

TL552

.R94

.F314

no. 74-2

FLIGHT TRANSPORTATION LABORATORY
REPORT R94-2

A DYNAMIC APPROACH FOR AIR TRAFFIC FLOW
MANAGEMENT OF ARRIVING AIRCRAFT AT A
CONGESTED AIRPORT

FABIEN FEDIDA

1994

MIT

DEPARTMENT
OF
AERONAUTICS
&
ASTRONAUTICS

FLIGHT TRANSPORTATION
LABORATORY
Cambridge, MA 02139-4307

FLIGHT TRANSPORTATION LABORATORY
REPORT R94-2

A DYNAMIC APPROACH FOR AIR TRAFFIC FLOW
MANAGEMENT OF ARRIVING AIRCRAFT AT A
CONGESTED AIRPORT

FABIEN FEDIDA

1994

**A Dynamic Approach for Air Traffic Flow Management of Arriving
Aircraft
at a Congested Airport**

by

Fabien Fedida

Submitted to the Department of Aeronautics and Astronautics in Partial Fulfillment of
the Requirements for the Degree of
Masters of Science in Aeronautics and Astronautics

Abstract

Both the airline industry and air travelers have been pummeled by increased delays experienced at major airports and, as a result, rising operating costs. In this thesis, we focus on the dynamic Arrival Flow Management sub-process of the more general Congestion Management process at a given airport. We show the inefficiencies of a current approach, Miles-In-Trail, and present and evaluate a new approach which we have called Integrated Interactive Dynamic Flow Control (IIDFC). IIDFC produces a set of Traffic Flow Management Advisories which are dynamically updated. It integrates all types of Traffic Flow Advisories and is interactive in the sense that the set of advisories generated and actually issued can be modified by Traffic Flow Managers. Given the complexity of the overall flow management problem, a Traffic Flow Management Simulator was implemented as part of this thesis in order to evaluate various dynamic flow control strategies.

Thesis Supervisor: Robert W. Simpson

Title: Professor of Aeronautics and Astronautics, Director of the Flight Transportation
Laboratory.

Contents

1	Introduction	8
1.1	Description of the Problem	8
1.2	A Current Approach: Early Descent/Miles-In-Trail (MIT)	9
1.3	Approach and Organization	11
1.3.1	Goals, Guidelines and Approach	11
1.3.2	The Simulator	12
1.3.3	Organization of the Thesis	12
2	Integrated Interactive Dynamic Flow Control	13
2.1	The Concept	13
2.1.1	Flow Control	13
2.1.2	Integrated Flow Control	15
2.1.3	Dynamic Flow Control	22
2.1.4	Interactive Flow Control	23
2.1.5	Summary	24
2.2	Heuristic Approach	25
2.2.1	Definitions	26
2.2.2	Establishing the Exit-Holding-Fix Time Sequence	28
2.2.3	Ensuring Time Separation Requirements	29
2.2.4	Heuristic Algorithm Defaults	32

2.3	Optimal Approach	33
2.3.1	Model Overview	33
2.3.2	Selecting the Best Flight Profile	39
2.3.3	Arc Costs	42
2.4	Extended Optimal Approach	47
3	The Traffic Management Simulator	51
3.1	Overview of the Simulator	51
3.2	Modeling and Design Considerations	53
3.2.1	Aircraft Generation	53
3.2.2	Air Holding	55
3.2.3	The “Freeze Time”	56
3.3	Simulator Algorithm	57
3.3.1	A Fast-Time Simulator	57
3.3.2	Initialization	59
3.3.3	New Enter	61
3.3.4	New Take Off	63
3.3.5	System Update	63
3.3.6	New Arrival	65
3.3.7	Exit from the Holding Fix	66
3.3.8	End Of Ban	68
4	Case Studies	70
4.1	Approach and Input File	70
4.1.1	Approach	70
4.1.2	Input File Analysis	71

4.2	Scenarios	74
4.2.1	Scenario 1: Tactical Air Holding -No IIDFC	74
4.2.2	Scenarios 2,3,4,5: IIDFC	75
5	Conclusion	80
5.1	Main Results	80
5.2	Future Directions	81
5.2	Personal Experience	82
Annex		83
A.1	Scenarios Notations	83
A.2	Scenario 1 Data and Plots	86
A.3	Scenario 2 Data and Plots	89
A.4	Scenario 3 Data and Plots	93
A.5	Scenario 4 Data and Plots	96
A.5	Scenario 5 Data and Plots	101
Bibliography		105

Chapter 1

Introduction

1.1 Description of the Problem

Both the airline industry and air travelers have been pummeled by increased delays experienced at major airports and, as a result, rising operating costs. Capacity improvements of the Air Traffic Control system, such as building new airports or expanding existing ones, have now been outpaced by an ever increasing Air Traffic demand, boosted at major airports by the airlines' use of "Hub and Spoke" networks following deregulation. As a result, the 30-40 nm. radius zones surrounding those airports (the Terminal Areas) have been running into saturation problems. At four major airports, Chicago O'Hare, Washington National, LaGuardia and JFK, hourly capacity control has been introduced to smooth the peaks in arrival traffic by scheduling landings months in advance in such a way that each aircraft is assigned to a landing slot. At other airports, the scheduled arrival rate often exceeds the best capacity rate. Moreover, the landing capacity rate of a given airport varies drastically due to changing weather and operational constraints. In that context, Traffic Flow

Managers *tactically* specify an Airport Acceptance Rate (AAR) at each major airport; the AAR is the common name used to designate the Terminal Area Acceptance Rate, that is to say the rate of aircraft entering the Terminal Area to proceed for landing.

In response to tactical flow management, Air Traffic Controllers have to apply tactical control over aircraft not only to ensure safe operations and ease en-route sector congestion, but also to execute dynamic Arrival Flow Management.

In this thesis, we focus on the dynamic Arrival Flow Management sub-process of the more general Congestion Management process at a given airport. This sub-process tries to ensure that the arrival rate at the Terminal Area is equal to or less than the Airport Acceptance Rate under time-varying conditions.

1.2 A Current Approach: Early Descent/ Miles-In-Trail (MIT)

The simplest way to achieve dynamic Arrival Flow Management for a given airport is to consider a simple queuing process where dynamic control of delays is achieved through air holding in a vertical stack upon arrival of aircraft at the Entry Fixes (located on the boundary of the Terminal Area). Thus, we would obtain a time-varying vertical queue at each of the Entry Fixes. Of course, such a procedure is unfeasible at many busy airports since the size of the queue is restricted, and it reserves large volumes of valuable airspace.

However, a subtler form of queuing, called Miles-In-Trail (MIT), can be considered and is currently implemented at some major congested airports such as Chicago O'Hare. Indeed, MIT resembles horizontal queuing: depending on the expected time variation of AAR, Air Traffic Flow Managers issue Miles in Trail constraints (i.e. distance

separation requirements between subsequent aircraft) along all arrival paths to the airport, and into sectors in different Centers.

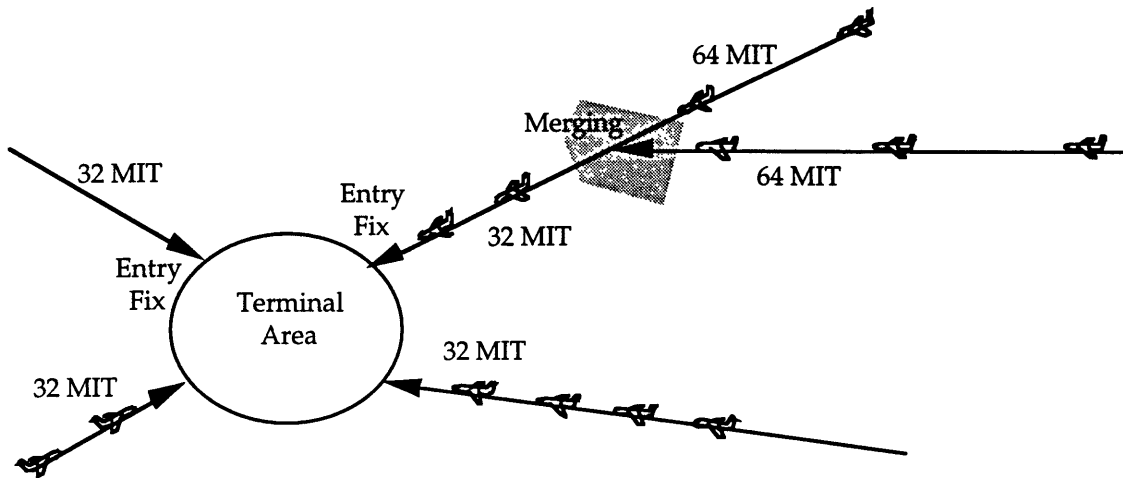


Figure 1-1: Miles in Trail (MIT) for a Terminal Area for Entry Fix Rates equal to 15 aircraft/hour.

For instance, Figure 1-1 illustrates the case where an arrival flow rate of 15 per hour is desired at any Entry Fix. One way to ensure that is to have a time separation between subsequent aircraft of $60/15 = 4$ minutes. Now, assuming that aircraft are flying at ground speed 8 nm/min., then this time separation requirement is translated into a MIT of $(4 \cdot 8) = 32$ nm on the final legs leading to the entry fixes. If two busy arrival airways are merging into a final leg, they are assigned a 64-nm MIT (corresponding to a flow rate of $15/2$ aircraft per hour).

Thus, we obtain a queuing-like system -not a true queuing system since the controller may speed up an arrival at the merge by asking for a higher speed, or by cutting a corner. Also, a controller might allow a pass if there is a gap ahead of a slower aircraft *and* if MIT can be achieved before handoff to the next center. In effect, controllers are expected to handoff with at least the MIT spacing.

This fixed assignment of MIT is inefficient since faster aircraft may join the end of the horizontal queue hundreds of miles away from the airport under consideration and then be forced to slow down and fly the remaining distance to go at reduced altitudes and at the lower speed of preceding aircraft types. By restricting passing in such a drastic way, MIT does not efficiently take advantage of the fact that today's jet transport aircraft have a range of feasible cruising speeds (from $M = 0.70$ or 460 knots, to $M = 0.84$ or 550 knots). One of the goals of our study is to eliminate these inefficiencies.

1.3 Approach and Organization

1.3.1 Goals, Guidelines and Approach

This research presents and evaluates, both theoretically and experimentally, a dynamic approach for Air Traffic Flow Management of arriving aircraft at a congested airport, which we have called Integrated Interactive Dynamic Flow Control (IIDFC). It eliminates the inefficiencies of Miles-In-Trail by assigning an entry slot to all arriving aircraft and issuing Traffic Flow Advisories for Ground Hold and Airspeed such that aircraft can arrive at Entry Fixes at desired spacing equivalent to MIT.

More specifically, our goal is to develop a "Dynamic" Flow Control Procedure in order to produce a set of Traffic Flow Management controls which would be dynamically updated every prespecified period of time (e.g. 20 or 30 minutes) to account for the actual evolution of variables affecting the system (such as weather or operational deviations...). IIDFC should be "Interactive" so that this set of Traffic Flow Advisories can be accepted or modified by Traffic Flow Managers; and IIDFC should be "Integrated" in the sense that it will include all forms of Traffic Flow Advisories: it

should not only perform Arrival Metering by using Cruise Speed Changes, Stack Air Holding, or Early Descent into MIT, but also Arrival Demand Management by using Ground Holding at the origin airports.

Furthermore, IIDFC should be able to trade off the costs incurred to the aircraft operators and the workload incurred to ATC in performing Dynamic Flow Control.

1.3.2 The Simulator

Given the complexity of the overall flow management problem, we have built a simulator as part of this thesis in order to develop, test and evaluate various dynamic flow control strategies. In designing our simulator, we emphasized its maintainability and flexibility to model the ATM system with various levels of details and perform a wide variety of air traffic flow management experiments and strategies. The simulator can account for real world deviations from the Traffic Flow Plan due to uncertain winds, changes in actual landing rate, etc.

1.3.3 Organization of the Thesis

In Chapter 2, we first explain the concept of Integrated Interactive Dynamic Flow Control and then present two Dynamic Resolution Algorithms which can be used at every system update: a Heuristic Approach, which was first implemented, and then the Optimal Approach. It also describes an Extended Optimal Approach which, in particular, allows to take into account specified Entry Fix Acceptance Rates (EFAR) as well as the Airport Acceptance Rate (AAR).

Chapter 3 describes the Traffic Management Simulator in brief detail.

Finally, several case studies are presented in Chapter 4 to evaluate IIDFC.

Chapter 2

Integrated Interactive Dynamic Flow Control

2.1 The Concept

2.1.1 Flow Control

For each major airport, Traffic Flow Managers specify an Airport Acceptance Rate which depends on weather and operational constraints. In fact, the Airport Acceptance Rate (AAR) is the common name used by Traffic Managers to designate the Terminal Area Acceptance Rate, where the Terminal Area is a zone centered on the airport and bordered by the holding fixes. Thus, the AAR is really the rate of aircraft exiting the holding fixes to proceed toward the runways.

Congestion Management for a given airport is superfluous when the traffic flow of aircraft proceeding inbound is very light in the sense that most time gaps between two successive arrivals meet or exceed the time separation requirements imposed by the

FAA and the arrival rate is much less than the Airport Acceptance Rate. However, such a situation is ideal and occurs very seldom. Today, Air Traffic Control (ATC), the major element of the National Airspace System (NAS), will usually have to ensure congestion management of traffic flows by exerting dynamic control over aircraft inbound to major airports.

Within the framework of this thesis we will focus on the congestion management process and more specifically on its metering sub-process.

The purpose of flow control is to smooth the peaks in the arrival rate at the airport so that it matches the Airport Acceptance Rate. A very simple way of approaching the problem would be to consider our system as a mere queuing process. This process would not issue any type of control to aircraft proceeding inbound before their arrival at the holding fix. Aircraft would then enter a vertical holding stack (equivalent to the queue in a queuing process) and be held until they receive clearance for landing. Thus, for purpose of congestion management, the ATC metering system would behave like a queuing server which would provide the required time separations of aircraft at the airport and a rate of aircraft exiting the holding stack air equal to the AAR.

However, such a process neglects the stack holding capacity and, as discussed in Section 2.1.2, is cost inefficient for the aircraft's operators.

In that context, a metering process would smooth the peaks in the arrival rate at the holding fix, but not only at the airport: controls over the arrival traffic streams would be exerted well before the holding fix. For instance, the ATC controller might decide to accelerate an aircraft to place it between two aircraft separated by a large gap and thus avoid delaying subsequent aircraft. However, one has to take into account the "costs" associated with such an action. Indeed, accelerating an aircraft would require the air traffic controller to issue a new speed and path to that particular aircraft and then to monitor it until it fills the assigned gap. Such a move would thus increase the air traffic controller workload.

We can now see that flow control must take into account several factors. In particular it should perform congestion management in such a way as to reduce aircraft operating costs, delays and holding delays at any particular Entry Fix, but it should also give the possibility to control the ATC workload associated with it. In that context, we introduce the concept of Integrated Interactive Dynamic Flow Control (IIDFC).

2.1.2 Integrated Flow Control

In this section, we discuss the “integration” concept of IIDFC.

In order to control the arrival flow at a single airport, Integrated Interactive Dynamic Flow Control integrates all types of Traffic Flow Advisories:

- Cruise Speed;
- Air Stack Holding;
- Early Descent/Miles-In-Trail;
- Ground Holding.

2.1.2.1 Cruise Speed Advisory

Each type of aircraft has a range of cruise speeds in terms of Indicated AirSpeed (IAS) within which it wants to operate.




Figure 2-1: Fuel Burn Rate versus Indicated Airspeed (given weight, altitude, temperature).

If we plot the fuel burn rate versus the IAS as illustrated on Figure 2-1, we indeed obtain a U-shaped curve and the minimum indicated airspeed is then declared at the point where the fuel burn rate is minimum. This minimum IAS is sometimes known as the “holding” airspeed and will vary somewhat with aircraft weight. The maximum IAS depends on the maximum thrust available, or other factors.

Now, since pressure (or air density) decreases with altitude, there are different True AirSpeeds (TAS, which really measures how fast a plane is moving) at each altitude level which correspond to the same IAS. Therefore, a given aircraft type will have different ranges of True AirSpeeds within which it wants to operate depending on its altitude and weight. We assume this range of speeds is known to Traffic Flow Managers.

Let us now see how we can define the optimum cruise speed within that range. This speed should minimize the “Cruise Cost” (CC) for a given flown distance. It amounts to minimize the Cruise Cost per Mile (CCM, in \$/nm) defined as the ratio of Cruise Cost Rate (CCR, in \$/min) and Ground Speed (GS, in nm/min):

$$CCM = \frac{CCR(\$ / \text{min})}{GS (\text{nm} / \text{min})} \quad (1)$$

The Cruise Cost Rate, or Cruise Cost per minute of flight, is a function of the fuel cost and the marginal operating costs of the aircraft operator in terms of time operating cost:

$$CCR = c_1 \cdot \text{FBR} + c_2$$

where FBR is the Fuel Burn Rate (gallon/min)

c_1 is the cost of fuel (\$/gallon)

c_2 time operating cost per minute of flight

Using the variations of the Fuel Burn Rate versus IAS and the correspondence between Indicated AirSpeeds and True AirSpeeds, we can then plot the Cruise Cost Rate versus the TAS for a given aircraft type and altitude.




Figure 2-2: Cruise Cost Rate versus True AirSpeed (for a given aircraft, altitude)

As shown on Figure 2-2, we correct the TAS by the wind speed to obtain Ground Speeds. Ratio (1) is then the slope of a straight line going from the origin of Ground Speeds and intersecting the U-shape curve (for feasibility) within the aforementioned range of desirable true airspeeds. Thus, the cruise cost for a given distance will be minimized when that slope is minimum; that is to say, when this straight line is tangent to the U-shaped curve (see Figure 2-2). The corresponding point of contact gives us the current value for Economic Cruise Speed. It now varies with weight, altitude, and windspeed.

Note that typically the cruise speeds of current jet transport aircraft can be controlled by over $\pm 15\%$ which provides similar control over the time-to-go.

In Cruise Speed Traffic Advisory, the pilot is asked to slow down or accelerate within the range of acceptable cruise speeds. He may receive this advisory by radio directly from the air traffic controller or we can imagine that in the future there would be a data link between the aircraft and the ground to transmit the speed advisories. The data link would strongly alleviate controller workload while using Cruise Speed Advisory under IIDFC.

Of course, the air traffic controller should not issue small changes of cruise speed, say less than 0.02 Mach. Furthermore, we can impose that the changes of cruise speed for a given aircraft be chronologically monotonic so that it is not accelerated and then decelerated (or vice versa) as Advisories are revised over time.

2.1.2.2 Air Holding Advisory: Vertical & Horizontal

There are two forms of air holding: classical “vertical air holding”, or Stack Air Holding, and “horizontal air holding” achieved by starting an Early Descent (ED) into the airport.

- *Stack Air Holding:*




Figure 2-3: The Operation of an Air Holding Stack

If aircraft have not received clearance for landing by the time they arrive at their holding fix, they enter the Stack Air Holding Pattern. Stack Air Holding occurs when

peaks in the arrival flows cannot be smoothed enough or when the Airport Acceptance Rate (or landing rate) is unexpectedly reduced due to bad weather or operational incidents.

As discussed in section 2.1.2.1, there is an Indicated AirSpeed (IAS) which minimizes the fuel burn rate; it will then be used as the IAS for holding. However, for the same IAS, true holding airspeed increases with altitude. As a result, the geometric pattern flown will increase in size as stack holding increases (see Figure 2-3).

Note that even though this operation is called stack holding, it is not always implemented as a First Come First Serve system. Indeed, ATC might decide to assign a lower level to an aircraft which is already late on schedule when it arrives at the holding fix (it is always possible to pass any late arriving aircraft directly through a stack at its assigned altitude).

The paradox about stack air holding is that it is both necessary and undesirable.

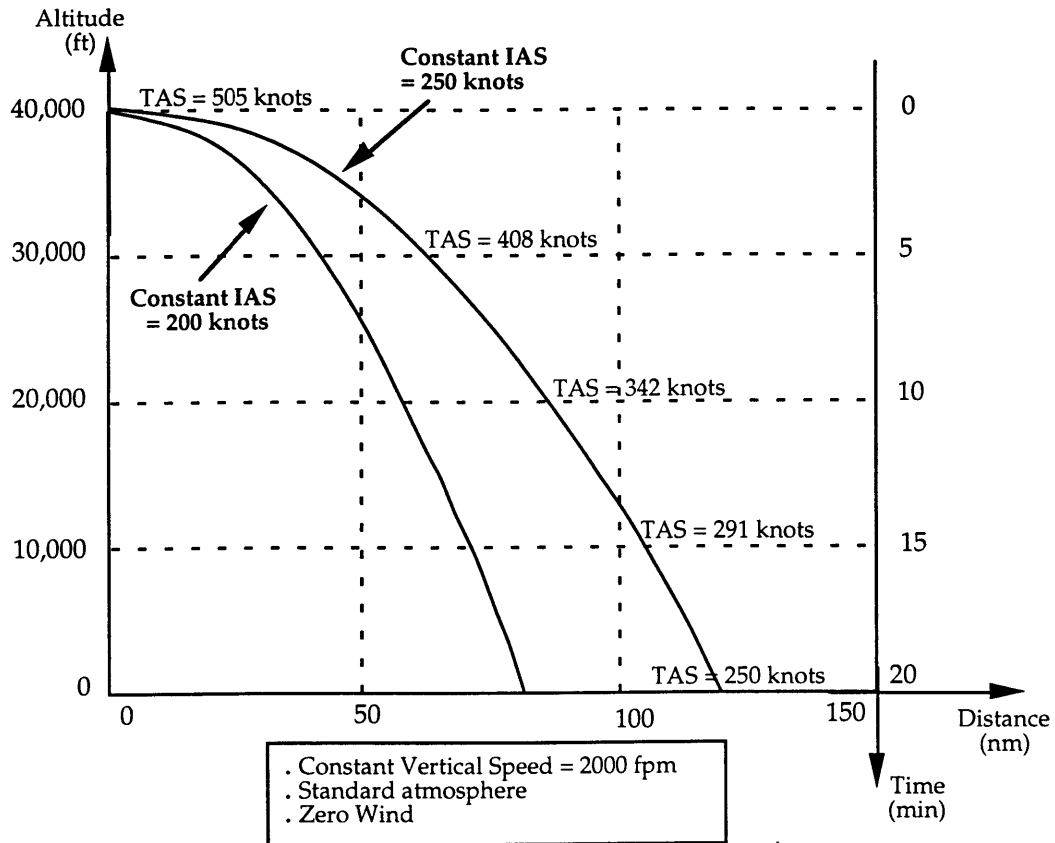
It is necessary because it allows an efficient use of the airport capacity by reducing unnecessary gaps between aircraft. In effect, in a dynamically changing environment, air holding aircraft provides a buffer against operational deviations and unexpected decreases in AAR. Furthermore, their exit time of the holding pattern and then their flying time to the runway can be controlled accurately as compared to a more uncertain arrival time at the holding fix.

