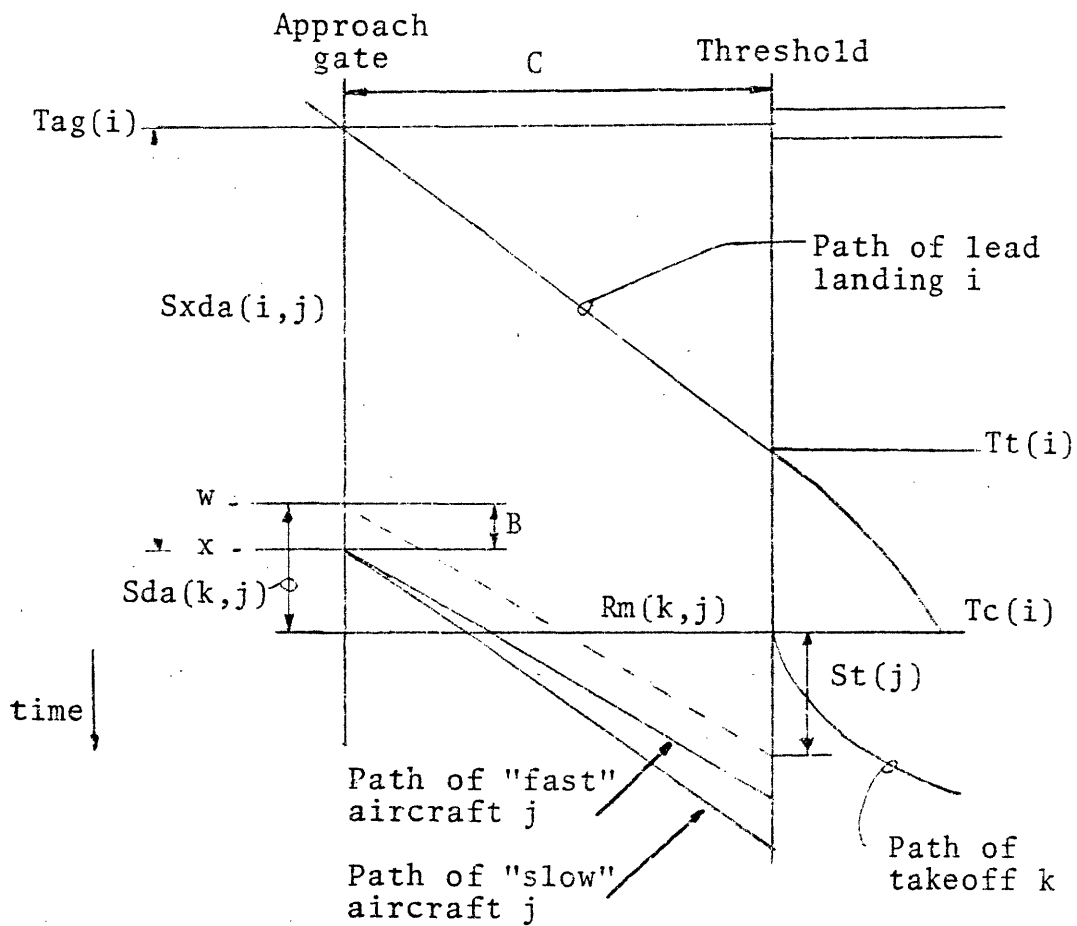
Data from HAI 78 p. 3-5 (values in seconds)

Table II-10: Departure/Departure Separations:
Typical Values

permitted between aircraft on divergent routes (ATC p. 95). FLAPS therefore allows the submittal of two sets of D/D separations: one for use when successive departures fly on the same course (as indicated by having the same departure fix number) and one when successive departures have diverging courses (different fix numbers).

4.e Mixed Operations

Mixed operation separations are by far the most complicated of single runway separations. There are two problems to be considered: 1) The rules governing the release of one or more departures between landings. 2) The spacing of arrivals to permit at least one departure to leave between successive landings. Figure II-15 is the basic time/distance diagram for this case. Several terms



In general $Tag(i) + Sxda(i,j)$ need not occur at time x as shown here.

FIGURE II-15. Separations for Mixed Operations on a Single runway.

used in the derivations of mixed operation separations are

- $R_m(k, j)$ - The closest distance to the threshold an arrival of class j may be when a departure of class k begins a takeoff roll (miles). Used to determine if a takeoff may roll and to set $S_{da}(k, j)$.
- $S_{da}(k, j)$ - The maximum interval allowed between the beginning of final approach of arrival class j and the beginning of a takeoff roll of class k . Derived from $R_m(k, j)$. Applied only when scheduling arrivals when the runway is operated in alternating mode using calculated D/A separations.
- $S_{xda}(i, j)$ - The minimum time separation imposed on class j aircraft at the approach gate when class j aircraft follows a class i aircraft. Applied only when the runway is operated in alternating mode using specified separations.
- Z_{x1a} - Number of standard deviation of the arrival error distribution used to set the length of buffer associated with $S_{xda}(i, j)$
- B_{sxda} - Standard deviation of error in arrival position at approach gate associated with $S_{xda}(i, j)$ (seconds).
- $T_c(i)$ - Time that landing aircraft i clears the runway (clock time).

Table II-11: Terms Used for Mixed Operation Separations

defined in table II-11.

Release of departure

Any time a runway is being used for both arrivals and departures, the takeoffs on the runway must not interfere with any arrivals that have already begun their descent. Thus, this type of separation must be met regardless of the operating mode for this runway (i.e. arrival priority, alternating, departure only).

The basic rule governing takeoffs in this situation is (HOR 75 p. 145-7) that a takeoff cannot begin to roll unless 1) there are no aircraft on the runway, and 2) the next arrival is farther from the threshold than some critical distance. We refer to this minimum critical distance as $R_m(k,j)$, where j is the class of the next aircraft to land and k refers to the class of the next takeoff. $R_m(k,j)$ is typically 2 nautical miles for all k and j , however we allow the case where $R_m(k,j)$ varies across aircraft classes.

Spacing of arrivals

The normal A/A separation often allows sufficient time between successive landings that one, and sometimes more than one, departure can takeoff between the landings. However, this is not the case for all A/A pairs and is a particular problem under VFR conditions where A/A separations are shorter than under IFR rules. To rectify the imbalance in capacity between landings and takeoffs and attendant accumulation of departures queuing for takeoff, runways are often operated to allow at least one departure between every pair of landings (WEI 80 p.61-7). This is accomplished by applying an additional separation between arrivals. Some pairs of landings will not be affected by this, as the A/A separation is already adequate for a departure to roll.

In the standard treatment of this problem (used, for example in the ASM model), arrivals are scheduled when they

reach the threshold. Thus, in order to calculate the earliest time that the arrival could cross the threshold, a minimum time separation may be directly applied from the time the previous departure began its roll. This time is marked as $St(k,j)$ in the figure for mixed operations. The set of $St(k,j)$ is generally referred to as the D/A separations. This separation would probably have been originally stated as a distance, $Rn(k,j)$, and converted to $St(k,j)$ by the formula:

$$St(k,j) = Rn(k,j)/Va(j)$$

However, FLAPS does not schedule arrivals at the threshold but at the approach gate - closer to where they are in fact sequenced. One aspect of this change is that now aircraft are no longer scheduled in the sequence they will use the runway. An arrival is scheduled and has separations applied to it before the takeoff is scheduled - even though the takeoff will use the runway before the arrival. The mixed operations figure (II-15) shows this. This means that any separation intended to allow alternation of landings with takeoffs cannot be applied when the arrival reaches the threshold. Rather we must work backwards from $St(k,j)$ or from $Rn(k,j)$ to find a separation that can be applied at the approach gate. We continue to term this derived separation a D/A separation, at the risk of some confusion with conventional nomenclature, because its use is the same as the standard D/A separation: to separate an

arrival from a previous takeoff when an alternating operations mode is being used.

Obviously, no more spacing should be added than is necessary. Any additional spacing beyond the minimum requirement is wasted time. Therefore, it would seem that the model should calculate the D/A separation based on $R_m(k,j)$, so that when the first landing of the pair clears the runway, the second landing will just be $R_m(k,j)$ miles out and the intervening takeoff can begin its roll. However, runways are not always operated in such an efficient manner. Often the "gap stretching" procedure is a blanket increase of the shorter A/A separations from 3 miles to some longer distance. Therefore FLAPS allows the D/A separations to be specified in either of two ways: 1) Specified D/A - the D/A separation is given directly and is applied in the same manner as the A/A separation. 2) Calculated D/A - the model calculates the minimum time internally based on $R_m(k,j)$. The specified D/A separation is given as $S_{xda}(i,j)$ in the figure. It may have its own buffer, b , to be drawn from the truncated normal distribution $Nt(Z_{xda}*B_{xda}, B_{xda})$, where Z_{xda} and B_{xda} are analogous to Z_{aa} and B_{aa} . It may appear contradictory to label a separation applied between two arrivals a "D/A separation". We label it $S_{dax}(i,j)$ because of its function: to delay an arrival a sufficient time behind the previous landing to allow a departure to leave between landings.

The calculated separation is represented by $S_{da}(k,j)$ in the figure. Note that neither the subscripts nor the point of reference of S_{da} are the same as for S_{xda} . The best way to understand the $S_{da}(k,j)$ may be to state the resulting rule for scheduling of arrivals and then explain how it is derived.

Rule: Under alternating operations using calculated D/A separations, landing aircraft (j) may begin a descent only after $T_c(i) - S_{da}(k,j) + b$, where $T_c(i)$ is the time the previous landing aircraft (i) clears the runway and b is the buffer associated with this case. In the mixed operations figure (II-15) the small x denotes this time.

The rule is derived in the following way. The earliest time that a takeoff may roll is the time the previous landing clears the runway, $T_c(i)$. If k is to depart, the subsequent landing, j , must be at least $R_m(k,j)$ miles out on final approach. If the length of the final approach is C miles, this implies that landing j cannot have flown more than $C - R_m(k,j)$ miles on final at $T_c(i)$. The time that aircraft j uses to fly this distance is $(C - R_m(k,j))/v(j)$. Thus, it would seem that aircraft j cannot be allowed to start down the final approach within $(C - R_m(k,j))/v(j)$ of $T_c(i)$.

However, this time cannot be used as just stated. At the time the controller is positioning the arrival of aircraft j at the approach gate, the controller does not yet

know the exact speed that the aircraft will fly on final. The controller only knows the typical speed of the aircraft. That is, $v(j)$ is unknown but $V_a(j)$ and $V_{as}(j)$ are known. Therefore it would be erroneous to project ahead in the model and take advantage of the exact approach speed to schedule the landing precisely. A more valid procedure would be to schedule aircraft j according to a worst case rule - allow j to go only when it cannot violate $R_m(k,j)$ regardless of $v(j)$. Figure II-15 shows the "cone" of possible flight paths for an aircraft that leaves at time w . From this it can be seen that, if we chose w so that the fastest possible j arrives at the point $R_m(k,j)$ miles out from the runway threshold just at $T_c(i)$, then slower aircraft that began their descent at w will be even farther away from the threshold at $T_c(i)$. Therefore, the time $T_c(i) - w = S_{da}(k,j)$ should be

$$S_{da}(k,j) = (C - R_m(k,j)) / (V_a(j) + Z_t * V_{as}(j)) \quad (19)$$

where Z_t is a tolerance value indicating how safe the separation should be. Table II-12 summarizes the scheduling rules for the mixed operations case.

Aircraft being scheduled	Mode of operation	Rule - commit if, and only if the following are met:
k - takeoff	(all)	1) D/D separation met 2) No aircraft on runway 3) Next landing is at least $R_m(k,j)$ miles from threshold
j - landing	Arrival Priority	1) A/A separations met
	Alternating Priority calculated separation	1) A/A separations met 2) Time is after $T_c(i) - S_{da}(k,j) + b$ where S_{da} is found from (19)
	Alternating Priority specified separation	1) A/A separations met 2) time since last commitment of a landing is greater than $S_{xda}(i,j)$

Table II-12: Mixed Operations: Summary of Rules

E DEPENDENT RUNWAY OPERATIONS

1 Introduction

This section discusses how aircraft separation movements on dependent runways are modeled. The procedures that are used are similar to those discussed above in section D for the single runway case, but there are a number of complications when the two operations being separated are on different runways. The same basic types of conflicts occur - A/A, A/D, D/A, D/D - but the values of the separations and the point of application vary. Where parallels to the single runway analysis exist, the procedure for dependent runways is not repeated, rather the single runway procedure is referenced.

There are several ways that two runways may depend on each other. They may 1) be physically intersecting, 2) be parallel or 3) intersect along the projection of the runway centerlines. We will discuss the first two cases in detail as they represent the bulk of airport conditions. The third case will be discussed only briefly. We will show that a number of configurations can be reduced to one of the first two cases, but we will not attempt to show that all possible configurations can be so modeled.

Table II-13 gives nomenclature used in this section. Some items are repeated from the table giving single runway nomenclature (Table II-8).

C	- Length of common approach path of the runway used for arrivals (miles). Parallel runways used for arrivals are assumed to have common approach paths of equal length.
Raat(i,j)	- Minimum distance separation imposed between a landing of class i and a following landing of class j on a parallel runway (miles).
Saat(i,j)	- Minimum time separation to be imposed on class j aircraft at the approach gate when a class j follows class i on final approach to a different runway (minutes).
Sadt(k,j)	- Maximum interval prior to a takeoff of class k clearing the intersection with the arrival runway that the next arrival of class j is allowed to begin final approach (minutes). May be specified directly or derived from Rmt(k,j).
Tag(i)	- Time when aircraft i crosses the approach gate (minutes).
Ti(i,l)	- Minimum of the time when aircraft i clears the runway or crosses the intersection with runway l (minutes). If aircraft i is a takeoff, Ti is always the intersection crossing time.
Tc(i)	- Time when aircraft i clears the runway (minutes).
Rmt(k,j)	- The closest to the runway threshold an arrival of class j can be when a takeoff of class k clears the intersection or takes off on a different runway (miles).
Zt	- A measure of tolerance for the "worst case" situation (number of standard deviations).
Va(i)	- Average approach speed, class i aircraft (knots).
Vas(i)	- Standard deviation of approach speed, class i aircraft (knots).
Tto(i)	- Average takeoff roll time to 6000 feet for a takeoff of class i (seconds).
Ttos(i)	- Standard deviation of takeoff roll time to 6000 feet for an aircraft of class i (seconds).
Xi(l,m)	- The distance from the threshold of runway l to the intersection of runway m (feet).

Table II-13: Dependent Runway Nomenclature


Each type of operation and geometry is discussed independently. It must be remembered that, as in the single runway case, a number of separation types may be applicable to the given situation and they must all be met before an operation may be committed.

2 Intersecting Runways

Almost every major airport in the United States operates some or most of the time with intersecting runway movements. Boston Logan Airport (see Chapter IV) and New York La Guardia Airport (see chapter III.B) are two examples.

2.a Arrival/Arrival separations

As far as can be determined, airports are not currently operated with arrivals on two intersecting runways. This could only be done if controllers could project the position of aircraft accurately enough to ensure separation at the intersection. Evidently this ability does not yet exist. Although counterexamples to the statement above could be cited, it turns out that, on closer look, our statement is still true. Consider, for instance, Dulles airport in Washington, figure II-16. Although it might appear initially that two runways are operated there simultaneously for landings, this is not really true as landings are not



The taxiway system has been simplified

FIGURE II-16. Dulles International Airport, Washington D.C.

assigned to the secondary runway unless the pilot can assure the controller that the aircraft will stop short of the primary runway. This is a "hold short" situation and will be analyzed in part 6 below. As explained there, no real separation is imposed between the two operations. Given that landings on the secondary runway hold short, these are de facto two independent runways. Accurate modeling of this case would depend on correctly estimating the percentage of aircraft (primarily general aviation at Dulles) assigned to the secondary runway. The same situation occurs at Boston Logan airport as at Dulles and is modeled in chapter IV below.

Since A/A separations will be needed for the parallel runway case, it is simple to extend the same capability to the intersecting runways case. This is accomplished by imposing a time separation using the same procedure described below for the parallel runway case. It may be that arrivals on intersecting runways could be studied in this way with a proper selection of separations, but since no data exist for this case, it is not possible to confirm this.

2.b Departure/Departure separations

The basic rule for the intersecting case (FAA 79 para. 1111, 1121) is that aircraft may not begin a departure if there is a takeoff rolling on the other runway but not yet

across the intersection. Effectively, the area marked "a" in figure II-17 is viewed as having a single aircraft occupancy rule. As the two runways point in different directions there is no need for any in-trail separation of the kind used when both operations were using the same runway.

2.c Mixed operations

Separation of landings and takeoffs from each other when they are on different runways involves the same two problems as the single runway case: release of a departure and gap stretching for arrivals to allow departures. Again these separations are considered independently of any other separations needed such as is required, for example, if the landing runway is also used for takeoffs.

Release of departures

Figure II-18 gives the time-distance diagram for this case. Aircraft in part a of the figure is the first to land. When this aircraft clears the intersection of the two runways at time $T_i(i,m)$, a decision is made whether or not to release a takeoff. Takeoff k is able to go if the next landing, j , is at least $R_{mt}(k,j)$ miles out on final and will be so when the takeoff k crosses the intersection. Note that should landing i exit the runway prior to crossing the intersection, the departure may roll as soon as the landing clears the runway. For simplicity we use $T_i(i,m)$ as the






FIGURE II-17. Intersecting Runway Departure Operations.



a) Time - Distance Diagram



b) Position of Aircraft at Time = $Ti(i,1)$

For clarity, buffer is omitted

FIGURE II-18. Mixed Operations on Intersecting Runways.

minimum of the time across the intersection or the time to clear the runway of aircraft i . Figure II-18b shows the position of all active aircraft at time $T_i(i,m)$. The rule used for releasing takeoffs means that the area marked "d" must not contain any aircraft at any time while the departure is in the area marked "e".

This procedure is complicated in practice by the fact that the future position of aircraft must be estimated in order to determine the adequacy of separation. The controller cannot know exactly how long it will take the takeoff to roll down to the intersection. The controller must estimate the distance from the threshold of the landing when the takeoff reaches the intersection. As before, we assume a worst case procedure is used.

If, for instance, the departure has a takeoff runway occupancy time of 34 seconds (5 seconds standard deviation) to reach 6000 feet then, using equation (13), its time to reach an intersection 3000 feet down will be 24 seconds with a standard deviation of 3.5 seconds. The uncertainty in the position of the arrival is much less than for the takeoff, even if it is assumed that the only information the controller has is $V_a(i)$, $V_{as}(i)$ and the present position. A 120 knot arrival with a standard deviation of 4 kts. currently 3 miles out will reach a point 2 miles out in 26 seconds with a standard deviation of 1 second. The uncertainty which exists in the landing's time to travel a

given distance is probably less than the controller is able to estimate but the major uncertainty is still with the takeoff. Thus we neglect any effect of the variation in arrival speeds and assume that the position of the arrival 20 to 30 seconds ahead can be estimated accurately enough to decide if the separation is met. We then compensate for this by using a more generous estimate for the value of Z_t , the "worst case" parameter, when computing how long the takeoff will need to get to the intersection. A more detailed study of this case should also consider how precisely controllers judge the 2 mile cutoff (R_{mt}). In other words, will they prohibit a takeoff if the arrival is projected to be 1.98 miles out when the takeoff crosses the intersection? The procedure for releasing a departure is then merely to check the location of arrivals and insure that the separation is met and will be met before allowing a takeoff to roll. The rule is: do not release a departure on runway l unless the next landing, j , on runway m will be at least $R_{mt}(k,j)$ miles out at a time

$$T = (T_{tc}(k) + Z_t * Z_{tos}(k)) * \text{SQRT}(X_i(m,l)/6000) \quad (20)$$

from the present, where Z_t should probably be 1.65 or more.

Providing space between arrivals to allow alternating takeoffs and landings is accomplished by the same general procedure as with the single runway case. We permit both the calculated and the specified gap stretching procedure. The specified separation is, as in the single runway case,

merely a second time separation imposed between arrivals. The maximum of the two separations is used to control arrival/arrival spacing.

The calculated separation for intersecting runways is more complicated than for single runways because the point of reference is changed. Using the notation introduced in the last section, aircraft j on runway m must be at least $R_{mt}(k,j)$ miles out on final at $T_i(k,m)$. This means that it cannot have flown more than $C - R_{mt}(k,j)$ miles by $T_i(k,m)$. However, $T_i(k,m)$ will be in the future at the time aircraft j is being scheduled. Thus, we use the same "worst case" formula to estimate $T_i(k,m)$ as we did in studying the procedures for the release of takeoffs.

The time that aircraft j will take to fly $C - R_{mt}(k,j)$ miles was given in the single runway case. In the notation for the intersecting case, this is now:

$$S_{dat}(k,j) = (C - R_{mt}(k,j)) / (V_a(j) + Z_t * V_{as}(j)) \quad (21)$$

The complete rule is therefore: Under alternating operation mode (calculated separation) with intersecting runways do not release a landing j until the time:

$$T_c(i) + T - S_{dat}(k,j) + b \quad (22)$$

when i is the previous landing on this runway and where T is given by (20), $S_{dat}(k,j)$ by (21) and b is the buffer associated with this case.

3 Parallel Runways

The air traffic control rules for parallel runways are very complicated and not easily codified. The air traffic control handbook (FAA 79) does not discuss rules for parallels in any one place or in a unified framework but presents the rules in several contexts. Two summaries of rules available to the author are neither complete nor entirely in agreement (HOR 75 p.96-7, KAY 79). The appropriate separation can be affected by not merely the distance separating the two runways and the weather conditions (IFR or VFR) but by the amount of offset of the two thresholds and the degree of divergence in the missed approach paths of the two runways (FAA 79 para. 744). Further complications are induced by the fact that local controllers use "marginal weather" rules for situations between IFR and VFR that are not official rules, let alone published or even consistent from airport to airport. We will briefly outline the range of possible dependencies and then consider in detail the rules necessary to support the analysis of all of the possibilities. The primary variable affecting the degree of dependence is distance separating the two runways:

4300 feet or more

Runways that are separated by this distance or more are completely independent under all operating conditions (FAA 78 para. 1103c).

2500 to 4300 feet

In this situation the runways are independent under VFR conditions. Under IFR arrival/arrival separations are imposed. Departures are independent from operations on the other runway (SIN 79 p.4-6).

700 to 2500 feet

Under VFR conditions the runways are dependent to the extent that an arrival on one runway must be over the threshold before a departure may be released on the other runway. Under IFR conditions the runways are essentially a single runway in terms of separations: operations on one runway must clear operations on the other runway by the same amount they would have to clear the operation if it were on the same runway.

Less than 700 feet

The two runways are essentially a single runway under all conditions.

There are additional breakpoints at various distances where the type of dependence does not change but the value of the separation imposed does (MAY 79 pp. 4-7). In view of the number of ways that two parallel runways can be dependent and of the fact that the type of separation (not merely the length of the separation) can vary, it would be very difficult to construct a model that, on the one hand automatically imposed the correct set of separations while, on the other hand, was flexible enough to permit

experimentation. Instead we will establish a procedure that can model the possibilities mentioned above.

3.a Arrival/Arrival separation

The only two possibilities mentioned in the available literature for A/A separations on parallel runways are 1) no dependency and 2) a minimum time between arrivals on adjacent runways. Thus, this case can be modeled by having the user designate whether two parallel runways have an A/A separation imposed or not. If they do, then we impose an A/A separation that is calculated and applied in the same manner as the A/A separation used when the two arrivals were on the same runway. May indicates that the value of this separation will vary from the standard values given above in table II-9 to a simple 2 mile separation between all pairs of arrivals depending on distance (MAY 79 p.6). Note that the "single runway" case can be modeled by imposing the identical A/A separation on both 1) arrivals on the same runway and 2) arrivals on different parallel runways. If the separation is specified in units of distance ($R_{aat}(i,j)$), then this procedure assumes that the length of the common approach paths (C) are equal. If the separation is specified as a time ($S_{aat}(i,j)$), then the length of the approach paths may be unequal.

3.b Departure/Departure separations

Departures on parallel runways are either 1) independent or 2) have a time separation between successive operations. The only example of the latter (MAY 79 p.5) used the same values of D/D separations as given in table II-10 for single runway separations. As with A/A separations these cases are modeled by providing a flag to indicate if a D/D separation is to be applied, and, if so, it is imposed in the manner of the single runway case.

