DYNAMIC SCHEDULING IN AIRLINE OPERATIONS

O. J. Akel

MIT

RW Simpson
(personal file)

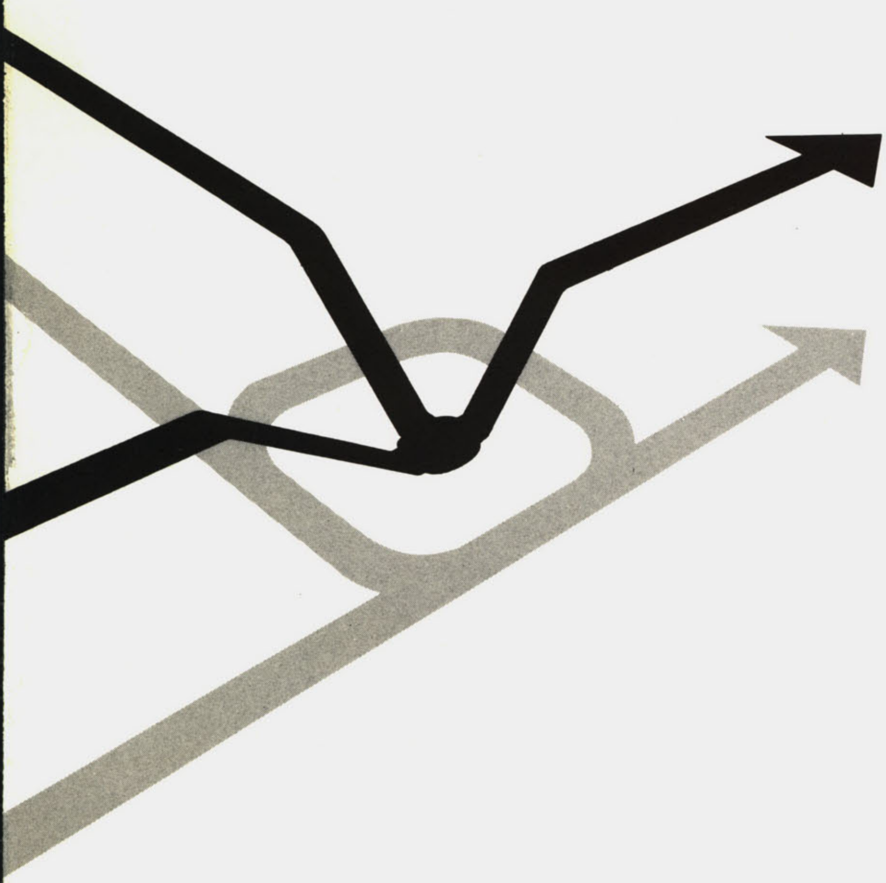

DEPARTMENT OF AERONAUTICS & ASTRONAUTICS

FLIGHT TRANSPORTATION LABORATORY
Cambridge, Mass. 02139

Prepared for the U.S. Department of Transportation Under Contract C-136-66

R-67-1

December 1967

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

Flight Transportation Laboratory

Technical Report FTL-R67-1

December 1967

DYNAMIC SCHEDULING IN AIRLINE OPERATIONS

O.J. Akel

This work was performed under Contract C-136-66 for the Office of High Speed Ground Transport, Department of Transportation, Washington, D.C.

TABLE OF CONTENTS

TABLE OF CONTENTS (Continued)

LIST OF FIGURES

LIST OF FIGURES (continued)

LIST OF TABLES

CHAPTER I

INTRODUCTION

1.1 Background

Although air transportation has been characterized by rapid development in vehicle design and performance, methods of airline management in the area of vehicle scheduling and control have advanced at a much slower pace.

Because of high costs of operation and the pressures of current competition and government controls, effective and efficient use of aircraft is becoming an increasingly essential objective. The goal is to achieve an optimal balance between net revenue to the airline and improved level of service to the customer. Improved return implies higher load factors and aircraft utilization whereas improved passenger service necessitates reduced waits and increased frequencies. These are often conflicting aims. New techniques must be mobilized to give management more useful and adaptive methods of operating and controlling an air transportation system. Perhaps the particular requirements of a very short-haul high density transportation system will lead to more demand responsive approaches. It is with this motivation that this study of dynamic dispatching strategy is undertaken.

1.2 The Scheduling Process

The decision-making process by which the system operates is called a scheduling strategy. Given the present system state in terms of accurate real time information concerning demands, passenger waiting time, vehicle availabilities, etc., and some short term expectations of future system states, a set of operating rules is established which determines the transportation system response. This set of rules, or strategy, always exists, either explicitly in the form of management policy directives, or implicitly in the form of the experience and intelligence used by the dispatcher. There are a wide variety of strategies available, each of which uses certain information about the system state.

1.3 The Scheduling Environment

In scheduling, it is convenient to distinguish between short-haul and long-haul type operations. Generally, circumstances inherent in longer haul operations make scheduling a simpler problem. These inherent features include the lower volumes encountered and the greater degree of advance preparation for the trip on the part of the passenger.

Short-haul traffic, on the other hand, usually involves larger volumes and little or no advanced preparation by the passenger. Included in the latter is the commuter service to and from metropolitan centers and the general collection and distribution problem between adjacent communities. Of course, a critical element in the scheduling environment is the uncertainty associated with demand. The uncertainty exists with respect to the total volume of traffic as well as the arrival distribution of individual passengers throughout the day.

Although the passenger places emphasis on speed, he is not only interested in short flight times but also in a short total elapsed time, i.e. the time from being ready to leave his home or office until the time when he arrives at his ultimate destination. Infrequent departures can result in unacceptable waiting times for many passengers who will consequently seek alternative modes of transportation. Excessive frequencies, on the other hand, burden the airline with unnecessary operating costs and over-capacity.

1.4 Types of Schedules

Basically, service schedules may take one of three forms:

a) Fixed Timetables

b) Dynamic or demand schedules

c) Mixtures of (a) and (b)

The fixed schedule, or timetable, is most commonly used, not only by airlines but also by intercity bus lines and passenger trains. Fixed timetables may be desirable for the passenger if it is compatible with his travel plans and he is able to secure a reservation.

From the operator's standpoint, a fixed timetable allows advanced planning and scheduling of resources. Through the use of historical data and knowledge of competitor action, departure times can be established which result in better operating efficiencies over the system. However, once scheduled, a flight must be operated regardless of adverse economic implications. Uncertainties in demand, therefore, can cause poor load factors and sub-optimal operation. While the passenger may want to be assured of departure times, perhaps less rigidity in times would permit an overall superior service.

At the other extreme lies the pure dynamic or demand type schedule where departures are governed by some function of the current state of the system. In its pure and unrealistic form dispatches would be made by a decision rule based only on some economic function of the number of passengers and their waiting times.

It is in the third type where the greatest interest lies. An example is the Eastern Airlines shuttle, where aircraft departures are scheduled at fixed times with supplementary aircraft accomodating the overflow. The popular appeal of this type of service is evidenced by Eastern's success in the shuttle. In a sense this is a demand schedule with a guarantee that passenger waiting time will not exceed some maximum value. With an uncertain operating plan of this type, the system must have excess resources in order to meet the guarantees during peak or above average traffic. Thus, lower efficiencies may result with correspondingly higher costs. However, better service for the traveller is presumably assured.

1.5 Purpose

The purpose of this study is:

1. To construct simulation models to represent several typical airline situations.
2. To formulate various dispatching criteria compatible with the environments modelled.
3. To demonstrate the use of the simulation models in:
    (a) Evaluating the formulated decision criteria
    (b) Isolating critical system variables.

(c) Determining capacity and aircraft requirements for a given system.

4. To explore other uses for the models.

1.6 Content of Study

In this study primary attention is focused on the dynamic schedule. In effect no prior schedule is assumed to exist and departures are governed basically by some economic consideration. Numerous such criteria are examined and tested on various model situations. In all cases there are overriding upper and lower bound heuristics which serve to limit passenger delays to a predesignated range.

Initially a mathematical approach was investigated with an attempt to apply decision trees and dynamic programming to the problem. These did not appear to be feasible for problems of practical size, however, and so attention turned to simulation. Simulation permitted the evaluation of different decision rules under different conditions. Further, it simultaneously created a good timetable. That is, for an assumed travel demand, use of the model generated a departure schedule conforming to the policy objectives built in. The demand assumption is reasonably valid as there usually exists a host of data on which to base expectations. Policy objectives are real factors in the scheduling decision but may vary considerably as the situation and relative competitive position changes. The sensitivity of these heuristics on overall system performance was examined by evaluating controlled changes. Three individual situations were modelled.

1. A two station shuttle problem
2. A nine station problem
3. A twenty station air taxi problem

1.7 Preview to Chapters

In Chapter 2, the various factors which enter into the decision to dispatch an aircraft are considered. From these factors several criteria evolve which are later used in the simulation models.

In Chapters 3, 4, and 5, the three problems listed above are discussed. Pertinent features of the structure and operation of the models are explained. The Chapters conclude with a discussion of the runs and the observations of significance.

In Chapter 6, the 'adaptive' approach to aircraft dispatching is considered. In essence the air taxi model of Chapter 5 was reprogrammed to permit past experience with respect to demand to be absorbed and later used to update the dispatching criteria.

DYNAMO and GPSS II simulation languages were used and runs were made on the M.I.T. Time Sharing System.

1.8 Validation

The two station and nine station problems are completely hypothetical. However, costs, fares, flight times, etc. used are considered compatible and realistic for such operations. Though no formal validation was possible the results seemed to conform with expectations of experienced individuals.

The twenty station air taxi problem was modelled from an operating helicopter service company in the area. Although here, too, no formal validation was conducted, various company executives were unanimous in asserting that the model conformed very closely to their actual operations. Again, it may be stated that the results, in general, were both reasonable and realistic.

CHAPTER 2

DEVELOPMENT OF DISPATCH STRATEGIES

2.1 Motivation

In the operation of an industrial enterprise, be it factory or airline, it is typical to set a specific profit objective as a corporate goal. Whereas price, or fare, is usually fixed by either supply-demand or else by government regulatory agencies, cost of operations, as influenced by quality of service and type of equipment, is in the realm of management policy and control.

Operating at lower echelons there may be other non-economic criteria, represented in different dimensions, but, nevertheless, contributing in some way to the specified overall economic objectives. In an airline these sub-objectives may be expressed as passenger delay or goodwill loss minimization, or as aircraft utilization and load factor maximization. We may enumerate many others while never losing sight of their overall economic implications.

In general the dispatching rules to be considered may be termed "heuristic" to differentiate them from the true optimal set of rules. The heuristic approach does not guarantee an optimal departure time but rather aims at achieving 'improved' performance. Through the facility of a simulation model, heuristics may be continually modified and observed until a practical and usable rule is evolved. Further simulation can be used to test the sensitivity of the decision variables of a given rule and thereby assess the value of added refinements.

The process of making a dispatching decision can be considered as a sequence of fundamental questions and a deduction based upon the answers. Therefore the development of a useful heuristic involves breaking down the process into basic elements and applying them systematically and consistently. In effect the attempt is to simulate the thought processes of an experienced dispatcher, the contention being that he will make good decisions most of the time. Poor performance is largely attributed to pressures which interrupt usual methodical and systematical reasoning, and to masses of unassimilated data which tend to be more confusing than helpful. With the development of logical heuristics and their diligent application by computers immune to such pressures, consistently superior performance should be realized.

In the simulation models of this paper, a number of different binary-type (go - no go) dispatching rules are developed and applied, more or less simultaneously. That is to say that only one of a sequence of rules need be satisfied to authorize a departure. Most of these may be classified as upper or lower bound constraints such as rules which cause an immediate dispatch when the maximum aircraft capacity is reached or which prevent an absurd delay time when only a few passengers are on hand.

However the major criterion is an economic objective function and embraces three components:

1. The accumulated fare
2. The fixed cost of operating the flight
3. A self-imposed penalty for delaying passengers.

2.2 The Fare

A fare structure is defined for each flight network considered. This is taken here to be a constant for a particular sector (or city-pair); special group excursion, mixed classes, children rates as well as round trip discounts are not considered. It would not be difficult to build fare variations into the models, but these would not contribute significantly to the conclusions and are omitted in the interest of simplicity.

2.3 Fixed Operating Cost

For each sector a direct cost of operating a flight is incurred regardless of the passenger load. This cost reflects the trip expenses of fuel, crew, fees, etc., a maintenance and a depreciation allowance. For each sector it may be considered to be a fixed quantity which we will take as some multiple of the passenger fare. This cost may not always be fixed, as for example when the models are being used to determine fleet size. Typically the same maintenance facility and crew can handle a range of aircraft.

Therefore the larger the fleet size, the smaller the maintenance burden each aircraft is called upon to bear. Maintenance costs constitute about 15% of the total fixed cost. Of course the option to subcontract maintenance nearly always exists and this would validate the assumption of fixed operating cost.

2.4 Passenger Delay Penalty

Passenger delay penalty is an essential factor in the dispatch decision criteria. Were it not included, the dispatching decisions would be nearly trivial, since with no penalty associated with passenger delay, the optimum strategy would be to dispatch only when full, except insofar as the aircraft might be needed elsewhere.

The delay penalty is a realistic factor implying both a short range and a long range consideration. In the short range, delaying a passenger excessively may mean a cancellation, thus losing him to a competitor or possibly to an alternate mode of transportation. The longer range consequence results from loss of goodwill. This may cause the passenger to intentionally avoid the airline sometime in the future because of some unpleasant previous experience.

Quantifying this function is no trivial task. Different people view this with varying degrees of importance. Indeed, the same person may well assess it different values depending on such variables as the time of day, the particular station location, or according to the activities of the competition.

Looked at from the passengers viewpoint, it is generally agreed that the penalty is not linear in waiting time but varies as some positive power (greater than one). It is assumed here that a passenger detained for two hours is likely to be more than twice as disturbed as the passenger held up only one hour. Secondly, it seems likely that the penalty should be different depending upon the projected length of the journey. A passenger who must wait one hour for a flight of ½ hour duration is apt to be more upset than a passenger who waits the same time for a four hour flight.

These facts are combined into a function by having the delay penalty vary inversely as the fare and directly as the square as the waiting time. A constant weighting factor, K_D, reflecting subjective values is included to complete the function. The delay penalty expression used in the simulation models is:

$$\text{Delay penalty} = \frac{n \cdot K_D}{f} \cdot \left(\frac{t}{2}\right)^2 \qquad (2.1)$$

where n = number of passengers waiting

t = longest waiting time

f = fare

Of course, it may also be argued that delay is relatively insignificant up to a given threshold or a point of discontinuity, beyond which further delay has extremely high cost. This threshold would represent the point at which the delay causes the passenger to miss a business appointment or a flight connection, etc. Clearly, this varies considerably and is most difficult to predict. Therefore, for the purpose of this thesis, the above expression will be used.

2.5 Selecting a Departure Time

For a given objective function, optimization is a relatively simple task when the function is continuous and differentiable. What we would really like to have is a continuous function with a single global peak. However in reality, there is a stepwise discontinuous curve since the function is made up of components shown in Figure 2.1. As discrete arrivals occur, a fare is collected. This is represented as a step function with the unequal intervals reflecting the randomly occurring interarrival times. Simultaneously with the arrival of the first passenger a delay penalty starts building up. Each arrival is a point of discontinuity and the slope of each segment becomes progressively more negative, reflecting the cumulative effect of an increased number of passengers waiting. Superimposed on these curves is a third signifying the fixed cost of operating the flight. These three components are combined in Figure 2.2 to show the single cumulative return criterion, or objective function.

Even if examination of this function is restricted

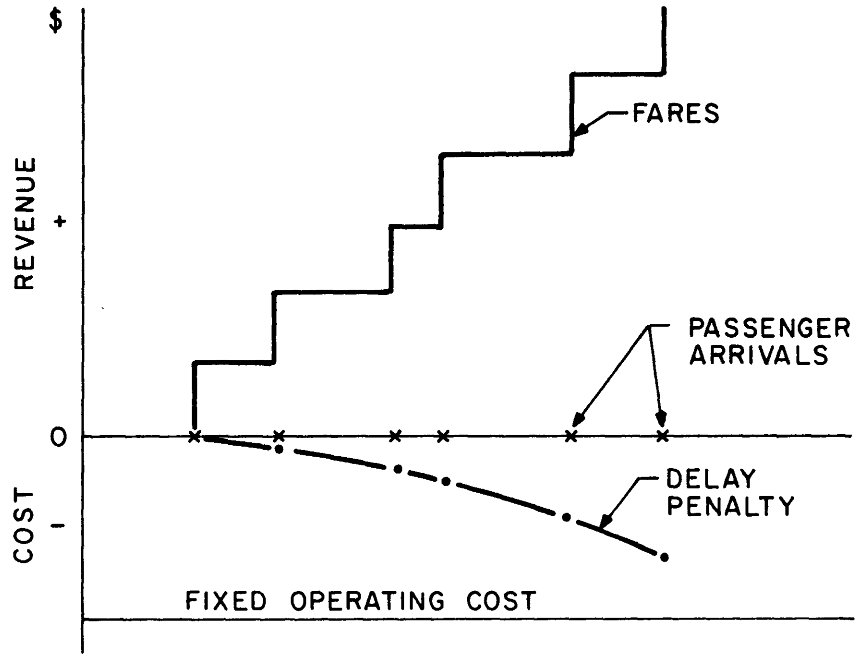

FIGURE 2.1: RETURN FUNCTION COMPONENTS

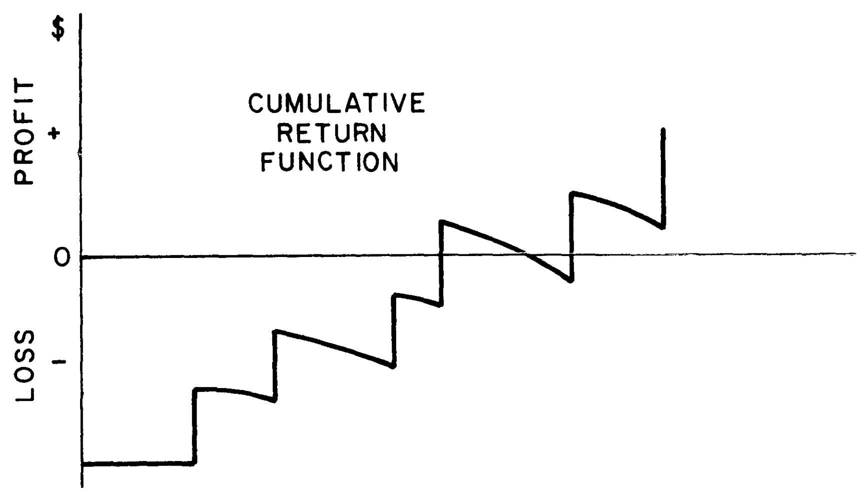

FIGURE 2.2: CUMULATIVE RETURN FUNCTION

to the discrete points of arrival, it may be easily concluded that more than one maximum is possible. This would be conceivable, for example, when a large number of passengers are waiting (and therefore the slope of the penalty curve is steep) and a relatively long period elapses before the next passenger arrives. Dispatching at this point would yield a sub-optimal result if the departure was followed by several arrivals in quick succession. Some expectation of future traffic is evidently required to assist in making optimal decisions.

One may choose to examine two forms of the objective function:

1. Cumulative return per flight
2. Average return per passenger

The following sections show some analytical considerations of these two cases in determining optimal dispatches when the arrival rate is uniform.

