THE DEPARTURE PLANNER:

A CONCEPTUAL DISCUSSION

December 1997

by

Eric Feron
R. John Hansman
Amedeo R. Odoni
Ruben Barocio Cots
Bertrand Delcaire
Xiaodan Feng
William D. Hall
Husni R. Idris
Alp Muharremoglu
Nicolas Pujet

Massachusetts Institute of Technology
Cambridge, MA 02139

TABLE OF CONTENTS

1. INTRODUCTION

This paper will discuss fundamental issues related to the development of the Departure Planner (DP) and will present some initial ideas concerning it. Its primary purpose is to serve as a means of stimulating comments and a focussed discussion among all the parties involved. The paper is also intended to serve as a "living document", i.e., it will be updated periodically to incorporate the most current thinking and findings related to the DP.

The contents of the paper are as follows: Section 2 identifies what are understood to be the fundamental objective, premises and scope of DP. It is suggested that DP offer a major opportunity to think in "systems" terms about increasing the efficiency of terminal area operations, an approach which could provide important benefits to users and providers of ATM services. Section 3 elaborates on DP's objectives and concludes that it may be impossible to identify a consensus "objective function" for such a tool in view of (a) the presence of many stakeholders with occasionally-conflicting objectives and (b) the complexity of the interactions that take place in the terminal area. The important implications of this observation for the fundamental design characteristics of DP are then explored. Section 4 reviews, at varying levels of detail, the principal types of constraints that the departure process must deal with at major airports and associated terminal airspace. Five types are examined: constraints at the runway system; at the gate/apron complex; in the taxiway system; in terminal airspace; and constraints due to environmental considerations. Section 5 discusses the pervasiveness of uncertainty in airport and terminal area operations and its impact on the design requirements of automation aids, such as DP. Section 6 describes information flows within airline departments, between airlines and between airlines and the ATM system as they affect airport operations and DP, in particular. Section 7 provides a brief description of data needs for DP and of plans to assemble the required data. Finally, Section 8 offers some preliminary thoughts on the principal features of the DP and discusses in some depth the idea of the "virtual queue".

Four appendices accompany this report. Appendix A describes the methodology of Activity Cycle Diagrams as a means of modeling ground flows at airports. Appendix B similarly describes the Petri Net approach as an alternative methodology. An extensive set of Petri Net models for the various types of airport operations is given. Appendix C presents in detail a heuristic algorithm for sequencing and scheduling departures from a runway under quite general conditions. This algorithm is analogous to the one used by CTAS to schedule arrivals. Finally, Appendix D presents an initial review of optimization approaches to the problem of scheduling and sequencing both arrivals *and* departures on a single runway.

2. FUNDAMENTAL PREMISES

It is important at the outset to place the Departure Planner into proper perspective as to its scope and potential contributions to ATM capacity and efficiency.

The DP is viewed here as a *real-time*[1] computer-based tool that would support *planning and operations* at major commercial airports. In the temporal sense, its scope (or "time horizon") for planning purposes would be of the order of a few hours (about 3 or 4, at the outside) but most of its emphasis would be on providing timely support to the management of air traffic operations over the next 15-30 minutes from current time. In the spatial sense, the DP's scope includes the entire airfield (including gate/apron) and the associated terminal airspace.

Overall, DP is primarily a *tactical* tool: *given current weather and a short-term (a few hours) weather forecast and given current traffic conditions and anticipated, short-term traffic demand, the tool provides advice aimed at increasing the efficiency of departures.*

While it may appear from the above that the effort to develop DP would be well-defined and would be narrowly focussed on terminal area departure operations, the truth is that DP offers an important opportunity to take a more "holistic" (or "systems") approach to terminal area operations and, as a result, possibly achieve significant improvements in the efficiency of the entire ATM system.

This opportunity arises from the current circumstances of the system. First, it is (or should be) clear that terminal areas and airports are the principal, by far, loci of ATM inefficiencies. This is borne out by any careful review of ATM delay statistics. For example, in an analysis of ATM-related delays and inefficiencies done for NASA prior to the launching of the AATT Program, ODONI [1995] estimated that at least 90% of these are associated with terminal area and airport constraints (and, correspondingly, less than 10% with en route constraints). Thus, any meaningful improvements in the efficiency of terminal area and airport surface operations would have a much more important impact on user costs and inefficiencies than any equivalent (in terms of percent change) improvements en route. (It also seems that the indirect "dogleg" routing of aircraft along the system of en route airways, which have been the focus of much attention recently, are even less of a factor, as a fraction of overall inefficiency costs.)

---

[1] "Real time", as used here, means essentially instantaneous response times.

Second, most efforts on developing automation aids for the terminal area (notably CTAS and, less prominently, the Terminal Area Productivity [TAP] program) have concentrated (quite properly) to date on increasing the efficiency of airport arrivals. However, because of the often-complex interactions between arrivals and departures that take place in practically every important terminal area, the addition of automation aids, such as DP, for departures, may have a profound impact on automation aids for arrivals, as well. For instance, the task of simultaneously managing departures and arrivals presents a new, complex challenge for the existing CTAS. This complexity immediately becomes apparent when one examines operations at such airports as Chicago's O'Hare and Boston's Logan where arrivals and departures typically "compete" for access to the same runway. The added complexity, moreover, becomes enormous in multi-airport terminal areas, such as that around New York City.

In fact, DP cannot be developed independently of CTAS or other arrival automation aids. Rather DP provides an excellent opportunity for launching a "second-generation CTAS", which would take a broader view of optimizing terminal area operations. This enhanced CTAS would encompass the DP as an integral part of its algorithms. To accomplish this integration, many of the algorithms, procedures and processes contained in the current CTAS would have to be revised, in some instances very significantly.

Third, the FAA and the airlines are currently entering a phase of a far more collaborative relationship with regard to the exchange of "real time" data and information about flight status, ATM status and mutual near-term intentions. Coupled with the ability to determine accurately aircraft positions on the airport surface and in airspace, this collaborative relationship opens possibilities for automation aids (such as CTAS and a future CTAS that would include DP) which simply did not exist even a couple of years ago.

Because of these important opportunities and implications of DP, it is essential, before beginning any actual development of tools, to investigate a number of fundamental issues, as well as to carry out a thorough dialogue among ATM researchers and professionals. This paper is intended as a first step in this direction.

3. OBJECTIVES OF DP AND RELATED IDEAS

Possibly the most critical question about DP is the definition of its objectives. It is easy to state, as we have done already, that the objective is to "increase the efficiency of airport departures", but this is too vague for the purpose of designing an actual decision-support system. When trying to become more specific, one must recognize at least two important facts:

(1) Many stakeholders are involved and their objectives are not necessarily the same: Thus there can be several different definitions of "efficiency" depending on each stakeholder's position. For instance, the FAA may define efficiency in terms of cumulative ("national") cost, i.e., total delay costs to users plus cost of providing the ATM service. By contrast, the airlines may define efficiency by comparing actual performance to the best case time/fuel trajectory in the terminal area and on the airport's surface. Moreover, each airline is concerned only about its own costs.

(2) Objectives must take into consideration the entire range of the impacts of DP: This can best be explained through an example. A simple objective that has been suggested occasionally for DP in the past is to "minimize the time from leaving the gate to takeoff (wheels up)". (The time to be minimized can be the average over all flights, a total over some period of time, etc.) Unfortunately, one obvious way to achieve this objective is through a bad strategy, namely by having only one moving departing aircraft on the airfield at any given time. This ensures that each aircraft will move unimpeded from the gate to the takeoff runway and that it will take off in minimum time; but the strategy will, at the same time, undoubtedly cause heavy congestion at the gates and can also grossly underutilize the capacity of the departure runway(s). The underlying difficulty, of course, lies in the fact that the stated objective addresses only one part of the impacts of DP, namely those on the departure-taxiing phase of a flight. It disregards completely the questions of availability of gates and of delay relative to scheduled (or preferred) takeoff time. This illustrates the point that, if the objectives stated for the DP do not take into consideration the full range of impacts of expediting departures, the DP may solve a problem by simply creating another (possibly worse) problem or bottleneck somewhere else.

Some *initial* thoughts and observations follow from the above:

1. It is probably impossible to identify a single objective function that would be fully satisfactory to all stakeholders. It is quite conceivable,

however, that a *somewhat informal* objective function that minimizes a properly selected "total cost" or "total delay " quantity (see below for more details) coupled with some constraints that ensure a *reasonably equitable distribution of delays* among airlines, might be acceptable to the great majority of stakeholders, including most importantly the FAA.

2. The somewhat informal objective function in question must be able to account not only for delays associated with departures *per se*, but also for any possible "side effects" that the DP may have on delays to arrivals and on delays caused by reduced gate availability. (This point follows from observation (b) above.)

   For instance, this objective function might take a form like

   $\sum a_i \cdot f_i(\mathrm{DELAY}_i)$

   In the summation above the a_i are user-specified coefficients, while the f_i are various simple functions of various types of delay (DELAY_i). Note that different functional forms may apply to each type of delay. Examples of delay types that would need to be captured in some way by this approximate objective function include:

   DELAY_1 = additional arrival delay, prior to landing on an arrival runway, caused by DP
   DELAY_2 = additional arrival taxi delay due to DP;
   DELAY_3 = additional arrival gate delay due to DP;
   DELAY_4 = departure delay taken at the gate (flow control or other causes);
   DELAY_5 = departure taxi delay (from gate to departure runway queue);
   DELAY_6 = delay waiting for departure runway (in queue for the runway);
   DELAY_7 = departure delay in terminal airspace.

3. Considerable effort and research should be directed toward determining the most appropriate functional forms of the f_i, as well as toward defining more precisely the quantities DELAY_i. Examples of questions that must be addressed include: is the "cost" associated with each particular DELAY_i non-linearly increasing with DELAY_i and, if so, in what way? What are the nominal times against which delays are measured? How easy is it to make these measurements in the field, so that performance can be monitored?

4. It should be noted that, depending on the situation, some of the terms in 2 above might be unnecessary. For example, the DELAY_1 term will be needed only in airports where departures interfere significantly with arrivals on the runway system or in terminal airspace (interference on the airport surface or at the gates is addressed by the DELAY_2 and

DELAY$_3$ terms) and thus may cause additional airborne delays to arrivals or even delay arriving flights at their originating airports, through flow control. Note also that it will be important to account only for delays caused by congestion and inefficiencies attributed to DP in the particular airport and terminal airspace under consideration. Delays due to "upstream" or "downstream" causes should not "count". For example, a gate delay of a departing aircraft in BOS which is caused by a ground delay program triggered by bad weather at ORD should not be attributed to inefficiencies in the departure process in Boston.

5. The term "somewhat informal objective function" has been used a couple of times above. This is because it is unlikely that any optimization performed by the DP will be carried out in any formal sense, e.g., by running a single, grand optimization algorithm. Instead the DP will probably consist eventually of a set of heuristic or exact optimization algorithms and tools, each of which will address specific aspects of the departure process and of interactions of that process with arrivals. The *overall target* of this system will be to reduce as much as possible a quantity such as the one given under 2 above --and it is in this sense that this quantity is called an "objective function".

6. In relation to 5 above, it can also be stated that, perhaps more than anything else, the DP's eventual success will be determined by its *modularity* and *flexibility*. This stems from the facts that: (1) the DP's objectives are quite "fuzzy", as already noted, and thus the system should be designed so that it will be able to adopt to a wide variety of perceptions as to what is and is not important to achieve (for example, to different settings of the coefficients a_i in our objective function in 2 above); and (2) the DP would have to be adjusted all the time so that it could operate in a great variety of local airport environments, each one presenting different challenges. The need for maximum modularity and flexibility is one painfully learned through the CTAS experience.

7. Similarly, there is a wide range of conditions, with respect to severity of congestion and other disruptions, that DP would have to contend with. There is a major difference between conditions when delays build up in a relatively smooth and straightforward fashion (primarily because capacity falls mildly short of demand over peak-demand periods of one-to-three hours) and so-called *irregular conditions* when airport operations can be brought to a virtual standstill by such events as thunderstorms or very low visibility. The approaches and tools that DP would utilize in the former cases may be entirely different from those in the latter. (Under irregular operations, a successful DP may be one that manages to get even a modest fraction – e.g., 20%, instead of 10% – of scheduled departures off the ground.) If DP is to remain operative under such a wide range of conditions, it may, in fact, be necessary to develop two or more versions of it.

8. An essential condition for gaining acceptance of DP by the ATM user community is the "reasonably equitable distribution of delays among airlines" and other users mentioned under 1 above. One simple way to ensure equitability and "fairness" in the distribution of delays is to maintain *approximately* a First Come, First Served (FCFS) discipline such as the one currently in use in airport departure (as well as arrival) operations. There is, however, a lot of room for interpretation here. On what basis is "first come" determined? Is it on the basis, for instance, of the departure times indicated in the most current version of filed flight plans or on the basis of the time when an airline declares an aircraft to be ready to leave the gate? What about the potential for "gamesmanship" by airlines during periods of congestion (i.e., declaring a flight to be ready for departure earlier than it actually is, so that it can obtain a high place in the queue)? Thus, questions related to the determination of the "baseline priorities" of airport operations must be addressed in connection with the DP.

9. Another consideration related to 8 above is the following. It is well understood that some benefits can be obtained from sequencing of departures to take advantage of diverging departure routes in terminal airspace. The DP may then employ some limited deviations from the FCFS sequence based on the concept of "Constrained Position Shifting" (See DEAR [1976] and ERZBERGER [1995]) which has already been used in the design of CTAS.

10. It should finally be noted that part of the reason "approximately First Come First Served" works today (i.e., deviations from it are acceptable to the users) is that we have multiple interacting queues at some phase of the arrival or the departure process. There is considerable room for exercising human discretion and judgement in merging these queues (e.g., the merging of standard arrival routes at the gate to the final approach path) which promotes efficiency while not inconsistent with notions of fairness. The DP needs to retain such opportunities to exercise judgement (either algorithmically or through human decision-making). If it is defines rigid priorities, just to optimize some specific objective function, then it may not gain user acceptance.

4. FLOW RESTRICTIONS

There is a large variety of flow restrictions that affect airport operations – and departures, in particular – and with which DP will have to contend. It is possible to develop various types of conceptual representations of these restrictions. Two examples of such representations, one using the methodology of *activity cycle diagrams* and the other, the older notation of *Petri Nets* are respectively given in Appendices A (p. 50) and B (p. 54). Representations of this type can be very useful for two reasons: (1) they are extremely helpful in communicating one's understanding of the relationships among the various facilities, services and processes involved in airport operations and, thus, in reaching consensus among experts regarding these relationships. (2) They can provide the starting point for models and analyses which quantify these relationships and help identify bottlenecks, scarce resources and ways to improve operating efficiency at an airport or at one or more of its parts.

Such representations, models and analyses must be an integral part of any future work on DP. In this section we shall discuss in some detail five types of flow restrictions, namely those associated with:

(1) Runway systems
(2) Gate/apron areas
(3) Taxiway systems
(4) Terminal airspace
(5) Environmental considerations

Various combinations of these five types of restrictions play a dominant role in limiting throughput and causing delays at most major airports in the United States.

4.1 Runway Capacity Restrictions

The most obvious and, in most cases, most important flow restrictions are those caused by inadequate capacity of the runway systems of airports. In many cases, it is also at the runway system where the most severe interactions between arrivals and departures take place.

It is essential to realize that *the DP should not consider capacity at the departure runways as simply an externally specified input. Instead, there may often be room for the DP to influence significantly the departure capacity of the runway system,* through (1) the allocation of total available runway capacity between arrivals and departures, (2) judicious assignment of traffic and of aircraft types to departure runways (if two or more departure runways are in use) and (3) improved sequencing of departures. These points will be discussed in this section, which will introduce the important concept of the *runway capacity envelope* for this purpose.

We use the term "runway system capacity" throughout this report to denote "maximum throughput capacity", i.e., the maximum number of operations (landings and takeoffs) that can be performed in one hour at a system of runways in the presence of continuous demand and without violating any ATC separation requirements between successive operations. It is well known that runway system capacity at most major U.S. airports is highly variable and depends on:

(1) The weather conditions (ceiling, visibility, precipitation, wind direction and strength);
(2) The ATC separation requirements;
(3) The aircraft mix;
(4) The operations mix (arrivals vs. departures);
(5) The runway configuration in use at any given time; the assignment of aircraft types and of operations (arrivals and departures) to the active runways;
(6) Human factors (level of performance of traffic controllers and pilots).

Of the above, factors (1), (2) and (3) are beyond the control of the ATM system at any given time. In other words, a primarily tactical tool such as DP (see Section 2) must consider these factors as exogenous. For example, the distribution of the types of aircraft that will use the runway system (aircraft mix) and the operations mix (e.g., 70% of operations demanded over the next 30 minutes are arrivals and 30% departures) are determined by the demand schedule for the current time interval, not by the ATM system. Thus, the DP must work primarily with factors (4), (5) and (6) to minimize, to the extent possible, the impacts of restrictions due to runway system capacity:

- It can manipulate, up to a point (see below for a more detailed discussion) the mix of arrivals and departures handled during successive blocks of time;
- It can select the most appropriate runway configuration at any given time (among those configurations eligible for use under the prevailing weather conditions); and
- It can assign aircraft and operations judiciously to active runways (if the configuration in use has more than one active runway).

Furthermore, it must provide controllers and pilots with the right information, guidance and tools so that they can achieve, in practice, a runway system capacity that is close to the best which is theoretically possible.

It is useful to discuss this topic by starting with restrictions associated with individual runways and then proceeding to considerations involving airports with multiple runways.

4.1.1 Single-Runway Capacity Restrictions

We begin by reviewing briefly the fundamental of determining the capacity of a single runway. This capacity can be approximated by a *capacity envelope* whose typical shape takes a form similar to that shown in Figure 1 and indicates the maximum capacities that can be achieved at the runway, under the entire range of possible arrival and departure mixes. The capacity envelope can be drawn, as shown in Figure 1, by interconnecting four points defined as follows:

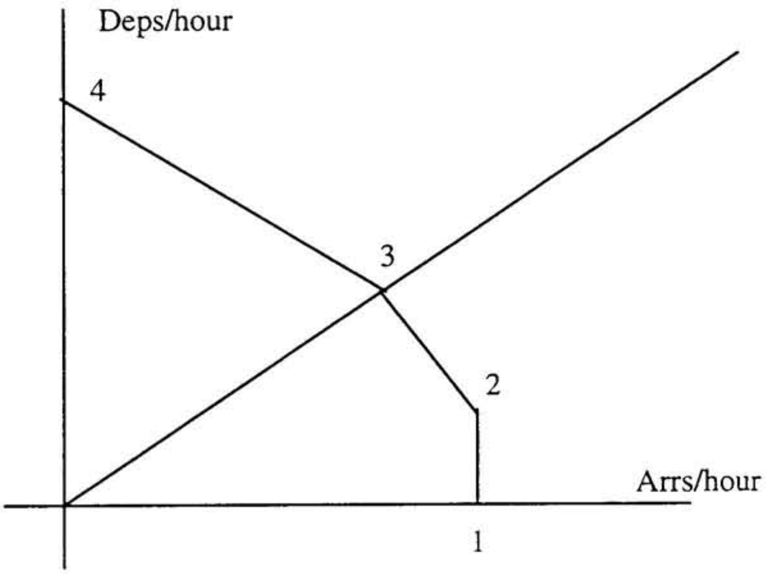

Figure 1: Capacity velope for a single runway

Point 1: *The "all arrivals" point, i.e., the capacity of the runway when it is used for arrivals only.*

Point 2: *The "freely inserted departures" point which has the same arrivals capacity as Point 1 and a departures capacity equal to the number of departures that can be inserted into the arrival stream "for free", by exploiting large interarrival gaps, i.e., without increasing the separations between successive arrivals.*

Point 3: *The "alternating arrivals and departures" point, i.e., the point at which an equal number of departures and arrivals is performed through an arrival-departure-arrival-departure-... sequencing, implemented by "stretching", when necessary, the interarrival gaps, so that a departure can always be inserted between two successive arrivals. (Note that Point 3 is on the 45° line emanating from the origin.)*

Point 4: *The "all departures" point, i.e., the capacity of the runway when it is used for departures only.*

The principal determinants of Points 1 and 4 (all-arrivals and all-departures, respectively) are:

(1) The ATC separation requirements on final approach and on the runway itself;

(2) The runway's physical characteristics (location and geometry of runway exits, length, location of the runway and its exits relative to the terminal building and its gates, etc.); and

(3) The mix of aircraft (i.e., the types of aircraft using the runway and the associated aircraft performance characteristics.

In addition to these three items, the sequencing of operations (arrivals vs. departures) on the runway is another important factor in determining Points 3 and 4.

Example: Logan Airport

At Logan International Airport, as in several of the busiest airports in the United States, the separation requirements between *successive arrivals* are:

able 1: Separation uirements on final approach between cessive arrivals on the same runway (nm)

Trailing a/c Leading a/c	H	L/M	S
H	4	5	6*
L/M	2.5	2.5	4*
S	2.5	2.5	2.5

(Applies only when leading aircraft is at the runway threshold)*

The symbols H, M/L and S stand respectively for "heavy" aircraft (maximum take-off weight, MTOW, of 350,000 lbs. or more), "large/medium" aircraft (MTOW between 12,500 and 350,000 lbs.) and "small" aircraft (MTOW of 12,500 or less). The exception is B757 aircraft which normally would belong to the L/M class but which, due to their wake vortex effects, are treated as a separate category requiring 4, 4 and 5 nautical mile separations if followed, respectively by H, L/M and S aircraft.

In addition to the in-air separation requirements of Table 1, a runway surface separation requirement also applies to arrival operations: a trailing arriving aircraft cannot touch down on a runway unless the preceding landing aircraft is either clear of the runway or, if still on the runway, is at least 8000 feet from the runway threshold.

For departures from the same runway, a set of separation requirements analogous to those of Table 1, but specified in seconds, applies as shown in Table 2 (the B757 is treated as a "Heavy" aircraft for departure separation purposes).

able 2: Separation requirements etween successive take-offs from the same runway (in seconds)

Trailing a/c Leading a/c	H	L/M	S
H	90	120	120
L/M	60	60	60
S	60	60	60

For cases in which arrivals and departures may use the same runway (for instance, in the cases of Points 3 and 4) the separation requirements between an arrival and an immediately following departure and between a departure and an immediately following arrival are also of interest. In the former case, the arriving aircraft must be clear of the runway before the (immediately following) departure can begin its take-off roll. In the second case, the arriving aircraft must be at least 2 nautical miles away from the runway threshold at the instant when the (immediately preceding) departing aircraft begins its take-off roll.

It should be noted that the 4, 5 and 6 nautical mile separation requirements for successive arrivals in Table 1 (and the separation requirements behind B757s) as well as the 90 and 120 second requirements for successive departures in Table 2 are wake vortex separation criteria. Observe that (1) runway occupancy times on landing are of the order of 30 to 70 seconds, (2) times required for take-off rolls are also of the order of 30 to 60 seconds and (3) the amount of time needed to cover 2.5 nautical miles on final approach is also of the order of 60-75 seconds for commercial jets. It then follows that both on departure (90 and 120 seconds) and on arrival (it takes anywhere from 100 seconds to more than 3 minutes to cover 4 to 6 nautical miles on final approach, depending on aircraft type) the wake vortex separation criteria impose major flow constraints. This is because they define the maximum guaranteed flow rate for a specified queue of aircraft. However, in some cases, the pilots can elect to waive the wake vortex separation protection and takeoff at shorter intervals provided they are sufficiently warned. Some potential for optimization is apparent in Tables 1 and 2 if the mix of aircraft on a given runway and their sequencing can be managed to minimize the required intervals between operations.

Another factor that can restrict runway utilization is the need to maintain and inspect the runways. In poor weather conditions, plowing and runway breaking tests may need to be sequenced into the departure flow process. Even in good weather, the runway must be periodically inspected for foreign objects

At least in theory, a single runway can serve, within one hour, an amount of traffic corresponding to any point within the polygon defined by Points 1, 2, 3, 4 and the origin. Each "demand point" is, of course, defined by its co-ordinates consisting of

x = no. of requested arrivals for the hour
y = no. of requested departures for the hour

Note that Figure 1 indicates clearly that there may often be a severe trade-off between the number of arrivals and the number of departures performed at an airport.

It should also be noted that the capacity level indicated by Point 3 may be difficult to achieve in practice, because it corresponds to a strictly "alternating" mode of operations (arrival-departure-arrival-departure...) which requires excellent skills in spacing successive arriving aircraft, so that a departure can be inserted between them. This strategy is often used for relatively short intervals of time during periods of high traffic at many U.S. airports. It is almost never used at European airports where, whenever arrivals and departures operate from the same runway, the usual mode of operations is to use the runway for a string of several arrivals, followed by a string of several departures, then back to arrivals, etc. In this mode of operation the runway capacity curve can be approximated by the straight line that connects Points 1 and 4 in Figure 1, i.e. the polygon of feasible capacities becomes a triangle, defined by Points 1 and 4 and the origin. Note that this triangle lies fully within the polygon of Figure 1, i.e., overall capacity gains can be achieved when interarrival gaps can be "stretched" to insert a departure.

4.1.2 Multiple-Runway Capacity Restrictions

For configurations involving two runways or more, the analysis becomes more complicated, if the capacity envelope is to be derived theoretically, but the basic concepts remain the same.

Consider, for example, the case of a configuration involving two parallel independent runways. This is a configuration which is prevalent in many large airports outside the United States (London Heathrow, Paris CDG, Amsterdam, Munich, New Milan Malpensa, New Athens, Singapore, New Hong Kong, to name but a few) as well as inside. At many of these airports outside the United States, it is standard practice by ATM authorities (environmental reasons are often a factor, as well) to dedicate one of the two runways exclusively to arrivals and the other exclusively to departures. In such cases, the runway capacity envelope for the entire airport (i.e., for the two runways used simultaneously) is as shown in Figure 2. Note that the polygon of feasible capacities is now a rectangle, whose sides are equal to the "all arrivals" (Point 1) and "all departures" (Point 4) capacities of Figure 1.