3.c Mixed operations

Mixed operations on dependent parallels are the most complicated situation, as there are three types of dependencies that can exist. In addition, there are two types of priorities to be modeled (arrival priority and alternating operations). The possible cases of dependency are 1) no dependency, 2) single runway case and 3) touchdown by the arrival necessary before release of the takeoff allowed.

The first situation, of course, poses no problem. The second case is functionally identical to the single runway case and the analysis and formulas used there apply directly. The third case requires an additional rule but is actually simple to analyze. This case is limited to conditions when the controller is able to see that the arrival has touched down in order to permit the release of

the departure. This occurs primarily with VFR rules, but it can happen under IFR rules with a low ceiling but good visibility at ground level. Under existing rules, all VFR A/A separations are 1.9 miles or more and all IFR A/A separations are 3 miles or more. When the buffer is added to these values, it means that the second landing of a pair will be at least 2 miles out on final when the first landing touches down. Therefore, under existing rules, it would be redundant to check the position of the next arrival. Despite this fact, the mechanism for checking is included in the model to permit the analysis of closer A/A separations should they be used in the future.

The procedure used for alternating operations on parallel runways is exactly the same as for the single runway case. Additional separation is provided either by applying a user specified separation or by calculating additional spacing needs, as in the single runway case.


Should the arrival/departure dependency be of the third form mentioned above (departure may leave when the next arrival is over the threshold) then, as explained, under current rules adequate spacing already exists to alternate operations. Were shorter A/A separations to be used, it would be straightforward to provide for alternating operations. The analyst can use the specified D/A separation and set the values for all closing cases equal to R_{mt} used in this case. This will insure that arrivals are at least R_{mt} miles apart at the threshold.

4 Projected Intersecting Runways


A number of airports have runways that are not parallel but do not intersect physically. Wichita, Ks., Kansas City, Mo., and Chicago O'Hare are three examples. A number of these situations can be modeled by the procedures discussed in the two sections above. When the two runways diverge slightly (see figure II-19a, Wichita) they are independent under certain operating conditions: 1) arrivals on one runway, departures on the second; 2) mixed operations on one and departures on the second. If the runways are operated in the converging direction they may be modeled as if they did intersect. This may be valid as a departure on one, for example, may be held until the departure on the other runway flew past the projected intersection point. This procedure would be more likely to be true when the projected intersection point is close to the real ends of the runway, as it is at Kansas City (see part b of the figure) than when the intersection is far from the runway end.

In other situations it may be more appropriate to regard the two runways as if they were parallel. An arrival on one may have to have cleared the runway before a departure could be released on the second.


When the projected intersection is beyond the end of one runway but on the other runway (see O'Hare example in part c of the figure, as well as the Dulles geometry



a) Wichita, Kansas
Airport



b) Kansas City, Missouri
Airport



c) Chicago O'Hare Airport

FIGURE II-19. Airports with Projected Intersecting Runways.

referred to above), then a "hold short" situation may be in effect and the two runways are independent.

It may not be possible to model all projected intersecting runways in the manner outlined above. However, it is clear from this discussion that in the overwhelming majority of cases with projected intersecting runways, the intersecting or parallel models discussed above can be brought to bear by making the appropriate adjustments in the model's parameters.

5 Assignment of Aircraft to Runways

5.a Discussion

Whenever more than one runway is in use for either landings or takeoffs, aircraft must be assigned to one of the active runways. In the demand section (II.C) reference was made to runway assignment as being made dynamically and the FAA model was criticized for prespecifying runway assignments. Consideration of airport operations shows that there are several factors that can influence runway assignments. Therefore, FLAPS provides several modes of runway assignment from which one of the modes may be selected as appropriate for a given case. After these modes are identified, several problems in modeling runway assignment will be discussed.

Direction of approach or departure

When runways are available for operation in different directions, then the choice of runway may be determined by the direction from which the aircraft approaches the airport (for landings) or the direction in which the aircraft wishes to depart (for takeoffs). Chicago O'Hare is often operated as a North airport and a South airport, each with landing and departure runways. Traffic is routed on to one or the other half of O'Hare in order to avoid intersections with traffic coming from the opposite direction.

Arrival and departure fixes provide the directional information in FLAPS. We model this type of assignment by providing two two-dimensional matrices, one for landings and one for takeoffs. Each matrix provides a set of multinomial probabilities giving the probability of using a runway, conditional on the relevant type of fix (arrival fix for landings, departure fix for takeoffs). When an aircraft becomes eligible for a runway operation, the appropriate set of multinomial probabilities is sampled to obtain the runway to be used.

As an example of how this procedure would be used, consider the following example. Suppose there are three arrival runways A, B and C and two approach fixes 1 and 2. All arrivals from fix 1 will use runway A. Half of fix 2 arrivals use A, 20% B and 30% C. In order to affect this in

FLAPS the following matrix, $X(f,r)$ would be specified:

		r - runway		
		A	B	C
f - fix	1	1.0	.0	.0
	2	.5	.2	.3

Then each aircraft that arrived would first have an arrival fix, f , assigned to it. Then row f of the above matrix would be used to set the odds in a multinomial draw which will randomly determine runway assignment.

Apron area

As the direction out of the airport from the runways can influence runway selection, so too can the destination within the airport. Pilots may choose between runways based on their intended apron to avoid excessive taxiing. This is a less significant factor than the others mentioned here and was not needed for any of the validation or application runs. Therefore this option was not modeled. If necessary, it could be added by providing a matrix of probabilities of runway assignment as a function of apron area for landings and takeoffs.

Aircraft requirements

Aircraft vary greatly in their performance and runways vary in their length. Not all of the largest aircraft can use all runways that might be available. To model this FLAPS allows a runway assignment mode based on the class of

the aircraft. This operates in an analogous fashion to the fix-based assignment above. For each aircraft class a set of probabilities giving runway to be used as a function of aircraft class is given.

Noise abatement

In order to reduce the impact of noise on residents under the flightpaths, aircraft may be directed away from "noise-sensitive" runways. An attempt may also be made to "balance" operations to expose residential areas more equitably to noise. This could be modeled using the aircraft class mode. Noisy jets, for instance, might be assigned entirely to a runway which avoids extensive exposure of people to noise. All classes may be given assignment percentages to balance the number of aircraft going to any one runway.

A second assignment mode is provided by FLAPS and may also be useful here. This is a simple multinomial assignment. If three runways are eligible for a particular operation, then a draw from a set of three multinomial probabilities can be made to determine the runway to be used.

Operational considerations

If airports are to utilize their limited runway capacity in an efficient manner, then it would seem that traffic should be assigned to runways as to fully utilize them. In particular, if one runway has aircraft waiting to

move and another runway is vacant, then any new aircraft should be sent to the empty runway. To model this procedure FLAPS has a runway assignment mode that directs each aircraft to the runway with the shortest queue waiting to use the runway.

The shortest queue assignment procedure is likely to be combined with some limitations on the ability of aircraft to use runways. Assignment may be to the runway with the shortest queue of any of the runways which a particular aircraft class may use. Therefore, the shortest queue assignment is implemented in FLAPS by specifying for each class of aircraft a set of flags indicating which runways that class of aircraft may use. Then, when the aircraft becomes eligible for a runway operation, the model scans the set of runways that are both 1) active and 2) usable for this class of aircraft. The aircraft is assigned to the runway with the smallest queue of aircraft awaiting service.

In summary, FLAPS allows four modes of runway assignment. Landings and takeoffs are considered separately and may use different modes for assignment. The modes allowed are:

- 1) Multinomial assignment
- 2) Multinomial as a function of aircraft class
- 3) Multinomial as a function of fix
- 4) Shortest queue of acceptable runways where runway acceptability is a function of aircraft class.

These capabilities represent an advancement in the modeling of airport operations over the prespecified assignment method used by the ASM model. They permit a wider range of real-world situations to be modeled. Mode 4, in particular, makes it possible for the airport planner to evaluate and investigate alternative runway assignment policies with respect to their impacts on airport delays.

5.b Sequencing of operations

As discussed in the previous parts of this section, once aircraft are assigned to a runway they do not move onto the runway until all relevant air traffic control separations are met. If these separations involve aircraft on runways other than the runway of the aircraft being scheduled (i.e., when two or more runways are dependent), then a problem in properly sequencing or alternating these operations can arise.

We illustrate the potential difficulty through the following example. Suppose arrivals are being fed to dependent parallel runways. An arrival of class 3 (small jet) has begun landing on runway 1. A class 1 (heavy jet) arrives and wishes to land on runway 2. Because class 1 is the fastest class of aircraft, all separations with it as a trailing aircraft are imposed at the threshold. Thus its needed approach gate separation is relatively large. If another, non-class 1 aircraft arrives on the first runway,

its required A/A separation is smaller than that of the class 1 aircraft, so it becomes eligible to land and begins its operation. Now the class 1 aircraft must be separated from this new landing on runway 1 and its landing is deferred again. In fact, should there be a continuous stream of arrivals on runway 1, the aircraft on runway 2 needing a long separation will not be allowed to land until an arrival appears on runway 1 which requires an even longer A/A separation.

This is not a realistic pattern of use. Runways when busy are much more often operated in an explicitly alternating mode. In order to force this to occur a mechanism is included in FLAPS which does not allow two successive landings to begin on a given runway unless no other runway has aircraft waiting to land.

The same difficulty logically exists for two dependent takeoff runways but is not often a problem in practice and so no specific mechanism for alternating takeoffs is included. The reason it is not a problem is that seldom are there two dependent runways used only for takeoffs. Usually one or both of the runways will have landings as well. The presence of landings and the attendant separation requirements (A/D) act to prevent the lockup described above.

6 Special Problems

This section considers three special situations that can arise and how they are modeled within the framework that has been set up for runway operations. They are covered here because they involve both aircraft performance and air traffic control procedures.

6.a Displaced runway threshold

A runway may be operated with a displaced threshold, that is, the aiming point for landings is moved in from the physical end of the runway. See figure II-20 for a description. This technique may be used to raise the flight path of landings for noise abatement. The full length of the runway is available for takeoff rolls from that end of the runway. Occasionally the threshold for takeoffs is moved in as well. Operations in the other direction may also use the full length out to the physical end of the runway. Arrivals may exit beyond the far threshold (if an exit is located there). Departures may "use" the full length of the runway in the sense that the entire length is used in the calculation of whether the runway is long enough to be safe for this kind of aircraft. Logan airport (see chapter IV) has several runways with displaced thresholds, including some with displaced thresholds on both ends of the runway.

Direction of operation:

Direction 1 \longrightarrow

Direction 2 \longleftarrow


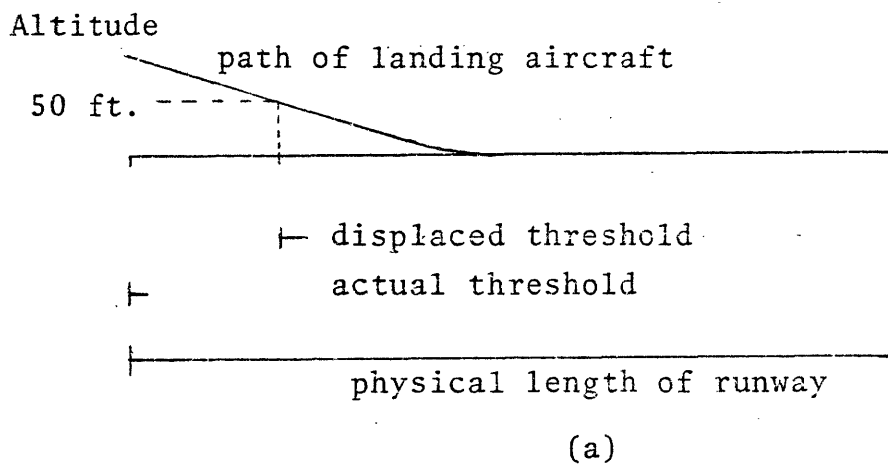


FIGURE II-20. Displaced Runway Thresholds.

This feature can be modeled by the procedures described in this section. If the threshold for takeoffs has also been moved in (to Y), then the runway can be modeled as if it actually began at Y rather than X by placing the runway end keypoint at Y. When the runway is modeled in the other direction (direction 2 in the figure) and if it is necessary to use the full distance for landings, an exit may be located in the area beyond the end of the runway.


If the landing and takeoff thresholds are in different places, then the analyst can model the situation by viewing the one real runway as two runways, one for landing and one for takeoffs. The runway is described to the model as two runway modules. Each module is on the same centerline but shifted (part b of the figure) so that one module ends where takeoffs begin and the other ends where the threshold is. The two are operated as dependent parallels. The separations imposed are those described in part 3 (parallel runways) so as to make the two modules a single runway operationally. This procedure will result in correct behavior. Before a departure can roll on the takeoff runway module any arrival must have cleared or be at least $R_{mt}(k,j)$ miles away from the threshold. If alternating operations are in effect, the procedures described in part 4 will allow space for departures. By describing the one runway as two modules, the performance of aircraft will be correctly modeled. Landings will touch down in the correct place on

the runway and roll across intersections at the appropriate time. The shifted takeoff roll point will ensure that takeoffs will be modeled as taking the correct amount of time to cross various intersections. When operating the runway in the other direction both landings and takeoffs should be directed onto whichever of the two runways is placed in the right position.

Should it be the case that both ends of the runway have displaced thresholds, the runway may still be modeled by the split runway procedure. The opposite ends of the two runway modules should be positioned to have one runway module end where takeoffs begin in the second direction and the other runway module end where landings touch down in the second direction.

6.b Hold short arrivals

A hold short arrival procedure is a situation where arrivals are directed onto a runway under the restriction that they be able to "hold short" of some point, usually the intersection with another active runway. See figure II-21 for an example. This is done to increase capacity by opening up another runway. This procedure is used at Dulles (discussed above) and at Logan (see chapter IV). This situation is trivial to model as the result is two independent runways for arrivals. As an arrival on one runway is assumed to stop short of the intersection, the



Landings on 1 exit at A or before.

FIGURE II-21. Hold Short Arrivals.

operation on the other runway proceeds independently of the first arrival. Correct modeling involves assigning arrivals to the proper runway in a correct manner.

The runways will not be independent for takeoff operations. Using the example of the figure, departures on 2 must clear departures on 1, and departures on 1 must clear arrivals and departures on 2. These particular separations are the same as the ordinary dependent runway separations discussed above.

This hold short procedure is sometimes used on a conditional basis. In other words, not all the arrivals will be able to hold short. This is the case in the Logan example below. Small aircraft hold short, but large jets need the full length of the runway and thus interfere with operations on the other runway. In the figure example heavies on 1 may not be able to stop short of runway 2.

This can be modeled in an approximate way by imposing an arrival separation on the two runways. The value of the separation will be zero for all pairs of arrivals which will not conflict because the arrival on the hold short runway does hold short. The value of the separation will be greater than zero for those pairs of aircraft where the arrival on the hold short runway does not stop short of the intersection. For the example of the figure, the table below demonstrates how this would work. The two matrices are the arrival/arrival separation imposed on each landing.

The entries marked X in the table are set to an appropriate value greater than zero. A landing on runway 2 (top matrix) will ignore operations on runway 1, unless the landing on runway 1 is a class 1 and thus will not stop before the intersection. If the runway 1 landing is a class 1, the non-zero separation provides adequate separation from other aircraft crossing the intersection.

	Class		
Trailing a.c. on runway 2:	1	2	3
Lead a.c. on runway 1:	1	X	0
	2	X	0
	3	X	0
Trailing a.c. on runway 1:	1	2	3
Lead a.c. on runway 2:	1	X	X
	2	0	0
	3	0	0

6.c Intersection departures

Runways for commercial jet departures usually need to be 7000 or more feet in length. Aircraft used for business or private flying need only a small portion of this distance to takeoff. If they start in the same location as larger aircraft, then they will take a long time to fly down the length of the runway, blocking other movements. To reduce the runway occupancy time on takeoff, small aircraft may

depart at some intermediate point a good distance down the departure runway.

This can be modeled in FLAPS as an extreme case of the displaced threshold. In this case the second runway is much shorter and begins at the point from which the departure will roll.

7 Comparison of separation methodologies

Consideration of the methods used in the ASM model for separating aircraft was deferred until after the completion of the analysis of aircraft movements on runways and separation procedures between conflicting operations. This was done because proper comparison of ASM and FLAPS procedures involves the issues of aircraft performance and, in particular, dependent runway separation rules. Now that this has been presented, the methods used by ASM to separate aircraft will be described. The reasons why the ASM method cannot accurately represent certain cases modeled by FLAPS are discussed.

The method used by ASM to separate aircraft relies on two principles: 1) single runway occupancy rule and 2) separations imposed as a minimum time between crossings of the runway thresholds. The first point simply means that no more than one aircraft is ever allowed to occupy any one of the runways. Crossing a runway does not count as occupying

it. The second point is implemented in the following way: If operation Y is to be separated from a previous operation X by an interval t or more, the procedure is to hold Y at its runway threshold until t seconds after X has crossed its runway's threshold. This procedure is used between landings, takeoffs and between operations of different types. It is used when the two operations are on the same or on different runways. The result is that each possible combination of a lead operation type (landing, takeoff), a trail operation type, a lead operation runway, and a trail operation runway is considered a different type of separation. Each of these separation types is specified by a matrix of separation values (times or distances) for each combination of lead aircraft class and trail aircraft class. For example, a matrix of separations could be specified as applying between a leading arrival on runway 1 followed by a departure on runway 2. An independent set of separations could be specified to apply between arrivals on runway 2 that are followed by departures on runway 3.

The ASM procedure has certain strong points. It is easy to describe and permits modeling of a wide range of situations in a unified framework. The defect of this procedure relative to the methods discussed above for FLAPS is that, as is the case with many other aspects of the ASM model, it sometimes ignores important interactions between the various parameters of the airport. As a result of this, certain cases can only be approximately modeled by ASM.

One such situation involves the separation of departures on one runway from arrivals on an intersecting runway. As has been stated, the rule is that the departure must be held until the arrival has either crossed the intersection or exited, whichever occurs first. This is modeled in ASM by the analyst calculating the average time each class of arrivals will need to travel from the threshold to the intersection and using this time as the A/D separation.

We refer to this average time to the intersection as S^* . This will be applied in the model as follows: When a departure is to be scheduled, it will not be allowed to begin until a time S after the previous arrival has crossed its runways' threshold. This differs from the actual rule of holding the departure until after the minimum of when the arrival crosses the intersection or clears the arrival runway. If most or all arrivals exit after the intersection, then the error introduced is, in practice, insignificant. The error would only be the absolute value of the difference between the actual time to crossing and average time to crossing for the landing, perhaps 3 or 4 seconds.

* S is actually a function of the arrival class. The following discussion does not depend on this, so, for simplicity we assume S is independent of the arrival class.

The difficulty with the ASM procedure comes when a significant fraction of arrivals exit before the intersection. If S , as defined above, is used as the A/D separation, then serious errors will result. When an arrival exits early, before the intersection, the departure will continue to be held for the full time S after the arrival crosses the threshold. This error could be 15 or 20 seconds and, in some situations, could be enough to mean that there will not be sufficient time between arrivals to permit a takeoff. In this situation the analyst could elect to reduce S to the average time the arrival will block the departure. That is, S would now be set to the probability that the arrival would cross the intersection times the average time to the intersection plus the probability that the arrival would exit early times the average length of time the arrival would be on the runway given that it exited early.*

This value for S may also cause erroneous behavior. It is true that this method will delay, on the average, each departure the correct interval after the arrival touches down. Nonetheless, this method may introduce a significant bias to the results. Under the arrival priority mode of operations, some departures will be able to go while others will not have sufficient time between arrivals to do so. This decision will often depend on whether the arrival exits

* The latter term might itself be a weighted average of times to several exits before the intersection

early or late. Taking the average arrival time on the runway may significantly change the fraction of interarrival gaps that can accommodate a departure.

For example, suppose that the time between the touchdown of a class 3 arrival and a class 4 arrival is 90 seconds and that the average time the class 3 arrival blocks the departure is 45 seconds if the arrival crosses the intersection and 30 seconds if it does not. Then, if the requirement for releasing a departure is that there be 50 seconds from the moment the takeoff begins to roll until the next arrival touches down, the actual behavior will be to allow the departure to go almost everytime the first arrival exits early (as $90 - 30 > 50$) but almost never* when the arrival exits late and crosses the intersection (as $90 - 45 < 50$). If the average time for the arrival to block the runway turns out to be below 40 seconds, then using the average will bias results towards letting more departures takeoff. If this average is over 40 seconds, then using the average time will bias results towards letting fewer departures takeoff.

This same problem occurs with parallel runways. When parallel runways are used with the separation requirement that a departure must hold until the arrival exits, the identical difficulty with using average arrival runway occupancy time will occur.

* The exact fraction would depend on the standard deviations of buffers and approach speeds and related factors.

An additional problem can occur when the analyst is interested in studying the effects of changes in aircraft performance or new runway exits on delay and capacity. Changing runway occupancy times changes the separations imposed in the cases described above. In the ASM model this interaction can be modeled only by the analyst recalculating the average runway occupancy times. The interaction between aircraft performance or type of exits on delay and congestion can take place, for the cases described, only outside the model with analyst intervention. This is a problem not only for the cases described, but also in other situations where the separation depends on runway occupancy time. In the intersecting runways case, D/A separations will be specified to prevent an arrival from being too close to the threshold when a departure crosses the intersection. This time will be specified as a minimum between the time when the departure leaves its threshold and when the arrival crosses its threshold. Thus, it is based, in part, on the time the arrival needs to cross the intersection with the arrival runway.

The separation methodology described in this and the previous sections does not suffer from these problems. FLAPS uses the exact time an arrival crosses the intersection or clears the runway, as appropriate, to control the release of departures. This prevents any biasing of results and also explicitly couples aircraft

performance to the determination of separations. This procedure will facilitate experimentation with FLAPS to assess the significance of proposed changes in runway exits on aircraft performance. A further advantage of the procedure in FLAPS derives from the inclusion of the landing performance model of section II.C. This means that the entire set of responses induced by, for example, changing a 90 degree exit (exit velocity 0 kts.) to a high speed exit (exit velocity 30 kts.) can be tracked by changing one parameter, rather than by first manually recalculating runway occupancy times and then manually recalculating separations.

F CONTROL OF THE AIRPORT

1 Introduction

The previous sections of this chapter have examined how the various portions of an airport operate. At several points we have seen how a given airport facility can be operated in a number of different modes. For example, runways may be in an alternating priority mode or an arrival priority mode. We did not consider at length the motivations for changing these policies.

The purpose of this section is to analyze the causes of changes in the way airports are operated. The frequencies of changes will also be examined. Our method of approach is to identify the major causes of changes and the resulting effects. We term these causes of change "factors" and the set of possible parameters they can alter the "operating policies" of the airport. We will see that two major factors (weather and congestion) and some minor factors can cause a complex set of responses in operating policies. These policies include runway assignment, aircraft performance, and air traffic control separations.

The frequency with which changes in operating policies occur means that models for analyzing airport operations should have a dynamic capability to follow changes in operating policies. Correct modeling of the transitions

from one set of operating rules to another will be discussed.

It will be shown that the restricted ability of existing models to analyze situations involving changes in operating rules limits their applicability and increases the difficulty of their validation.

We will propose capabilities that are a significant advance over present models. We will show that the type of changes in operating policies encountered in practice can be described as one of two kinds - changes triggered at a certain time and changes triggered in response to a certain condition on the airport. We develop each of these specific procedures in response to a specific factor and then generalize them for use in modeling changes caused by other factors.

2 Weather

2.a Summary of effects

Weather conditions have a profound effect on the operation of an airport, changing a number of operating policies.

Separations

As indicated in section II.C, different separations apply under VFR and under IFR conditions for arrival/arrival spacing. The difference in capacity and, therefore, delay

is significant. A typical reduction in landing capacity would be from 40 landings per hour in VFR conditions to 30 per hour in IFR conditions. Transition between types of separations can probably be modeled as occurring instantaneously. That is, the first aircraft to be scheduled after the time of the transition merely uses the new separation.

Dependencies among runways

As seen in section II.D (also see chapter IV) reductions in visibility and ceiling can make a pair of previously independent parallel runways inter-dependent. This will also reduce the capacity of the airport and increase delays. As with separations, changes in dependency can probably be considered instantaneous. The full impact of these changes will require several minutes to take effect and will not do so until all aircraft scheduled under the old procedures have cleared the runway.