However, stack air holding is also undesirable. In effect, from the point of view of ATC, it blocks valuable airspace around the Terminal Area. Furthermore, it is always more fuel efficient for an aircraft operator to reduce Cruise Speed towards the holding airspeed when the assigned arrival delay becomes known, than to fly at high Cruise Speed towards the airport and then be air held at the entry fix.

- *Horizontal Holding:*

Under Integrated Interactive Dynamic Flow Control, the arrival time at the holding fix can also be controlled by choosing an early top-of-descent point (i.e. doing an Early

Descent, ED). This occurs today around major airports where aircraft are spaced at certain “miles-in-trail” along arrival paths. To match the slower ground speeds of low-altitude aircraft close to the airport, aircraft are brought down to lower altitudes as they join the end of a long horizontal “traveling” queue.



Source: TATCA Working Paper Number 3, July 1989, R. W. SIMPSON, Flight Transportation Laboratory, MIT

Figure 2-4: Typical Descent Profile for Jet Transports

Thus, instead of an obvious vertical queue at the Entry Fix, there is a long horizontal queue approaching the Entry Fix at the holding speed. It is an obvious queue when visualized on the Aircraft Situation Display used by today’s Traffic Flow Managers. In effect, for the same Indicated AirSpeed (IAS), the True AirSpeed (TAS) will be much higher at higher altitudes (see Figure 2-4) because of the gradient of pressure with altitude. Let us imagine a situation where there is no wind then, in order to maintain a

constant IAS, which is sensitive to the pressure encountered when moving, the aircraft will have to fly at a faster speed (in fact true airspeed) as shown in Figure 2-4.

For instance, let us consider the situation where controllers bring down jet aircraft which are flying at IAS = 250 knots from 40,000 feet to 10,000 feet. At 40,000 feet of altitude the TAS is 505 knots whereas it is 291 knots at 10,000 feet. Now suppose that the controller decides to delay the arrival of an aircraft by bringing it down 10 nm earlier. In other words, this 10 nm segment will be flown at a TAS of 291 knots (at 10,000 feet) instead of a TAS of 505 knots (at 40,000 feet). The plane will be then delayed by:

$$\Delta t = \frac{d}{TAS_{10,000}} - \frac{d}{TAS_{40,000}} = \frac{10}{291} - \frac{10}{505}$$
$$\Delta t = 1.456 \times 10^{-2} = 0.87 \text{ minutes}$$

Thus the ratio of time delay per distance of an Early Descent is 0.87 minutes/10 nm. Note that in real situations, the corresponding TAS and groundspeeds will depend on wind variations over space, time and altitude.

2.1.2.4 Ground Holding Advisory

Holding an aircraft on the ground allows the Traffic Flow Manager to control the departure time, and thus the arrival time. The main advantage of such Traffic Advisories is that it allows aircraft operators to save fuel and reduces the number of aircraft in enroute and airborne holding.

However, one must be aware that it can only be used as a gross control over the arrival time. For instance, the taxi times to the end of the runway vary by as much as 15 minutes from one day to the next depending of which runway is used for departure. Furthermore, it is all the more difficult to pre-specify take off times at busy airports since departing aircraft are often held at the runway before take off. Therefore, ground

holding should be used only if a significant delay for that aircraft is forecasted or only if this delay cannot be “absorbed” by speed control.

Two main constraints in airline operations should also be taken into consideration when issuing dynamic Ground Hold Advisories:

- Within some time period of an airline departure (e.g. 30 minutes), it is crucial that ATC remains committed to that particular departure time: i.e. the aircraft departure process is considered to have started and cannot be revoked.
- Departure Times issued by IIDFC must always be greater than the original scheduled departure time i.e. airline aircraft cannot depart earlier than scheduled.

Note that dynamic control of the departure time may be particularly effective when the Airport Acceptance Rate suddenly increases after being low for some period of time. IIDFC would then ask a significant proportion of aircraft which were ground held to depart as soon as possible. We see here that operating Ground Holding assumes that departure time updates can be communicated promptly to airlines and the towers of originating airports.

2.1.3 Dynamic Flow Control

In this section, we discuss the “dynamic” concept of Integrated Interactive Dynamic Flow Control.

The main idea expressed by “dynamic” is that a flow control plan will be regularly updated every pre-specified time period called *Tupdate* (e.g. 30 minutes, or once per hour). In particular, it means that controllers can *only* issue Traffic Flow Advisories every *Tupdate* minutes to airborne and on-the-ground aircraft which have filed a flight plan. Thus, ATC will allow an airborne aircraft which newly “enters” the flow control system to fly at its current cruise speed until the next update. Likewise, an aircraft at

an originating airport making a request for a certain departure time will not learn of its planned departure time until the next system update. This helps to control the Traffic Management workload, and ensures that up-to-date information is being used throughout the day.

In a dynamically evolving environment, such a strategy allows flow control managers to account for all dynamic parameters which affect the system such as:

- actual position, cruise speeds and predicted delays of airborne aircraft;
- ground holding, predicted delays of aircraft on the ground;
- actual arrival delays and actual air holding at the airport;
- cancellations, delays, new flight plan filings or new aircraft (airborne or on the ground) which have entered the system since the last update;
- revised weather forecasts along the route and at the airport as well as operational incidents or constraints (e.g. snow-plowing) which translate into revised forecasts of Airport Arrival Rate over the next few hours.

Note that all the above information is known to Traffic Flow Managers, currently, via the Aircraft Situation Display.

However, traffic advisories will not need to be issued at every update. Indeed, an aircraft might be assigned the same set of controls for several updates in a row, and it is a desirable characteristic of IIDFC that it can be made to minimize the number of changes in Traffic Advisories.

2.1.4 Interactive Flow Control

In this section, we discuss the “Interactive” concept of IIDFC.

Flexibility of such a Traffic Flow Management System should be one of the main concerns in the implementation. The system should behave well whatever the values of

the aforementioned parameters (“inter-update time”, time for commitment to a departure time before actual departure, etc.).

Furthermore, we want this system to be interactively responsive to various constraints placed dynamically by flow managers or even by airline operational control personnel. Such possible constraints are numerous.

For instance, controllers might decide not to use a type of Flow Advisory (e.g. Ground Holding), or limit the number of Traffic Flow Advisories that they have to issue at each update time .

Terminal area congestion sometime drastically impacts the holding area of a particular entry fix. Controllers might then decide to limit or eliminate airborne holding at that entry fix. They may also want to limit the arrival rate at any holding fix independently of AAR.

Furthermore, delay and ATC workload should not be the only parameters taken into consideration before issuing controls to the aircraft. Fuel costs should also play a major role in the decision. These can be specified by the airline operator along with the upper and lower limits on Cruise Mach Number for his aircraft.

The ultimate goal of IIDFC is that, once controllers have defined their constraints, IIDFC will give them the *best assignment of Traffic Advisories* to issue to the fleet currently proceeding inbound or to aircraft still on the ground at originating airports.

2.2.5 Summary

Integrated Interactive Dynamic Flow Control is a new concept for Traffic Flow Management aimed at controlling the arrival flow at a single airport to ease congestion. In order to account for the actual evolution of variables affecting the system -such as weather, actual delays at the airport, new flight plan filings, etc.-, the Traffic Flow Management controls will be dynamically updated every T_{update} (e.g. 15 minutes) and

only every T_{update} . Thus a plane entering the system between two updates will not be issued a congestion management control until the next update. At each update, the resolution part decides where the delay of aircraft in the system will be occurring. It creates a set of "Traffic Flow Advisories" which can be accepted, modified by Traffic Flow Managers and then executed by ATC controllers at diverse locations in the field. Those advisories must be simple and feasible to execute, and they must be communicated in a timely fashion.

Furthermore, IIDFC integrates all types of traffic flow advisories to achieve efficiencies (Air Holding, Cruise Speed, Ground Holding and Early Descent). It can use all, or only a selection of these advisories.

Eventually, IIDFC should be interactively responsive to various constraints placed by flow managers or even by airlines such as limiting the controller workload, or limiting airborne holding at any entry fix.

2.2 Heuristic Approach

As discussed earlier, a plan for optimal assignment of delay will be produced by a Dynamic Resolution logic at regular interval (e.g. every 15 minutes). Traffic Flow Advisories will then be extracted from this plan in order to achieve it.

In this section, we present a heuristic sequencing algorithm to be used as a Dynamic Resolution Logic. This algorithm was first developed before adopting the optimal approach described later.

2.2.1 Definitions

All quantities described hereafter are presented for one aircraft at the system update time: i.e. the time at which the Dynamic Resolution is executed. For didactic purposes, Figure 2-5 illustrates those definitions for an airborne aircraft in a fictitious situation.




Figure 2-5: Arrival Times and Exit Time from the holding fix for an airborne aircraft currently scheduled to fly at cruise speed V and be air-held.

- The current value of the *Scheduled Arrival Time (SAT)* at the holding fix is based on the controls that the aircraft has been previously issued, the current wind forecast and the current aircraft location.
- The current value of the *Scheduled Exit Time (SET)* from the holding fix is based on the same parameters but it further includes planned air holding: i.e. it is equal to SAT when there is no planned air holding.
- The *Earliest Exit Time from the holding Fix, EET*, of an airborne aircraft is the time at which it could exit the holding fix if it were to fly its current distance-to-go at maximum cruise speed and not be air held (see Figure 2-5). For an on-the-ground aircraft, it further assumes that the plane would take off at its Earliest possible Take Off Time (ETOT). Note that the ETOT can be different from the current (i.e. last issued) Take Off Time (TOT).

- We then define $TAT_{left} = (SAT - EET)$ as the Total Amount of Time we can “push” earlier the aircraft’s SAT on the timeline (i.e. to the *left* if the timeline is oriented from left to right) by using *all* Traffic Flow Advisories: Ground Holding, Cruise Speed and Air Holding.
- AT_{left} is the Amount of Time we can “push” earlier the aircraft’s SAT on a timeline by asking it to fly the distance-to-go at its maximum cruise speed i.e. by using only Cruise Speed Advisory. Thus, an airborne aircraft flying at its maximum cruise speed from its current position to the holding fix, would arrive at a time $t = SAT - AT_{left}$ at the holding fix. Likewise, an on-the-ground aircraft taking off at its current (last issued) Take Off Time (TOT) and then flying at its maximum cruise speed would arrive at $(SAT - AT_{left})$ at the holding fix.

Note that for an airborne aircraft, $TAT_{left} = AT_{left}$.

- $Tright$ is the latest time at which an aircraft would arrive at the holding fix if it were to take off at its current Take Off Time (for an aircraft on the ground) and fly until the holding fix at its minimum cruise speed. Note that $Tright$ is neither the latest time at which the aircraft can arrive at the Entry fix, since the aircraft can also be delayed by ground holding; nor the latest exit time from the holding fix, since it can also be delayed by air holding.
- On the timeline of Exit Times from the holding fix, we call a gap between two planes which does not satisfy the time separation requirement as a *too-tight gap*. Likewise, gaps larger than the separation requirement are called *slack gaps*, and gaps equal to the separation requirement, *tight gaps*. This time separation requirement can be either set to $1/AAR$ to achieve a rate less than AAR, or it can be derived from the in-trail separation requirements imposed by the FAA to prevent collision and wake-vortex turbulence effects.

2.2.2 Establishing the Exit-Holding-Fix Time Sequence

This heuristic dynamic flow control algorithm operates on an exit-holding-fix time sequence of aircraft. It then produces SETs such that the planned Airport Acceptance Rate is not exceeded by ensuring that the time separations between successive SETs are greater than $1/\text{AAR}$. However, it does not change the sequence once it has been established. Setting up the sequence is therefore crucial.

In a first step, we must then sequence the fleet along a timeline using some “sequencing criteria” such as:

- the Original Nominal Exit Time from the holding fix (ONET) which is the original requested exit time from the holding fix. For an airliner, it would be computed by subtracting the flying time between the holding fix and the runway, from the original requested time of landing.
- the current Nominal Exit Time from the Holding fix (NET) which is the time at which the aircraft would leave the holding fix if it were to take off at its current TOT (for an on-the-ground aircraft), fly the distance-to-go at its optimum cruise speed and not be air held;
- the Earliest Exit Time from the holding fix (EET) defined earlier.

For each aircraft, the SET is then temporarily set to its initial sequence time. Furthermore, TAT_{left} and AT_{left} are changed so that they become relative to this initial sequence time; for instance, TAT_{left} becomes the maximum amount of time we can push the aircraft to the left (i.e. earlier) on the timeline *from its initial sequence time* (as opposed to its SAT before). Figure 2-6 illustrates that case when we use ONET as the sequencing criteria.

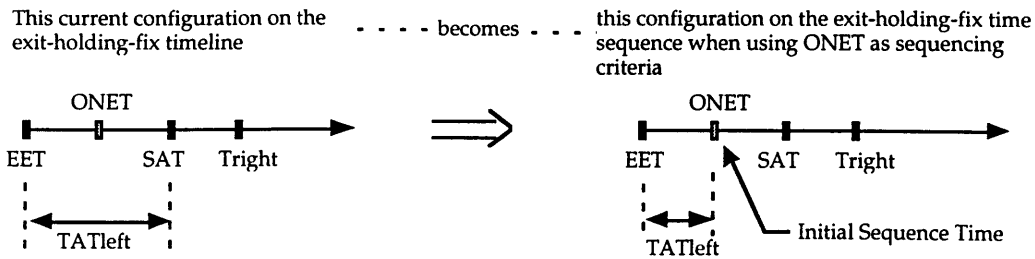


Figure 2-6: Change of TATleft when sequencing criteria used is ONET (for a given aircraft and Dynamic Resolution Time).

Note that, at some point in time in the system, it is possible for an aircraft to have an ONET less than its EET; that is to say that it can only exit the holding fix later than its ONET even if it flies its distance-to-go at maximum speed and is not air held. For instance, this could happen to an aircraft which has been ground held for a long time. When using ONET as a sequencing criteria, we will in that case insert the aircraft in the exit-holding-fix time sequence according to the closest possible exit time to its ONET i.e. its EET; TATleft will be then set to zero.

2.2.3 Ensuring time separation requirements

Once the exit-holding-fix sequence has been established according to some criteria, the heuristic proceeds from left to right on this timeline (i.e. from earlier to later) and ensures iteratively that any gap between two successive aircraft on the timeline satisfies the time separation requirement.

When the algorithm finds a too-tight gap, the main idea to make it a tight gap is to reduce the slack gaps which are on its left on the timeline (i.e. push aircraft earlier) if possible. If it is not sufficient, the trailing aircraft (of the too-tight gap) is then pushed to the right. This technique calls for higher speeds by aircraft to fill any gaps in the arrival flow and reduces delay significantly.

- Reducing the slack gaps on the left of a too-tight gap

Suppose we have the situation illustrated on Figure 2-7, where aircraft are ordered according to some sequencing criteria:

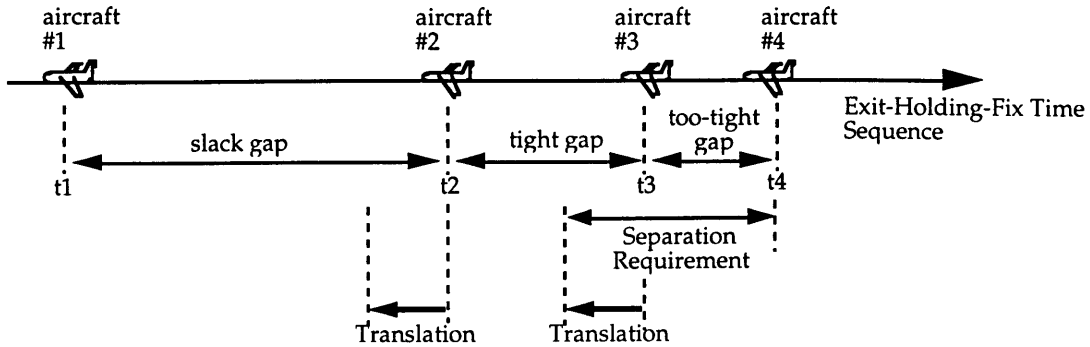


Figure 2-7: Reducing a slack gap

In order to ensure the separation requirement between aircraft #3 and #4, we would reduce the slack gap between planes #1 and #2 by translating earlier (i.e. to the left) both aircraft #2 and #3 by the minimum of:

- . TAT_{left} of aircraft #1, TAT_{left} of aircraft #2; (1)
- . $(\#1_ \#2 \text{ slack gap}) \text{ minus } (\#1_ \#2 \text{ separation requirement})$; (2)
- . $(\#3_ \#4 \text{ time separation requirement}) \text{ minus } (\#3_ \#4 \text{ too tight gap})$; (3)

If the minimum is (3) then reducing this slack gap is sufficient to ensure the necessary time separation between aircraft #3 and #4 as shown in Figure 2-7. If the minimum of those three quantities is not (3), the separation requirement between aircraft #3 and #4 is still not met after “translation”. We then find another slack gap on the left of the too-tight gap and proceed the same way until the separation criterion is satisfied or until there are no more reducible slack gaps.

For every too-tight gap, it is now clear that we have to establish some sort of priority in choosing the slack gaps (on its left) that we will reduce. To do so, each of the slack gaps relative to a given too-tight gap (i.e. each of the slack gaps on the left of that too-tight gap) is attributed a number “S”; S is equal to the sum of the differences between

the current position on the time sequence and the ONET for each aircraft following the slack gap but preceding (and including) the leading plane of that particular too-tight gap.

In the above example, the “S” of the slack gap between aircraft #1 and 2 relative to the aircraft #3-4 too-tight gap would be:


$$S = (t_2 - ONLT_2) + (t_3 - ONLT_3)$$

Since all aircraft following the slack gap and preceding the too-tight gap are moved earlier on the timeline when reducing the slack gap, we will first reduce the slack gaps with the largest “S”. So doing we implicitly recognize that there is a much smaller penalty in accelerating an aircraft than in delaying it. When there is a tie between the “S” of two slack gaps, we first reduce the closest one to the too-tight gap.

Eventually, whenever an aircraft is translated along the timeline, its TATleft and ATleft must be updated accordingly.

- Pushing the trailing aircraft to the right

If the initial too-tight gap still does meet the time separation requirement after reducing its associated slack gaps then the trailing aircraft is delayed accordingly.



On the above example, aircraft #4 is delayed to a position such that the separation requirement between aircraft #3 and 4 is met. As shown in this case, if it has to be delayed more than the position of its trailing plane (plane #5), the latter should then

also be delayed at a position after the new position of aircraft #4 and such that the separation requirement between aircraft #4 and 5 is met. These actions cause the delay in the arrival flows associated with queuing operations.

When there are no more too-tight gaps, the aircraft positions on this timeline are the new values for Scheduled Exit Times. Controllers will then issue the necessary set of Traffic Advisories to the fleet so that aircraft can achieve those SETs.

2.2.4 Heuristic algorithm defaults

- As discussed earlier, the algorithm always looks for possibilities to ensure a time separation requirement by accelerating aircraft rather than delaying them. As a consequence this algorithm should tend to accelerate aircraft. If it is reasonable to consider that there is a smaller penalty associated with accelerating an aircraft than with delaying it, we would like to control more efficiently the relative levels of those penalties.
- The heuristic approach does not handle satisfactorily the case where the Airport Acceptance Rate varies in time due to the weather and operational deviations: i.e. it does not allow to integrate efficiently the AAR forecast in issuing the SETs.
- Furthermore, if the fleet consists of very different types of planes (as far as ranges of cruise speeds are concerned), we can expect a decrease in the heuristic performance in minimizing delay since the algorithm cannot swap aircraft positions on the exit-holding-fix time sequence.
- Generally speaking, the heuristic lacks flexibility. For instance, the number of moves of airplanes along the timeline, which reflects the number of speed change advisories, cannot be controlled efficiently. In the same order, it would be poor at

ensuring that all speed changes must be monotonic. In a few words, the heuristic approach is not interactive.

2.3 Optimal Approach

In this section, we present another Dynamic Resolution logic which would be executed at every system update (i.e. at regular intervals of, say, 15 minutes) in order to produce a plan for assignment of delay. This solution is called optimum solution because it uses an “optimal assignment” technique, and fulfills all the requirements of Integrated Interactive Dynamic Flow Control discussed in section 2.1. In particular, it solves all the aforementioned problems encountered by the heuristic approach.

2.3.1 Model Overview

Modeling our problem by an “optimal assignment” network (see Figure 2-8) suits particularly well all IIDFC goals.

The initial goal was to set up a network whose minimum cost flow solution would minimize the aircraft delay. Later, the costs (and, as a result, the minimum cost flow solution) came to reflect both operator and traffic management costs.

The network can be preconstructed and updated at each iteration or we can choose to create a new network at each Dynamic Resolution call. As discussed later in Section 3.3.5, we chose the second solution in the implementation of our simulator.

Note that such a network model satisfies the requirement of a fast dynamic resolution. Indeed, many fast codes exist to solve the Minimum Cost Flow problem in seconds using today’s workstations or PCs. Moreover, as shown in Figure 2-8, the “tree-like

structure” of our network model is synonymous with low complexity which enables those algorithms to run even faster.