2.6 Cumulative Return per Flight

First consider the relation for total return per flight

$$\text{Return } R = nf - cf - \frac{nK_D}{f} \cdot (\bar{T})^2 \qquad (2.2)$$

where f = fare per passenger

n = number of passengers = λt

c = fixed cost as a multiple of fare

K_D = delay penalty constant

$\bar{T}$ = average waiting time = $t/2$

t = passenger arrival rate, assumed constant over t for this analysis

$$R = tf - cf - \frac{\lambda t K_D}{f} (t/2)^2$$

Therefore

$$R = tf - cf - \frac{K_D}{4f} t^3 \qquad (2.3)$$

Differentiating

$$\frac{dR}{dt} = f - \frac{3K_D}{4f} \cdot \lambda t^2 = 0$$

Therefore optimal departure time t_o , for a constant arrival rate is

$$t_o = \sqrt{\frac{4f^2}{3K_D}} = \frac{2f}{\sqrt{3K_D}} \qquad (2.4)$$

Note that in optimizing total return, the fixed cost is irrelevant. In effect the revenue from fares is being off-set by the increasing value of the delay penalty. When the

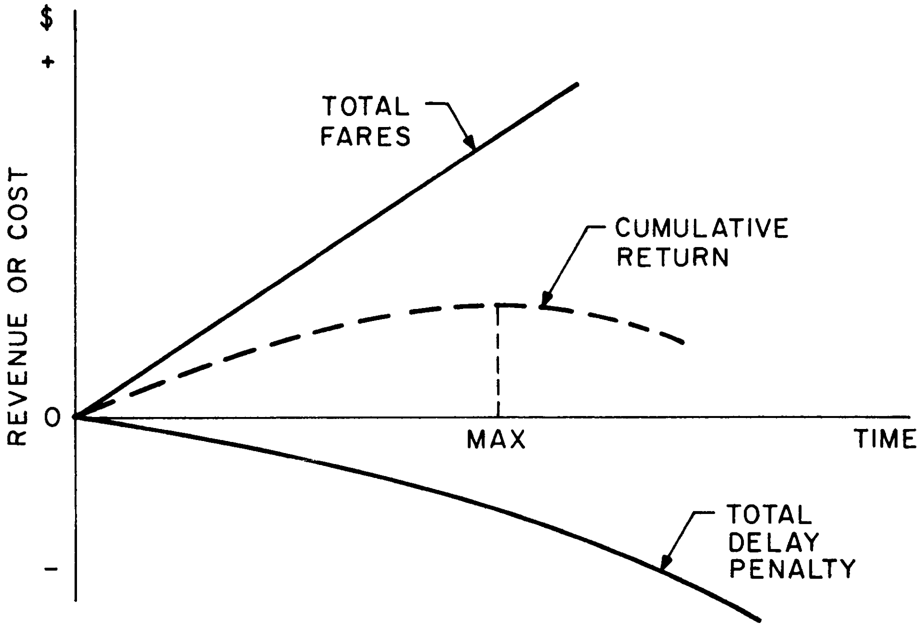

FIGURE 2.3: OPTIMIZATION OF TOTAL RETURN PER FLIGHT

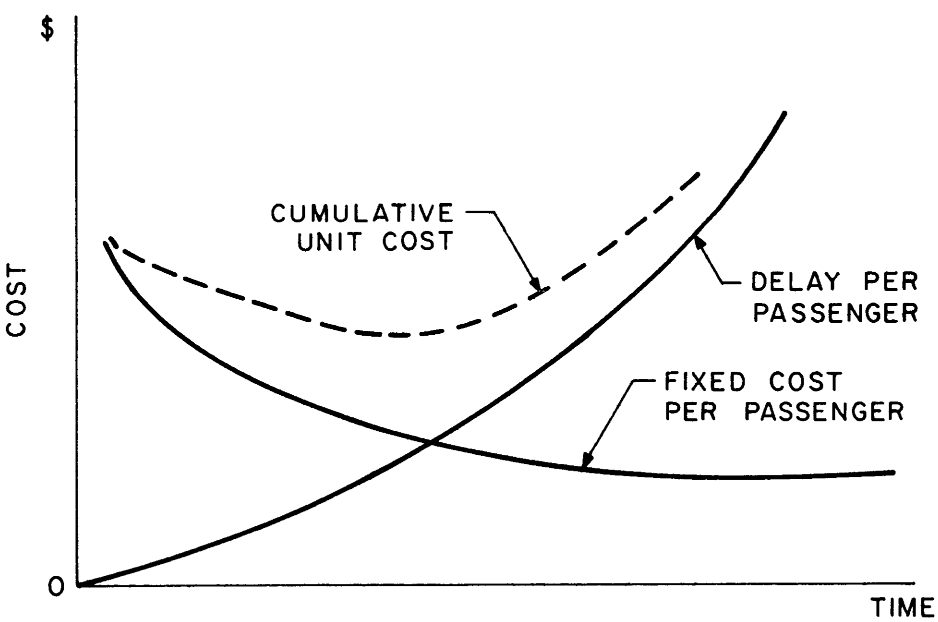

FIGURE 2.4: OPTIMIZATION OF AVERAGE RETURN PER PASSENGER

total fare accumulated per unit of time is equal to the delay penalty build up per unit of time, the optimal departure time is at hand. See Figure 2.3.

2.7 Average Return per Passenger

From equation 2.2, the average return per passenger is derived:

$$\text{Return/pax} = \bar{R} = R/n = \frac{nf}{n} - \frac{cf}{n} - \frac{nK_D(\bar{T})^2}{nf}$$

$$= f - \frac{cf}{\lambda t} - \frac{K_D t^2}{4f} \qquad (2.5)$$

Differentiating

$$\frac{dR}{dt} = \frac{cf}{\lambda t^2} - \frac{2K_D t}{4f} = 0$$

Therefore the optimal departure time

$$t' = \sqrt[3]{\frac{2cf^2}{\lambda K_D}} \qquad (2.6)$$

Note that in optimizing average return per passenger the fare is irrelevant. We are offsetting the decreasing fixed cost per passenger as the number of passengers increases with the increasing delay cost as the time (and passengers) increase. When these two factors are equal, the minimum cost per passenger obtains. See Figure 2.4.

In the absence of a delay penalty, the return per passenger may be expressed as

$$R' = f - \frac{cf}{n} \tag{2.7}$$

Therefore as passengers increase, the average profitability or profit margin increases. Now consider the effect of a delay penalty. Assuming a constant arrival rate, the delay penalty and the total curve must be represented as shown in Figure 2.5.

2.8 The Marginal Concept

It seems clear that the immediate objective of an airline should be the maximization of total return, within the resource constraints, as opposed to maximizing average return per passenger. The latter concept is analogous to the "full costing" averages considered in various areas of economics, and as such may be a perfectly valid parameter in long run considerations.

However in the short run when facilities and capacities are fixed, and in keeping with the spirit of dynamic scheduling, a variation of the total return criteria seems more appropriate. Here again from

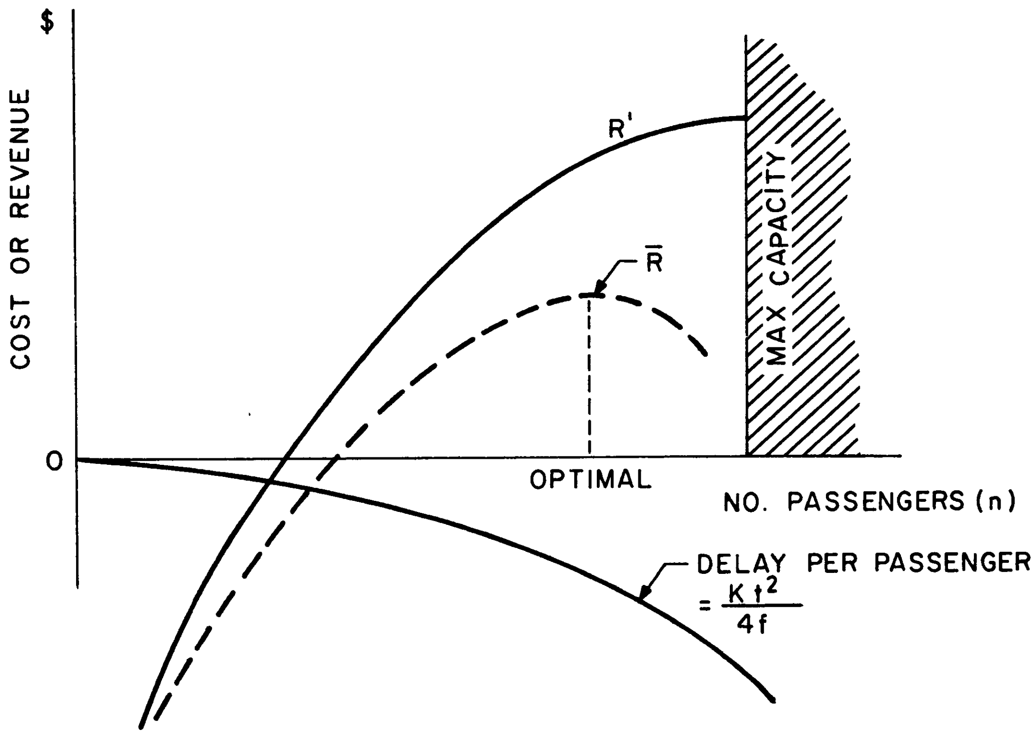

FIGURE 2.5: AVERAGE RETURN PER PASSENGER

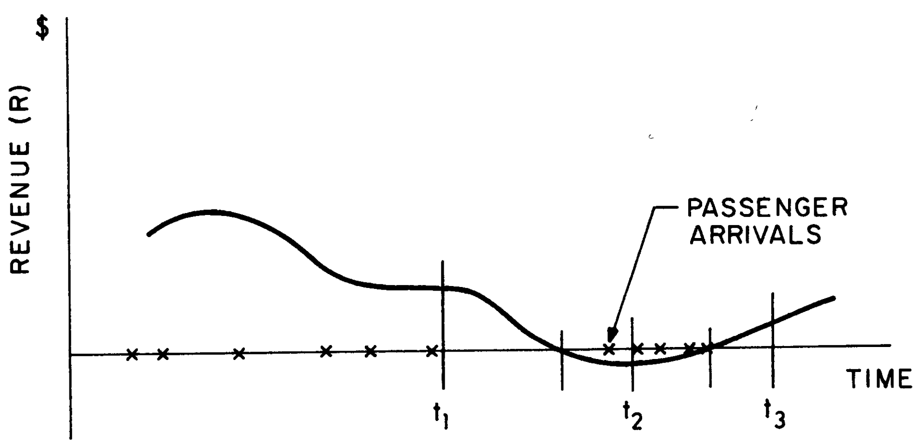

FIGURE 2.6: EFFECT OF NON-UNIFORM ARRIVAL RATES OF THE RETURN FUNCTION

economics we may draw the concept of "marginal" returns and apply it to the aircraft dispatching problem.

In considering whether an aircraft should depart at a given time t, or should wait an interval of time Δt to depart at a time t_2, the net contribution to overhead is the difference between the fares derived from passengers arriving in Δt, and the additional delay penalty incurred by the present passengers waiting until time t_2. An incremental penalty can be included for the waiting incurred by the arriving passengers in Δt. The marginal contribution to overhead then becomes:

$$MCON = mf - \frac{K_D}{f}\left[n\left(\frac{t_2}{2}\right)^2 - n\left(\frac{t_1}{2}\right)^2 + m\left(\frac{\Delta t}{2}\right)^2\right]$$

$$= mf - \frac{K_D}{4f}\left[n(t_2^2 - t_1^2) + m\,\Delta t^2\right] \qquad (2.8)$$

where m = expected number of passengers arriving in Δt

n = present number of passengers at t_1

If the arrival rate is uniform, the obvious criterion for dispatch would be a zero or negative value of MCON. (This would occur at intervals of time different from t_o of Section 2.6. since the incremental delay penalty of the arriving passengers is computed in a different manner.)

However, with non-uniform interarrival times this may not be the case since a large number of arrivals in a short period of time may quickly reverse the trend of the curve as demonstrated in Figure 2.6.

Characteristically demand for commercial air travel fluctuates throughout the day describing in essence a two-peaked curve. Therefore the total daily demand, whatever it may be, will be distributed in the manner suggested by Figure 2.7.

Arrival rates (and interarrival times) are consequently functions of the time of day. This point is significant since it invalidates the previous assumption that λ is a constant. We must therefore consider it as a time dependent, 'controlled variable' in that it changes throughout the day but in a predictable fashion.

In essence the problem reduces to one of selecting an appropriate time horizon over which to examine the the marginal returns of revenue and penalty. During

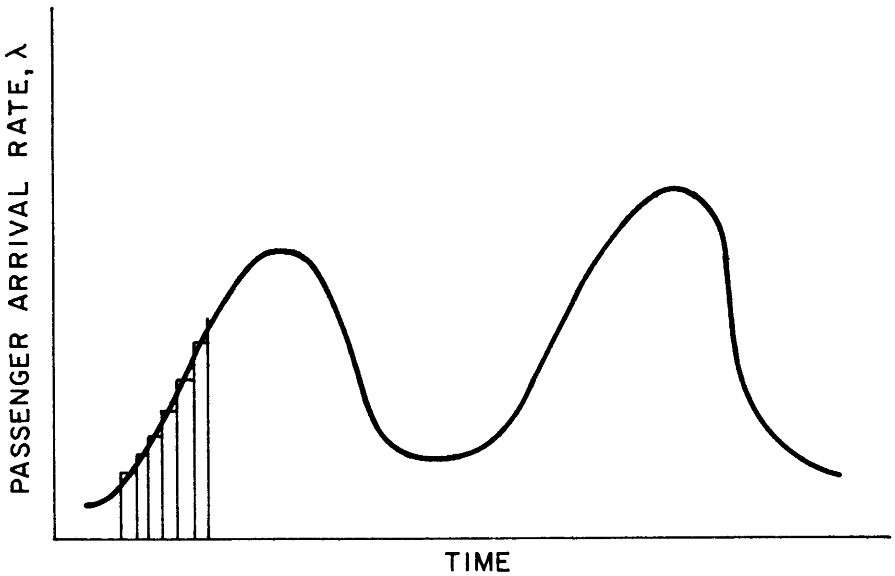

FIGURE 2.7: DAILY ARRIVAL RATE DISTRIBUTION

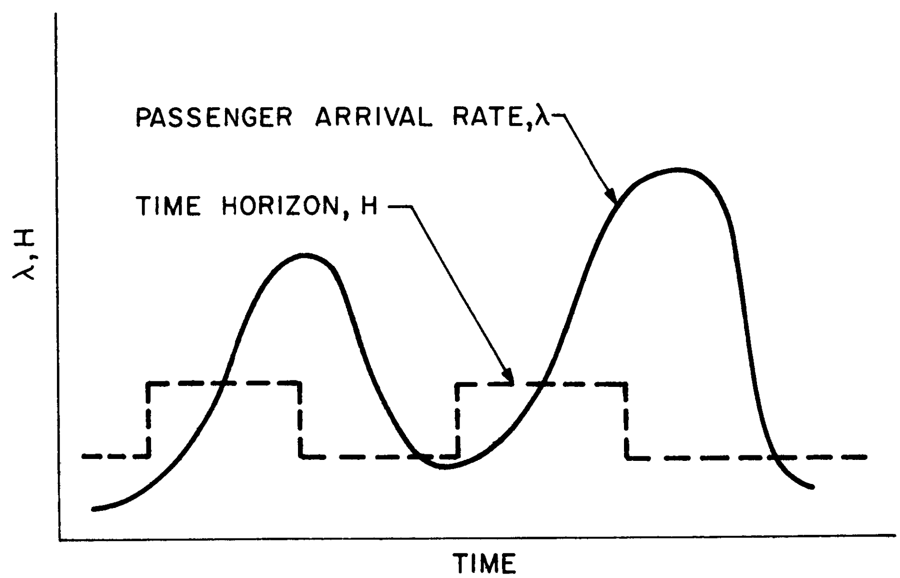

FIGURE 2.8: TIME HORIZON FOR THE MARGINAL RETURN CRITERION

peak periods, when there is a high probability of acquiring additional passengers by delaying the flight further, to extend the horizon would seem reasonable. Conversely, during slack periods, the value of looking beyond the next few time periods is correspondingly reduced. It is therefore feasible to define a variable horizon whose value changes as a function of the time of day distribution. It may be a simple two valued function (Figure 2.8) or a more complex infinitely variable relation. Therefore in equation 2.8, Δt becomes a function of time, H(t).

$$MCON = mf - \frac{K_D}{4f}\left[n\,(t_1 + H(t))^2 - t_1^2 + m\,H(t)^2\right] \quad (2.9)$$

This expression is in terms of quantities readily available to the decision rule at any time in any simulation. If MCON ≤ 0, then the decision to depart is made at t_1.

2.9 Coupling Considerations

In addition to the marginal contribution considerations of operating a flight from a particular station, it is often of equal importance to look ahead to the projected use of the vehicles further down the line.

The question is when and to what extent does the state at other stations become significant in the decision to dispatch a vehicle from this station?

There are basically two separate considerations of relevance.

1. The first pertains to the current disposition of the aircraft in the network
2. The second pertains to the passenger/waiting time states at the stations involved.

For example, if a decision is to be made whether to dispatch an aircraft from station A to station B, the passenger state at B becomes pertinent to the decision at A so long as the subsequent dispatch of B's passengers depends upon the arrival of this particular aircraft from A. That is to say, there are no aircraft at B or on route to B which could satisfy the expected demand at B. Clearly with unlimited aircraft in the system, the dispatch decisions at a particular station are independent of the system state at other stations. However, when the number of aircraft are limited, relative location status becomes increasingly important. Consider two stations in a network. From the standpoint of aircraft disposition, the wait time prior to dispatching from A to B may be considered as the sum of two components as determined by Figures 2.9a and b. Assuming equal demand expectations at the two stations, the lesser the proportion of aircraft at A the longer the wait time desirable at A. In this case waiting longer may result in more passengers being carried on the flight, without endangering the subsequent departures from B. On the other hand the greater the number of aircraft at A relative to the fleet size, the wiser the choice to depart early and thereby make available sufficient aircraft at B to handle the expected demand without suffering an unnecessarily high delay penalty. These two components may be weighted unequally if considered necessary (e.g. unequal demands at the two stations).

The second factor affecting the dispatch decision concerns the relative buildup of passengers at the stations involved. Important here are expectations

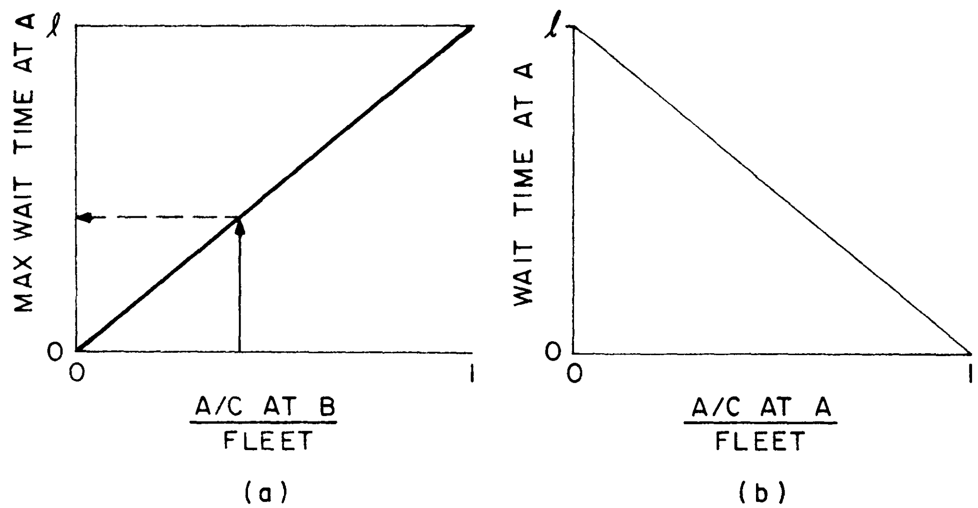

FIGURE 2.9: AIRCRAFT DISPOSITION COMPONENTS FOR DECISION COUPLING

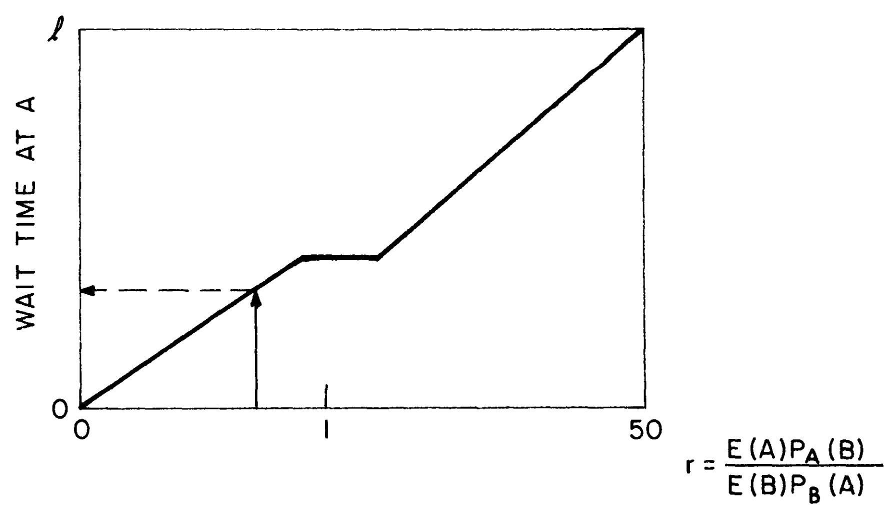

FIGURE 2.10: RELATIVE PASSENGER DEMAND COMPONENT FOR DECISION COUPLING

of daily demand and the relative time of day variations at each station. Consider impending flights from A to B and B to A. Passengers arrive at some expected rate with the probability of a passenger going from A to B being $P_A(B)$ and from B to A, $P_B(A)$. Let E(A) be the expected number of passengers originating at station A during a given time period

As the ratio of $\frac{P_A(B)\quad E(A)}{P_B(A)\quad E(B)} = \frac{\text{Expected no. of pax A to B}}{\text{Expected no. of pax B to A}} = r$

varies, the desirable predeparture waiting time for flight A-B varies from 0 to some number greater than one which is governed primarily by policy considerations and published guarantees. See Figure 2.10. This says that if there is a greater volume of potential passenger flow from A to B than from B to A in the given time period, all other things being equal, it is in the interest of the airline to delay the flight at A as long as practically possible. Conversely, if the ratio r is very small, i.e. there are a relatively larger number of passengers wishing to travel from B to A, then it may be desirable to dispatch the flight immediately in spite of the apparent dictates of the economic criterion. This is particularly significant in the absence of reliable data on passenger demands and when a given level of service is guaranteed.