Runway configuration in use

By runway configuration we mean the choice of which of the airport's runways are active (have traffic assigned to them), for which operations they are active (landings, takeoffs or mixed operations) and in which direction they are to be operated.

Aircraft must land into the wind and the direction and velocity of the wind can dictate which runways are usable. The runway configuration in use often changes several times during the course of a day at a typical airport.

Information on configurations in use at La Guardia for one week in December, 1978, was made available to the author. During seven days an average of 2.4 changes in configuration were made each day between 7 A.M. and 10 P.M. On no day was the airport operated entirely with one configuration.

A second way that weather affects runway use is that during low visibility a number of lighting and landing aids (visual approach slope indicator, instrument landing system) may become required for landing. Each runway is equipped with various combinations of these aids, and this can govern which runways are open for use.

Transition between runway configurations is fairly complex and cannot be considered to take place instantaneously. Once a change is decided upon, aircraft must begin to be directed to the new runways to be used. For landings this may mean that, for some time, aircraft may continue to operate on the old runways, as they may be too close to the approach gate to be diverted. Any departure that has begun to taxi to a takeoff runway or is at the runway endpoint will often be allowed to continue to the runway and takeoff. Once no more aircraft remain to use the old configuration, operations will begin to use the new configuration. Separations can be provided in the manner described earlier in this section. If the direction of use of a particular runway has changed then, at a minimum, all

aircraft on the runway will have to be cleared before operations can begin in the opposite direction. This would mean that any takeoff operating under the old regime would have to have climbed out before any landings could use the runway in the new direction. There would likely be a delay of several minutes between the last takeoff in one direction and the first landing in the other direction.

To model the transition the following procedure is used. At the designated time of the changeover the runway assignment procedure is changed to assign any new departures from the apron to the new runways. At changeover plus C_x minutes the runway assignment procedure is changed to assign any new arrivals to the new arrival runways. Operations are scheduled to meet any relevant separation, the user having specified proper separation between the old and new sets of runways. If a landing on a runway will use the runway in a different direction from the previous landing, then it is not allowed to begin its descent from the approach gate until C_a minutes after the previous takeoff has cleared the runway. If a takeoff on a runway will use the runway in a different direction from the previous takeoff, then it is not allowed to begin its operation until all aircraft currently on the runway have cleared. Times C_x and C_a are set by the analyst.

Aircraft mix

Landings in bad weather also require both specific pilot skills and on-board navigation aids. Pilots who might be permitted to takeoff may still cancel their flights due to bad weather in surrounding area. Small aircraft tend to have limited capabilities for bad weather operations and to be flown more often by pilots without the training for bad weather flying. Therefore, light aircraft will probably constitute a smaller percentage of the mix during bad weather. Procedures for changing the aircraft mix over time were given in the demand section, II.C, above.

Runway assignment

The result of the last several factors will be shifts in the overall runway assignment for landings and departures. This may imply either changes in the percentage assignment or a change to a new mode of operations.

Aircraft performance

Rain can reduce the effectiveness of braking and change exit selection for landings. Taxiing speeds are probably lower in bad weather, as pilots tend to be more cautious.

2.b Implications for analysis

Being able to follow changes in weather is a necessary capability for airport models. This is because weather conditions can change in several different ways over the course of a day. Changes in visibility may mandate a change

in flight rules (VFR to IFR or reverse). Low visibility is often but not always associated with rain or snow. Wind direction can change several times a day in any kind of weather. The result is that it is rare to find an airport operating in the same configuration, using the same rules for an entire day.

If we wish to fully analyze how airport operations, our models must have an ability to alter the conditions in effect during the course of a run and to model the transition from one configuration to the next.

2.c Evaluation of existing models

The ASM airfield model has a very limited ability to follow weather-induced changes in operating policies. No changes in separations, dependencies or aircraft performance are allowed. One change in the active runways is allowed during any run. However, this is done at the expense of halving the number of runways that may be active at any one time. Because of the schedule techniques used by the ASM model, changes in aircraft mix and minor changes in runway assignment (i. e. relative probabilities of assignment) can be simulated by the user suitably constructing the base schedule. The MITASIM model has no ability to model this type of changes.

The limitations on the ability of existing models to analyze this type of change result in a significant

constraint on the data that can be used to validate the models.

A natural procedure for validating such a model would be to take measurements of flow rates, delay, aircraft characteristics, etc. at an airport for one or two weeks, such as was done at O'Hare in the summer of 1977. This record will typically show some change in operating policies every 2 to 4 hours. If a capacity analysis is being done one may combine the records of a given configuration if it appears more than once during the study period. Any comprehensive validation will involve a sequence of comparisons between the model and data obtained under each operating policy.

Analysis of the models' performance is foreclosed as a practical matter by the existing models' lack of ability to simulate alterations in operating rules. If matching of delay performance is desired, no combination of data from different periods with constant operating policies is appropriate due to the difference in initial conditions for each period.

Modeling of changes as a sequence of simulations of successive time periods would be difficult. An example offers the best explanation of why this is the case. Suppose an airport uses one set of operating rules before 12 noon, and a second from noon onward. Assuming the airport was essentially empty at some early morning time, the period

before noon can be simulated without difficulty. Simulation of the period after noon cannot be performed accurately unless 1) the condition of the airport at noon can be used to start the second run of the model and 2) transition between the first and second operating rules can be represented. The first problem may be approximated by making the second run of the model with a warmup period. The demand rates in this warmup period would be adjusted until the set of conditions at noon in the second run closely approximate the model generated conditions at noon at the end of the first run. The more congested the airport really was at noon, the more important is this matching. The second problem, transition time, may or may not be significant depending on the type of change in operating policy being made. The lull in operations caused by changes in runway configurations will be hard to model.

The problems just discussed in terms of validation also appear when existing models are used for airport analysis. Investigations of the effects of changes in operating policies is an important issue for controllers. These effects cannot be analyzed using existing models.

2.d Modeling of policy changes

Changes in weather conditions are external to the airport. They occur at a certain time, rather than being dependent on certain airport conditions. In order to model

them, we need a procedure to permit arbitrary changes in operating rules at any desired time. This turns out to be very easy to do. FLAPS allows the input of a file of orders, each consisting of a time and a set of parameter names and values. During each replication* of a run each order is read at its specified time, and the changes in parameter values are implemented. Most policy changes involve simple changes in command values. Changes in runway configurations involve a set of orders to implement the change. See chapter IV for several examples.

3 Congestion

3.a Summary of impacts

The buildup in the number of aircraft waiting to land or takeoff can trigger changes in operating policies. The most obvious way that this occurs is in a switch from an arrival priority mode to an alternating priority mode for a mixed operations runway when the queue of departures rises above some critical value. The alternating mode is then maintained until the number of aircraft in the departure queue has been reduced to some acceptable value.

There are a number of possible "second order" policy shifts that could occur as a reaction to congestion. First, if the departure queue continues to grow despite the imposition of alternating priority rules, then the

* See section II.G for definition.

controller may elect to change to the departure only mode by halting all arrival operations on the runway.

A shift might be made to an entirely new runway configuration. Airports are normally run using the configuration with the highest capacity that weather conditions make available. However, if, due to noise abatement considerations, for example, a lower capacity configuration is being used, then congestion may induce a change to increase capacity.

Even more subtle changes in operating policies are possible if controllers can anticipate the traffic demands immediately ahead. Then, for example, they may elect not to change operating policies if they expect a lull in the flow of arrivals. It may be that some responses to congestion are best described as changes in runway assignment. There may be a runway nominally in use but, during severe congestion some aircraft may be directed from it to another runway. Congestion may induce changes in aircraft performance, as controllers urge pilots to expedite their movements. The extent to which such procedures are used is unknown.

A wide range of capabilities to alter operating policies in response to congestion is necessary for studying which policies are most effective. It may be that certain commonly used policies are not optimal.

3.b Evaluation of existing models

Both the ASM model and the MITASIM model (NOR 79) have a capability to automatically alter the operating mode for a mixed operations runway. The ASM model does this by specifying a critical queue length and a time interval. Then, each time an arrival is scheduled, the departure queue for the runway is checked. If it is above the critical value, then the specified time interval is added to the A/A separation.

The MITASIM model employed a method that improves on the ASM procedure in two respects. First, two critical queue lengths are specified: one for the transition "up" from arrival priority to alternating priority and a second for the transition "down" from alternating priority to arrival priority. By specifying the latter value to be several aircraft lower than the first, rapid alternation between the priorities is prevented. The second change is in the manner in which the amount of extra space between arrivals is provided. MITASIM uses the "calculated separation" procedure described in section II.D. Neither model permits the more complicated changes discussed above. Lack of more complicated policy alteration options reduces the ability of the models to be used to explore optimal policies.

3.c Modeling of congestion induced changes

Changes caused by congestion cannot be modeled through the same procedures used for weather-induced changes. This is because the triggering factor is not external to the airport (weather) but is one aspect of how the airport is operating. What is external is the definition of the condition. That is, the exact situation that causes the changes, queue length, for example.

Modeling of these changes is done in FLAPS by defining a set of conditions (such as takeoff queue on runway 3 greater than X aircraft) that will trigger an event. The analyst designates what conditions are to be set and at what value (X). When they occur, a list of parameters is automatically read and changes are made. Currently FLAPS allows only two general types of conditions: 1) takeoff queue on runway X has more than Y aircraft, and 2) takeoff queue on runway X has less than Y aircraft.

There are no restraints on what changes in parameters can be made when the condition is true. Each time the takeoff queue is added to (condition 1) or subtracted from (condition 2), the relevant condition is checked and triggered if true.

Extensions to FLAPS can be made to enlarge the above set of conditions.

4 Other factors

A number of additional factors may occur that result in changes in operating policies. Certain runways may be subject to a curfew to control noise exposure. Accidents may block a runway. Perhaps an extreme example is provided by Logan Airport where several runways overfly shipping lanes. When a liquid natural gas (LNG) tanker enters or leaves Boston Harbor, it may not be overflown. While the LNG ship makes its transit, several runways are rendered inactive for intervals of 10 to 20 minutes at a time. The entire process may take several hours. Fortunately, this does not occur often enough to be a major issue to be analyzed, but it does illustrate the limiting case.

Whatever motivates the changes in operating policies, the two procedures described above provide a wide ranging capability to model them. As long as the instigation of the change can be described as occurring at a specific time or under a specific set of airport conditions, and the desired policy change is describable in terms of the runway and aircraft parameters specified so far, then the procedures set up above can be used to model it.

G STATISTICAL ISSUES IN MODELING AIRPORT OPERATIONS

1 Introduction

This section will discuss statistical considerations which are pertinent to the analysis of airport operations. FLAPS is implemented as an event paced digital simulation, and we will first briefly survey the current state of statistical analysis of such simulations. Certain aspects of current work relate directly to the type of situation we wish to analyze, but the airport environment poses a number of significant problems for the statistician. In each area we draw conclusions about the statistical features relevant for airport simulation. The ASM model is briefly evaluated with respect to each of these topics. Secondly, we will consider what output statistics are required by the cases we wish to analyze with FLAPS.

2 Survey of Statistical Issues Relevant to Simulation

2.a Simulation context

In order to understand the application of what follows, it is necessary to discuss the basic concept of a simulation from a statistical point of view.

A simulation such as FLAPS is based on a stochastic description of how the airport operates. The computers used

to implement simulations are completely deterministic. Thus, some way must be found to simulate randomness. This topic is discussed below under section 2.b.

This probabilistic base to simulation means that the outputs of a simulation are random variables and must be analyzed as such. Each "trip" through the simulation is but one realization or replication of the simulation. Running the identical case with different random numbers will produce a different realization. Methods of reducing the significance of variation among replications are discussed in section 2.c

Of major significance to the analysis of simulation is that its outputs are inherently autocorrelated. For example, if the airport is very congested, then the delay experienced by a typical aircraft is likely to be large. It is very likely that the delay experienced by following aircraft will also be large. This autocorrelation complicates application of estimation procedures considerably.

There are two fundamental approaches that are used to obtain the true average behavior in the presence of autocorrelation. One approach is to run the simulation for an extended period of time. Average behavior is obtained by computing estimators over the entire duration of the run. The second procedure is to make many shorter replications, each starting at some initial condition (usually empty and

idle). Average behavior is obtained by computing estimates across the several replications. There are a number of statistical tradeoffs between the two methods. Each method creates certain problems for the analyst. Some of the more significant issues are reviewed below in section 2.d.

2.b Random number generation

The subject of how to generate random numbers is vast and complicated. Fishman cites (FIS 78 pp.350) a 1972 bibliography with 491 entries on this subject. The pace of developments is such that much of the best work from 1972 and before has been surpassed. It is not our purpose to attempt to contribute to the theory of random number generators (RNG), but anyone who constructs a simulation must be sensitive to the ways the use of an RNG affects the model.

Choice of random number generator

Much has been written (FIS 78 pp. 345-91) on the correct choice of RNGs. This choice poses difficulties for the analyst, as RNGs may be machine dependent, are usually written in assembler language and exceed the abilities of most application oriented users to meaningfully test them. For example, even the GPSS simulation language package employs a RNG with so many problems that it leads Fishman to this conclusion: "Anyone who contemplates using GPSS cannot expect to defend his results successfully on a truly scientific basis." (FIS 78 p. 166)

The majority of RNGs in use today are multiplicative congruential (MC) generators which generate a stream of numbers, $Z(i)$, which are uniform on the interval from 0 to the largest integer capable of being stored in the machine. The $Z(i)$ are divided by this maximum size to obtain $r(i)$, a number uniformly distributed on the interval 0 to 1. The formula used by all MC generators is of the form:

$$Z(i) = A * Z(i-1) \text{ (modulo } M)$$

The user supplies a starting seed for $Z(0)$ and the RNG recursively generates the series $Z(i)$.

The ASM model uses an undocumented MC generator with $A = 2051$ and $M = 4194304 = 2^{22}$. This form restricts the initial seed $Z(0)$ to odd numbers, and the FAA routine has a mechanism to insure that, if the user inputs an even seed, it is "bumped" by 1 unit to the next odd value.* This provision is valuable as the sample inputs given (PMM 77 p.87) include 10 random number seeds (for use in 10 replications), the numbers 1001 to 1010 inclusive. It may therefore be the case that sample outputs and validation runs reported by the FAA are really based on 6 different replications, not 10. Four replications may be doubled, those replications beginning with the seeds (1002,1003), (1004,1005), (1006,1007) and (1008,1009) being the same.

* These observations are based on a listing of the code of the ASM model.

The routine used in FLAPS was developed by Lewis (LEW 69). It was used in the FORTRAN IMSL statistical package. It uses $A = 16807$ and $M = 2147483647 = 2^{31} - 1$. This is a variety of MC generator known as a prime modulus multiplicative congruential generator. Fishman applies a large number of tests to several generators of this form and his conclusion is that this one is acceptable (FIS 78 p.369, 382). Users concerned about the quality of the RNG may elect to change A to one of the other values that Fishman tests which perform somewhat better.

Use of random number generators

Several authors (FIS 78 pp.35-6, SCE 74 pp. 144-7) emphasize that the correct use of a RNG in a simulation involves not only choosing an acceptable algorithm, but also 1) having each process in the model use its own independent stream of random numbers and 2) explicit identification in the simulation outputs of the first and last seed used. The first point allows for more reproducible experiments. The user may change one process in the simulation, without disturbing the sequence of random numbers used in the other processes. This allows a clearer examination of the effects of the change in the outputs than would be the case where all the random number streams were changed. Seeing the final seed used allows the analyst to prove that two cases used exactly the same random variables except for the process that was altered. This also allows the running of

independent experiments by using the last seed in the first experiment as the first seed in the second.

The ASM model provides neither of these procedures. FLAPS allows the printing of beginning and ending seeds if the analyst so requests. FLAPS uses three streams of random numbers, one for calculation of separations, one for the generation of schedules and one for general use. The user may specify any seeds for these streams, but default values are provided from the table in FIS 78 p. 486 to be 1 million apart in the sequence generated by this RNG. Preliminary runs of FLAPS indicate that each stream may be called on the order of 100,000 times in a "moderately large" run.

Generation of non-uniform numbers

There are needs for random numbers with distributions other than uniform. There are a wide variety of techniques catalogued for generating a wide variety of probability density functions. See, for example, FIS 78 pp. 392-480. FLAPS allows the user a choice of six probability density function forms whenever a random variable is to be used:

- 1) Deterministic
- 2) Normal - using the central limit theorem with user specified truncation
- 3) Normal - using a formula from Box & Muller (BOX 58, also see FIS 78 p. 410)
- 4) Triangle - convolution of two uniform numbers
- 5) Erlang - with user controlled order (k)
(k = 1 is negative exponential)
- 6) Uniform

This allows the user to experiment with the sensitivity of results to different probability density function forms. This may prove useful in view of the scarcity of information about the exact probabilistic form of many key airport parameters.

2.c Variance reduction techniques

There has been considerable interest in developing ways to control stochasticity in simulations in order to obtain more accurate estimates of simulation outputs (KLE 74 pp.105-263, FIS 78 pp.114-126). Techniques such as stratified sampling, control variates and importance sampling are not directly applicable for use in a simulation that has both multiple inputs and multiple outputs, as is the case here.

One method which may be applied is the technique of antithetic variates. The basic idea is that the variance of the simulation outputs may be reduced by having each pair of replications use streams of random numbers that are negatively correlated. One way to achieve this is as follows: If replication i uses a stream of 0 to 1 random variables $r(j)$, then replication $i + 1$ should use stream $s(j)$, where $s(j) = 1 - r(j)$. For replication $i+2$ the procedure is started again with a new seed for the stream $r(j)$. With the MC type of random number generator, this method proves easy to implement. It has been shown (KLE 74

pp.254-6) that streams $r(j)$ and $s(j)$ will have this property if the seeds used to start the second stream $s(j)$ is $M - Z(0)$, where $Z(0)$ is the seed used to start the first process and M is the modulus of the RNG, as indicated above.

When analyzing the outputs from this situation, each replication can no longer be considered as independent. However, each pair of replications can be so considered. Thus, the degrees of freedom in any average across n replications is no longer $n - 1$ but $n/2 - 1$. The source of this technique (KLE 74 p. 199) also reports its successful application to a number of simulations.

The procedure was not included in FLAPS but could be easily added as an extension to the model.

2.d Output analysis

There is much work reported on analysis of the outputs of a simulation. However, few advances have been made in analysis of the particular type of simulation represented by FLAPS. We will first discuss why FLAPS must be run as a multiple replication type of simulation and how estimators are produced for this situation. We will then discuss a number of problems and techniques used to analyze this type of simulation.

Overview of analytical framework

The basic complicating factor in analyzing simulation outputs is the presence of autocorrelation. There are two

distinct approaches to dealing with this problem. One method operates the simulation as a single long replication but attempts to manipulate the outputs to extract meaningful estimators. Analysis focuses on initial condition bias, final condition bias and various methods of estimating mean values from autocorrelated data. Among the latter techniques are skipping data points or blocks of data, using the regenerative nature of queueing systems and performing time series analysis (GOR 69 pp.277-95, FIS 71, KLE 74 pp.87-90, FIS 78 pp.219-273).

Implicit in all of these procedures, however, is the assumption that the system being modeled is in a steady state condition. Since a significant reason for constructing an airport simulation in the first place is the fact that the airport is seldom in steady state we cannot use a single long experiment for our analysis. If accurate estimates of the performance of an airport under time varying loads is to be obtained then the model must be formulated and run to make multiple replications of the same time period. Only in this way can the average performance and the range of typical performances be determined. Unfortunately, little analysis has been done on the multiple replication case. Even the infrequent paper studying of the multiple replication case, such as TUR 76, is often concerned with steady state situations.

The standard treatment for multiple replication simulations is to assume that each replication is independent and that any set of observations of an output across replications will be normally distributed, permitting use of the classical estimation process (KLE 74 p. 85-6). For example, if $X(i,j)$ is some performance statistic, such as average delay to all landing aircraft in period j of replication i , and I is the number of replications to be run, then the estimate of the true average landing delay in period j would be:

$$X(j) = (1/I) \sum_{i=1}^I x(i,j) \quad (23)$$

the standard deviation of $x(i,j)$ would be:

$$Xs(j) = (1/(I-1)) * \sum_{i=1}^I x(i,j)^2 - (1/I) * \left[\sum_{i=1}^I x(i,j) \right]^2 \quad (24)$$

Finally, the 95% confidence interval of the mean would extend from $X(j) - Xc(j)$ to $X(j) + Xc(j)$ where:

$$Xc(j) = Xs(j) * t(.025, I-1) / \text{SQRT}(i) \quad (25)$$

The term $t(.025, n)$ represents the t statistic for a 95% two tail confidence interval with n degrees of freedom.

Assumption of normality

An assumption of the classical estimators given above is that the underlying process, $X(i,j)$, is normally distributed. The question naturally arises of whether this model will generate normally distributed data, and if it does not, how dependent our estimators are on the assumption of normality.

The central limit theorem is general enough that it is not implausible that estimators across replications would be normally distributed. Kleijnen discusses (KLE 74 p. 87, 454-7, also see GOR 69 p. 280) a form of the central limit theorem called the "stationary r-dependent central limit theorem" which applies specifically to the case of interest to us, where we want to compute estimates of $X(j)$ and the individual terms $x(i,j)$ are themselves averages of successive points of a time series. The conclusion is that the distribution of $X(j)$ is asymptotically normal. The only caveats would be that the theorem applies to stationary processes and that we have no idea how fast it converges.

Reports on the distributional form of data generated from actual simulations are rare. FLAPS supports investigation of the question of distributional forms of the simulation output by printing replication by replication values, should the analyst so elect, for the majority of the statistics that are collected. We report here results for one such case as a preliminary investigation into the question of normality. Thorough investigation of this topic would be a major research undertaking.

The case that was used is described in section III.B as a part of the New York La Guardia airport capacity runs. We use the single runway configuration. The IFR capacity was found by FLAPS to be approximately 26.8 arrivals and 26.8 departures per hour. For the tests of normality this

configuration was run for 50 replications of 4 hours. Both the arrival and departure streams had an average demand rate of $0.8 * 26.8 = 21.4$ aircraft per hour. Note that due to gaps in the arrival process, the departure process is operating at less than 80% of capacity.

Table II-14 presents summary statistics for the run. The normality of 3 variables was investigated:

- 1) Total average landing delay
- 2) Total average takeoff delay
- 3) Number of landings in the 2nd period

Period	Landings	Landing Delay	Takeoffs	Takeoff Delay
1:00 to 2:00	21.0 (4.2)	4.1 (4.9)	21.9 (4.1)	3.2 (2.5)
Entire run	82.4 (7.7) 2.2	4.5 (2.9) 0.8	83.9 (9.5) 2.7	3.2 (1.5) 0.4

Results given as: mean
 (standard deviation)
 half width of 95% confidence interval

Arrival and departure flows:
Negative exponential interaircraft times
Average of 21.4 operations/hour

Remaining parameters:
Described in section B.1 of chapter III


Table II-14: Normality Test: Summary Results

The test for normality used was the Shapiro-Wilk test (SHA 65), described as among the most sensitive for a number of different types of departures from normality (FIS 78 p.72). The Shapiro-Wilk test statistic varies from 0 (all data grouped at one point) to 1 (perfectly normal data). For $n=50$ as used here, critical points (FIS 78 p.319) are:

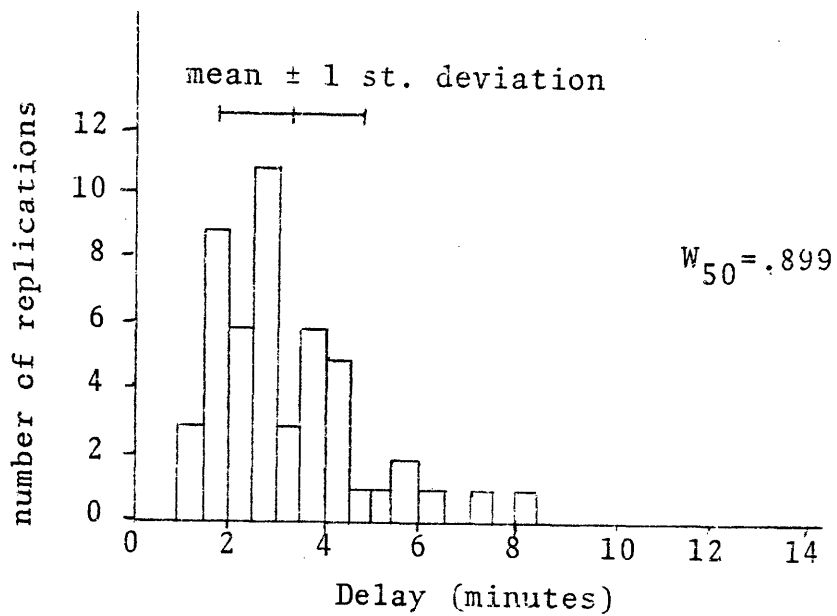
<u>Level</u>	<u>Critical Value of W</u>
.01	.930
.10	.955
.50	.974
.90	.985
.95	.988

Figures II-22 and II-23 present graphs of the statistics and the values of the Shapiro-Wilk test statistic, W_{50} , for each one.