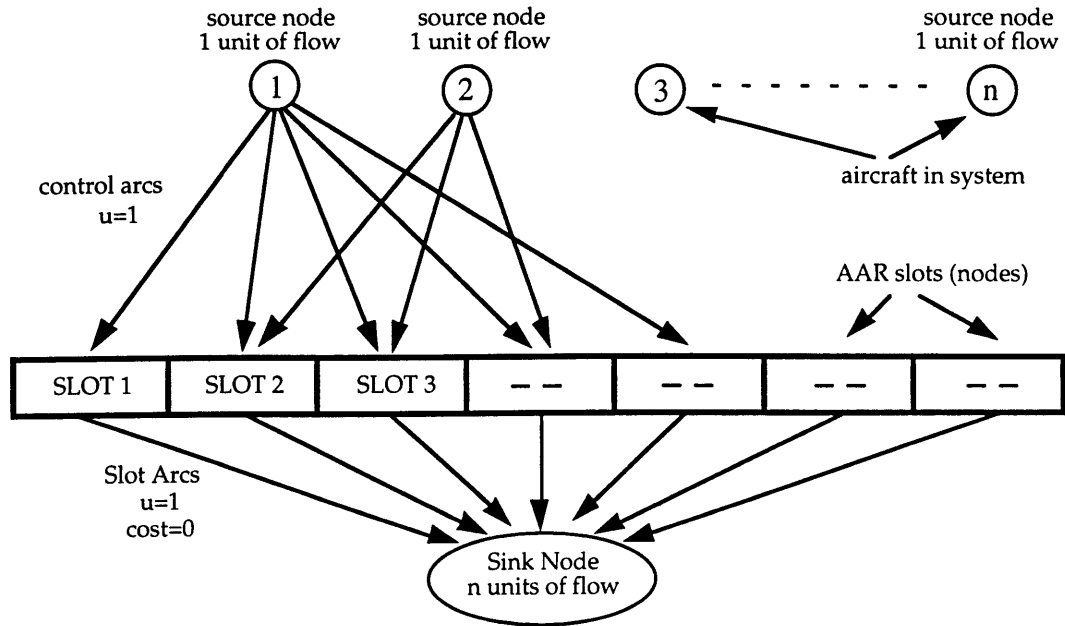


Figure 2-8: Network for Optimal Assignment of Delay

There are three kinds of nodes in the Network for Optimal Assignment of Delay:

- The *aircraft nodes* represent aircraft which are in the system: that is to say, aircraft which already made a request for landing at the airport. Recall that there is a requested exit time from the holding fix corresponding to this requested landing time. Those planes can be either airborne, flying toward the airport under consideration, or still on the ground at their originating airport. A third category consists of aircraft scheduled to land first at an intermediate airport (i.e. different from the one under congestion management control) but are then continuing toward the airport under congestion management control. Those aircraft will be considered as on-the-ground aircraft at the intermediate airport and will therefore be part of the system.

From the point of view of the network representation, each of those nodes is a 1-unit flow source node. Thus the number of flow units circulating in the network is the number of aircraft under consideration.

- The *AAR nodes* or *AAR slots*. Under congestion management control, a certain Airport Acceptance Rate (AAR) is specified depending on various factors such as the weather or operational deviations. Recall that the AAR is in fact the Terminal Area Acceptance rate. It can be easily converted in exit-holding-fix slots (or AAR slots). For instance, an Airport Acceptance Rate of 30 aircraft/hour would mean that there are 30 AAR slots of 2 minutes each in each hour at the entries to the Terminal Area. However, the forecasted AAR is only rarely constant over a long time period (more than a few hours). The size of the AAR slots will then vary along the time line to match the variations of the forecasted AAR. That case is illustrated on Figure 2-9.




Figure 2-9: AAR Slots Size Matching the Airport Acceptance Rate Forecast.

- There is also one *Sink Node* to ensure feasibility of the problem. In effect, the number of flow units emitted by the network sources (the aircraft) must be equal to the number of flow units “absorbed” by the sink.

The goal of our resolution is to assign each aircraft represented in the network to a unique slot and to ensure that the AAR is respected. Thus, since each aircraft is a source of one flow unit and the smallest quantity which can circulate on any link is one unit (it is an integer problem), we are sure that only one of the “control arcs” originating from a given aircraft node will carry a flow; since there is only one arc per feasible aircraft-slot pair, we are then guaranteed that each aircraft will be assigned to some slot. Furthermore the “Slot Arcs”, which go from each slot to the sink node, have a flow upper bound of one which ensures that at most one aircraft is assigned to any slot. This guarantees that the actual flow rate at the Terminal Area will be less or equal to the AAR.

For implementation purposes, note that an arc from an aircraft to a slot is constructed as long as the aircraft can make at least one point in that slot i.e. as long as there exists an acceptable set of controls (Ground Holding, Cruise Speed, Air Holding) to make the aircraft exit the holding fix (or enter the Terminal Area) at that point in time. But, if an aircraft can make the middle of that slot then it is that middle-slot time which will be assigned as a Scheduled Exit Time (SET) from the entry fix. We then implicitly assume that ATC is able to apply dynamic spacing to ensure that the time separation requirement between subsequent aircraft at landing is met.

Thus, a particular control arc represents the decision to assign an aircraft to an AAR or Terminal Area Entry slot. More specifically, it corresponds to a given aircraft delay. Thus, the network would solve for minimum delay if the delays were used as costs on these arcs. For a given slot, the best flight profile (see section 2.3.2) determines the assignment of delay between Cruise Speed, Ground Holding and Air Holding. The

actual cost associated with that arc combines the aircraft operator's trip cost (both in time and fuel), and the ATC workload costs involved in making the slot (see section 2.3.3).

The slots coverage by a given aircraft will be limited on the left by the earliest time at which it can enter the Terminal Area and on the right by the on-board fuel. One might also decide to reduce this range to account for maximum holding or put a very high cost on arcs pointed to slots later than a given exit time from the holding fix (e.g. if the aircraft has an connection to an international flight). This illustrates how flexible our network model is and how it particularly adapts well to constraints placed by Traffic Flow Managers and airline operational control personnel.

Let us mention the practical constraint which is that ATC often considers that aircraft "close" to the holding fix, i.e. scheduled to arrive at the holding fix within some time period called Tfreeze (e.g. 30 minutes), should not be controlled anymore by IIDFC. Such aircraft will still be represented in the network, but each of the associated aircraft nodes will only have one control arc terminating on the slot which contains the aircraft current Scheduled Exit Time (SET) from the holding fix.

Another important feature of our model is that it also allows us to account for current holding delays at the airport which could be the result of operation deviations. In effect, the earliest slot which is represented in the network at each update (that is the earliest exit slot from the Entry Fix which can be assigned to an aircraft) corresponds to the time at which all the currently air-held aircraft will have exited the Entry Fixes to proceed for landing.

If the Airport Acceptance Rate is very low then the size of the AAR slots is large and dividing each slot into "sub-slots", as illustrated on Figure 2-10, might be wise. In effect, when the size of an AAR slot is large, the algorithm should be able to assign the "best" exit time over the sub-slots of that AAR slot to a given aircraft. One way to do so is to have several control arcs originating from the same aircraft-node (each one

corresponding to a “sub-slot” or SET) and “pointing” to a given AAR slot. A best flight profile (in the sense of minimum trip cost) is associated to each one of those arcs. Figure 2-10 shows how the network is modified: it gets more dense (there are more nodes and arcs) but the complexity remains low.

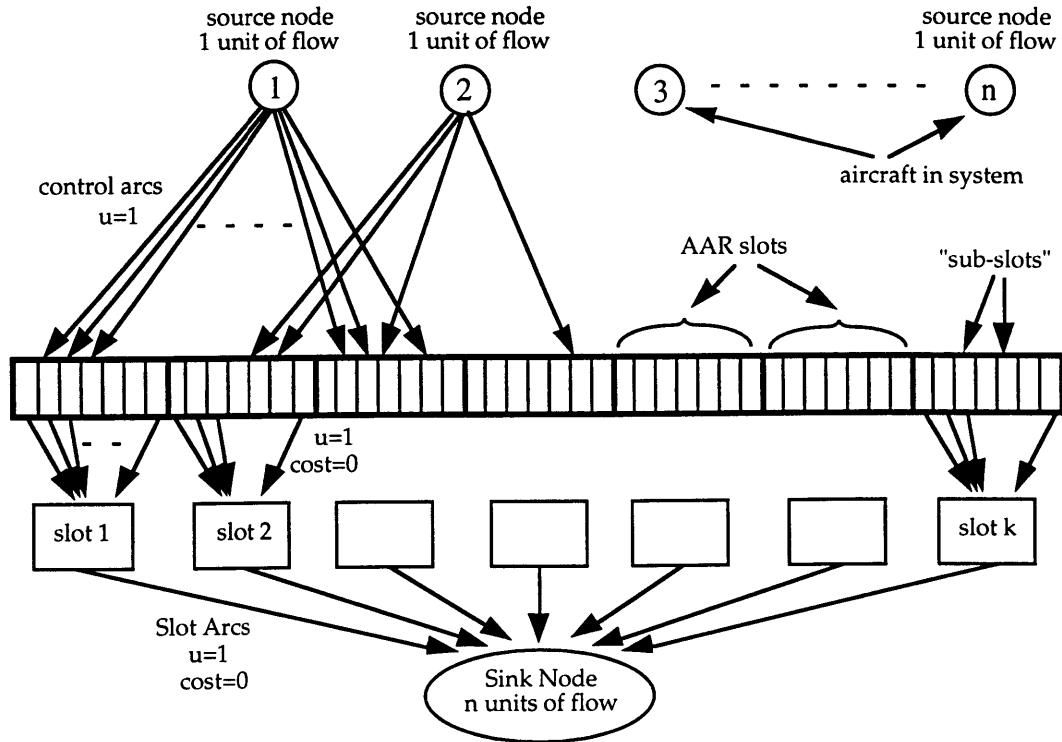


Figure 2-10: Network for Optimal Assignment of Delay with sub-slots

Note that the minimum cost obtained from solving this network is guaranteed to be less than the minimum cost which would have been obtained from solving a network without sub-slots, provided that cost are assigned on control arcs the same way. However, this idea has not been implemented. It is considered to be a “local” metering process best implemented by local controllers, not Traffic Flow Managers.

2.3.2 Selecting the Best Flight Profile

From the aircraft's operator side, there is a best flight profile for each control arc; that is to say, for each aircraft-slot pair. Indeed, one must be aware that there may be various possible sets of controls which could enable a given aircraft to make a given slot. For instance, an aircraft might be able to enter the Terminal Area at the same time if it flies at its current cruise speed and is then air held, or if it only slows its cruise speed. In that case, we will select the best flight profile; that is to say the set of controls which minimizes the aircraft's operator cost.

This user cost depends on the flight operating costs -fuel burnt and time operating cost-, and on the delay from original scheduled landing time. However, for the slot under consideration (i.e. for a given landing time), the delay cost is fixed. The best flight profile associated with a control arc should then minimize the aircraft's operator cost by minimizing its flight operating cost.

2.3.2.1 Case of an Airborne Aircraft

The case of an Airborne Aircraft is the simplest one. We must tradeoff two controls: Cruise Speed and Air Holding. Furthermore, for a given aircraft-slot pair, the flying time is fixed which means that the time operating cost is fixed. Thus, minimizing the user cost amounts to minimize the fuel burnt during a given time period; that is to say that we must minimize the fuel burn rate.

We have seen in section 2.1.2.1 that the fuel burn rate rises when the cruise speed is increased from the minimum acceptable cruise speed, called the holding airspeed, to the maximum cruise speed. Thus, *during a fixed period of time*, an aircraft will always save fuel by slowing down its speed within that range.

Now for a given aircraft-slot pair, *the flying time is fixed*. Therefore if the aircraft can make the slot under consideration by slowing down its cruise speed within the

acceptable range it should do so. This means that we will always use speed control in order to achieve a slot; it is only when such a control is not sufficient that the aircraft should be assigned air holding. That situation arises when the slot is later than the time at which the aircraft would enter at the airport by flying at the minimum cruise speed (i.e. Tright).

Thus, for a given aircraft and using Figure 2-5, the best flight profiles are defined as follows:

- slots between EET and Tright should be made by adjusting the cruise speed;
- slots after Tright should be made by flying at minimum cruise speed (and so doing the aircraft arrives at Tright at the holding fix) and being air held until the Scheduled Exit Time from the holding fix.

2.2.3.2 Case of an Aircraft on the Ground

We have to take another control, ground holding, into account. Our goal is still to minimize fuel burnt and time operating cost but it is not equivalent to minimizing the fuel burn rate anymore since, now, the flying time is not fixed. For a given aircraft-slot pair (i.e. for a given delay), there is no operating cost for being ground held and, as a result, the user cost is minimized by keeping the aircraft on the ground and then let it fly at its optimum speed -which minimizes flight operating cost (i.e. fuel burnt and time operating cost) as shown in Section 2.1.2.1- toward the airport.

Furthermore, we must consider three additional constraints:

- Recall that some time period before departure (e.g. 30 minutes), we are committed to the *current* Take Off Time (TOT). It may be the Original Take Off Time (OTOT, originally requested Take Off Time) or the last TOT issued.
- ATC cannot ask an aircraft to take off before its OTOT.
- We will not bother to issue a new departure time if it is too close to the previously issued departure time. For instance, we do not want to issue a new departure time

of 11:23 am to an aircraft previously scheduled to depart at 11:25. The main reason is that, as seen earlier (Section 2.1.2.4), the ground holding control should only be used as a “gross” control. Thus we will not issue a new TOT if the aircraft can make the slot only by adjusting its cruise speed once airborne.

Therefore, given an aircraft-slot pair, we will have various Earliest Take Off Times (ETOTs) depending on the relative time positions of the OTOT, the current TOT and the time in minutes at which the Dynamic Resolution is implemented (t_{now}):

- (a) If $OTOT \leq t_{now} + 30 \leq TOT$ then the Earliest Take Off Time (ETOT) is $t_{now} + 30$.
- (b) If $t_{now} + 30 \text{ min} \leq OTOT \leq TOT$ then the $ETOT = OTOT$.
- (c) Whatever $OTOT (\leq TOT)$ is, if $TOT \leq t_{now} + 30$ then we are committed to the TOT and we cannot change it. In particular, $ETOT = TOT$.

An aircraft in situation (c) can only cover slots to the right of $SAT - AT_{left}$. We will construct the corresponding arcs but the aircraft will not be issued a new cruise speed at that point.

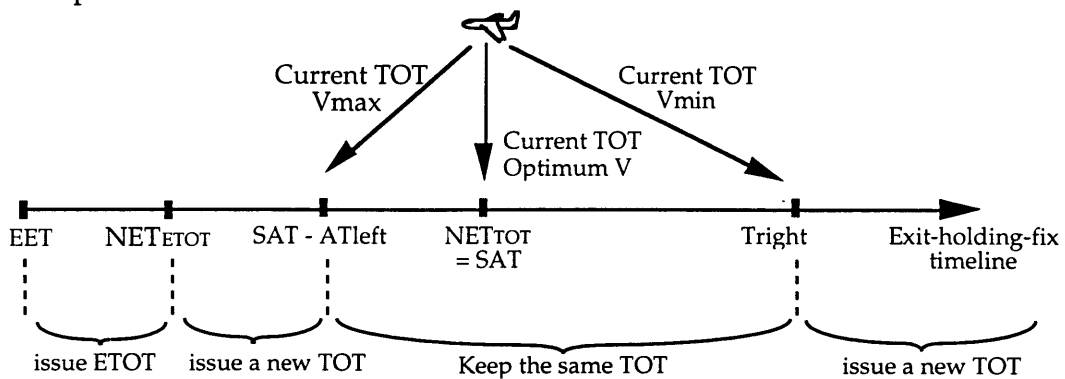


Figure 2-11: Feasible slots and corresponding ground holding policies for an on-the-ground aircraft if ATC can issue a new TOT.

Call NET_t the time at which the aircraft would exit the holding fix if it were to take off at “t”, fly at its optimum speed and not be air held. Figure 2-11 illustrates the case of an aircraft in situation(a) and (b): slots to the right of the Earliest Exit Time (EET)

from the holding fix are feasible slots. For those slots the best flight profile are defined as follows:

- The aircraft can make slots between SAT - ATleft and Tright by taking off at its current TOT and then adjusting its cruise speed. However, we will not issue a new cruise speed at that point. In effect, the next update after take off will happen even before the aircraft has time to reach its cruise altitude and adjust its speed. Thus we implicitly consider that, once airborne, this aircraft will then be treated as any airborne aircraft and any desired speed changes will then be issued.
- Slots from EET to NET_{ETOT} will be made by taking off at ETOT and then increasing the cruise speed. For the same reasons as above, we will not issue the new cruise speed at that time.
- Eventually, slots to the right of Tright, and between NET_{ETOT} and SAT - ATleft, will be made by issuing an increased Ground Hold or later TOT such that the aircraft can then fly at its optimum speed.

2.3.3 Arc Costs

As discussed earlier, each control arc of the network represents a decision to assign an aircraft to a Terminal Area slot. Furthermore, each one of those arcs has its *own* non-zero cost per unit of flow carried. This cost depends on two sets of costs: the User Cost and the Air Traffic Management Cost.

2.3.3.1 User Cost

For a given control arc, the aircraft operator's cost -the User Cost- depends itself on two sets of cost:

- The Delay cost, D

The delay is simply the difference between the Scheduled Exit Time associated with the slot under consideration, SET_{slot} , and the Original Nominal Exit Time (ONET) from the holding fix or original requested exit time of the aircraft. The corresponding delay cost depends on the aircraft operators' on-time performance cost.

Note that we express this cost as a quadratic function of the delay. As far as delay cost is concerned, this ensures that it is more expensive (i.e. less desirable) to increase the delay of an aircraft already late on schedule, than to assign that same incremental delay to an aircraft which is not as late on schedule. Also, when aircraft in a sequence have the same planned delay (e.g. zero when they enter), it guarantees that the algorithm will prefer delaying all aircraft from one slot rather than assigning all the delay to one aircraft in the sequence.

Thus, the delay cost, D, is:

$$D = k_D \cdot (SET_{slot} - ONET)^2 \quad \text{where:}$$

- . $SET_{slot} > ONET$;
- . k_D is the cost per squared minute of delay.

In general, the cost for being early is zero. However, when fuel cost is not taken into consideration, there is no penalty for being accelerated, i.e. for being early on schedule; in that case, the cost for all arcs leading to slot earlier than the original scheduled exit time are equally cheap and the resolution then chooses at random one of them. As result, a tendency to accelerate aircraft will be observed even when it is not needed. To avoid this tendency, a very small cost for being early might be assigned the same way as above if fuel costs are disregarded.

• The aircraft incremental flight Operating Cost, OC:

This cost is the sum of the incremental Fuel Cost, FC, and the incremental Time Operating Cost, TOC:

$$OC = FC + TOC$$

- *Case of an airborne aircraft:*

For an airborne aircraft, we have seen in Section 2.1.2.1 that, at any time, the minimum Operating Cost for the remaining distance to go is obtained by flying at optimum (or nominal) speed toward the Entry Fix and not being air held. At each update, for a given aircraft-slot pair (i.e. for a given arc), we define the incremental operating cost as the difference between the Operating Cost associated with the remaining portion of the flight and the aforementioned minimum operating cost.

The incremental Fuel Cost is then:

$$FC = c1 \cdot (FB_{slot} - NFB)$$

where:

- $c1$ is the cost per pound of fuel (\$/lb);
- NFB is the Nominal Fuel Burn, i.e. the fuel which would be burnt in the remaining portion of the flight if the aircraft were to fly at its nominal speed from its current location;
- and FB_{slot} is the fuel which would be burnt on the remaining portion of the trip if the aircraft were to fly the Best Flight Profile (defined in Section 2.3.2.) associated with the slot under consideration. For instance, for a slot later than the aircraft's Tright, it means that the aircraft will fly at its minimum speed (or holding speed) and be air held. The FB_{slot} would then be expressed as:

$$FB_{slot} = FBR_{hold} (SET_{slot} - t)$$

Note that it is possible to have a negative incremental fuel cost FC when FB_{slot} is less than the NFB. Indeed, the nominal speed is the speed which minimizes the user operating cost, not the fuel cost. Thus, an aircraft burns more fuel (see Figure 2-1) when it flies faster than the nominal speed but it burns less fuel when it flies slower than the nominal speed.

Now, let us call the Nominal Arrival Time, NAT, the time at which the aircraft would arrive at the Entry Fix if it were to fly at its nominal speed from its current location. Note that it is also the Nominal Exit Time, NET, since this “nominal” flight profile assumes no air holding. The incremental Time Operating Cost is then given by:

$$TOC = c_2 \cdot (SET_{\text{slot}} - NAT)$$

Note that the TOC can be negative when the slot under consideration is earlier than the NAT; in effect, in that case, the aircraft would not fly as long as in the aforementioned minimum operating cost flight profile which results in a saving as far as Time Operating cost is concerned.

- *Case of an aircraft on the ground:*

Recall that, for a given exit slot, we define the “best flight profile” such that any on-the-ground aircraft flies at its nominal speed toward the airport. During the flight the Operating Cost is minimum. Also, there are no Fuel Cost or Time Operating Cost for remaining on the ground; in fact, the only penalty associated with being in ground hold comes from the increase of planned delay. Thus the incremental user Operating Cost is zero for an aircraft on the ground, whatever the feasible slot under consideration:

$$OC = FC + TOC = 0 \text{ for an aircraft on the ground.}$$

2.3.3.1 Air Traffic Management Workload cost, W

When setting the Air traffic Management Workload Cost associated with a given arc, we must consider two kinds of costs: the Traffic Flow Advisory cost, TFA, and the Air Holding Delay cost, HD.

- The Traffic Flow Advisory, TFA:

IIDFC assumes that ATC determines Ground Hold (GH) and Cruise Speed (CS) traffic flow advisories to the fleet of aircraft proceeding inbound at each update but there are only a small number of *new* Advisories to be issued. There is an ATC workload cost associated with issuing a revised Cruise Speed to an airborne aircraft and a cost associated with issuing a revised Ground Hold to aircraft on the ground. For a given aircraft-slot pair, the ATC workload cost will then be the sum of two terms which depend on the pre-selected best flight profile for the slot (see section 2.3.2):

$$TFA = GH \text{ Change Cost} + CS \text{ Change Cost}$$

Controllers can themselves set the values of those costs and thereby determine the nature and volume of Traffic Flow Advisories. For instance, a cost of 1 can be issued for each action, or a high cost can be assigned to Cruise Speed to reduce the number of Cruise Speed Advisories issued at each update.