The above relative-aircraft and relative-passenger criteria may be combined into some single equation form with equal or unequal weighting, and applied simultaneously with the economic and upper-lower bound criteria.

Depending upon the problem at hand, it may often be found that much simpler heuristics are useful. In a two station shuttle, one criteria in the decision rule might be governed by the number of aircraft available to the station and the total number of passengers waiting in the system. A limit for the product of these two measures is established, and, once exceeded, the aircraft is dispatched. In Figure 2.11 any point outside the shaded area would signal a departure. In effect, this states that if there is a high concentration of aircraft at the other station, and thus fewer at this station, a greater passenger demand is required for a departure. The establishment of the limit might depend upon the number and capacity of the serviceable aircraft in the system, route competition, management policies, etc.

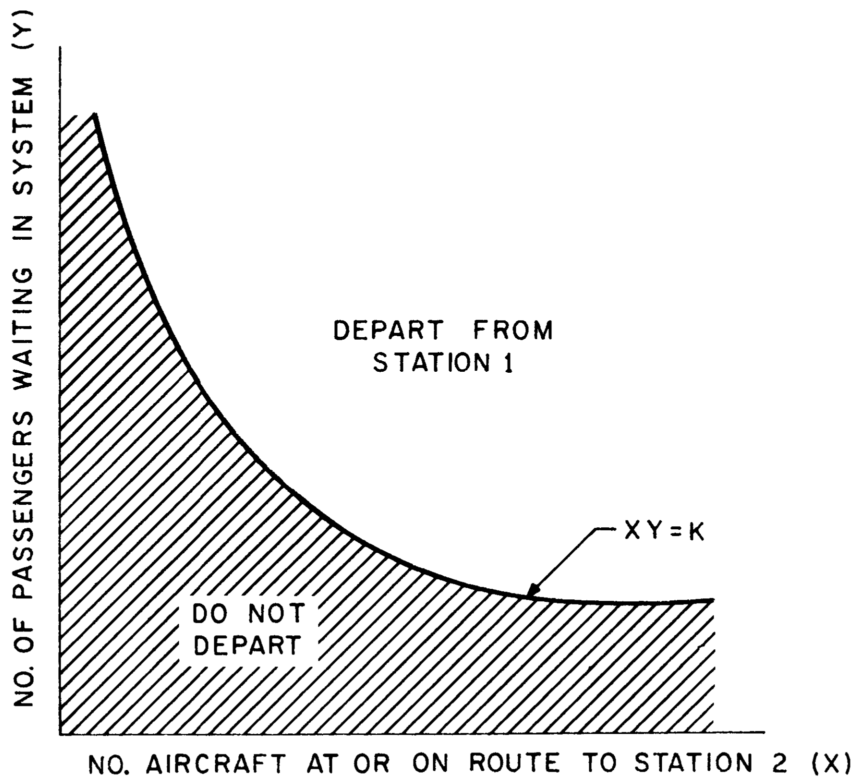

FIGURE 2.11: DISPATCHING HEURISTIC FOR DECISION COUPLING — TWO STATION SHUTTLE.

CHAPTER 3

THE TWO STATION PROBLEM

3.1 Dynamo Model

Among the initial simple problems considered was the two station shuttle with a fixed number of aircraft pre-positioned at the two stations. The average value of daily passenger demand was considered to be normally distributed. Time of day hourly variations were assumed known, given the daily average and the model was run using three different simple dispatching strategies.

Dispatch an aircraft:

Strategy 1) When demand is 55 passengers or more (60 for station 2)

Strategy 2) Every three hours after 9:00 AM provided that at least 55 passengers are on hand (60 at station 2)

Strategy 3) Every three hours and anytime passengers waiting exceed the set limit, 55 or 60.

The limit of 55 passengers was determined by carrying out an indirect optimization of the return equation.

$$\text{Return} = (P_W)\ (\text{Fare}) - K(P_W)^3 - \text{Fixed cost} \quad (3.1)$$

where P_W is the number of passengers waiting and the second term is an approximation of the delay penalty.

This model was written in Dynamo, a time driven simulation language. Outputs consisted primarily of a departure schedule with certain accrued parameters of interest carried along.

Of significance are:

TRET Total accrued returns from flights

PAXW Number of passengers waiting

LOAD Number of passengers carried on the flight

NACG Number of aircraft on the ground

NACF Number of aircraft flying

TPCAR Total accrued passengers carried

Aside from the schedule generated, the parameters tabulated serve a useful purpose. NACG gives a good indication of any over-capacity. Further it might be feasible to adjust departures slightly to realize an additional saving of one aircraft at a small cost in waiting time. The number of passengers waiting at any particular time, PAXW, is a useful indicator of the level of service being achieved.

3.1.1 Discussion of Runs

Results of several runs are tabulated in Table 3.1. Runs 1, 2, and 5 are identical with the exception of the strategy used. From the overall return standpoint, Run 5 (strategy 3) is slightly better than run 2 even though four additional flights were required to carry approximately the same number of passengers. The higher delay penalties involved in run 2 offset the fixed cost savings realized in higher load factors. Strategy 2 (run 1) is more difficult to satisfy than the others, consequently less flights were operated and longer passenger delay times were suffered. This detrimental effect of strategy is measured by the low associated total return. Runs 2 and 3 examine the sensitivity of the "optimal" passenger load to the overall performance of the system. In run 3 the load limit

TABLE 3.1 RUN SUMMARY

TWO STATION MODEL - DYNAMO

Run Number	1	2	3	4	5
Decision Rule Load Limit 1 Load Limit 2	2 55 60	1 55 60	1 55 55	1 55 60	3 55 60
Aircraft @ 1 @ 2	3 2	3 2	3 2	5 0	3 2
Pax Carried 1 2 Total	248 320 568	248 368 616	248 364 612	248 240 488	246 382 628
Revenue from 1 2 Total	1535 308 1843	1535 1073 2607	1535 987 2522	1535 231 1766	1091 1653 2744
No. Flights 1-2 2-1 Total	4 4 8	4 5 9	4 5 9	4 3 7	6 7 13
Average Load 1 2 Overall	62 80 71	62 73.5 68.2	62 72.8 68.0	62 80 69.5	41 54.7 48.2
Pax Waiting @ 1 (End of Day) 2	8 66	56 14	8 23	56 143	8 4

required at station 2 was reduced from 60 to 55. As a result the same number of flights were operated with the same average load factor. However total return dropped 3%.

In all runs discussed thus far three aircraft were pre-positioned at station 1 and two at station 2. In runs 2 and 4 the effect of pre-assigning all aircraft to station 1 was examined. As no ferry flights are permitted, passengers at station 2 must wait for a revenue flight from 1 to arrive before being accomodated. As a result heavy delay penalties were suffered at station 2 as evidenced by the low return ($1766 versus 2607).

While much can be learned from this model, its limitations are clear. First, there is little room for any degree of sophistication in the decision rules. Secondly, it is difficult to incorporate waiting times as an element in the decision process. Further it is difficult to expand the number of stations to a realistic level.

These and other considerations prompted a decision to move to a more powerful and flexible simulation language. Thus the remaining models were written in GPSS II and run on the MIT Time Sharing System.

3.2 GPSS II Model

This model represents essentially the same situation as that in Section 3.1, but possesses greater flexibility in the demand function as well as the decision criteria. Several working versions of this model were written. The one included herein is the most advanced from the standpoint of comprehensiveness. However, from the standpoint of computer time, it is relatively inefficient; having small time intervals and a large number of transactions (i.e., unit passengers).

3.2.1 Model Structure and Operation

In this model passenger arrivals are considered to be poisson with a known variable mean. A passenger interarrival time is calculated for each station in accordance with the distributions shown in Figures 3.1 and 3.2.

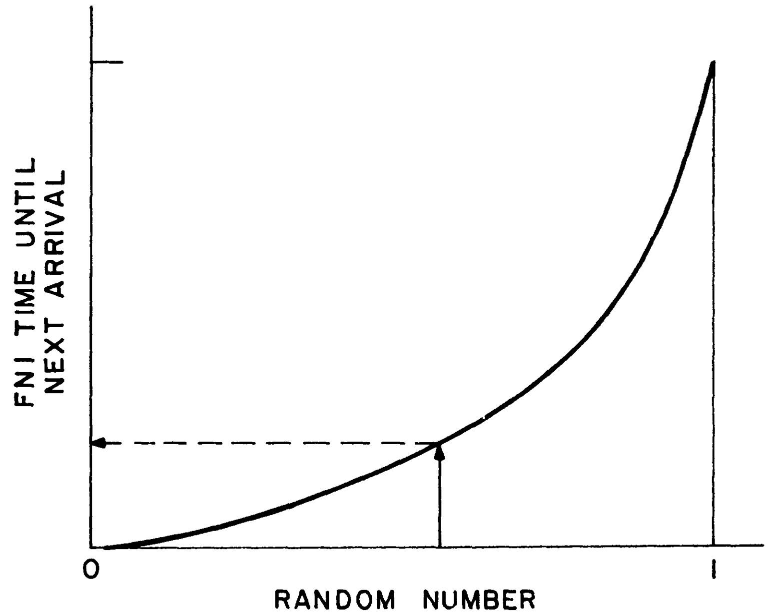

FIGURE 3.2: PASSENGER INTERARRIVAL TIMES—POISSON PROCESS.

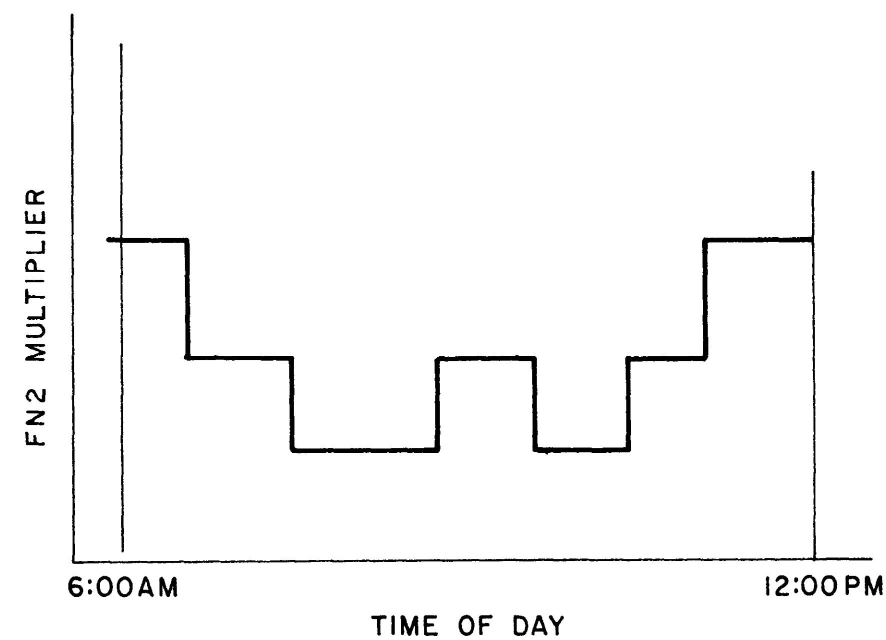

FIGURE 3.1: DAILY VARIATION IN AVERAGE INTERARRIVAL TIME.

Specifically,

$$IA = FN1 \cdot FN2 \tag{3.2}$$

where IA = Interarrival time

FN1 = the selected interarrival time

FN2 = time of day variation multiplier

Arriving passengers queue up to await a dispatch decision. Their arrival times and subsequent waiting times are maintained by the program. Flight times, fares and operating costs are assumed constant and the same in either direction.

Aircraft may be pre-positioned where desired, however the model in its present form does not possess a capability to handle ferry flights from one station to another.

3.2.2 Dispatch Strategy

At each station there are three separate rules which are examined prior to a departure decision. Only one need be satisfied.

1. Capacity decision: This permits a departure when a full vehicle load is on hand.
2. Economic decision: This is the "marginal cost equals marginal revenue" concept. When the expected revenue from a further delay is less than the incremental delay penalty expected, then a departure is initiated.
3. Coupled decision: This rule is intended to tie in the state at station 1 with the state at station 2. In this case the state is a function of the number of passengers waiting and the disposition of aircraft. Waiting times per se do not enter into this decision.

Critera 1 is self-explanatory; criteria 2 is explained in Chapter 2. The third criteria, however is a pure heuristic. Basically it states that the more aircraft at, or on route to, station 1, the fewer total number of passengers at 1 and 2 that

TABLE 3.2 - RUN SUMMARY

TWO STATION MODEL - GPSS II

Run Number	1	2	3	4	5	6
Decision Rule	Marginal M	M	M	M	M	M
Coupled	No	No	No	No	(inactive) yes	Yes
Delay Penalty K_D	$\frac{1}{2000}$	$\frac{1}{1000}$	$\frac{1}{1000}$	$\frac{1}{2000}$	$\frac{1}{2000}$	$\frac{1}{2000}$
Time Horizon H	30 min	30 min	f(T.O.D)*	f(T.O.D)	f(T.O.D)	f(T.O.D)
Run Time (min)	1000	1000	951	1000	1000	1000
Total Pax @ 1	256	241	240	230	230	230
@ 2	186	169	162	128	128	160
Total Return 1	797	-1725	3490	4415	4415	3451
2	-183	-2418	1160	1613	1613	1879
Total Delay Costs 1	142	159	2208	1476	1476	981
2	225	161	1057	1055	1055	1060
Flights From 1	9	12	4	3	3	4
2	7	10	4	2	2	3
Pax Left @ 1	11	26	6	37	37	37
2	1	18	9	59	59	27
Ave. Delay @ 1	35.7	27.2	84.5	91.5	91.5	75.2
2	47.9	30.8	69.5	113.6	113.6	91.5
Ave Load	27.6	18.6	50.3	71.5	71.5	56

* A function of time of day

must be available before an aircraft can be dispatched from 1. In Figure 2.11 any point outside the shaded area permits a departure.

This is a relatively simple attempt to link the decisions at the two stations. This rule may be further refined by considering the expected passenger build-up at the other station during the sector flight time, and some quantification of passenger waiting times.

3.2.3 Output Format

1. Schedule: The simulation generates a departure schedule showing

   a) time of departure or timetable

   b) passenger load

   c) net revenue for the flight

   d) the number of aircraft remaining at the point of departure.

2. Queues: Statistics on queues at stations 1 and 2 are tabulated showing maximum and average queue length, average waiting time, etc.

3. Savex: The system maintains a count on a number of parameters of interest, e.g.

a) Total accrued delay penalty at each station

b) Total fares collected at each station

c) Total return per station

d) Number of arrivals and departures

e) Number of aircraft flying, etc.

4. Tables: Statistics on passenger arrival rates showing the means and standard deviation as well as distributions are tabulated.

3.2.4 Discussion of Runs

Selected runs are tabulated in Table 3.2. The utility of the variable time horizon in the marginal return dispatch criterion is seen in a comparison of runs 1 and 4, and of runs 2 and 3. In 1 and 2 a time horizon of 30 minutes is employed while in runs 3 and 4 the horizon is either 30 or 60 minutes depending on the time of day. Note the substantial improvement in total return realized in the latter case. The improvement is attributed primarily to the large reduction in the number of flights operated. Higher passenger average waiting times are reflected in the large delay penalties. However, these large penalties were more than offset by the cost savings resulting from higher load factors.

Runs 3 and 4 may be compared to determine the effect of changes in the weighting, or importance, assigned to the passenger waiting times. In run 3 the delay is weighted twice that of run 4. With the greater emphasis on passenger delays, flights depart earlier and with less passengers. Average waiting times decrease correspondingly while the return function suffers moderately.

Runs 5 and 6 were intended to evaluate the effect of the coupling rule. In essence the coupling rule will tend to authorize a departure when a higher percentage of the aircraft are at the station in question and when there are many passengers in the system, i.e., at both stations.

With a coupling parameter of 500 in Run 5, i.e.

$$(X) \cdot (Y) \text{ greater than } 500$$

where X = number of aircraft at or en route to the station

Y = number of passengers at the station

the rule was always inactive and performance was identical to that of Run 4. In Run 6 a constant of 250 was tested in the rule. As might be expected better customer service resulted with a 19½% reduction in delay penalties. This is also seen in the reduction in average passenger waiting times. However, net revenues suffered a decline of 11.6% due to an extra two flights operated to accomodate an additional 32 passengers. As they now stand, these rules are by no means in their most desirable form. Many more runs would be required and basic policy guidelines of the airline would necessarily enter into consideration, especially with respect to the treatment of passenger delays.

CHAPTER 4

THE NINE STATION PROBLEM

The nine station model was built to depict a multi-station network problem with full interaction between city-pairs. It was designed primarily to evaluate the relative success of certain decision criteria and heuristics in realizing an improved level of operation, and to that end, to reflect the state of the system at any point in time.

As full interaction implies, both revenue and ferry flights are permitted from all stations to any of the other stations. Only non-stop service is considered, however, and there is no capability for multiple sector flights. This latter feature is a logical and important extension to the current problem.