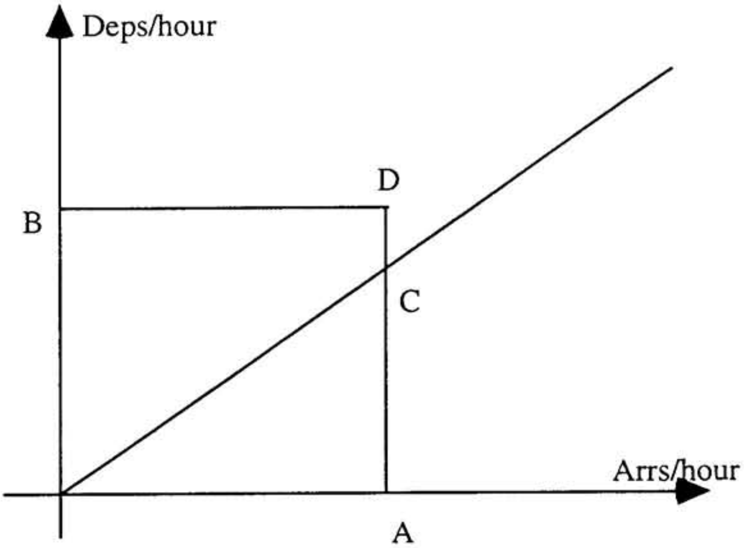

Figure 2: Capacity envelope for two ependent runways, one used only for ivals and the other nly for departures

It is also interesting to note that these airports typically "declare" (for slot allocation purposes) a capacity corresponding to Point C on Figure 2 (i.e., a number of operations equal to twice the length of AC) under the implicit assumption that demand in each hour will consist of 50% arrivals and 50% departures. In truth, these airports can accommodate more operations [= (length of BD) + (length of AD)] corresponding to Point D during hours when more departures than arrivals are demanded.

However, at most major U.S. airports with multiple runways, the situation is not nearly as simple as the one shown in Figure 2. The reason is that, at many of these airports, runways are not "dedicated" to arrivals only or to departures only, but may accommodate both arrivals and departures during any given block of time. And even in cases where arrivals and departures may not be "mixed" on any of the runways, the assignment of individual runways to arrivals or to departures may change over time in response to traffic requirements.

Example: Logan Airport

These points are illustrated well by Logan International Airport. Table 3 lists the seven "basic" runway configurations at Logan as identified by the FAA. A schematic layout of the runway system at Logan is provided in Figure 3. Note that configuration 7 includes all cases in which a single runway is in use.

Table 3: Basic Runway Configurations Identified by the FAA at Boston's ogan International Airport

	Landing Runways		Departure Runways	
ID	Primary	Secondary	Primary	Secondary
1	04L & 04R	--	09	04L & 04R
2	22L & 27	--	22R	22L
3	33L & 33R	--	27	33L
4	04L & 04R	--	04R	04L
5	22L	22R	22L & 15R	--
6	15R & 15L	09	15R	09
7	ANY SINGLE RUNWAY			

The meaning of "primary" and "secondary" in Table 3 is somewhat arbitrary. "Primary" means that the runway is being used (for arrivals and/or for departures) to the fullest extent possible under the prevailing operating conditions and rules, while "secondary" implies use as a back-up or as needed for specific needs. For example, consider the basic configuration 1: In this configuration runway 04L and 04R are designated as the primary arrival runways. There is a "hidden constraint", however, in that the use of 04L for jet arrivals is strongly discouraged for noise reasons (see also Section 4.5 below). Similarly, runway 09 is designated as the primary departure runway. However, jet departures that require a long runway, such as long-range flights may select or be assigned to runway 04R (approximately 10,000-foot long vs. 7,000 for runway 09). In addition, many departures by non-jets are assigned to secondary-departure runway 04L to relieve runway 09.

A further complication is that, as weather changes, "sub-configurations" of the basic configurations emerge. For example, in the case of the basic configuration 1, the use of the three runways described above applies only as long as weather conditions are BVFR or better (ceiling at 1,000 feet or more and visibility at 3 statute miles or more). This is shown at the upper-left-hand part of Figure 4, prepared and provided by Flight Transportation Associates, Inc.[2] However, as weather conditions become more difficult, the use of runway 04L changes (diagrams at upper-right-hand and middle-left-

---

[2] The numbers shown at each arrival and departure runway (e.g., "Arrive 4R: 3 through 9") denote types of aircraft assigned to that runway; for the purposes of these diagrams, aircraft using Logan are subdivided into nine types, e.g., type 8 means "wide-body jets requiring a long runway (> 7,000 feet) for take-off.

hand of Figure 4) due to the proximity of that runway to 04R (runway centerlines separated by only 1,600 feet). First, (upper-right-hand diagram) the use of 04L is reduced to only departures by non-jets and, when the weather is in Category 2 or 3 (middle-left-hand), the runway is not used at all. The final two diagrams on Figure 4 (middle-right-hand, lower-left-hand) refer to the use of basic configuration 4 of Table 3. The use of runway 04L again changes from "arrivals and departures of non-jets" in good weather (middle-right-hand) to "only departures of non-jets" (lower-left-hand) in lower ceiling and visibility. Note that, once again, in Category 2 and 3 conditions, runway 04L is not used and basic configuration 4 turns into basic configuration 7 of Table 3 ("ANY SINGLE RUNWAY") with runway 04R the only one in use.

Overall, it turns out that Logan Airport can be operated in more than 30 different runway system configurations, if one counts all possible sub-configurations (including counting separately each configuration involving a different single runway under basic configuration 7). This gives a good indication of the immense complexity of the task of operating this particular airport.

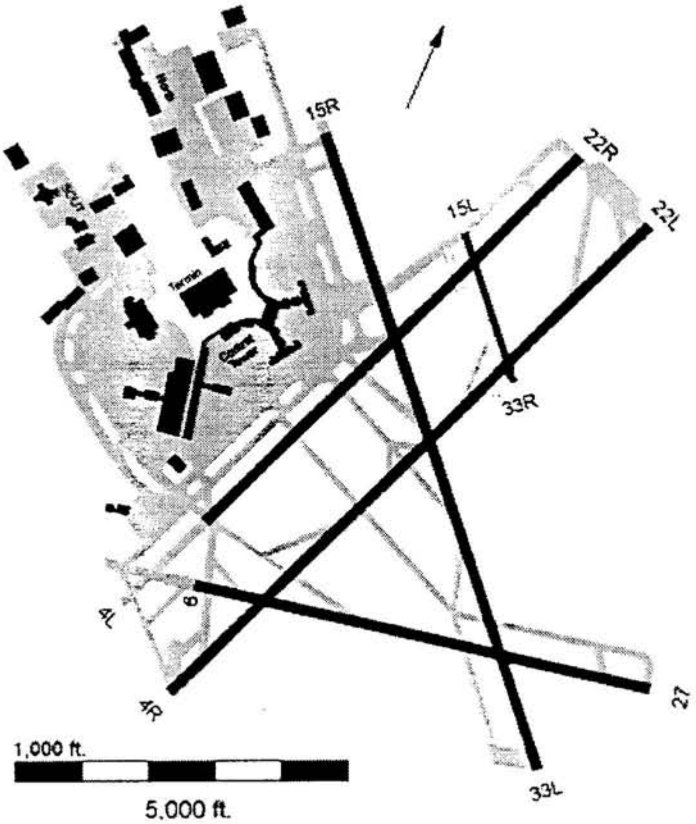

Figure 3: Logan International Airport

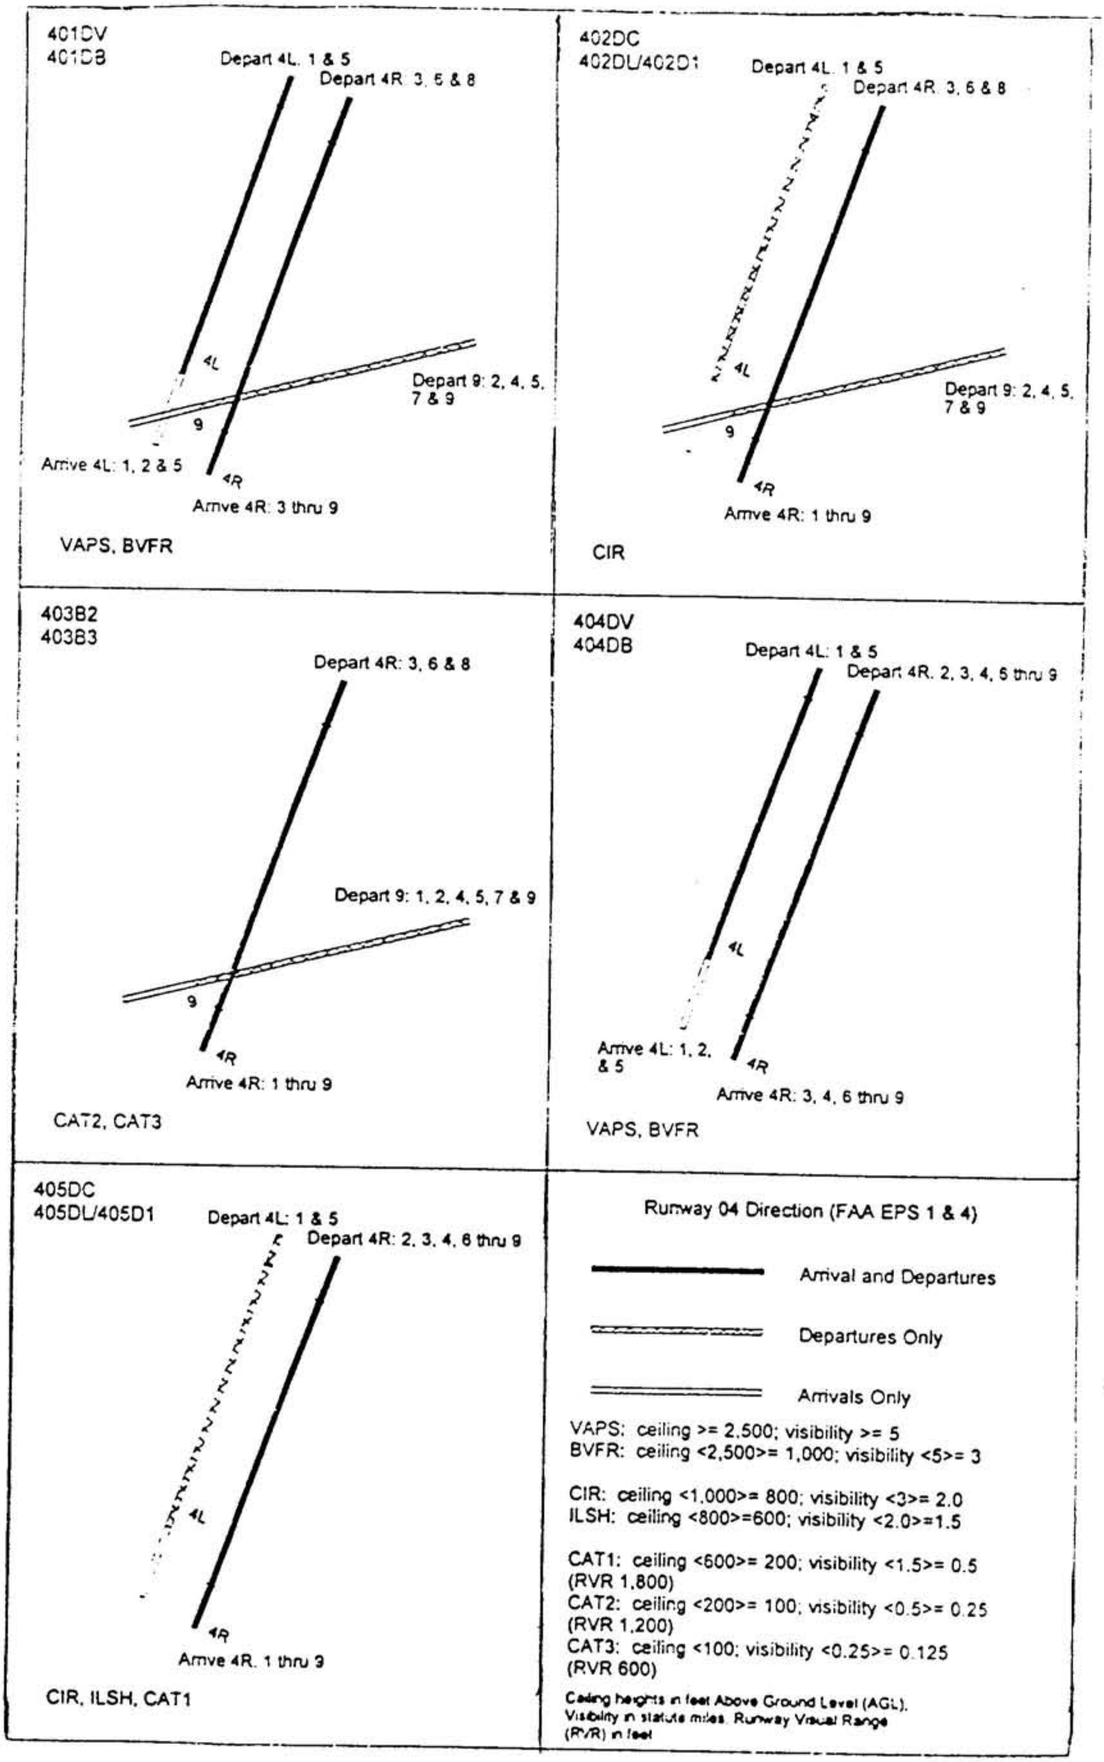

Figure 4: "Sub-configurations" of basic configuration 1

In the most general case, the runway capacity envelope for a multi-runway configuration takes roughly the shape shown in Figure 5. The capacity envelope can now be approximated by the piece-wise linear function ACDEFGB, a more general form of the envelope shown in Figure 1.

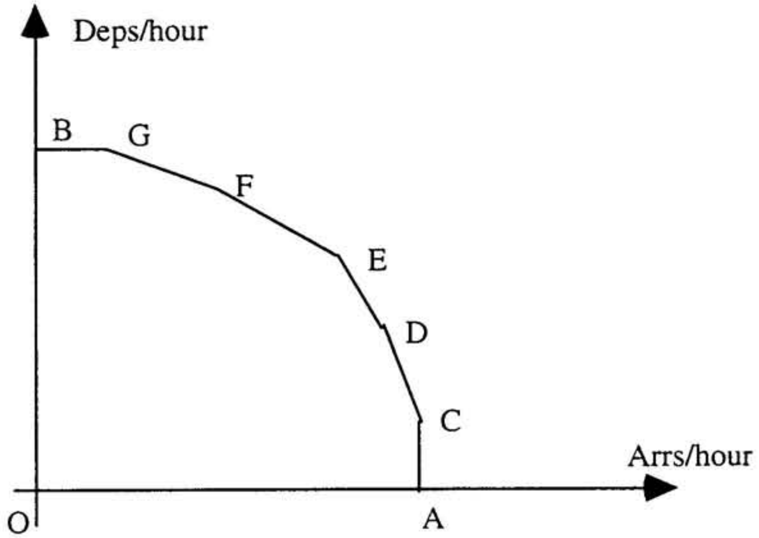

Figure 5: A typical capacity envelope for a multi-runway configuration

Capacity envelopes such as the one shown in Figure 5 can be obtained in two ways: (1) empirically, by taking multiple observations at an airport in peak traffic periods and recording the number of landings and takeoffs that take place during continuously busy time intervals, under different mixes of arrivals and departures (see GILBO [1991]; or (2) theoretically, through the use of analytical or simulation models (see LMI [1995] and STAMATOPOULOS *et al.* [1997]). Both of these approaches require significant effort and attention to some potential pitfalls.

Among the several comments that could be made about Figure 5 one is particularly relevant to the discussion of the DP: multi-runway systems offer the opportunity to control the assignment of arrivals and departures to runways, as well as the assignment of aircraft types to runways. Thus, a multi-runway system "may be better than the sum of its parts". A simple example will again suffice to explain this point. Suppose that Figure 5, represents in fact the capacity envelope of a configuration consisting of two independent runways. Then, Point A corresponds to the capacity of the runway configuration, call it Z, when both runways are used for arrivals only. The point here is that Z may be *more than two times the "all arrivals" capacity of a single runway*, indicated by Point 1 of Figure 1. This is because Point 1 of Figure 1 is computed for the aircraft mix that uses the entire airport. However, with two runways, it is possible to assign certain types of aircraft (e.g., all non-jets) to one runway and the remaining types (e.g., all jets) to another and by so segregating traffic (i.e., by having more homogeneous populations of aircraft on each runway) achieve a higher capacity, Z, than if both runways were serving all types of aircraft. The same observation applies not only to Point B (the "all departure point of the envelope of Figure 5) but also to all the other points on the envelope. In fact, for all points of the envelope other than A and B, the airport operator can control not only assignments of aircraft types, but also assignments of types of operations to the runways, as already noted.

We conclude from this discussion that one of the principal opportunities available to DP for improving the efficiency of airport operations is due to the control it could exercise on factors (5) and (6) that affect the flow restrictions caused by the runway system. The development of algorithms for doing so could be an important area of future research. Earlier work by VRANAS [1992], BERTSIMAS and STOCK [1994, 1997] and, especially, by GILBO [1993] provides a good starting point for this purpose.

4.2 Gate/Apron Restrictions

The gate is one of the turning points of airport operations. Indeed, the airport gates stand on the critical path of the airline business process and operational dynamics:

- All the revenue-generating loads of the flight (passenger, cargo and mail) meet the aircraft at the gate.
- Gates are expensive and are therefore usually a scarce resource - which implies that flights will compete for gate assignments and that some flights will incur delays when no gate is available.
- Gates play a critical role in hub-and-spoke operations. A bank of flights arriving at an airport will use many gates simultaneously to exchange connecting passengers, bags, and crewmembers. In practice, limited gate resources can introduce bank delays and misconnections, which can disturb the airline operations for several hours.

Flow restrictions at the gate/aprons level may be due to several factors, including gate allocation, individual gate operations, gate-to-gate interactions, or interaction of gate operations with other airport operations. We review in this section (1) the individual gate dynamics and (2) the overall organization of the gate/apron area

4.2.1 Individual gate dynamics

Figure 6 shows a PERT-CPM diagram of the turnaround process (i.e. the typical sequence of events between the arrival of an aircraft at the gate and the departure of the next flight) for a major airline.

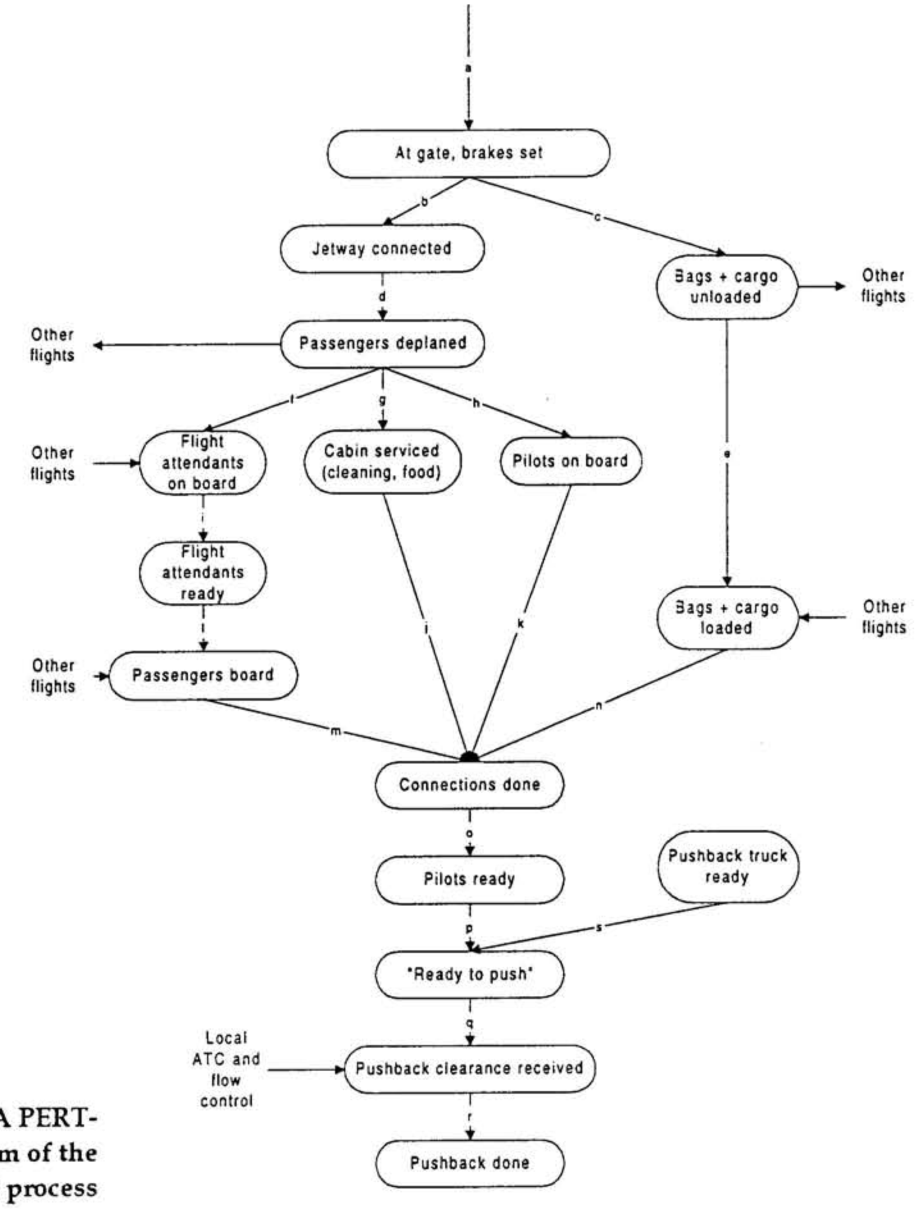

a: taxi in
b: connect jetway
c: install loading equipment and unload
d: deplane passengers
e: load cargo
f: flight attendants board
g: clean cabin, food
h: pilots board
I: flight attendants do cabin check
j: disconnect catering equipment & close catering doors
k: pilots initial checks
I: passengers board
m: close cabin doors
n: close cargo doors
o: pilots finish check-up
p: disconnect jetway & ground equipment
q: ask for pushback clearance
r: push back
s: connect truck

Figure 6. A PERT-
PM diagram of the
urnaround process

This diagram leads us to note that:

- Many operations are performed at the gate: Although delays may be rare for most of the gate operations, this diversity offers much more potential sources of delays than on other segments of the departure process. These operations must be performed with limited ground crew and equipment, which is often shared by several gates.
- The gate is not an independent system: Many events in this diagram are strongly coupled with the rest of the airport and airline systems, which implies that the gate departure delay is not only influenced by the individual turnaround process ("turnaround delay") but also by delays occurring in the rest of the system ("system delays"). In particular, if connections take place, the "passengers board" event is

strongly coupled with the rest of the airline system. The "pushback clearance received" event is strongly coupled to the ATC system.

From this point of view, "shuttle" operations are quite remarkable: special measures are taken (1) to reduce turnaround times to achieve higher service frequencies and aircraft utilization and (2) to decouple these turnaround operations from the rest of the airline system:

- Point-to-point service (no connecting passengers);
- Single aircraft type (usually 727 or 737), dedicated fleet and crewmembers;
- Dedicated gates, ground equipment and crew, complemented by mainline equipment and crew in peak operations;
- Crew and aircraft pairing (a crew flies the same aircraft all day long) to eliminate crew change delays;
- Crewmembers review and sign all flight plans at the beginning of the day;
- Weight and balance information and flight plan updates sent directly to the cockpit printer via ACARS (eliminates the need to print the paperwork at the gate podium printer);
- Reduction of boarding time ("zone boarding");

Turnaround times of 20 to 25 minutes (for a Boeing 737-300) can be achieved consistently using these dedicated turnaround procedures.

4.2.2 Overall organization of the gate/apron area

On most American airports, domestic airlines lease clusters of gates that they control entirely. International flights use a specific terminal where gates are shared among all users. These leases have long duration and lead to major investments within the related terminals. Airlines are often not able to lease more gates. Gates are thus a scarce resource for airlines. This explains why some airlines try and avoid keeping an aircraft at the gate longer than necessary. Assigning efficiently aircraft to gates is an important task of each airline's station. Under irregular operations, the gate assignment plan must be amended in real-time throughout the day to take into account delays and cancellations; individual gate dynamics then needs to be taken into account. In particular, changing a gate assignment less than 20 minutes before arrival usually results in a longer turnaround (need to transfer outbound bags and passengers, ground crew customer service agents, etc. to the new gate). Finally, it is clear that the principle of strict gate ownership leads to additional delays when operations are perturbed. For instance, at Logan Airport, we have observed a full hour of delay between the landing and the arrival of passengers inside the terminal for a late charter flight that could not find an international gate available!

When deciding on which gate an aircraft will park, an airline must know the geometrical constraints that apply. These constraints are of two types: (1) aircraft/gate compatibility and (2) gate combination conflicts. The first are caused by the large differences in aircraft size and geometry; for instance, a "wide-body" aircraft (B747, A340) may not fit in a small gate or the jetway may be too high for aircraft with low doors. This also means that there are significant gate cross-coupling operational effects. For example, at Logan airport, if gate B5 hosts a 757, gate B3 can only accommodate a 737, and vice-versa. These considerations lead to some complex gate management, especially if some flights are delayed and thus block a gate longer than expected. At Logan airport, this situation really occurs for 6 airlines:

- US airways: 8 aircraft types on 16 gates
- American Airlines: 8 aircraft types on 10 gates
- Delta Airlines: 7 aircraft types on 11 gates
- United Airlines: 6 aircraft types on 7 gates
- Continental Airlines: 4 aircraft types on 6 gates
- Northwest Airlines: 5 aircraft types on 6 gates

At a more macroscopic level, terminal geometry and ownership often constrain ramp movements. Jets are not allowed to pushback at the same time from two adjacent gates. Terminal-induced constraints may sometimes be very strong. At Logan Airport, this is the case in the "horseshoe" between terminal B and C (see figure 7 below).

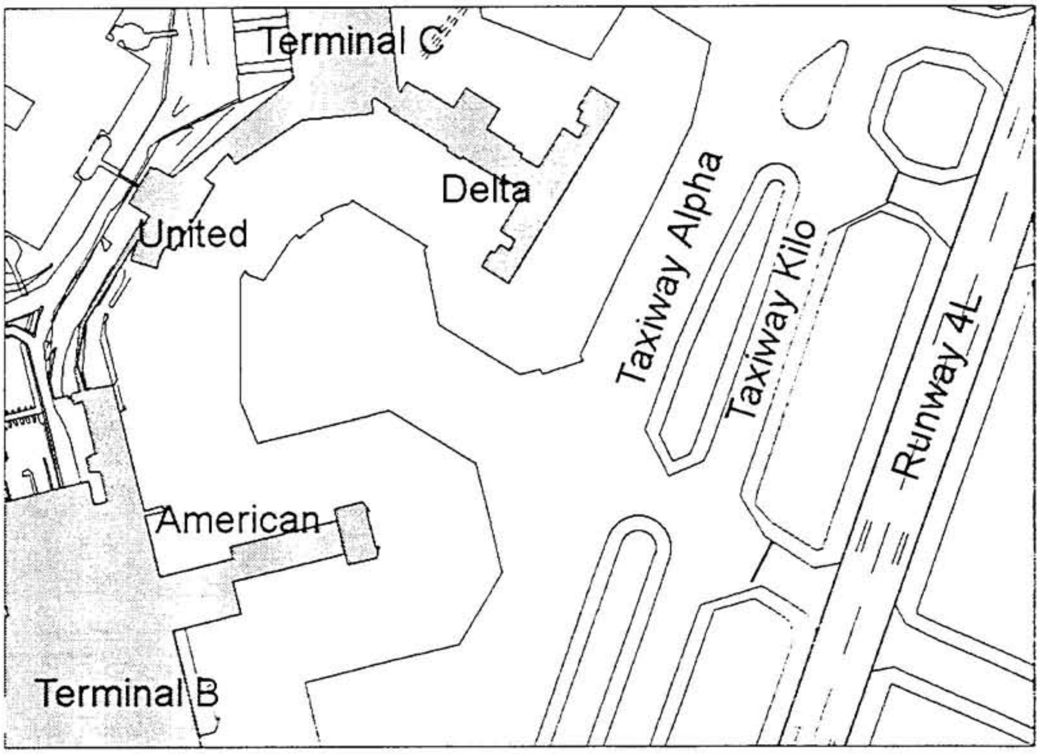

Figure 7: "Horseshoe" at Logan Airport

This apron area is under the control of three different airlines. Only one aircraft movement is allowed at a time. During rush hours, it was observed that such a situation could lead to a large excess of demand over movement capacity. For example, several aircraft can be ready to push back at the same time. In other situations, an inbound aircraft arrives in front of the apron and is not allowed to enter because another aircraft is pushing back.