- Air Holding Delay Cost, HD:

At a system update, there is no immediate Traffic Management Cost associated with assigning an air holding delay to an aircraft since we do not issue a Holding Traffic Flow Advisory. However, when this aircraft arrives at the Entry Fix, it must then be monitored in the holding area. Thus, there is workload cost associated with planning an air holding delay for an aircraft. Thus when the Scheduled Exit Time associated with the slot under consideration, SET_{slot} is greater than T_{right} , we express the holding delay cost, HD, as:

$$HD = c3 \cdot (SET_{slot} - Tright)$$

where $c3$ is the cost per minute of planned holding delay.

Independently from the Air Traffic Management workload, this cost can be used to reduce or minimize air holding delay. Also, we may want to associate a higher coefficient $c3$ for holding at a particular entry fix when we desire to avoid air holding at that location.

The total cost for a given control arc will be a weighted function of those four costs:

$$\text{Cost} = w1 \cdot D + w2 \cdot HD + w3 \cdot OC + w4 \cdot TFA$$

Setting such a network model enables us to obtain an interactive flow control resolution. Indeed, by changing those weights, controllers can change the nature of the solution proposed, and they are always sure that the resolution generated is the best possible answer to the problem given the weights. Furthermore, for any aircraft assignment to a slot, it always minimizes the cost to the operators since it uses the Best Flight Profile. Furthermore, since the optimal logic evolves quickly, Traffic Managers can examine several new resolutions and control the number of Traffic Advisories issued at each update.

2.4 Extended Optimal Approach

In order to take into account a system where Air Traffic Flow Managers also specify Entry Fix Acceptance Rates (EFARs) as well as the Airport Acceptance Rate, we build another network model illustrated on Figure 2-12.

At each entry fix, we convert the entry fix acceptance rate to *entry slots*, which are in fact slots for departure from the holding fix, and we still create a set of slots for the

Airport Acceptance Rate. In fact the system is divided into as many “sub-system” as there are entry fixes. Each entry slot is then connected to an AAR slot.

Thus, each aircraft is then part of the sub-system corresponding to the entry slot at which it is supposed to arrive. It is then assigned to a unique entry fix slot by *control arcs* which reflect how Ground Holding, Cruise Speed control and Air Holding can change its holding fix arrival time. The corresponding arc costs are set the same way as explained in Section 2.3.3.




Figure 2-12: Network for Extended Optimal Approach (two-entry system with specified acceptance rates at each entry).

When an aircraft exits its air holding stack at which it was waiting for landing clearance, the controller can still divert it from its “normal” path between the holding

fix and the runway. This process, called spacing, allows to delay an aircraft up to two or three additional minutes. As shown in Figure 2-12, our model represents this process by *spacing arcs* which go from entry slots to airport slots.

Therefore, the set of spacing arcs originating from a given entry slot cover all airport slots from the time at which it would arrive at the runway without being spaced to the maximum time at which it would arrive by spacing controls.

A capacity of one unit is imposed on each entry slot (accomplished by splitting its node) so that only one aircraft can be assigned to an entry slot: this guarantees that the EFAR will be respected. By imposing an flow upper bound of one on AAR-slot arcs, we then ensure that at most one aircraft will be assigned to each AAR slot so that the rate of aircraft entering the terminal area is also less than the AAR.

Note that Air Traffic Managers are now able to control the Terminal Area spacing workload by putting a cost on the spacing arcs. For instance, let us consider the situation where we have two holding fixes and two aircraft arriving at the same time at a different holding fix. Both can be assigned their own EFAR slot (which have the same time position) but cannot be assigned the same AAR slot. If there is no cost on the spacing arcs then the planes will be directed to different AAR slots by spacing. However, if the spacing arcs cost are sufficiently high then those will be directed to different AAR slots by changing their EFAR slots, thus “transferring” the flow control workload from Terminal Area controllers to enroute controllers.

Another extension of the model can be mentioned here. In certain circumstances, some of the arrival aircraft may be qualified to proceed to land while others are not. Today, the weather at an airport is classified as Category 1,2 or 3 depending on ceiling and visibility for its approaches. If the aircraft is only qualified to land under Category 1, then it cannot proceed to land when the weather is Category 2 or Category 3. Thus, for the model described in Section 2.3.1, there will not be feasible control arcs from it to

AAR slots at such times. Similarly, a Category 2 aircraft cannot land under Category 3. In the extended model discussed above, we will use the time correspondence between AAR slots and EFARs slots to translate this “Category labeling” of the AAR slots to a similar “Category labeling” of EFARs slots; that is to say that there will not be feasible control arcs from a Category 1 aircraft to EFAR slots corresponding to times at which the weather is Category 2 or 3.

So, whenever weather forecasts for Category 1,2,3 weather exists, not only will the AAR be reduced but certain traffic will become unable to land. Both the network model presented in Section 2.3.1 and the above “extended model” will be able to give Traffic Flow Managers a best Traffic Flow Plan under such conditions, and show the planned number of holding aircraft over time at each entry fix.

Chapter 3

The Traffic Management Simulator

3.1 Overview of the simulator

The purpose of the Traffic Management Simulator is to model dynamically the main features of the system in order to test various Air Traffic Flow Management strategies. Within the framework of this thesis, we will more specifically confine ourselves in using it as a tool to investigate and evaluate dynamic flow control of arriving aircraft at a congested airport through the development of a set of dynamic algorithms.

The execution speed of the simulation as a whole was not of primary concern. Instead, in designing our simulator, we paid constant attention to enhance its maintainability and flexibility to model the system with various levels of details and perform a wide variety of air traffic flow management experiments and strategies. However, time efficiency of our dynamic flow algorithms was certainly an issue. Indeed, the

prospective users (the Traffic Flow Managers) would expect to get their strategy output fast enough to allow for examination of several options before issuing various new Traffic Flow Advisories.

With that in mind, the simulator was implemented in ANSI C language (18,000 lines of code, 179,508 bites) which ensures portability on DOS, UNIX and Macintosh platforms. The user is only limited by how much memory is available to run the simulation. The program includes, with the minimum necessary modifications to ensure portability, the Network Simplex Code for Minimum Cost Flow by Yusin Lee from the MIT Operations Research Center. This code proved to be particularly time efficient for our purpose, but all codes for Minimum Cost Network Flow problems are extremely fast for problems of the nature described here.

Traffic is generated by random landing requests from aircraft at different origins, along different arrival paths. Aircraft are of different types and, at any time, are on the ground or in the air proceeding inbound. Along the airways, there are time varying forecast winds. Eventually, there is a time varying Airport Acceptance Rate.

Dynamic flow is exercised every Tupdate (e.g. 15 min.) of simulator time. Certain aircraft are then issued a set of controls (ground hold, speed change, air holds) according to the dynamic flow control algorithm. Those controls are recorded for statistical review.

3.2 Modeling and design considerations

The purpose of this section is not to describe thoroughly our model but rather present its main features.

3.2.1 Aircraft generation

- Presently, the simulation can take into consideration three types of planes: heavy, large and small. Heavy aircraft would be of the B747, B767, A300, DC10, DC8 or L1011 type whereas the large class would include aircraft ranging from the turboprop like the ATR42 to the B757 and the DC9. Eventually, small aircraft would be the small piston engine aircraft and the smallest turboprops. We decided to limit ourselves to those three classes but the simulation user is free to define as many classes as he wants by just changing the value of the constant NTYPES in the header file and modifying the input file accordingly.

As a result, depending on the class of the trail and lead aircraft, different time separations are required for landing approaches to the same runway. If the simulator is used to analyze the efficiency of a metering process governed by dynamic flow algorithms, then we are only interested in the separation between exit times from the holding fixes. In that case, we only take into account the Airport Acceptance Rate in determining the time separation (separation = $1/\text{AAR}$).

- Furthermore, each aircraft is associated with a weather category (1,2 or 3) as explained in Section 2.4; for instance, a Category 2 aircraft can land by Category 1 and Category 2 weathers but is not qualified to land by Category 3 weather.

- For each holding fix and its associated routes, requests for landing from airborne or on-the-ground aircraft follow a Poisson process and can be generated at a time varying rate.

One could object that aircraft entering the system filed a flight plan and that, by doing so, they become in fact part of a schedule known by the ATC system well before their arrival at the airport. As a result, the arrival rate should vary throughout the day but only on a known (and repetitive over a month for instance) daily cycle. However, one cannot ignore the uncertainties in aircraft operations due to winds, runway configurations and departure times. For instance, the runway in use for departure at a given airport depends on the winds direction and intensity; upon leaving the gate, the taxi times to the end of the runway may therefore vary by 5-15 minutes; True cruise times may vary by ± 0.5 hours from day to day. Thus, even a deterministic schedule of arrivals resembles a Poisson arrival flow if there are large enough errors in maintaining the schedule. It would be possible to extend the generation of aircraft traffic to include input schedules.

- The original requested time of arrival at the holding fix of an airborne aircraft is computed from its location, the wind forecast along its route and assuming it desires to fly at its optimum cruise speed toward the airport. This time is also the Original Nominal Exit Time (ONET) from the holding fix since the ONET assumes that the aircraft will not be air held.
- An on-the-ground aircraft makes a initial request for an Original Take Off Time (OTOT). The pilot is restricted to announce his requested Take Off Time no less than Treqmin (e.g. 30 minutes) before departure. The ONET is computed from the originating airport location, the wind forecast and assuming that the aircraft will

take off at its OTOT and then fly at its optimum cruise speed toward its destination without being air held.

- For testing purposes we included the possibility for the aircraft to fly at a precise decimal airspeed (e.g. $M = 0.7204567$) as long as it is in the range of feasible cruise speeds. For the same reasons, the air traffic controller can also issue such a speed. However, there is also the much more realistic option to issue discrete cruise speeds i.e. speeds from a pre-defined set (e.g. $M=0.72$ or $0.73 \dots$).

3.2.2 Air Holding

Because of weather, operational incidents and uncertainties, the arrival flows often cannot be smoothed enough to match exactly the landing rate. In such an event, aircraft must enter the holding pattern.

In our simulation, we will consider the air holding to be *continuous*. If controllers are using horizontal holding (using the Early Descent control) then it is a reasonable assumption. However, this simulation feature may appear questionable when they only use vertical air holding. Indeed, in reality, aircraft hold in a stack by following the racetrack pattern shown in Figure 2-3. Typically, each turn or each straight leg requires one minute. However, the length of the straight legs can be reduced to zero. In that case, the holding pattern is reduced to a continuous turn. Thus, if an aircraft is to be held 6 minutes then it can follow one “normal” holding pattern (4 minutes) and one “reduced” holding pattern (2 minutes). Furthermore, if asked, the pilot can deviate from the racetrack pattern to be even more accurate and achieve a specified Exit Time from the holding fix. This shows that air holding can be considered continuous.

If the simulation is to be used to track landings then we integrate air holding and spacing in our “simulator air holding”. We can then also consider our “simulator air holding” to be continuous.

Furthermore, it is assumed there is no practical limit on the stack holding capacity. In other words, we assumed that there can be an infinite number of holding aircraft waiting for landing clearance. Instead, Air Traffic Flow Managers can impose a limit on the planned air holding for any aircraft by bounding the coverage of exit-holding-fix slots (“times” in the heuristic approach). The latest time at which the aircraft can exit the holding fix is then set to ($T_{right} + \text{maximum air holding}$).

Eventually, several strategies of insertion in the simulator holding queue can be considered. For instance, we could choose the First Come First Serve “server” or, instead, allow insertion of an aircraft according to the delay it has already endured. For instance, let us imagine the situation in which an aircraft has been held for 4 minutes when an aircraft, which is already 15 minutes late, arrives at the holding fix. In the second strategy, the last aircraft to arrive will be inserted before the holding aircraft.

3.2.3 The “freeze time”

At any time t , all aircraft which are within a flying time called “ T_{freeze} ” of their holding fix will not be issued new controls anymore. When t is the time at which the dynamic flow algorithms are executed, only aircraft which will arrive at the holding fix later than $t + T_{freeze}$ can be issued a command (Ground Hold, Cruise Speed change, Early Descent). This model reflects the actual “freeze zone” defined by air traffic flow managers in which they constrain themselves not to issue any new advisories to aircraft except for Air Holds. Aircraft arriving at a Terminal Area will begin descent

during T_{freeze} , and Congestion Management processes will shift to local metering and spacing.

Note that we chose to exercise the Dynamic Resolution only on aircraft which are scheduled to arrive at their holding fix later than $t + T_{freeze}$. The fact that there are other aircraft in the system -which could be holding at the entry fix- is taken into account by stating that the earliest SET that we can issue must be greater (by at least $1/AAR$) than the latest SET among the fleet of aircraft scheduled to arrive within T_{freeze} . At each update, this allows to account for actual holding at the Entry Fixes due to operational deviations. Furthermore, it reduces the complexity of the Dynamic Flow Control problem since we only represent aircraft whose SAT is greater than T_{freeze} in our network.

3.3 Simulator algorithm

3.3.1 A fast-time simulator

Within the scope of this thesis, we implemented a fast-time simulator or, more specifically, an event-paced simulator.

For the purpose of simulating dynamic Air Traffic Management, the choice to use an event-oriented versus an interval-oriented model is obvious. Indeed, in the situation where it takes on average two hours for a plane to fly from its originating airport to its destination, it would be very time inefficient to follow the entire progression of the aircraft in real time (say every minute) when we are only interested in the delay it underwent. A better approach would be to keep a list of forthcoming events (in this case, times of System Updates, Arrival at the holding Fix, and Landing) in order of occurrence and then step the clock to the time of the next event directly.

However, since requests for arrival at the holding fix are generated randomly, we cannot predict the future that is to say that we cannot keep a list of forthcoming events. Of course, one can object that there is no such thing as computer randomness and that, knowing the seed of our random generator, we can always pre-compute the times of the next events. In that case, pre-computing and then storing all those events would still be a tremendous waste of computer resources.

The solution we implemented was to have each event generate its time of next occurrence. Such a method is usually referred as *bootstrapping*.

The general algorithm of the simulator is presented in Figure 3-1.

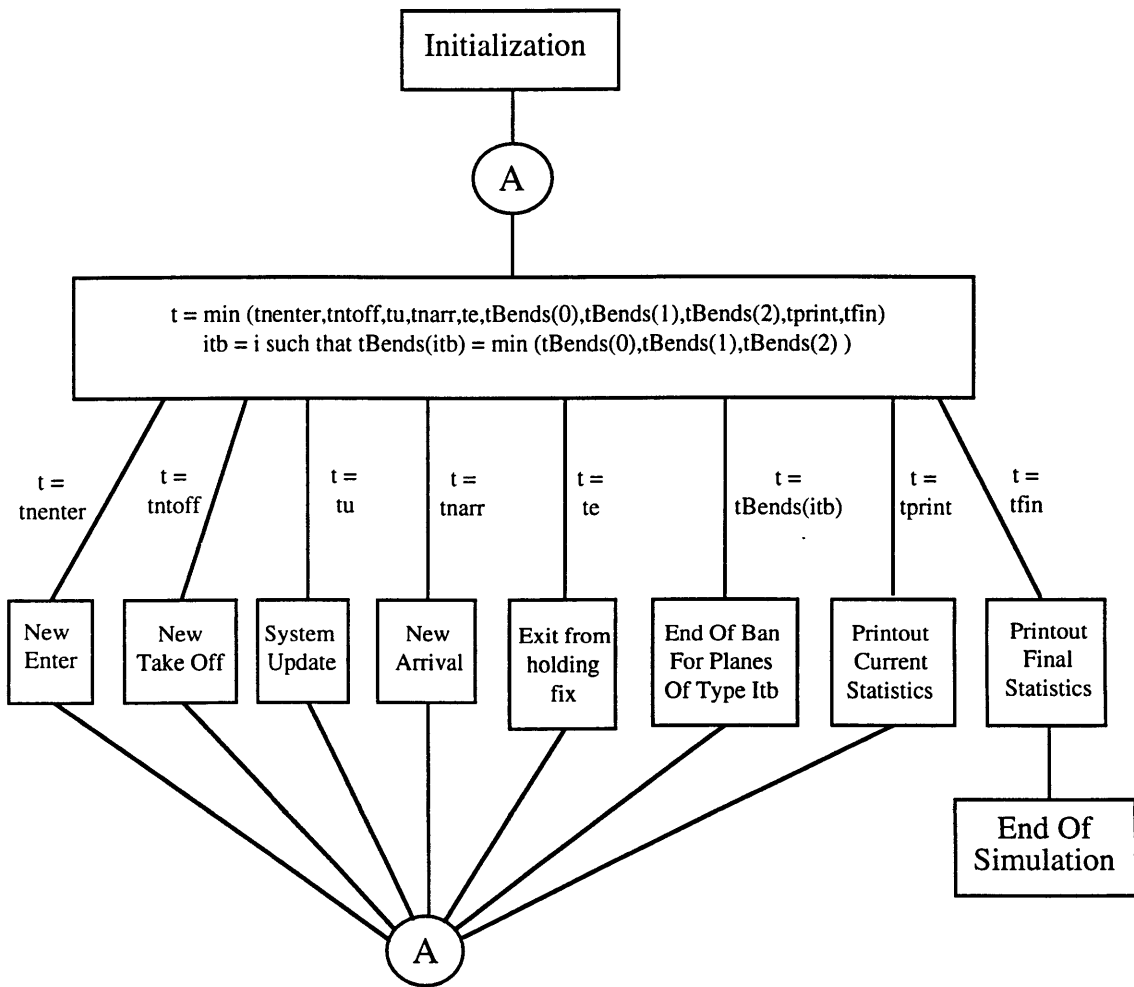


Figure 3-1: Overview of the Simulator Algorithm

As shown in Figure 3-1, each time an event is completed (i.e. algorithm returns to A), the clock will then jump to the time of the next event through the “minimum” operator. This will trigger the processing of the corresponding event.

Event “New Enter” is simply the event of receiving a request for landing from a new plane in the system. This plane can either be airborne or on the ground.

“System Update” is the event of updating the SETs by executing the dynamic flow control algorithms and issue the corresponding Traffic Flow Advisories.

An “New Arrival” is an arrival at the holding fix. And “End Of Ban For Planes Of Type itb” marks the clearance for exiting the holding fix for any plane of type “itb” (if we have three types of planes, then itb could be 1, 2 or 3)

Hereafter, we present the main functions of this code.

3.3.2 Initialization

As shown in Figure 3-2, the initialization consists mainly of six steps.

The first step initializes the statistics then $N[i]$ is set to zero for all aircraft, that is to say, all types of aircraft receive clearance for landing ($N[i] = 1$ means that aircraft of type i are not allowed to land at that point in time: i.e. there is a ban on landing for aircraft of that type).

We then create the heads of the three following queues: the SAT Queue which is the queue of planes ordered by their Scheduled Arrival Time at their holding fix; the Air Holding Queue consisting of airplanes which are air held; and the TOT-Queue or queue of planned Take Off Times of on-the-ground aircraft.

The next step is to initialize the ATC system by reading the input file. It defines the wind forecasts along the routes, the traffic characteristics (request rate for each arrival stream of aircraft, proportion of requests from on-the-ground aircraft, characteristics

of each type of plane..etc...), and the operational constraints: T_{freeze} , the freeze time and T_{update} , the fixed time period between two dynamic flow control executions.

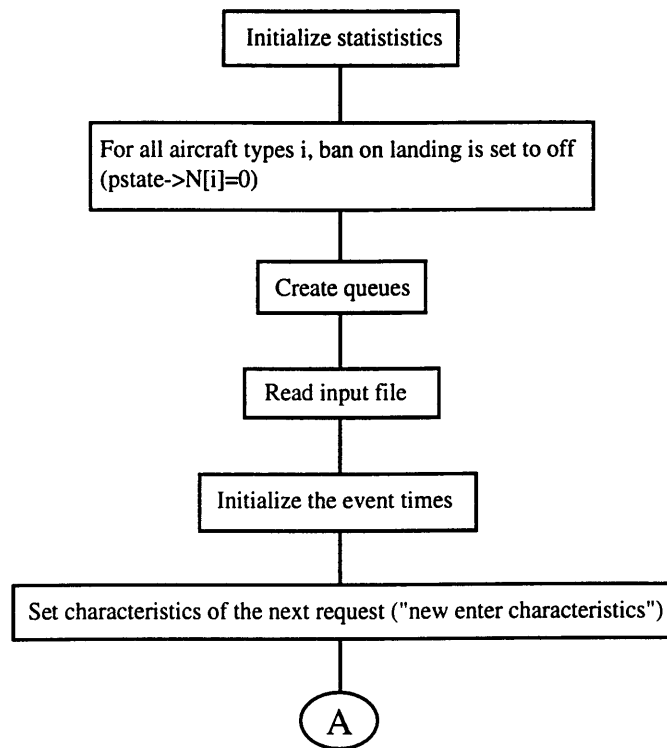


Figure 3-2: Initialization of the Simulator

We then initialize the event times:

- the simulation time is set to zero: $t = 0$.
- the time of the next enter, $t_{nent} = t + \text{random interarrival time}$. This interarrival time follows an exponential pdf with a parameter which is the sum of the request rates for all arrival streams so that in the end we obtain a Poisson process.
- the time at which dynamic flow control is executed, t_u is set to T_{update} if dynamic flow control using Ground Hold, Cruise Speed change and Air Holds, is selected. If the experimentator wants to see the results given by using only Air Holds, he can set the t_u to infinity.

- the time at which the first statistics are printed out in the output file, t_{print} is initialized by the user.
- t_{fin} is equal to the time at which simulation ends; it is entered by the user.
- all other event times shown on Figure 3-1 are set to "infinity". This guarantees that the first event will be either a New Enter, a System Update, a Statistics Printout or the End Of the Simulation

Eventually, we set the characteristics of the newly entered aircraft such as its location, if it is an airborne or on-the-ground aircraft ..etc...

3.3.3 New Enter

Event "New Enter" handles the case when an airborne aircraft enters the Traffic Management system or rather when it is first known to the Traffic Flow Manager in the dynamic flow control system. For an on-the-ground aircraft, "New Enter" is when the Flow Manager first receives its demand for take off. The algorithm of that function is outlined on Figure 3-3.

The Scheduled Arrival Time at the holding fix of this "newly entered" plane is set to the Original Nominal Landing Time which assumes that the aircraft flies at its optimum cruise speed from its location to the holding fix (see section 2.2.1). We also compute several quantities among which the distance from the airport at which the aircraft will be at the next System Advisory Update; we assume here that the plane will fly at its optimum cruise speed until that time.