4.1 Model Structure and Operation

An interval of time is taken to be 15 minutes. At each interval a transaction is created, representing some variable number of passengers wishing to travel from point of origin A to destination B. This input transaction may be split into any number of similar transactions to permit simultaneous creation of passenger demands for a number of sectors in the system. For the system described here 16 such transactions were introduced at each time interval. The origination-destination parameters (city-pairs) may be assigned values in a number of random or biased-random ways. Also the number of passengers associated with a transaction may be determined in any desired way -- poisson, random, etc. In the model constructed, a control loop selects a random number for each city-pair from an assumed distribution curve. This is taken to be the mean daily demand for that sector. This number, biased by a time of day variation parameter, (Figure 4.1) establishes the number of arrivals in the given time period. Further a multiplier is included to compensate for the number of time periods that are likely to be missed over the course of the day. The missed time periods occur since there are 72 city-pair combinations but only 16 assignments made per time

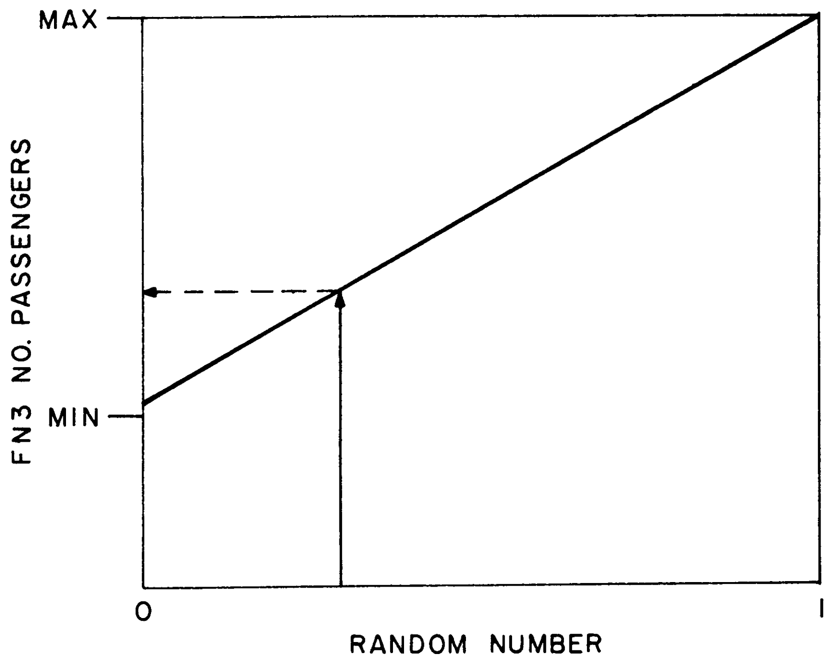

FIGURE 4.2: DAILY TOTAL DEMAND DISTRIBUTION

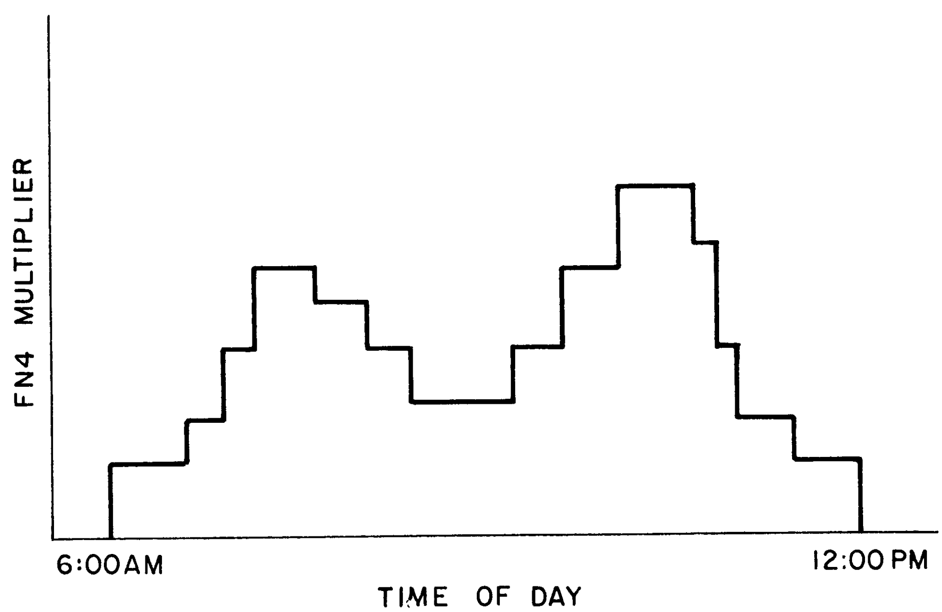

FIGURE 4.1: DAILY DEMAND VARIATION MULTIPLIER

period. For this case a multiplier of 72/16 or 5 would be in order. If warranted greater accuracy could be obtained by making a definite assignment for each city-pair combination at each time interval. The method used, however, reduces model size and introduces an additional element of randomness which is acceptable under such conditions of demand uncertainty.

Arriving passengers queue up by city-pairs. In addition to the passenger state for each city-pair, a fare, operating cost and flight time structure for the associated sector is included. These are derived from the matrix shown in Figure 4.3. The 9 x 9 matrix is partitioned into nine 3 x 3 sub-matrices. All similarly labelled sub-matrices have common fare, operating cost and flight time structures determined by:

Sector fare per passenger = 10 x L ($)

Sector operating cost = 250 x L ($)

Sector flight time = 2 x L (15 min. intervals)

where L is the label value A, B, or C

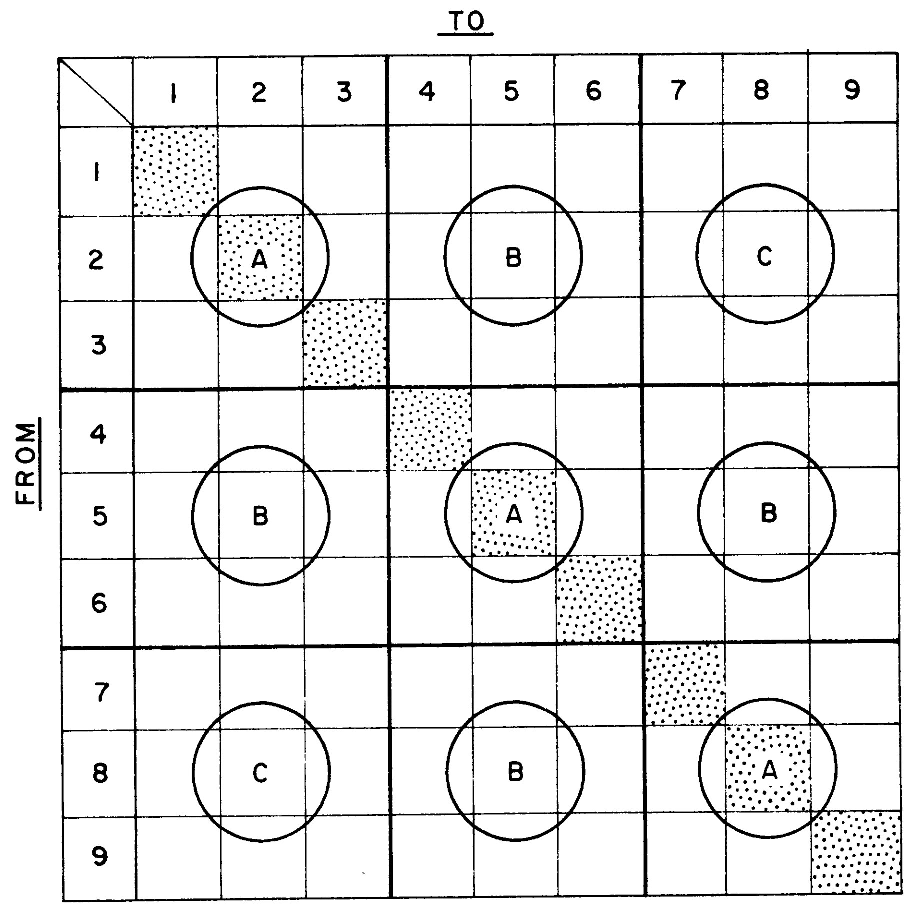

ALL COMMON LABELS INDICATE SIMILAR FARE, OPERATING COST, AND FLIGHT TIME BETWEEN CORRESPONDING CITY-PAIRS

FIGURE 4.3: FARE/COST/FLIGHT TIME STRUCTURE —NINE STATION NETWORK

Initially the number of aircraft in the system is specified. Initial aircraft disposition may be set deterministically or in a random fashion by a pre-position control loop. After simulation has commenced, the program routes the aircraft in accordance with the dispatch and ferry decision and aircraft availability.

4.2 Dispatch Decision Rules

A number of dispatching criteria were tested on this model. The following demonstrated the most promising performance.

Dispatch when:

1. Aircraft capacity: A capacity load is available
2. Economic: The incremental passenger delay penalty to be experienced is greater than the marginal revenue expected during the next time period. The time period is considered to be a variable function of the time of day.

3. Time limit: A parameter which is a function of the number of passengers waiting and the longest passenger waiting time exceeds a preset limit. This is intended to prevent extremely long delays when only a few passengers are on hand.

Of course in all cases an aircraft must be available at the station in question.

In addition to the decisions governing revenue flights, rules controlling the dispatch of ferry flights are also included. This rule is similar to criterion 3 above with the added condition that there must be a free aircraft somewhere in the system. In particular:

1. A parameter, which is a function of the number of passengers in the queue and the longest waiting time, exceeds a prespecified limit.
2. There are no aircraft at the station of origin.

3. There is at least one free aircraft on the ground in the system and as yet undesignated for service.

4. There have been no prior ferry calls which have not yet arrived.

5. There are no revenue flights on route to the station.

With these conditions satisfied, the program searches all stations starting at the nearest and will requisition the first free aircraft found. If there are passengers queued up for the particular sector to be ferried, the flight will be regarded as a revenue flight and dispatched immediately. Otherwise, the sector will be operated empty as a pure ferry. The total number of ferry flights flown and the associated costs are recorded at the completion of the run.

4.3 Output Format

For all revenue flights a departure schedule is generated showing:

1. Time of departure
2. Origin-destination code
3. Passenger load
4. Flight revenue (fares less costs)

Statistics are also tabulated on queues, revenue, and ferry flights, aircraft location, demand distributions, etc.

4.4 Discussion of Runs

The runs for this model were intended to evaluate the relative merits of various decision criteria and the sensitivity of other determinants operating in the system. Consequently, the measurements are not of particular interest for their absolute values, but rather for the magnitude and direction of the deviation effected by a controlled change somewhere in the system.

Overall performance is measured by two parameters:

1. Total profit for the period (fares less operating costs.)

2. Average passenger queue lengths for the period of the run.

The first is a measure of the overall economic success from the airline's standpoint. The second is a measure of the level of service provided to the customer.

In the interest of computer time economy, simulation runs were allowed to run for 7½ hours, from 6:00 AM through the peak period of the morning and into the slack of the afternoon, until 1:30 PM. These were considered representative for the purpose of this report, however, for more accurate comparisons, longer runs would be required, extending not only through the full day but over several days. This is of particular importance in view of the random selection techniques employed by the model.

A summary of some of the runs are tabulated in Table 4.1.

TABLE 4.1 SELECTED RUN SUMMARY
NINE STATION MODEL

Run	58	59	63	68	72	74	75	76	77	78	79	80
Passenger Demand	(Low) L	L	(High) H	H	L	L	L	L	L	(Med) M	M	M
Delay Pen. K_D	1	1	1	3	3	3	3	5	1/4	1/4	5	1
Time Horizon H	* 3V	3V	3V	3V	3V	3V	3V	3V	3V	3V	3V	2 V
Guarantee K	5	10	25	25	25	25	25	25	25	25	25	25
Run Duration	450	450	450	450	450	450	450	450	450	450	450	450
Demand Expect.	1/6	1/6	3/2	3/2	1/6	1/6	1	1/6	1/6	1/3	1/3	1/3
No. Aircraft	10	10	15	5	20	5	10	10	10	10	10	10
Net Revenue	-7960	-5820	9220	9130	-4000	-3010	-4180	-4410	-3800	-1090	- 2330	-1520
No. Rev. Dep.	26	23	26	24	21	17	19	20	17	20	23	21
No. Ferries	1	1	0	0	0	3	0	0	0	0	1	1
Cost of Ferries	250	500	0	0	0	1500	0	0	0	0	250	250
No. Pax Carried	215	238	1146	963	297	260	237	239	225	414	433	422
No. Times Criteria Used: Rev.Ferry	9	5	3	15	3	6	6	6	6	6	6	9
Capacity	0	0	11	2	0	0	0	0	0	0	1	0
Economic	6	12	6	5	15	10	12	14	9	10	13	8
Guarantee	12	7	7	3	4	2	2	1	3	5	4	5
Total Delay Cost	-	-	-	14.431	3798	10.463	6357	10.621	688	455	14.238	3747
Ave. Load	8.3	10.3	41.1	40	14.1	15.3	12.5	11.9	13.2	20.7	18.8	20.1
Ave. No. Pax Waiting	5.2	5.0	55.8	62.4	6.4	7.15	4.9	4.8	5.6	10.3	10.1	10.5

* Variable time horizon with a max. of three periods

4.4.1 Effect of the Delay Penalty

The assignment of values to the passenger waiting times is believed to be an important factor in any economic dispatching criterion. To determine just how important and through what mechanism it operates, several identical runs were made with only the subjective delay penalty constant changed.

Runs were made at two levels of demand, the results of which are presented in Figure 4.4. The general trend for both demand levels is a deterioration in economic performance as the delay penalty is increased. Notice that there exists a range in which no change is experienced. The measure of passenger service, average queue length, runs contrary to the economic trend and improves with increasing delay penalty.

Increasing the delay constant induces earlier departures. However, in the range between K = 2 and 4, at least for the conditions of this case, the delay penalty is insufficient to overcome the fare increment

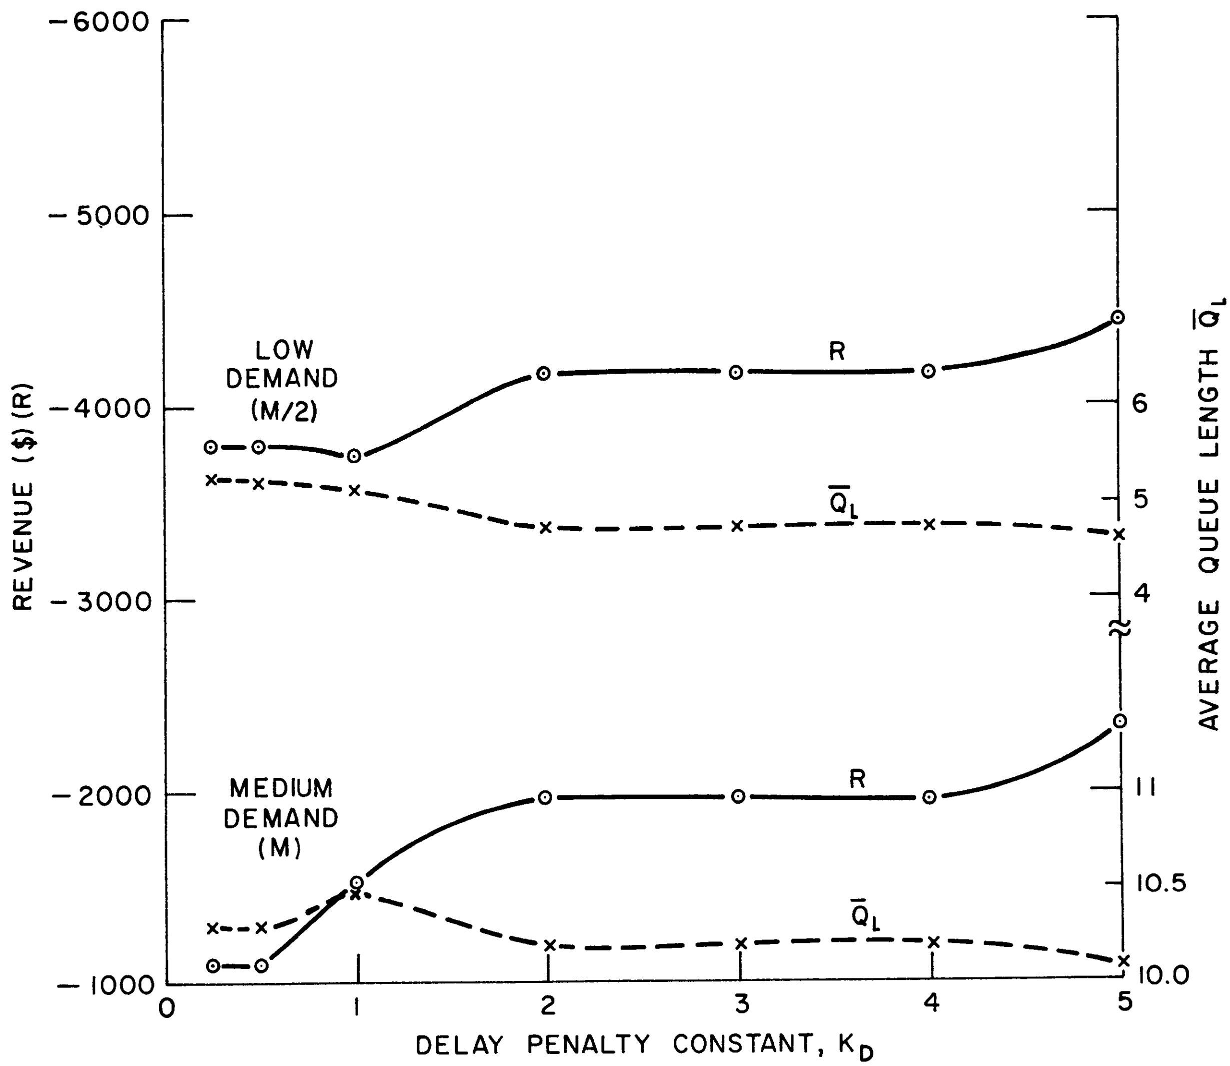

FIGURE 4.4: EFFECT OF DELAY PENALTY CONSTANT

expected. Beyond 4, however, this is reversed, and a flight is dispatched earlier. This later precipitated a ferry flight which was not required for the case of the lower delay constants. This higher penalty cost thus resulted in fewer passengers carried (by 5) and one extra flight. For other fare/cost structures and passenger arrival rates, a different 'break-point' would most probably exist.

If no passengers are expected over the time horizon, the economic decision criterion would automatically dictate a departure since the fare increment (=0) would always be less than the delay increment. This may be detrimental and suggests the desirability of coupling this criterion with an added requirement for minimum number of passengers. This is especially significant in a low demand situation when the probability of any passengers arriving in the next time interval is very small.

Average loads in the low demand case range from 12 to 13. In the higher demand case (twice the lower) the average load varies between 20 and 21, i.e., a factor of about 1.6. Runs were also conducted for an even higher demand case (ten times the lower). For this run, average loads were at about 54, up slightly over four times. With more efficient use of the aircraft implied by higher average loads, substantial improvements in gross or net revenues result. However, as seen, load factors do not automatically increase proportionately with customer demand. Perhaps a useful by-product of this model is in its ability to indicate reasonable aircraft capacities for a given fleet size and demand expectation. With slight modification it may also be used to determine fleet size for given capacity aircraft.

Ferry flights can be a source of distortion in trends. When called, a ferry is drawn from available aircraft in the system. If one is nearby, a cost of \$250 is incurred, but if it is far away, as much as \$750 surcharge must be paid for the same service. This ferry, being further away, also takes longer to arrive. Therefore, in addition to the extra cost, longer waiting times are suffered. Longer periods of running the model are required to average out the distortions from such effects.