The interarrival process to this simulation has the time between aircraft arrivals distributed negative exponentially. The service time distribution, as controlled by the distribution of the buffer on A/A separations, is normal. Therefore, we would expect the distribution of landing delays to be somewhere between those two distributions. Takeoffs are also generated with negative exponential interdeparture times. However, the service time function cannot be characterized in any direct way. D/D separations are deterministic for any single combination of aircraft classes but take on several values across various combinations of classes. The arrival process interferes



a) Landing Average Delay



b) Takeoff Average Delay

FIGURE II-22. Distribution of Average Delay.




FIGURE II-23. Distribution of Number of Landings.

with the landing process, further complicating the distribution. The sum of all these factors might be expected to result in a more normal-like distribution for takeoff delay than landing delay.

The statistic for average landing delay to all aircraft is graphed in figure II-23.a. There is much less than one chance in a hundred that the underlying process could be normal and generate this distribution of results. The distribution produced could be better described as gamma than Gaussian. The distribution of average takeoff delay, in part b of the same figure, is closer to being symmetrical than landing delay. The value of the test statistic, .899, is still below the .01 possibility of being generated by samples from a normal population. It is unclear whether takeoff delay appears more normal than landing delay because the fundamental process is more normal, or because the average value of takeoff delay is lower than landing delay. The lower overall average may mean that the takeoff delay distribution produced here is merely a scaled down version of the landing delay distribution.

Delay to an aircraft cannot, of course, be negative. It may be that at low values of average delay this fact will lead naturally to the one-tailed form of delay distributions we see here. Higher average delay values might result in the distribution converging to a more Gaussian form.

Only one distribution of the number of operations was plotted. The one chosen was the distribution of the number of landings in period 2, shown in figure II-23. In spite of the fact that the underlying distribution for this statistic is discrete and that the Shapiro-Wilk test is the second most sensitive among tests for normality for this specific problem (FIS 78 p. 73), the distribution is the most normal of the three. We can reject the null hypothesis of normality at a very low significance level (between .10 and .50). However, the fact that this test covers only one period, as opposed to the full length of the simulation, while scoring higher on W50, may be taken to indicate that this process is more normal than either delay process.

The tentative conclusion which emerges from this preliminary analysis is that more attention should be given to the question of the distributional form of outputs of complex simulations. The assumption of normality may not always be justified.

In view of the non-normality of these outputs, it is reassuring that the classic estimators for X_s and X_c turn out to be robust under departures from normality. Non-normality does not bias the estimate of X_s , nor does the test using the t statistic depend on normality, so the estimate of X_c may be used as presented (BOX 53, SCH 59 p.335, 346). Both the cited authors make the point that the same is not true for any inferences about the distribution

of X_s . Were confidence intervals for X_s to be calculated, their validity would depend critically on the underlying $X(j)$ being normal. In summary, it can be concluded that we may use the procedure of equations (23) to (25).

Time series analysis

Time series analysis was mentioned as a technique appropriate for single simulation experiments. However, it may also have utility as a way of validating outputs of a multiple replication simulation. Using the earlier example, the set of $X(j)$ would be considered a time series and compared with the time series of real world observations, $Y(j)$, of the performance measure. The reason for considering such a test is, again, the high probability that adjacent elements of both of these series are autocorrelated. The procedure would be to fit a Box & Jenkins moving average, autoregressive model to both $X(j)$ and $Y(j)$ and examine the resulting models for differences. Fishman was an early advocate of the use of time series and spectral analysis procedures (FIS 67). Hsu and Hunter have recently (HSU 77) performed such a comparison of actual and estimated time series on a model of air traffic control communications. They derive an estimator for evaluating the comparison.

Time series analysis is conceptually very appealing as a method of examining simulation outputs. However, there are a number of difficulties with applying it to the airport

problem. First of all, it does assume the underlying system is covariance stationary, i. e. in steady state. If there is any trend or cycle in the original data, it must be removed. The demand function driving the simulation may provide one possible estimate of the trend to use in extracting the steady state time series. A more serious problem is the short length of any time series to be compared. Even if the $X(j)$ and $Y(j)$ are sampled as frequently as every 15 minutes, a 12 hour simulation will produce only 48 data points, insufficient for reliable estimates. Further limitations on the use of time series methods arise from the fact that data on the overall performance of airports is, as we will see in chapter III, both rare and unreliable. It is unusual enough to find mean values of performance measures, let alone time series of such measures for an entire day. So, while the method is undoubtedly applicable to this case, there are severe practical limits on its usefulness.

3 Selection of Simulation Outputs

An important design issue for FLAPS is what outputs the model should produce. We need to consider both what actual performance measures should be calculated and also how they should be aggregated over time and space.

3.a Statistics for capacity and delay

We have often referred to the model as one oriented to capacity and delay issues. At the heart of the outputs are the estimates of delay to aircraft and the number of movements that take place.

Capacity is assessed by counting the number of aircraft movements. Three measures are often suggested for assessing delay: 1) the average delay, 2) the fraction of aircraft delayed more than some value and 3) the length of the queue of aircraft waiting to use the runway.

To enable analysis of transient conditions, each statistic must be measured over smaller periods of the run in addition to averages being computed for the entire run. As mentioned in the demand generation section II.C above, the model divides the run into a series of equal length time periods. Each capacity and delay statistic is measured for each time period.

Additional breakdowns are needed to assess performance. Each runway's activities are reported separately as well as the entire airport's. Landings and takeoffs are tabulated separately.

Significance of delay statistics

To accurately interpret the outputs, it is important to understand exactly when and what is being measured. Each time an aircraft begins a landing or a takeoff movement, the difference in time between the scheduled and actual start of

the movement is counted as the delay experienced by the aircraft. This measure corresponds to W_q , waiting time in the queue in queueing notation. The period in which the aircraft is counted is the period in which it completes the runway operation.

The meaning of this delay measure for landing aircraft may be questioned as, in fact, aircraft do not queue up at the approach gate. Any arrival delay encountered is probably taken in vectoring or in holding stacks many miles from the runway. The point, however, is that this delay occurs because of the limitations on the acceptance rate of the runway. To the extent that FLAPS correctly models this process, it correctly measures the delay induced by it. An additional point about the collection of average delay statistics will be made below after the discussion of extreme delay statistics.

Analysts are often interested in the fraction of aircraft that experience extreme delay, where the definition of extreme may be set by the user. An often used value is 20 minutes. We must consider carefully the construction of this statistic. One method would be merely to maintain a running count for each period, runway and type of operation of the total number of operations and of the number experiencing extreme delay. At the end of the set of replications the ratio of the two would be an estimate of the true proportion.

The only difficulty with the procedure is that each replication will have different numbers of operations. We would expect that replications with more operations will tend to have a higher proportion with extreme delays. This, moreover, might be reversed for operations of low priority in the model where high congestion might actually mean fewer operations of low priority. The procedure suggested above would tend to give greater weight in the overall average to replications with higher flow rates. Yet each replication in some sense, is an equally plausible realization of the system. Therefore, there seems no reason why each replication should count differently in the final average.

The following measure is proposed for assigning equal weight to each replication. For each replication a proportion, $p(i)^*$, of aircraft delayed more than the critical amount is calculated. Then, after the simulation is over, the average of $p(i)$, standard deviation of $p(i)$ and confidence interval of the average of $p(i)$ is calculated in the same manner as X , X_s and X_c are calculated. The only difference is that no observation of $p(i)$ is collected for any replication where no operations took place. The assumption that the t statistic may be used on this procedure can be defended on the grounds that the distribution of $p(i)$ is binomial (with a change of scale) and that the binomial distribution converges to the normal

* $p(i)$ will have other dimensions for period, runway and type of operation, here omitted for clarity.

as the number of observations increases.

The same point about the weight given to each replication made in connection to extreme delay measures applies to average delay. Average delays are estimated in the same manner. After each replication an estimate is made of average delay by dividing total minutes of delay by number of operations. This is then averaged across replications to give the final estimate. As with extreme delay, the possibility exists for bias in either direction, if each replication is not given equal weight.

The queue of aircraft waiting to land and takeoff on each runway is measured at the end of each period. This is not an actual average of the queue length during the period but a "snapshot" of the queue length at the end of each period. This is much easier to calculate than the average queue length during the period. The "snapshot" procedure will have a larger variance than the period average measure. Finding the average queue length requires that an observation of the queue be taken every time an aircraft enters or leaves a queue. However, the "snapshot" procedure provides an unbiased estimate of the true length at the end of the period.

3.b Landing performance statistics

In the landing aircraft performance section, II.D.3 above, a number of modeling issues were raised regarding the

use aircraft make of runways. As discussed, we may want to assess the impact of new exits or new aircraft on the distribution of exits used or on runway occupancy time. In order to do this, statistics on landing performance are needed. Because of the nature of takeoff operations (see II.D.4) no statistics on takeoff performance are collected.

Each time an aircraft exits a landing runway the runway occupancy time and exit used are recorded. Average runway occupancy time and percentage use of each exit for each class of aircraft and each runway are calculated. A runway used in both directions has data collected separately for each direction. As the performance of each aircraft is independent of the performance of other aircraft and independent of the amount of delay at the airport, there is no need to obtain statistical estimates on a replication by replication basis, as was done with delay and extreme delay.

III MODEL VALIDATION

This chapter discusses the question of how a large-scale complex model such as FLAPS is validated. We will see that validation involves a number of different tests, including, but not limited to, comparisons between outputs of the model and observed data. This chapter presents several such comparisons of FLAPS to both observed data and the results of other models. However, the scarcity of reliable data and the incompleteness of other models means that confidence in the validity of FLAPS cannot rest entirely, or even primarily, on such tests.

A THE PROBLEM OF VALIDATION

This section discusses how large-scale models can be validated. The concept of "validation" has been described as the one major remaining methodological problem in simulation (NAY 67 pB-92). A major reason for this is that validation "involves a host of practical, theoretical, statistical, and even philosophical complexities." (NAY 71 p.21) At its most fundamental level, validation involves unresolved epistemological questions of how we can "know" something to be true. In the context of complex models, a number of conclusions have come to be generally accepted. We review them here.

A consensus seems to be emerging to divide what has so far been loosely called validation into three stages (NAY 67, LAW 80, GAS 80):

- 1) Verification - establishing the extent to which a computer implementation of the model is equal to the description of it in formulas, on paper or in the analyst's mind,
- 2) Validation proper - establishing the extent to which a model description, such as chapter II of this work, is in fact an accurate representation of how the real world functions, and
- 3) Certification - establishing that the model is appropriate to answer the questions being studied.

Speaking colloquially, verification is proving that you did what you thought you did, validation is proving that the world operates the way you think it does, and certification is proving that you are studying the right thing to begin with.

Verification involves techniques such as modular design of the computer program, detailed walk-throughs of outputs, deterministic runs, runs equivalent to simple queueing systems, and so forth. Much verification work was done for FLAPS, but it is a separate issue from the concern of this chapter and is not described here.

Certification involves topics such as identifying critical questions about the real world system and finding

future problem areas. The work that was done in this area for FLAPS is contained at various points in Chapter II where the importance of analyzing capacity and related questions for airports is established. The criticism of the ASM model's attention to ground operations at the expense of runway operations, its schedule generation methodology and its inability to model changes in operating rules are primarily criticisms of its appropriateness as an analytical tool. The development of FLAPS to overcome these limitations is an attempt to create a more appropriate analytical tool. The results presented in Chapter IV also demonstrate the appropriateness of the model for studying airport operations.

We turn now to the consideration of validation as defined narrowly in point 2 above.

First, validation is a question of degree and not an absolute yes or no decision. The more tests performed, the more confidence one can have in the model, but it is neither cost effective nor possible to pursue complete validation (NAY 67 p.B-93, LAW 79 p.8). If complete and total validation could be achieved, there would in fact be little purpose to building a simulation, as no uncertainty in knowledge of the real world system would remain to be investigated with the model (CON 59 p.104).

A number of approaches to validation were identified by Naylor and Finger in their paper (NAY 67). These were: 1)

Rationalism, that views models as built up logically from a set of axioms that are themselves not subject to proof. Verification is here seen as "the problem of searching for a set of basic assumptions underlying the behavior of the system of interest" (NAY 67 p. B-93). 2) Empiricism, where each of the assumptions of a model must be supported by experimental evidence and 3) The methods of Positive Economics where ability to predict correctly is the sole indicator of validity. It is claimed that each approach is by itself sufficient to establish a model's validity, but Naylor and Finger contend that, when operating as we usually are, with a scarcity of information, we cannot completely verify a model with any one method and must use any of the three approaches as appropriate to the information we do have. Thus, we might empirically test a uniform random number generator to see if it produces well distributed numbers, rationally argue that certain mathematical transformations on these uniform numbers will produce random numbers with other desired probability density functions and test an entire model by having it predict a certain case for which we do have data.

A recent survey of validation techniques (LAW 80) develops the philosophies identified by Naylor to catalogue possible tests for validity into three general techniques.

(LAW 80 p.10-14):

- 1) Face validity - how reasonable are the assumptions
- 2) Tests of assumptions
- 3) Tests of results

All three methods may be used to assess the degree of validity of FLAPS. The reasonableness of the assumptions was defended as they were developed in Chapter II. Tests of assumptions were performed for the important landing performance model in section II.D. At a number of points FLAPS uses methods and/or data that have gained general acceptance among airport analysts.

Validation of results poses its own difficulties. Both Naylor and Law point out (NAY 67 p.B-93, LAW 80 p.10,15) that the interpretation of a result of a statistical test on a simulation can be ambiguous. A test which results in a rejection will be taken to mean that there is something wrong with the model. A test which does not result in a rejection is, of course, preferable, and taken as confirmation of the accuracy of the model, but of what is it a confirmation? How much does it enhance our confidence in the model? Obviously answers to these questions will depend on the difficulty of the test, how much is being tested, and the possibility of offsetting errors among other factors. Resolving these questions is a much more subjective process than a statistical test.

There are other problems in testing results of simulations in general and of an airport simulation in particular. We have seen (section II.G) that simulations pose a considerable strain on the classic theory of hypothesis testing. This is exacerbated by the scarcity of reliable data on events at airports. Data on delay are collected on an ongoing basis by a number of airlines and, to some extent, by the FAA. Such data can be helpful in making general assessments of delay conditions but are not sufficiently accurate and detailed to be used to validate a model. Delay data are typically collected by pilots (for airlines) or by controllers (for the FAA). Both groups are very busy with other duties and tend to look upon recording delay as a secondary task. Delays are usually given in multiples of 5 or 15 minutes. Most importantly such data rarely include information on conditions in effect at the airport at the time the data are collected.

The result of all of these problems is that few classic statistical tests will be performed and presented in the sections below. Instead, we will present a number of comparisons involving a variety of models and airports. The intent is to establish that the model described in Chapter II gives reasonable results over a range of conditions, and it provides options and insights that other models cannot match.

B TEST CASE: LA GUARDIA


La Guardia airport in New York (see figure II-1) has been the subject of a number of studies in the past. This section reports a series of comparisons between FLAPS and other reported results on La Guardia. The geometrical information was assembled as described in section II.B. Figure III-1 shows the airport geometry elements for La Guardia used in FLAPS.

1 Comparison to FAA capacity analysis

1.a Introduction

A study has been conducted by the FAA (FAA 77), estimating capacities for four La Guardia runway configurations. The study used a capacity model developed by Peat, Marwick and Mitchell*. Two of the four configurations were run on FLAPS, using the same values of parameters which were reported in the FAA study and assuming standard values for the parameters not reported by the FAA. FLAPS was found to exhibit slightly lower capacities for reasons probably related to A/D separations.

* Different from the airfield simulation model developed by PMM which has been reviewed in this work.




Keypoint	x	y	Keypoint	x	y	Keypoint	x	y
811	37	6	821	90	178	212	130	155
812	170	221	822	314	43			
111	37	6	121	90	178			
112	53	32	122	143	148			
113	77	70	123	179	126			
114	92	92	124	208	108			
115	100	108	125	218	103			
116	110	122	126	248	84			
117	117	136	127	311	46			
118	124	145	128	153	140			
119	147	173						
120	170	221						

FIGURE III-1. La Guardia: FLAPS Geometry Representation.


1.b Case description

The study (FAA 77) gives the aircraft mix and arrival/arrival separations that were used. The values are shown in table III-1. Other separations (A/D, D/D) were not supplied. Values assumed for FLAPS were those presented in section II.D as typical. Aircraft performance information was not specified in the study, however, the aircraft mix percentages are given in terms of four classes and the size of each class is given. They correspond closely to classes 1 to 4 as described in II.D, and the aircraft performance parameters described there were used.


Two configurations were used as shown in figure III-2. The first configuration is a single runway which represents the actual use at La Guardia of three different single runways, as listed in the figure. This case was also used in the test of normality presented in section II.G. Configuration 2 is an intersecting runways case with the intersection located close to the start of the arrival runway. There are two ways in which such a configuration can be used: 1) arrivals on 22, departures on 31; 2) arrivals on 13, departures on 4. Each case was run for both IFR and VFR separations with aircraft mixes varying slightly under IFR and VFR conditions.



Configuration 1



Configuration 2.a



Configuration 2.b

FIGURE III-2. La Guardia Configurations.

Aircraft mix:

Class:	D	C	B	A
VFR:	.03	.72	.15	.10
IFR:	.03	.76	.16	.05

Arrival/Arrival Separations:
(miles at threshold)

IFR:

Trailing Aircraft Class-	D	C	B	A	
Leading Aircraft Class-	D	4.5	6.0	7.3	8.4
	C	3.2	3.2	4.9	6.1
	B	4.2	3.0	3.2	3.5
	A	3.0	3.0	3.2	3.5

VFR:

Leading Aircraft Class-	D	4.1	5.1	5.6	5.0
	C	3.1	3.0	3.9	3.7
	B	2.5	2.3	1.1	1.5
	A	2.5	2.3	1.1	0.8

Data from FAA 77 p. A1-16

Table III-1: La Guardia Validation: Capacity Parameters

1.c Discussion of results

Table III-2 presents the results of the runs. The FAA results, as given in FAA 77, do not distinguish among the three single runway situations or between the two intersecting cases. The FAA results shown in table III-2 are the capacities for 50% arrivals. FLAPS results come from a run in which the runways were saturated with aircraft. Ten replications of four hours each were run, and the results shown are the averages of the last three hours of each replication. Confidence intervals at the 95% level on the FLAPS estimates vary among the cases from plus or minus 1 to 2 aircraft.

	IFR			VFR		
	A	D	T	A	D	T
Configuration 1:						
FAA	26.2	26.2	52.5	28.3	28.3	56.7
FLAPS(AP):	26.8	26.8	53.6	31.0	21.9	52.9
FLAPS(Alt):				28.5	30.3	58.8
Configuration 2:						
FAA	34.1	34.1	68.2	37.6	37.6	75.2
FLAPS(2.a,Rmt=2.0):	26.8	30.9	57.7	31.2	33.6	64.8
FLAPS(2.a,Rmt=1.0):				31.3	38.5	69.8
FLAPS(2.b,Rmt=2.0):	26.9	33.6	60.5	31.3	37.5	68.6
FLAPS(2.b,Rmt=1.0):				31.8	42.9	74.9

FAA results from FAA 77
 FLAPS results from run of 10 replications for 4 hours,
 average of last three hours.

A - Arrivals
 D - Departures
 T - Total Operations
 AP - Arrival Priority
 Alt - Alternating Mode

Table III-2: La Guardia Validation: Capacity results

The single runway results (configuration 1) of the two models are very close for IFR. Under VFR the arrival priority scheme reduces takeoffs more than arrivals are increased. The alternating mode produces a more equal split between arrivals and departures and an overall capacity very close to the FAA results.

FLAPS results for configuration 2 show somewhat lower numbers of operations per hour than the FAA figures. The FAA numbers imply a significantly higher number of arrivals than for configuration 1. However, this appears unrealistic

as the same A/A rules apply to both cases. FLAPS arrival capacities, on the other hand, are the same for all cases. The increment in capacity by going from IFR to VFR is seven aircraft/hour for FLAPS (8 in configuration 2.b), the same as the increment given in the FAA results. As mentioned previously, several parameters are unspecified in the FAA report and, therefore, may not coincide with the corresponding parameter values used in the FLAPS runs. One particularly significant parameter is the minimum distance to the threshold of an arrival when a departure rolls across the threshold. This parameter was termed Rmt in section II.E. Two miles is the standard value for this parameter. However, La Guardia is noted for its tightly run procedures, and this value may well be relaxed. FLAPS applies Rmt, as was discussed in II.E, when the departure crosses the intersection of the runways. If this separation were assumed to apply only when the departure begins to roll or if it were less than 2 miles, it would raise capacity significantly. Results are shown in Table III-2 for the VFR cases with an Rmt of one mile. Note that a value of 1.0 mile applied when the departure crosses the intersection, as in FLAPS, is roughly equivalent to an Rmt of 2.0 miles applied when the departure begins to roll. This is because a departure takes approximately 30 seconds to travel to the intersection in this case, and an arrival will travel about 1 mile on final in 30 seconds. An arrival that was two

miles out on final when the departure begins to roll will be approximately 1 mile away from the runway threshold when the departure crosses in front of it. This single change allows an additional 5 departures per hour under VFR conditions.

It was mentioned that the FAA capacity figures in FAA 77 do not distinguish between the two possible runway alignments for configuration 2. As expected FLAPS produces higher capacity estimates for 2.b than for 2.a. The two cases, while very similar, are not identical since the intersection is closer to the arrival end of the runway in 2.b than in 2.a. This means that an arrival blocks a departure for a shorter period in 2.b than in 2.a. Additionally, the intersection is closer to the departure end of the runway in 2.b than in 2.a. This means that a departure needs a shorter gap (see Section II.E) between arrivals to depart in 2.b than in 2.a. The net effect of the difference is 4 additional departures per hour.

It is important to emphasize that the results for configurations 1, 2a, and 2b were obtained by FLAPS without any modifications to program inputs, once the La Guardia geometry was provided (Figure III-1). The only change in parameters for the two cases (2.a, 2.b) was a change in runway assignment policies to redirect operations. The two cases of configuration 2 could be run on the ASM model as well, but the analyst would be required to recalculate the separations for each case. As discussed in part E.7 in

chapter II, ASM applies separations as a time between threshold crossings. There are no internalized rules (i.e. "hold arrivals until departure crosses the intersection") in ASM, as there are in FLAPS. Thus, to run these cases on ASM the analyst would have to explicitly calculate the time required for each class of arrival to cross the intersection and submit these as inputs. These separations are different for the two cases of configuration 2.

2 Comparison with ASM and MITASIM

2.a Background to the comparison

As part of the validation process that was conducted for the ASM model during the period 1977-78, MITRE/Metrek proposed a series of sensitivity runs to be made, using La Guardia Airport as a test case. A set of 14 runs was suggested using the data base supplied in the PMM users' manual (PMM 77), see KUL 78b, p.2-1 to 4-1. Several of those cases have been run and results reported in KUL 78a and KUL 78b*. The MITASIM airport simulation (NOR 79) was developed, in part, as a way of testing the ASM model on these cases. A number of sensitivity cases were run on MITASIM in late 1978.

* Note that only the cases labeled "revised sensitivity" runs in the cited documents are from the 14 cases.

In this part of the validation we used FLAPS to replicate those sensitivity runs. Comparisons are made among ASM, MITASIM and FLAPS for 3 cases and between MITASIM and FLAPS for 2 others. In several cases ASM's performance was in substantial disagreement with FLAPS and MITASIM, due, we suspect, to partially incorrect separations. FLAPS and MITASIM are in general agreement. MITASIM tends to give slightly lower delay estimates, due to different procedures used to define separations on intersecting runways and different D/D separations.