Note that this aircraft is then always included in the SAT-queue, whether it is airborne or not. This is to be sure that the dynamic flow algorithm, which takes only into account aircraft which are more than T_{freeze} away from their entry fix in the SAT-Queue, will also apply to on-the-ground aircraft. Furthermore, let us emphasize that,

whereas the Air Holding Queue and the SAT-Queue are queues of aircraft, the TOT-Queue is just a queue of Take Off Times.

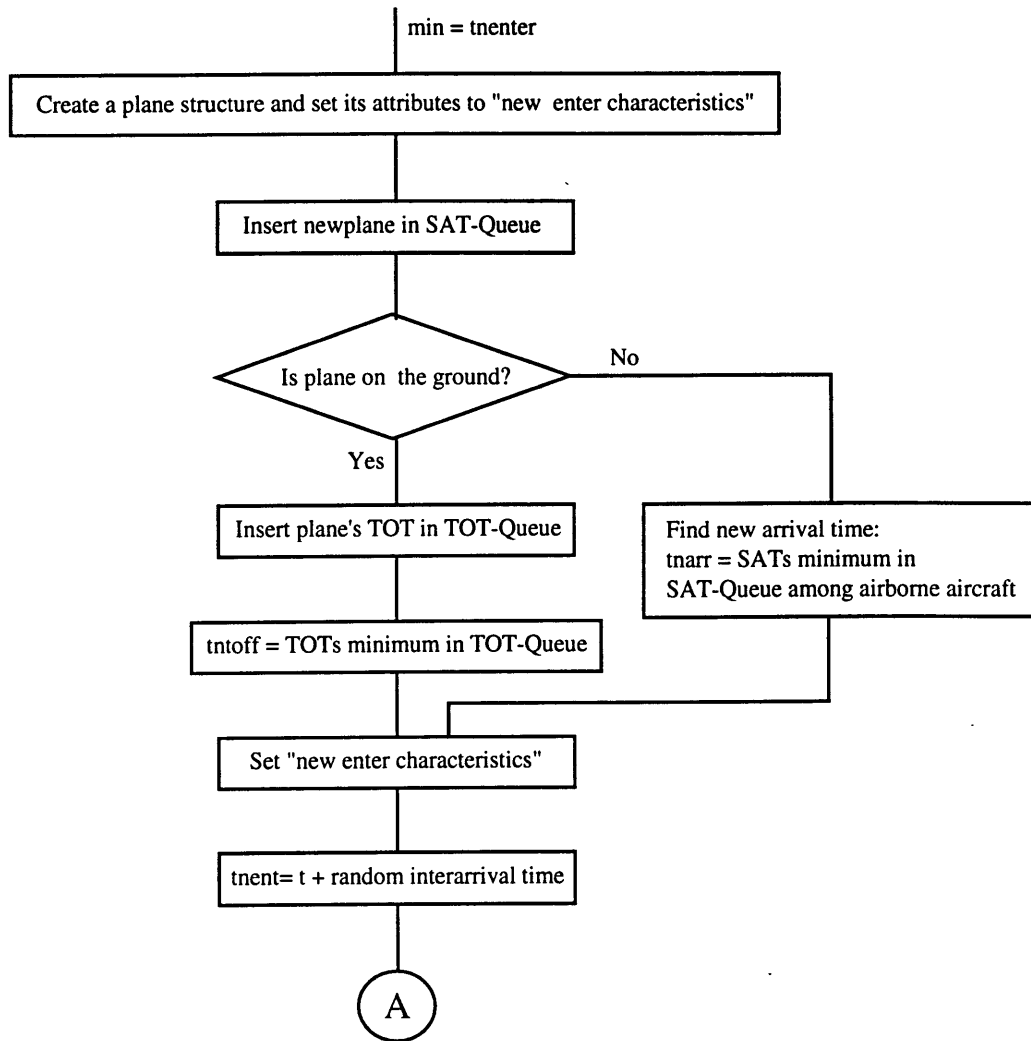


Figure 3-3: New Enter Function of the Simulator.

Thus, a given plane cannot be in the Air Holding Queue and the SAT-Queue at the same time but can be in the SAT-Queue and also be represented in the TOT-Queue.

The rest of the algorithm showed on Figure 3-3 is pretty self explanatory. Let us just note the bootstrapping method which not only sets the time of the next-system-enter but also its characteristics.

3.3.4 New Take Off

When this function is triggered, the first aircraft in the TOT-Queue, which is then the on-the-ground aircraft whose TOT equals the current simulation time, takes off. In other words, its TOT is deleted from the TOT-Queue and its SAT-Queue status is changed from “on-the-ground” to “airborne”.

The t_{ntoff} , time of the next take off, is then updated to the new first TOT in the TOT-queue, or infinity if the queue is empty.

Eventually, note that the t_{narr} , time of new arrival at an entry fix, should also be updated in this function. Indeed, the fact that a plane in SAT-queue turned from on-the-ground status to airborne status might change the t_{narr} . For instance, let us imagine the case where the SAT-Queue consist only of one aircraft, which is on the ground. In such a situation the t_{narr} would be infinity. But, if the new take off event is triggered, then this aircraft becomes airborne and the t_{narr} must be set to the SET of this newly airborne aircraft.

3.3.5 System Update

This event consists in executing the dynamic flow control algorithm, find the new SETs and issue the Traffic Flow Advisories. Either the heuristic or the optimal approach can be used at that point. It is therefore the step involving the most computations and as a result the speed limiting factor of the simulator. When one is using the optimal approach, there are several way of proceeding. Indeed, we could choose to preconstruct a network and then update it at each Dynamic Resolution by setting the new costs and flow bounds of the arcs. However, we have to be aware that there is a cost of keeping the same network: its complexity will be in fact very high since it

should be able to fit all situations which could occur. Thus, we decided to construct a new network at each step.

As explained earlier, dynamic flow control is exercised only on aircraft which are within T_{freeze} of the holding fix. Of course, if there are no such aircraft then the time at which the next dynamic flow control occurs is updated and the function exits (see Figure 3-4).

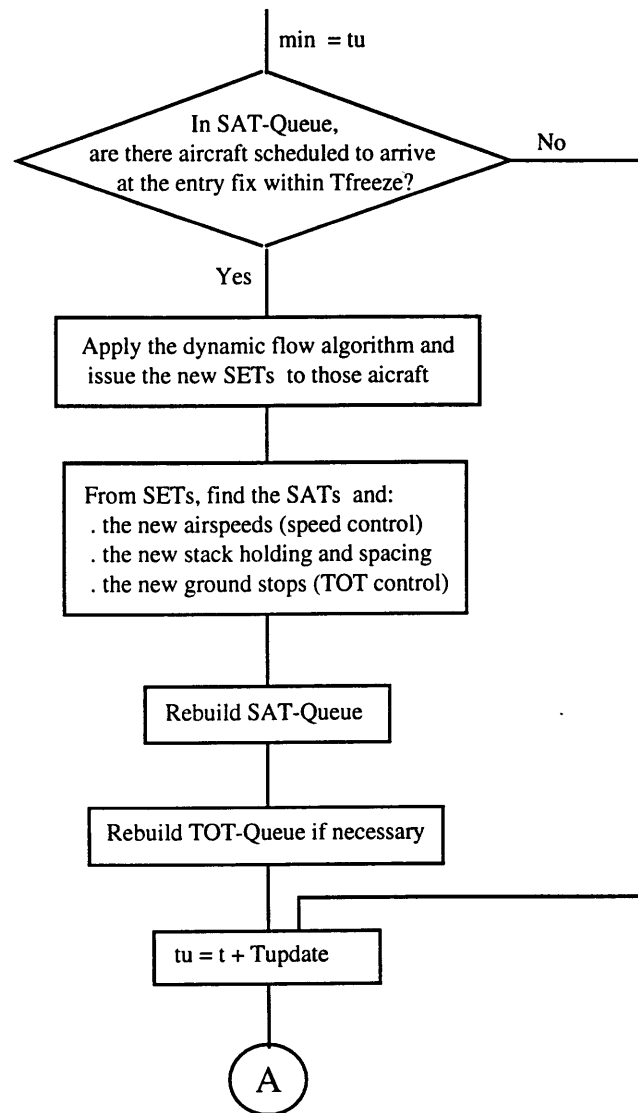


Figure 3-4: Dynamic Resolution Logic of the Simulator

Note that the dynamic flow control algorithm produces a plan and we first assign the corresponding SETs to the fleet proceeding inbound. It is only then that, depending on the associated best flight profiles (see Section 2.3.2), the aircraft is attributed a SAT and the corresponding controls: Cruise Speed change, new Take Off Time or Air Holds.

Of course, following this step, the TOT-Queue and the SAT-Queue must be rearranged and the time of the next System update set to $t + T_{update}$.

3.3.6 New Arrival

This event is triggered when an aircraft arrives at the holding fix i.e. when the simulation time is equal to the SATs minimum among airborne aircraft in the SAT-Queue. Its algorithm is mapped on Figure 3-5.

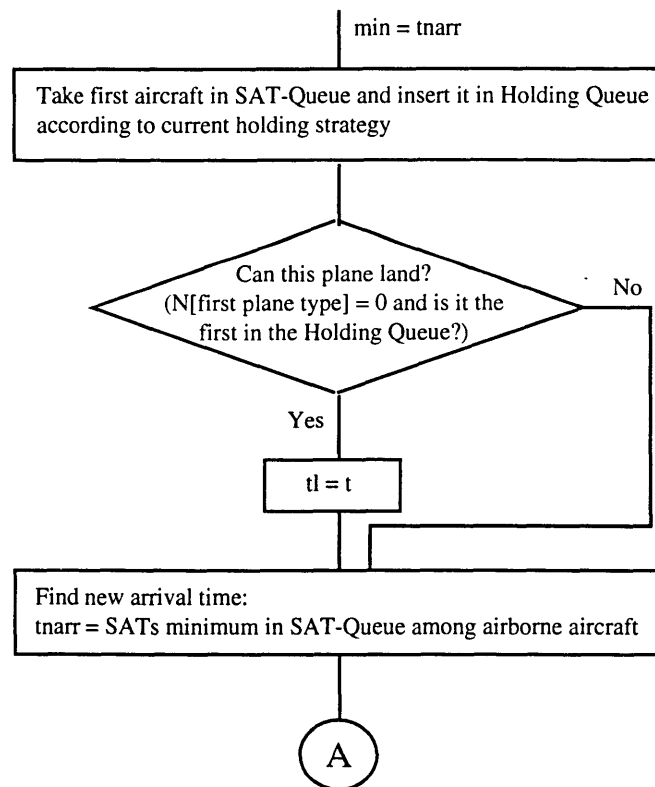


Figure 3-5: New Arrival Function Algorithm

At that point in time, there may be only a certain type of aircraft allowed to exit the holding fix to proceed for landing. For instance, the weather might only allow landing of Category III aircraft. Or, if an aircraft has previously exited the holding queue, we do not want to have another one exiting the queue before some time has elapsed.

If there is no ban on exiting the holding fix for this aircraft (that is to say if there is clearance for exit for this type of aircraft) and if, after insertion in the Air Holding Queue, it is the first plane in queue, then t_e , the time of next exit from the holding fix, is set to the current simulation time. By doing so, the simulator will jump to the "Exit from the holding fix" event at the next time the minimum operator is called. This will for instance apply to the first aircraft which arrives at the holding fix.

Note that in either case (whether there is or there is not an exit clearance for this aircraft), we insert the aircraft in the Air Holding Queue. This trick allows to reduce significantly the complexity of the simulator algorithm.

3.3.7 Exit from the Holding Fix

When the current simulation time is equal to the Exit-from-the-Holding-Fix time, t_e , the event to be processed is an exit from the holding fix.

As shown in Figure 3-6, note that it is always the exit of the first aircraft in the simulator Air Holding Queue. The main reason is that we always insert an arriving aircraft at the holding fix in the Air Holding Queue, even if the queue is empty (see Figure 3-5). However, it is easy to modify this part of the algorithm (and the test in Figure 3-5) such that the exit of an aircraft does not require that the plane is the first one in the simulator Air Holding Queue.

Statistics are updated every time an aircraft exits an holding fix to take into account its attributes we are interested in (total, air holding or ground holding delay; type;

number of speed changes..etc..). Note that it is the difference between the ONET and the actual exit time which measures the delay of the aircraft.

At that point, there is the option to print the characteristics of this "newly exiting" plane into an output file.

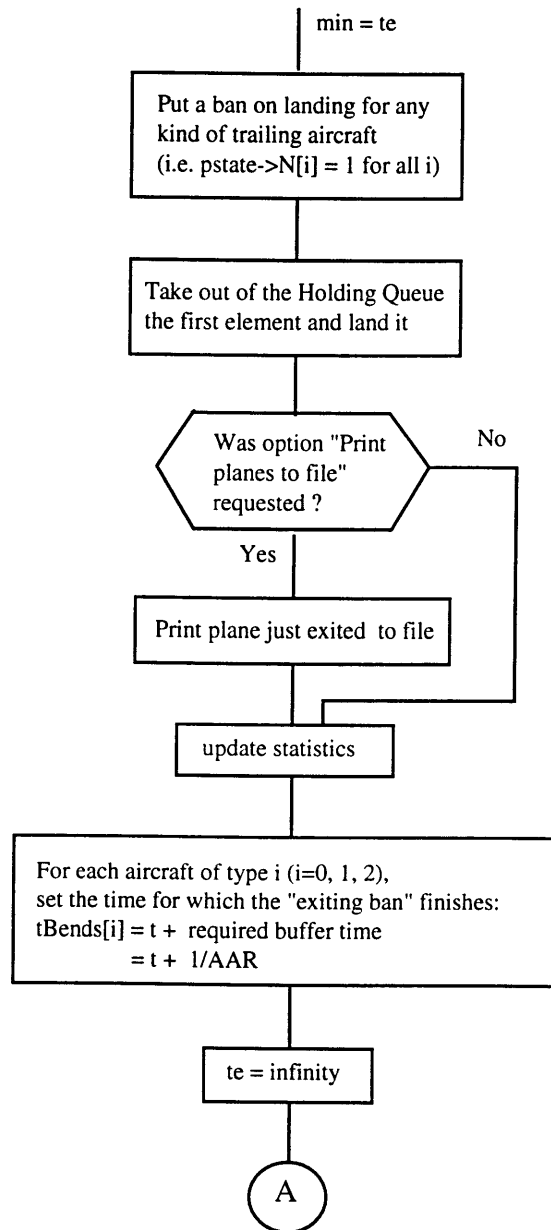


Figure 3-6: Handling a "Exit-from-the-holding-fix" Event within the Simulator

We then put a "ban on exiting the holding fix" for any kind of subsequent aircraft that is to say we prohibit the subsequent aircraft exits as long as the time separation requirement is not satisfied. The only way a given type i_0 of aircraft can then get a clearance for exiting is when $t = t_{\text{Bends}}[\text{type } i_0]$ (see section 3.3.8). Note here that the t_{Bends} depends on the current time and the Airport Acceptance Rate. There is a possibility to take into account the wake vortex time separation requirements. In that case the separation would depend on the type of the aircraft which exits and the type of the subsequent aircraft. Thus, for one aircraft which exits, there would be as many t_{Bends} as there are aircraft types.

3.3.8 End Of Ban

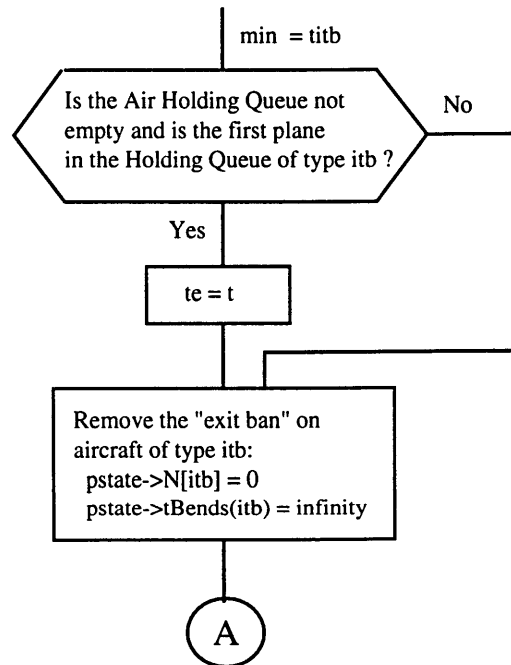


Figure 3-7: Handling an "End Of Ban" Event within the Simulator

This function and the $t_{\text{Bends}}[i]$ were introduced to model a situation in which we take into account the wake-vortex effects which result in time separation requirements

depending on the weights of the leading and trailing aircraft. Of course, it also suits the situation where we only consider $1/AAR$ as a separation requirement.

At any time, if an aircraft is in the Air Holding Queue, the only way it can exit it is in the situation where it is the first aircraft in this queue and where there is no ban on landing for aircraft of its type.

The ban on aircraft type itb ends when the current simulation time is the exit-time of the last aircraft which exited plus the minimum time separation requirement for a trailing aircraft of type itb (see Figure 3-7). At that time, if the first aircraft in the holding queue is of type itb , it can then proceed to landing ($t_e = t$). Otherwise it will have to wait for the end of ban on aircraft of its type to be triggered.

Chapter 4

Case Studies

4.1 Approach and Input File

4.1.1 Approach

The purpose of this section is to evaluate some features of IIDFC through data obtained by running the simulator. Our goal is to study and evaluate the efficiency of IIDFC at handling a shortage of capacity at the airport, that is to say a period where the rate of requests for landing at the airports exceeds the Airport Acceptance Rate (AAR). It is important to note that the Traffic Flow Managers have a forecast of the AAR but, since requests are generated randomly (at the arrival traffic rate known from the airline schedule) in the simulator, they do not know the actual arrival schedule in advance. For instance, they know about a request of an airborne aircraft only once it has been “generated” that is to say when it enters the Flow Management System; and they know about the request of an aircraft on the ground when its flight

plan (perhaps revised) has been filed. This request information corresponds to that known to the current Aircraft Situation Display in the USA.

4.1.2 Input File Analysis

We include most of the parameters which affect the system and the dynamic resolution in the input file so that the user of the simulator can vary all of them without having to recompile any program file. As illustrated on Figure 4-1, the input file consists of four parts:

- The *Airport* part defines the forecast of Airport Acceptance Rate. In the case of figure 4-5, the AAR is forecasted to be 60 aircraft/hour from simulation time $t=0$ to $t=7$ (in hours) and then to drop to 30 aircraft/hour for three hours (from $t=7$ to $t=10$); it then goes back to normal (i.e. 60 aircraft/hour). We consider that there are two arrival paths at the airport ("Stream 1" and "Stream 2"). Along those paths, it is possible to define wind magnitudes but for this study we put a zero wind. In effect we can read that, on Stream 1, the magnitude of the wind is zero in the three "wind zones": between 0 and 300 nm from the airport (the first "wind zone"); between 300 and 600 nm (the second "wind zone") and between 600 and 900 nm.
- The second part of the input file, *Traffic*, first defines the rate of request for landing from aircraft flying along Stream 1 (20 requests/hr) and 2 (20 requests per hour). As indicated, 90% of those requests come from aircraft on the ground between 45 minutes and 3 hours before planned departure time.

Requests are randomly generated at three locations along each stream defined by their distance from the airport. Figure 4-1 illustrates the case where requests are generated at 400 nm (30% of the requests), 600 nm (30%) and 800 nm (40% of the

requests) along both streams. This information determines how long the aircraft is handled by Traffic Flow Managers before its arrival at the airport.

```

- AIRPORT:
  AAR Forecast:
    Number of different AAR forecasted = 3
      t      AAR
      0.0    60
      7.0    30
      10     60

    Three wind zones on each stream:
      Locations (nm):
        Stream 1=  0 300 600 900
        Stream 2=  0 300 600 900
      Wind magnitudes:
        Stream 1 = 0 0 0
        Stream 2 = 0 0 0

- TRAFFIC:
  Request Generation rate for the two streams (#/hr):
    stream 1 = 20
    stream 2 = 20

  Fraction of requests from on-the-ground aircraft = 0.9
  Minimum flight plan filling time before departure = 0.75
  Maximum flight plan filling time before departure = 3

  Requests Locations:
    Stream 1:
      Three Locations (nm) = 400      600      800
      Breakdown (%)       = 0.30     0.30     0.40
    Stream 2:
      Three Locations (nm) = 400      600      800
      Breakdown (%)       = 0.30     0.30     0.40

  Percentage of small aircraft = 0.25
  Percentage of large aircraft = 0.25
  Percentage of heavy aircraft = 0.50

  Speed Ranges (knots):
      min.    nom.    max
  small      400    450    500
  large      450    500    550
  heavy      500    550    600

- IIDFC:
  Tupdate (hr.) = 0.25
  Tfreeze (hr.) = 0.5
  Commitment to current departure time = 0.5

  Arc Cost:
    weight of Delay cost = 50.0
    weight of Holding Delay Cost = 50.0
    weight of Flight Operating Cost = 0.0
    weight of Traffic Flow Advisory Cost = 0.0

  Delay Cost:
    cost per squared minute of late delay ($/min) = 0.33 0.33 0.33
    cost per minute of early delay ($/min) = 0 0 0.0 0.0

  Holding Delay Cost
    cost per minute of holding delay ($/min) = 30 0 30.0 30.0

  Incremental Flight Operating Cost:
  Operating Time Cost:
    Operating Time Cost per minute of flight ($/min) = 1.0 1.0 1.0
  Fuel Cost:
    Fuel Burn Rate per knot above nominal speed (lb/hr/knot) = 30 30 30
    Fuel Burn Rate per knot below nominal speed (lb/hr/knot) = 15 15 15
    Nominal Fuel Burn Rate (lb/hr) = 4000 6000 8000
    Cost of fuel per lb ($/lb) = 0.1

  Traffic Flow Advisory Cost:
    cost of issuing a new Cruise Speed ($) = 1.0
    cost of issuing a new Ground Hold ($) = 1.0

- SIMULATION ALGORITHM:
  Maximum additional ground hold issued at every update to each aircraft (hr.) = 4.0
  Maximum air holding issued at every update to each aircraft (hr.) = 2.0

```

Figure 4-1: Input File Format

Thus, those are crucial parameters of the Dynamic Flow Control System since they determine how much information Traffic Flow Managers have about the future planned arrival rate at the airport and also determine how much they can control an aircraft. For instance, an airborne aircraft “popping up” at 300 nm from the airport cannot be speed controlled very much.