A much more favorable situation occurs when a major airline has control over a full terminal and part of the ramp area around it, as is the case for United Airlines in Chicago O'Hare (see Figure 8).

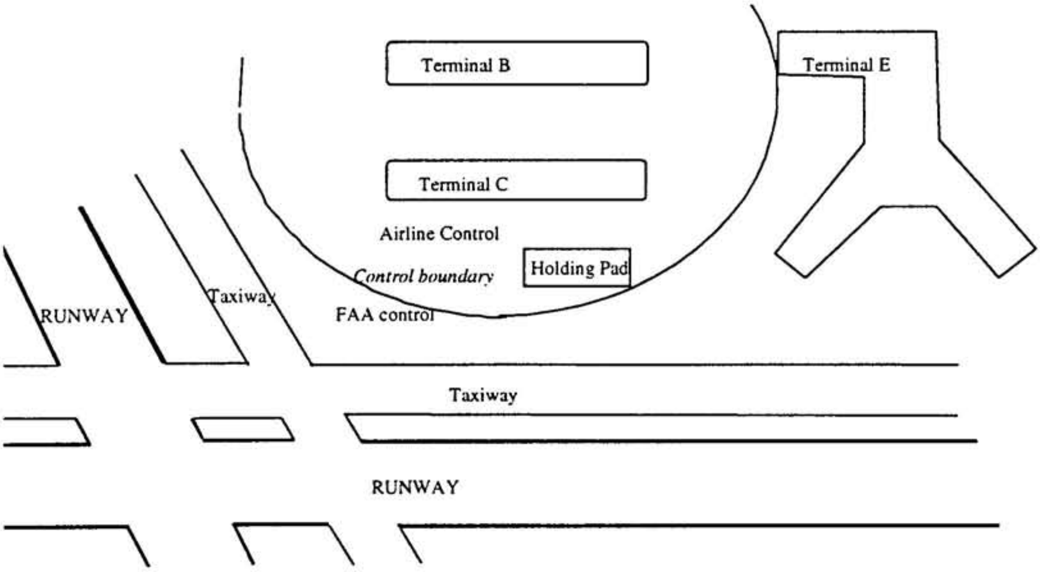

gure 8. Terminal layout and areas f control (United Airlines at ORD)

This layout has important consequences on the dynamic interactions of the airline and the airport:

- The airline operates a ramp control tower that can deliver harmonized pushback clearance to the aircraft parked at its terminals. Aircraft are handed off to the FAA ground control as they taxi towards the control boundary next to the departure runway.
- Arriving aircraft are handed off from the FAA ground control to the airline ramp control as they approach the control boundary. The ramp control tower then directs the aircraft to their assigned gates.

The airline's Station Control Center optimizes in real time the gate assignments to minimize delays and bank misconnections. The airline can also use a large holding pad to:

- Allow an early arriving aircraft to wait until its assigned gate becomes available.
- Store a departing flight delayed by flow control to free its gate for other flights.

Hence the holding pad can be used by the airline as an arrival and departure buffer when gate space is tight. As a result, the airline has complete flexibility on gate assignments and aircraft flow around its terminals. The ramp tower can sequence the gate allocation and pushback according to airline-specific preferences which would be currently difficult to communicate in real time to the FAA - for example giving priority to "bank driving" flights, to business market flights or to late flights.

4.3 Taxiway System Restrictions

In addition to gate interactions, taxiways experience three principal types of restrictions. These may be due to crossing active runways, to intersection with other taxiways, and to queueing phenomena (single-file processions and inability to pass other traffic or to operate 2-way taxiways).

Figure 9 below shows Logan's Southwest corner, where all restrictions occur interdependently.

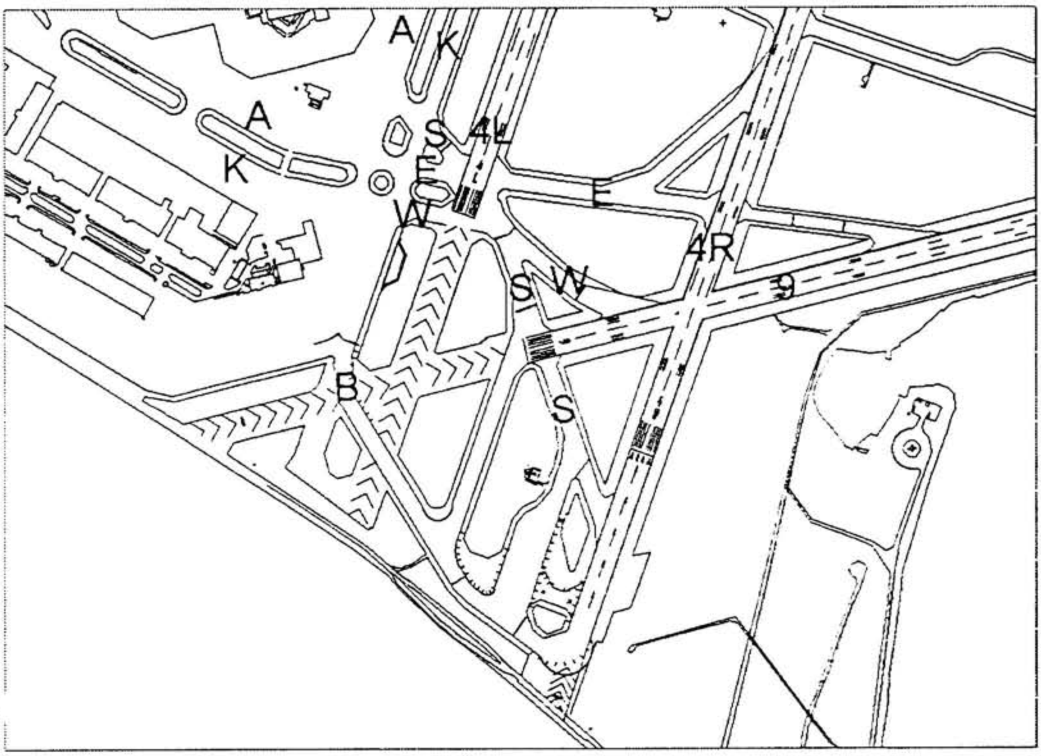

Figure 9: Taxiway yout on the South West corner at Logan Airport

When in Northeast configuration (runways 4L and 4R used for departure and landings, runway 9 used for take-off only), active runway crossing occurs when aircraft cross runway 4L from taxiways Sierra, Echo or Whiskey to reach taxiway Sierra. Likewise, Logan airport controllers will restrict movement on taxiway Bravo when crossing Runway 4L. The effect of active runway crossing is to effectively bind the flow of aircraft on the taxiway. The safety issues involved with active runway crossing are of major importance (runway incursions).

Intersections with other taxiways occur throughout the southwest corner. For example, taxiway Kilo intersects taxiway Sierra, Echo and Whiskey, respectively. The effect of taxiway intersection is to limit the flow capacity in each taxiway. Unlike runway crossing problems, traffic on two crossing taxiways is not constrained as much by safety reasons, and controllers are free to direct the flows.

Queueing phenomena occur on a taxiway network when the flow of incoming aircraft is higher than the maximum flow a given taxiway can handle. Sometimes, queue length is limited. For example, the portion of Taxiway Sierra between runways 4L and 9 can hold a maximum of either two jet aircraft or 1 jet and 2 turboprops. In comparison, Taxiways Alpha and Kilo have a large queueing capacity prior to intersecting Sierra.

The Southwest corner at Logan illustrates that the interaction of two or more taxiway restrictions may actually result in serious performance loss in the departure process. For example, an overload of aircraft waiting on Sierra, Echo and Whiskey for take-off on 4L may slow choke traffic departing on runway 9.

4.4 Terminal Area Airspace Restrictions

The ability of the terminal area airspace to accept departing aircraft can also act as a constraint to the departure process. Terminal area restrictions can be separated into those that are the result of local constraints in the immediate departure area and those that are due to constraints further downstream in the aircraft's trajectory.

Local restrictions can result from a variety of sources. If any of the immediate TRACON sectors on the departure path are restricted due to capacity overload, the departure flows may be limited or stopped. Restrictions can also occur due to in-flight emergencies, Navaid, computer, communications or surveillance equipment difficulties or "pop-up" traffic from secondary airports. Traffic restrictions can also occur due to interference between the arrival and departure flows. For most configurations, the arrival/departure flows are specifically separated to minimize this interference. However, during configuration changes, the arrival departure flow streams must be re-directed which can result in restrictions to both the departure and arrival streams.

Convective weather or icing conditions can also result in a re-routing or disruption of the departure and arrival flows. This effect is particularly problematic when the weather prevents the overflight of key transition points, such as sector boundaries or arrival and departure fixes. In addition, the weather can block entire departure routes or even close the airport entirely. Other weather phenomena that can limit the utilization of departure runways are windshear and very low visibility that may be below the Runway Visual Range (RVR) limits for Part 121 carriers. Moreover, in snow and icing conditions, aircraft must depart within their de-icing fluid Hold Over Time (HOT) limitations or they will need to be re-routed to a de-icing location.

Environmental factors, in particular noise over the communities, can also restrict or influence the departure flows (see Section 4.5, p. 29). Noise abatement departure procedures are common at many high-density airports. In addition to specific departure procedures, agreements with local community groups, such as those described for Logan International Airport in the next section, may limit the times of operation or frequency of use of specific runway configurations.

Other local restrictions to the departure flows include the 250-knot speed limit below 10,000 ft. This restriction has been lifted on an experimental basis at Houston with apparently beneficial results. The higher speeds on departure seem to clear the aircraft out of the departure airspace more swiftly and to reduce both the flow restrictions and the need for the controllers to estimate the acceleration point on the climb.

There are also some very unique local restrictions on departures. For example, at Boston, Runway 22R has a climb gradient restriction when there is a ship in the channel to Boston harbor.

In addition to local restrictions, *downstream* factors can reduce the departure flows. In most cases these result from metering restrictions or flow control restrictions. Metering restrictions, such as in-trail separations (miles-in-trail or minutes-in-trail) applied at downstream points, or restrictions on sector boundary acceptance rates can result from sector overload, airport congestion, weather or equipment failure at some downstream location. Restrictions of these types are complex to manage in that they may not be applicable to all departure aircraft.

Flow Control restrictions can also result in restrictions to the departure process. Typically flow control restrictions are imposed before the aircraft leave the gate. However, in some cases and locations, the Flow Control system specifies a precise "wheels up" time which is an additional constraint on the departure scheduling process. Occasionally, flow control may also be imposed while aircraft are in the taxi process. This is normally due to development of convective weather or emergencies involving runways that suddenly close an airport or severely restrict flows in a specific direction. In such cases all aircraft filed for the impacted destination will be held on the ground.

4.5 Restrictions Due to Environmental Concerns

Environmental considerations are playing a growing role in restricting airport operations throughout the United States and in most developed nations. The restrictions are primarily noise-related, but concerns about air pollution and water pollution are also beginning to have an impact in a number of locations, especially in Western Europe, through imposition of constraints on the total number of operations an airport can perform per year, restrictions concerning engine run-ups and reduced-engine taxiing and regulations regarding de-icing fluids, disposal of waste fluids and other materials.

An important characteristic of restrictions caused by environmental concerns is that they are very *location-specific* . This has major implications for the design of decision-support systems, such as the DP, because it makes it difficult to develop general-purpose algorithms for such systems. The following discussion of environmentally-related restrictions at Logan International Airport illustrates this point and underscores the complexity of some of the technical issues involved.

Example: Logan Airport

Logan Airport is built mostly on landfill and is surrounded by several densely populated areas. Noise pollution is a major concern and has led to an acrimonious relationship over more than 40 years between the Massachusetts Port Authority ("Massport"), which owns and operates the airport, and the surrounding communities and political jurisdictions. Tensions continue today and, in fact, may be exarcebated in the near future, as a result of the pressure that Massport faces to improve the efficiency of Logan operations in order to deal with growing demand and delays.

The principal environmentally-related constraints that aircraft operations at Logan face today are driven by Massport's desire to meet a number of goals aimed at mitigating the exposure of surrounding communities to noise. These goals were arrived at jointly by Massport and community representatives during the early 1980s. While the goals are not mandatory, Massport has pledged its best efforts toward helping achieve them as closely as possible. The degree to which the goals are being achieved is closely monitored and reviewed on a continuous basis by both Massport and the community representatives.

The goals of the noise-mitigation program, known as the Preferential Runway Assignment Program consist of three parts, corresponding to an annual goal, a 72-hour goal and a short-time horizon (a few hours) goal. To help ATM reach these goals, Massport commissioned in the mid-1980s the development of special-purpose software (known as the Preferential Runway Assignment System or PRAS). The specific goals are as follows:

1. **Annual Goal:** Logan Airport has 4 runways (04R/22L, 04L/22R, 09/27 and 15R/33L) which have sufficient length to handle (most, but not necessarily all) jet operations. The annual goals agreed to by Massport and neighboring communities set targets on the usage over the entire year of each end of these runways, for arrivals and for departures by *jets*. Specifically, a target is set for the percent of jet arrivals and jet departures over the year that will operate from/to each runway end. In computing these percentages, each night-time jet operation (defined as an operation taking place between 10 p.m. and 7 a.m.) is multiplied by a factor of 10. "Bad weather" periods, when the selection of the active runway configuration is limited by strong winds, poor visibility, snowstorms, thunderstorms, etc. are not included in the computation of the percentages.

   Table 4 shows the annual targets ("PRAS effective usage goals"). They are intended to strike a compromise between efficiency of operations and noise exposure of the various populated areas around the airport.

   Table 4: PRAS goals on effective runway use by jets vs. actual use in 1996

Runway End	PRAS Effective Usage Goals (%)		1996 Effective Usage (%)	
	Arrivals	Departures	Arrivals	Departures
4R/L	21.1	5.6	32.1	7.2
9	0.0	13.3	0.0	25.3
15R	8.4	23.3	0.8	11.6
22L/R	6.5	28.0	16.3	35.9
27	21.7	17.9	23.8	14.3
33L	42.3	11.9	27.0	5.7

   A number of points should be noted about Table 4. First, runway 33L, when used for arrivals, and 15R, when used for departures, constitute the only two cases where runway usage at Logan generates minimal noise impacts. This explains, the relatively high targets (42.3% of all arrivals, 23.3% of all departures) specified for them. Second, there are additional "hidden" constraints that affect the potential for achieving the goals: the use of runway 04L/22R for jet arrivals and of 04L for jet departures is "strongly discouraged". This means that the 21.1% and 5.6% targets for runways 04L and 04R shown in Table YY must be achieved, in practice, by 04R alone (and, similarly, the arrivals target for runways 22L and 22R must be achieved by 22L alone). Third, the use of runway 09 for jet arrivals is essentially prohibited for reasons of both safety and noise --hence the 0% target and actual use shown in Table 4. Finally, as the actual statistics for 1996 suggest, success in meeting the annual goals is limited – 1996 statistics are not atypical of recent performance.

The reasons for the significant differences between annual goals and actual runway usage are many and their discussion is beyond the scope of this paper. We note, however, a particularly important fact: the targets were set by Massport and the communities, but the organization which has responsibility for runway usage and assignments at Logan is, of course, the FAA through the Chief of Tower Operations. It is unclear whether ATC considerations have been adequately included in setting the annual targets and whether the objectives of Massport and, especially, the airport's neighbors are sufficiently consistent with those of the FAA.

2. **Dwell Time:** "Dwell time" refers to the total amount of time for which a runway is used for jet arrivals or jet departures during a day (excluding the hours between midnight and 7 a.m.). The dwell-time goal for Logan Airport is 7 hours, i.e., no runway end should be used for more than 7 hours during each 17-hour day for operations involving jet arrivals or departures. Periods of bad weather, when the selection of the active runway configuration is limited by strong winds, poor visibility, snowstorms, thunderstorms, etc. are not included in the computation of the dwell times.

3. **Persistence:** "Persistence" refers to the total amount of time for which a runway is used for jet arrivals or jet departures during a 3-day period (excluding the hours between midnight and 7 a.m.). The persistence goal for Logan Airport is 23 hours, i.e., no runway end should be used for more than 23 hours during each 72-hour period (excluding the aforementioned night hours) for operations involving jet arrivals or departures. Once again, periods of bad weather are not included in the computation of persistence. Please note that, if the dwell-time goal is met on three consecutive days, the persistence goal is also automatically met for the same three days. Thus, the persistence goal is somewhat redundant with the dwell goal and is designed to discourage violation of the dwell goal over a period of several days in succession.

It should be noted that, in response to the perceived shortcomings of the initial PRAS software, Massport has commissioned the development of entirely new software ("PRAS 2" or "ENPRAS" --for Enhanced PRAS) to support Preferential Runway Assignment. PRAS 2 has been prepared by Flight Transportation Associates, Inc., a consulting firm, and is a decision-support system installed in the Logan TRACON that makes recommendations to ATC regarding selection of runway configurations over the next several hours in a manner intended to improve compliance with annual, dwell-time and persistence goals.

In addition to the three goals set by the Preferential Runway Assignment Program, a number of additional noise-mitigation practices are currently in place at Logan and may have an impact on runway operations. Massport, for example, recommends the use of runway 33L for late night (midnight to 6 a.m.) arrivals and of 15R for late night departures, with a 10-minute separation required between a departure from 15R and the next arrival (in the opposite direction) on 33L. About 70% of all such late night arrivals did indeed take place on 33L in 1995 and 1996 (as well as about 50% of all late night departures on 15R). Similarly, Massport recommends the practice of reduced-engine taxiing for certain parts of the airport and enforces restrictions on night-time engine run-ups and use of auxiliary power units (APUs). Certain other restrictions on Logan airside operations, such as limiting the number of aircraft on taxiway November to a maximum of 5 can also be ultimately traced to noise concerns.

5. UNCERTAINTY AND ITS ROLE

Uncertainty is pervasive in airport and terminal area operations. It is a factor that leads to fundamental limitations in the efficiency of airport operations. It may have significant implications for several concepts put forward in this paper. Any automation tool, including obviously the DP, that does not take into account this uncertainty and is not able to contend with it, will be doomed to failure.

There are numerous sources of uncertainty that affect the departure process. Some of the most important (this is only a partial list) include:

- Weather (e.g., as it affects airport capacity or as it causes changes to the runway configurations in use);
- Airline operations (e.g., the exact time when an aircraft will be ready to leave its gate or the amount of time a push-back operation will take);
- Air traffic operations (e.g., the time when an arriving aircraft will actually touch down or exit from the arrival runway, or the exit selected by these arriving aircraft, or the time when an adequately long "gap" between arrivals will be found to release one or more departures from a runway);
- Human factors (e.g., reaction times, decisions and actions by pilots, airlines and air traffic controllers).

An essential DP-related task is therefore to study, quantify and model the sources of uncertainty in the current system. It should be recognized, however, that because of local factors and the broad range of conditions that are encountered at major airports, this quantification can only be done in an approximate way. To this end, one should begin by obtaining a good sense of the *relative magnitude* of the uncertainties involved and of the range of values that various random variables may take in practice. For example, the time required for a pushback operation is probably subject to much higher uncertainty than runway occupancy times on take-off and may, thus, need to be addressed at a high level of detail in designing the DP. It is also possible that even a small amount of uncertainty in other variables (e.g., the accuracy of final approach spacing between consecutive arriving aircraft at the same runway) may have very significant implications for the DP (in this example, because it affects in important ways the co-ordination between arrivals and departures). This kind of information can be obtained through a carefully designed combination of direct observations and of contacts with airlines, ATM operators, etc.

The dependence of the uncertainty "on time before the fact" is also a major consideration. In some cases, uncertainty decreases as the time for the initiation of a departure operation approaches, as for example in the case of the length of the departure queue at the runway. (It is far more difficult to predict this length 30 minutes in advance than 5 minutes in advance.) In other cases, uncertainty persists up to the end (e.g., in the case of pushback times from some gates).

Simple probabilistic models of the most important types of uncertainty need to be developed, such as approximate probability distributions, when possible, or approximate expected values and standard deviations.

A number of related observations can also be made:

1. Reduction of uncertainty may be a worthwhile objective in itself, independently of any effects such reduction may have on airport delays or other inefficiencies. For example, an airline may prefer to be told that a given flight will be cleared with certainty to leave the apron area in 10 minutes, than that it will be cleared sometime between 1 and 8 minutes from now.

2. Virtual queues (see Section 8.3 below), if not cleverly designed, may add to uncertainty in ramp operations; this has to be avoided if airlines are to accept virtual queueing.

3. A most important source of uncertainty for airport surface operations is associated with combinations of movements in apron areas. This was illustrated by the "horseshoe" example at Logan Airport (Section 4.3) where only one aircraft is allowed to move at a time, effectively blocking access to numerous gates. Because such events are very difficult to forecast accurately, even a few minutes in advance, this type of uncertainty may prove hard to deal with.

4. The impacts of uncertainty may be highly nonlinear. "Uncertainty" is typically perceived as a measure of the deviation of a variable quantity from some nominal value. That nominal value is usually an expectation (e.g., the expected value of pushback time or the expected distance from the runway threshold to the point where landing aircraft actually touch down). However, some of these deviations may have small consequences for a range of values and much more important consequences for slightly larger values. Stated differently, quite often a slight increase of the uncertainty associated with specific parameters may have a number of major consequences for operations.

Example: Logan Airport

This last point can be illustrated through an example based on experience at Boston's Logan Airport. Consider the situation shown in Figure 10 where two active crossing runways are being used. Aircraft arriving on runway 1 may exit either on Exit Ramp 1 or 2. In the former case, the aircraft may then taxi directly back to the terminal. In the latter case, however, aircraft will have to cross runway 2 to get back to the terminal. However, Runway 2 is being used for departures, and the aircraft from Runway 1 may then have to wait before crossing Runway 2 --sometimes for several minutes. The difference between always being able to take Exit Ramp 1 vs. sometimes taking Exit Ramp 1 and sometimes 2, may be a matter of a few hundred feet in the variability ("uncertainty) associated with the touchdown point on Runway 1. But a slight increase in this uncertainty may result in tremendous variability of taxi times between Runway 1 and the airport's terminal.

ure 10: Uncertainty ue to exit locations

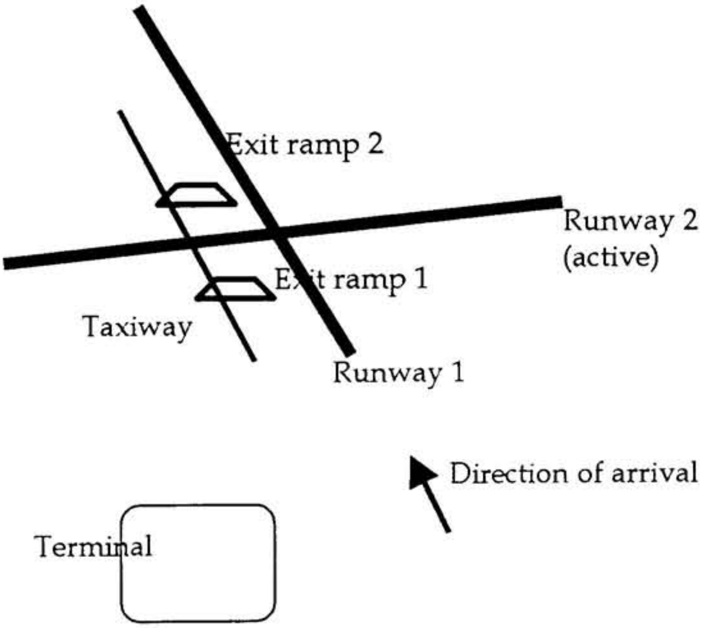

6. INFORMATION FLOWS

Much of the success of a properly designed decision aid depends on its ability to gather sufficient information about the state of the airport and the parts of the air traffic system that affect the airport. The ability of decision-makers to make good decisions is directly related to the accuracy and completeness of the information available to them.

Currently, each airport partner holds a significant amount of information, yet that information often forms an incomplete picture of the situation. There is often little or no overlap in the information held by the various airport partners. Much of the success of the future integrated airport management tool will depend on its ability to gather this information effectively from its various sources and on its ability to distribute it among the airport partners. Past experience with SMA has shown that the appropriate packaging and distribution of information between the airport partners to enhance situational awareness may be perceived as more important than decision making, although complete separation of information packaging from decision making is probably not possible. The organization of information flows across an airport is strongly dependent on the airport configuration and the identity of the major partners; this fact must be accommodated by any useful decision aid.

6.1 Current information flows

The current information flow of interest to airport operations is summarized in Figure 11 (next page).

Proceeding in chronological order, the primary source of information for departure planner is the operational schedule drawn from the Official Airline Guide (OAG) by the airlines. This schedule is continually updated to include previously uncertain information such as fleet status and positioning, passenger bookings, etc.