4.4.2 Effect of the Time/Passenger Waiting Constraint

Aside from the economic and capacity criteria, the third criterion in the decision rules, an upper bound on the delay time, also plays a significant role. This criterion essentially places a limit or threshold on the number of passengers waiting and the time they have been waiting. The more passengers waiting, the less waiting time required to trigger a dispatch (or call for a ferry). Conversely, the fewer the passengers, the longer the time they are required to wait before the criteria is satisfied. That is, the longest waiting time (in 15 minute intervals) plus number of passengers waiting must be greater than the threshold specified. Of course, this criterion may be overridden by the capacity and economic criteria. Thresholds of 5, 10, 15, 20, and 25 were examined under two levels of demand, one twice the other. All other variables in the system were held constant. As seen in Figure 4.5, and as might be expected, better economic

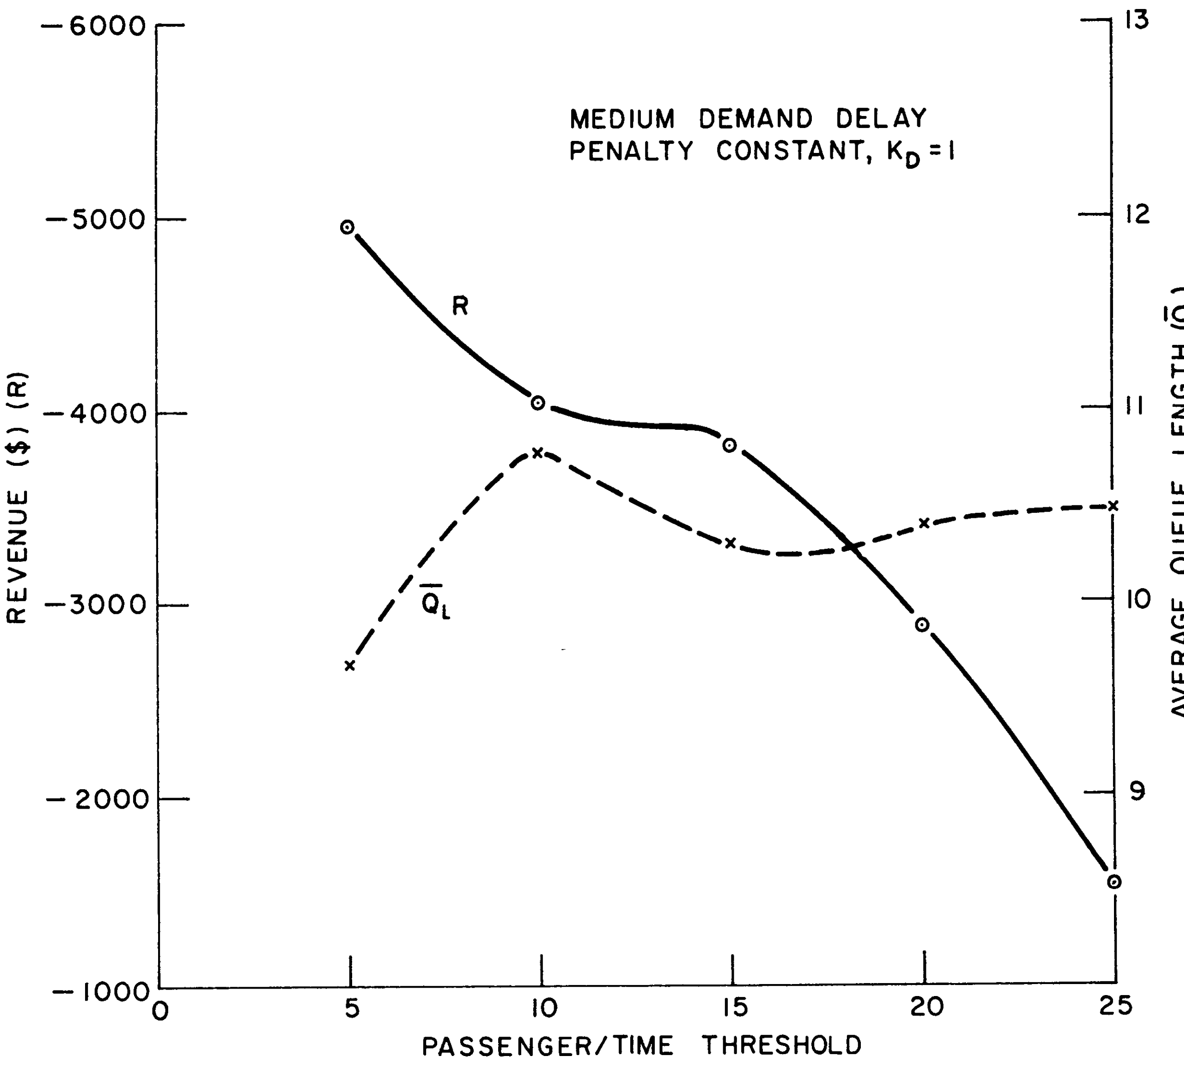

FIGURE 4.5: EFFECT OF PASSENGER GUARANTEES

performance is realized with higher thresholds. This is true since the higher the threshold, the fewer times it will become active. As this criterion is inherently uneconomic, less frequent resort to its use would be desirable. From the level of service viewpoint, less concrete statements can be made. For both the higher and lower demand there are pronounced fluctuations. These are attributed primarily to the randomness of the passenger generation mechanism of the model described in Section 4.1. The lower thresholds are conducive to better service since more flights are operated to service the same demand. The economic criterion is based upon expectations for the immediate future whereas the waiting criterion is a function solely of the current state at the station in question. If a station loses its only aircraft just prior to the arrival of a large number of passengers, it may be sometime before another aircraft can be made available to accommodate them. Thus large queues may result throughout the system, directly stemming from the early dispatches with small loads. To investigate this more thoroughly, passengers must be generated for each feasible sector at each time

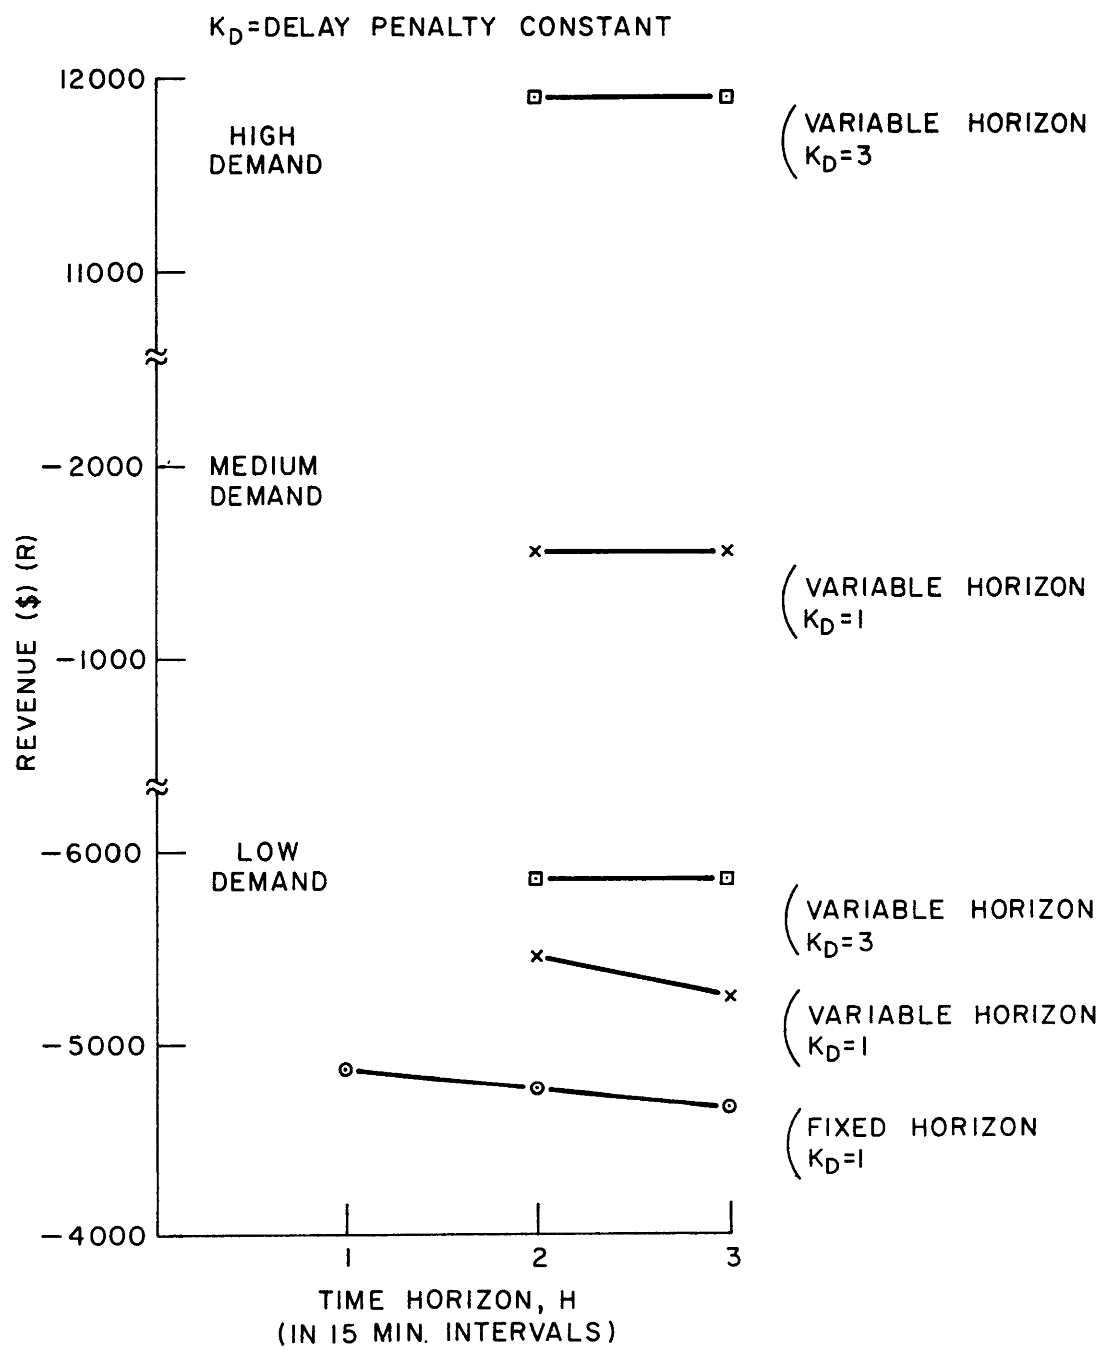

FIGURE 4.6: EFFECT OF THE TIME HORIZON

period and more numerous and longer runs made. From this, more reliable trends should be observable.

4.4.3 Effect of Time Horizon

Several runs were made to determine the sensitivity of the time horizons over which the economic decisions are evaluated. In runs 6 and 7, all parameters are identical except for the time horizon used during the peak periods of the day. In run 6, this is three periods (45 minutes); in run 7, it is two periods (30 minutes). The longer horizon resulted in a 4.4% improvement in revenues, precipitated by a slight improvement in load factors and an increase in the number of passengers carried. The reason for this improvement is not particularly obvious. Looking ahead further resulted in more flights being operated than was the case with the shorter view. The extra 27 passengers carried more than offset the cost of the extra flight and contributed to a reduction in losses.

A fixed (as opposed to a variable) time horizon was also tested. Figure 4.6 confirms the desirability of a longer horizon for the economic criterion. It appears that the longer horizon makes a smaller, yet real, contribution during slack times as well.

For comparison purposes runs were also made with the variable rule using a higher delay cost, i.e., three times that used in the runs described above. These too may be seen in Figure 4.6. In both the low and the high demand case there was no perceived change in performance with the use of the longer horizon. The horizon affects only the applicability of the marginal cost criterion. When the delay cost or demand exceeds a certain level, the relative effects of the longer horizon are masked by the magnitude of the delay penalty.

4.4.4 Effect of Number of Aircraft in the System

Previously, all runs were conducted with ten serviceable aircraft available in the system. To examine the effect of introducing additional aircraft, all other conditions remaining the same, several runs were made with 5, 15 and 20 aircraft randomly distributed among the nine stations. Comparisons were made at three levels of demand and are shown in Figure 4.7.

In the high and medium demand situations, low delay constant, a marked deterioration in performance resulted. With more aircraft distributed through the system, the aircraft in general departed with less passengers on board. This reduction in average load factors decreased net revenue per flight from $570 to $354 for the high demand case, and from -$72 to -$132 in the medium demand case. Level of service also deteriorated between 4 and 6%. In all cases only one additional flight was operated by the larger fleet, yet fewer passengers were carried. In general, however, as the number of aircraft increase to 20, revenues improve, but better service (lower average queue lengths) are not necessarily assured. One explanation of this rather surprising result is that with no 'coupling' of decisions between stations, the extra aircraft were very inefficiently operated. Further, with fewer aircraft, many dispatching decisions are barred due to unavailability of a vehicle. With more aircraft, more flights can be operated sooner. Lower load factors inevitably

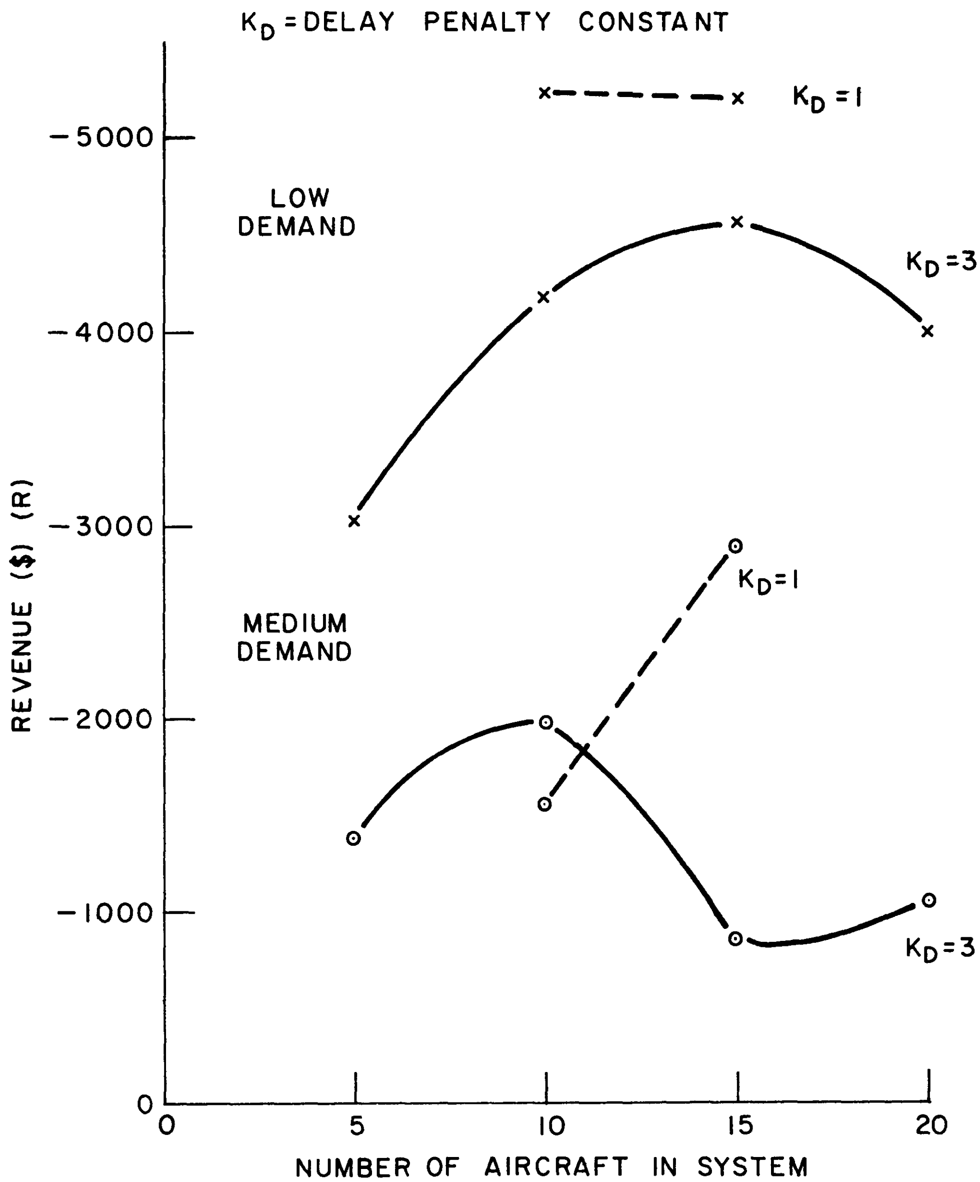

FIGURE 4.7: EFFECT OF FLEET SIZE

result. What is not seen, however, is the overall return function, i.e., fares less operating costs less delay penalties. As the dispatching decision is based upon this, the lower load factors and net return must be considered together with passenger delay cost -- a long term element not reflected in immediate revenues.

In the low demand case a slight improvement in performance was experienced with the larger fleet. This improvement is attributed primarily to less ferry requests and more use of the economic criterion.

With a higher delay constant, a substantially different performance is observed as seen in Figure 4.7. The higher delay penalty constant encourages earlier departures. This may improve or deteriorate performance depending upon the state of the stations involved. Few conclusions can be drawn on this without longer and many more runs.

4.4.5 Effect of Errors in Expectations

Attempts were made to determine the effect of gross errors in the passenger demand expectations.

From the few runs made it appeared that performance was relatively insensitive to errors in daily demand expectations, and that errors in the expected timing of passenger arrivals was much more serious. These findings, however, are considered inconclusive and considerably longer runs are needed.

In order to determine more conclusive trends in this model, it is essential that passenger arrivals be assigned to each city-pair for each time interval. Further, it is equally important that much longer runs be made to smooth out random effects and rounding off errors.

CHAPTER 5

THE AIR TAXI PROBLEM

The air taxi problem represents a different environment from those models discussed previously. The situation is typically characterized by a large number of stations served from some central point in the system. Although aggregate demand figures may be known with a reasonably high level of confidence, there exists more uncertainty with respect to demand at a particular station. Consequently fixed time schedules are generally replaced by maximum waiting time guarantees. A passenger, therefore, may be guaranteed a pickup within 30 minutes of his call or within 15 minutes of his appearance at the central point. These guarantees are quite vital for this type of operation, and it is essential that the operator be able to achieve the level of service expected by the passengers.

This problem is modelled after an existing air taxi company in the Boston area which serves suburban communities from a major central airport. Although

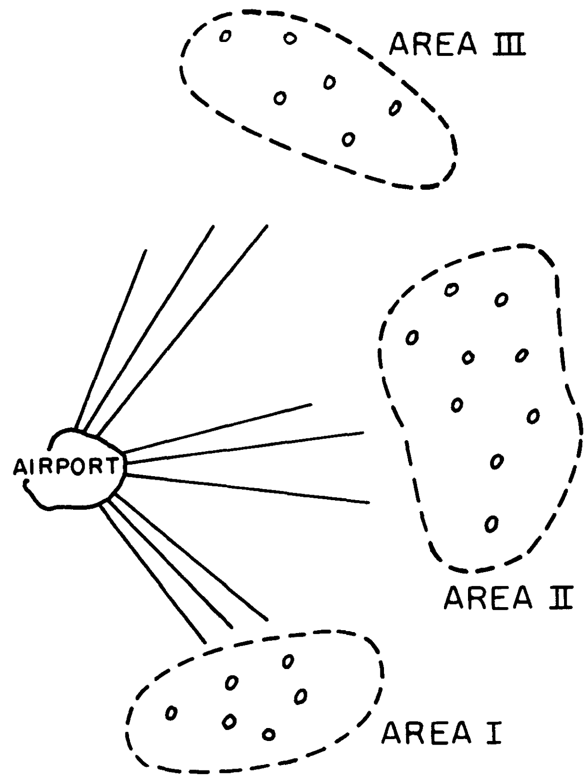

FIGURE 5.1. HELICOPTER AIRTAXI SERVICE

expandable into other uses, it has been designed primarily as a capacity planning model. The question explored is one of fleet size required to ensure a given level of service. The simulation approach is extremely useful in such cases where there is stochastic demand and a complex interaction of decisions. In addition, the sensitivity of the various decision variables to errors and random fluctuations can be readily examined. Relative profitability of any type of operation can be gauged.

5.1 Model Structure and Operation

The network consists of twenty outstations or pickup points, served from a central airport. For convenience, these stations are grouped into three distinct geographic areas as shown in Figure 5.1. This permits certain simplifying yet realistic assumptions to be made.

1. A vehicle dispatched to or from a given station will serve all stations in the common area, within its capacity constraint.

2. Flights to any station within an area have

the same fare and flight time structure.

3. Station to station movement within an area will charge an additional five minutes to the flight time for each such move.

4. No inter-area flights are permitted.

Passenger arrivals are considered at two points.

a) At the Airport with the destination being an outstation, determined at random

b) At an outstation, also determined in a random manner, with destination being the Airport.

There is assumed to be no inter-station traffic. Passenger arrivals were generated in several different ways; first with a uniform arrival rate and, in later runs, with the time of day variations included (Figure 5.2). The effect of demand shifts and general uncertainty were examined by varying the method of generating traffic. In general

$$\text{Interarrival Time} = \text{mean} \times \text{Bias} \qquad (5.1)$$

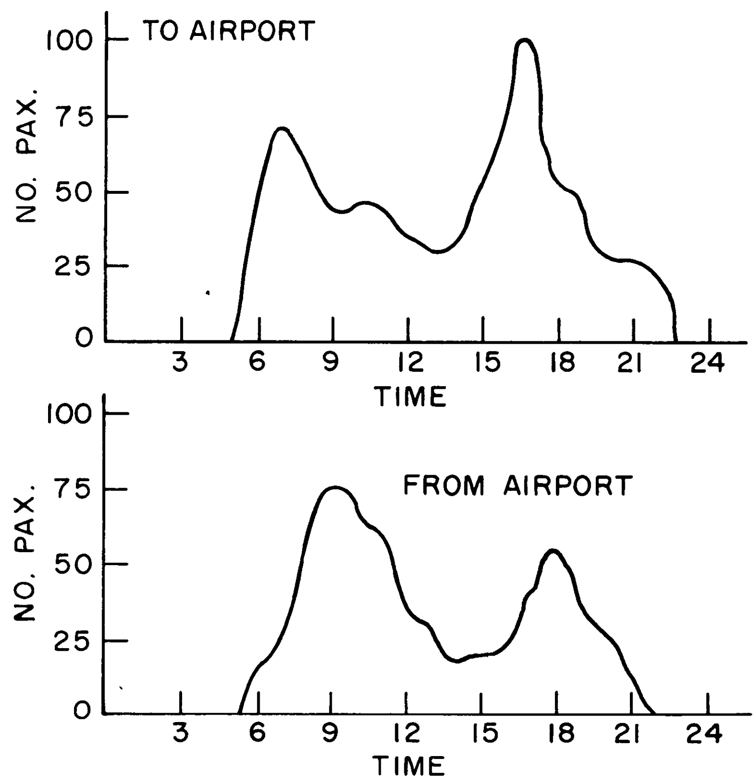

FIGURE 5.2 DAILY VARIATION IN DEMAND TO AND FROM AIRPORT

the bias being a function of the time of day variation distribution.