This part also defines the fleet characteristics; that is to say, the breakdown between small, large and heavy aircraft (see Section 3.2.1) and their minimum, nominal and maximum cruise speeds. It is expected that each operator will supply this information on the range of possible cruise speeds for each flight plan.

- The third part of the input file, IIDFC, defines how Dynamic Flow Control is exercised when selected (the program prompts the user to ask if tactical Air Holding only or IIDFC -heuristic or optimal- should be exercised). It first defines T_{update} , T_{freeze} (see Section 3.2.3) and the time period before current take off time of an on-the-ground aircraft when we commit to that Take Off Time, i.e. we do not issue it a new ground hold anymore. Then the cost structure to be used for computing the cost of a given arc in the network is described.
- The last part of the input file, Simulation Algorithm, is directly related to the way the network is built by the Dynamic Resolution Logic. In the case illustrated by Figure 4-1, it means that, for an on-the-ground aircraft, the program will construct arcs from its current Earliest Exit Time, EET , to $T_{right} + 3hrs$, and, for an airborne aircraft, between EET and $T_{right} + 2hrs$. Those two constants then determine how many arcs are constructed in the network and, thus, the execution time for the dynamic algorithm.

For instance, if each aircraft can cover 3hrs of airport slots when the AAR is 60 aircraft/hour, it means that a set of $3 \cdot 60 = 180$ possible control arcs originates from each aircraft. When there are 200 aircraft in the system it means that 3200 control arcs are constructed! However the execution remains very fast.

The construction parameters should depend on how long the shortage of capacity is lasting. Indeed, we should at least be able to assign an on-the-ground aircraft which is currently scheduled to arrive at the holding fix at the beginning of the capacity deficit period (e.g. here $t=7$), to several slots after the end of that period (e.g. here $t=10$).

In the following studies, we confine ourselves to analyze the effect of changing the cost structure, i.e. the only part of the input file which changes is the part concerning IIDFC. However, we do not pretend to carry out a detailed sensitivity analysis. There are many more experiments to undertake in later studies to examine the performance of the IIDFC concept.

4.2 Scenarios

In the following scenarios, the *Airport*, *Traffic* and *Simulation Algorithm* parts of the input file are identical. Recall that AAR is 60 aircraft/hr from $t=0$ to $t=7$, 30 aircraft/hr from $t=7$ to $t=10$, and then 60 aircraft/hour from $t=15$; the arrival rate is 40 aircraft/hour. Thus we have a shortage of capacity during 3 hours, from $t=7$ to $t=10$. In each of the following sections, we change the parameters of the IIDFC part. *For each run, data, plots and corresponding notations are given in Appendices.*

4.2.1 Scenario 1: Tactical Air Holding -No IIDFC

This run is a benchmark run where IIDFC is not exercised and only tactical air holding is used: aircraft are neither issued Ground Holds (GH) nor Cruise Speed Changes (CS). Thus all the delay comes from being held at the Entry Fix. As seen on Figure 6-3, the

delay of landed aircraft is small as long as the AAR is greater than the arrival rate, that is to say until $t=7$. It then builds up very fast, up to 46 minutes for the last aircraft which landed at the end of the shortage of capacity (see Figure 6-1). Since this delay comes only from holding, we see in Figure 6-2 that the number of aircraft holding the Entry Fixes also increases accordingly during that period to reach 16 aircraft at Entry Fix 1 and 18 aircraft at Entry Fix 2 before the AAR returns to 60 aircraft per hour. There are 33 aircraft in air hold when the system reaches its busiest state at $t=10$.

When AAR goes back to normal (60 aircraft/hr), air holding delays return to zero in an hour and 45 minutes.

At $t=15$, 480 aircraft have landed. Among those, were 423 aircraft which were on the ground when they made their original request for landing ("the ground-start aircraft"). For all aircraft which landed at $t=15$, the total delay amounts to 5233 minutes (10.9 mins per aircraft); 4867 minutes (11.5 mins per aircraft) came from the ground-start aircraft and 366 minutes (6.4 mins per aircraft) from the airborne-start aircraft.

4.2.2 Scenarios 2,3,4,5: IIDFC

In all these Scenarios, IIDFC is used, and thus controllers issue Traffic Flow Advisories (Ground Hold Advisory, GHA, and Cruise Speed Advisory, CSA) to aircraft proceeding inbound to the airport under Dynamic Flow Control.

4.2.2.1 Scenario 2

All parameters are identical to those defined in Figure 4-1 except the weights in the arcs cost and the delay cost structure. Indeed, in this scenario, we assign a weight of 50% to both the Delay and Air Holding Delay, and zero percent to the Flight Operating Cost and the Traffic Flow Advisory Cost. Thus, we are trying to use IIDFC

in order to minimize a weighted sum of Delay and Holding Delay. Also, the cost for being late is a linear function of the delay with coefficient \$10/min, and we put a non-zero cost for being early of \$1/min. This is to reduce the tendency that the algorithm would have to speed up the planes when it is not necessarily needed.

Still, we observe some "negative delay" before $t=7$ when landings are early (see Figure 6-4). Indeed, since the penalty associated with an early delay is not high, the Dynamic Flow Control Algorithm will prefer to move 9 planes 1 slot ahead than to delay one aircraft by 1 slot.

Recall that, from $t=7$ to $t=10$, the AAR is 30 aircraft/hour whereas the requests for arrival rate is 40 aircraft/hour (on average equal to 36 ground-start aircraft/hour plus 4 air-start aircraft/hour). In order to make the actual arrival rate match the AAR, we see on Figure 6-5 that Ground Hold Advisories are being issued from $t=5.5$; that is to say, an hour and a half before the shortage of capacity. Note that it is approximately the average flying time to the airport. Thus, it means that ground start aircraft which ask to arrive at the beginning of the "capacity deficit" period are issued Ground Holds. This demonstrates that the Dynamic Flow Control algorithm prefers to use Ground Holds.

Figure 6-6 shows that, effectively, the maximum number of aircraft in Air Hold is now 5 aircraft at $t = 9.75$. Thus, we have been able to reduce the maximum number of aircraft in Air Hold by 28 (as compared to Scenario 1), from 33 to 5 ! Consequently the total air holding delay is now only 287 minutes as compared to 5233 minutes when only tactical air holding is used.

Furthermore Figure 6-8 illustrates the fact that most ground-start aircraft which were issued a ground held, land after the capacity deficit period -when Air Holding Delay decrease (see also Figure 6-7).

IIDFC also greatly reduced overall delay (see Figure 6-3 and 6-7). The cumulative delay for all aircraft which landed at $t=15$ is 3594 minutes (=7.4 mins/aircraft) as

compared to 5233 minutes (10.9 mins/aircraft). That is, total delay has been reduced by 30%! This is because aircraft are being asked to speed up to fill gaps instead of taking delays. Note that both the delay for ground-start and for air-start aircraft were reduced; it has been reduced by a factor of 5 for air-start aircraft and by a factor of 1.4 for ground-start aircraft. There is 259 mins (=0.61 mins/aircraft) of air hold as compared to 6.1 mins/aircraft of ground holding for ground-start aircraft. The average number of Cruise Speed Advisories issued to each aircraft is 1.8 for an average flying time in the system of one and a half hour.

4.2.2.2 Scenario 3

Scenario 3 has also the parameters given on Figure 4-1 except that, as done in Scenario 2, the cost for (positive) delay is a linear function with coefficient \$10/minute. There is no cost for “negative delay”. The weights of the costs differ from Scenario 2 since we now assign an almost equal weight to Delay, Air Holding Delay and Flight Operating Cost. This reduces the importance of minimizing both Delay and Air Holding Delay.

As shown on Figure 6-10, the pattern of Ground Hold Advisories and Speed Advisories is the same as in Scenario 2.

This results in an increase by 10% of the cumulative delay for all aircraft which landed at $t=15$ (from 3594 to 3942 minutes) as compared to Scenario 2 (see Figures 6-4 and 6-9). Of course, we still have a cumulative delay much less than in Scenario 1 where only tactical air holding was used. Air Holding Delay has a similar increase from Scenario 2, from 287 minutes to 301 minutes.

However, we can see that adding a weight on the Flight Operating Cost has almost totally eliminated the “negative” delay (i.e. aircraft arriving ahead of schedule) that we observed in Scenario 2 (without this time putting a cost for “negative delay”).

4.2.2.3 Scenario 4

This Scenario has the same parameters as Scenario 2 and 3 except that this time, we take into account the cost for issuing Traffic Flow Advisories: we assign an equal weight of 25% to each one of the costs (Delay, Holding Delay, Flight Operating Cost, Traffic Flow Advisory).

This results in a new increase of 3% in cumulative total delay from Scenario 3 to 4073 minutes (= 8.5 mins/aircraft) as shown on Figure 6-14. However, IIDFC still saves 20 hrs of cumulative total delay as compared to tactical air holding (Scenario 1). Cumulative air holding delay is substantially up from 301, in Scenario 3, to 1270 minutes (422%) as shown by comparing Figures 6-9 and 6-14. Here again, we are still far below the 5233 minutes of cumulative air holding delay observed when IIDFC is not used; i.e., in Scenario 4 we are saving $5233 - 1270 = 3963$ minutes of Air Holding Delay as compared to Scenario 1.

Comparison of Figures 6-5, 6-10 and 6-15 shows that the number of Cruise Speed Advisories has diminished as compared to Run 3. The total number of Cruise Speed Advisories which were issued during the simulation was reduced by 18% (from 744, in Scenario 3, to 595 here; see Figure 6-9 and 6-14). We also observe that the total number of Ground Hold Advisories has literally plunged down to 37 from 431. However, comparing Figures 6-13 and 6-18 shows that, since the number of Ground Hold Advisories was considerably reduced from Scenario 3 to Scenario 4, aircraft which were ground held spent much more time in Ground Hold. See Figure 6-18 where aircraft landing around $t=10$ have averaged 90 minutes in ground delay as compared to 30 minutes in Scenario 3. Indeed, consider the situation where aircraft A has already been ground held 1 hour at the time where an aircraft B, which has not yet been held, asks for take off clearance. Because the delay cost is a linear function of delay, the extra cost for holding on the ground for the same amount of time each of those aircraft is the same. Scenario 5 shows how to solve that problem.

4.2.2.4 Scenario 5

In Scenario 5, we assign an equal weight of 25% to the Delay, Holding Delay, Flight Operating, and Traffic Flow Advisory costs as done in Scenario 4. However, the cost for being late on the original schedule is now a quadratic function of the Delay; i.e., for each slot and corresponding Scheduled Exit Time, SET_{slot} , the delay cost is:

$$D_{cost} = k \cdot (SET_{slot} - ONET)^2 \text{ when } SET_{slot} > ONET,$$

and we choose $k = \$0.33/\text{min}^2$

This quadratic cost for delay eliminates the problem mentioned in the above section. Indeed, comparison of Figure 6-18 and 6-23 shows that the maximum ground delay experienced is now less than 40 minutes versus more than 90 minutes in Scenario 4. In effect, at a given system update, aircraft which have already been in Ground Hold for some time are given priority for take off clearance over aircraft which have just made their first request for take off (and which have not been in Ground Hold yet). Thus, there is now a First Come First Serve system for aircraft on the ground which results in a rotation of aircraft which are ground held. Of course, it increases the cumulative number of Traffic Flow Advisories; and particularly the number of Ground Hold Advisories increased from 37 in Scenario 4 to 199 in Scenario 5 (see Figure 6-14 and 6-19). However it still remains 50% less than in Scenario 3 (where it was 451 since there was no cost for Traffic Flow Advisories). Also, as shown in Figures 6-10 and 6-20, the maximum number of ground hold advisories issued at an update is now 38 as compared to a maximum of 56 in Scenario 3.

Delay (and above all Holding Delay) remain well below their level of Scenario 1 where only tactical air holding was used (not IIDFC). Cumulative delay is reduced by 15% for the whole fleet, by 55% for air-start aircraft and by 12% for ground-start aircraft. Cumulative holding delay has almost been eliminated: it is approximately down 95% for air-start and ground-start aircraft.

Chapter 5

Conclusion

5.1 Main results

A new approach for handling Air Traffic Flow Management of arriving aircraft at a congested airport has been presented and investigated.

At this stage, we can claim that Integrated Interactive Dynamic Flow Control, IIDFC, has a great potential to achieve efficiencies in Dynamic Arrival Flow Management at a given airport. In preliminary testing, it is behaving as we expected: not only can it trade off User Costs and Traffic Flow Management Costs at levels which can be modified at any time during the process, but it can also account for the evolution of critical parameters (e.g. the Airport Acceptance Rate, actual holding delays, wind variations, etc.) in real time.

We have built a powerful Traffic Management Simulator which allowed us to test various “responses” of IIDFC to typical situations where the Traffic Flow Managers have a forecasts of the Airport Acceptance Rate. This tool should be very useful in subsequent experimentation and development work for IIDFC.

5.2 Future Directions

Integrated Interactive Dynamic Flow Control for congestion management at a given airport seems to be a promising direction. A number of issues are left to be explored:

- First, there is considerable **experimentation** to be undertake. This could include:
 - *A Detailed Sensitivity Analysis* which should investigate the influence of key parameters such as the time between subsequent system updates, the commitment time of ATC before departure of aircraft, or the costs of using various forms of controls - air hold, ground hold and speed control. There are a number of typical scenarios where the capacity deficit, capacity margins, length of capacity deficit (etc.) need to be investigated.
 - *Hedging*: when a forecast of Airport Acceptance Rate (AAR) is known but Traffic Flow Managers are uncertain about the time when it will go down, they can “distribute” this decrease over some time period. For instance, a decrease of the AAR which is supposed to go down between 6pm and 7pm and then reach its low (or high) value, can be modeled by a linear decrease between 5pm and 7pm so that Traffic Flow Managers can *hedge* against possible uncertainties in the forecast.
 - *The Response time to Actual Evolution*, when, for instance, the AAR unexpectedly goes down because of operational deviations. Furthermore, the robustness of IIDFC to errors in surveillance, tracking, winds, expected departure times (etc.) as a function of the time period between subsequent system updates should be examined.

- Secondly, there are **extensions of the IIDFC concept**. The Extended Optimal Approach presented in Section 2.4, which takes into account specified Entry Fix Acceptance Rates, is yet to be investigated more thoroughly, both theoretically and experimentally.

5.3 Personal Experience

This research was the occasion for me to discover what I consider a fascinating and promising field: Dynamic Arrival Flow Management at a congested airport. On a technical perspective, I believe I learned tremendously about operational constraints in Air Traffic Control and Airline Operations. While I was unable to carry a detailed sensitivity analysis using the Traffic Management Simulator, I also learned about “simulation-based” studies.

On a more general perspective, this study was the occasion to taste the joys and difficulties of research. I have learned a lot about methodology in research. This experience should definitively be part of any graduate scientific education. It is the source of my motivation to come back to MIT to enter the Ph.D. Program in Air Transportation.

Annex

A.1 Scenarios Notations

In this section, we present the data which was obtained from running the simulator under the scenarios described and analyzed in Chapter 4.

For each scenario, we present the statistics which are currently tracked within the simulator in figures entitled "*Tab of Statistics vs. Time*". Let us explain, for one row -i.e. at a given time t - what they mean:

- t is the simulation time in hours.
- E is the number of aircraft which *exited* the Entry Fix, that is to say which entered the Terminal Area, between $t - 0.25$ (i.e. $t - 15$ minutes) and t .
- Ea is the number of "air-start" aircraft which exited the Entry Fix in the same period. An air-start aircraft entered the system while airborne.
- Eg is the number of "ground-start" aircraft which exited their Entry Fix between $t-0.25$ and t . A ground-start aircraft first made its request for arriving at the airport under congestion management as it was flying toward, or when it was already on the ground at an intermediate airport.
- D is the delay averaged over all aircraft (in min.) which entered the Terminal Area between $t - 15$ minutes and t (that is to say averaged over E aircraft). This

delay is the total delay over the originally requested time; i.e. it is the difference between the Actual Exit Time (AET) and the Original Nominal Exit Time (ONET) from the Entry Fix.

- Da is the averaged delay (AET - ONET) in minutes over all air-start aircraft (Ea) which entered the Terminal Area between $t - 15$ minutes and t .
- Dg is the averaged delay (AET - ONET) in minutes over all ground-start aircraft (Eg) which entered the Terminal Area between $t - 15$ minutes and t .
- AHD is the Air Holding Delay (in min.) averaged over all aircraft which entered the Terminal Area between $t - 15$ minutes and t (that is to say averaged over E aircraft). For each aircraft, the holding delay is the difference between the Actual Exit Time (AET) and the Actual Arrival Time (AAT) at the Entry Fix.
- $AHDa$ is the averaged holding delay (AET - AAT) in minutes over all air-start aircraft (Ea) which entered the Terminal Area between $t - 15$ minutes and t .
- $AHDg$ is the averaged holding delay (AET - AAT) in minutes over all ground-start aircraft (Eg) which entered the Terminal Area between $t - 15$ minutes and t .
- Egd is the number of ground-start aircraft which were issued a *ground delay* at their originating airport, and which *exited* the Entry Fix of the airport under congestion management between $t - 15$ minutes and t .
- $GDgd$ is the averaged Ground Delay (or ground hold) in minutes that those Egd aircraft endured.
- SC is the averaged number of *speed changes* (or speed advisories) that all aircraft which entered the Terminal Area between $t - 15$ minutes and t were issued during their inbound flight.
- T gives an indication of the average time each of the E aircraft spent in the system, air holding not included. It is given in minutes.

- N is the number of aircraft in the system at update time t . It gives us an idea of the size of the problem which must be solved by the Dynamic Resolution Logic which is used.
- $Nh1$ is the number of aircraft in air hold at Entry Fix 1 at update time t .
- $Nh2$ is the number of aircraft in air hold at Entry Fix 2 at update time t .
- Ng is the number of aircraft on the ground awaiting takeoff at update time t .
- Ngd is the number of aircraft with an issued ground delay at time t (we keep track of Ngd only in Scenario 5).
- GHA is the number of Ground Hold Advisories which were issued to the fleet when $Tupdate = t$. Recall that IIDFC is exercised every 15 minutes in all those scenarios.
- CSA is the number of Cruise Speed Advisory which were issued to the fleet at time t .

The last row of the tab "Fleet Sum" gives the sum over time of E , Ea , Eg ; the cumulative values (over time) of D , Da , Dg , AHD , $AHDa$, $AHDg$; the sum of all Egd ; the cumulative value of $GDgd$ (over all Egd aircraft); and the total number of GHA and CSA which were issued during the simulation. Thus, this line is used to give an overall rating on the scenario under consideration..

This tab is followed by several plots:

- "*Traffic Flow Management Advisories vs. Time*" plots show GHA and CSA versus time.
- Plots entitled "*Number of Holding Aircraft*" show the time variation of the number of aircraft in air hold at entry fix 1 ($Nh1$), Entry Fix 2 ($Nh2$) and in ground hold (Ngd) versus time.

- *“Average Delay for Landed Aircraft”* plots show the evolution of D, AHD and GHD versus time. GHD is the averaged Ground Hold Delay for all aircraft which landed between t - 15 minutes and t. Thus, it is given by:

$$\text{GHD} = \frac{\text{Egd} \times \text{GDgd}}{\text{E}}$$

- Plots entitled *“Average Ground Delay of Landed Aircraft which were Ground Held”* show the variation of GDgd versus time.