2.b Case description

Each of the 14 sensitivity cases involves running La Guardia with landings on runway 31 and departures on runway 4. The landing runway has two exits prior to the intersection and one exit after it. Arriving aircraft land and taxi across the departure runway to a gate. Through aircraft spend no time at the gate, but begin an immediate taxi out to the departure runway where they queue for takeoff. The schedule for all cases uses the base schedule method without any lateness distribution being applied before the schedule's use in each replication. The schedule consists of 208 aircraft. Almost all of the aircraft are through aircraft. There are four classes of aircraft, corresponding to those used in the capacity studies in part 1 above, save that the approach speed of the class 4

aircraft is set at 110 kts. As will be seen, all cases have an oversaturated arrival process. All cases are run for four hours and each of the three models used 10 replications.

As the entire set of inputs for the base case has been published elsewhere (KUL 78b Appendix B), we will report only those inputs that vary across the cases of interest or are of special interest in comparing results. Table III-3 lists these parameters. The five cases chosen for analysis are listed below. We follow the numbering of the cases and the names for the separations given in KUL 78a and KUL 78b.

- 1A) Base Case - IFR separations, departure trigger at 12 (Q=12)
- 2A) Group 2 IFR separations, Q=12
- 3A) Group 4 IFR separations, Q=12
- 7A) IFR separations, Q=24
- 8A) IFR separations, Q=1

The three sets of separations (IFR, Group 2 IFR, Group 4 IFR) involve changes in both A/A and D/D separations. They are given in table III-3. The various values of the departure queue trigger refer to the size of the departure queue at which a changeover from arrival priority to alternating operations mode is made. Thus, Case 8A is effectively always in alternating operations mode. This type of procedure was discussed in section II.F above.

The inputs are defined in KUL 78b in terms of runway occupancy time to an exit and specific exit probabilities.

Note, however, that these inputs to ASM and MITASIM are in

Replications: 10

Hours: 4

Closest arrival allowed when departure crosses intersection, Rmt: 0 miles.

Separations (at threshold, minutes):

	A/A (miles)				D/D (minutes)			
	1	2	3	4	1	2	3	4
IFR:								
1	5	6	7	7	1.5	2	2	2
2	4	5	5	5	1	1	1	1
3	4	4	4	4	1	1	1	1
4	4	4	4	4	1	1	1	1
Group 2 IFR:								
1	4	5	6	6	1	1.5	1.5	1
2	3	4	4.5	4.5	1	1	1	1
3	3	3	3	3	1	1	1	1
4	3	3	3	3	1	1	1	1
Group 4 IFR:								
1	3.5	4	5	5	1	1	1	1
2	3	3	3.5	3.5	1	1	1	1
3	3	3	3	3	1	1	1	1
4	3	3	3	3	1	1	1	1

Table III-3: Sensitivity runs: Parameters

fact computed as results in FLAPS. The results were duplicated in FLAPS by adjusting exit locations, deceleration speeds and other parameters until the detailed output of landing statistics showed agreement between FLAPS' performance and the inputs to the other models.

2.c Discussion of results

Summary statistics for the five cases are presented in table III-4. Graphs of the landing and takeoff delay for the cases are given in figures III-3 through III-7.

The base case, 1A, results are very similar among the three models. In each case the ASM takeoff delay results are much lower than the other two models. Based on all available information, it is the opinion of the author that the apparent reason for this difference is the improper specification of A/D separation and some internal problems in the ASM model concerning application of the additional gap in the alternating operations mode. The value of A/D separations involving a class 4 aircraft as the arrival is specified as zero (PMM 77 p. 88, KUL 78b p. B-22). This implies that, if a departure is ready to take off when a class 4 aircraft is on final, then the departure will not be held until the arrival clears the landing runway. This means that ASM is adding about 30 seconds (the arrival runway occupancy time of a class 4 arrival) to the "window" that the departure could use for taking off. Similarly D/A separations with an arrival of class 4 are zero. This implies that an arrival of class 4 may cross the landing runway threshold simultaneously with a departure beginning to roll. Returning to A/D separations, there are other suspect values in ASM. An arrival of class 2 or 3 crossing the threshold can be followed by a departure of any class as




FIGURE III-3. Sensitivity Run 1A: Base Case.




FIGURE III-4. Sensitivity Run 2A: Group 2 Separations.

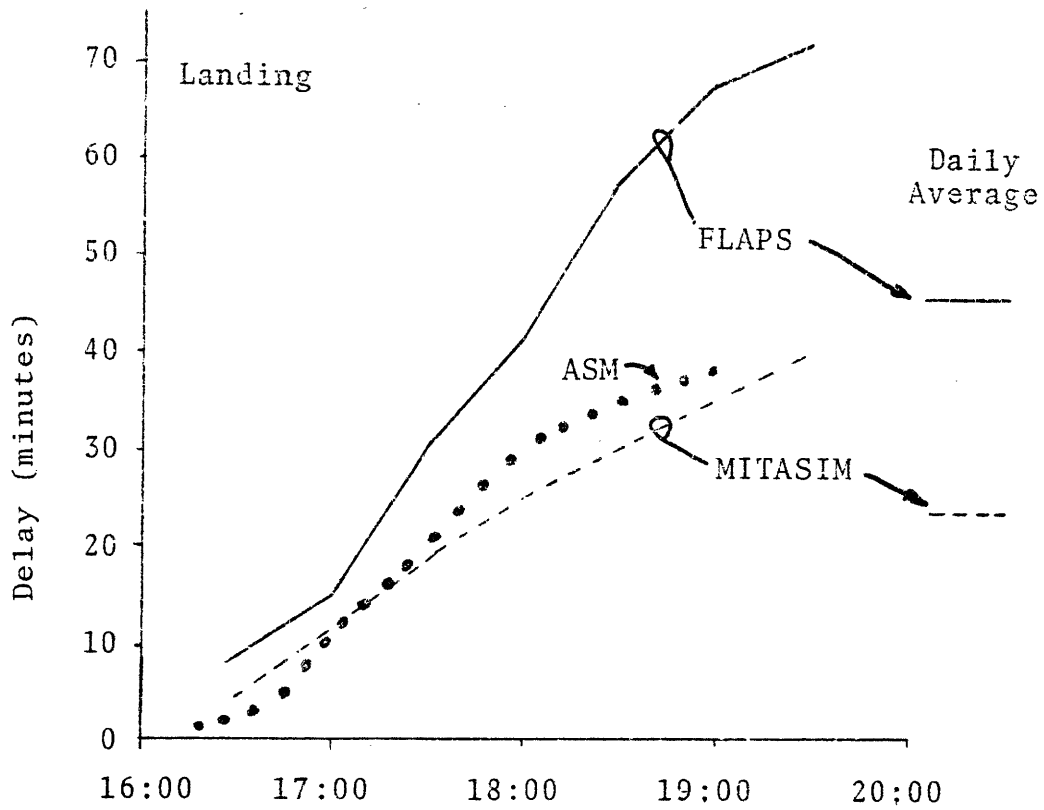



FIGURE III-5. Sensitivity Run 3A: Group 4 Separations.




FIGURE III-6. Sensitivity Case 7A: $Q = 24$




FIGURE III-7. Sensitivity Case 8A: Q = 1.

<u>Case</u>	<u>Model</u>	<u>Landings</u>	<u>Average Landing Delay</u>	<u>Takeoffs</u>	<u>Average Takeoff Delay</u>
1A	FLAPS(2)	88.7	57.4	142.5	12.9
Base	MITASIM	88.8	57.4	142.6	11.3
	ASM	83*		126	
2A	FLAPS	101.0	50.1	139.8	23.8
Group 2	MITASIM	107.8	45.8	155.4	10.2
	ASM	84*		130	
3A	FLAPS	108.8	45.8	132.0	30.5
Group 4	MITASIM	131.0	22.8	155.4	29.8
	ASM	107*		141	
7A	FLAPS	88.6	57.3	142.6	12.4
Q = 24	MITASIM	105.0	47.7	152.2	10.3
8A	FLAPS	88.5	57.3	142.4	13.8
Q = 1	MITASIM	104.9	47.6	152.4	9.6

* ASM results from KUL 78b p. C-1 for first three hours only. No results have been reported for ASM on cases 7A and 8A.

Table III-4: Sensitivity Runs: Summary Statistics

soon as 20 seconds later. This is a much smaller time than the arrival will in fact require to cross the intersection or exit the landing runway. The net effect of these changes is to greatly increase the departure capacity by eliminating a significant degree of interaction between arrivals and takeoffs. This problem was noted in KUL 78b. As both MITASIM and FLAPS have internal logic to hold departures until an arrival either clears the runway or crosses the

intersection (whichever comes first), these errors cannot be duplicated to see if they account for the difference in results.

The disagreement of FLAPS and MITASIM is probably due to several minor differences in the way the models treat separations. FLAPS applies the "worst case" analysis to separations (as described in section II.D) more consistently than does MITASIM. The MITASIM run was also made with D/D separations smaller than the standard case. FLAPS was run once again using the D/D separations used by MITASIM and the result is shown in figure III-3 as FLAPS(2). This is much closer to the MITASIM values in the second half of the run. Despite the disagreement in level of delay, there is excellent agreement between the two in the pattern of results.

The set of cases involving changes in separations (1A: Base, 2A: Group 2, 3A: Group 4) shows considerable differences in the response of the models. Differences in separations may account for a significant fraction of the differences in results. With all models there is a general trend toward reducing landing delay and increasing takeoff delay as A/A and D/D separations are reduced. For average landing delay this is 57.4, 50.1, and 45.8 minutes for FLAPS (case 1A, 2A, and 3A respectively), and for MITASIM 77.7, 45.8, and 22.8 minutes. The trend on takeoff delays is again more consistent (12.9, 23.8, 30.5 minutes) for FLAPS than for MITASIM (10.3, 10.2, 29.8 minutes).

Part of the explanation may lie in the different, smaller D/D separations used by MITASIM. This allows more instances of 2 or more departures between arrivals if interarrival gaps permit. The departure queue length in case 2A is essentially unchanged from case 1A in MITASIM. This is unlike case 3A, where the departure queue length is considerably above 12 for almost the entire run of MITASIM. The gap provided for departures to leave between arrivals is only made long enough for one departure. The subsequent takeoff must therefore wait for the next landing before it can takeoff. As this will be on the order of 1.5 minutes later, D/D separations will not apply in this case. Thus, for case 3, MITASIM and FLAPS agree closely on takeoff delay, despite having different D/D separations. The difference in landing delay is harder to explain, but MITASIM's abrupt reduction in landing delay from case 2A to 3A is harder to explain than FLAPS' more even reduction.

The three queue trigger cases (8A: $Q = 1$, 1A: $Q = 12$, 7A: $Q = 24$) each have the same basic pattern of results. The three models are in close agreement in regard to landing delay. Takeoff delay results vary significantly, but the ordering of the models is constant. ASM always has very low takeoff delays, while MITASIM and FLAPS are in reasonable agreement, with FLAPS showing the higher values. Table III-5 shows the departure queue length results for each half hour in FLAPS and MITASIM. This information is not

available for ASM, but the very low takeoff delays imply that queue lengths were not significant.

The most apparent point is the small impact a change in

Time	8A(Q = 1)		Case			7A(Q = 24)	
	FLAPS	MITASIM	1A(Q = 12)	FLAPS(2)	MITASIM	FLAPS	MITASIM
16:30	10.4	7.0	10.6	10.4	7.2	10.5	7.2
17:00	8.0	1.3	8.0	8.0	1.5	7.7	1.5
17:30	8.9*	2.7	8.7*	7.5	2.7	8.2*	2.7
18:00	15.7*	11.1*	14.5*	13.4*	11.5*	14.1*	11.5*
18:30	13.9*	9.9*	12.6*	11.4*	10.2*	12.3*	10.2*
19:00	11.5*	6.3	10.4*	8.8*	6.7	10.2*	6.7
19:30	5.5	2.3	4.3	2.3	3.6	4.7	3.6
20:00	2.0	0.6	1.7	1.6	0.8	1.9	0.8
C. I.	1.2	1.0	1.4	1.6	1.0	1.5	1.2

C.I. - Average half width of 95% confidence interval on half-hour estimate of queue length.

* Maximum observed queue length equaled or exceeded 12 in one or more replication.

Table III-5: Average Departure Queue Length Statistics

the value of Q has on the results. This is primarily due to the use of large A/A separations (4 miles and above) which means that, most of the time, sufficient space is provided already by the A/A separation to allow one or more departures between landings. Thus, when both FLAPS and MITASIM calculate D/A separations, this does not result in any additional constraint on arrivals. The ASM model procedure of adding a uniform interval to each A/A separation, when the departure queue is above the critical

value, ordinarily results in more responsiveness to changes in the trigger value. However, the suspected mistake in A/D separations probably means that the takeoff queue did not reach either 12 or 24 to trigger the changeover in the ASM model runs. This is speculation, since no results for cases 7A and 8A have been reported for ASM.

C TEST CASE: BROMMA AIRPORT

1 Case Description

1.a Introduction

Bromma airport in Stockholm was the subject of an extensive analysis of delays under current and projected demands in 1979 (ODO 79). FLAPS was compared with the results in ODO 79 for 3 cases. There was very good agreement between the two studies.

1.b Aircraft Separations

Bromma airport has a single runway that is used for both takeoffs and landings. Odoni (in ODO 79) did not employ the standard separations between aircraft classes which have been used for various analyses in this chapter. Nor did Odoni use one set of operating rules for all aircraft. Under VFR weather conditions, there are instead some aircraft which use IFR separation standards and other aircraft which use VFR separation standards. All IFR operations use the same set of separations and all VFR operations use a different set of separations. Operations of the two types are intermixed. All separations apply at the runway threshold and are deterministic, i.e., all buffer widths are assumed zero. Table III-6 gives the complete set of separations. which were modeled in FLAPS by treating

each type of operation as a class. Class I represents IFR

Leading a/c		Trailing a/c		Average separation (sec)
Arriving	Departing	Arriving	Departing	
IFR		IFR VFR		120
			IFR VFR	45 60
	IFR	IFR VFR		105
			IFR VFR	60 120 55
VFR		IFR VFR		45
			IFR VFR	45 45 45
	VFR	IFR VFR		60
			IFR VFR	45 60 45

From ODO 79 p. C-2. Incorporates changes described in text of ODO 79.

Table III-6: Practical Separation Minima at Bromma Airport

movements; class V represents VFR movements. Separations were derived from those listed in table III-6 and are given in table III-7.

To model A/A separations, both classes were given deterministic approach speeds of 120 kts.* Then the A/A separations, specified as times, can be taken directly from

table III-6. They were imposed at the approach gate but,

Arrival/Arrival Separations (seconds):

Trailing Class -	I	V
Leading Class -	I 120	45
	V 45	45

Departure/Departure Separations (seconds):

Trailing Class -	I	V
Leading Class -	I 120	55
	V 60	45

Arrival/Departure Separations (seconds):

			Effective Lead Aircraft
			Class Arrival Runway
			Occupancy Time
Trailing Class -	I	V	
Leading Class -	I 45	60	52.5
	V 45	45	45

Departure/Arrival Separations (seconds):

Trailing Class -	I	V
Leading Class -	I 105	60
	V 60	45

Converted Departure/Arrival Separation (miles):

Trailing Arrival(j) -	I	V
Lead Takeoff(k) -	I 3.5	2
	V 2	1.75

Table III-7: FLAPS Equivalent Separations

because approach speeds are the same for each class, the same separation also hold at the threshold. Departure/Departure separations can be taken directly from table III-6

Time values for A/D separations are taken from the Bromma values and presented in matrix form in table III-7. Since the only A/D separation logic in FLAPS is to hold

* However, any other approach speed could also have been used.

departures until the arrival clears the runway, these separations were simulated by adjusting arrival runway occupancy times (a.r.o.t.). One exit was defined on the runway and a deterministic set of landing aircraft performance parameters was chosen to result in a 45 second a.r.o.t. for class V aircraft. Class I aircraft were given a 52.5 second a.r.o.t., midway between the two separation values.* This is the only approximation to the published separations needed to prepare inputs for FLAPS.

Departure/Arrival separations are given in matrix form in table III-7. Given the 120kt. approach speed, these can be converted directly into miles and used as the values of $R_m(k,j)$, presented at the bottom of the table.

1.c Demand

A number of daily demand profiles were tabulated and reported for Bromma in ODO 79. The "weekday" demand function was chosen for the comparison. A separate demand profile is given for landings and for takeoffs for each of the following: 1) Linjeflyg (LIN) Airline IFR operations; 2) general aviation flights under IFR operations; and 3) general aviation flights under VFR rules. These were converted to meet FLAPS' requirements by the following procedure. Figure III-8 gives a schematic of the process. Fifty per cent of IFR flights were LIN flights, as used in ODO 79. The

* The cases we examine have a 50% VFR, 50% IFR mix, see below.

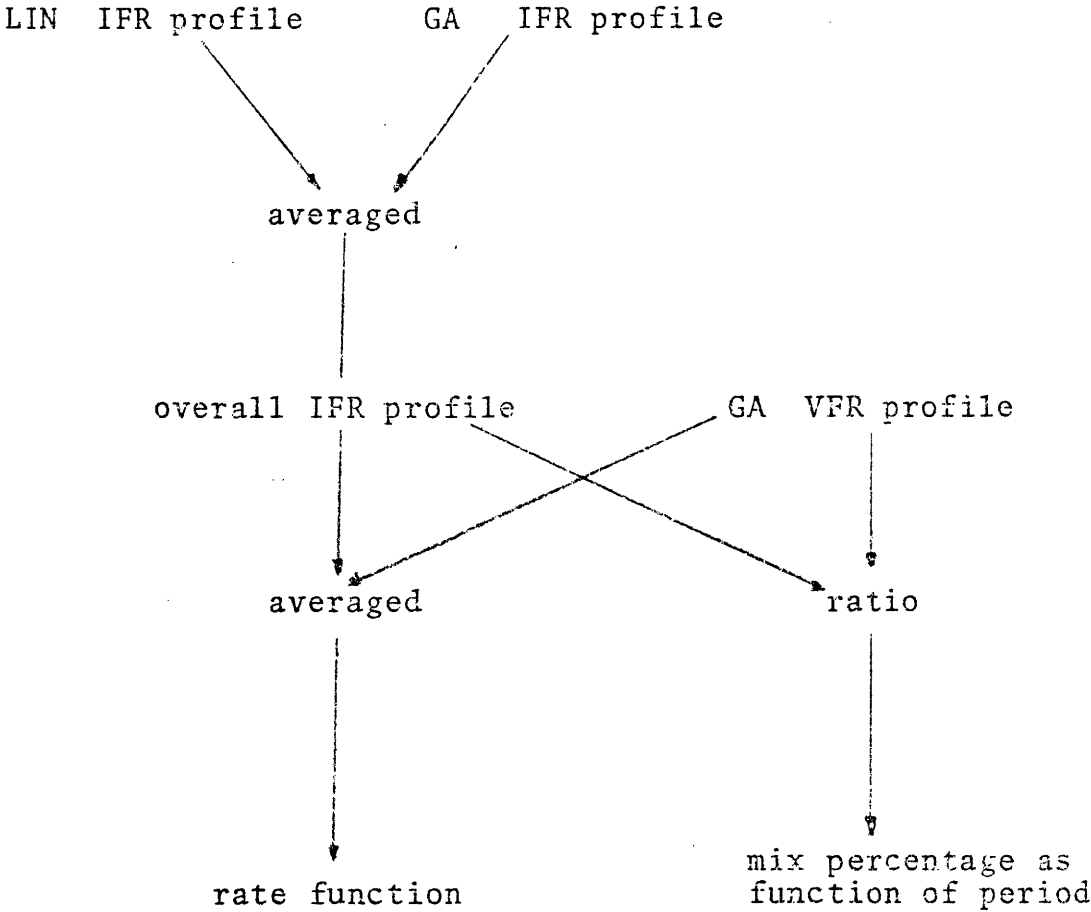


FIGURE III-8. Bromma Demand Profile.

average of (1) and (2) above was taken to obtain an overall IFR profile for both landings and takeoffs. ODO 79 reports delays for a number of different IFR/VFR splits. This example used 50⁺ IFR flights. The overall IFR profile and (3) above were averaged to obtain an arrival rate function and a departure rate function. These rate functions were used to control a generated schedule procedure (see section II.C) with all aircraft as either arrivals or departures (i.e., no through aircraft). Interaircraft times were distributed negative exponentially. For both arrivals and departures, a mix percentage was created as a function of the period of time involved. The mix percentage for class V was taken from the ratio of VFR operations to total operations for each period. The resulting demand profiles are given in figure III-9. The rates in the figure are given as "fraction of landings (takeoffs) per hour". This fraction is multiplied by total landings (takeoffs) desired over the entire run to obtain arrival (departure) rates per hour. Three cases were run, at 400 operations/day (200 landings, 200 takeoffs), at 600 operations/day (300 landings, 300 takeoffs), and at 800 operations/day (400 landings, 400 takeoffs).





FIGURE III-9. Bromma: Demand Functions

2 Discussion of results

The delay estimates in ODO 79 are not produced through a simulation. They come from an analytical model, DELAYS, that treats the runway as an M/G/1 queue. The state equations for this queue are solved using a time-dependent demand rate. This produces an estimate of the waiting time in the queue (delay), and other queueing statistics, for each time interval during the day. Explanation of the theoretical background for this model can be found in ODO 79 Appendix C and HEN 75.

Results for the three cases run on Bromma are given in

Model	Delay (minutes) mean 95% C. I.		Fraction of aircraft Delayed > 20 minutes
400 Operations/day:			
FLAPS	0.7 0.1		0.00
DELAYS	0.6		0.000
600 Operations/day:			
FLAPS	2.6 0.6		0.01
DELAYS	2.3		0.012
800 Operations/day:			
FLAPS	15.0 2.8		0.27
DELAYS	14.7		0.315

DELAY results from ODO 79, Appendix E

Table III-8: Bromma: Summary Comparison of Results

table III-8. Comparisons of the pattern of delay produced during the day are given in figures III-10 to III-12. The models seem to be in good agreement. This is particularly significant in view of the fundamentally different nature of the two models. One difference is the timing of the peak delay. The FLAPS peak delay occurs later than DELAYS, particularly on the 800 operations/day case. This is probably due to FLAPS collecting delay statistics when aircraft complete landings or takeoffs. DELAYS estimates are for delay expected as of the moment an aircraft arrives. Thus, the difference is more significant at higher average delays.




FIGURE III-10. Bromma Airport: 400 op./day




FIGURE III-11. Bromma Airport: 600 op./day




FIGURE III-12. Bromma Airport: 800 op./day

D CONCLUSIONS

Several general points can be made about the tests presented in this chapter. FLAPS has been used to estimate both capacities and delays and has produced reasonable estimates in all situations. Comparisons have been made using data from two airports and against several different models. We have shown that FLAPS can be used successfully to analyze questions of interest to airport analysts. Together with the material in chapter II, the results presented here establish considerable confidence in the validity of FLAPS, although further tests should most definitely be conducted in the future.

IV EXAMPLE OF MODEL APPLICATION

A INTRODUCTION

This chapter* will demonstrate the capability of FLAPS to analyze significant questions concerning airport operations. This will be done by performing a number of capacity and delay analyses for Logan Airport in Boston, Massachusetts. The process of assembling input data for FLAPS will be described. Capacity estimates for two runway configurations under several weather conditions will be presented. A demand rate function will be hypothesized for a third configuration. Delay estimates for this configuration will be obtained for several runs that demonstrate the dynamic capabilities of FLAPS. By demonstrating that FLAPS can analyze the questions that arise in performing a comprehensive analysis, this chapter contributes to the certification of the model.

A second objective of this chapter is to discuss a number of issues connected with the use of the model. Several definitions of capacity are used in airport analysis. We will discuss the appropriate techniques for obtaining these estimates with a simulation. The relative sensitivity of the model to certain parameters, its input

* The author wishes to thank William Hoffman and Steven Aschkenase of Flight Transportation Associates, Cambridge, Massachusetts, for assistance in preparing this particular case study.

requirements and its computational efficiency will also be discussed briefly.

It should also be emphasized that this chapter is not intended as a comprehensive analysis of operations at Logan. Performing a full analysis would be a major undertaking. It would involve identification of a complete set of configurations as well as collection of data and numerous runs of the model, all beyond the scope of this work. Rather, the intent here is to establish that FLAPS could be used for such a comprehensive analysis. This chapter is not offered as further evidence of the model's validity, and no comparisons to existing models are made. Sufficient capacity and delay data were not available for Logan to perform a meaningful validation.