The Dispatch department of the Airline Operating Center (AOC) then uses this information to prepare flight plans. In addition, the AOC gathers information about general airport conditions and about conditions in other parts of the airspace, including en route sectors, terminal areas and other airports from Air Traffic Flow Management (ATFM) and occasionally Air Traffic Control (ATC). Most of the information that is given to ATFM by the AOC consists of flight plans sent out by Dispatch about two hours prior to departure. Some of the information going from ATM to AOC is made available via the repetition of ETMS (Enhanced Traffic Management

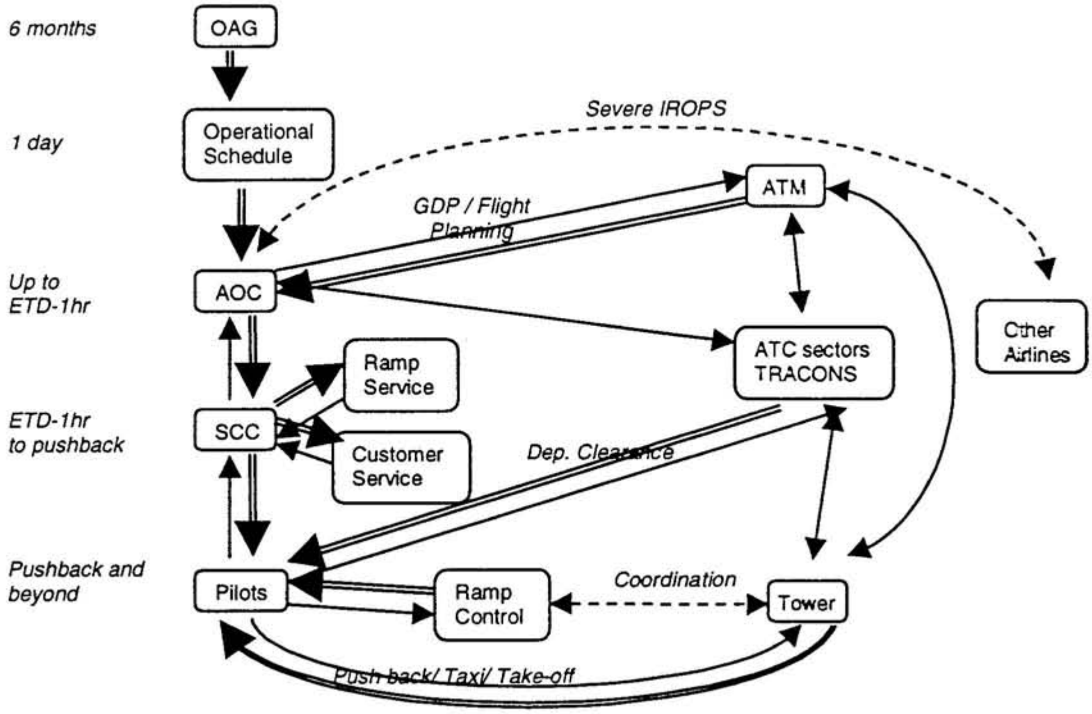

Figure 11: Information flows across ATM (as of today) *Notes: Double arrows indicate information reporting and order passing. Single arrows indicate information reporting only. ETD= Estimated Time of Departure. SCC= Station Control Center.*

System) data in the AOC. Additional information can be gathered from or transmitted to the AOC (especially during the occurrence of ground delay programs) via the airlines' ATC coordinators. In extreme occasions, airlines may also communicate directly and informally with each other in an attempt to mitigate the impact of severe irregular operations such as bad weather or any other major event on operations safety and schedule integrity.

About one hour prior to departure, the airline Station Control Center (SCC), which is located at the airport, takes control over the departure process. The SCC combines flight plan information and local airport conditions to orchestrate the departures. Depending on the airline-specific managerial strategies, the station and the AOC operate more or less independently. At that point, the AOC has effectively lost control over the departure process. SCC supervises boarding and loading operations via information exchange with the customer service agents and ramp personnel.

When the aircraft is ready for pushback, SCC transfers information and aircraft management responsibility to the pilots. Information transferred at that time usually includes updates to the flight plan as well as weight and balance information. Pilots then may interact with the Ramp Tower for pushback and other apron operations, and then with ATC (tower and TRACON) for eventual taxi and take-off.

One global characteristic that may be drawn from current information flows is that in many parts of the system, information tends to follow the same path and pace as operations do: For example, the tower becomes aware of a possible aircraft departure only when the pilot asks for push-back, whereas it could have known about it in principle earlier.

6.2 Future information Flows

Information flows will be further improved by the introduction of Collaborative Decision Making (CDM) and new information sharing via the connection of AOCs to a single information server ("AOCnet"). The station is usually very well aware of the local airport conditions, especially the availability of ramp equipment including tugs, gates and other resources, while the AOC is not.

In many respects, airlines already have the potential to manage significant amounts of information related to the best departure time for their aircraft, including en route congestion and/or ground delay program at destination. They may be able to condense this information to make it available to Airport operators. This information is an attractive complement to what may be gathered by airport operators themselves.

Currently, accurate information from the airlines about their future departure demand, including the amount of demand and its timing, is not made adequately available to the ATM system. Yet, for any type of virtual queue to be implemented, this information must be available. Furthermore, this information must be available to ATM decision-makers if they are to consider trade-offs of the sort discussed in Section 4.2.

However, as was discovered by the CDM group when they tried to make similar arrival demand information available to the ATM system, there are several complex issues related to obtaining this data. In particular, in order for an airline to provide accurate departure demand information, it must be in its best interest to do so. Careful consideration must be given to how potential uses of the airline-supplied data might provide incentives or disincentives to the airlines to report the data accurately. Furthermore, it should be in an airline's best interest to report information promptly, since the timeliness of information has a large impact on its usefulness to decision-makers.

The approach used by the CDM group is to allow the airline information to be conveyed indirectly by allowing the airlines to take part in the decision-making process. Through the decisions made by an airline in pursuit of its own interest, the arrival demand information is made available to the ATM system. It is quite likely that a similar arrangement could be devised to provide accurate departure demand information, allowing the tradeoffs of Section 4.2 to be made and thus allowing the airport resources to be used as effectively as possible. In such an arrangement, each airline could provide not only the demand data that could be derived from an accurate schedule, but also information about the airline's prioritization over any of its resource requirements that might conflict with each other.

The process of providing this information would most likely be iterative; an airline would provide information about its priorities and make decisions, then the ATM system would update its model of the system resources and provide updated system state information to the airlines. The airlines would then use this updated information as the basis for further decisions and information about those decisions, and so forth. Through such a process, the airlines should benefit not only through the additional options available to them and the improvement in system-wide performance due to better information available to the ATM decision-makers, but also through a better ability to predict system performance based on more accurate and complete information provided by the ATM system. Better prediction of ATM departure delays should allow the airlines to schedule their flight crews, ground crews and ground equipment more efficiently, as well as allow them to gauge more accurately the effects of delays on the down-line operations early enough to take proactive action.

The implementation of such information exchange capabilities will probably impact the communication structure between the airport partners. While currently most of the communication load between airlines and the airport tower is ensured via the pilot, the implementation of a virtual taxi queue will probably require enhanced, formalized communication structures between the airport tower and the Station Control Center. Indeed, pilots may be too busy or may not be best informed to communicate and discuss the above-mentioned information with the tower.

Other areas of potential for improved information flow exist between elements of the ATM system. For instance, predictions about sector capacity problems might be used to influence departure sequencing -- a flight scheduled to depart through a sector that will be loaded to capacity for a few minutes might be dropped back in the departure queue until its route is clear, preventing costly en route delay and allowing other aircraft with no such en route problems to use the runway earlier. Similarly, departure queue sequencing could be used to improve estimates of en route sector loading, allowing detection of en route capacity problems far enough in advance to re-route flights with little cost impact.

6.3 Conclusions and recommendations

In a future collaborative ATM system, airlines could contribute significantly to the efficiency of the departure process, in several ways. This could be accomplished through the following steps (listed in increasing order of difficulty in implementing them):

(1) Individual airlines constantly update information concerning the status of each flight preparing for departure, e.g., the expected time until ready to leave the gate, expected time when push-back could begin, etc.

(2) Individual airlines make known to the system provider (e.g., the FAA) their preferences and priorities regarding their own departures.

(3) Airlines collectively, and with support from the ATM system provider, establish a real-time, market-based mechanism (not necessarily involving actual dollars) for exchanging departure slots at times of scarce departure capacity.

The potential of these mechanisms (for reducing uncertainty in the departure process, in the case of the first one, for maximizing utility, in the case of the second, and for resolving conflicting objectives and preferences in the case of the third) will be studied. In particular, the potential benefits from reduced uncertainty as a result of step 1 will be assessed to the extent possible.

7. DATA REQUIREMENTS

The development of DP must be informed by extensive data, some of general nature and some strictly local. The importance of the latter has been underscored by experience with CTAS to date.

Relevant data, some descriptive and others strictly quantitative, must be obtained in every one of the areas discussed earlier in this paper:

- Physical and other constraints to traffic flows on the runways and taxiways, in the apron and gates area and in terminal airspace;
- Level of uncertainty associated with every aspect of airport and terminal area operations;
- Details on operational procedures followed by both users and providers of ATM services and the content, timing and quality of the information flows among them.

Equally important, the "objective functions" and priorities of each of the many stakeholders in terminal area operations must be understood as well as possible, including the criteria and metric each uses to assess the performance of the ATM system.

Example: Logan Airport

A data collection program at Logan Airport has been initiated in connection with the DP project. Its components include the following:

(1) Visits to the Logan tower, TRACON and Boston ARTCC (Nashua, NH): Informal arrangements have been made with FAA personnel to obtain easy access to these facilities by faculty and students.

(2) Informal continuing discussions with ATC staff: An experienced, active air traffic controller is interacting on a volunteer basis with the project team and provides ongoing consultation with regard to Logan terminal area procedures and controller practices. Through this contact, detailed discussions are arranged with flow management, tower and TRACON specialists at Logan.

(3) Interviews with former Logan tower chief: A former Logan FAA tower chief, employed at MIT Lincoln Laboratories, is also providing advice on a continuing basis to the project.

(4) Focused interviews: A series of focused interviews is about to be initiated. The interviews are intended to elicit comments about where perceived constraints and bottlenecks are (cf. Section 4 above) and about the "objective functions" of the various "stakeholders". These include:

- Controllers (tower, ground, departure, arrival, PDC, flow management);
- Pilots;
- Airline dispatchers (AOCs);
- Airline local station managers;
- Airport Authority (Massport) managers.

(5) Airline delay data: Two major airlines have agreed to provide data, based on their respective, pilot-reporting programs for all their flights into Logan. Analysis of the data has commenced.
One airline, collects data on 22 delay categories for airport surface operations, as tabulated below, and on 7 delay categories for the airborne portion of flights. This airline has provided delay data for operations in the first 10 months of 1997 at Logan and, for comparison purposes, at three other major US airports.

A. Out to Off Delays
Airplane de-icing
Airplane systems operational check, cabin check, MEL confirmation, etc.
ATC hold for departure control
Awaiting ATC en route clearance
Awaiting radio closeout information
Awaiting takeoff weather minimums
During pushback due to tug malfunction, equipment problems, or traffic congestion in ramp area
During taxi due to traffic congestion in ramp area
Field traffic vehicular or airplane
Loading additional bags
Other flights landing or departing
Recalculation of takeoff performance data
Runway change
B. On to In Delays
Awaiting gate assignment
Awaiting gate when assigned gate is occupied
Awaiting guideman to park at assigned gate or ground equipment interference
Awaiting Jet Bridge or Passenger Service Personnel.
During tow-in, due to tug malfunction, equipment problems or traffic congestion in ramp area.
Field traffic vehicular or airplane
Mechanical malfunction, airplane towed from runway, etc.
Ramp congestion either airplane or ground equipment
Reduced visibility, snow, ice

The second airline has a different set of categories emphasizing gate delays. Data collected include:

A. Turnaround delays
Delay due to aircraft servicing Delay to complete loading of bags Delay due to cockpit checks
B. System delays
B.1 Airline delays
Delay due to ramp congestion Delay holding for connecting passengers or passengers from canceled flights Delay holding for connecting cockpit crew members Aircraft delivered late due to aircraft substitution.
B.2 ATC delays
Delay due to ATC flow control (esp. ground delay programs) Delay due to local ATC clearance delivery

(6) ATC-derived data: The Volpe National Transportation Systems Center is making available Enhanced Traffic Management System (ETMS) data on Logan terminal airspace operations.

(7) ASQP data: The ASQP data on which the so-called "Airline On-Time Statistics" are based are readily available for Logan airport, among many others. The project team is also trying to arrange with the FAA for early access to the more detailed CODAS data. (The CODAS database has not yet been released for general use; according to the FAA, release may take place by March 1998.)

(8) Field data collection: Some additional data and information, beyond items (1) through (7) above, will undoubtedly have to be collected by the project team, possibly by means of physical observation. For example, the probability distribution of the time from initiation of push-back to "ready for taxi" seems to be particularly important for the DP due to the apparently large variability of this time in practice. Data on this quantity will probably have to be collected at Logan Airport in the near future.

8. A PRELIMINARY CONCEPTUAL OUTLINE OF THE DP

This section presents a conceptual outline of the DP, taking into consideration the discussion and observations of Sections 2 - 7. Clearly this outline is a very preliminary one and is expected to undergo major changes in the future.

8.1 Objective

The objective of the DP could be stated as:

"To minimize the expected delay cost associated with processing aircraft from the moment when they are actually ready to leave the gate to the moment when they leave the terminal area, while ensuring equitable treatment to all users of the airport."

This statement is consistent with the requirements we set in Section 3:

(1) It considers the entire range of impacts of the DP, including its potential impact on arrivals.[3] In fact the delay costs that would be considered would consist of costs due to:

- Additional delays to arrivals due to interference from departures;
- Gate delays on departure
- Departure taxiing delays
- Delays due to waiting for the use of the departure runway
- Delays of departures in terminal airspace

(2) It recognizes explicitly the importance of "fairness" without, however, committing to an entirely rigid FCFS discipline.

(3) It allows for differentiating among different types of delay by attempting to minimize delay cost instead of delay time. This would mean, for example, that, everything else being equal, one minute of delay before leaving the gate would be preferred over one minute of delay in terminal airspace or on the taxiways.

(4) By taking into consideration when an aircraft is "actually ready to leave the gate", it avoids two potential pitfalls. First, it distinguishes between the time when an airline declares a flight ready to leave the gate and the time when the flight is in fact ready to leave the gate. The two are not necessarily the same (see also Section 8.3 p. 46). Second, it also distinguishes between the time when a flight is ready to leave the gate

---

[3] For example, "expected delay cost associated with processing departing aircraft" can include delay costs to arrivals due to processing departures.

and the time when it is instructed by the ATM system to do so. Again the two are not necessarily the same.

(5) By considering "expected delay cost" rather than simply "delay cost", it recognizes explicitly the presence of uncertainty, i.e., the fact that delay is a random variable as a result of its dependence on a number of other random variables.

8.2 Overall Configuration of DP

On the basis of the discussion in Sections 2-7, it would seem that DP should consist of two principal parts: a *Configuration Planner* that would have an approximately 3-4 hour time horizon; and a *Tactical Planner* with an approximately 30-45 minute time horizon. The tasks addressed by each of these two parts can be outlined as follows:

1. **Configuration Planner**: The principal objective of the Configuration Planner will be to schedule the use of an airport's runway configurations over a planning horizon of the next 3-4 hours. This scheduling will take advantage of the observations regarding capacity that were presented in Section 4.1 (p. 10) in order to maximize the efficiency of runway operations in the sense described in Section 8.1 above. The Configuration Planner would thus consider:

   - Short-term weather forecasts for the airport and terminal area; anticipated demand levels
   - Anticipated arrival and departure rates over successive intervals of roughly 15-30 minutes during the planning horizon
   - Environmental constraints on the use of the runways, such as the ones described in Section 4.5 for Logan Airport.

   On the basis of this information, the Configuration Planner will advise ATC with respect to (1) selecting a near-optimal future "schedule" of runway configurations over the planning horizon and (2) assigning arrivals and departures, as well as mix of aircraft, to the active runways for each configuration in that "schedule".

2. **Tactical Planner**: The Tactical Planner will assist ATC in managing the departure process over a planning horizon of the next 30-45 minutes. To this effect the Tactical Planner will:

   - Maintain and update the Virtual Queue (see Section 8.3 below);
   - Determine projected take-off times for each flight in the virtual queue;
   - Consider trade-offs, when feasible, between "taking" expected departure delays at the gate vs. on the taxiway system (after leaving the gate area);

- Assign departing aircraft to runways, if more than one departure runway is available;
- Sequences departures to increase efficiency, while maintaining a reasonably fair allocation of delays among the various airport users.

8.3 The Virtual Queue

The concept of *virtual queue* is one that definitely merits serious exploration in connection with the Tactical Planner in the DP. A virtual queue in the DP context can be defined as a *notional waiting line* of departing aircraft arranged, at any instant of time, according to the order in which the aircraft would be expected to take-off from a runway. (If two or more departure runways are currently in use [or are expected to be shortly] then multiple virtual queues – one for each departure runway will be in use. As an alternative, in such cases there might be a single virtual queue with each aircraft in the queue being "tagged" to indicate which departure runway it will use.) The "virtual" designation stems from the fact that some of the aircraft in the virtual queue may not be physically present on the taxiway system. The virtual queue may have a "tentative" part (i.e., the scheduled departure time and the sequence of some aircraft may be subject to change due to the fact that there is still considerable time to go – e.g., more than 15 minutes until the actual departure event) and a "fixed" part (e.g., the departure time and sequence may be "frozen" 10 or 15 minutes before the assigned time for take-off of each flight).

The virtual queue may offer important benefits to the DP. Consider two scenarios sharing the same flight arrival sequence, arrival times, departure sequence, and departure times. Under scenario 1, aircraft must wait in a physical queue on the airport taxiways to reserve their departure slots, whereas in scenario 2 the aircraft enter a virtual queue, so that each aircraft will taxi to the runway only shortly before its takeoff time.

When the airport is congested, the airlines have the option under the virtual queue arrangement, scenario 2, of keeping the aircraft at the gate for much of the time that would be spent in a physical queue on the taxiway system under scenario 1. There may be significant advantages to doing so. First, and most obviously, scenario 2 saves fuel that would be consumed by aircraft idling in queue under scenario 1, saves time on the engines, and may achieve savings in crew costs, as well. Furthermore, passengers and baggage arriving late enough to miss the flight under scenario 1 might still make the flight under scenario 2. This could have a significant effect at hub airports during periods of congestion, when many arriving flights are late and departure queues tend to be long.

Note that, in principle, the virtual queue is a generalization and extension of the notion of a physical queue. We are free to define it to have or not have various features. In the worst case the virtual queue would be defined to be identical to the physical queue(s) of aircraft waiting for departure and nothing would be lost compared to current practice. If carefully defined and managed, however, the virtual queue may be used to convert taxi delay to less costly gate delay on a one-on-one basis and to increase operational flexibility for the airlines without sacrificing fairness.

Regarding this last point, the virtual queue provides an ideal environment in which to implement an "approximately First Come, First Served (FCFS) discipline" in sequencing departures, which was postulated as a highly desirable attribute of the DP in Section 3. As noted then, one reason "approximately FCFS" works today and deviations from FCFS are acceptable to the users is that multiple queues exist at some phases of the arrival and of the departure process and, thus, there is considerable room for exercising human discretion and judgement in merging these queues. The virtual queue, with decisions on departure sequencing being made while aircraft are still sitting at the gate stands, offers ample opportunity for exercising such discretion for the purpose of promoting efficiency while deviating only mildly from FCFS and not penalizing any user or class of users systematically. Methods for doing so have been studied already extensively (DEAR [1976], PSARAFTIS [1978], VENKATAKRISHNAN *et al.* [1993]).

An additional way in which the airlines could derive much value from the virtual queue is by having a reasonably accurate estimate of the time when the aircraft will depart the gate. If an airline could negotiate a pushback time in advance, this would probably help in many ways, e.g., in scheduling ground crews, baggage handling, etc. A pre-negotiated pushback time would also give the airline the opportunity to make better decisions on how to handle last-minute passengers and on whether to try to get late bags on a plane. It should also be possible to change the pre-negotiated time in light of new information that becomes available before pushback.

The heuristic algorithm presented in Appendix C provides the basic elements of an approach for implementing the virtual queue by identifying favorable sequences of departing aircraft and subsequently scheduling their time of departure form the gate and from the runway.

The principal impediment to implementing a virtual queue may be limitations on physical access to the runway system due to airport taxiways' geometry. In other words, the taxiway network at many airports often limits the amount of re-ordering of the departure sequence, once a departing airplane leaves its gate. If an airplane is to be, for example, the sixth one from now to take-off, it is often the case that the only way this can be done is to have the aircraft physically enter the (departure) highway route immediately after the fifth airplane in the sequence and immediately before the seventh. The heuristic algorithm of Appendix C can be used to take into consideration some of the constraints imposed by such physical limitations.

Potential complications that may arise in connection with the virtual queue – some of which may have already been treated under NASA's SMA research – must also be studied carefully. Perhaps the most crucial of these concerns the time when an aircraft becomes eligible to join the virtual queue. The current system apparently considers a flight to have entered the active departure stream when the corresponding aircraft pushes back from the gate, signaling that it is ready for takeoff. However, a system that allows aircraft to enter the virtual queue for takeoffs before pushing back may allow aircraft to enter the queue before they are really ready for departure. In other words, aircraft that could not enter the departure queue under the current system might be able to enter the virtual queue, if the associated airline exercises "gamesmanship" to reserve an earlier slot than it would otherwise obtain. From the perspective of one interested in minimizing the sum of delays in flight takeoff times, the current system and the virtual queue system would still perform identically, assuming that the virtual queue does not adversely affect flight readiness and that, in the physical queue, gridlock does not prevent the ground-based queue from filling departure capacity. However, an individual flight may see more or less delay in its takeoff time under virtual queueing than under the current system, depending on the rules governing when and how a flight may enter the virtual queue and the degree to which airlines would be willing to bend these rules.

It may therefore be important to think of mechanisms that would ensure fairness in the virtual queue and would enforce desirable approximately FCFS departure sequences without requiring aircraft to queue on the taxiways. Such rules may include a specification of what "ready to leave the gate" means. For example, a plane can be eligible for entering the queue as soon as it declares that it can push back within a pre-determined short period of time, e.g., 15 minutes from now.[4] Should it fail to do so (if asked to) a penalty might be imposed that could take various forms, such as a fine to the airline, or loss of this airplane's position in the queue, so that it has to reenter the queue and wait for its turn once again. The question of how much of a penalty there should be is in itself an interesting one. Too high a penalty (or too short a period of time in the above scenario) would be counterproductive, since it might place too many constraints on the airline.

---

[4] This would also help set the planning horizon for the virtual queue.

9. REFERENCES

BERTSIMAS, Dimitris & Sarah STOCK [1994], *The air traffic flow management problem with en route capacities,* Operations Research Center Working Paper # 3726-94, Massachusetts Institute of Technology (also to appear in Operations Research, 1997).

DEAR, Roger G. [1976] *The Dynamic Scheduling of Aircraft in the Near Terminal Area,* Ph.D. thesis, Flight Transportation Laboratory Report R76-9, M. I. T., Cambridge, MA

ERZBERGER, Heinz [1995] *Design Principles and Algorithms for Automated Air Traffic Management,* presented at Mission Systems Panel of AGARD, Report AGARD LS-200, pp. 7.1-7.31

GILBO, Eugene [1991] *Airport Capacity Estimates for Various Weather Conditions,* ATMS/ETMS Project Memorandum, Unisys Corporation and Volpe National Transportation Systems Center, Cambridge, MA

GILBO, Eugene [1993] *Airport Capacity: Representation, Estimation, Optimization,* IEEE Transactions on Control Systems Technology, 1, pp. 141-154

LOGISTICS MANAGEMENT INSTITUTE [1995] *Estimating the Effects of the Terminal Area Productivity Program,* (Wingrove, W. E., David Lee, P. F. Kostiuk and R. V. Hemm, co-authors) Report NS301R3, McLean, VA

PSARAFTIS, Harilaos N. [1978] *A Dynamic Programming Approach to the Aircraft Sequencing Problem,* Ph.D. Thesis, Flight Transportation Laboratory Report R78-4, M. I. T., Cambridge, MA

STAMATOPOULOS, M. A., K. G. ZOGRAFOS and Amedeo R. ODONI [1997] *An Analytical Model for Runway System Capacity Analysis,* Proceedings of IFAC International Transportation Symposium, Chanea, Greece

VENKATAKRISHNAN, C. S., Arnold BARNETT and Amedeo R. ODONI [1993] *Landing Time Intervals and Aircraft Sequencing in a Major Terminal Area,* Transportation Science, 27, pp. 211-227

VRANAS, Peter [1992] *The Multi-Airport Ground Holding Problem in Air Traffic Control,* Ph.D. Thesis, Operations Research Center, M. I. T., Cambridge, MA

APPENDIX A: ACTIVITY CYCLE DIAGRAMS

A.1 Using activity cycle diagrams to model ground flows

Many modeling techniques are available to describe flows and their restrictions. This appendix presents one of them: *activity cycle diagrams*[5]. Their main application lies in the discrete-events simulation of *strong queueing structures*. They can not model systems where the interruption of an active state may occur before it reaches its scheduled termination.

These diagrams help to consider on the same level all the resources used in the system and to clearly define the *delayed* situations that are the focus of this project. By providing a simple and compact representation of the system, they may ease discussion on the issues we have identified in this project.

An activity cycle diagram considers various objects:

(1) Entities, permanent or temporary;
(2) Attributes possessed by entities;
(3) States entities may be in; these can be regarded as attributes of entities;
(4) Events corresponding to significant state change;
(5) Activities initiated at each event, transforming the state of entities;

In our case, permanent entities are aircraft, runways, gates, ramp staff and pushback trucks, etc. Temporary entities are flights, passengers, bags, cargo, pilots and flight attendants, etc.
An entity is in an active state when it interacts with others and in a dead state when it waits before an activity starts. In our diagram, boxes represent active states while dead states are circles. Arrows represent flows of entities between dead states and activities.
For instance, if we consider the basic operations involving a runway, we note that it is either:

- used by an aircraft taking off;
- used by an aircraft landing;
- idle;
- used for a runway crossing;

The runway alternates between these four states. When it is idle, it can start the activity "landing" as soon as desired. This activity involves both the

---

[5] The principles of these diagrams are explained in Michael PIDD, 1992, *Computer Simulation in Management Science*, Wiley, Chichester, 3rd ed. or in James MCDONALD, Jan SZYMANKIEWICZ and Keith TURNER, 1988, *Solving business problems by simulation*, 2nd ed., McGraw-Hill, Maidenhead, England.

aircraft and the runway. Once the aircraft has landed, the runway is still blocked for some time (due to wake turbulence) whereas the aircraft can proceed to another activity, such as taxiing.

A.2 An example: a detailed activity cycle diagram of ground operations

In this diagram, we have limited ourselves to the main operations performed on the ground and we have shown only the *critical path* along which the temporary entities circulate. This model provides a list of 31 *dead states* that should embrace most sources of delay. It reads the following way:

1. The first source of delay is the dead state No. 1 that corresponds to aircraft waiting for the runway to become available for landing.

2. Once on the ground, an aircraft may face three sources of delays on the taxiway network: it can wait before being able to taxi (2), before crossing a runway (3) or before accessing a holding pad (staging area) (4). The state No. 2 must be understood as the time an aircraft waits before taxiing *because taxiways are used by other entities.* At checkpoints, which are events of the system, the aircraft is either allowed to move on to its gate or it has to go to a staging area.

3. Some delays may come from the unavailability of the gate; they must be put in dead state No. 5 even if they happen while taxiing.