A.2 Scenario 1 Data and Plots

(See next page)

Figure 6-1: Tab of Statistics vs. Time

t	E	Ea	Eg	D	Da	Dg	AHD	AHDa	AHDg	Egd	GDgd	SC	T	N	Nh1	Nh2	Ng	GHA	CSA
0.25	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	0.0	7	0	0	5	0	0
0.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	0.0	21	0	0	18	0	0
0.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	0.0	28	0	0	25	0	0
1	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	53.3	35	0	0	32	0	0
1.25	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	65.5	43	0	0	38	0	0
1.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	0.0	57	0	0	47	0	0
1.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	0.0	69	0	0	55	0	0
2	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	96.0	82	0	0	67	0	0
2.25	1	0	1	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	107.4	93	0	0	69	0	0
2.5	2	2	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	87.3	97	0	0	72	0	0
2.75	2	1	1	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	113.4	111	0	0	78	0	0
3	7	1	6	0.1	0.0	0.1	0.1	0.0	0.1	0	0.0	0	136.9	111	0	0	75	0	0
3.25	9	2	7	0.5	0.0	0.7	0.5	0.0	0.7	0	0.0	0	132.7	110	0	0	76	0	0
3.5	4	1	3	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	138.8	114	0	0	74	0	0
3.75	8	2	6	0.2	0.4	0.2	0.2	0.4	0.2	0	0.0	0	143.1	118	0	0	71	0	0
4	6	0	6	0.2	0.0	0.2	0.2	0.0	0.2	0	0.0	0	184.8	130	0	0	81	0	0
4.25	11	1	10	0.7	1.6	0.6	0.7	1.6	0.6	0	0.0	0	163.6	128	0	0	78	0	0
4.5	6	1	5	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0	135.9	128	0	0	73	0	0
4.75	14	2	12	1.5	2.3	1.4	1.5	2.3	1.4	0	0.0	0	165.0	129	1	0	70	0	0
5	13	1	12	0.7	0.6	0.7	0.7	0.6	0.7	0	0.0	0	192.2	125	0	1	72	0	0
5.25	6	1	5	1.1	0.0	1.4	1.1	0.0	1.4	0	0.0	0	208.9	124	0	0	58	0	0
5.5	8	1	7	0.5	1.0	0.4	0.5	1.0	0.4	0	0.0	0	158.8	129	0	0	58	0	0
5.75	10	4	6	0.8	1.0	0.6	0.8	1.0	0.6	0	0.0	0	146.8	134	0	0	67	0	0
6	10	2	8	1.0	1.2	0.9	1.0	1.2	0.9	0	0.0	0	151.3	132	5	1	70	0	0
6.25	15	3	12	5.6	5.4	5.7	5.6	5.4	5.7	0	0.0	0	177.4	133	2	2	72	0	0
6.5	15	1	14	5.5	4.2	5.6	5.5	4.2	5.6	0	0.0	0	174.5	131	6	2	77	0	0
6.75	15	0	15	6.5	0.0	6.5	6.5	0.0	6.5	0	0.0	0	207.8	130	0	2	81	0	0
7	12	4	8	1.2	1.2	1.2	1.2	1.2	1.2	0	0.0	0	146.9	129	0	0	82	0	0
7.25	8	1	7	5.0	1.7	5.4	5.0	1.7	5.4	0	0.0	0	175.8	134	5	4	85	0	0
7.5	7	2	5	12.4	13.0	12.2	12.4	13.0	12.2	0	0.0	0	154.0	133	5	1	82	0	0
7.75	8	2	6	13.3	18.2	11.7	13.3	18.2	11.7	0	0.0	0	167.1	133	4	1	77	0	0
8	7	0	7	13.3	0.0	13.3	13.3	0.0	13.3	0	0.0	0	191.5	133	2	5	74	0	0
8.25	8	0	8	14.7	0.0	14.7	14.7	0.0	14.7	0	0.0	0	198.6	137	4	5	75	0	0
8.5	7	0	7	14.9	0.0	14.9	14.9	0.0	14.9	0	0.0	0	202.6	137	6	7	70	0	0
8.75	8	0	8	18.9	0.0	18.9	18.9	0.0	18.9	0	0.0	0	189.5	137	10	5	68	0	0
9	7	0	7	23.8	0.0	23.8	23.8	0.0	23.8	0	0.0	0	186.5	139	12	7	68	0	0
9.25	8	0	8	30.1	0.0	30.1	30.1	0.0	30.1	0	0.0	0	197.8	144	8	12	67	0	0
9.5	7	1	6	32.9	34.6	32.6	32.9	34.6	32.6	0	0.0	0	153.6	142	12	18	59	0	0
9.75	8	1	7	37.7	41.1	37.2	37.7	41.1	37.2	0	0.0	0	142.7	146	15	17	62	0	0
10	7	0	7	42.1	0.0	42.1	42.1	0.0	42.1	0	0.0	0	173.0	147	16	17	60	0	0
10.25	15	0	15	46.1	0.0	46.1	46.1	0.0	46.1	0	0.0	0	178.3	139	16	17	61	0	0
10.5	15	2	13	42.8	45.4	42.4	42.8	45.4	42.4	0	0.0	0	176.4	133	14	13	66	0	0
10.75	15	1	14	32.5	34.9	32.3	32.5	34.9	32.3	0	0.0	0	184.4	126	13	10	57	0	0
11	15	0	15	29.8	0.0	29.8	29.8	0.0	29.8	0	0.0	0	191.1	118	9	7	55	0	0
11.25	15	2	13	24.2	23.6	24.3	24.2	23.6	24.3	0	0.0	0	176.9	112	4	4	55	0	0
11.5	15	1	14	12.3	6.4	12.7	12.3	6.4	12.7	0	0.0	0	182.7	103	3	2	57	0	0
11.75	12	0	12	2.6	0.0	2.6	2.6	0.0	2.6	0	0.0	0	169.2	99	0	1	59	0	0
12	7	0	7	0.6	0.0	0.6	0.6	0.0	0.6	0	0.0	0	188.8	104	0	0	60	0	0
12.25	11	0	11	1.3	0.0	1.3	1.3	0.0	1.3	0	0.0	0	203.5	101	2	1	57	0	0
12.5	15	2	13	3.5	2.4	3.6	3.5	2.4	3.6	0	0.0	0	149.8	91	0	0	53	0	0
12.75	5	0	5	0.4	0.0	0.4	0.4	0.0	0.4	0	0.0	0	193.5	97	0	0	54	0	0
13	5	2	3	0.2	0.0	0.3	0.2	0.0	0.3	0	0.0	0	148.8	101	1	2	56	0	0
13.25	11	0	11	1.6	0.0	1.6	1.6	0.0	1.6	0	0.0	0	170.7	97	0	0	55	0	0
13.5	9	3	6	0.1	0.2	0.1	0.1	0.2	0.1	0	0.0	0	162.2	96	0	0	55	0	0
13.75	9	1	8	0.3	0.0	0.3	0.3	0.0	0.3	0	0.0	0	160.3	94	1	0	52	0	0
14	11	1	10	0.4	0.3	0.4	0.4	0.3	0.4	0	0.0	0	192.7	89	0	0	52	0	0
14.25	9	1	8	0.2	0.0	0.2	0.2	0.0	0.2	0	0.0	0	160.6	90	0	0	49	0	0
14.5	4	1	3	0.1	0.6	0.0	0.1	0.6	0.0	0	0.0	0	130.0	92	0	0	45	0	0
14.75	9	0	9	0.4	0.0	0.4	0.4	0.0	0.4	0	0.0	0	180.1	92	0	0	49	0	0
15	6	0	6	0.2	0.0	0.2	0.2	0.0	0.2	0	0.0	0	176.1	95	0	0	49	0	0
Fleet Sum	480	57	423	5233.2	365.8	4867.3	5233.2	365.8	4867.3	0	0.0	0						0	0

- Scenario 1 -

Figure 6-2: Number of Holding Aircraft





Figure 6-3: Average Delay for Landed Aircraft



- Scenario 1 -

A.3 Scenario 2 Data and Plots

(See next page)

Figure 6-4: Tab of Statistics vs. Time

t	E	Ea	Eg	D	Da	Dg	AHD	AHDa	AHDg	Egd	GDgd	SC	T	N	Nh1	Nh2	Ng	GHA	CSA
0.25	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	7	0	0	5	0	1
0.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	21	0	0	18	0	1
0.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	28	0	0	25	0	0
1	1	1	0	-0.9	-0.9	0.0	0.0	0.0	0.0	0	0.0	1.0	52.5	35	0	0	32	0	1
1.25	1	1	0	-0.8	-0.8	0.0	0.0	0.0	0.0	0	0.0	1.0	64.7	43	0	0	38	0	1
1.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	57	0	0	47	0	1
1.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	69	0	0	55	0	3
2	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	96.0	82	0	0	67	0	8
2.25	1	0	1	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	107.4	93	0	0	69	0	9
2.5	2	2	0	-0.1	-0.1	0.0	0.0	0.0	0.0	0	0.0	1.0	87.1	97	0	0	72	0	8
2.75	2	1	1	-0.7	-0.6	-0.8	0.0	0.0	0.0	0	0.0	1.0	112.7	111	0	0	78	0	9
3	7	1	6	-0.4	-0.8	-0.3	0.0	0.0	0.1	0	0.0	1.0	136.5	111	0	0	75	0	12
3.25	9	2	7	-0.5	-0.3	-0.6	0.1	0.2	0.1	0	0.0	1.8	132.1	110	0	0	76	0	10
3.5	5	1	4	-0.6	-0.8	-0.6	0.0	0.0	0.0	0	0.0	0.8	142.7	113	0	0	74	0	20
3.75	7	2	5	-0.7	-0.5	-0.8	0.0	0.1	0.0	0	0.0	2.0	139.7	118	0	0	71	0	24
4	7	0	7	-0.9	0.0	-0.9	0.0	0.0	0.0	0	0.0	1.6	183.4	129	0	0	81	0	19
4.25	10	1	9	-0.8	-1.9	-0.7	0.1	0.0	0.1	0	0.0	1.2	160.9	128	0	0	78	0	22
4.5	8	1	7	-1.4	-0.8	-1.5	0.1	0.0	0.1	0	0.0	1.4	148.4	126	0	0	73	0	14
4.75	15	2	13	-0.7	-1.2	-0.7	0.2	0.3	0.2	0	0.0	1.6	162.2	126	0	0	70	0	19
5	12	1	11	-0.8	-0.4	-0.8	0.2	0.1	0.2	0	0.0	2.8	205.2	123	0	0	72	0	22
5.25	4	1	3	0.0	-0.9	0.3	0.1	0.0	0.1	0	0.0	1.0	189.4	124	0	0	58	0	33
5.5	8	1	7	-0.6	0.0	-0.6	0.1	0.0	0.1	0	0.0	1.9	158.2	129	0	0	58	0	36
5.75	11	4	7	-1.1	-0.3	-1.5	0.1	0.3	0.0	0	0.0	1.6	145.8	133	0	0	67	2	26
6	14	2	12	-1.6	-0.6	-1.8	0.3	0.2	0.3	0	0.0	2.1	173.0	127	0	0	70	9	18
6.25	15	4	11	0.9	0.6	1.0	0.2	0.2	0.3	0	0.0	1.8	162.1	128	0	0	72	20	19
6.5	14	0	14	2.1	0.0	2.1	0.2	0.0	0.2	0	0.0	2.7	193.7	127	0	0	77	28	17
6.75	13	2	11	2.0	-0.2	2.4	0.2	0.1	0.2	0	0.0	2.0	184.7	128	0	0	81	51	18
7	15	3	12	-1.9	-1.4	-2.0	0.2	0.2	0.2	0	0.0	1.9	167.5	124	0	0	83	49	25
7.25	7	0	7	0.7	0.0	0.7	0.5	0.0	0.5	0	0.0	1.3	176.1	130	1	0	88	60	19
7.5	8	4	4	5.1	6.4	3.8	1.1	1.7	0.6	0	0.0	1.9	137.8	128	0	0	89	47	13
7.75	7	0	7	4.0	0.0	4.0	0.6	0.0	0.6	0	0.0	2.3	191.6	129	1	0	89	42	18
8	7	0	7	5.1	0.0	5.1	1.0	0.0	1.0	0	0.0	2.0	203.5	129	1	0	86	53	24
8.25	8	0	8	5.4	0.0	5.4	0.4	0.0	0.4	0	0.0	2.9	208.7	133	0	0	91	8	16
8.5	7	0	7	8.0	0.0	8.0	0.3	0.0	0.3	2	3.8	2.1	202.2	133	0	0	90	43	15
8.75	8	0	8	12.2	0.0	12.2	0.4	0.0	0.4	7	9.6	2.0	201.9	133	1	0	91	28	11
9	7	0	7	18.9	0.0	18.9	2.5	0.0	2.5	7	12.6	2.3	202.3	135	2	0	86	18	18
9.25	8	2	6	15.4	7.8	17.9	3.2	1.7	3.7	6	7.8	2.4	155.6	140	0	2	82	17	32
9.5	7	0	7	26.7	0.0	26.7	3.3	0.0	3.3	7	17.4	1.9	204.8	138	1	2	69	11	45
9.75	8	1	7	13.8	12.2	14.0	3.1	1.6	3.3	7	4.0	2.0	180.3	142	1	4	72	4	29
10	7	1	6	32.7	12.6	36.1	6.8	8.4	6.5	6	22.5	1.3	192.8	143	1	1	71	0	19
10.25	14	1	13	35.9	5.8	38.2	0.6	1.2	0.5	12	36.3	1.6	198.9	136	1	0	65	0	28
10.5	15	0	15	23.8	0.0	23.8	1.1	0.0	1.1	12	22.7	2.7	219.3	130	2	0	68	0	17
10.75	15	0	15	43.7	0.0	43.7	1.6	0.0	1.6	15	35.5	2.4	228.3	123	1	1	57	0	24
11	15	2	13	26.0	8.8	28.6	0.5	0.0	0.6	10	29.9	2.5	214.0	115	0	2	55	0	24
11.25	14	0	14	37.4	0.0	37.4	0.5	0.0	0.5	10	44.7	2.4	222.6	110	0	0	55	0	17
11.5	15	1	14	13.2	6.2	13.7	0.5	0.0	0.5	6	21.1	1.9	178.9	101	0	1	57	0	19
11.75	11	0	11	0.1	0.0	0.1	0.3	0.0	0.3	0	0.0	2.0	172.5	98	0	0	59	0	10
12	6	0	6	-0.2	0.0	-0.2	0.0	0.0	0.0	0	0.0	1.0	187.3	104	0	1	60	0	10
12.25	15	1	14	-2.1	-2.6	-2.1	0.3	0.8	0.3	0	0.0	2.3	183.7	97	0	0	57	0	6
12.5	11	1	10	0.3	0.1	0.3	0.2	0.7	0.2	0	0.0	2.1	154.0	91	0	0	53	0	14
12.75	5	0	5	-0.8	0.0	-0.8	0.0	0.0	0.0	0	0.0	1.4	192.6	97	0	0	54	0	12
13	6	2	4	-0.7	-0.7	-0.8	0.0	0.0	0.0	0	0.0	1.2	152.2	100	1	1	56	0	14
13.25	10	0	10	0.5	0.0	0.5	0.4	0.0	0.4	0	0.0	1.2	170.5	97	0	0	55	0	9
13.5	10	3	7	-0.6	-0.3	-0.7	0.2	0.2	0.2	0	0.0	1.3	167.4	95	0	1	55	0	10
13.75	9	1	8	-0.5	-0.3	-0.5	0.1	0.0	0.2	0	0.0	0.8	148.1	93	0	0	52	0	7
14	10	1	9	-0.2	0.3	-0.2	0.2	0.6	0.1	0	0.0	1.5	199.8	89	0	0	52	0	13
14.25	9	1	8	-0.4	-0.9	-0.4	0.1	0.0	0.1	0	0.0	1.1	160.1	90	0	0	49	0	12
14.5	5	1	4	-0.5	0.0	-0.6	0.0	0.0	0.0	0	0.0	1.4	123.8	91	0	0	45	0	16
14.75	8	0	8	-0.3	0.0	-0.3	0.2	0.0	0.2	0	0.0	1.3	189.5	92	0	0	49	0	16
15	6	0	6	-0.6	0.0	-0.6	0.0	0.0	0.0	0	0.0	1.2	175.5	95	0	0	49	0	6
Fleet Sum	480	57	423	3594.0	73.1	3521.0	287.5	27.9	259.3	107	2606.1	884.4					490	939	

- Scenario 2 -

Figure 6-5: Traffic Flow Management Advisories vs. Time





Figure 6-6: Number of Holding Aircraft



- Scenario 2 -

Figure 6-7: Average Delay for Landed Aircraft





Figure 6-8: Average Ground Delay per Landed Aircraft which were Ground Held



- Scenario 2 -

A.4 Scenario 3 Data and Plots

(See next page)

Figure 6-9: Tab of Statistics vs. Time

t	E	Ea	Eg	D	Da	Dg	AHD	AHDa	AHDg	Egd	GDgd	SC	T	N	Nh1	Nh2	Ng	GHA	CSA	
0.25	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	7	0	0	5	0	0	
0.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	21	0	0	18	0	0	
0.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	28	0	0	25	0	0	
1	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	53.3	35	0	0	32	0	0
1.25	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	65.5	43	0	0	38	0	0	
1.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	57	0	0	47	0	0	
1.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	69	0	0	55	0	1	
2	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	96.0	82	0	0	67	0	3	
2.25	1	0	1	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	107.4	93	0	0	69	0	6	
2.5	2	2	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	87.3	97	0	0	72	0	8	
2.75	2	1	1	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	113.4	111	0	0	78	0	3	
3	7	1	6	0.3	0.0	0.3	0.0	0.0	0.0	0	0.0	1.0	137.2	111	0	0	75	0	7	
3.25	8	2	6	0.4	-0.3	0.6	0.0	0.1	0.0	0	0.0	1.0	126.9	111	0	0	76	0	8	
3.5	5	1	4	0.2	0.0	0.2	0.0	0.0	0.0	0	0.0	0.2	147.6	114	0	0	74	0	8	
3.75	8	2	6	0.4	0.6	0.3	0.1	0.2	0.0	0	0.0	0.9	143.4	118	0	0	71	0	14	
4	6	0	6	-0.1	0.0	-0.1	0.0	0.0	0.0	0	0.0	0.5	184.7	130	0	0	81	0	13	
4.25	11	1	10	0.7	1.4	0.6	0.1	0.1	0.1	0	0.0	1.5	164.2	128	0	0	78	0	12	
4.5	7	1	6	-0.6	0.0	-0.7	0.1	0.0	0.1	0	0.0	0.4	147.2	127	0	0	73	0	6	
4.75	15	2	13	0.3	0.8	0.2	0.2	0.1	0.2	0	0.0	1.3	164.0	127	0	0	70	0	15	
5	12	1	11	0.3	0.8	0.3	0.1	0.1	0.1	0	0.0	0.8	195.7	124	0	1	72	0	22	
5.25	5	1	4	0.6	0.0	0.8	0.1	0.0	0.2	0	0.0	1.8	208.0	124	0	0	58	0	32	
5.5	8	1	7	0.1	0.7	0.1	0.1	0.4	0.0	0	0.0	1.4	158.9	129	0	0	58	0	37	
5.75	10	4	6	0.2	0.4	0.1	0.0	0.1	0.0	0	0.0	0.9	147.0	134	0	0	67	1	26	
6	13	2	11	-0.6	-0.1	-0.8	0.2	0.1	0.3	0	0.0	1.3	166.6	129	1	0	70	12	17	
6.25	15	3	12	2.7	2.7	2.7	0.3	0.3	0.3	0	0.0	2.0	175.9	130	0	0	72	19	22	
6.5	15	1	14	3.4	0.6	3.6	0.2	0.3	0.2	0	0.0	2.4	180.4	128	0	0	77	38	20	
6.75	14	2	12	3.4	-0.4	4.0	0.1	0.1	0.1	0	0.0	2.6	192.5	128	0	0	81	37	16	
7	15	3	12	-2.3	-2.3	-2.3	0.1	0.1	0.1	0	0.0	1.5	161.3	124	0	0	83	44	20	
7.25	7	0	7	0.5	0.0	0.5	0.2	0.0	0.2	0	0.0	1.9	188.7	130	0	1	88	62	19	
7.5	8	4	4	4.8	6.0	3.6	0.8	1.4	0.2	0	0.0	1.8	137.9	128	0	0	88	48	21	
7.75	7	0	7	3.7	0.0	3.7	0.3	0.0	0.3	0	0.0	1.6	191.7	129	1	0	91	14	9	
8	7	0	7	4.9	0.0	4.9	0.8	0.0	0.8	0	0.0	2.1	203.5	129	1	0	91	56	15	
8.25	8	0	8	5.1	0.0	5.1	0.4	0.0	0.4	1	2.7	3.1	208.5	133	0	0	95	19	16	
8.5	7	0	7	7.7	0.0	7.7	0.7	0.0	0.7	0	0.0	2.4	201.6	133	1	0	93	43	14	
8.75	8	0	8	8.4	0.0	8.4	0.4	0.0	0.4	6	4.7	1.8	202.2	133	0	0	96	28	16	
9	7	1	6	11.8	6.4	12.7	0.8	0.5	0.9	6	7.2	1.9	175.8	135	0	0	86	14	27	
9.25	7	1	6	13.3	3.1	15.0	0.1	0.2	0.1	5	11.6	2.4	162.9	141	0	0	84	11	32	
9.5	8	0	8	28.6	0.0	28.6	0.2	0.0	0.2	7	27.8	2.4	215.1	138	0	0	66	3	44	
9.75	7	1	6	29.4	13.1	32.1	0.7	2.1	0.5	6	26.5	2.1	208.2	143	1	2	68	0	30	
10	8	1	7	34.2	7.5	38.0	5.3	3.3	5.6	6	32.0	1.5	181.0	143	2	0	64	2	16	
10.25	14	1	13	33.9	6.7	36.0	2.3	2.1	2.3	13	28.3	1.9	194.4	136	1	0	63	0	14	
10.5	14	0	14	41.8	0.0	41.8	1.9	0.0	1.9	13	36.9	1.6	220.3	131	2	2	67	0	15	
10.75	15	0	15	33.7	0.0	33.7	2.5	0.0	2.5	14	27.3	1.8	210.9	124	4	1	57	0	12	
11	15	2	13	32.3	12.2	35.4	3.4	4.1	3.3	12	27.0	2.5	221.9	116	1	2	55	0	10	
11.25	15	0	15	24.9	0.0	24.9	1.4	0.0	1.4	11	22.3	2.5	227.3	110	1	1	55	0	10	
11.5	15	1	14	11.0	5.1	11.4	0.6	0.0	0.6	5	18.6	1.8	176.0	101	0	0	57	0	14	
11.75	11	0	11	1.5	0.0	1.5	0.2	0.0	0.2	1	8.4	1.4	179.3	98	0	0	59	0	8	
12	6	0	6	0.2	0.0	0.2	0.1	0.0	0.1	0	0.0	0.2	187.5	104	0	0	60	0	5	
12.25	13	0	13	-0.4	0.0	-0.4	0.2	0.0	0.2	0	0.0	1.2	199.3	99	0	0	57	0	3	
12.5	13	2	11	1.9	1.2	2.1	0.2	0.3	0.2	0	0.0	1.5	146.9	91	0	0	53	0	9	
12.75	5	0	5	0.4	0.0	0.4	0.1	0.0	0.1	0	0.0	0.8	193.8	97	0	0	54	0	9	
13	5	2	3	0.0	-0.1	0.2	0.1	0.0	0.2	0	0.0	0.4	148.8	101	1	2	56	0	8	
13.25	11	0	11	1.6	0.0	1.6	0.6	0.0	0.6	0	0.0	0.6	171.7	97	0	0	55	0	4	
13.5	9	3	6	0.4	0.9	0.2	0.0	0.1	0.0	0	0.0	0.9	162.5	96	0	0	55	0	3	
13.75	9	1	8	0.6	0.0	0.7	0.1	0.0	0.1	0	0.0	0.6	156.7	94	0	1	52	0	4	
14	11	1	10	0.5	0.7	0.4	0.3	0.0	0.3	0	0.0	1.1	196.2	89	0	0	52	0	8	
14.25	9	1	8	0.5	0.5	0.5	0.1	0.5	0.1	0	0.0	0.3	161.0	90	0	0	49	0	7	
14.5	4	1	3	-0.1	0.0	-0.2	0.0	0.0	0.0	0	0.0	0.3	129.9	92	0	0	45	0	10	
14.75	9	0	9	0.3	0.0	0.3	0.1	0.0	0.1	0	0.0	0.4	180.3	92	0	0	49	0	7	
15	6	0	6	0.5	0.0	0.5	0.1	0.0	0.1	0	0.0	1.0	176.6	95	0	0	49	0	9	
Fleet Sum	480	57	423	3941.9	103.7	3838.3	301.4	26.5	275.2	106	2578.2	710.1					451	744		