B THE AIRPORT

Boston's Logan Airport is one of the busiest in the United States. In 1979 more than 15 million passengers and 226,000 metric tons of cargo were served by 319,000 aircraft movements (landings and takeoffs) (ATA 80 p.5).

The airport has been the subject of several capacity and delay studies. Many of the specific assumptions used in this analysis come from OSE 78, a report by the Office of Systems Engineering Management of the FAA. The report will be referred to as the OSEM study.




FIGURE IV-1. Boston Logan Airport.

Figure IV-1 is a map of Boston Logan Airport. The airport is surrounded on three sides by Boston Harbor, making physical expansion difficult. Residential areas are under the flight paths of all runways save the approach to 33L, making noise reduction a major area of concern. Runways 4R, 15L, 22L, 27 and 33L are equipped with instrument landing systems.

The geometry information was collected as described in section II.B. An arbitrary origin due south of the southern end of runway 4L was chosen so that all keypoints assumed positive coordinate values. The scale of the map used and the scale of the grid used to overlay the map resulted in a scale of 42.5 feet per coordinate unit. It was found that, with careful use, a map of the scale of figure IV-1 was sufficient to define geometry features to within 50 to 75 feet of published values. Modules for runways 4L/22R, 4R/22L and 9/27 were defined. Runways 13R/33L and 13L/33R were not analyzed in this case study and were not included in the model. The keypoints and coordinates are listed in table IV-1.

A number of Logan runways have displaced thresholds. The split runway technique (see part E.6 of chapter II) was used to model this feature for runway 4R/22L. When the runway is used in the 4R direction, the landing threshold is at the point marked "L4R" in figure IV-1. Takeoffs, however, may begin farther back at point "T4R". In the

<u>Keypoint number</u>	<u>X coordinate</u>	<u>Y coordinate</u>	<u>Purpose of keypoint</u>
811	0	80	runway end point
812	99	216	runway end point
110	0	80	exit
111	6	88	exit
112	30	120	exit
113	42	138	exit
114	58	159	exit
115	81	190	exit
116	100	217	exit
821	18	47	runway end point
822	124	189	runway end point
121	12	38	exit
122	38	74	exit
123	48	87	exit
124	58	98	exit
125	64	108	exit
126	76	123	exit
127	98	154	exit
128	134	200	exit
831	7	75	runway end point
832	171	54	exit
131	7	65	exit
132	65	61	exit
133	97	59	exit
134	150	56	exit
135	171	54	exit
223	31	63	intersection

Table IV-1: Logan Airport: Geometry Elements

opposite direction, this runway, 22L, also has a displaced threshold. The landing threshold for 22L is designated "L22L" and the takeoff threshold "T22L". Two runway modules were defined for the physical runway 4R/22L. The first of the two is used for landing only and extends from L4R to L22L. The second module is used for takeoffs only and extends from T4R to T22L. The runway modules have the same

centerline, with the takeoff runway module overlapping the ends of the landing runway module. All actual runway exits on 4R/22L were defined as being on the landing runway module; the takeoff runway module needs no exits. Both of these modules intersect with 9/27. To insure proper operation of the split runway modules as a single runway, they are specified as dependent parallels. Departures on the takeoff "runway" (T4R/T22L) must clear arrivals on the landing runway (L4R/L22L) for this split runway case. Of course, the two modules must always be operated in the same direction.

Table IV-2 lists the three runways in use and the keypoint numbers and distances for each exit and intersection. These distances are derived by the model from the keypoint coordinates given in Table IV-1. Values in parentheses are the corresponding published values of these runway features when available.

A final approach path of 6 nautical miles (6.9 statute miles) was used (OSE 78 p. 11). Exit speeds were estimated from visual examination of the airport map. The resulting map of keypoints and modules is shown in figure IV-2.

C DESCRIPTION OF OPERATING POLICIES

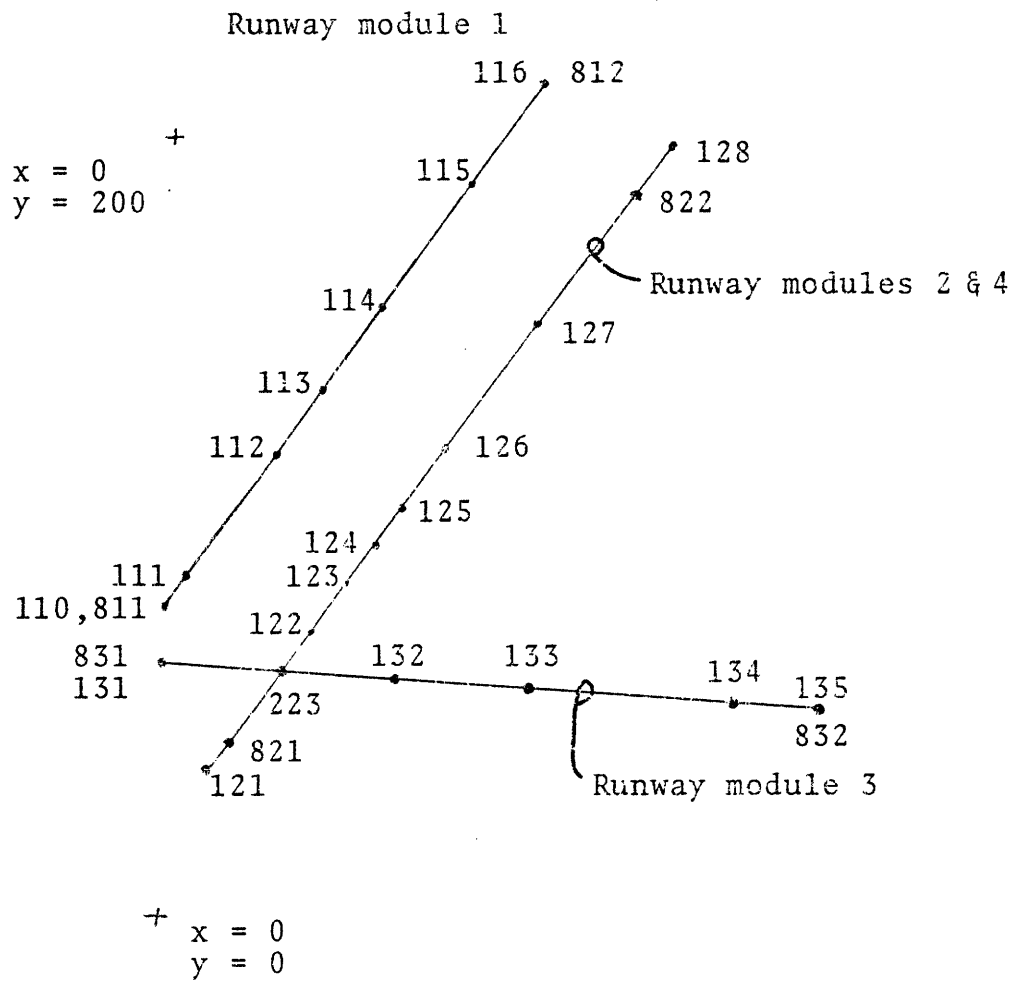
We will describe the analysis of the capacity of one runway configuration under a variety of weather conditions

Runway:	4L/22R	4R/22L	9/27
Runway Module Numbers:	1	2 and 4	3
Length:	7149 (7032)	7531 (7494)	6986 (7000)
Length of common approach path:	6	6	6
Direction 1 is	4L	4R	9
Direction 2 is	22R	22L	27

Runway	Keypoint	Exit parameters			
		Direction 1 Length	Velocity	Direction 2 Length	Velocity
4L/22R	110	0 (0)	N. U.	7149 (7032)	0
	111	425	N. U.	6724	0
	112	2125	0	5025	0
	113	3043	0	4106	0
	114	4165	0 + 2	2984	15
	115	5806	15	1344	0 + 2
	116	7209	0	60	N. U.
4R/22L	121	-460	N. U.	7990	0 + 2
	122	1428	0 + 2	6103	5
	123	2125	0 + 2	5406	30
	124	2755	0	4778	0
	125	3247	15	4284	N. U.
	126	4063	N. U.	3468	15
	127	5678	15	1853	0 + 2
	128	8160	15	-632	N. U.
	223	876	N. A.	6656	N. A.
	9/27	131	0 (0)	N. U.	6986 (7000)
132		2471 (2500)	0 + 2	4515 (4500)	30
133		3833	0 + 2	3152	10
134		6090	0 + 2	897	10
135		6986 (7000)	0	0 (0)	N. U.
223		1024	N. A.	5962	N. A.

Numbers in parentheses are published distances.
 N. U. - Exit not used for traffic in that direction.
 Exit velocity in knots.
 Number after plus sign is number of seconds aircraft needs to clear runway after coming to complete stop.
 N. A. - Not applicable.

Table IV-2: Logan Airport: Runway Modules



See Table IV-1 for coordinates of keypoints.
 See Table IV-2 for exit distances and velocities.

FIGURE IV-2. Logan Airport: FLAPS Geometry Representation.

and of a second runway configuration under a single weather condition. The two configurations chosen are among the most complicated in use at Logan Airport. We will briefly describe them here. In following sections additional details for each configuration will be presented.

The first configuration is shown schematically in figure IV-3, parts a, b, and c.

Case a

This is used under conditions labeled VFR-1 (cloud ceiling 2500 feet or higher, visibility 5 miles or more). Aircraft land on 4R and 4L, takeoff on 4R, 4L and 9.

Case b


This is used under conditions labeled VFR-2/IFR-1 (ceiling 800 to 2500 feet, visibility 2 to 5 miles). Landings approach in a single queue headed for 4R. Once below the cloud ceiling (i.e. in view of the runway) some aircraft execute a "sidestep" to land on 4L. The process for modeling this procedure will be discussed below. Takeoffs use 4R, 4L and 9.

Case c


This is used in IFR-2 (ceiling 200 to 800 feet, visibility 1/2 to 2 miles). Landings are limited to 4R. Takeoffs can use 4R, 4L and 9.

Case d


Case d uses the same runways as case c but has a different aircraft mix. Case d is used in IFR-3 conditions



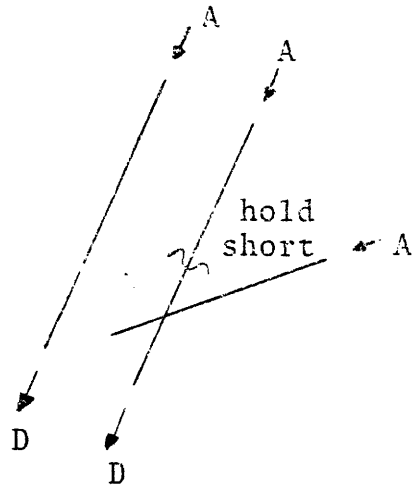
a) Case a



b) Case b



c) Case c & d



d) Case e

FIGURE IV-3. Logan Configurations.

when the ceiling is less than 200 feet or visibility is below 1/2 mile.

The second configuration, case e, is shown schematically in part d of the figure. Here the runways are operated in the opposite direction. Case e is used only under VFR-1 conditions. Landings use 22L, 22R and 27. Landings on 22R must be able to hold short of 27. This process will be discussed below. Takeoffs use 22R and 22L.

Each of the five cases has different aircraft mixes and separation rules. These are discussed under the appropriate section below.

D CHARACTERISTICS OF DEMAND

1 Aircraft classes

Four aircraft classes are identified in the OSEM report. For this analysis this mix was altered by breaking the class for twin-engine propeller aircraft into three subclasses. This was done because of the considerable diversity in the performance characteristics of this class of airplanes. The resulting 6 classes are:

- Class A - Single-engine piston, Gross Takeoff Weight (GTW) less than 12,500 lb.
- Class B1 - Twin-engine piston, GTW less than 12,500 lb.
- Class B2 - Twin-engine turbo prop, GTW less than 12,500 lb.
- Class B3 - Twin-engine turbo jets,

GTW less than 20,000 lb.

Class C - Four-engine propeller and non-heavy jet between 12,500 lbs. and 300,000 lbs.

Class D - Heavy jet capable of 300,000 lbs. or greater takeoff weight.

2 Aircraft Performance

Table IV-3 lists the assumed aircraft performance characteristics for Logan. The starred entries in the table come directly from the OSEM report. The remaining entries have the same values as discussed in the landing performance chapter. Approach speeds were supplied by Flight Transportation Associates (FTA). OSE 78 provides data for arrival runway occupancy times. After the outputs of the model were examined with an initial set of parameters, the deceleration rate and coasting speed were adjusted to the values shown in table IV-3 to produce a closer match with the OSEM data on arrival runway occupancy time. The resulting parameters values are still within the reasonable range. This process is discussed further in section F below.

3 Aircraft Mix

The five cases use four different aircraft mixes. The overall percentage of each class changes because most small aircraft cannot fly during bad weather, due to the

Attribute	A	Aircraft Class			C	D
		B1	B2	B3		
Approach speed (knots):						
mean, Va	80	100	115	130	130	140
st. dev., Vas	6	4	4	4	4	4
p. d. f.			Triangle			
In-air deceleration (knots):	14	12	10	8	8	8
Float distance (feet):						
mean	700	1000	1200	1400	1500	1500
st. dev.	250	300	400	450	450	450
p. d. f.			Uniform			
Deceleration rate (f/s/s):						
mean	5.75	5.75	5.75	5.75	5.75	5.75
st. dev.	.75	.75	.75	.75	.75	.75
p. d. f.			Normal			
Runway coasting speed (knots):			25			
Departure runway occupancy time (seconds):						
mean*	29.	32.	34.	36.	39.	39.
st. dev.	5.	5.	5.	5.	5.	5.
p. d. f.			Normal			

Starred attributes are from OSE 78.
 st. dev. - standard deviation
 p. d. f. - probability density function

Table IV-3: Assumed Aircraft Performance Parameters

requirement of instrument ratings for the pilot and of sophisticated navigational equipment for the aircraft. The OSEM study assigns aircraft to runways on basis of the class of aircraft (OSE 78 pp. 14-17). The corresponding assignment by class procedure in FLAPS was used for the capacity runs. As the set of active runways varies over the

five cases, we need aircraft assignments to runways for each case. These are also available in the OSEM report.

Since 1978, when the OSEM report was prepared, conditions at Logan have changed somewhat. Flight Transportation Associates has made available revised aircraft mix percentages for the 6 classes used here. Table IV-4 presents these percentages. The overall mix was also provided by FTA. The runway assignment percentages (to the left of the vertical lines) come from OSE 78. As will be seen in the section on "discussion of results", these runway assignments create problems in performing a capacity analysis. This is because they do not correspond to the traffic assignment that results in the highest airport capacity. We will also run the "shortest queue" assignment procedure for these cases and compare the results. The shortest queue assignment was used to allow aircraft to use any runway where the OSEM assignment percentage was non-zero. The assignment for any particular aircraft is made to the shortest queue at the time the aircraft is ready to use the runway.

E AIR TRAFFIC CONTROL RULES

A number of air traffic control parameters must be specified. IFR rules are used for cases c and d. VFR rules are used for cases a and e. An intermediate type of A/A

Class	Arrivals				Departures			
	4R	4L	9	Mix	4R	4L	9	Mix
Case a:								
A	.00	1.00	.00	.06	.00	1.00	.00	.06
B1	.17	.83	.00	.22	.00	1.00	.00	.22
B2	.17	.83	.00	.15	.00	1.00	.00	.15
B3	.17	.83	.00	.03	.00	1.00	.00	.03
C	1.00	.00	.00	.44	.41	.00	.59	.44
D	1.00	.00	.00	.10	.66	.00	.34	.10
Case b:								
A	.00	1.00	.00	.03	.00	1.00	.00	.03
B1	.17	.83	.00	.21	.00	1.00	.00	.21
B2	.17	.83	.00	.16	.00	1.00	.00	.16
B3	.17	.83	.00	.03	.00	1.00	.00	.03
C	1.00	.00	.00	.46	.41	.00	.59	.46
D	1.00	.00	.00	.11	.66	.00	.34	.11
Case c:								
A		none		.00		none		.00
B1	1.00	.00	.00	.15	.00	.80	.20	.15
B2	1.00	.00	.00	.17	.00	.80	.20	.17
B3	1.00	.00	.00	.04	.00	.80	.20	.04
C	1.00	.00	.00	.52	.53	.00	.47	.52
D	1.00	.00	.00	.12	.85	.00	.15	.12
Case d:								
A		none		.00		none		.00
B1	1.00	.00	.00	.02	.00	.53	.47	.02
B2	1.00	.00	.00	.04	.00	.53	.47	.04
B3	1.00	.00	.00	.04	.00	.53	.47	.04
C	1.00	.00	.00	.70	.42	.00	.58	.70
D	1.00	.00	.00	.20	.82	.00	.18	.20
	22L	22R	27		22L	22R	27	
Case e:								
A	.00	1.00	.00	.06	.00	1.00	.00	.06
B1	.10	.78	.12	.22	.00	1.00	.00	.22
B2	.10	.78	.12	.15	.00	1.00	.00	.15
B3	.10	.78	.12	.03	.00	1.00	.00	.03
C	.60	.00	.40	.44	.17	.83	.00	.44
D	.75	.00	.25	.10	.83	.17	.00	.10

Runway assignment percentages (to left of vertical lines) from OSE 78 p. 14-17.

Aircraft mix percentages (to right of vertical lines) provided by FTA.

Table IV-4: Logan Airport: Aircraft Mix and Runway Assignment

separation is used in case b. A/A and D/D separations for all cases are from OSE 78 and are shown in table IV-5. The values used are very similar to the ones given in tables II-8 and II-9 earlier. For mixed operations, no separations are given in the OSEM study, so we assume 2 miles for

	Arrival/Arrival (miles)				Departure/Departure (miles)			
IFR rules (cases c and d):								
Trailing class --	D	C	B	A	D	C	B	A
Leading class: D	4	5	6	6	90	120	120	120
C	3	3	4	4	60	60	60	60
B	3	3	3	3	60	60	60	60
A	3	3	3	3	60	60	60	60

Intermediate rules (case b):

Trailing class --	D	C	B	A	D	C	B	A
Leading class: D	3.2	4.3	5.2	5.2	90	120	120	120
C	2.5	2.5	3.2	3.2	60	60	60	60
B	2.5	2.5	2.5	2.5	60	60	60	60
A	2.5	2.5	2.5	2.5	60	60	60	60

VFR rules (cases a and e):

Trailing class --	D	C	B	A	D	C	B	A
Leading class: D	2.7	3.6	4.5	4.5	90	120	120	120
C	1.9	1.9	2.7	2.7	60	60	45	45
B	1.9	1.9	1.9	1.9	50	45	40	35
A	1.9	1.9	1.9	1.9	50	40	30	30

Data from OSE 78 p. 12.

B - represents classes B1, B2 and B3

Table IV-5: Logan Airport: Separation Parameters

$R_m(k, j)$ and $R_{mt}(k, j)$.

The extent to which operations on various runways are interdependent varies over the 5 cases. For each individual runway, the separation procedures used in the example runs are a straightforward application of the procedures given in section D of chapter II and will not be described here. The following discussion concerns only separation procedures for dependent runways. Table IV-6 summarizes these procedures and indicates which separations are applied for the various cases.

Arrival/Arrival separations are applied 1) between 4L and 4R only in the VFR-2/IFR-1 case b, and 2) in the conditional hold short case e, between 27 and 22L. Arrivals in case b use a "side-step" procedure. Aircraft approach in one line to 4R. Near the runway certain class A and class B landings execute a side-step and land on 4L. (Class C and Class D aircraft at Logan are not allowed to use runway 4L for landings due to noise considerations.) This procedure is modeled by assigning landing aircraft to both 4L and 4R, in accordance with the mix percentages indicated in Table IV-4. The values of A/A separations imposed between landings on the two different runways are the same as those used between landings on the same runway. In case e, runways 27 and 22L operate in a conditional hold short mode. (The example of this "hold short" procedure, discussed in part 6 of section II.E, was, in fact, the situation at Logan in case e.) This situation is modeled, as was suggested, by

Separation	a	b	Case c	d	e
A/A					
2/1	No	Yes	N.A.	N.A.	No
1/2	No	Yes	N.A.	N.A.	No
3/2	N.A.	N.A.	N.A.	N.A.	Yes(a)
2/3	N.A.	N.A.	N.A.	N.A.	Yes(a)
D/D					
1/4	No	Yes	Yes	Yes	No
2/1	No	Yes	Yes	Yes	No
3/4	Yes	Yes	Yes	Yes	N.A.
4/3	Yes	Yes	Yes	Yes	Yes
A/D					
1/4	No	Yes	Yes	Yes	No
2/1	No	Yes	Yes	Yes	No
2/3	Yes	Yes	Yes	Yes	N.A.
3/4	N.A.	N.A.	N.A.	N.A.	Yes
2/4	Yes	Yes	Yes	Yes	Yes

- Yes - Movements on the indicated runways are checked for nonviolation of the given separation requirements.
- Yes(a) - Conditional hold short separation, see text.
- No - The movements are independent.
- N.A. - Not Applicable.
- x/y - Runway modules of leading, x, and following, y, operation. Note that modules 2 and 4 both refer to runway 4R/22L.

Table IV-6: Air Traffic Control Dependencies

holding all landings on runway 27 for 90 seconds after a heavy aircraft (class D) has begun landing on runway 22L, and by preventing a heavy from landing on runway 22L for the first 90 seconds after a landing has begun on runway 27. Other separation values, i.e., for: 1) a non-heavy on 22L following all landings on 27, and 2) all landings on 27 following a non-heavy on 22L, are set to zero. Note that

these separations apply to the module used for landing on 22L, not the runway module used for takeoffs on 22L.

Departure/Departure separations are applied between the parallel runways in the two VFR cases (a and e) and in the VFR-2/IFR-1 case b. As runways 9/27 and 4R/22L intersect, D/D separations between these intersecting runways are always applied.

Departures on one of the parallel runways must be held for arrivals to clear on the other runway in the IFR cases only. Departures on either 9/27 or 4R/22L must be held until arrivals on the other runway clear the intersection. The line of A/D separations in table IV-6 marked "2/4" is included to make the two runway modules for 4R/22L operate as a single runway.

F INPUT REQUIREMENTS

1 Summary of Input Requirements

This section summarizes the various kinds of inputs necessary for preparing a case for FLAPS. Figure IV-4 is a copy of the entire input file for the VFR capacity cases. For each item below we refer, where appropriate, to tables that present the specific values of the input. Letters in brackets after an item below refer to the portion of the input file in Figure IV-4 where the same item is defined. A

number of parameters with internal defaults were not respecified in the input file.

A complete input file for FLAPS contains the following categories of data:

- 1) Simulation control parameters
length of run, length of periods, number of replications, starting time, etc. {A}
- 2) Output control parameters {B}
- 3) Geometry parameters
 - a. Module definitions {Figure IV-2, Table IV-2, C}
 - b. Keypoint definitions {Table IV-1, D}
- 4) Aircraft attributes {Table IV-4, E}
- 5) Demand generation parameters {Table IV-4, F}
- 6) Aircraft separations
 - a. Types of dependencies {Table IV-6, G}
 - b. Values of separations {Table IV-5, H}
- 7) Random number seeds {default values used}
- 8) Runway assignment parameters {I}
- 9) Dynamic operating rule parameters {J}
Controls changes in operating rules during a run.