4. As soon as an aircraft occupies a gate, the flight is separated into various entities that follow different paths. Airline operations aim at synchronizing these flows so that they complete their ground activities at the same time. Many operations of fixed duration have to be performed while some may or may not be needed, such as connections or maintenance. For instance, dead state No. 6 corresponds to delays before the unloading of bags happens. This activity is not starting because the ramp crew is not ready: some equipment may be missing or not working properly, etc. Then, there may be some delays in bag connections, for instance because an arrival is late (No. 7). Later, the actual loading can also be delayed (No. 8). Once all these loading operations are finished, the bags enter in the *dead state No. 10* where they wait for other operations to process: passenger boarding, pilots check-up, etc.

5. On the departure side of ground operations, the key event is the **pushback clearance** obtained in most airports from an airline controller and sometimes, from a FAA ground controller (e.g. Logan Airport). In the first case however, the tower delivers a clearance when the aircraft leaves the airline-controlled aprons to enter the network of taxiways. In this model, we have not considered the time spent waiting for this clearance as a separate dead state. This is because we have focused on the *physical flows*. Many other

information flows happen and they would require a much more complex representation.

6. Then, an aircraft may be impeded in its taxiing (27), when it goes to a departure staging area (28), crosses a runway (29) or de-ices (30).
7. It finally gets to its departure queue that corresponds to dead state No. 31.

The table on the next page lists all the *dead states* of the system and their probable sources. In any case, these states are *always due to the unavailability of one or another entity.*

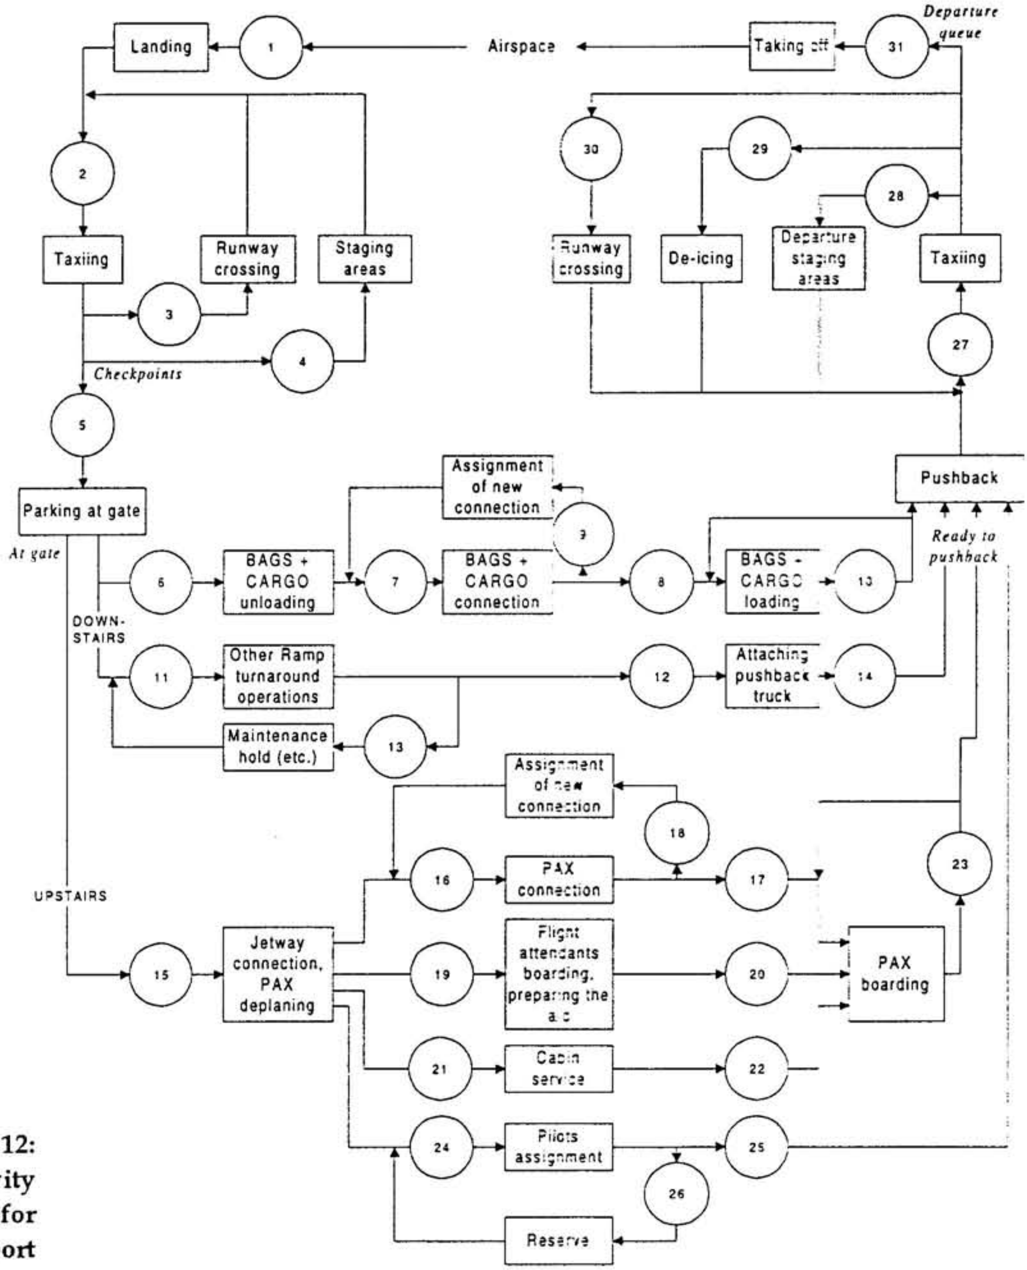

Figure 12: Schematic Activity Cycle diagram for an airport

D.S.	Entity	Probable source
On the runway		
1	Flight	Landing runway unavailable
On the taxiway system		
2	Flight	Taxiways unavailable (a priori because of congestion)
3	Flight	Runway used for an other operation
4	Flight	Staging areas full (rare)
5	Flight	Gate unavailable
Downstairs, bags and cargo		
6	Bags and cargo	Ramp crew/equipment unavailable for unloading
7	Bags and cargo	Delays in the connection of bags and cargo
8	Bags and cargo	Ramp crew/equipment/aircraft unavailable for loading
9	Bags and cargo	Delays in the assignment of a new connection once they have missed their flight
10	Bags and cargo	Waiting for 1) final boarding of passengers 2) Pilots to be ready 3) the pushback truck to be attached
Downstairs, other ramp operations		
11	Aircraft	Ramp crew not ready for other turnaround operations
12	Aircraft	Pushback truck/crew not ready
13	Aircraft	Maintenance equipment/crew not ready
14	Aircraft	Waiting for 1) final boarding of passengers 2) Pilots to be ready 3) loading of bags and cargo
Upstairs		
15	Passengers (PAX), flight attendants, pilots, cabin	Delays in the connection of the jetway
16	Passengers	Delays in passenger connection
17	Passengers	Waiting for 1) end of cabin service 2) flight attendants boarding
18	Passengers	In case of missed connection, delay in the assignment of a new connection
19	Flight attendants	Delays in flight attendants boarding (a priori due to a late assignment or late arrival)
20	Flight attendants	Waiting for 1) passengers connecting from other flights 2) end of cabin service
21	Cabin	Cabin service crew/equipment not ready
22	Cabin	Cabin waiting for 1) passengers connecting from other flights 2) flight attendants boarding
23	Flight attendants, passengers, cabin	Waiting for 1) loading of bags and cargo 2) Pilots to be ready 3) the pushback truck to be attached
24	Pilots	Delays in the assignment of pilots to flights
25	Pilots	Waiting for 1) final boarding of passengers 2) loading of bags and cargo to finish 3) the pushback truck to be attached
26	Pilots	Delays in the assignment to the reserve/ a later flight
On the taxiway system		
27	Flight	Taxiways unavailable (a priori because of congestion)
28	Flight	Departure staging areas full
29	Flight	De-icing pad unavailable
30	Flight	Runway used for an other operation
On the runway		
31	Flight	Take-off runway unavailable: departure queue

Table 5: Dead states for an airport activity cycle diagram, and their probable source.

APPENDIX B: PETRI NET REPRESENTATIONS

B.1 Introduction

This report shows a first attempt to model the airside operations at an airport using Petri Nets. Petri Nets are a very simple modeling tool. Its descriptive value is probably higher than its analytical value. It is very suitable however, for simulation type analysis.

The intention of this exercise is to create a unified mental model of the airside operations. It will be very helpful as a communication tool throughout the project. It can be carried out at any level of details and modularity. It will certainly point out the critical features of the operations, like bottlenecks, deadlocks, interdependencies, etc.

First Petri Nets are described briefly. Then the modeling exercise is carried out for a single aircraft, from on-gate to takeoff, and from landing to on-gate. When interaction with other aircraft, traffic, or components of the system is required, this is shown as a box that could be modeled by another Petri Net. Arrivals and departures are then combined by introducing Petri-Net models for the gate and the runway cycles. This first attempt is intended for feedback. Therefore it is certainly not complete.

B.2 Petri Nets

Petri Nets model both the states and the actions of a system or process. When an action is performed the system changes from one (or more) state to another. A state is represented with a circle or an ellipse, and in called a place. An action is represented with a rectangle or a bar, and is called a transition. Arcs connect between places and transition. Arcs cannot connect two places or two transitions since a transition is always needed to change the state from one place to another. The places, transitions, and arcs constitute the structure of the Net. To indicate that the system or process is in a given state, the corresponding place is marked with a token (black dot). The distribution of the tokens on the places at a certain time is called a marking. The marking indicates the state of the system or process at that time. The marking changes as the system changes its state and therefore it represents the dynamic behavior. Each transition has input places connected with inward arcs, and output places connected with outward arcs.

When each of the input places contains a token the transition is enabled, and the corresponding action can take place. When a transition is enabled it may occur. The occurrence of a transition removes tokens from each of the inputs places and adds tokens to each of the output places. The number of tokens removed and added is determined by an inscription that is attached to each arc.

B.3 Generalizations for Petri Nets

Petri Nets have been generalized in many ways in order to increase their modeling capabilities. Some of these generalizations are described here, and have potential for modeling the airport airside operations.

(1) **Colored Petri Nets (CPN)** allow each token to have an identity (color). The inscriptions on the arcs and transitions in this case determine what color tokens are removed and added to places when the transitions are enabled. Each place has a color set associated with it as well. With CPN it is possible to identify the aircraft flowing in the Petri Nets used to model the airport airside operations. Tokens could also represent different objects in the system, like aircraft, gates, runways and so on, and their interactions.

(2) **Timed Petri Nets (TPN)** allow the transitions to take some time to fire after being enabled. Instantaneous and timed transitions can be used in the same Petri Net. Usually the timed transitions are black and the instantaneous transitions are white.

(3) **Stochastic Petri Nets (SPN)** allow the assignment of probabilities to arcs in the Net. This allows, for example, one of several transitions enabled at the same time to fire according to the probability distribution. For example, when the aircraft is ready to take off, it could either take off or abort. This could be modeled stochastically.

(4) **Queueing Petri Nets (QPN)** allow the tokens in a place to model a queue. In this case when a transition is enabled, the decision on which token to remove from several tokens waiting in an input place can based on some queueing discipline, such as FCFS.

B.4 Potential for analysis

In this section, some notes are made about the potential of the Petri Net model in terms of analysis. No analysis is carried out here; simply some of the properties and tools of Petri Nets are described, which show how they can be analyzed.

Static properties are decided from the structure definition of the Petri Net without considering any possible occurrences. Static properties include for example subgraphs of the Net, which are formed from a subset of the nodes and arcs.

Dynamic properties describe the behavior of the Petri Net. These include:

- **Boundedness:** Upper and lower bounds on how many tokens of a particular color a place can have.
- **Home:** A marking or set of markings to which it is always possible to return.
- **Liveness:** Binding elements that are always active. A binding element is a transition with its surrounding variable tokens and inscriptions bound to some values. Live transitions can always occur, but don't necessarily occur.
- **Fairness:** Tells us how often the different binding elements occur.

Simulation is the most straightforward analysis tool for Petri Nets. There are also computer tools for the implementation of Petri Nets graphically. However, simulation allows understanding and debugging of the Net behavior, but it does not allow formal analysis of the Net properties.

Occurrence graphs show the sequence of occurrences of the transitions. It contains a node for each reachable marking, and arcs connecting the markings starting from the initial marking and following all possible sequences.

Invariants are equations that are satisfied for all reachable markings. They describe a set of sequences that have no total effect. Invariants constitute a very important analysis tool, they are less compact than the occurrence graphs that can become very large, but they require more mathematical analysis.

Reduction rules exist which allow reducing a large Petri Net into smaller and more manageable ones.

Finally, the performance of the Petri Net in terms of efficiency is another important property in some applications, especially if the Net is used for real time applications.

B.5 Petri Nets and the Departure Planner

In this section some ideas about how the Petri Nets models of the airport airside operations can be useful for the Departure Planner. These ideas are by no means exhaustive. They are simply the result of an ongoing process of literature review and modeling refinement, both of which are very important at this stage of the project.

B.5.1 Graphical and conceptual representation

There are mixed views about the power of Petri Nets as an intuitive graphical representation of the dynamics of a system. As systems become more complex the number of places and transitions increases, the arcs cross, and the picture becomes cluttered.

In order to keep the picture clear graphically however, many components and details of the system need to be kept out. This is achieved for example in the PERT-CPM chart of Figure 6.

This can also be achieved using Petri Net type representations, such as the Activity Cycle diagram in Appendix A. This activity cycle diagram is similar to a Petri Net by matching the dead states to places, the activities to timed transitions, the entities to tokens, and the attributes of the entities to colors of the tokens. By keeping many components and details out, and combining activities, the activity cycle diagram presents the whole airport airside system in one clear flow diagram.

The Petri Nets in this appendix model the different parts of the activity cycle diagram of Appendix A in more details. Graphical clarity is still achieved by keeping the individual graphs small and decoupled, and by decomposing the graph in stages. For example, in Figure B2 one transition is used to represent all the turnaround activities, to transform the flight from the "ready for turnaround" state to the "ready for pushback" state. This transition is then expanded in Figures B3 to B17, to represent each of the turnaround activities, some of which are concurrent and some are sequential. These Petri Nets show also the interactions between the flight and the different other resources such as the fueling, cleaning and catering crews, the passengers, the baggage and freight crews, the tug crew, the cockpit and cabin crews, the maintenance crews, and the ATC and ramp control. In Figure B26, when the aircraft and the gate cycles are combined, the turnaround activities are combined again into one transition between the on-gate and the ready-for-pushback states.

This ability to reduce a complicated Petri Net or to expand a simple one to any level of details is an attractive feature of Petri Nets. It allows a trade off between the graphical intuitive representation and a detailed complex representation depending on the purpose for which the Petri Nets are used. In fact the complex details are only needed for simulation and analytical purposes, and therefore, can always be hidden when the Petri Nets are used for representation and communication purposes.

B.5.2 Functional representation

The Petri Nets represent all the activities and functional requirements which transform the flight, and other agents in the system such as crews and gates, from one state to another. When stochastic timed Petri Nets are used, the times for the occurrence of the transitions are modeled by probability distributions. In this case it is possible to estimate a time window when an aircraft (or an agent) will be at a future desired state, given the current state. For example, it would be possible to estimate a time window for reaching the "ready for pushback'" state given the different states that the aircraft and the other agents and resources are currently in, in Figures B1 to B17.

This can be done by aggregating the time along the routes in the Petri Net that lead the aircraft from its current states to the desired one. This aggregation includes (in terms of estimates) the remaining time for the current activities, the waiting time before future activities, and the time for the future activities. The times of sequential activities are added, and the maximum of the times of concurrent activities is selected. The difficulty, however, is in coming up with the time distributions for the different activities, under different conditions.

In this sense the Petri Nets model the dynamics of the airside system, and can play a central role for the Departure Planner in this part of the system. As central for example, as the role of the trajectory synthesizer of CTAS, which is used to estimate arrival times at different points in the terminal area air space. The dynamics here however, is of a discrete event nature rather than continuous. This role is essential for all the tools that the Departure Planner will suggest, and is one of the main current concerns as mentioned in Section 6.3.

B.5.3 Fault diagnosis and constraint identification

As a graphical and a functional representation of the system, Petri Nets type models are very useful in identifying where problems may arise in the system and the sources of these problems. At least graphically it is possible to visualize the critical parts of the system where delays may occur. In the activity cycle diagram (Appendix A) many of these delays are pointed out. Both in the activity cycle and in the Petri Nets, whenever an agent, for example an aircraft or a crew, is in a place ready for the next transition or activity, and other places that are needed for this transition are still empty, the entity is waiting and incurring delay. Also whenever there are more that one aircraft in a place, and the next transition is enabled, only one aircraft is going to be active and the rest will wait in the queue.

For example, in the gate cycle in Figure B26 an aircraft can be "ready to park" while the gate assigned to it is not in the state of "free and available". This aircraft has to wait until its gate becomes available in order for the next transition (the parking event) to occur. The gate would still be in one of the other places in the cycle, either being occupied by another aircraft that in undergoing turnaround activities or holding on the gate, or is free but blocked by the an aircraft pushing back, or by ground equipment. If for example the gate is occupied by an aircraft undergoing turnaround activities, a look at the status of the aircraft in the on-gate Petri Nets (Figures B1 to B17) tells us what is holding it up.

Besides the graphical representation however, Petri Nets provide an abstraction tool that allows for the analysis and diagnosis of many structural and dynamic situations that are restrictive. Vast literature is available today that document many applications of Petri Nets models for identifying and diagnosing problems, especially in computer systems and manufacturing systems.

Two such problems are deadlocks and bottlenecks. Deadlocks occur because of the following conditions:

- Mutual exclusion: processes require the exclusive use of a resource.
- Hold while waiting: processes hold onto resources while waiting for other required resources to become available.
- No preemption: processes holding resources determine when they are released.
- Circular waits: closed chain of processes in which each process is waiting for a resource held by the next process in the chain.

The first three are most of the time inherent in the system. Circular waits need to be identified and avoided. In the example above an arriving aircraft is ready to park but its assigned gate is occupied. Following the Petri Nets of the turnaround activities of the parked aircraft, it may turn out that the gate is being held because the departing aircraft cannot pushback due to blocking by the arriving aircraft. Heuristic algorithms have been suggested in the literature to prevent and avoid deadlock situations.

The bottleneck of the system is the part of the system that is operating at the slowest rate, and hence is restricting the flow in the entire system. For example, Figure B26 shows the gate cycle, Figure B28 shows the runway landing cycle, and Figure B29 shows the runway departure cycle. There may be several of these cycles in an airport. These cycles feed aircraft to each other, each one at its cycle rate. The system as whole however, can only process aircraft at the slowest rate. If the runway cycle is slower that the gate cycle, the runway is going to build large queues, unless the gate cycle is slowed down.

(This is a very simplistic example, other components of the system, and other cycles need to be taken into account). Some heuristic algorithms have been documented in the literature, which identify the slowest cycle in a system and find optimal initial markings of Petri Nets accordingly.
It should be mentioned here also that because of the discrete nature of Petri Nets they could also model the logic and procedures that are currently used by the different controllers, in dispatching for example. These logic and procedures can be tested for faulty or problematic behavior as they interact with the rest of the system.

B.5.4 Analysis of new strategies

Just like the current system and procedures, many of the new tools and procedures that the Departure Planner will suggest can be modeled using Petri Nets. Petri Nets therefore, will also provide a test bed for analyzing how the new system will be able to resolve the different problems identified in the current system and procedures. This can be done through comparative analysis and simulation.

It is premature to elaborate on testing new strategies at this point. However, it is important to realize the potential for it as a motivation for undergoing such a modeling exercise. Petri Nets can be helpful at all the levels at which the Departure Planner tries to contribute. At the tactical and implementation level, the functional modeling is essential as mentioned above to ensure stable and smooth operations under a particular configuration. Also at the configuration planning level, the higher level mode switching logic can be modeled using Petri Nets. Often modeling techniques like Petri Nets have been used to model such higher level, supervisory, task switching decision-making.

For example, not only the stability and performance of the operations under each configuration should be ensured, but also the transition between different configurations needs to be stable and smooth. By modeling such switching logic (which may be suggested by the Departure Planner) using Petri Nets and integrating it with the lower level operations under each configuration, interesting analytical research is motivated.

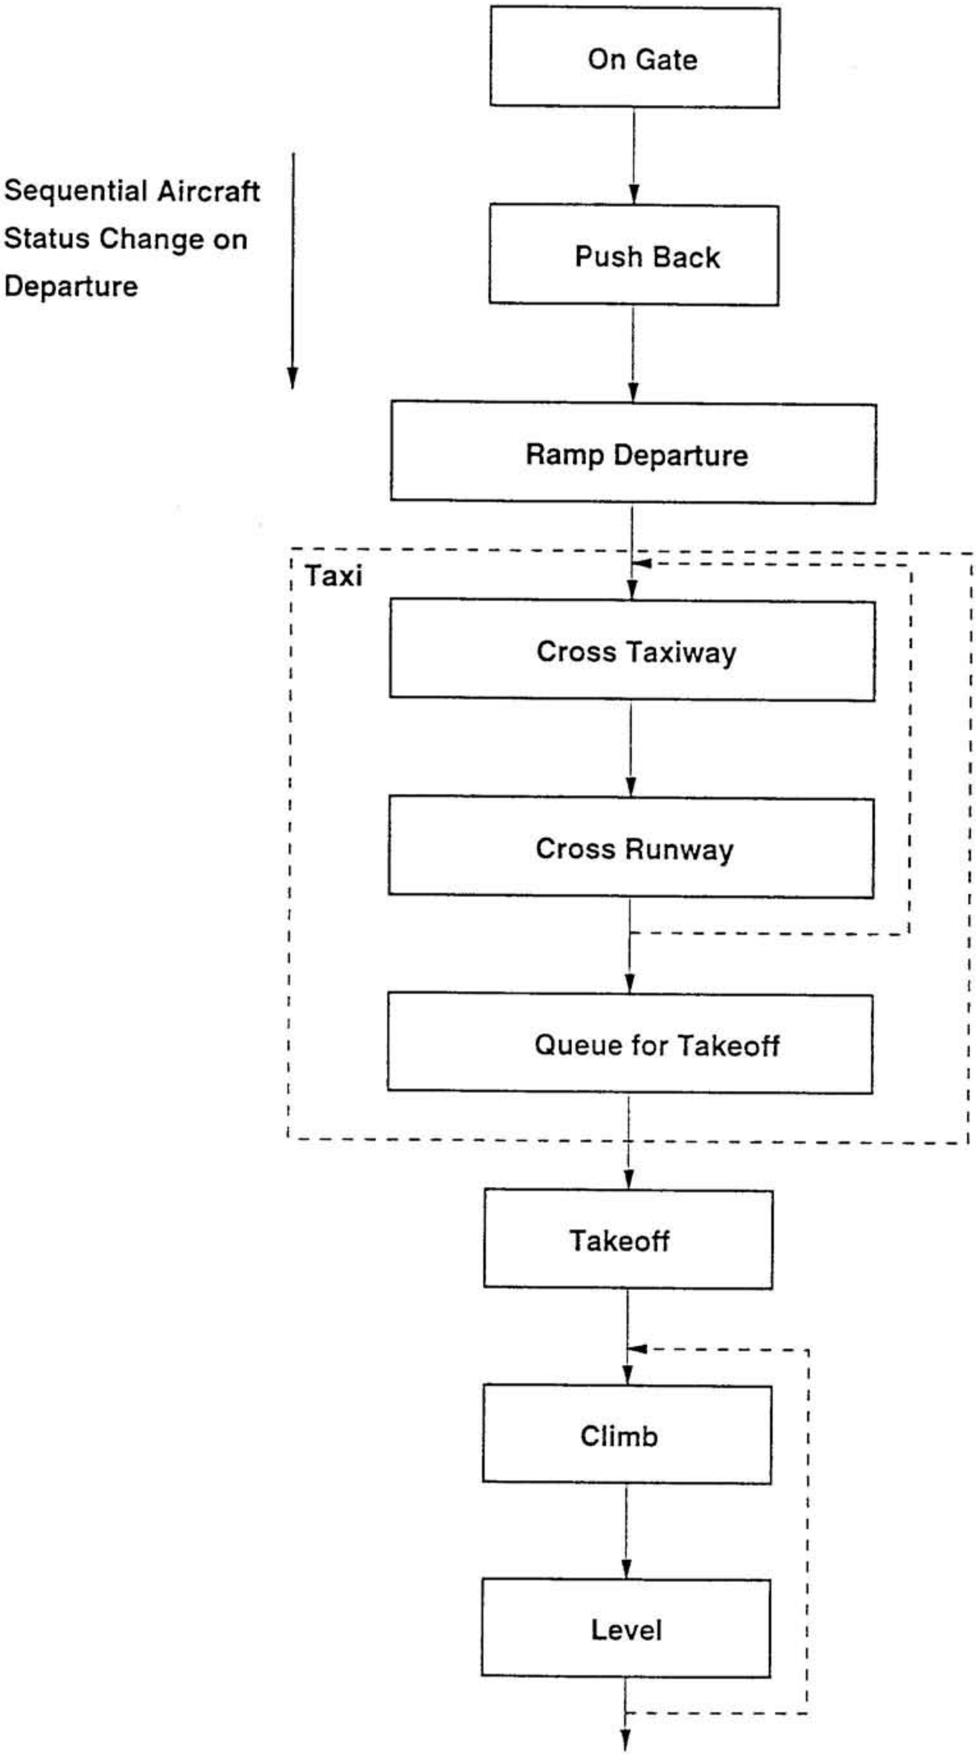
Sequential Aircraft Status Change on Departure
On Gate
Push Back
Ramp Departure
Taxi
Cross Taxiway
Cross Runway
Queue for Takeoff
Takeoff
Climb
Level

Figure B2: Turnaround Process (Preliminary)

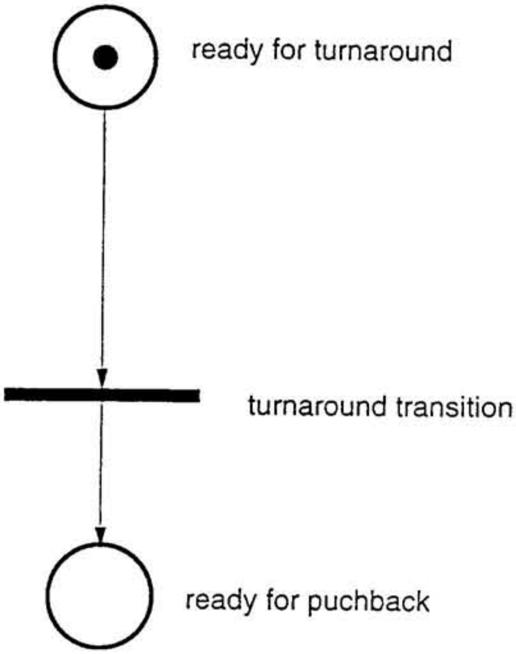

Figure B3 : Fueling Process (Preliminary)