Queues are designated by area and not by individual stations. Although statistics on queues, with average waiting times and delay distributions, are tabulated, no explicit delay penalty is included in the dispatching criteria. The nature of the model makes this unnecessary since a level of service is first postulated and the incremental capacity required to achieve this is, in essence, created. It is in this way that the capacity requirements are determined.

The fare, costs, and flight time structure assumed for this problem are shown in Figure 5.3.

5.2 Decision Rules

The following decision rules are built into the model.

1. Generally, every customer is made to wait at least 15 minutes but not more than 30.

2. A single aircraft should cater for all

Airport to ___	Units	Area 1	Area 2	Area 3
Passenger Fare	$	8	12	16
Block Time	Mins.	10	20	30
Operating Cost	$	1.25 x Block Time		
Intra-Area Transfer Time	Mins.	5 minutes/station		

Figure 5.3 Fare/Cost/Flight Time Structure

- Air Taxi Service

business in a particular geographic area up to its maximum physical capacity of 3. This small capacity vehicle is typical of a taxi system. It would be interesting to increase the scale of demands and vehicle size to see the effect on operating strategies.

3. In the peak hours between 6:30 A.M. and 8:30 A.M. and between 3:30 and 5:30 P.M., an aircraft will remain at its last outstation for 30 minutes in anticipation of one or more passengers materializing from somewhere in the area. Otherwise it will return immediately after discharging its passengers.

4. In all cases, when a full capacity load is available, the aircraft will depart immediately.

5. When a passenger arrives at an outstation with no aircraft available

in the area and no aircraft en route to the area, he will be held 15 minutes prior to calling for a ferry from the airport.

6. If an aircraft is not available at the outstation where the passenger is and several aircraft are on the ground at other outstations in the same area, one of these will be selected and ferried to the desired point of departure. The criterion for this selection is the one which has been grounded the longest.

7. If a flight (either revenue or ferry) is en route to an area, no further ferry calls will be made from that area.

5.3 Output Format

Operation of the model yields a departure schedule for the revenue flights and a count of the total number of aircraft in operation (i.e., in flight or at outstations) at the time of each departure. Auxilliary outputs include:

1. The number of revenue and ferry flights
2. The total accrued profit that day (fares less costs)
3. The total number of daily passengers arriving at the airport and at the outstations
4. The passenger interarrival time distributions
5. Queue statistics by geographic areas.
6. The location of aircraft at the end of the run.

Several runs were made varying the critical variables to determine the sensitivity of the solution to errors and normal variations. The critical variables considered here were the passenger interarrival times and certain decision time variables, such as waiting times prior to dispatching an aircraft or calling for a ferry.

TABLE 5.1 RUN SUMMARY AIR TAXI MODEL

Run No.	Time Decision Var.	Mean Inte Arri Time	No. Flts. Rev	No. Flts. Ferry	Pax. Carried L → o/s Outbound	Pax. Carried o/s → L Inbound	Pax. Carried Tot.	Load Factor w.r.t. Rev. Flts.	Load Factor over-all	% Ferry Flts.	Utilization Pax/AC	Max. # AC Req'd	Tot. Rev. $
1	Original WTL = 15 WTO = 15 WTC = 15	20 Uni.	82	37	45	48	93	.375	.260	31.0	15.5	6	Loss 2068
2	WTL = 20	20 Uni.	69	29	44	47	91	.440	.310	29.5	18.2	5	1529
3	WTL = 20	20	79	50	53	75	128	.540	.330	38.7	18.3	7	2046
4	Original	20	94	56	54	75	129	.457	.286	37.3	16.1	8	2472
5	Original	30	70	31	37	49	86	.410	.284	30.7	14.3	6	1682
6	Original	15	104	56	72	98	170	.544	.354	35.0	18.9	9	2420
7	WTL = 20 WTO = 20 WTC = 20	20	83	41	53	75	128	.513	.344	33.0	18.3	7	1892

*WTL = Waiting Time @ Airport Prior to a Departure.

WTO = Waiting Time @ Outstations During Peaks Prior to Ferry Return.

WTC = Waiting Time @ Outstations Prior to Requesting Ferry Dispatch from Airport.

5.4 Discussion of Runs

A summary of the simulation runs may be found in Table 5.1.

5.4.1 Effect of Demand

Any estimation of passenger demand is bound to be in error. To determine the sensitivity of the number of aircraft required to changes or errors in this factor, three runs were made using identical dispatching rules, but with the mean passenger interarrival times changed. Comparing runs 4 and 5, when the mean is increased by 50%, i.e., less frequent arrivals and fewer total passengers, the number of aircraft required drops 25%, from 8 to 6. Further, profit showed a 32% improvement (less loss). This improvement was the result of a reduction in the percent ferry flights required. Load factors did not change appreciably.

In runs 4 and 6 the effect of a 25% reduction in the mean interarrival time was investigated. One extra aircraft was needed to operate the service for this demand. The increase in the number of passenger arrivals in a unit of time permitted better load factors to be realized. The percent ferry flights showed a slight improvement resulting also in a marginal 2% improvement in profit. Compared to run 5, the difference in the number of ferries required is a bit surprising. One would think that with more passengers, a greater 'packing' of passengers might be achieved with corresponding less ferrying needed. On the other hand, however, less demand also implies less ferries. These two factors apparently offset one another.

In summary, over a fixed time span, as interarrival times decrease, the number of aircraft required increase in a nonlinear fashion. Also average load factors and aircraft utilization tend to improve. However the number of ferry flights necessitated depends upon the relative buildup of traffic at the various stations in the system. Of course, revenue is strongly affected by the number of ferries. More runs of this nature would be required to establish more concrete relationships.

5.4.2 <u>Effect of Time of Day Peaks</u>

Runs 1 and 4 show the relative effects of the time of day variations in demand on the number of aircraft required. The peaks experienced in a normal day's operation call for more aircraft than would otherwise be needed (8 Vs 6). More passengers are served in run 4 and the aircraft daily load (no. pax per aircraft) is roughly the same. However, about 20% more ferry flights are required to meet the quality of service desired, resulting in 19% more losses as compared to the uniform arrival rate case of run 1. Note that in this case the peaking of traffic increases the probability of improving load factors to or from any particular area. Again, however, it is the relative times of peaking at the two points with respect to the decision time variables that is important.

To illustrate this last point, consider runs 1 and 2. Both runs are with identical uniform arrival rates. In run 2 the decision time variable with respect to the waiting period at the airport was changed to coincide with the mean interarrival time. As might be expected this permitted the possibility that the departing flight for a particular area might carry two passengers instead of just one; that is if both passengers happened to be going to the same area. The result was a reduction by one of the number of required aircraft, increased aircraft daily load and a 26% improvement in revenue.

5.4.3 Decision Time Variables

Runs 1 and 2 and runs 3 and 4 demonstrate the effect of a five minute variation in the dispatching time criterion from the airport. Instead of delaying passengers a minimum of 15 minutes prior to departure (runs 1 and 4), this wait was increased to 20 minutes. Runs 1 and 2 are with a uniform arrival rate while runs 3 and 4 take into account the realistic daily demand fluctuation.

In both instances a savings of one aircraft results. In addition, at least a 13.5% improvement in aircraft average daily load results, with 17% improvement in revenues.

Runs 4 and 7 indicate the combined effect of changing three time variables:

1. Waiting time at the Airport prior to a departure = WTL
2. Waiting time at an outstation prior to a ferry return, during a peak period = WTO'
3. Waiting time at an outstation prior to calling for a ferry flight to accommodate passengers = WTC

In Run 7 the above were changed from 15 minutes to 20 minutes. It may be seen that a savings of one aircraft results. Since the same number of passengers is carried, a higher aircraft daily load is achieved, 18.3 pax/aircraft as compared to 16.1 in run 4. This improvement is also seen in the overall load factors (.344 vs. .286). Further a 17% improvement in revenue is realized.

It is interesting to compare these improvements with those achieved by changing only WTL as in run 3. In both 3 and 7, the same number of passengers was transported with the same size fleet. Although load factors were about the same, the number of ferry flights was reduced by 18% in run 7, resulting in an improvement in revenue of over 6%.

It may be safely said that increasing the waiting times does not necessarily always improve performance, even in the absence of a penalty for delaying passengers. For example, as seen in runs 3 and 7, delaying an aircraft at an outstation in anticipation of a demand building up in the area within the wait period could result in an adverse situation at another point. More aircraft would be out at any point in time, therefore, more aircraft would have to be activated at the Airport to accommodate incoming passengers.

5.4.4 Summary

It is the interaction of the decision time variables together with the timing of the passenger arrivals at the various points in the system that dictate the capacity requirements. Only a few of the variables have been examined here. There are many more.

Generally the model tends to verify what appears to be intuitive. Its value, however, is in applying a quantitative interpretation to various decisions, the interaction of which is not necessarily clear. Several rules of thumb are suggested, however, the impact of these too may best be evaluated by simulation.

Many more runs would be needed before a concrete recommendation can be made on the capacity question. Management must establish and evaluate a large number of policies. In certain cases, it may be necessary to compromise a quality of service, a waiting time, for the sake of a substantial savings in capital equipment, or for better utilization. The establishment of these policies must be made by management; the evaluation of these can be readily undertaken by a simulation model such as that described and used here.

The model may be easily stretched to include many more stations and each with different fare and flight time requirements. However, certain assumptions with respect to the area groupings should be retained for expediency. It is also possible to generate passenger arrivals in many fashions, poisson, monte carlo, etc.

To make the model more realistic, perhaps a passenger delay penalty should be included to reflect the natural consequences of indiscriminate delays, planned in the name of improved performance. Of course there are many other decision variables which may be of importance in different situations and may be incorporated. There appear to be no restrictions to including other decision rules in the model, within the physical limitations of the GPSS language.

The language of aircraft dictated by this model implies 'serviceable' aircraft. In reality, aircraft reliability and the maintenance schedule must be considered in interpolating between serviceable aircraft and actual fleet size.

CHAPTER 6

THE ADAPTIVE DECISION MODEL

6.1 Introduction

Historically aircraft dispatching policies have been administered by individuals who have repeated the task over and over again. In the course of this repetitious process they have become increasingly more proficient and sensitive to the critical variables of the problem. They have been further assisted by a large amount of historical data assimilated for eventual use in this task. However the limited ability of humans to rapidly absorb diverse data and to extract relevant statistics is well known. Typically by the time this data has been sifted and put into usable form, the true existing situation may well have changed to the point where this new information is no longer of significant value.

Characteristically the decision rules are formalized by specifying some numerical threshold which tends to remain relatively fixed over long periods of time. Pronounced deviations from expectations are probably noted and alternative action taken, but gradual changes may go largely undetected. Further, varying conditions at other stations which may bear upon the dispatch decision at this station are largely ignored. The ideal one would hope to achieve, therefore, is a dynamic threshold whose value would change in response to changes in system state and expectations.

Given perfect information on future demands -- time, place, and number, "optimal" fixed scheduling would be readily achievable. It is the uncertainty about the demand and inadequacies of available expectations that render the task 'sub-optimal' and dynamic. Though we cannot hope to know the future explicitly, perhaps by judicial assimilation and application of historical data, the optimal may be approached.

The object of this Chapter, therefore, is to incorporate an adaptive decision approach such that historical information, compiled in real time, can be used in specifying a truly dynamic rule. This heuristic would be sensitive to the critical system variables such as current aircraft disposition as well as to demand expectations for the following time periods. The task, therefore, involves building into the model a simple 'learning' capability, to capture, and have available for immediate use, pertinent assimilated data. Essentially a Bayesian approach is then applied to modify the current rule for tomorrow's use, based upon what the model has experienced today, and in previous days.

To exemplify this problem, the air taxi model of Chapter 5 was selected. Structurally all elements remain the same, however, considerable modification was required to build in the necessary changes in decision rules and information storage and manipulation.

6.2 Model Operation

Operation of the model commences with the simultaneous generation of passengers at the central airport and at the outstations. The interarrival times are established as described in Chapter 5. Passengers queue up by area to await satisfaction of one of the dispatching criteria. Here the adaptive decision model differs from the standard air taxi model of the preceding Chapter. Whereas in the latter a fixed maximum passenger waiting time threshold is set and applied throughout the day, in the adaptive model a dispatch threshold is recalculated each time a departure is considered. Its value depends upon certain expectations with respect to passenger demand at the stations immediately affected and to the current disposition of aircraft in the system.

Also the simulation time interval was changed from one minute to five minutes to make the program more efficient for the longer runs required. For the purpose of scheduling in this problem, a day was considered to be 16 hours, from 6:00 A.M. to 10:00 P.M. It is important to note that whereas the model runs for several days, conceptually, these days may well represent the performance of selected average days, e.g., a Monday over several consecutive weeks or months. This is relevant in that traffic patterns are likely to be cyclical over the course of a week, each day showing a distinctive pattern.

6.3 The Learning Process

During the course of the day, data on demand is compiled and stored away in computer locations or cells, which have been previously assigned. The address of these storage cells is a unique number indicative of the time of day (in ten minute intervals) and the direction of the demand, i.e., inbound or outbound from the airport. The contents of the cell is the number of passengers travelling in the indicated direction during the particular time interval. Thus, what is in effect a demand distribution is being generated throughout the course of the day. Simultaneously, a count is maintained of the total number of passengers travelling in each direction to or from the individual areas. This count is later used to establish relative probabilities.

At the commencement of the simulation run, the demand storage locations and probabilities are initialized with a priori values. These may be based either on past data or our own subjective expectations.

At the termination of the day, the distribution and probabilities are updated using, in essence, a

Bayesian approach. For example, at the end of the first day a demand storage location for outbound traffic between 7:00 and 7:10 A.M. would contain a number representing the sum of the a priori value and the number conforming to today's arrivals occurring during that time interval. Although the relative weightings assigned to the priors and the current measures may vary as we choose, here they were considered to be equally weighted. That is, the posterior value of the passenger function

$$= \frac{\text{a priori (or prior) + current}}{2} \qquad (6.1)$$

Therefore, updating at the conclusion of the day requires simply a division by two. The current passenger arrivals are being added to the prior in real time, but as the distribution is referred to only to determine expectations for future time periods, only the prior (yesterday's posterior) enters into today's calculations.

With the day's operations complete, the program also calculates six probability measures based upon the particular day's performance.

1. Prob(Inbound passenger originates in area 1)

2. Prob(" " " " " 2)

3. Prob(" " " " " 3)

4. Prob(Outbound passenger is destined for area 1)

5. Prob(" " " " " " 2)

6. Prob(" " " " " " 3)

Of course, the first three and the second three probabilities sum to one. These measures are the daily average probabilities of traffic flow to or from the individual areas. During the simulation they are used to update the prior (or a prior) probability measures. Here again, an equal weighting was assumed, thus attributing greater importance to the more recent information.

$$\frac{\text{Prior + Current}}{2} \doteq \text{Posterior}$$

The posterior becomes tomorrow's prior and is used in the dispatching decision.

6.4 Decision Heuristics

Basically there are three criteria comprising the dispatch rule at the airport.

1. A capacity load is on hand

2. A ferry call has been received from an outstation. In such cases, any waiting passengers for that area will be taken.

3. The calculated dispatching threshold has been exceeded as explained in 6.5.

For the outstations the same general rules apply except that an additional heuristic criteria is established to govern the decision to call for a ferry flight from the airport. This criterion states that the total number of aircraft at or en route to the area in question are insufficient to meet the current demand. If such is the case, another ferry will be called provided that this action does not deplete the airport of its last aircraft. If only one aircraft is available at the airport the ferry request will not be immediately satisfied.

6.5 The Dispatching Threshold

The major decision heuristic of this simulation consists primarily of determining a numerical 'threshold' for passenger maximum delay time. When a passenger's waiting time exceeds this value, the flight will be dispatched. The threshold consists basically of a weighted average of two components which are functions of:

1. Relative demand expectations
2. Aircraft disposition

It is in the first component where the 'learned' distributions and probabilities are used. The second is a function of the current state of the system with respect to aircraft location.

6.5.1 Relative Demand Expectation

A fundamental and obvious objective in aircraft scheduling is to have the aircraft available where the demand is most likely to occur. In a stochastic system with many stations and limited aircraft a calculated risk is involved in achieving this objective. Better knowledge of demand expectations (where, when and how many) would presumably result in improved scheduling performance. In the decision to dispatch from the airport to area 2, for example, neglecting all other factors, two demand expectations are of interest.

a) the expected number of passengers arriving at the airport for travel to area 2 during the next time period (here next ten minutes)

b) the expected number of passengers arriving at area 2 for travel to the airport during that time period in which this flight will arrive if dispatched immediately. The passengers arriving in intermediate time periods are assumed to have been accommodated by other aircraft in the area.

Clearly there is much room for refinement in these rules, particularly with respect to the size of the time period considered. This could be made a variable function of time of day; however for the purpose of this exploratory study, the above simpler version is considered adequate.

Let the expected number of total outbound passengers be E(O) and the expected number of total inbound passengers be E(I). Further, let P_I (A2) be the probability for this time of day that an inbound passenger originates in area 2 and P_O (A2) be the probability that an outbound passenger is destined for area 2. Then, as the ratio

$$r = \frac{E(O)\ P_O\ (A2)}{E(I)\ P_I\ (A2)} = \frac{\text{expected \# of pax outbound to A2}}{\text{" " " " inbound from A2}} \tag{6.2}$$

varies, the corresponding component of the waiting time threshold varies. As seen in Figure 6.1 the greater the value of r, all other things being equal, the more desirable it is to delay the departure of the flight, since we are expecting relatively more arrivals at the airport. Therefore, the relative demand expectation component is allowed to increase to a maximum value ℓ which is an upper limit for passenger waiting times at the airport.

6.5.2 Aircraft Disposition

In a decision to dispatch an aircraft from the airport to area 2, the percent of the serviceable fleet at or en route to the airport and the percent at or en route to area 2 are both significant. More explicitly, if a very small percentage of the fleet is at the airport, we might wish to delay the flight. If the percent of the fleet at or en route to area 2 was high we might also wish on this count to delay the flight at the airport to avoid a high concentration of aircraft in area 2. These two factors of the aircraft disposition component, Figure 6.2, are weighted equally here but need not be.

6.5.3 Calculation of the Waiting Time Threshold

The two components of the threshold, aircraft disposition and relative expected demand, may be weighted differently depending primarily upon the criticality of the aircraft capacity available to the airline.