2 Comparison of Input Requirements

Some indication of the performance of FLAPS with regard

```

PARAM: MKPER=6, MKREP=10, PERLEN=60.0, STTIME=0.0,           {A}
IREPORT(1)=1, IREPORT(2)=0, IREPORT(3)=0,                   {B}
IREPORT(4)=0, IREPORT(5)=0, IREPORT(6)=0,
IRERESULT(1)=1, IRESULT(2)=1, IRESULT(3)=1,
IRERESULT(4)=1, IRESULT(5)=1,
IDCNTR(1)=2, IDCNTR(2)=0, IDCNTR(4)=1, IDCNTR(5)=0,
ATC.IA2A1(2,3)=2, ATC.IA2A1(3,2)=3,                         {G}
ATC.IA2D1(4,3)=1, ATC.IA2D1(3,2)=1,
ATC.ID2A1(3,4)=2,
ATC.ID2D1(4,3)=1, ATC.ID2D1(3,4)=1,
ATC.IA2D1(4,2)=1, ATC.ID2A1(2,4)=2,
ATC.ITYPE=2, ATC.IDIROP(1)=1, ATC.IDIROP(2)=1,
ATC.IDIROP(3)=1, ATC.IDIROP(4)=1,
ATC.MDATC(1)=2, ATC.MDATC(2)=2,
ATC.MDATC(3)=2, ATC.MDATC(4)=2,
TITLE='LOGAN AIRPORT CAPACITY ANALYSIS CASE A';
ATC:
ACTIVE: 1 1 1 1
ASSIGNRW: 2 2                                               {I}
ASSIGN: 2 A 0 .75 1 1 0 .6 1 1 3(.73 .83 1 1)
          1 1 1 1 *
ASSIGN: 2 D .167 .167 .167 1 .83 .83 .83 1
          4(1 1 1 1) *
ASSIGN: 4 A 2(0 1 1 0) 3(1 1 1 0) 1 0 0 0 *
ASSIGN: 4 D 2(1 0 0 1) 4(1 0 0 0) *
ASSIGN: 2 A 2(0 1 1 1) 3(.83 1 1 1) 1 1 1 1 *
ASSIGN: 2 D 0 0 .34 1 0 0 .59 1 4(1 1 1 1) *
ASSIGN: 4 A 2(0 1 0 0) 3(1 1 0 0) 1 0 0 0 *
ASSIGN: 4 D 2(0 0 1 1) 4(1 0 0 0) *
ORDER: 0:00 IDCNTR(1)=0;                                     {J}
ORDER: 0:00 ATC.MDASGN(1)=2, ATC.MDASGN(2)=2;
ORDER: 1:30 IDCNTR(1)=2;
ORDER: 1:50 IDCNTR(1)=0;
ORDER: 3:00 ATC.MDASGN(1)=4, ATC.MDASGN(2)=4;
END

```

FIGURE IV-4. Logan VFR Capacity Input File.
(continued on next page)

```

AIRCRAFT:
APPRSPEED: 4 2 140 4 130 4 130 4 115 4
           100 4 80 6 *
FLOATDV: 1 1 8 1 8 1 8 1 10 1
           12 1 14 1 *
FLOATDIS: 7 2 1500 450 1500 450 1400 450 1200 400
           1000 300 700 250 *
LANDDECR: 2 3 6(5.75 .75) *
VCOAST: 1 1 6(25 0) *
SKDTIMEGATE: 3 2 55 20 45 15 3(35 15) 60 15 *
MINTIMEGATE: 1 1 5(30 10) 45 15 *
TAKEOFFROLL: 4 1 39 5 39 5 36 5 34 5 32 5 29 5 *
END
SCHEDULE:
ARATE 1 1 1.0 85 3*60 21*65
DRATE 1 1 1.0 85 3*60 4*60 17*50
APRONAREA: 1 1 1 1 1 1 1 1 1 *
ARIVFIX: 1 1 1 1 1 *
DEPIFIX: 1 1 1 1 1 *
AIRCRAFTMIX:
ARRIVE .10 .54 .57 .72 .94 5*1 *
DEPART .10 .54 .57 .72 .94 5*1 *
SCHEDPARM: 'GENER' 'NEW' 0.0
END
SEPARATION:
AA: 1 MILES 18. 1.65 2 2
4 5 6 6 6 6 3 3 4 4 4 4 24*3 *
AA: 2 MILES 18. 1.65 2 2
2.7 3.6 4.5 4.5 4.5 4.5 1.9 1.9 4*2.7 24*1.9 *
AA: 3 MILES 18. 1.65 2 2
3.2 4.3 3*5.2 5.2 2.5 2.5 4*3.2 24*2.5 *
DA: 1 MILES 0. 1.65 2 2
6*2. 6*2 6*2 6*2 6*2 6*2.0 *
XDA: 1 MILES 18. 1.65 3 2 4.0 5*6.0 30*4.0 *
XDA: 2 MILES 18. 1.65 3 2 4.0 5*4.5 30*4.0 *
DD: 1 MINUTES 0. 1.65 1 3 1.5 5*2 30*1.00 *
DD: 2 MINUTES 0. 1.65 1 3 1.5 5*2
1.0 1.0 0.75 0.75 0.75 0.75
3(.833 .75 .667 .667 .667 .5333) .833 .667 4*0.5 *
D2: 1 MINUTES 0. 1.65 1 3 1.5 5*2 30*1.0 *
COPY1INTO2
TAA: 1 MINUTES 18. 1.65 3 2 6(1.5 0 0 0 0 0) *
TAA: 3 MINUTES 18. 1.65 3 2 6*1.5 30*0 *
END

```

FIGURE IV-4, continued. Logan VFR Capacity Input File.
(continued on next page)

```

MODULE:
RUN3 1 811 812 EXIT: 110 111 112 113 114 115 116          {C}
      VEXIT: -20 0 -20 0 0 0
      0 0 -2 15 15 2 0 -20 FAPDIS: 6. 6. *
RUN3 2 821 822 EXIT: 121 122 123 124 125 126 127 128
      VEXIT: -20 -2 -2 5 -2 30 0 0 15 -20 -20 15
      15 -2 15 -20 FAPDIS: 6. 6. INTR: 223 224 *
RUN3 3 831 832 EXIT: 131 132 133 134 135 VEXIT: -20 5
      -2 30 -2 10 -2 10 0 -20 INTR: 223 FAPDIS: 6 6 *
RUN3 4 121 128 INTR: 224 *
END
KEYPOINT: 42.5          {D}
811 0 30 821 13 47 831 7 65
223 31 63 224 31 63
812 99 216 822 124 189 832 171 54
110 0 30
111 6 33 121 12 38 131 7 65
112 30 120 122 38 74 132 65 61
113 42 138 123 48 37 133 97 59
114 53 159 124 58 98 134 150 56
115 81 190 125 64 108 135 171 54
116 100 217 126 76 123
127 93 154 128 134 200 9999
END

```

FIGURE IV-4, continued. Logan VFR Capacity Input File.

to the workload that it imposes on its users can be obtained by comparing the input requirements of FLAPS with those of ASM. Unfortunately, such a comparison is made difficult by at least two factors. First, the models do not have equivalent capabilities. FLAPS does not model ground movements, and this is a major portion of the ASM input requirements (see section II.B). On the other hand, FLAPS has added capabilities in the area of demand generation, specification of dependent runway separations, and specification of changes in operating rules, all of which require additional inputs. Second, we do not have input files for the two models for identical cases.

Despite these cautionary points, it is rather obvious that setting up a case for analysis using FLAPS is a much simpler task than setting up a case for ASM. We saw earlier (section II.C) that several hundred or even several thousand lines could be required to prepare a case for ASM. The 100 line file of FLAPS for a moderately complex airport compares very favorable with that.

G CAPACITY ANALYSIS

1 VFR Results

Table IV-7 presents the capacity analysis results for the VFR cases. Results are reported for arrivals and

departures separately. In two cases, results are also listed for each active runway. The values reported are estimates obtained from averaging flow rates produced by FLAPS with the various streams of aircraft saturated. Ten replications of 3 hours each were used. The first hour is used as a warmup period, and results are averaged over the last two hours. A typical 95% confidence interval on a flow rate was plus or minus one aircraft per hour. The size of the confidence interval is the reason results are reported in Table IV-7 using at most two significant figures. Results under the heading "Percent Arrivals" are the capacity estimate for case a for the entire configuration under sustained operation at the given arrival-departure ratio. The method used to obtain these estimates is described below. The results for the VFR cases (and to a lesser extent the IFR results in the next section) illustrate some of the difficulties that can arise in performing capacity analysis with complicated configurations.

There are, in fact, at least three different "capacities", depending on how the problem is defined. This section discusses the proper modeling technique for each type of capacity in the context of discussing the reported results for case a.

Capacity estimates using FLAPS were reported in chapter III. In those cases, only one runway was in use for each

Assignment Method	Runway	Arrivals	Departures	Total	Percent Arrivals 40%/50%/60%
Case a: Arrival Priority					
OSEM	4L	30	22	52	124/120/113
	4R	34	18	52	
	9	0	20	20	
	Total	64	60	124	
SQ	4L	34	15	49	
	4R	34	4	38	
	9	0	32	32	
	Total	68	51	119	
Case a: Alternating Operations					
OSEM	Total	55	68	123	
SQ	Total	56	66	122	
Case e: Alternating Operations					
OSEM	22R	26	32	58	
	22L	22	8	30	
	27	17	0	17	
	Total	65	40	105	
SQ	22R	26	32	58	
	22L	21	16	37	
	27	21	0	21	
	Total	68	48	116	

OSEM - Runway assignment using OSE 78
SQ - Shortest queue assignment

Table IV-7: Logan VFR Results

type of operation. Thus there was no question of deciding how to assign aircraft to runways. The Logan cases have, typically, multiple runways in use for each type of operation, so a policy for assigning aircraft to runways must be provided. The OSEM report provides one possible method - a certain percentage of each aircraft class must go

to a given runway. The capacity estimates resulting from this method of assignment are reported as "OSEM" in table IV-7. This represents the first of three possible capacity estimates, capacity when runway assignment is constrained, or given in advance.

When the airport is simulated under this constrained assignment policy in case a, runway 4R becomes saturated with landings at a significantly lower overall arrival rate than does 4L, as more traffic is assigned to 4R than to 4L. If the flow rates are increased further, then, at some higher value, both runways will operate at saturation. The queue of aircraft waiting to land will be growing faster on 4R than on 4L. At the end of the simulation, there will be a short queue on 4L and a very long queue on 4R. With both runways saturated, the arrival capacity will then be close to equal on the two runways, as the same A/A separations are used and the runways differ only in aircraft mix. However, if flow rates were increased so that both runways were saturated, the mix percentage of actual landings in this situation will be slightly different than the mix percentage of arrivals. The reason for this is that arrivals assigned to 4L will have a better chance of having actually landed by the end of the simulation run than aircraft on 4R.

If, then, the objective is to determine the capacity of the airport, while maintaining the constraint on assignments, then the proper technique is to raise flow

rates until one of the runways (in this case 4R) becomes saturated. This was done to derive the entries in Table IV-7. Using this procedure will mean that there is additional time available for operations which is not being utilized. However this additional time cannot be utilized without violating the assignment constraint.

The second of the three methods of assignment can be termed a flexible, or shortest queue, SQ, assignment. In this method, it is assumed that some classes of aircraft may be limited to operating on certain runways, but there is no preordained constraint on how aircraft within each class will be assigned to the runways which that particular class is eligible to use.

With the shortest queue assignment, the mix percentage for actual landings will always be the same as the given mix percentage for arrivals. However, the assignment of aircraft to runways is determined by the computer in the course of the simulation (subject, of course, to any constraints that may exist on the use of runways by some aircraft classes). Thus, the runway assignment percentages for arriving aircraft of classes A and B in case a do not turn out to be the same under the SQ procedure as the OSEM assignments. As one might expect, the SQ procedure leads to both arrival runways being used about equally for case a. Consequently a larger fraction of class A and class B aircraft are assigned to runway 4L under SQ than under the

original OSEM assignment. The equal capacities, 34 aircraft per hour, for 4R and 4L may seem surprising since all heavy jets (with the longer separations they require) are assigned to 4R. However, this effect is apparently counteracted by the fact that the aircraft assigned to 4R have slower approach speeds (80 to 130 kts.) than do the aircraft assigned to 4L (130 to 140 kts.). The SQ assignments also differ from OSEM assignments for takeoffs. However, even with SQ, runways 4R and 9 are saturated with departures at flow levels higher than those at which 4L is saturated. The reason for this is that, in case a for Logan, there is little flexibility allowed in the assignment of departures. Classes A, B1, B2, and B3 must be assigned to 4L. Classes C and D can be assigned only to 4R or to 9. This means that, in effect, each aircraft is deterministically assigned to either 4L or to the 4R/9 set of runways. The heavy load on 4L causes the runway to saturate at flow rates below those required to saturate 4R and 9. Increasing flow rates for departures will eventually saturate 4R and 9 as well, but, again, only at the cost of causing the mix of aircraft that actually takeoff to deviate from the given mix of aircraft.

The allocation of departures between 4R and 9 changes from nearly half and half in the OSEM assignment to almost entirely all to runway 9 under the SQ procedure. This occurs because arrivals on 4R block departures on 4R for a considerably longer time than they block departures on intersecting runway 9.

The third type of capacity estimate often computed is a capacity for cases when the entire airport is restricted to operate for a sustained period with a given ratio of arrivals to departures. Typical ratios used are 40:60 (A:D), 50:50 and 60:40. Ratios more extreme than these are not commonly observed at peak hours. The ratio of arrivals to departures at which a configuration operates will vary in response to changes in separations and runway operating mode (arrival priority, alternating operations).

There are two techniques for obtaining estimates of aircraft capacity with a simulation for given A:D ratios. The first method uses simple extrapolation and interpolation, and gives approximate results. The second is a more time-consuming trial-and-error method which yields exact results. These methods will be described briefly below.

Both methods, however, are subject to a limitation. A/A and D/D separations have certain minimum values, so there is an ultimate maximum arrival capacity and an ultimate maximum departure capacity. In case a, for example, the maximum possible arrival capacity is 68 aircraft per hour. Since there are three departure runways in this case, even with maximum arrivals, departures under some saturation conditions are more than 40% of the total. The 60% arrival capacity is then found simply by reducing departures to 40% of the total. This means that total capacity is $68/.6 = 113$ aircraft per hour at 60% arrivals.

Simple interpolation-extrapolation was used to obtain an approximate estimate of airport capacity for the 50% arrival and 40% arrival situations. Case a was run using the alternating operations mode and the results are as given in table IV-7. The simple method assumes that the tradeoff between arrival and departure capacity can be adequately approximated as linear over the range 40:60 to 60:40*.

Figure IV-5 shows how the technique is used for case a. The arrival priority result gives capacity at $68/119*100 = 57\%$ arrivals. The alternating operations result gives capacity at $56/122*100 = 46\%$ arrivals. The arrival and departure capacities for these two are plotted at their respective A:D ratios. Lines are drawn connecting the arrival capacities and the two departure capacities. The lines cross at the point of a 50:50 ratio between arrivals and departures. The number of operations per hour of each type at this point is 60. Thus a 50% arrival capacity of $60 + 60 = 120$ operations per hour is found. To find the 40% arrival capacity figure the lines are projected from the alternating operations point towards a lower arrival ratio (to the left of the figure). A 40% arrival capacity of 50 arrivals plus 74 departures equals 124 operations per hour. The projection can be justified in this case since the ultimate departure capacity is very high, perhaps 100 operations per hour or more.

* Subject to the ultimate capacity constraint for arrivals only and for departures only.




FIGURE IV-5. Interpolation-Extrapolation Method for "Percentage Arrival" Capacity Estimation.

The estimation process described above was used here in order to obtain a simple approximate estimate of capacity. A more detailed capacity estimation study could easily investigate the exact tradeoff between arrivals and departures. A trial-and-error method could be used by slowly increasing A/A separations from the VFR minimum and observing the resulting flow rates. The capability of FLAPS to modify operating rules would facilitate combining several of these experiments into each run.

The results for case e for alternating operations are also included in Table IV-7. The results for the two assignment policies (OSEM and SQ) are very similar. The difference in arrival capacities is primarily due to random variations in the simulation runs. Additional departure capacity is provided by the SQ procedure, as compared to the capacity achievable by strictly adhering to the OSEM runway assignments, by reassigning aircraft from 22R to 22L. The resulting capacity estimate ($68 + 48 = 116$) is very close to 60:40 arrivals to departures ratio, and so may be used as an estimate of capacity for this situation. To obtain the estimates for 40% and 50% arrivals would require larger A/A separations. This could be done directly, as described above.

The complexities of obtaining a simple capacity estimate for a complicated configuration case that we have just discussed illustrate the fact that care must be taken

in the specification of what exactly is desired or what exactly has been estimated in such capacity analyses. If the airport is, in fact, constrained in certain ways (i.e. small aircraft really cannot use 4R or 9, or arrivals are 60% of total operations), then any capacity estimation procedure should reflect this fact by recognizing that there will be unused capacity, due to the constraint. On the other hand, if the purpose of the analysis is to estimate the maximum capacity for this case, then the constraints should not be applied at all. The model should be run without constraints to find what the optimal assignment policy is.

2 IFR Results

Table IV-8 gives capacity estimates for the three IFR cases. It uses the same format as table IV-7. Case b involves the side-step procedure described above. Cases c and d have all landings on 4R and differ only in terms of aircraft mix. The results for these three cases involve similar issues as under VFR results and, so, are not discussed in detail. Cases b and c as constrained by the OSEM assignment, have large, unused departure capacity. The SQ method results in significantly larger departure capacities. Regardless of assignment method, all cases have departure capacities large enough that, even when the

Assignment Method	Runway	Arrivals	Departures	Total	Percent Arrivals 40%/50%/60%
Case b:					
OSEM	4L	13	29	42	
	4R	14	13	27	
	9	0	23	23	
	Total	27	65	92	
SQ	4L	15	27	42	
	4R	15	14	29	
	9	0	40	40	
	Total	30	81	111	75/60/50
Case c:					
OSEM	4L	0	21	21	
	4R	27	13	40	
	9	0	27	27	
	Total	27	61	81	
SQ	4L	0	32	32	
	4R	27	8	35	
	9	0	36	36	
	Total	27	76	103	68/54/45
Case d:					
OSEM	Total	28	60	88	
SQ	Total	28	62	90	70/56/47

All capacities in aircraft per hour.
All cases are arrival priority mode.

Table IV-8: IFR Capacity Estimates

configuration is operated at the ultimate flow rate for arrivals (full arrival priority), arrivals do not constitute 40% of total operations. Hence, all of the "per cent arrival" capacity figures can be obtained directly using the

arrival flow rate.

H DELAY ANALYSIS

The cases presented in the validation chapter and in this chapter have so far been primarily capacity estimations. The only delay cases examined (La Guardia sensitivity runs, Bromma) have been cases where the operating rules remained constant over the duration of the run. This is partially a consequence of the fact that existing models have very little capability to alter operating rules. FLAPS has such capabilities, and this section discusses three hypothetical cases that illustrate a practical use of this capability.


1 Case description

We continue to use Logan Airport as our illustrative example. For this demonstration we limit aircraft to the use of just two runways: 9 and 4L/22R. The geometry information for 4L/22R is retained in the input specification of this case, but as no aircraft are assigned to 4L/22R, it was not a factor in these runs. As all landings (takeoffs) are confined to a single runway at any one time, the runway assignment problems discussed above do not arise. An eight hour case is run, for the 9AM to 5PM period. During this


time, the runway configuration in use changes twice. See figure IV-6 for configurations and times of change. These specific changes of configuration may not reflect actual operating practice. Note that this results in one runway being used in both directions at various times in the run.

A demand rate for arrivals and departures is hypothesized and is shown in figure IV-7. Inter-aircraft times for both landings and takeoffs were distributed as second order Erlang random variables. This was done, rather than using inter-aircraft times distributed negative exponentially, to reduce the stochasticity of the input process. This permitted the use of fewer replications to achieve acceptably narrow confidence intervals. The aircraft mix used was that of cases a and e. Classes B1, B2, B3 were collapsed into a single class B with the performance attributes of the B2 class in order to permit the addition of more classes as described below for case 3. All other parameters are the same as in the Logan capacity cases a to e, above.


For the base case (case 1) the airport operates under VFR-1 rules all day. Case 2 operates under VFR-1 conditions from 9AM to 9:30, IFR-3 from 9:30 to 12:30 and VFR-1 for the remainder of the day. See figure IV-8, for a summary of these changes. Case 3 is the same as case 2, except that aircraft performance parameters are changed during IFR conditions to simulate the slower movement of aircraft on



a) Configuration a, 9 AM to 11 AM



b) Configuration b, 11 AM to 1 PM



c) Configuration c, 1 PM to 5 PM

FIGURE IV-6. Logan Delay Configurations.




FIGURE IV-7. Logan Delay Cases: Demand Function.




FIGURE IV-8. Logan Delay Cases: Summary of Parameter Changes.

the runway in bad weather. Aircraft operating in bad weather were assumed to undergo the following changes with respect to aircraft performance parameters compared to aircraft in VFR weather (See also table IV-3): 1) float distances increase 20% for classes C and D, 15% for classes A and B; 2) deceleration rate for all classes decreases to 5 f/s/s from 5.75 f/s/s; 3) coasting speed for all classes is 20kts. rather than 25kts.; 4) takeoff runway occupancy time increases by 2 seconds for all classes; 5) the standard deviation of approach speed increases by 1kt.; 6) the standard deviation of takeoff runway occupancy time increases by 1kt.

These changes were implemented by defining four new aircraft classes, AI, BI, CI, and DI. These classes are the bad weather equivalents of classes A, B, C, and D. The aircraft mix was set so that during good (VFR) weather, only classes A, B, C, and D would use the airport. During bad (IFR) weather, a changeover is made to classes AI, BI, CI, and DI. The mix of the "bad weather classes" was the same as the mix of the "good weather classes".

2 Discussion of results

Average delay for the three cases is graphed in figures IV-9 (landing) and IV-10 (takeoffs). Each case was run for 16 replications. Selected 95% confidence intervals are shown in the figures.




FIGURE IV-9. Logan Delay Cases: Landing Delay.




FIGURE IV-10. Logan Delay Cases: Takeoff Delay.

As might be expected, in case 2, there are much larger landing delays than in case 1, due to the lower airport capacity during IFR operations. It should also be noted in case 2 that arrival delays continue to grow after VFR rules are imposed. This occurs for two reasons. First, the backlog caused by IFR conditions takes time to dissipate. The aircraft that land immediately after 12:30 (when VFR is reimposed) were already in the landing queue during IFR weather. Second, the demand rate increases after 12 o'clock. This illustrates the potential usefulness of being able to model operations during the entire day with a single run of the simulation. Even if the analyst is only interested in the period after 12 o'clock, the "initial" conditions at 12:00 can greatly influence delay estimates for a considerable period of time afterward. Note, as well, that landing delays in case 3 were unchanged from case 2. This is as expected, since all cases are run in an arrival priority mode and A/A separations are not affected by changes in runway occupancy time.

Takeoff delays decline from case 1 to case 2. This occurs because the larger D/D separations in use during IFR weather are more than offset by the additional opportunities to depart provided by the larger arrival/arrival spacing. The increase of delay from case 2 to case 3 is particularly significant. It illustrates the impact aircraft performance can have on airport delay. This impact was made clear by

the model and resulted directly from changing aircraft performance characteristics, not separation requirements. It was not necessary to first estimate changes in runway occupancy time caused by bad weather, and then estimate revised separations for use in the model.

I MODEL PROPERTIES

1 Sensitivity Analysis

An extensive sensitivity analysis of model parameters would be an integral part of using the model for a particular airport. As the primary purpose of this work is to develop a model rather than to perform an actual analysis, we have not performed such a sensitivity analysis. However, in the course of using the model for validation runs and for the examples of this chapter, a number of general conclusions about the relative sensitivity of the model results to changes in the values of certain parameters were made and are reported here.

Landing Performance Model

As reported in section II.D there is a significant range of uncertainty in several parameters of this part of the model. This means that, where runway occupancy time data exist, it can be useful to adjust parameters within the reasonable range to approximate more closely the data.

Arrival runway occupancy times (r.o.t.) for Logan are reported in OSE 78, though without any indication of the conditions under which they were obtained. No exit selection probabilities are reported. When FLAPS was first run it was apparent that it was producing shorter arrival r.o.t. than reported in OSE 78. To bring results into closer agreement with the OSE data, two parameters were adjusted. Coasting speed affects only r.o.t., it does not affect exit selection. As no data on the value of this parameter exists, it was chosen to be revised. It was adjusted from 30kts. to 25kts. Deceleration rate affects both r.o.t. and exit selection probabilities. The initial value of this parameter used for Logan was 6.00 f/s/s. This was the highest deceleration rate used for the validation runs, so its value was changed to 5.75 f/s/s. The combination of these two changes added approximately 2 seconds to average r.o.t. This still left FLAPS arrival r.o.t. below OSE 78 results. Since the latter tend to be longer than those reported in SWE 72 and KOE 78, we did not pursue closer agreement between FLAPS and OSE 78.

The impact on capacity of changing runway occupancy times will vary depending on several factors. No change in arrival capacity will occur in arrival priority cases as arrival capacity is controlled entirely by A/A separations. The impact on departure capacity will vary depending on the extent to which arrivals block departures. If departure

capacity is affected then this, in turn, can affect arrival capacity in a case that uses alternating operations. An example of the potential interaction of shifts in aircraft performance with capacity and delay estimates was provided by the delay cases run in the previous section.