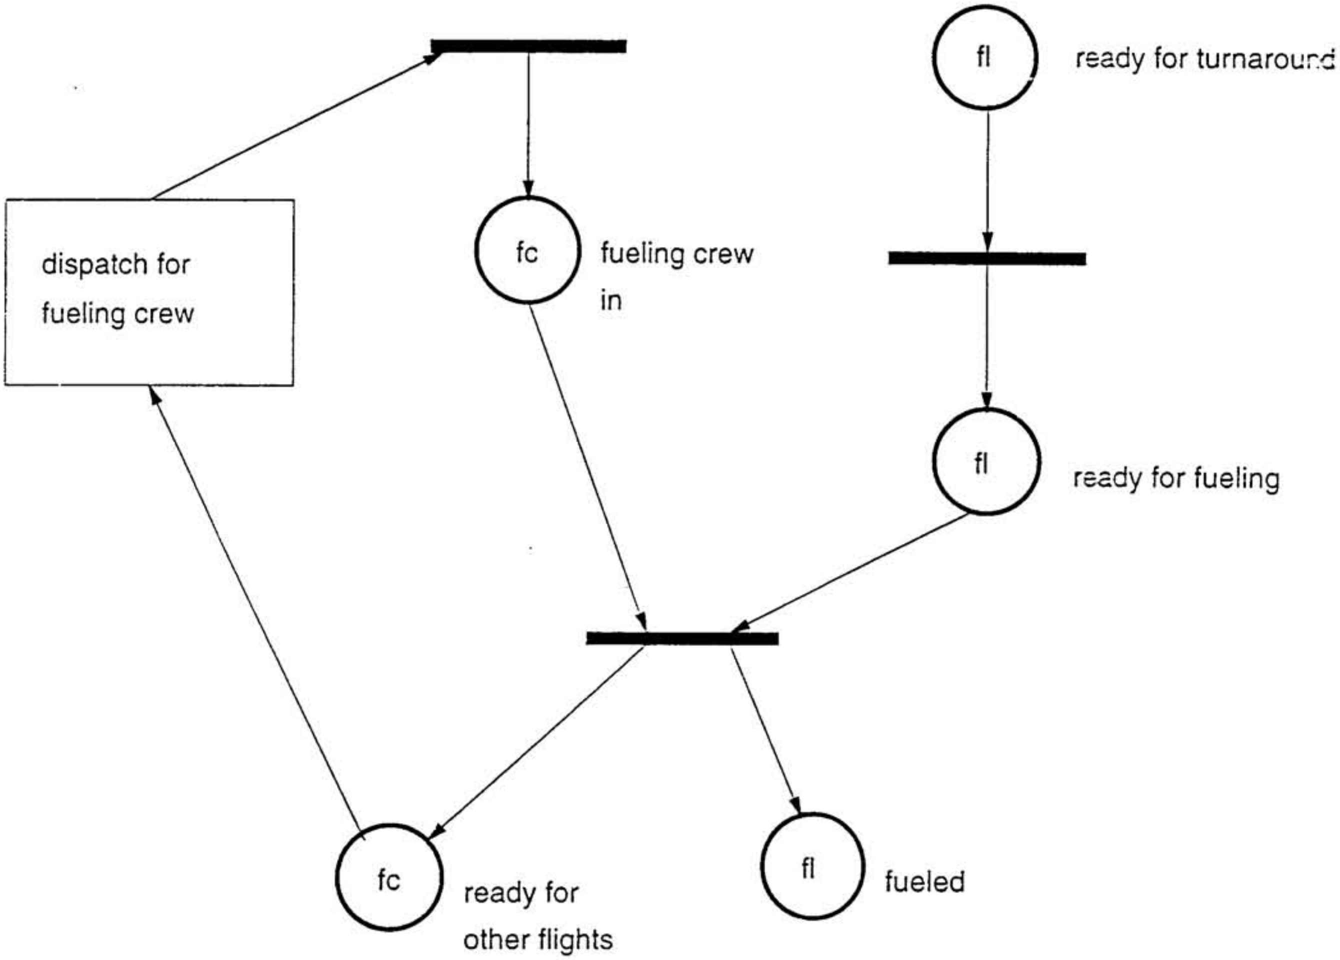

fl = flight
fc = fueling crew

Figure B4 : Deicing Process (Preliminary)

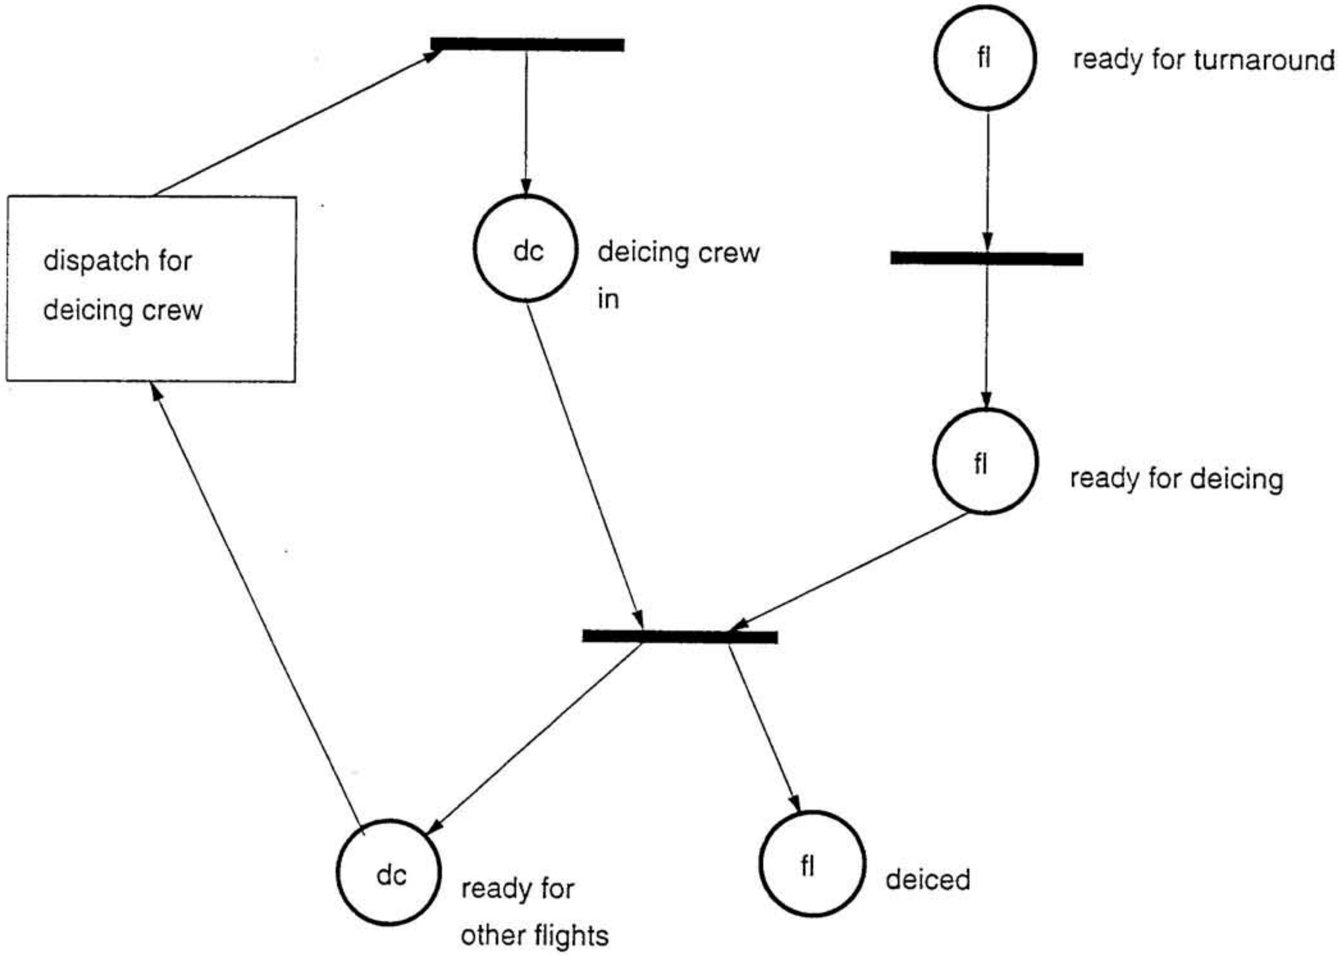

fl = flight
dc = deicing crew

Figure B5 : Catering Process (Preliminary)

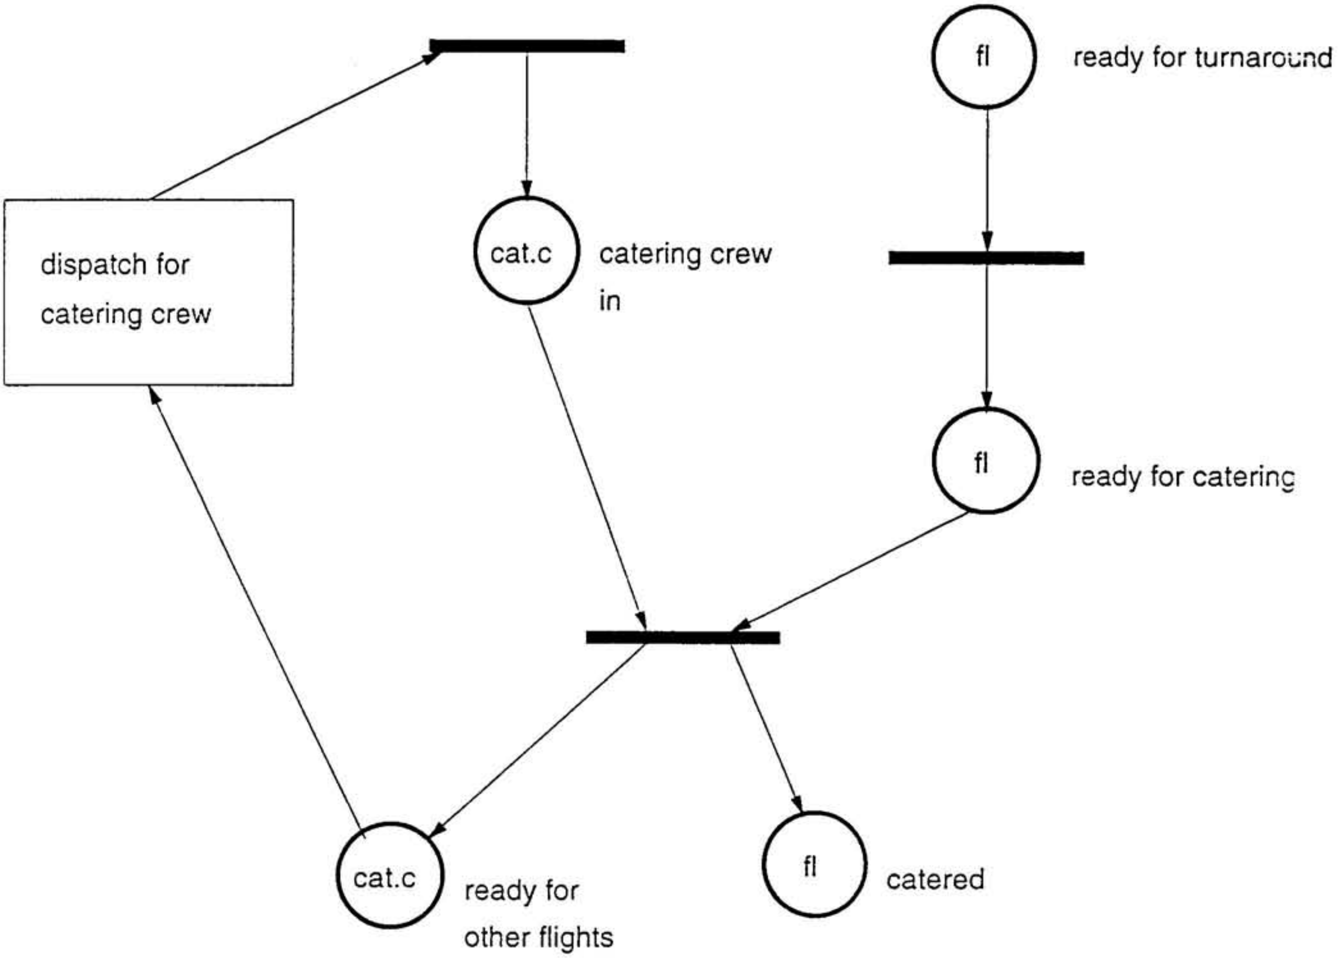

fl = flight
cat.c = catering crew

Figure B6: Cleaning Process (Preliminary)

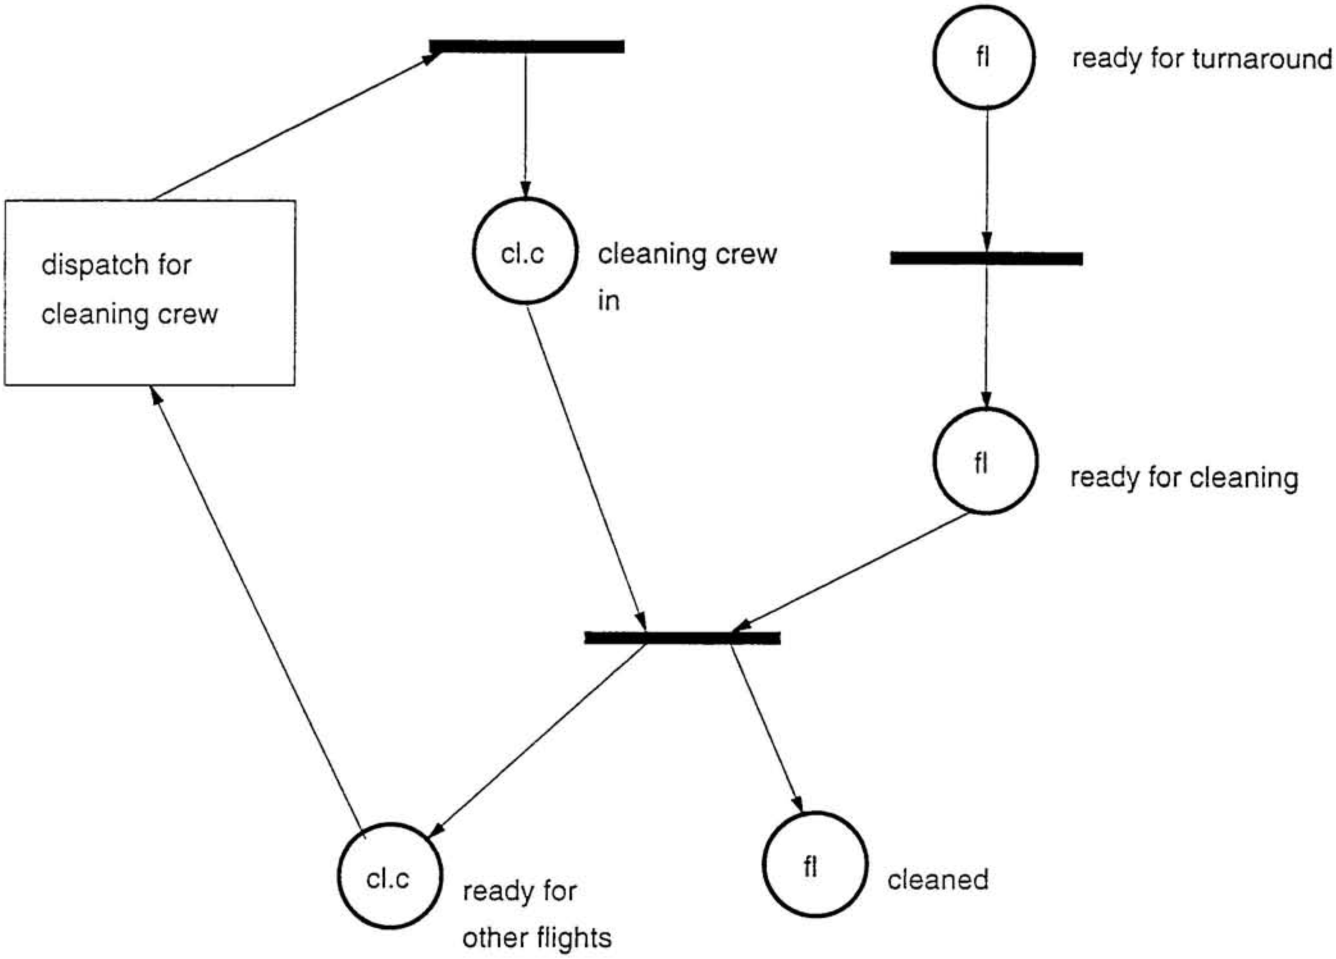

fl = flight
cat.c = catering crew

Figure B7 : Walkaround Process (Preliminary)

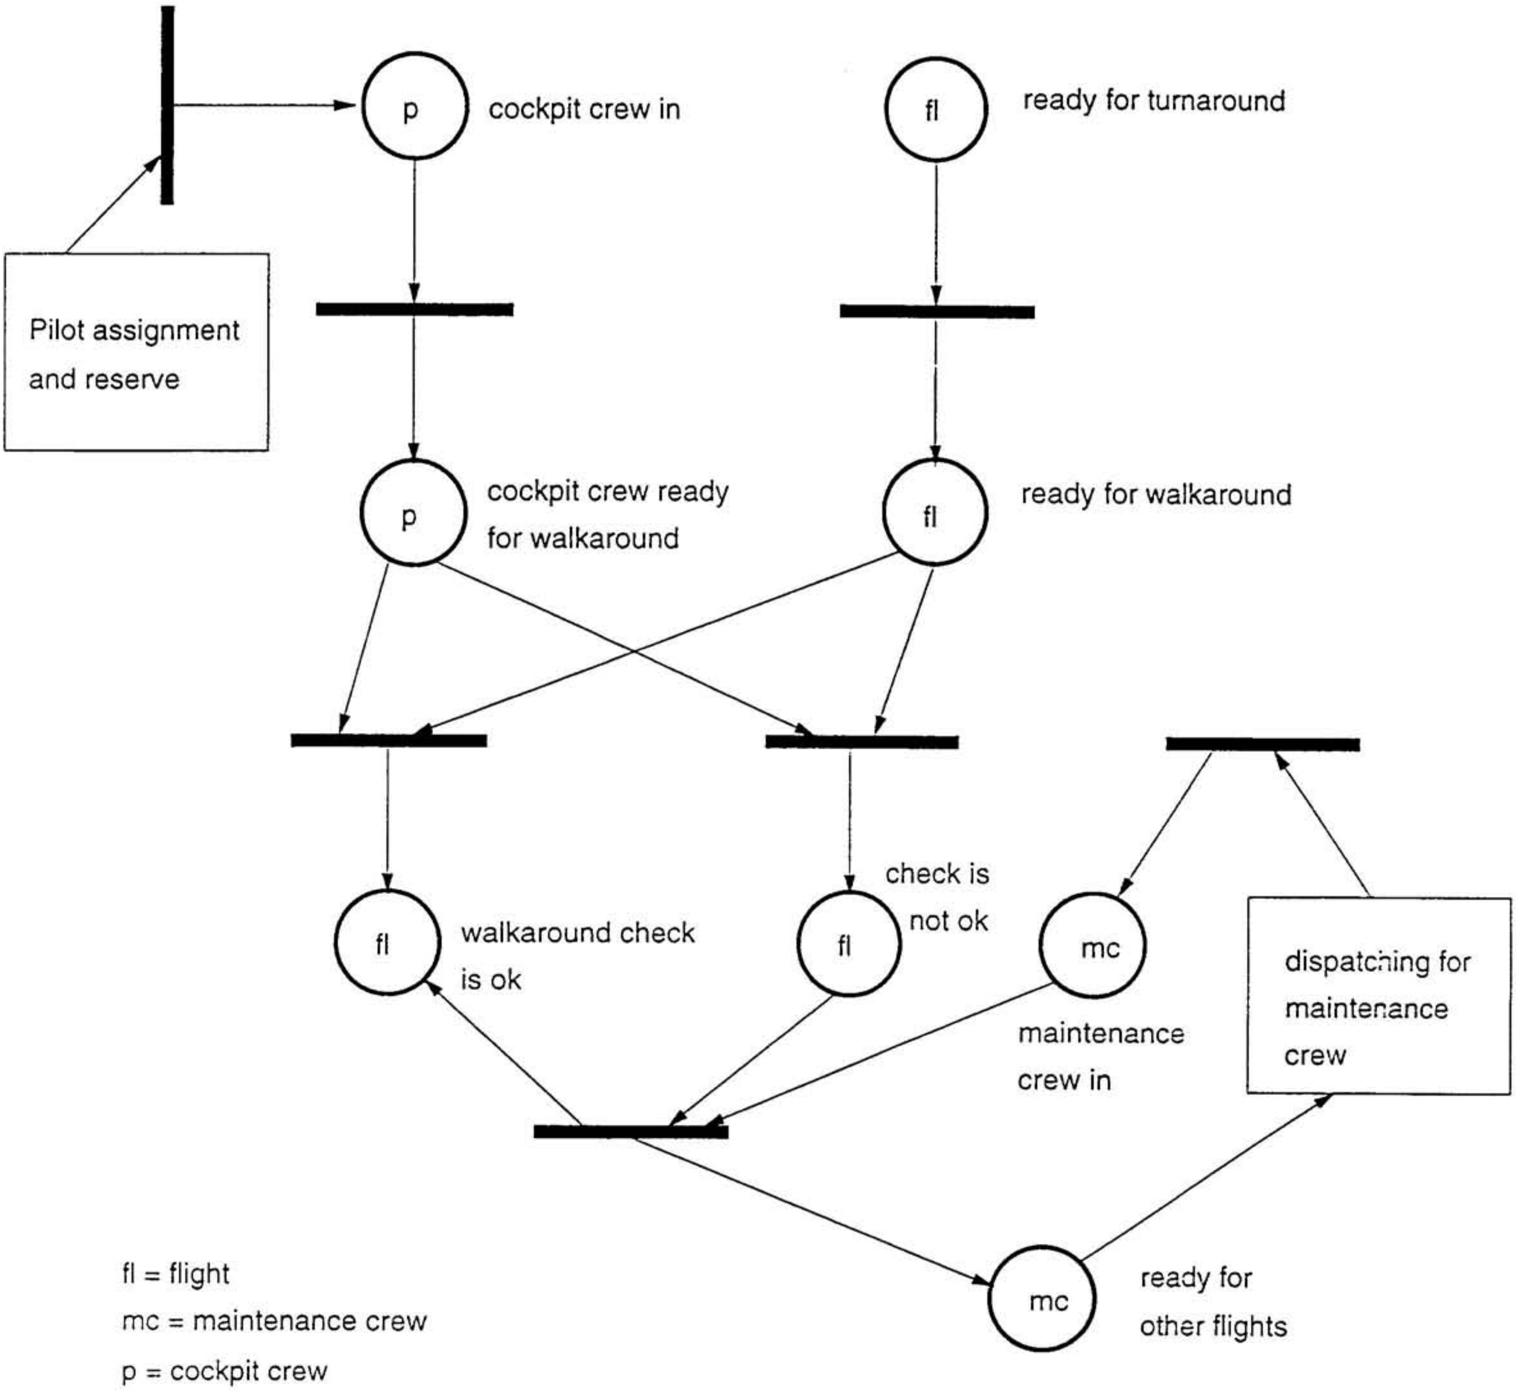

Figure B8 : Cockpit System Check Process (Preliminary)

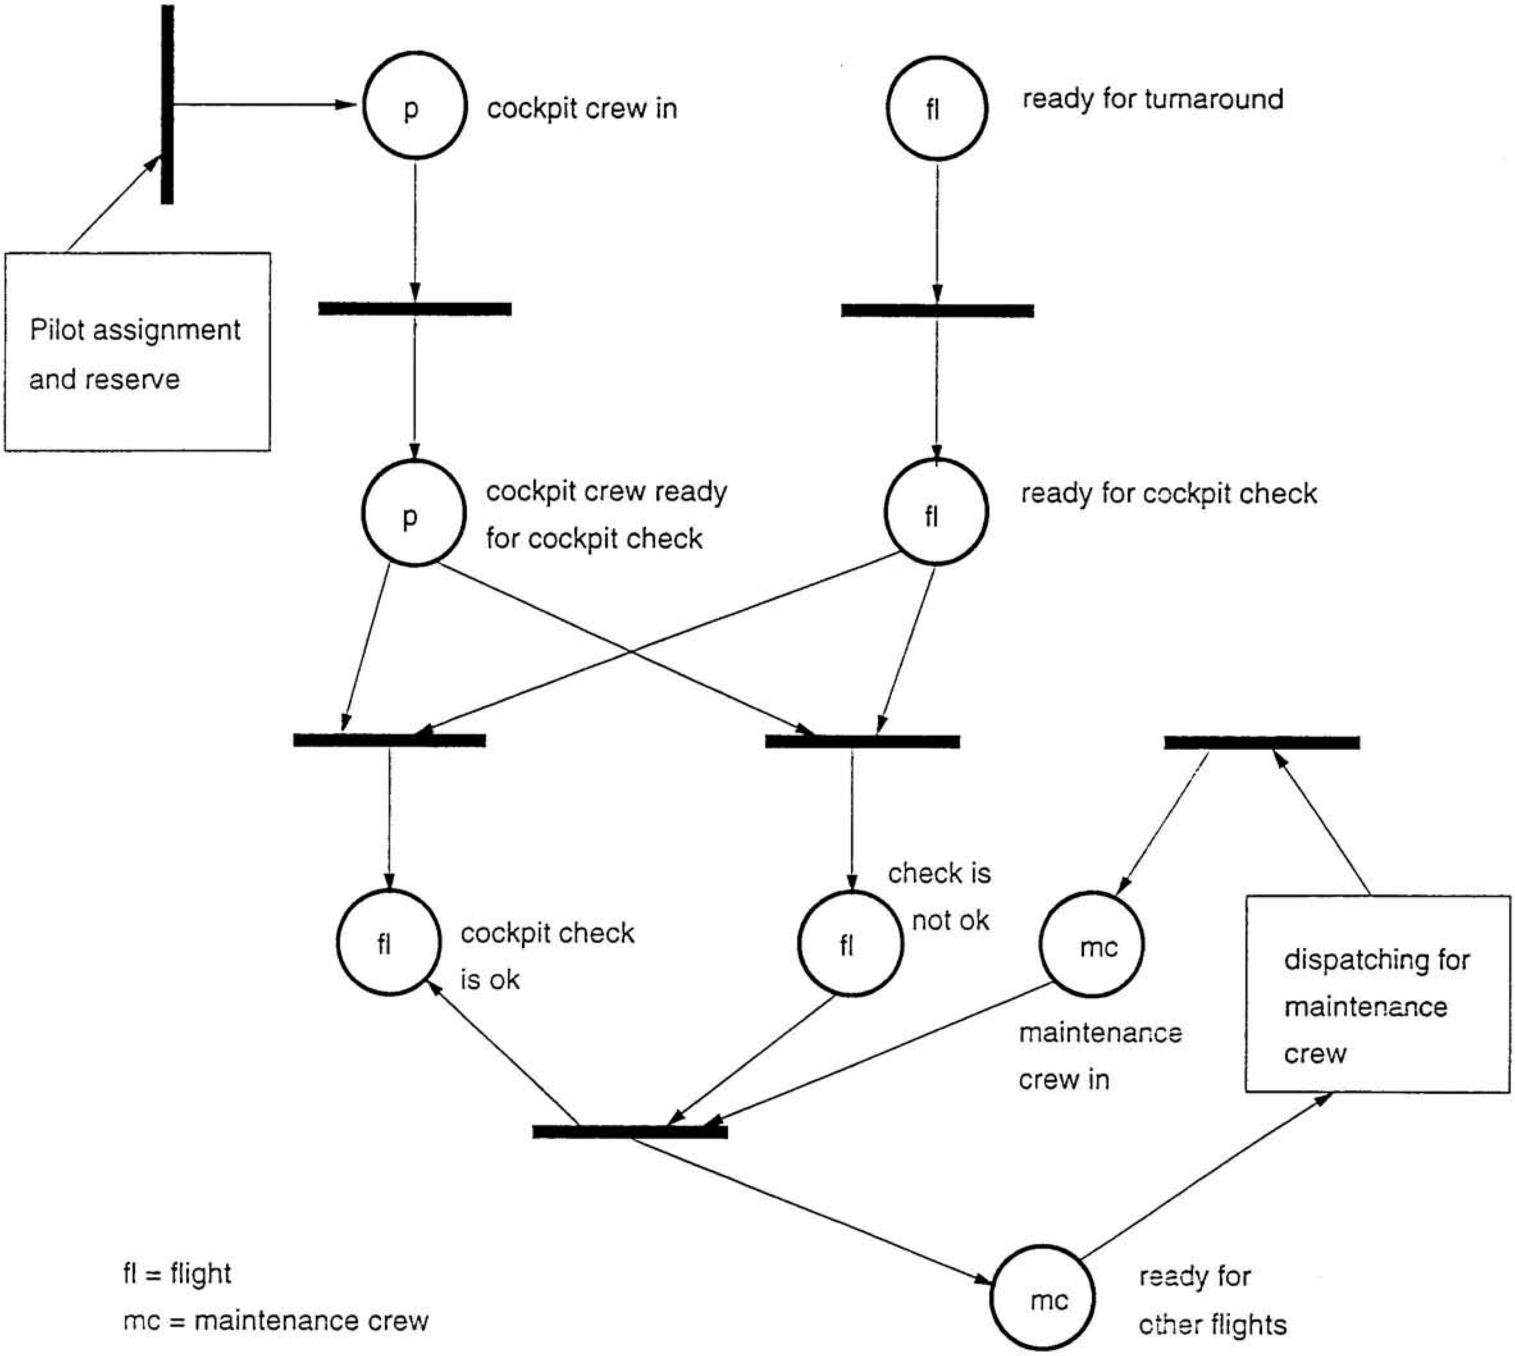

fl = flight
mc = maintenance crew
p = cockpit crew

Figure B9 : Preflight and Cockpit Setup Process (Preliminary)