Clearly the lower limit of this threshold is

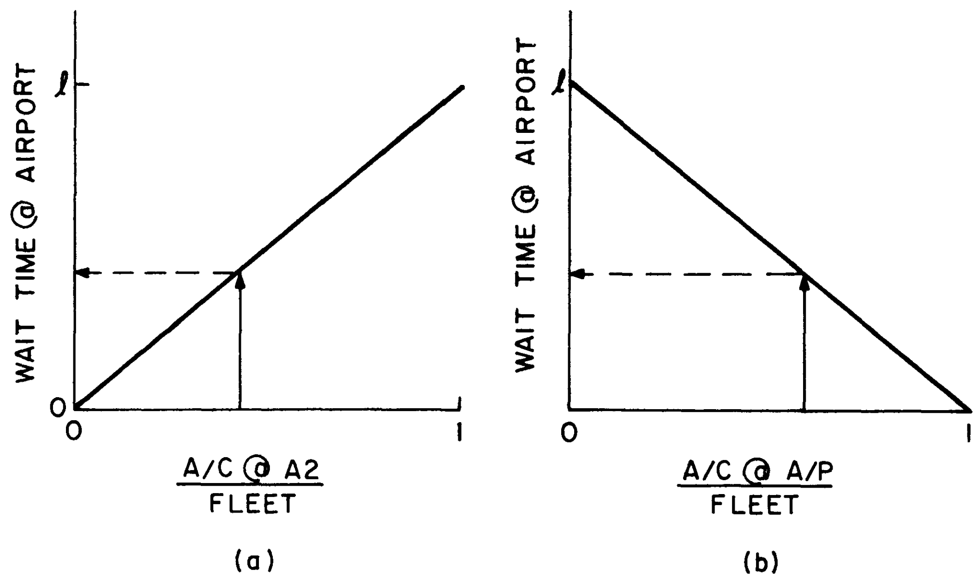

FIGURE 6.2: AIRCRAFT DISPOSITION COMPONENTS FOR DECISION COUPLING

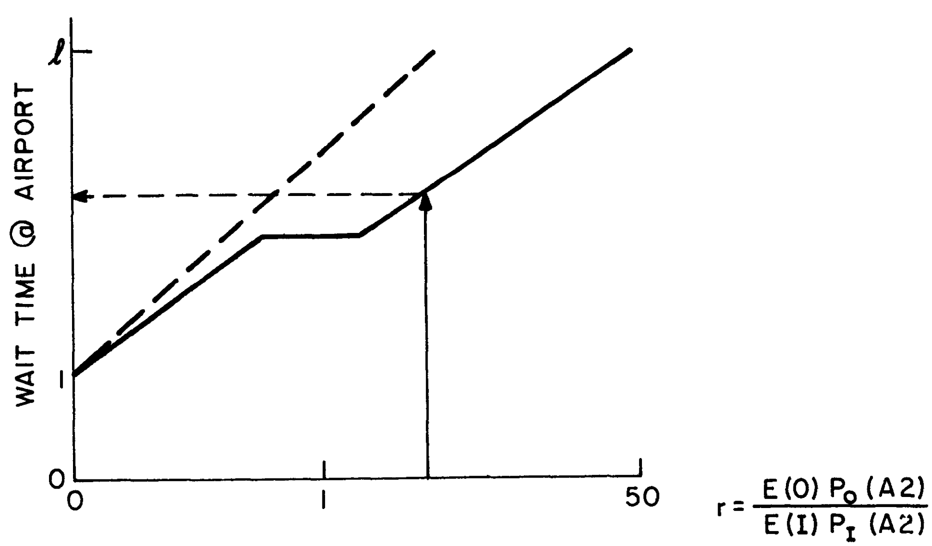

FIGURE 6.1: RELATIVE PASSENGER DEMAND COMPONENT FOR DECISION COUPLING

zero indicating an immediate departure. The upper limit ties in with the maximum waiting time guarantees, ℓ, specified by the company. These limits may be different for different stations. Here they are taken as 20 minutes at the airport and 30 minutes at outstations.

6.6 Discussion

Runs for this model were designed to reveal any advantage the adaptive decision criteria may hold over a conventional fixed threshold. In the evaluation two parameters were considered

1. A measure of economic performance -- net revenue (fares less operating costs)
2. Two measures of passenger service --
   a) Total passenger-minutes spent in waiting
   b) A calculated delay penalty

Runs were made with the guaranteed waiting times fixed at the maximum and other intermediate values. The system performance with these rules was compared with the performance where the threshold was allowed to adapt itself according to two components -- the past trends in demand experienced and the current system state with respect to aircraft location.

TABLE 6.1 SELECTED RUN SUMMARY

ADAPTIVE DECISION MODEL

Run		8	9	10	11	12	14	15	16	17
Airport	Demand	/	/	1/3	0	0	1/2	1/2	/	1/3
Threshold	AC Pos (a)	40	20	1/3	0	1/2	1/2	0	30	1/3
Comp. Wgt.	AC Pos (b)	/	/	1/3	0	1/2	0	1/2	/	1/3
Ferry	Demand	/	/	1/3	1	0	1/2	1/2	/	1/3
Threshold	AC Pos (a)	60	30	1/3	0	1/2	1/2	0	45	1/3
Comp.Wgt	AC Pos (b)	/	/	1/3	0	1/2	0	1/2	/	1/3
Outstation	Demand	/	/	1/3	1	0	1/2	1/2	/	1/3
Threshold	AC Pos (a)	60	30	1/3	0	1/2	1/2	0	45	1/3
Comp. Wgt.	AC Pos (b)	/	/	1/3	0	1/2	0	1/2	/	1/3
Tot. Rev. ($)		103	-2848	-3860	-4710	-2848	-4850	-4193	-1031	-4094
Tot. Delay (min)		4084	3609	3288	3362	3608	3623	3541	4034	3277
$\sum P_W T^2$		20020	15650	13150	13802	15650	15675	15790	19645	12952
# Ferry Flts.		75	116	126	137	116	142	132	85	130
No. Rev. Flts.	Outbound	225	263	273	274	263	262	267	242	273
	Inbound	204	227	225	223	227	211	221	211	225
No. Pax Carried		855	858	862	864	858	870	858	843	864
Remarks		fixed	fixed		Demand Comp. Only	A/C Pos Comp. Only			fixed	Comp. Slope Moded

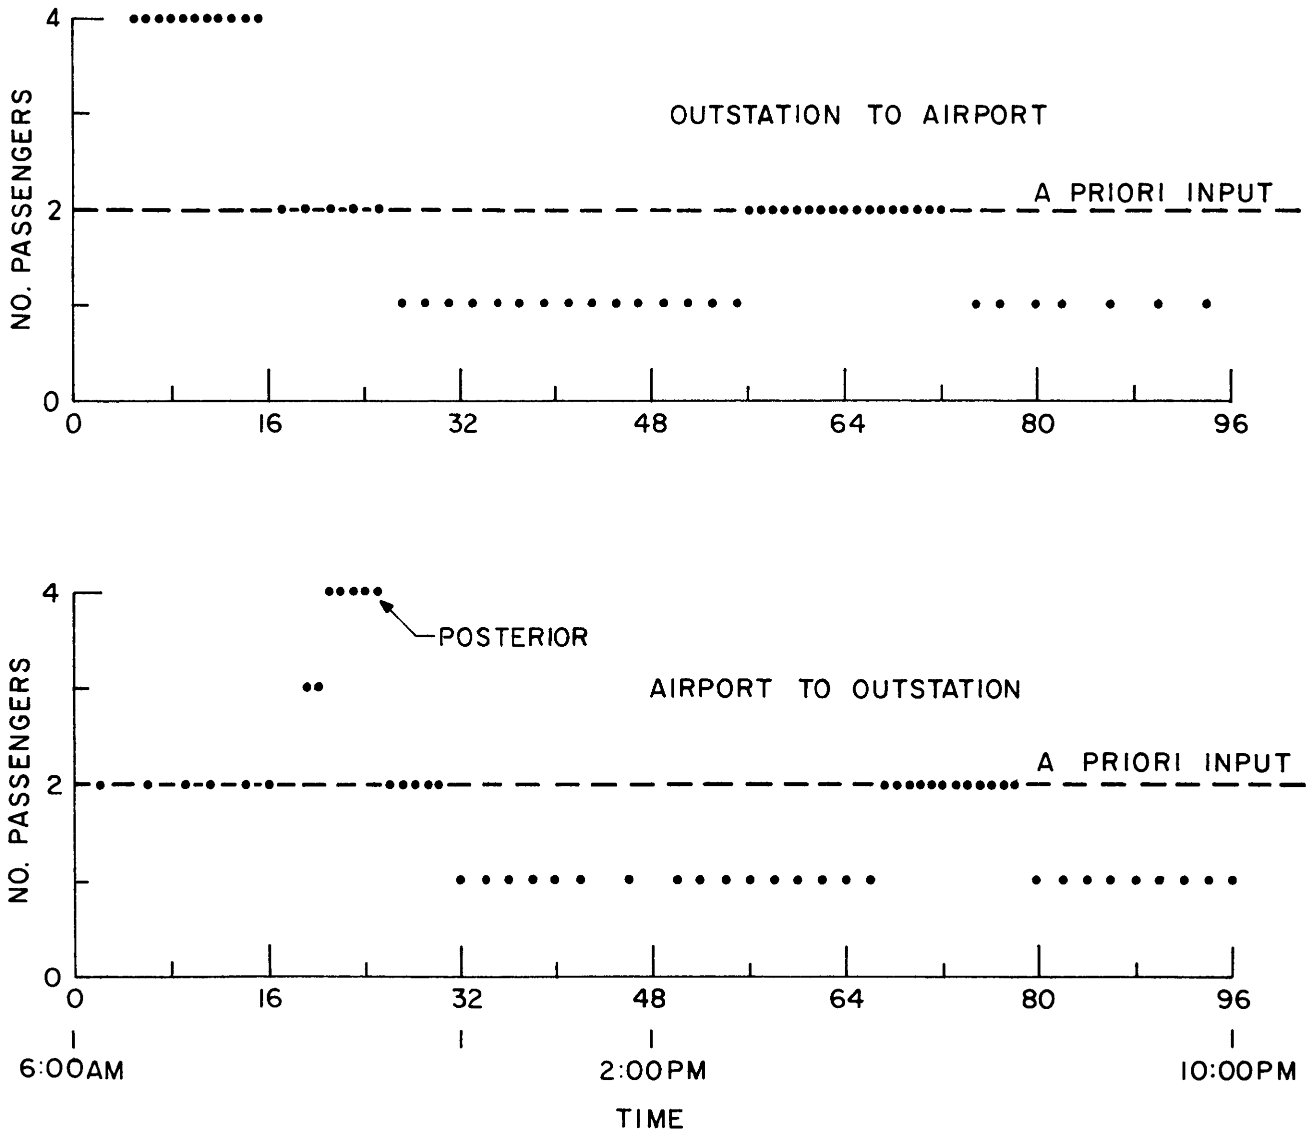

FIGURE 6.3: PASSENGER DEMAND DISTRIBUTION
AFTER 5 DAY RUN
—ADAPTIVE DECISION MODEL

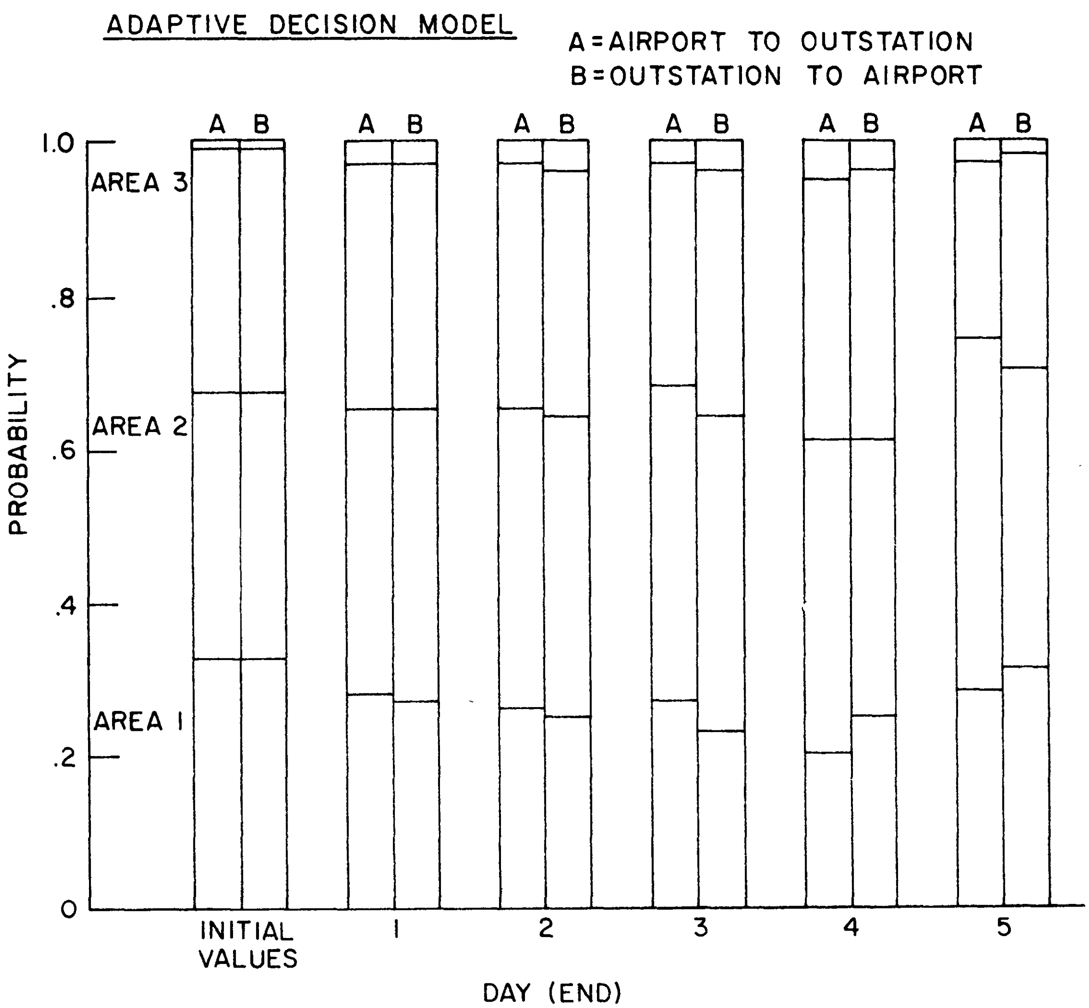

FIGURE 6.4: PROBABILITY MEASURE CHANGES RESULTING FROM PAST DEMAND EXPERIENCE.

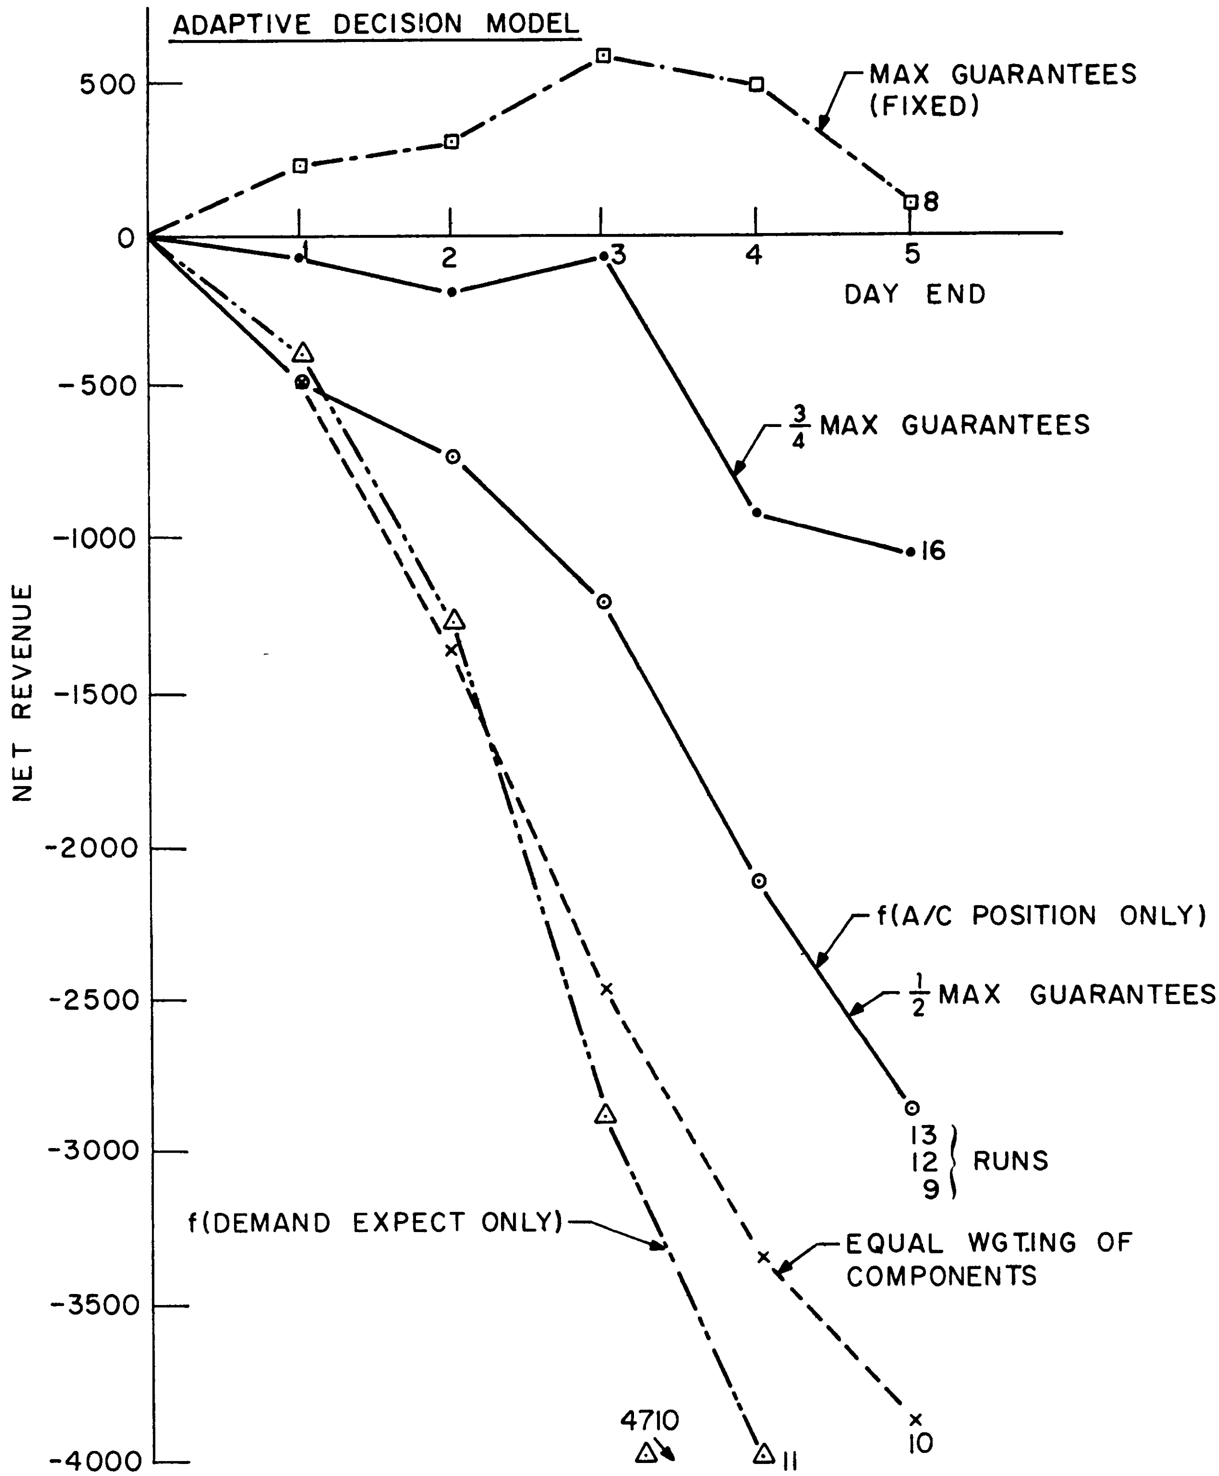

FIGURE 6.5: CUMULATIVE NET REVENUES

Several such runs were made placing different weightings on the two components. A summary of runs is tabulated in Table 6.1.

In addition to comparisons between fixed and adaptive criteria, the general capability of the system to adapt itself to changes was observed. In Figure 6.3 final demand distribution is compared with the a priori distribution which was preloaded. Figure 6.4 shows how the probability measures changed for each day of the five day run.

The results of some of the runs are plotted in Figure 6.5 (revenues) and Figure 6.6 (waiting times).