Aircraft Separations

Aircraft separations have a very significant impact on capacity. They must be specified with care to obtain accurate results. As has been mentioned, it is often the case that airport analysts report only A/A and D/D separations used in their work. The La Guardia capacity case provided an example of the importance of carefully specifying the A/D and D/A separations as well (Rmt in FLAPS notation).

Demand Generation

In a delay case, the methods used to generate aircraft can greatly influence not only the variability of the results but the mean values as well. This was discussed in section II.C. For capacity cases, due to the fact that the airport is saturated and the stochasticity of the demand is not a factor.

2 Computational Efficiency

FLAPS is written in PL/1 and contains approximately 4000 lines of code. It is currently operated from the CMS

interactive system on an IBM 370/168 machine using a virtual operating system. The program requires 720k bytes of storage to run. The 6 hour, 10 replication capacity cases (cases a to e) required between 35 to 45 c.p.u. (central processor unit) seconds to produce all the results (OSEM and SQ assignments) reported for each particular case. Each of the three 16 replication, 8 hour delay cases (cases 1 to 3) required 44 to 46 c.p.u. seconds. Under the rate structure in effect at the time these cases were run, one c.p.u. minute costs \$6.00.

Only limited information about the computational efficiency of ASM is available to the author. It is known that a five hour run of a simple 2 runway, minimum taxiway system airport cost \$63.00 (BAL 76 p. 4). A three hour, ten replication, run of O'Hare (a very complicated system) cost \$76.00 (BAL 76 p. 7). No indication of the cost structure or computer used is given. A three hour, ten replication, run of a LaGuardia example under congested conditions took 18 c.p.u. seconds on a CDC CYBER 70 and cost \$18.00 (PMM 77 p. 44). A two hour run of O'Hare with five runways operating is said to cost \$100.00.* These cost figures do not seem to be directly correlated with the complexity of the situation being simulated, but they give an idea of the execution costs of the model.

* Personal communication to Prof. Odoni by PMM staff, 1978.

As with the comparison of input requirements in section F above, a detailed comparison of ASM and FLAPS is, thus, not possible. Information on the cost of identical cases is not available. The models have different capabilities. FLAPS and ASM are also run on different computers using different rate structures. However as the most expensive FLAPS run reported cost less than one third of the most inexpensive ASM run, it would seem that the added capabilities of FLAPS (internal demand generation, essentially unlimited maximum replications, more extensive statistics, calculation of aircraft performance) do not result in a model too expensive to use extensively. The low average cost of a FLAPS run, \$3 to \$6, means that it is feasible to perform many more experiments with alternative airport configurations and operations conditions than has heretofore been possible.

J CONCLUSIONS

This chapter further demonstrates the model's flexibility and usefulness in airport analysis. Logan Airport is a challenging airport to model because of its numerous "special case" situations (intersecting runways, parallel runways, displaced thresholds, conditional hold short arrivals, multiple runways in use simultaneously, side-step arrivals, frequent weather changes, frequent changes in the runway configurations in use, among others).

The cases described in this chapter illustrate the model's flexibility, adaptability and efficiency. It should be emphasized again that this chapter is not offered as a comprehensive analysis of the operation of Logan Airport. Such an effort is a major undertaking and is beyond the scope of this work. Instead, we establish here that FLAPS is capable of performing such a comprehensive analysis of a major airport. FLAPS can be used for capacity and delay estimation. It can be used for analysis of dynamic environments. It can model a number of special case situations. The combination of all of the above capabilities is unique to FLAPS.

V SUMMARY AND CONCLUSIONS

A REVIEW

The aim of this work was to contribute to the analysis of airport operations by developing a unified modeling framework that can analyze the performance of airports under a wide variety of conditions. To accomplish this task a number of specific contributions have been made to modeling airport operations. The cumulative effect of incorporating these contributions in a single model is the construction of an analytical tool which significantly expands the scope of issues that can be analyzed.

A simple, efficient method of specifying airport geometry is a prerequisite for a usable model. We saw that existing methods are so cumbersome as to interfere with the usability of the models. A new method for representing geometry was developed. It permits rapid specification and alteration of airport geometry. A related issue is the definition of patas aircraft take in taxiing over the airport. We found that current methods were neither sufficiently flexible nor economical. We proposed and described a new approach which approximates taxiing delay in an economical fashion while allowing flexibility in the level of detail at which ground operations are modeled. This approach, however, has not been implemented and remains a potential area for further work (see section B).

Existing methods of specifying the number and type of aircraft which use the airport were reviewed critically. These methods, while conceptually sound, may lead to misleading results if applied to situations involving a considerable amount of uncertainty about aircraft demand. It was established that current demand generation methods failed to capture certain types of uncertainty in demand and that this results in significantly underestimating delay. A new demand generation method was developed which is free of this problem.

Several contributions were made in this work to the modeling of runway operations. The landing performance of aircraft is modeled according to the underlying physical process, rather than by specifying the resulting behavior as do existing models. The reasonableness of this landing performance model was verified by testing it against two data sets. It was found that the model gave good results.

The rules for separating aircraft were modeled differently in FLAPS than in existing models. FLAPS, where appropriate, uses separation rules based explicitly on the performance of individual aircraft. Existing models usually specify separations solely on the basis of average performance. It was demonstrated that the new methods can model a class of situations which existing models can only approximate. The direct connection between separations and aircraft performance in FLAPS was shown to facilitate

modeling of the impact on delay produced by changes in aircraft performance.

Airports operate in a dynamic environment. Changes in weather induce a complex set of changes in operating policies and runway configurations as well as in the demand for landings and takeoffs and in the performance of aircraft. The occurrence of congestion can trigger changes in operating policies. Existing models have almost no capability to model time triggered (e.g., weather induced) changes and a limited capability to model event triggered (e.g., congestion induced) changes. FLAPS is specifically designed to facilitate analysis of both of these situations.

The most significant of the contributions of this work can be placed into one of two categories. They either provide a dynamic capability, where previously the analysis was limited to the static case, or they model interactions between airfield components, where previously the interaction had to be modeled externally. The dynamic capabilities were discussed above, and in section II.F. Interactions modeled were: 1) the interrelationships among the several aircraft attributes which, in turn, determine the characteristics of the aircraft generated by the simulation; 2) the interaction of aircraft attributes with runway exit characteristics to determine aircraft performance, and 3) the interaction of aircraft performance and separation rules to determine when operations may be

committed to move. The potential significance of each of these interactions was established as they were introduced in chapter II.

The cumulative effect of incorporating interactions between parameters and of providing a dynamic modeling capability is to greatly expand the scope of questions which can be addressed in a single model. In addition to standard capacity and delay estimates, FLAPS can be used to directly analyze situations involving the interactions listed above and situations where operating rules change during the course of a day. This allows an analysis of a complete day's operation under realistic conditions in one model run. The model can be used to address a number of questions related to optimal design of the airport geometry and of operating rules. Some of these possibilities were explored in chapters III and IV. Further examples will be given in the following section.

The model was partly validated by testing its results against results obtained from four other models of two different airports. Agreement with the other results was good in most cases. A majority of the significant disagreements could be attributed to minor discrepancies in the problem definition. The utility of the model for airport analysis was established through sensitive demonstration of its applicability to the case of Logan Airport. The model was able to replicate many of the

special situations in effect at Logan. The capability to model changes in operating rules was also demonstrated in that environment.

The advances and capabilities that we have just outlined, have not resulted in a model which is difficult to use or computationally inefficient. Input requirements for running the model are simple and can be prepared with relatively little effort. Individual runs of the model are very inexpensive. The fact that, models of three airports could be developed and run within the time and monetary constraints of this work testifies to the efficiency and ease of use of the model.

B FURTHER WORK

1 Needed Research

Despite the fact that aircraft, airports and the air traffic control system are closely monitored on a minute-by-minute basis there are surprisingly many important gaps in our state of knowledge when it comes to developing and validating airport performance models. As a result of this study, a number of topics have been identified as important for future research.

Performance of landing aircraft

As discussed in section II.D, development of a landing

performance model was hindered by a lack of consensus on the values of a number of basic parameters of the process. Research in this area has specialized on measuring parameters related to some particular aspect of the landing procedure. In order to develop better models, there is a need for comprehensive data collection efforts. Only with this kind of research can interactions between parameters be understood and quantified. It may be possible for some of this research to be performed in cockpit simulators with pilots. Some specific issues for investigation are: 1) the timing and rate of deceleration from approach to touchdown; 2) the correlation between approach speed and float distance; 3) the value, variance, and functional form of deceleration rate; 4) the value, variance, and functional form of coasting speed and exit speeds as a function of exit type; 5) interactions between some or all of the above.

Performance of departing aircraft

We were not able to develop a model of takeoff performance in this work, due to an almost total lack of data on this procedure. As the takeoff procedure is significantly less complex than the landing process, less effort should be required to obtain data that would provide the background for developing a credible model of this process.

Delay statistics

The motivation for this work has been the need to

predict airport level of service and performance, the primary measure of which is average delay. Despite the interest of many parties in this measure and the extensive monitoring of aircraft movements, there does not yet exist any reliable system for measuring and reporting delays. Such a system is needed to, first, assess the main problem areas under current conditions, and second, to permit validation of models such as FLAPS against "real-world" data, rather than solely against other models. It must also be strongly emphasized that the value of accurate data on delay statistics is greatly diminished if these statistics are not accompanied by a full set of information on the operating rules and demands for service which produced the recorded delays.

Communication of results

Research in the field of airport operations is hindered by a lack of communication of the results of work performed. Much work on airport analysis is published as consultants' reports which receive very limited distribution, as working papers, as unpublished memoranda and the like. This is particularly true of data collection and of more practically oriented activities, as the accompanying bibliography to this work attests. An effort should be made to collate and abstract this literature.

The perceptions and experience of people who operate the airport on a day to day basis have not been

systematically incorporated into the theoretical analysis. The decision making process of pilots (e.g., selection of exits) and controllers (e.g., judgement of separations) has not been documented in such a way as to make it readily usable in an analytical framework such as FLAPS.

2 Potential uses for FLAPS

At several points in this work suggestions have been made concerning significant experiments which can be performed using FLAPS. These suggestions are collected and briefly summarized here.

Capacity analysis

Significant analytical work has been done in the area of predicting the capacity of a single runway used for either arrivals only or for mixed operations. For more complicated situations (e.g., intersecting runways, cases involving more than two runways simultaneously) fewer studies are available. While some of these situations may prove amenable to analytical methods, it may be easier to use a simulation to obtain capacity estimates for them. An adaptable simulation, such as FLAPS, can be used to conduct systematic experiments which investigate the effects of changes in specific aircraft parameters or operation procedures on overall capacity for complex configurations.

A practical example may be cited. Examination of the feasibility of shorter A/A separations currently a subject of intense study (KOE 78. SWE 79). When a departure runway intersects or is coincident with an arrival runway, shorter A/A separations will mean less time for departures to roll between arrivals. For some configurations involving intersections near the beginning of the runway, the impact of shorter A/A separations on departures would be small. For other configurations (departures a runway also used for arrivals, intersections far from the beginning of the runway) the impact would be more significant. It would be interesting to analyze the way in which the reduction in departure capacity is related to the position of the intersection of the departure and the arrival runway. Analysis of the sensitivity of the results to changes in aircraft mix or D/D separations could be done.

Management of airport operations

Most airport analyses to date have been static in nature. Their primary concern has been to estimate capacities for a given situation and runway configuration. Little study has been made of how to manage the airport in a dynamic sense. This is presumably because, until now, there were no models capable of analyzing changes in operating rules. The capabilities of FLAPS in this area make feasible a study of this topic.

Changes in runway configurations involve a transition period. The nature, duration and effects of this transition are largely unexplored. It may be that certain transitions result in major disruptions to operations. Knowing the significance of a transition between two configurations is important in determining the desirability of making the change. For example, if the airport were operating in configuration "a" and a change in wind or noise curfews would make possible a shift to a higher capacity configuration "b", the decision to shift from "a" to "b" may not always be the correct one. It depends, among other things, on an estimate of how long the conditions which permit the use of "b" will persist, the transition "costs" of changing from "a" to "b", and the transition "costs" of changing from "b" to some other configuration. It is likely that controllers have developed a good intuitive understanding of which transitions work for their own airport. With a simulation capable of analyzing these changes it may be possible to quantify the process involved or even to suggest better operating policies.

Sensitivity studies

As indicated in the discussion of the landing performance model in sections II.D and IV.I, work should be done to explore the sensitivity of this model to the value of a number of different variables. Some work in this area prior to further collection of data on landing operations

could identify the most important variables and help focus the data collection process.

Once confidence is established in the values and form of the landing performance model, it would be appropriate to make a systematic study of how changes in landing performance affect the capacity of various configurations. This would involve, for example, assessing the magnitude of performance changes in bad weather. Then FLAPS could be used to estimate capacity changes under these circumstances for a number of airport configuration.

Statistical tests

In section II.9 a number of statistical issues related to the interpretation of simulation results and the design of simulation experiments were discussed. Much further theoretical work should be done on these issues. Most of these statistical problems are not unique to airport simulation. Nonetheless, FLAPS may provide a testbed for persons interested in using a simulation of a complex system to generate data for this type of work. The ability of FLAPS to use a variety of probability density functional forms and to produce concise replication by replication outputs would facilitate such an experiment.

Interaction with other models

Even though the airside of airports is a distinct and relatively "isolatable" entity, it interacts in significant ways with the entire air network on one hand and with the ground network on the other.

As discussed in section II.D, the sequencing of arrivals to the airport is actually accomplished at a significant distance from the airport. The acceptance rate of the departures ATC sector may limit the number of aircraft the airport can allow to takeoff. Certain runway configurations may not be usable at certain times due to airspace constraints. Thus modeling of the terminal airspace is another distinct and important area of work. Were a good terminal airspace model available, it would then be possible to combine it with FLAPS in order to study a set of airports in an urban area (e.g., New York) jointly with the common terminal airspace. Such a study would certainly provide important insights into the operation of the entire terminal area system.

The understanding of significant interactions between the airside and groundside of airports would also benefit by a joint modeling effort. Patterns of arrival, choice of apron areas and gates, and taxiing congestion are factors that cansometimes influence the passenger load in the terminal building. Conversely, the ability of the terminal to accomodate the parking of cars, unloading of baggage, ticketing and boarding of passengers, can all influence aircraft movements.

3 Extensions to FLAPS

Several extensions to FLAPS could be made to enhance its usefulness for airport analysis.

Ground operations

The most significant extension would be the implementation of the ground operations model developed in section II.B. This would allow analysis of ground congestion, and of the interactions, if any, between the taxiway system and the runways.

Runway assignment rules

An extension to FLAPS could be developed to model non first-come-first-served (FCFS) runway assignments. Certain pairs of aircraft require large A/A and D/D separations. By shuffling the order in which aircraft are removed from the queue of those waiting to use the runway it may be possible to minimize the occurrence of long separations and thus increase capacity. This is a subject that has received considerable attention recently (DEA 76, PSA 73) and is certain to receive much more in the future.

Antithetic variates

As mentioned in the statistics section, II.G, an antithetic variate capability could be easily added to FLAPS and might result in a significant reduction in the

variability of the outputs.

Validation tests

The validation tests performed in chapter III were sufficient to establish the model's general credibility. Additional validation tests would be desirable to test various aspects of the model. The ideal situation would be to have data on actual delays collected at an airport for a period of several days. Naturally, all major characteristics of the aircraft that used the airport and the operating rules in effect would need to be recorded. If such tests were performed carefully, it would allow testing of several features of the model which, of necessity, have remained unvalidated to date. Of particular interest is the detailed handling of changes in operating rules and in runway assignment policies.

BIBLIOGRAPHY

- ATA 80 Air Transport Association of America, "Air Transport 1980: The Annual Report of the U. S. Scheduled Airline Industry", June 1980.
- ATS 80 Air Traffic Service Division, Delay Statistics, Circular, 1980.
- ATW 80 "U. S. Airline Traffic, Fuel Costs", Air Transport World, April 1980, p. 14.
- BAL 76 Ball, Carl, "Model Users' Manual for Airfield Capacity and Delay Models", Report FAA-RD-76-128, Book 2, Federal Aviation Administration, November 1976.
- BOE 69 The Boeing Company, Jet Transport Performance Methods, Boeing Document No. D6-1420, Sixth Edition, May 1969.
- BOX 53 Box, G. E. P., "Non-Normality and Tests of Variances", Biometrika, vol. 40, June 1953, pp. 318-335.
- BOX 58 Box, G. E. P., M. E. Miller, "A Note on the Generation of Random Normal Deviates", Ann. Math. Stat., vol. 29, 1958, pp. 610-1.
- CON 59 Conway R. W., B. M. Johnson, W. L. Maxwell, "Some Problems of Digital Simulation", Management Science, vol. 5, 1959, pp. 92-110.
- DEA 76 Dear R. G., "The Dynamic Scheduling of Aircraft in the Near Terminal Area", Ph.D. Thesis, Flight Transportation Lab Report R76-9, M.I.T., Cambridge, Ma, 1976.
- DIJ 59 Dijkstra, E. W., "A note on two problems in connection with graphs", Numerische Mathematik, 1, p.269, 1959.
- FAA 77 "New York La Guardia Airport", Unpublished FAA memoranda, Appendix 3: Methodology, Appendix 12: New York La Guardia Airport", 1977.
- FAA 79 "Air Traffic Control", FAA Order No. 7110.65A, January 1979.
- FIC 74 Fischer, Martin J., "The Waiting Time in the Ek/M/1 Queuing System", Operations Research, vol. 22:4 July-August 1974 pp. 893-902.

- FIS 67 Fishman, George S., Philip J. Kiviat, "The Analysis of Simulation-Generated Time Series", *Management Science*, vol. 13:7, March 1967 pp.525-557.
- FIS 71 Fishman, George S., "Bias Considerations in Simulation Experiments", *Operations Research*, vol.20:4, 1972, pp.735-90.
- FIS 78 Fishman, George F., Principles of Discrete Event Simulation, (1973, John Wiley & Sons., New York).
- GAS 80 Gass, Saul I., Bruce W. Thompson, "Guidelines for Model Evaluation: An Abridged Version of the U. S. General Accounting Office Exposure Draft", *Operations Research*, vol. 28:2, March-April 1980 pp. 431-9.
- GOR 69 Gordon, Geoffrey, System Simulation, (Prentice-Hall, Englewood Cliffs, N. J., 1969).
- HAI 78 Haines, Andrew L., "Parameters of Future ATC Systems Relating to Airport Capacity and Delay", Report FAA-EM-78-8A, Federal Aviation Administration, Department of Transportation, June 1978.
- HEN 75 Hengsbach, Gerd, Amedeo Odoni, "Time Dependent Estimates of Delays and Delay Costs at Major Airports", Flight Transportation Laboratory Report R75-4, Massachusetts Institute of Technology, January 1975.
- HOC 77 Hockaday, Steven, Don Maddison, "Development and Application of an Airfield Simulation Model", *TRB Record #655*, 1977, pp. 17-21.
- HOR 75 Horonjeff, Robert, Planning and Design of Airports, 2nd. Ed., (1975, McGraw-Hill, New York).
- HSU 72 Hsu, D. A., J. S. Hunter, "Analysis of Simulation-Generated Responses Using Autoregressive Models", *Management Science*, Vol 24:2, October 1977 pp. 181-90.
- KLE 74 Kleijnen, Jack P., Statistical Techniques in Simulation 2 vol., (Marcel Decker, New York, 1974).
- KOE 78 Koenig, Steven, "Analysis of Runway Occupancy Times at Major Airports", FAA report No. FAA-EM-78-9, May 1978.
- KUL 73a Kulkarni, A. N., "Validation of PHM & Co.'s Airfield Delay Simulation Model", MITRE-METREK memorandum number W47-MO905, August 3, 1978.

- KUL 78b Kulkarni, A. N., "Review of Diagnostic and Sensitivity Tests on PMM & Co.'s Airfield Delay Simulation Model", MITRE Corporation Working Paper WP-13351, September 1973.
- LAW 79 Law, Averill M., "Validation of Simulation Models, I: An Overview and Survey of Real-World Practices", Technical Report 73-14, Dept. of Industrial Engineering, University of Wisconsin, August 1979.
- LEF 30 Lefer, Henry, "Air Traffic Control System for the 1990s and Beyond is Gestating at FAA", Air Transport World, April 1980, pp.30-3.
- LEW 69 Lewis, P. A. W., A. S. Goodman, J. M. Miller, "A pseudo-random number generator for the System/360", IBM Systems Journal, Vol. 8:2, 1969, pp. 136-145.
- LEW 73 Lewis, P. A. W., G. S. Shedler, "Simulation of Nonhomogeneous Poisson Processes by Thinning", Report NPS55-73-14, Naval Postgraduate School, June 1978.
- NAY 67 Naylor, Thomas H., J. M. Finger, "Verification of Computer Simulation Models", Management Science, vol. 14:2, October 1967 pp. B-92 to B-101.
- NAY 71 Naylor, Thomas H., Computer Simulation Experiments with Models of Economic Systems, (John Wiley, New York, 1971).
- NOR 78 Nordin, John, "Principles of a Flexible Simulation Model of Airport Operations", M. S. Thesis, Dept. of Civil Engineering, Massachusetts Institute of Technology, August 1973.
- NOR 79 Nordin, John, Amedeo Odoni, Feeter Kivestu, "Users Guide for the M. I. T. Airfield Simulation Model", Unpublished working paper, January 1979.
- ODO 79 Odoni, Amedeo R., "Annual Capacity and Delay at Bromma Airport, Stockholm", Flight Transportation Associates, Cambridge, Massachusetts, March 1979.
- OSE 78 "Logan International Airport Runway Capacity", Office of Systems Engineering Management, Federal Aviation Administration, August 1978.
- PSA 78 Psaraftis, Hailaos N., "Dynamic Programming Algorithms for Specially Structured Sequencing and Routing Problems in Transportation", Ph.D. thesis, Department of Ocean Engineering, M.I.T., September, 1978.

- PMM 77 Peat, Marwick, Mitchell & Co., "PMM & Co. Airfield Simulation Model User's Manual", April 1977.
- SHA 75 Shannon, Robert, Systems Simulation, The Art and Science, (1975, Prentice-Hall, Englewood Cliffs, New Jersey).
- SCH 59 Scheffe, Henry, The Analysis of Variance, (John Wiley, New York, 1959).
- SCH 74 Schriber, Thomas J., Simulation Using GPSS, (John Wiley & Sons, New York, 1974).
- SHA 65 Shapiro, S. S., M. B. Wilk, "An Analysis of Variance Test for Normality", *Biometrika*, vol. 52:3-4, 1965 pp. 591-611.
- SIN 75 Sinha, A. N., A. L. Haines, "Longitudinal Separation Standards on Final Approach for Future ATC Environments", MITRE Technical Report MTR-6979, October 1975.
- SIN 79 Sinha, A. N., "Summary of IFR Parallel Runway Configuration Rules", MITRE Corporation Memo Number W47-M1155, January 25, 1979.
- SWE 72 Swedish, William, "Some Measures of Aircraft Performance on the Airport Surface", M. S. Thesis, Dept. of Aeronautics and Astronautics, Massachusetts Institute of Technology, FTL Report R-72-4, May 1972.
- SWE 79 Swedish, William, "Evaluation of the Potential for Reduced Longitudinal Spacing on Final Approach", Report FAA-EM-79-7, Federal Aviation Administration, U. S. Department of Transportation, August 1979.
- TEI 66 Teichroew, Daniel, John Lubin, "Computer Simulation - Discussion of the Technique and Comparison of Languages", *Communications of the ACM*, v.9 #10 October 1966, pp. 723-741.
- TUR 76 Turnquist, Mark, Joseph M. Sussman, "Toward Guidelines for Experiment Design in Queuing Simulation".
- WEI 80 Weiss, William, "A Flexible Model for Runway Capacity Analysis", M. S. Thesis, Dept. of Civil Engineering, Massachusetts Institute of Technology, June 1980.

BIOGRAPHICAL NOTE

John Nordin was born January 9, 1954 in Toronto, Canada. He attended Kansas State University, majoring in Computer Science. He received a B.S. degree in Systems Science from Michigan State University in July 1976 and a S.M. degree in Transportation Systems from the Massachusetts Institute of Technology in August 1978. He is a member of Tau Beta Pi Engineering Honorary Society and The Operations Research Society of America. After graduation he will be employed by Systems Applications Incorporated of San Rafael, California.