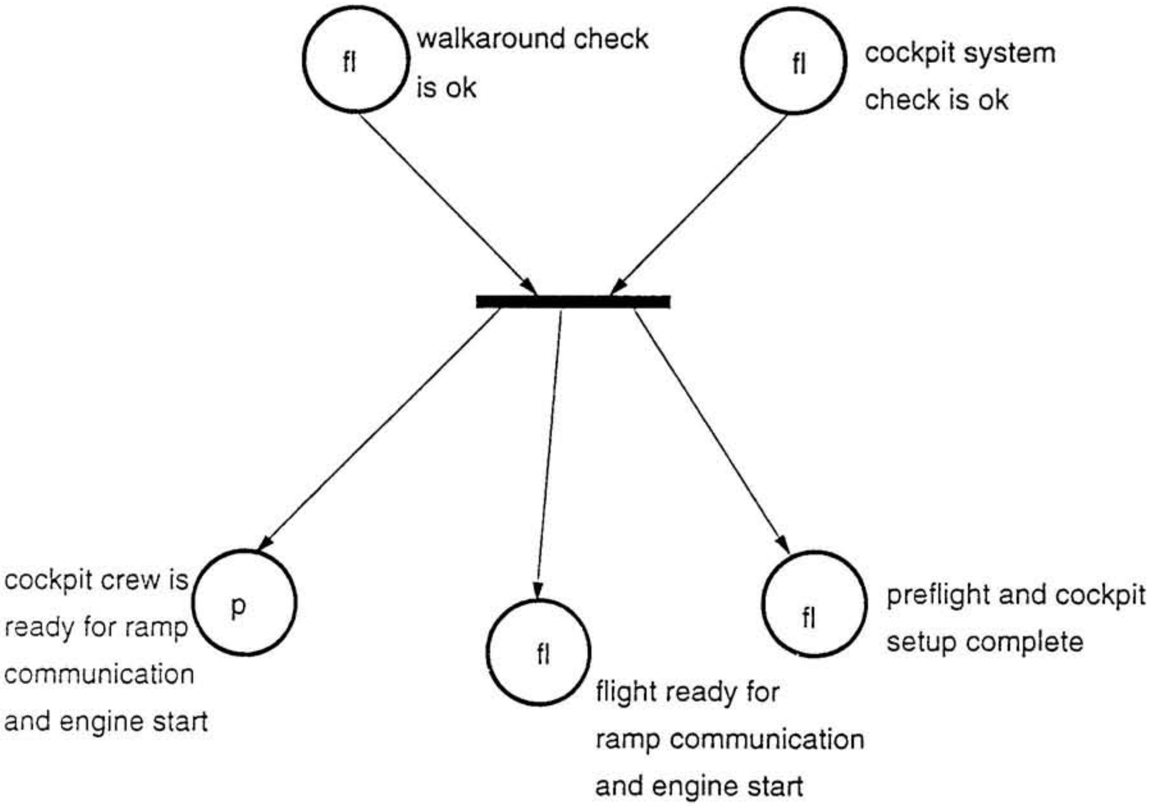

fl = flight

p = cockpit crew

Figure B10 : Passenger Boarding Process (Preliminary)

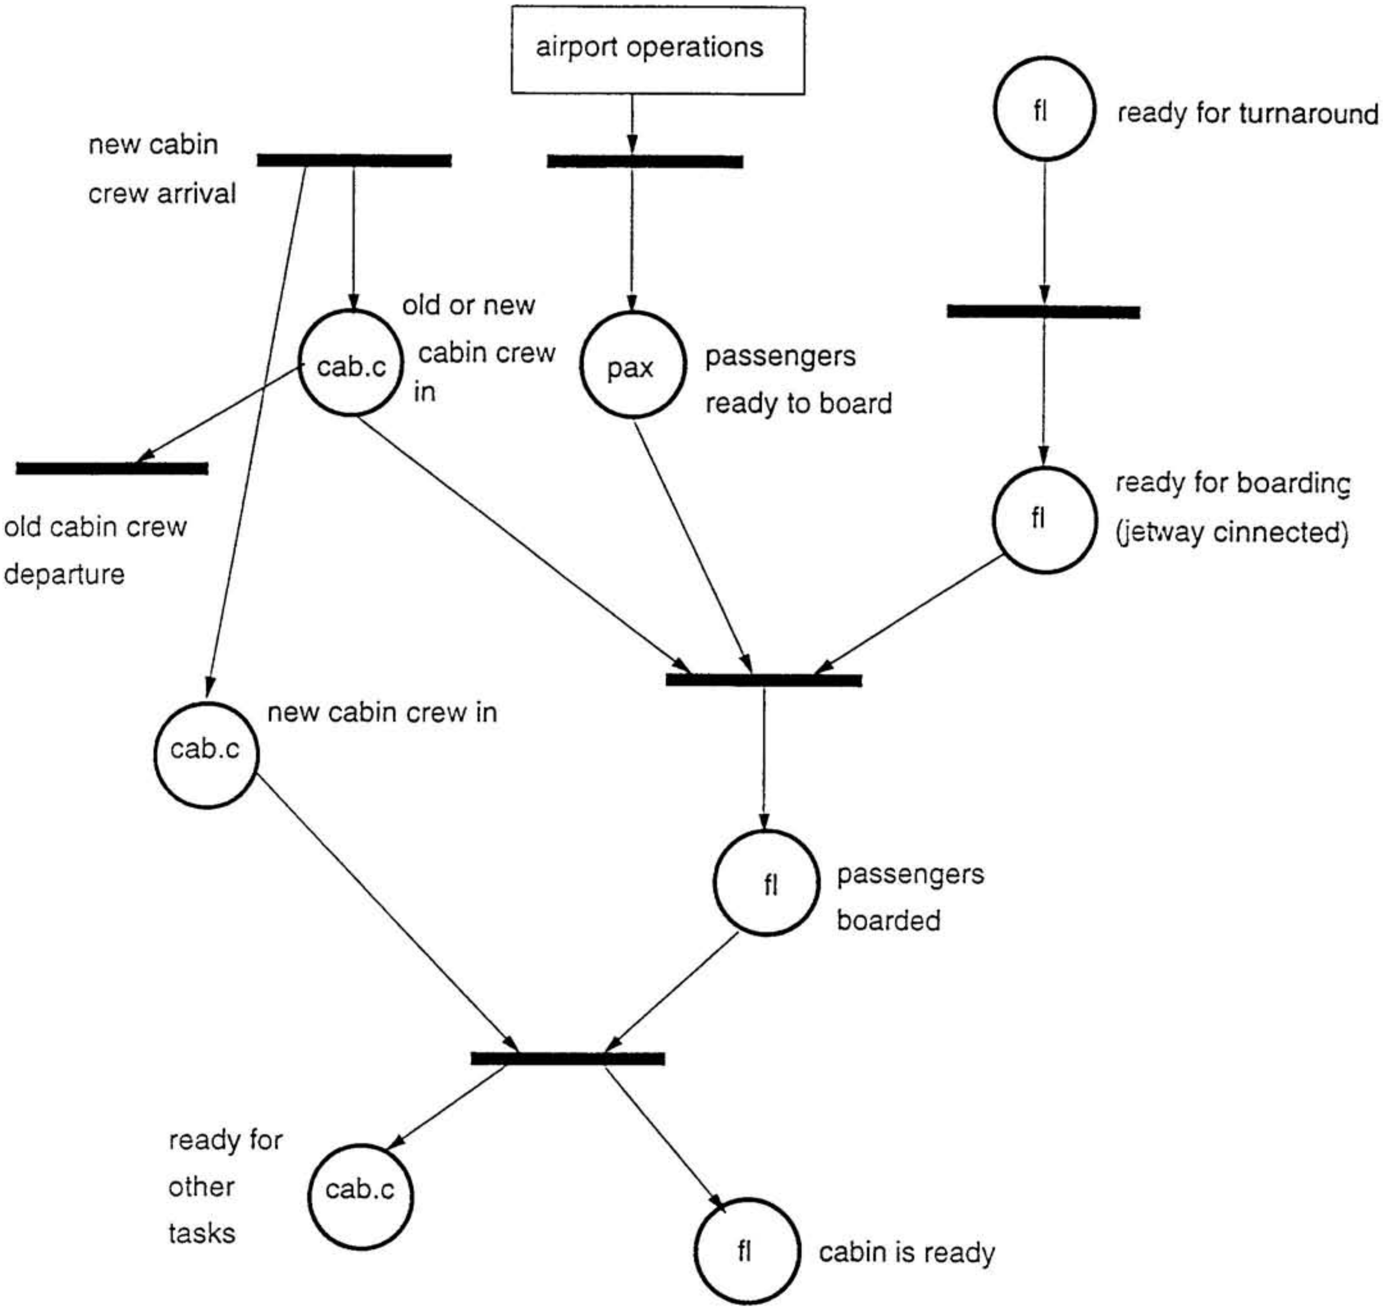

fl = flight

cab.c = cabin crew

pax = passengers

Figure B11 : Baggage and Freight Loading Process (Preliminary)

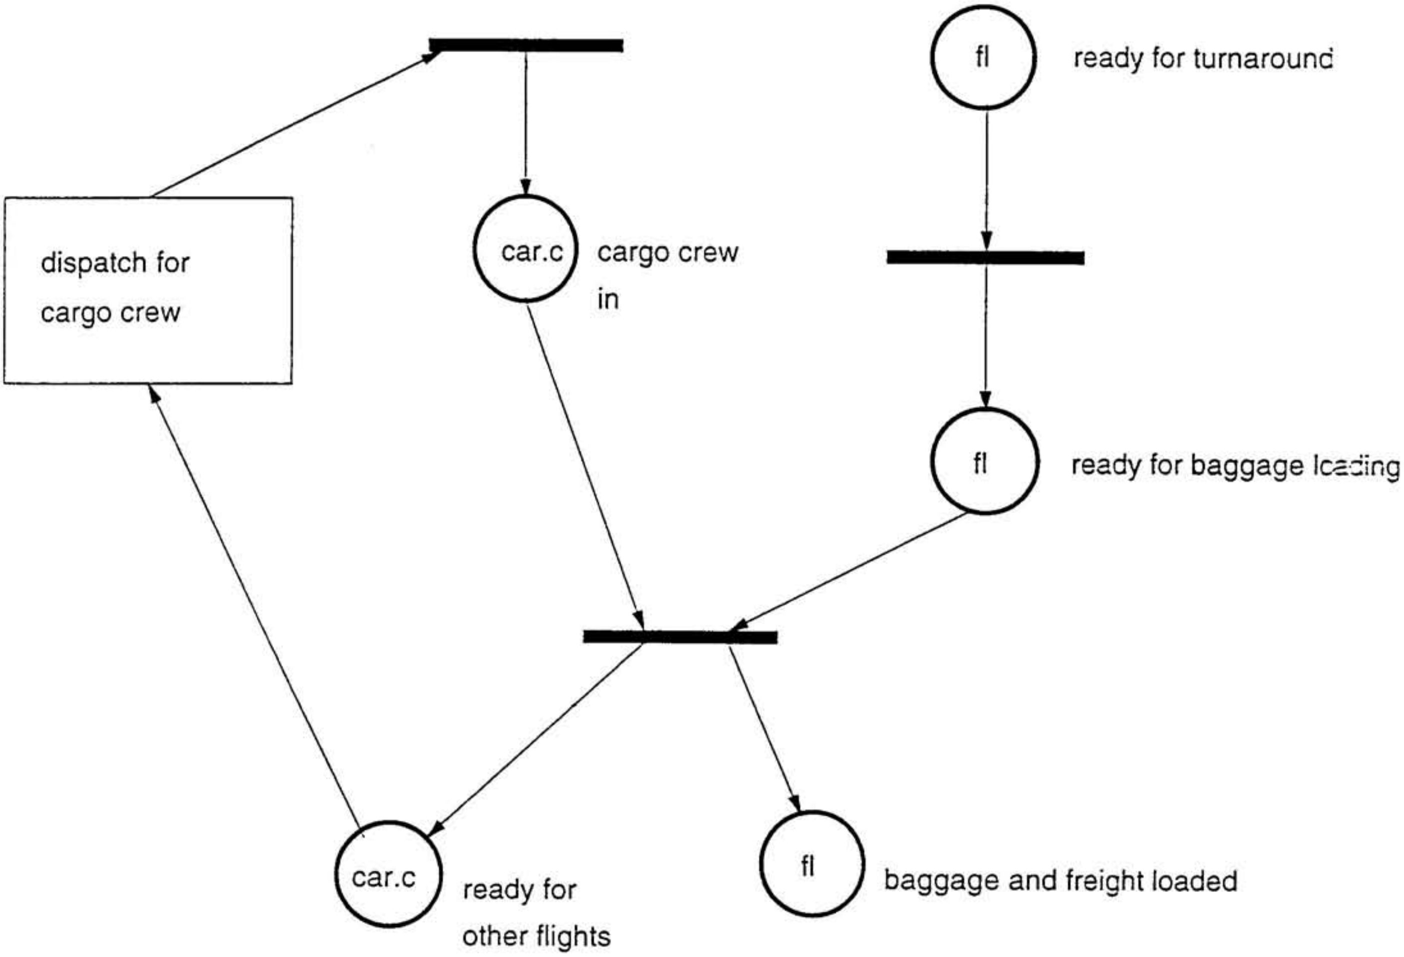

fl = flight
car.c = cargo crew

Figure 12: Aircraft Preperation Process (Preliminary)

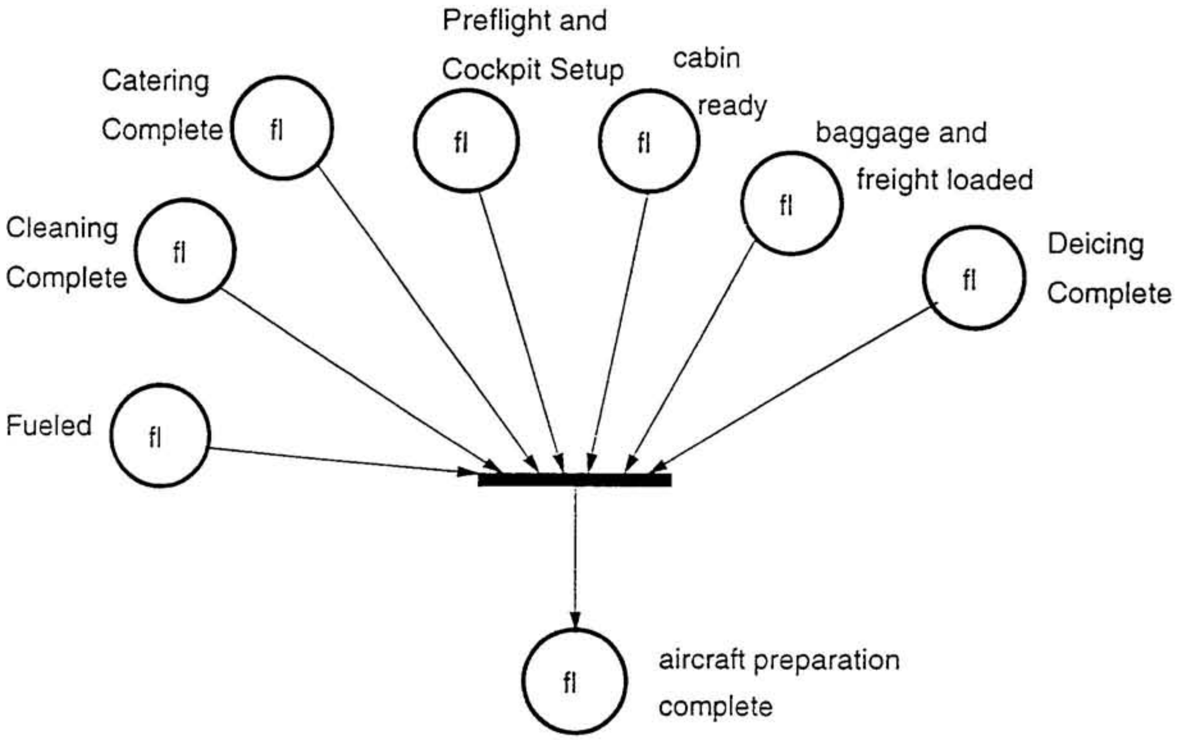

fl = flight

Figure B13 : Flight planning Process (Preliminary)

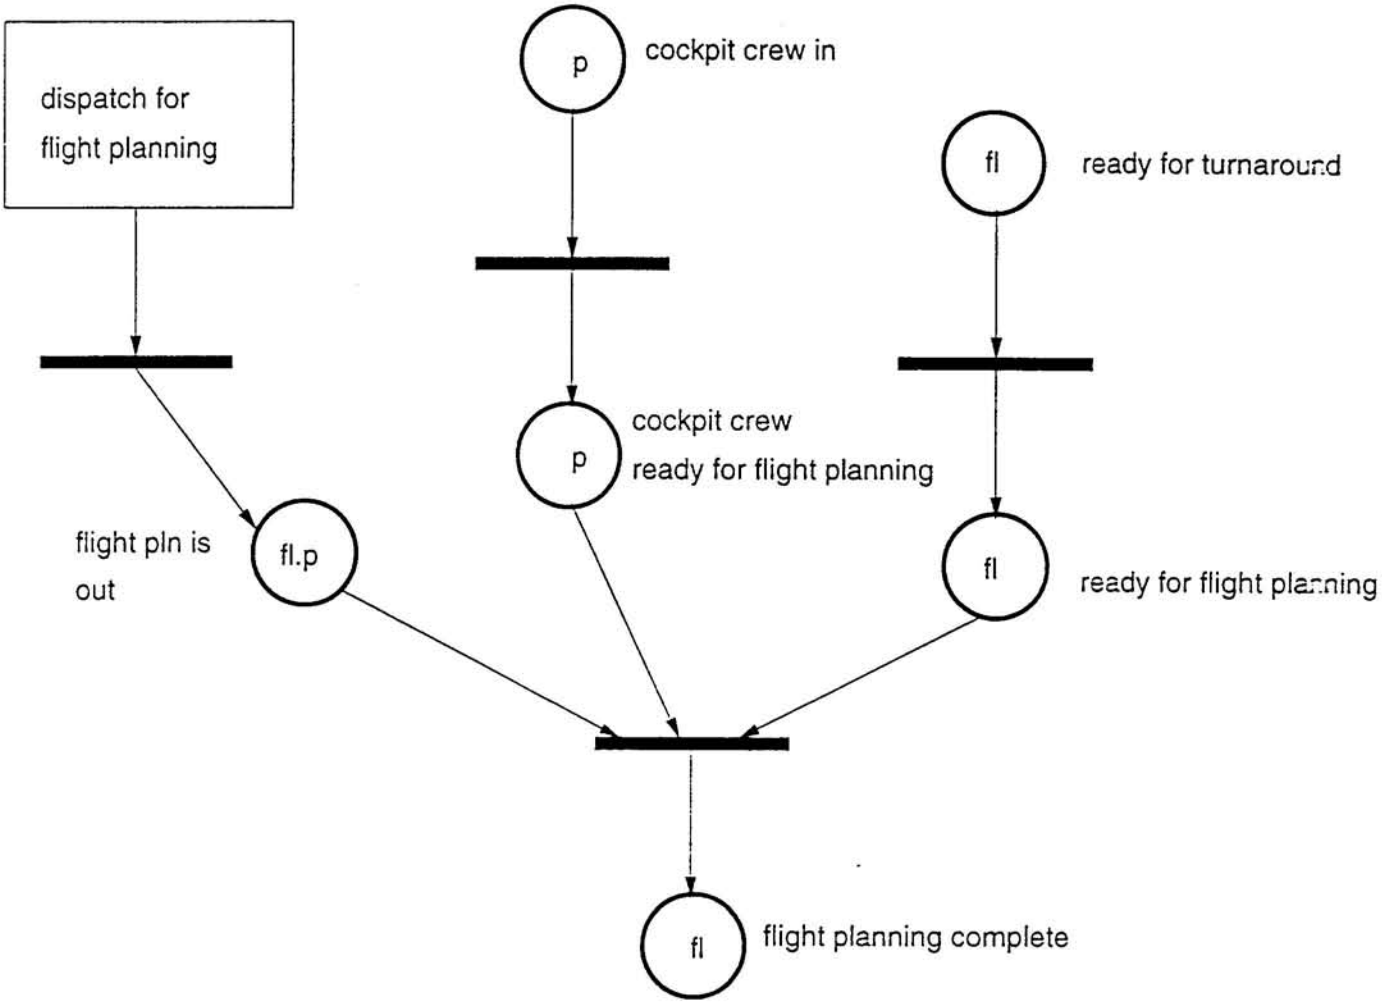

fl = flight
p = cockpit crew
fl.p = flight plan

Figure B14 : Engine Start Process (Preliminary)

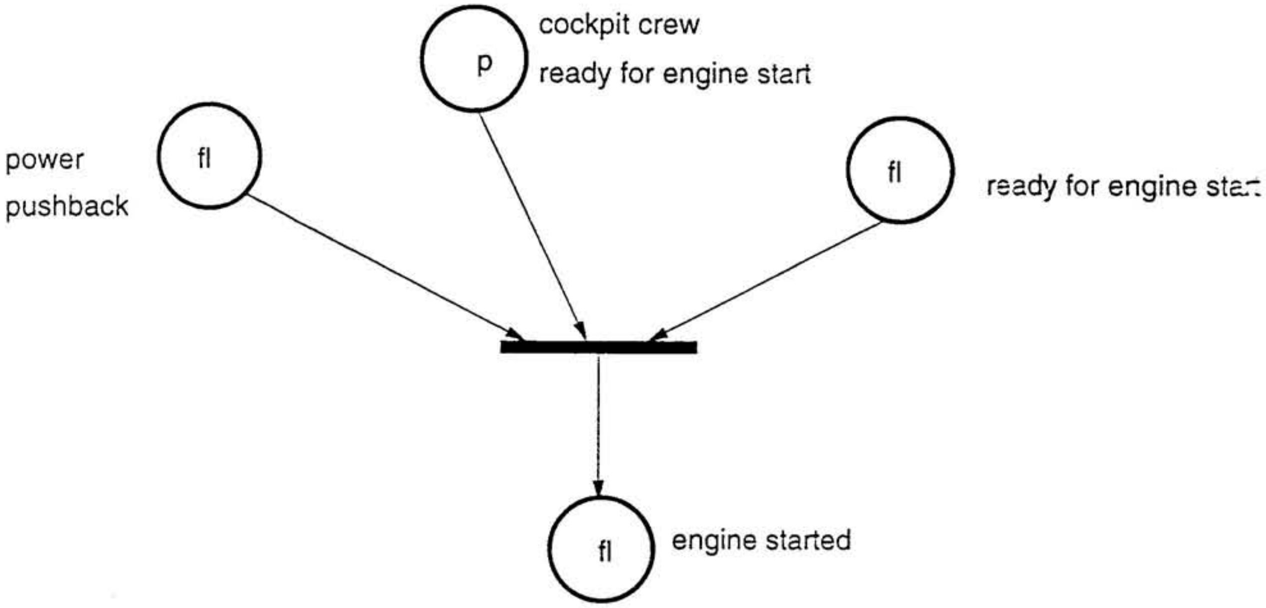

fl = flight
p = cockpit crew

Figure B15 : Tug Connection Process (Preliminary)