With the thresholds held constant at the maximum guaranteed values, the system requires substantially fewer ferry and revenue flights to accommodate the same volume of traffic. This apparent efficiency shows up in the accumulated revenues for the five day period. However, as seen in Figure 6.6 the fixed maximum threshold provides the worst service to the passenger. The reverse is true for the adaptive

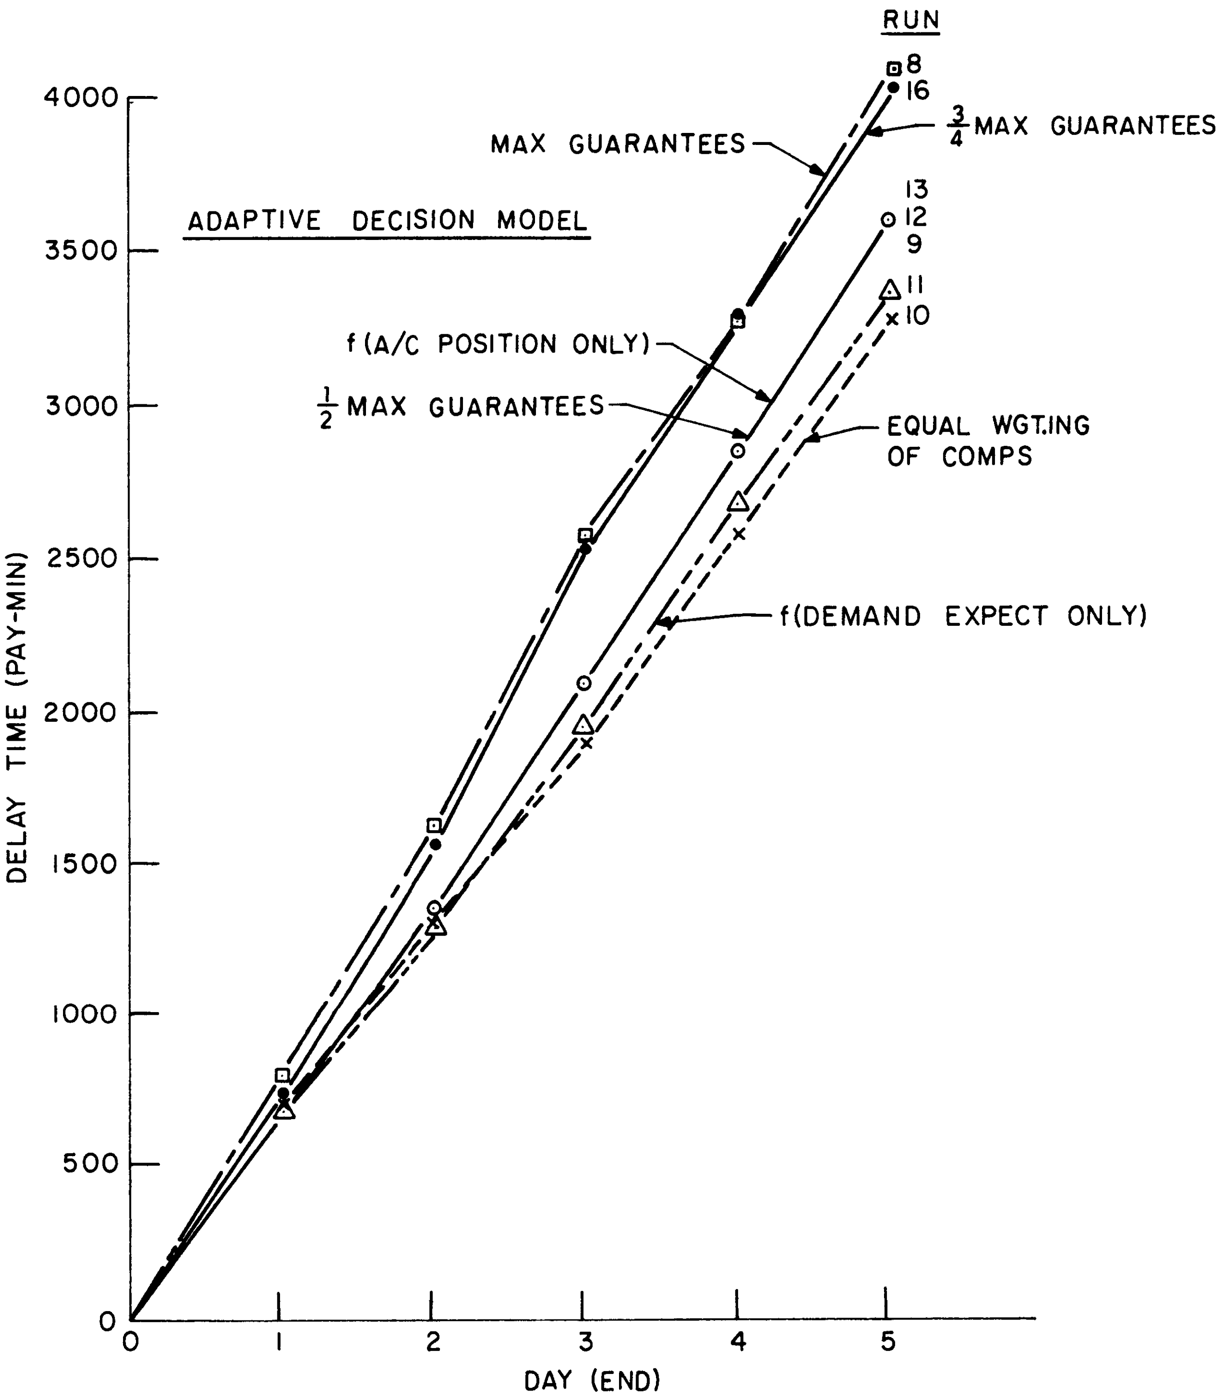

FIGURE 6.6: CUMULATIVE PASSENGER-MINUTES WAIT

wait time threshold, where revenues show a large loss but passenger service is improved.

To determine whether the adaptive threshold is really better than a fixed, it is again necessary to place a value on the wait time of the passenger. There appears to be sufficient justification to assume that this penalty should be rather heavily weighted. First, there is a relatively large fare differential between the air taxi and the conventional taxi or private automobile, whereas the difference in actual 'block time' may be only 15 to 20 minutes.

Consequently, a passenger delay of that order could be sufficient inducement to cause him to seek alternate means of transportation. Secondly, inbound passengers in most instances are destined to the airport to make onward flight connections and outbound passengers are often busy executives with tight schedules. In both cases excessive wait times may have adverse long-run and short-run effects for the company.

Consistent with our previous assumptions, the delay penalty is assumed to vary directly as the square of the waiting time and inversely as the fare. For simplicity the average fare is used and absorbed into the delay constant, K_D, to give

$$\text{Delay Penalty} = P_w \, K_D \, (\bar{T})^2 \qquad (6.3)$$

where P_W = Number of passengers waiting

$\bar{T}$ = Average waiting time in the queue

The total delay penalty for all queues is

$$\sum_{\text{all queues}} P_w \, K_D \, (\bar{T})^2 \qquad (6.4)$$

The net return to the airline may be taken as

$$R = \sum \text{fares} - \sum \text{operating cost} - \sum_{\text{all queues}} P_w \, K_D \, (\bar{T})^2 \qquad (6.5)$$

In Figure 6.7 this function is plotted against various values of the constant, K_D, for runs with different thresholds. The 'critical' delay constant shown is that value of K_D at which the adaptive threshold and the maximum fixed threshold are equivalent in overall performance. For this particular

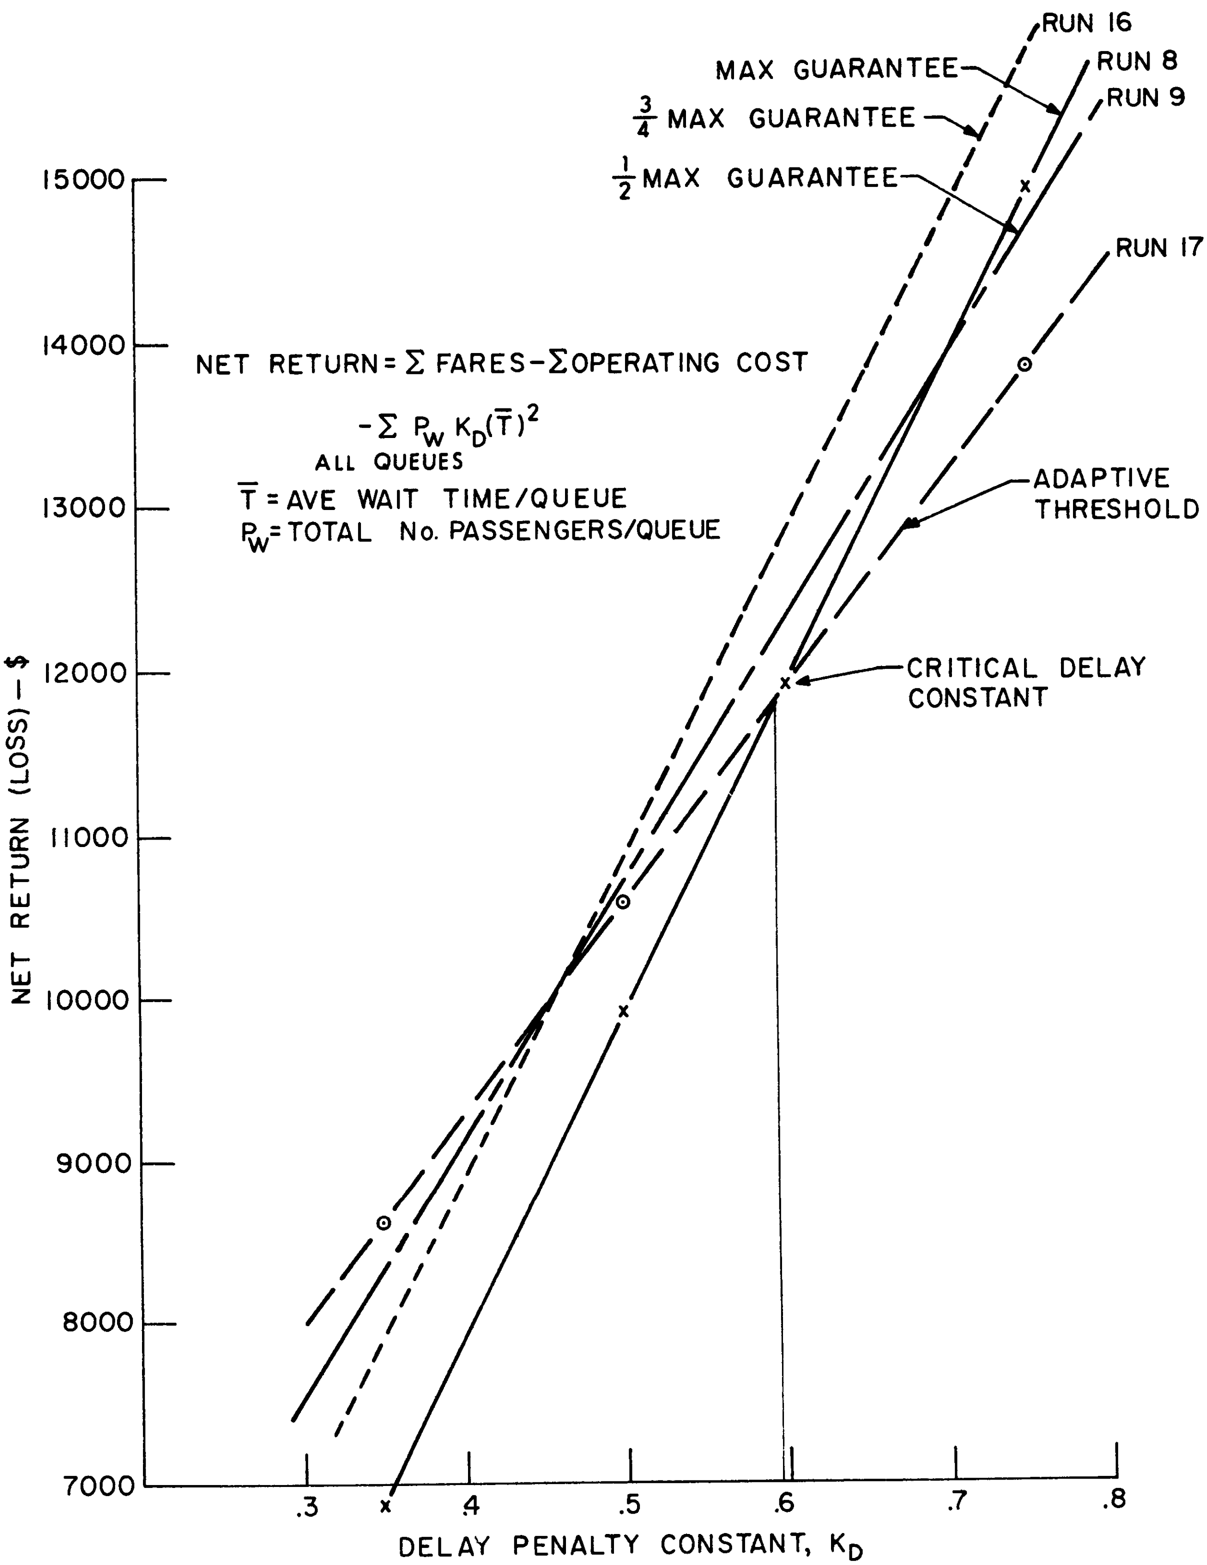

FIGURE 6.7: NET RETURN FUNCTION

case, if K_D is greater than .595 the adaptive threshold is superior to the fixed threshold. Note that it is also superior to other fixed thresholds which are less than the maximum (½ and 3/4 of the maximum value). This critical delay constant of .595 is equivalent to assessing a penalty of $2.40 for a ten minute wait or $5.40 for a fifteen minute wait. This may seem excessive for an average fare of $12. However, with flight time averaging only 20 minutes and for reasons mentioned above, perhaps this is indeed reasonable.

Improvements in the adaptive threshold may be possible through further analysis. For example, changing the shape of Figure 6.1 to the dotted line shown, resulted in a 1.65% improvement in net return (for $K_D = .6$)

In this model a fixed fleet size of eight aircraft was assumed. Often limited capacity results in an increase in the 'effective' threshold, since regardless of the value set on the threshold a flight cannot be operated unless an aircraft is available at the station. This 'effective' threshold is more apparent in situations calling for a low value, as may be detected from the average passenger wait time and maximum queue lengths. The overall effect is to smooth out or minimize differences and tends to mask the true effect of a specified threshold difference. To eliminate this influence runs should be made with increase in fleet size.

6.7 Summary

The adaptive decision approach to the air taxi dispatching problem is intended to provide a reasonable balance between economic returns and passenger service, within the maximum wait time guarantees. It is of value, therefore, only insofar as passenger waiting times are of substantial importance. If not, maximum wait times obviously yield the best overall returns.

In this study the adaptive decision technique has been used in a limited way, i.e., to set a threshold on passenger wait time. There are other applications which merit investigation. For example, the use of expected demand to pre-position aircraft at various points in the system, in the absence of any actual passenger calls, may be highly desirable. Additional investigation is also warranted into the shapes of the component curves in Figures 6.1 and 6.2. Although the trends are realistic, changes in slope may yield further improvement.

In this model it has been assumed that relative demand expectations and aircraft disposition are the only dynamic factors of importance in the decision to dispatch a flight. In reality there may be others. Further, the relative weighting of the components is of some importance, since with substantial over-capacity, the relative aircraft location component would be less important. The ideal weighting of components may well vary considerably with particular system attributes.

CHAPTER 7

SUMMARY AND CONCLUSIONS

7.1 Summary

It has not been the intention here to develop optimum rules for dispatching aircraft. Such attempts usually result in unwieldy or else very limited and specialized criteria. The object, however, has been to approach the problem of aircraft scheduling in a stochastic system through the use of simulation, and to explore typical dispatching strategies.

A number of different rules were formulated using a combination of simple mathematical criteria and heuristics. The overall effectiveness of these rules were then tested from both the operators and the customer viewpoint. The vehicle by which the criteria were evaluated was the simulation model. Three models, representing three characteristic networks were built:

1. The two station shuttle

2. The nine station fully interacting network
3. The air taxi service.

7.2 Conclusions

The models of this study proved successful in achieving the purpose for which they were intended. It seems reasonably certain that most existing situations of practical interest may be adequately modelled and operated on a real time basis to simulate an existing system. In general it may be concluded that simulation is a valuable tool in the establishment and evaluation of dispatching criteria. Policy guarantees may be given quantitative measures and critical variables isolated. It may also be used to advantage to create a fixed timetable in an uncertain environment, and to establish 'optimal' capacity and fleet size requirements for given operating conditions.

It is both difficult and dangerous to make generalizations. Each scheduling environment demands different considerations. What may be obvious in one case may well be masked by other more dominant factors active in a different case. However, certain specific conclusions may be drawn from the experience of the models.

1. The marginal economic criteria of the two and nine station problems is a reasonably successful approach to balancing company and passenger interests. The importance of the level of passenger waiting times varies with the level of demand. Specifically, in the low demand case when the probability of no passenger arrivals in the time horizon considered is very high, then the rule is satisfied by default. This suggests the desirability of coupling this rule with a minimum passenger requirement.

2. The time horizon over which the marginal rule is examined is important only in low demand situations with relatively low weighting on passenger waiting times. In other

cases it tends to become less important for the normal ranges considered.

3. The effect of maximum waiting time guarantees is as expected. In general as the maximum times are raised, the company profits through improved load factors, but service deteriorates as evidenced by longer average queues. The value here is in the quantification of such action.

4. In a large system with few aircraft, many decisions to depart are barred due to unavailability of aircraft. Therefore, introducing additional aircraft often improves the level of service to the passenger but results in poorer economic performance because aircraft depart earlier with smaller loads.

5. In capacity planning for the air taxi situation, decreasing interarrival times increased aircraft requirements in a non-linear fashion.

6. In all the models, the number of ferry flights necessitated depends upon the relative build-up of traffic at various points and the initial location assignment of the aircraft. To balance out the effect of randomness, much more and longer runs are required.

7. Demand variations throughout the day necessitate reserve capacity consistent with the degree of fluctuation. For the air taxi problem considered, this extra capacity requirement was on the order of 30%.

8. The 'adaptive' threshold approach is a potentially useful dispatch criterion in situations where passenger waiting time is a relatively important factor.

TABLE 7.1

COMPUTATION TIMES

Problem	Program	Run Duration	Time Incre- Ment	Max # Trans- actions	(sec.) CTSS Time	
					CPU	SWAP
Two Station	Dynamo	24 HR. 240 HR.	1 HR. 1 HR.		5.5 14	2 9
Two Station	GPSS II	1000 min.	1 min.	400	180	18
Nine Station	GPSS II	450 min.	15 min.	1000	41	12
Air Taxi	GPSS II	1320 min.	1 min.	300	31	16.5
Adaptive Air Taxi	GPSS II	4800 min.	5 min.	1050	125	20

7.3 Computation Time

Practically all runs were made on the M.I.T. Time Sharing System. This proved highly efficient for program development and debugging. However, the on line version of GPSS II requires all of the 7094 core memory, exclusive of the area reserved by the supervisor. Therefore, a severe penalty is paid each time the program is swapped in and out of core, and consequently is assigned a very low priority by the scheduling algorithm. Simulations, therefore, consume substantially more than would the same program on batch processing. Average computation times are shown in Table 7.1.

7.4 Suggestions for Further Study

Due to the random features of the models, a great many runs will be required prior to the establishment of confident quantitative measures and trends. Further extension and sophistication in the simulation models is a logical direction in which to move. To this end, it will be necessary to rewrite the existing programs in a more powerful and flexible language. SIMSCRIPT would be an improvement, but for much larger and complex problems, perhaps a general language such as MAD or FORTRAN would provide greater flexibility. The tradeoff here however is in programming and debugging time. List processing languages would be particularly useful in any attempt to extend the adaptive decision model.

Other features which must be incorporated into future models are

1. a multi-stop flight capability
2. a greater coupling of decisions between affected stations
3. a 'two class' aircraft
4. different aircraft types - capacities, speeds, etc.
5. maintenance and crew requirements
6. airport handling constraints

Expansion of the model to increase the number of stations and aircraft in the system and to accommodate a greater number of passengers is limited only by language and computer storage limitations. For most practical problems this should provide no obstacle.

APPENDIX I

BIBLIOGRAPHY

1. A Systems Analysis of Short Haul Air Transportation, Part III, M.I.T., August 1965.

2. J. W. Devanney III, Passenger Transportation Scheduling, S. M. Thesis, M.I.T. 1962.

3. O'Doherty, J. D., A Location Study for a System of V/STOL Airports in the Boston Metropolitan Area, M.I.T. Flight Transportation Laboratory, 1966.

4. An Integrated, Evolutionary Transportation System for Urban Areas, Project Metran, M.I.T. Report No. 8, Spring 1966.

5. Gunn, Wm. A., Airline System Simulation, IAS Paper No. 62-109, June, 1962.

6. Beckmann, M.J. and Bobroski, F., Airline Demand: An Analysis of Some Frequency Distributions, Yale University Paper, 1957.

7. General Purpose System Simulator II Manual, IBM B20-6346-1, 1963.

8. Pugh, A. L., III, Dynamo Users Manual, May 1963.

9. Wooten, C. M., Management Information System for a 1980 VTOL Transportation Network, M.I.T. S.M. Thesis, 1965.

10. Smith, N. P., Jr., Real Time Applications of Computers to Airline Schedule Control 15.95 Project Report Sloan School of Management, M.I.T., May 1966.

11. Simpson, R. W., Computerized Schedule Construction for an Airline Transportation System, MIT Flight Transportation Laboratory Report 8-T-66-3.