- Scenario 3 -

Figure 6-10: Traffic Flow Management Advisories vs. Time





Figure 6-11: Number of Holding Aircraft



- Scenario 3 -

Figure 6-12: Average Delay for Landed Aircraft

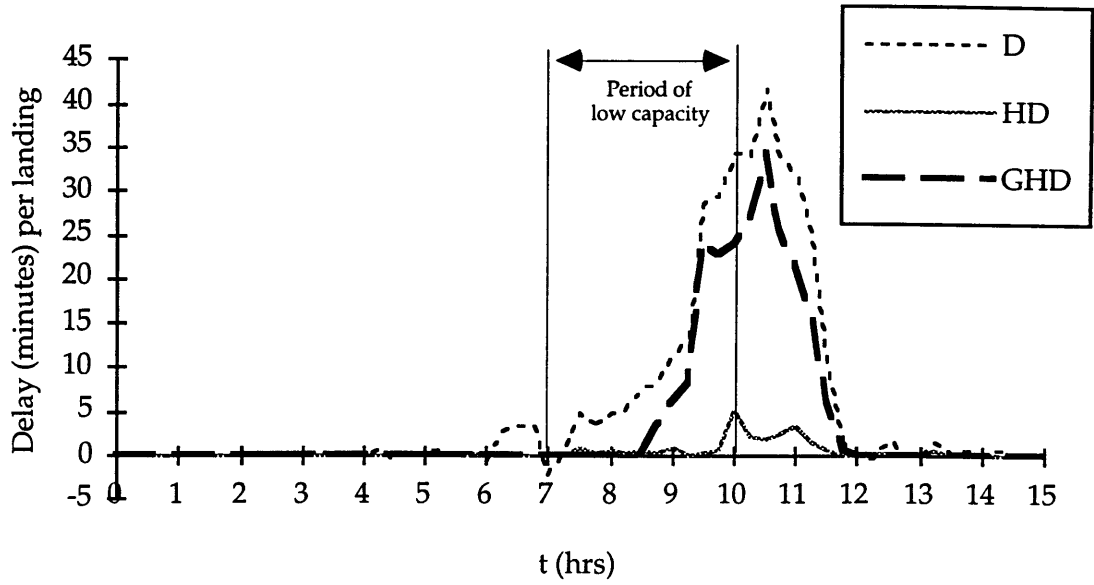



Figure 6-13: Average Ground Delay per Landed Aircraft which were Ground Held



- Scenario 3 -

A.5 Scenario 4 Data and Plots

(See next page)

Figure 6-14: Tab of Statistics vs. Time

t	E	Ea	Eg	D	Da	Dg	AHD	AHDa	AHDg	Egd	GDgd	SC	T	N	Nh1	Nh2	Ng	GHA	CSA
0.25	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	0.0	7	0	0	5	0	0
0.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	0.0	21	0	0	18	0	0
0.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	0.0	28	0	0	25	0	0
1	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	53.3	35	0	0	32	0	0
1.25	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	65.5	43	0	0	38	0	0
1.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	0.0	57	0	0	47	0	0
1.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	0.0	69	0	0	55	0	1
2	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	96.0	82	0	0	67	0	3
2.25	1	0	1	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	107.4	93	0	0	69	0	6
2.5	2	2	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	87.3	97	0	0	72	0	8
2.75	2	1	1	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0.0	113.4	111	0	0	78	0	3
3	7	1	6	0.3	0.0	0.3	0.0	0.0	0.0	0	0	1.0	137.2	111	0	0	75	0	7
3.25	8	2	6	0.4	-0.3	0.6	0.0	0.1	0.0	0	0	1.0	126.9	111	0	0	76	0	8
3.5	5	1	4	0.2	0.0	0.2	0.0	0.0	0.0	0	0	0.2	147.6	114	0	0	74	0	6
3.75	8	2	6	0.4	0.6	0.3	0.1	0.2	0.0	0	0	0.9	143.4	118	0	0	71	0	13
4	6	0	6	-0.1	0.0	-0.1	0.0	0.0	0.0	0	0	0.5	184.7	150	0	0	81	0	11
4.25	11	1	10	0.8	0.6	0.8	0.1	0.0	0.1	0	0	1.4	164.2	128	0	0	78	0	11
4.5	7	1	6	-0.6	0.0	-0.7	0.1	0.0	0.1	0	0	0.4	147.2	127	0	0	73	0	7
4.75	15	2	13	0.3	1.4	0.1	0.2	0.2	0.2	0	0	1.1	164.1	127	0	0	70	0	15
5	12	1	11	0.3	0.8	0.3	0.1	0.1	0.1	0	0	0.7	195.7	124	0	0	72	0	17
5.25	5	1	4	0.8	0.0	1.0	0.1	0.0	0.1	0	0	2.0	208.2	124	0	0	58	0	22
5.5	8	1	7	0.3	0.8	0.2	0.0	0.0	0.0	0	0	1.3	159.0	129	0	0	58	0	24
5.75	10	4	6	0.2	0.4	0.1	0.1	0.2	0.0	0	0	0.9	147.0	134	0	0	67	1	19
6	13	2	11	-0.7	-0.1	-0.8	0.2	0.1	0.3	0	0	1.2	166.5	129	0	1	70	3	17
6.25	15	3	12	2.7	3.4	2.5	0.2	0.3	0.2	0	0	1.3	170.7	130	0	1	72	2	21
6.5	14	1	13	3.3	1.7	3.5	0.2	0.2	0.3	0	0	1.6	183.8	129	1	0	77	4	19
6.75	14	1	13	4.1	0.0	4.4	0.2	0.0	0.2	0	0	1.9	201.5	129	0	0	81	3	11
7	15	4	11	-1.0	-0.3	-1.3	0.1	0.2	0.1	0	0	1.3	160.0	125	0	0	82	2	21
7.25	7	0	7	1.4	0.0	1.4	0.2	0.0	0.2	0	0	1.9	185.1	131	0	0	85	4	21
7.5	8	4	4	6.8	8.1	5.5	1.7	2.7	0.8	0	0	1.6	137.5	129	0	1	83	2	15
7.75	7	0	7	5.4	0.0	5.4	1.0	0.0	1.0	0	0	1.9	185.8	130	2	0	81	2	14
8	8	0	8	6.5	0.0	6.5	1.8	0.0	1.8	0	0	1.6	206.3	129	0	0	79	2	18
8.25	7	0	7	6.3	0.0	6.3	0.5	0.0	0.5	0	0	2.3	217.6	134	0	1	82	4	9
8.5	8	0	8	10.2	0.0	10.2	3.5	0.0	3.5	0	0	1.8	189.6	133	1	1	81	6	9
8.75	7	0	7	12.1	0.0	12.1	4.8	0.0	4.8	0	0	1.7	206.5	134	4	2	80	2	9
9	8	0	8	16.4	0.0	16.4	9.3	0.0	9.3	1	3.0	1.4	202.0	135	2	2	83	0	5
9.25	7	1	6	17.7	15.7	18.0	8.0	8.6	7.9	1	21.7	1.4	158.3	141	3	6	82	0	20
9.5	8	1	7	24.9	18.1	25.9	12.7	11.7	12.9	2	20.3	1.4	159.7	138	5	6	72	0	18
9.75	7	0	7	23.8	0.0	23.8	16.4	0.0	16.4	0	0	1.4	184.1	143	5	8	73	0	17
10	8	0	8	29.7	0.0	29.7	21.5	0.0	21.5	0	0	1.4	189.1	143	3	6	67	0	12
10.25	14	3	11	29.6	24.6	30.9	13.8	17.5	12.8	3	41.4	1.1	178.8	136	3	2	64	0	12
10.5	15	0	15	25.0	0.0	25.0	4.9	0.0	4.9	5	39.2	1.3	210.5	130	5	3	66	0	7
10.75	15	0	15	32.3	0.0	32.3	7.8	0.0	7.8	5	57.0	1.3	211.4	123	4	3	57	0	9
11	15	1	14	28.4	18.7	29.1	6.6	7.2	6.6	3	78.2	1.4	196.7	115	7	0	55	0	13
11.25	15	1	14	49.5	17.7	51.8	6.7	7.9	6.7	6	91.8	1.3	236.9	109	2	2	55	0	13
11.5	14	1	13	13.6	3.4	14.4	1.6	0.3	1.7	1	89.5	1.3	184.9	101	1	0	57	0	17
11.75	10	0	10	0.9	0.0	0.9	0.1	0.0	0.1	0	0	1.0	171.0	99	0	1	59	0	7
12	7	0	7	0.7	0.0	0.7	0.1	0.0	0.1	0	0	0.9	189.4	104	0	0	60	0	4
12.25	13	0	13	-0.2	0.0	-0.2	0.4	0.0	0.4	0	0	1.5	199.3	99	0	0	57	0	4
12.5	13	2	11	1.9	2.2	1.9	0.1	0.2	0.1	0	0	1.5	146.9	91	0	0	53	0	10
12.75	5	0	5	0.4	0.0	0.4	0.0	0.0	0.0	0	0	0.6	193.8	97	0	0	54	0	7
13	5	2	3	-0.1	-0.4	0.1	0.0	0.0	0.0	0	0	0.6	148.7	101	1	2	56	0	9
13.25	11	0	11	1.4	0.0	1.4	0.6	0.0	0.6	0	0	0.6	171.5	97	0	0	55	0	3
13.5	10	3	7	0.1	0.6	-0.1	0.0	0.1	0.0	0	0	0.9	160.3	95	0	0	55	0	3
13.75	8	1	7	0.7	0.0	0.8	0.2	0.0	0.2	0	0	0.8	158.3	94	0	1	52	0	4
14	11	1	10	0.4	0.7	0.4	0.2	0.0	0.3	0	0	0.9	196.2	89	0	0	52	0	5
14.25	9	1	8	0.5	0.5	0.5	0.1	0.5	0.1	0	0	0.3	161.0	90	0	0	49	0	7
14.5	4	1	3	-0.1	0.0	-0.2	0.0	0.0	0.0	0	0	0.3	129.9	92	0	0	45	0	6
14.75	9	0	9	0.3	0.0	0.3	0.1	0.0	0.1	0	0	0.4	180.3	92	0	0	49	0	10
15	6	0	6	0.5	0.0	0.5	0.1	0.0	0.1	0	0	0.7	176.6	95	0	0	49	0	8
Fleet Sum	480	57	423	4073.5	203.7	3869.8	1270.3	103.6	1166.7	27	1545.9	564.3					37	595	

- Scenario 4 -

Figure 6-15: Traffic Flow Management Advisories vs Time





Figure 6-16: Number of Holding Aircraft



- Scenario 4 -

Figure 6-17: Average Delay for Landed Aircraft

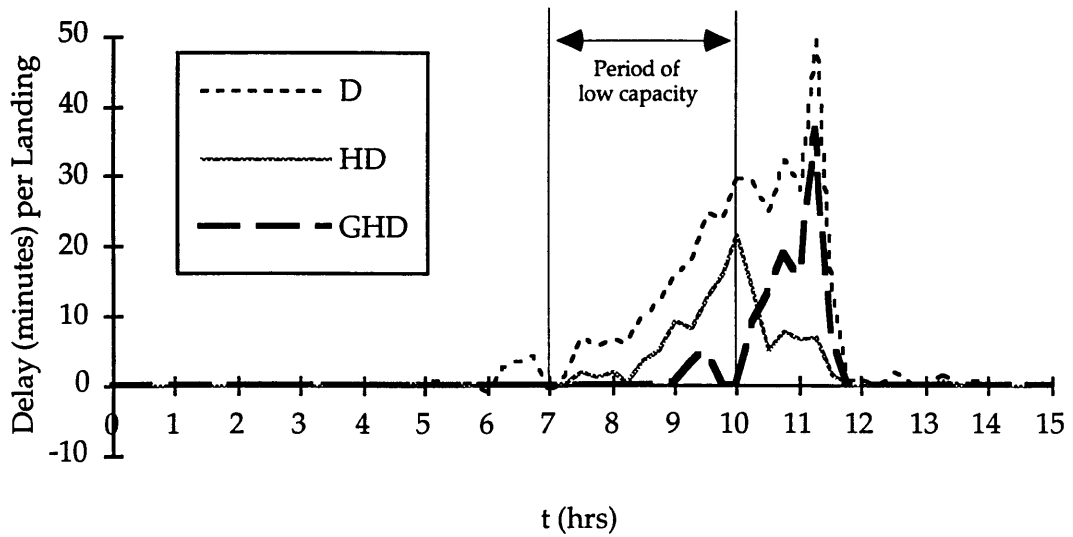



Figure 6-18: Average Ground Delay per Landed Aircraft which were Ground Held



- Scenario 4 -

A.6 Scenario 5 Data and Plots

(See next page)

Figure 6-19: Tab of Statistics vs. Time

t	E	Ea	Eg	D	Da	Dg	AHD	AHDa	AHDg	Egd	GDgd	SC	T	N	Nh1	Nh2	Ng	Ngd	GHA	CSA
0.25	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	7	0	0	5	0	0	0
0.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	21	0	0	18	0	0	1
0.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	28	0	0	25	0	0	1
1	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	53.3	35	0	0	32	0	0	1
1.25	1	1	0	1.1	1.1	0.0	0.0	0.0	0.0	0	0.0	2.0	66.5	43	0	0	38	0	0	0
1.5	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	57	0	0	47	0	0	1
1.75	0	0	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	0.0	69	0	0	55	0	0	3
2	1	1	0	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	96.0	82	0	0	67	0	0	6
2.25	1	0	1	0.0	0.0	0.0	0.0	0.0	0.0	0	0.0	0.0	107.4	93	0	0	69	0	0	7
2.5	2	2	0	0.7	0.7	0.0	0.0	0.0	0.0	0	0.0	1.5	87.9	97	0	0	72	0	0	9
2.75	2	1	1	0.8	1.6	0.0	0.0	0.0	0.0	0	0.0	1.0	114.2	111	0	0	78	0	0	11
3	6	1	5	1.4	3.0	1.1	0.1	0.0	0.1	0	0.0	1.2	134.8	112	0	0	75	0	0	8
3.25	9	2	7	1.9	2.1	1.8	0.2	0.4	0.1	0	0.0	1.2	131.7	111	0	0	76	0	0	8
3.5	5	1	4	1.4	1.0	1.5	0.2	0.1	0.2	0	0.0	1.2	148.6	114	0	0	74	0	0	7
3.75	8	2	6	1.0	1.5	0.8	0.1	0.1	0.1	0	0.0	1.0	144.0	118	0	0	71	0	0	13
4	6	0	6	0.4	0.0	0.4	0.1	0.0	0.1	0	0.0	0.7	185.2	130	0	0	81	0	0	15
4.25	11	1	10	1.5	3.5	1.3	0.1	0.3	0.1	0	0.0	1.4	165.0	128	0	0	78	0	0	11
4.5	6	1	5	0.7	1.7	0.6	0.0	0.0	0.1	0	0.0	0.8	136.6	128	0	0	73	0	0	11
4.75	14	2	12	1.6	2.3	1.4	0.1	0.1	0.1	0	0.0	1.4	166.4	129	0	0	70	0	0	12
5	13	1	12	1.2	2.4	1.1	0.2	0.2	0.2	0	0.0	0.8	193.3	125	0	0	72	0	0	15
5.25	6	1	5	1.7	1.2	1.8	0.1	0.0	0.1	0	0.0	1.8	210.5	124	0	0	58	0	0	28
5.5	8	1	7	0.8	0.0	0.9	0.0	0.0	0.0	0	0.0	1.0	139.5	129	0	0	58	0	0	29
5.75	10	4	6	0.9	1.0	0.8	0.1	0.2	0.1	0	0.0	0.7	147.6	134	0	0	67	1	1	22
6	13	2	11	0.0	1.2	-0.2	0.2	0.0	0.3	0	0.0	0.9	168.5	129	0	0	70	12	11	16
6.25	14	3	11	3.4	3.7	3.3	0.2	0.2	0.2	0	0.0	1.7	169.7	131	1	0	72	26	16	23
6.5	15	1	14	3.9	2.3	4.0	0.3	0.5	0.3	0	0.0	1.9	181.8	129	0	0	77	48	32	20
6.75	14	1	13	4.4	-0.2	4.8	0.3	0.0	0.3	0	0.0	2.2	204.1	129	0	1	81	58	19	13
7	15	4	11	-1.3	-0.8	-1.5	0.2	0.1	0.2	0	0.0	1.7	161.2	125	0	0	83	68	13	17
7.25	7	0	7	1.3	0.0	1.3	0.2	0.0	0.2	0	0.0	2.0	181.8	131	0	0	94	86	38	7
7.5	8	4	4	6.8	8.1	5.5	1.5	2.1	0.9	0	0.0	1.8	137.8	129	0	1	91	80	11	13
7.75	7	0	7	5.5	0.0	5.5	1.2	0.0	1.2	0	0.0	1.7	185.7	130	2	0	91	77	8	17
8	8	0	8	6.5	0.0	6.5	1.5	0.0	1.5	0	0.0	1.9	206.6	129	0	0	86	73	17	15
8.25	7	0	7	6.5	0.0	6.5	0.1	0.0	0.1	1	6.4	1.7	203.1	134	1	0	93	76	13	12
8.5	7	0	7	9.9	0.0	9.9	0.4	0.0	0.4	1	8.1	1.9	216.6	134	0	0	94	73	10	12
8.75	8	0	8	14.1	0.0	14.1	0.4	0.0	0.4	7	10.7	1.3	200.0	134	0	0	93	65	3	13
9	7	0	7	20.2	0.0	20.2	1.2	0.0	1.2	7	13.5	2.4	204.5	136	0	0	89	55	2	14
9.25	8	2	6	22.2	8.5	26.8	1.2	1.3	1.2	6	18.9	1.3	182.4	141	0	0	86	44	5	22
9.5	7	0	7	30.5	0.0	30.5	0.1	0.0	0.1	7	25.1	1.9	202.6	139	0	0	76	31	0	26
9.75	7	1	6	29.3	4.9	33.6	0.3	0.6	0.3	6	27.4	1.7	192.4	144	2	0	74	18	0	18
10	8	1	7	37.9	16.4	41.0	2.1	3.2	1.9	7	32.0	1.5	199.1	144	0	0	65	7	0	24
10.25	14	1	13	43.8	5.3	46.8	1.1	0.7	1.1	13	39.5	1.9	210.2	137	1	0	63	3	0	20
10.5	15	0	15	40.8	0.0	40.8	0.8	0.0	0.8	15	34.6	1.2	223.9	131	1	1	66	0	0	12
10.75	15	0	15	30.2	0.0	30.2	0.7	0.0	0.7	15	23.4	1.4	230.3	124	0	0	57	0	0	11
11	14	2	12	26.2	11.5	28.6	0.5	0.3	0.5	12	22.1	1.9	207.2	117	0	1	55	0	0	16
11.25	15	0	15	23.4	0.0	23.4	0.4	0.0	0.4	14	18.1	1.9	207.4	111	0	1	55	0	0	13
11.5	14	1	13	11.2	6.5	11.6	0.7	0.7	0.7	4	12.4	1.5	191.6	103	1	1	57	0	0	11
11.75	11	0	11	2.5	0.0	2.5	0.2	0.0	0.2	0	0.0	1.4	167.6	100	0	1	59	0	0	7
12	8	0	8	1.4	0.0	1.4	0.2	0.0	0.2	0	0.0	1.4	193.0	104	0	0	60	0	0	6
12.25	11	0	11	0.7	0.0	0.7	0.1	0.0	0.1	0	0.0	0.9	194.7	101	1	0	57	0	0	9
12.5	15	2	13	3.2	2.1	3.4	0.3	0.5	0.3	0	0.0	1.3	159.7	91	0	0	53	0	0	10
12.75	5	0	5	1.1	0.0	1.1	0.1	0.0	0.1	0	0.0	1.0	194.4	97	0	0	54	0	0	11
13	4	1	3	0.7	0.9	0.7	0.0	0.0	0.1	0	0.0	1.0	170.3	102	0	2	56	0	0	7
13.25	12	1	11	2.3	3.6	2.1	0.3	0.1	0.3	0	0.0	1.1	163.8	97	0	0	55	0	0	3
13.5	9	3	6	1.0	0.9	1.0	0.1	0.0	0.1	0	0.0	1.1	163.0	96	0	0	55	0	0	5
13.75	9	1	8	1.0	1.2	1.0	0.2	0.3	0.2	0	0.0	0.6	157.1	94	0	1	52	0	0	8
14	11	1	10	1.0	1.7	0.9	0.2	0.0	0.2	0	0.0	1.0	196.8	89	0	0	52	0	0	7
14.25	9	1	8	0.7	0.5	0.8	0.1	0.5	0.1	0	0.0	0.4	161.2	90	0	0	49	0	0	9
14.5	4	1	3	0.8	0.0	1.1	0.0	0.0	0.0	0	0.0	0.5	130.8	92	0	0	45	0	0	12
14.75	8	0	8	1.2	0.0	1.2	0.0	0.0	0.0	0	0.0	1.1	184.9	93	0	1	49	0	0	10
15	7	0	7	1.1	0.0	1.1	0.2	0.0	0.2	0	0.0	1.0	173.4	95	0	0	49	0	0	9
Fleet Sum	480	57	423	4440.0	165.4	4274.9	185.1	22.3	162.7	115	2814.0	658.9							199	697

- Scenario 5 -

Figure 6-20: Traffic Flow Management Advisories vs. Time





Figure 6-21: Number of Holding Aircraft



- Scenario 5 -

Figure 6-22: Average Delay for Landed Aircraft





Figure 6-23: Average Ground Delay per Landed Aircraft which were Ground Held



- Scenario 5 -

Bibliography

1. H. Erzberger and W. Nedell, "Design of Automation Tools for Management of Descent Traffic", NASA Technical Memorandum 101078, December 1988, .
2. R. C. Larson and A. R. Odoni, "Urban Operations Research", Prentice Hall, Inc., Englewood Cliffs, New Jersey 0732, 1981.
3. R. W. Simpson, TATCA Working Paper 3, Flight Transportation Laboratory, Massachusetts Institute of Technology, 1993.
4. B. Nedell, "A Discussion of Operational Traffic Management Procedures at the Denver and Forth Worth Air Route Traffic Control Centers", unpublished paper.
5. J. W. Andrews, "Impact of Weather Event Uncertainty upon an Optimum Ground-Holding Strategy", Air Traffic Control Quaterly, Volume 1, Number 1, 1993.
6. P. B. Vranas, D. J. Bertsimas, and A.R. Odoni, "Dynamic Ground-Holding Policies for a Network of Airports", Working paper, Operations Research Center, Massachusetts Institute of Technology, 1992.
7. M.C. Chiang, "A Study of Errors in Predicting Arrival Fix Times in Air Traffic Control", Masters Thesis, Massachusetts Institute of Technology, June 1993.