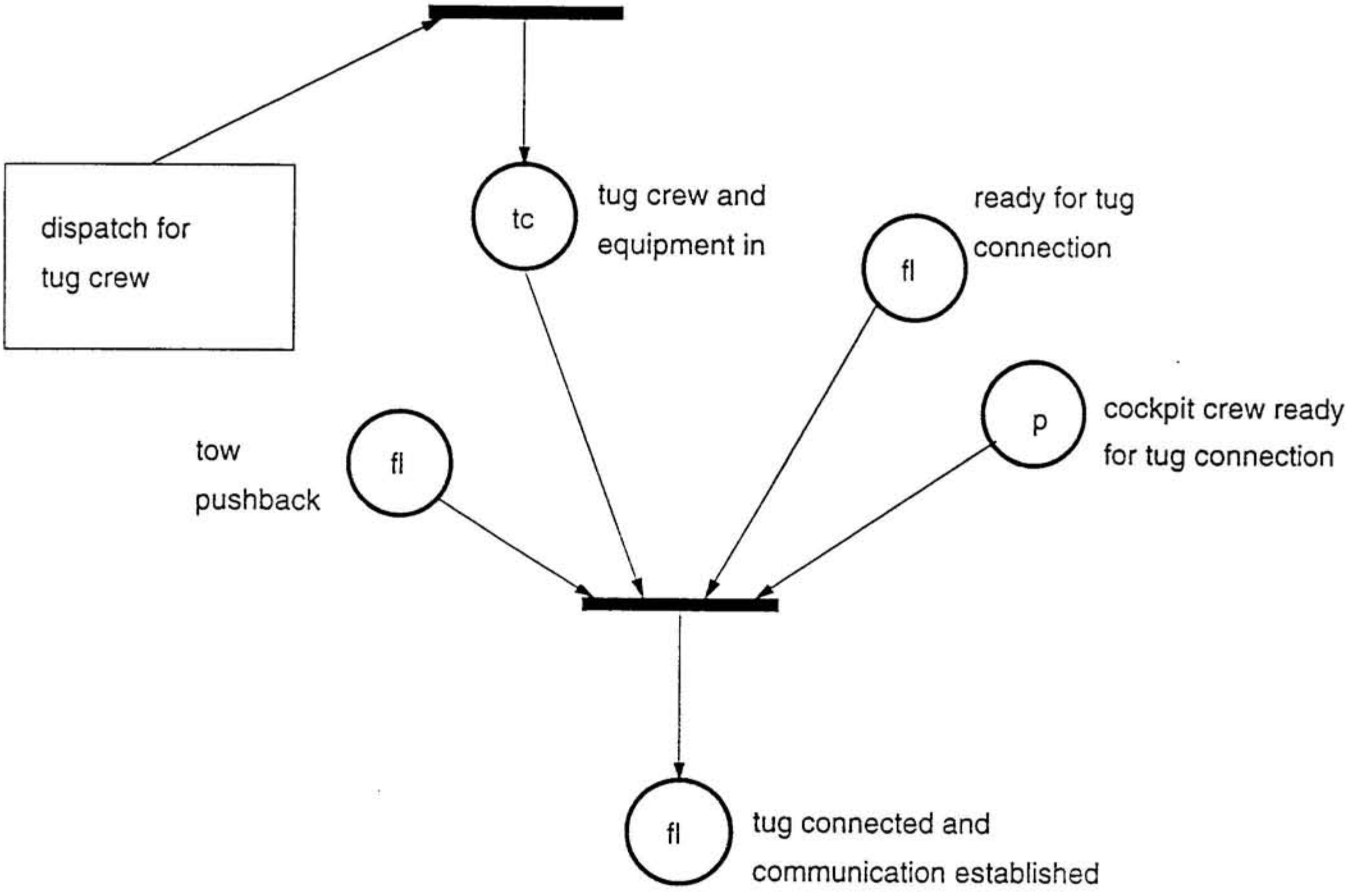

fl = flight

tc = tug crew

p = cockpit crew

Figure B16 : Ramp Control Communication Process (Preliminary)

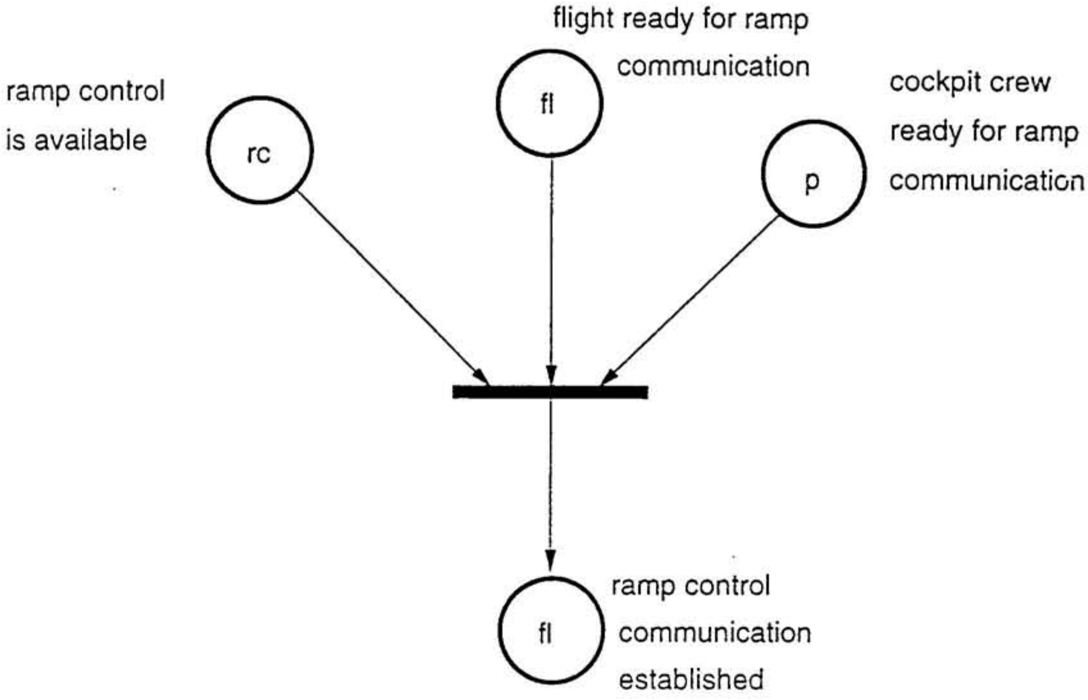

fl = flight

p = cockpit crew

rc = ramp control

Figure B17: Ready For Pushback Process (Preliminary)

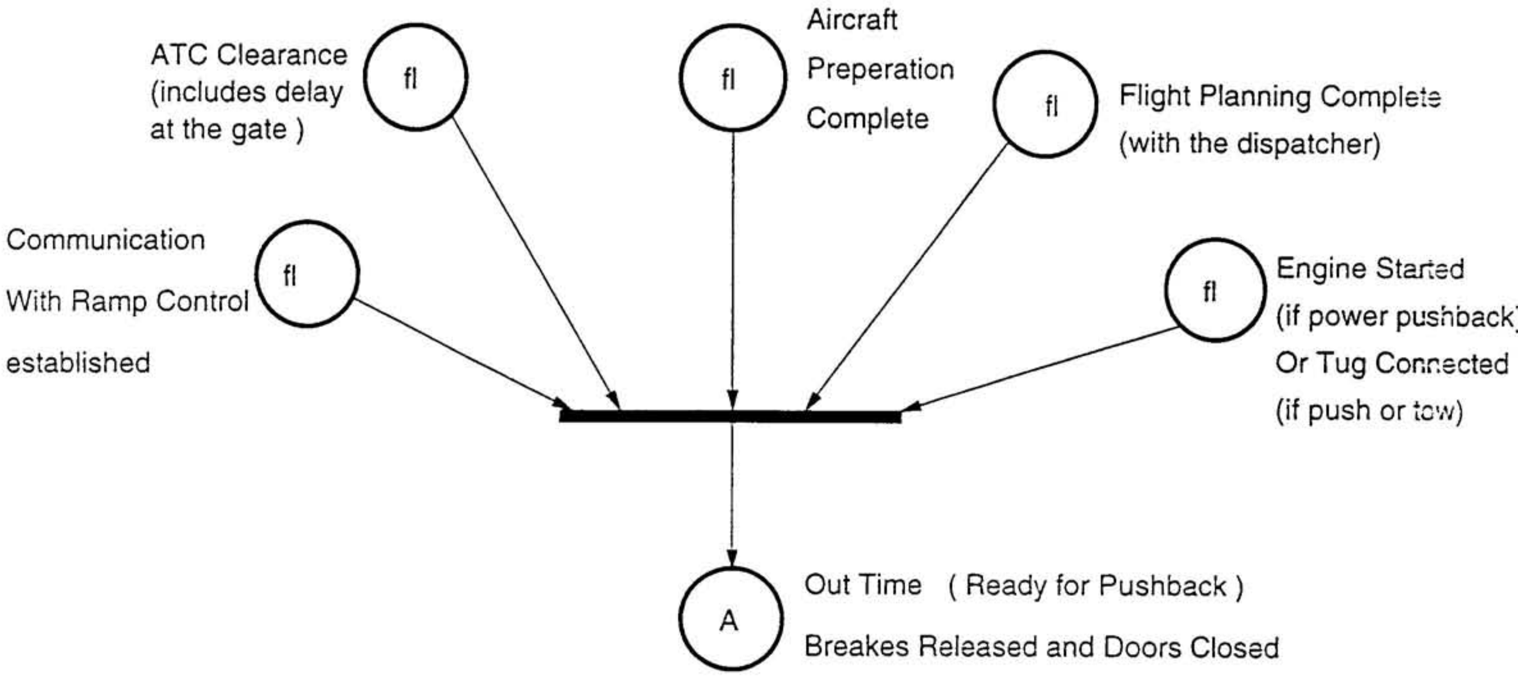

fl = flight

A = indicates connection to the next Petri Net

Figure B18: Pushback Process (Preliminary)

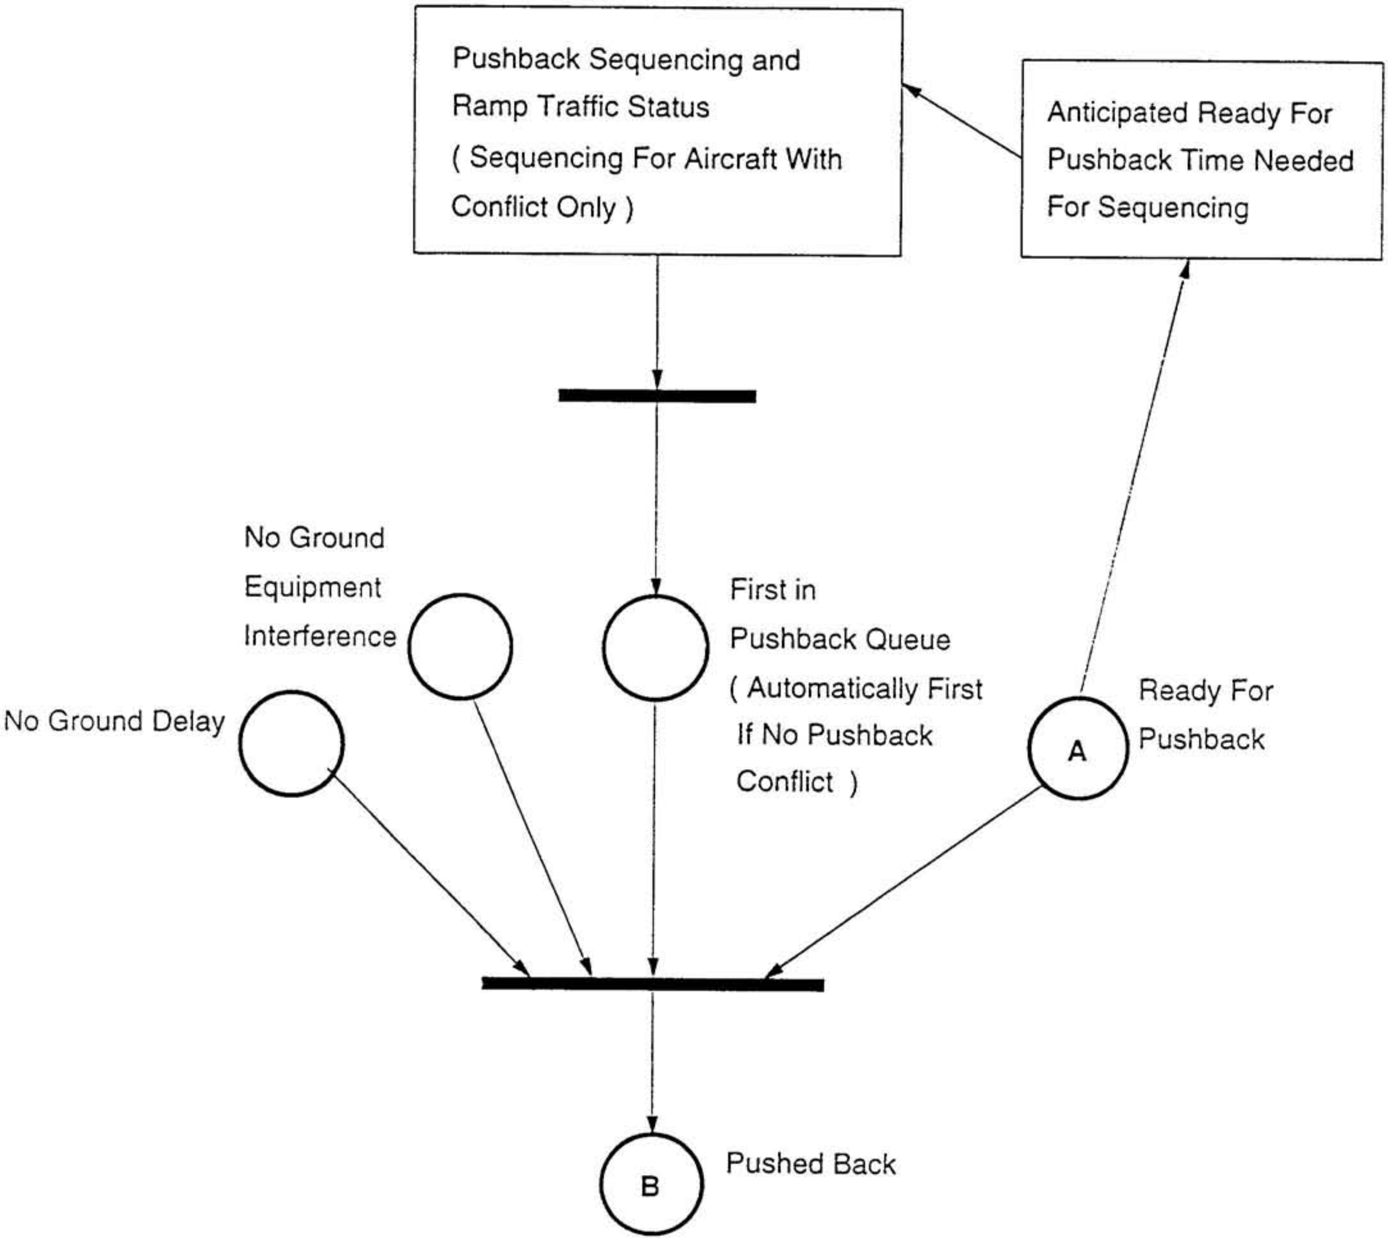

Figure B19: Ramp Exit Process (Preliminary)

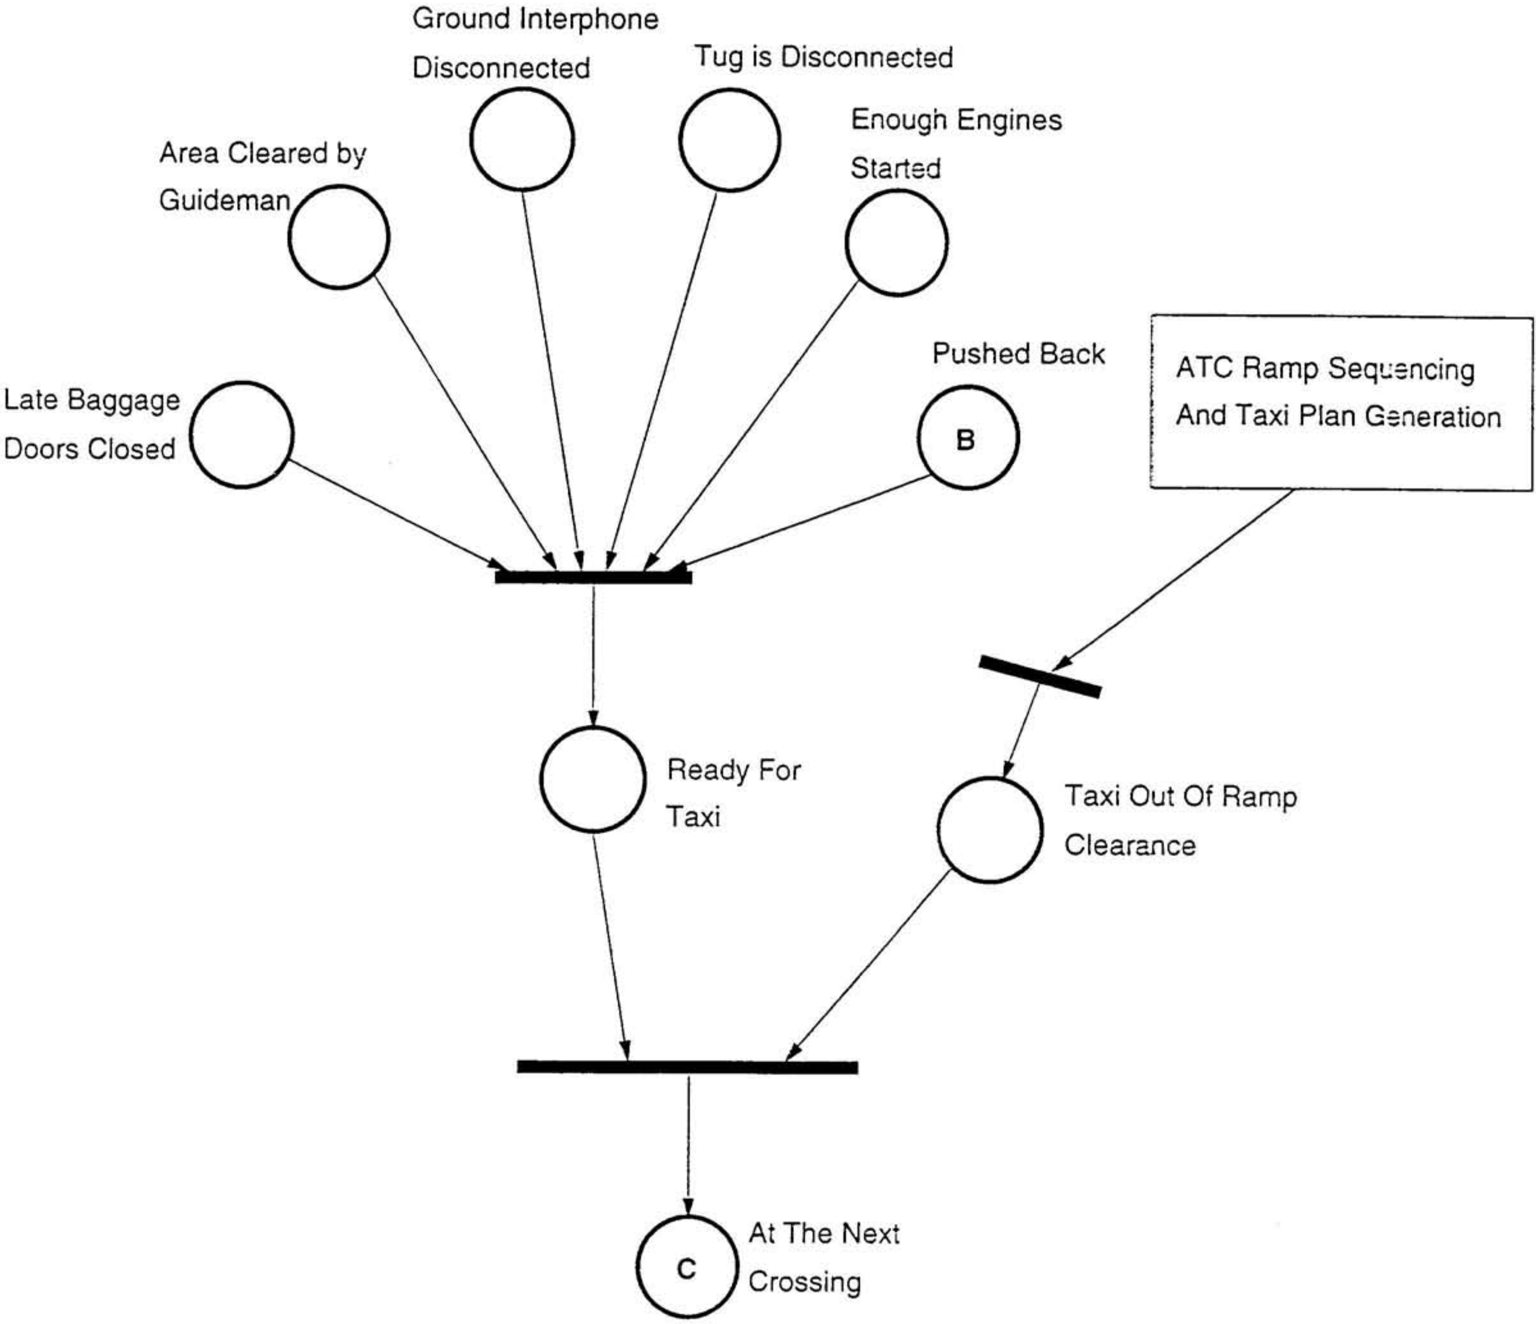

Figure B20: Taxiway Or Runway Crossing Process (Preliminary)

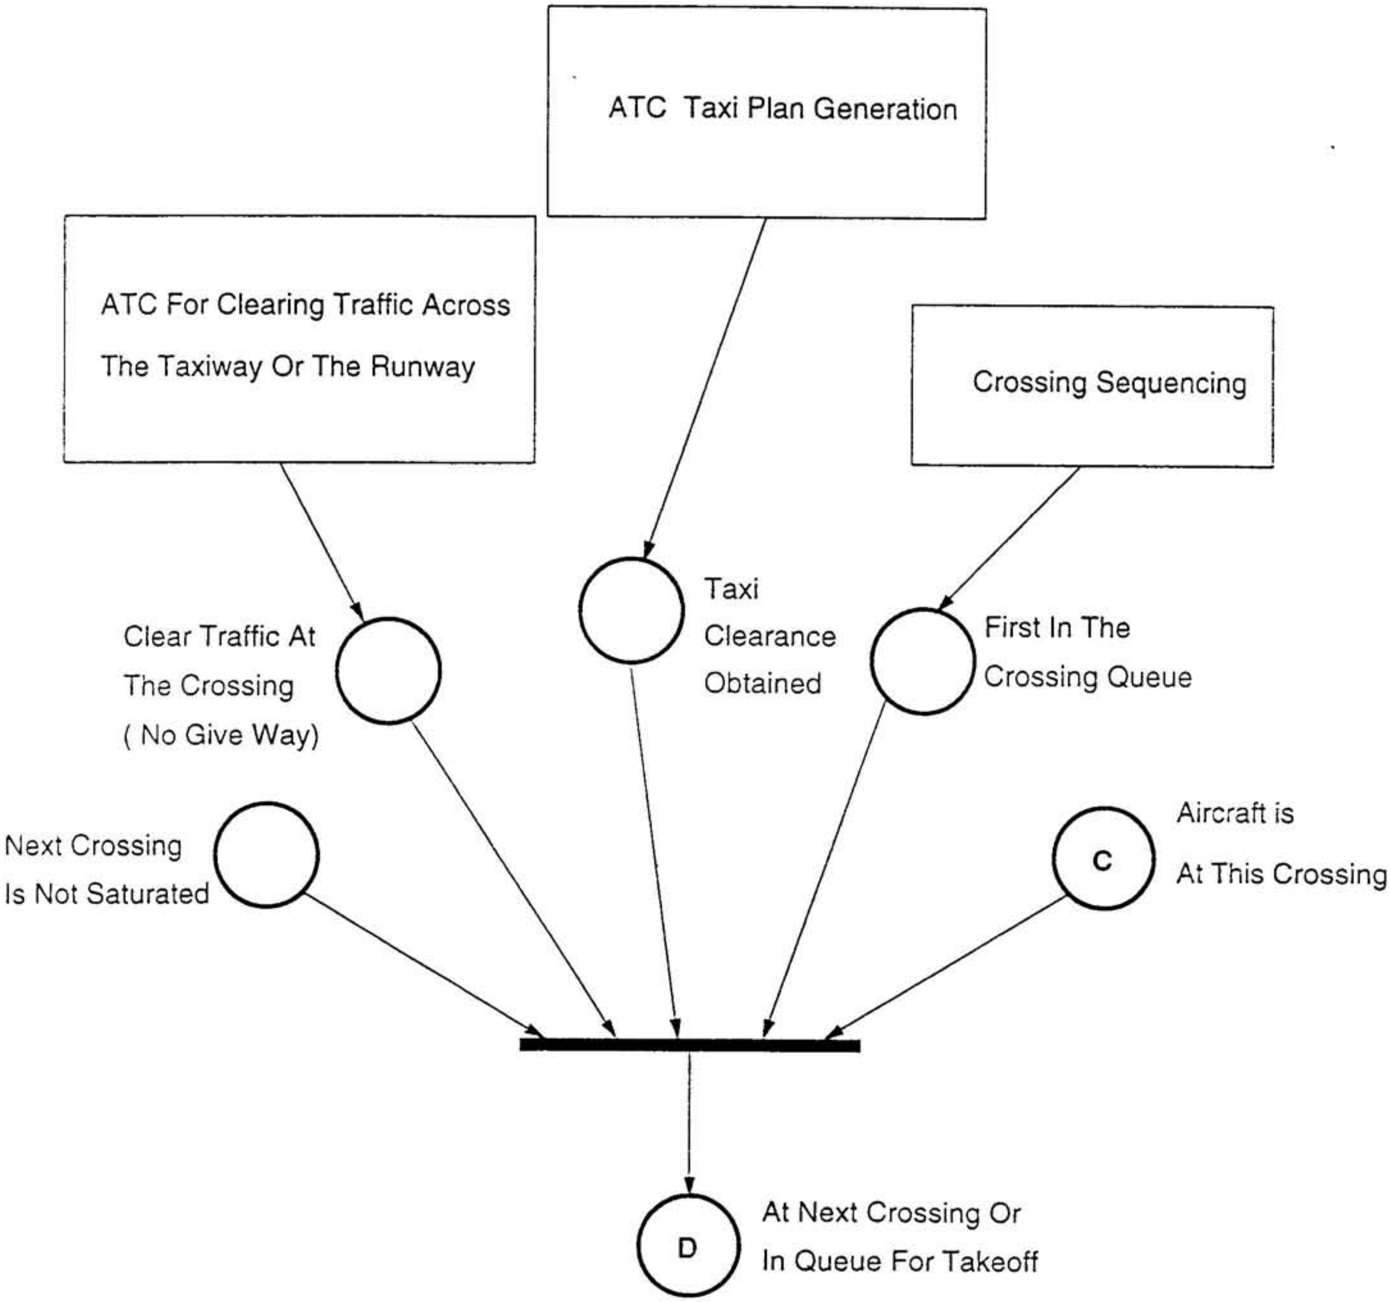

Figure B21: Queuing For Takeoff Process (Preliminary)

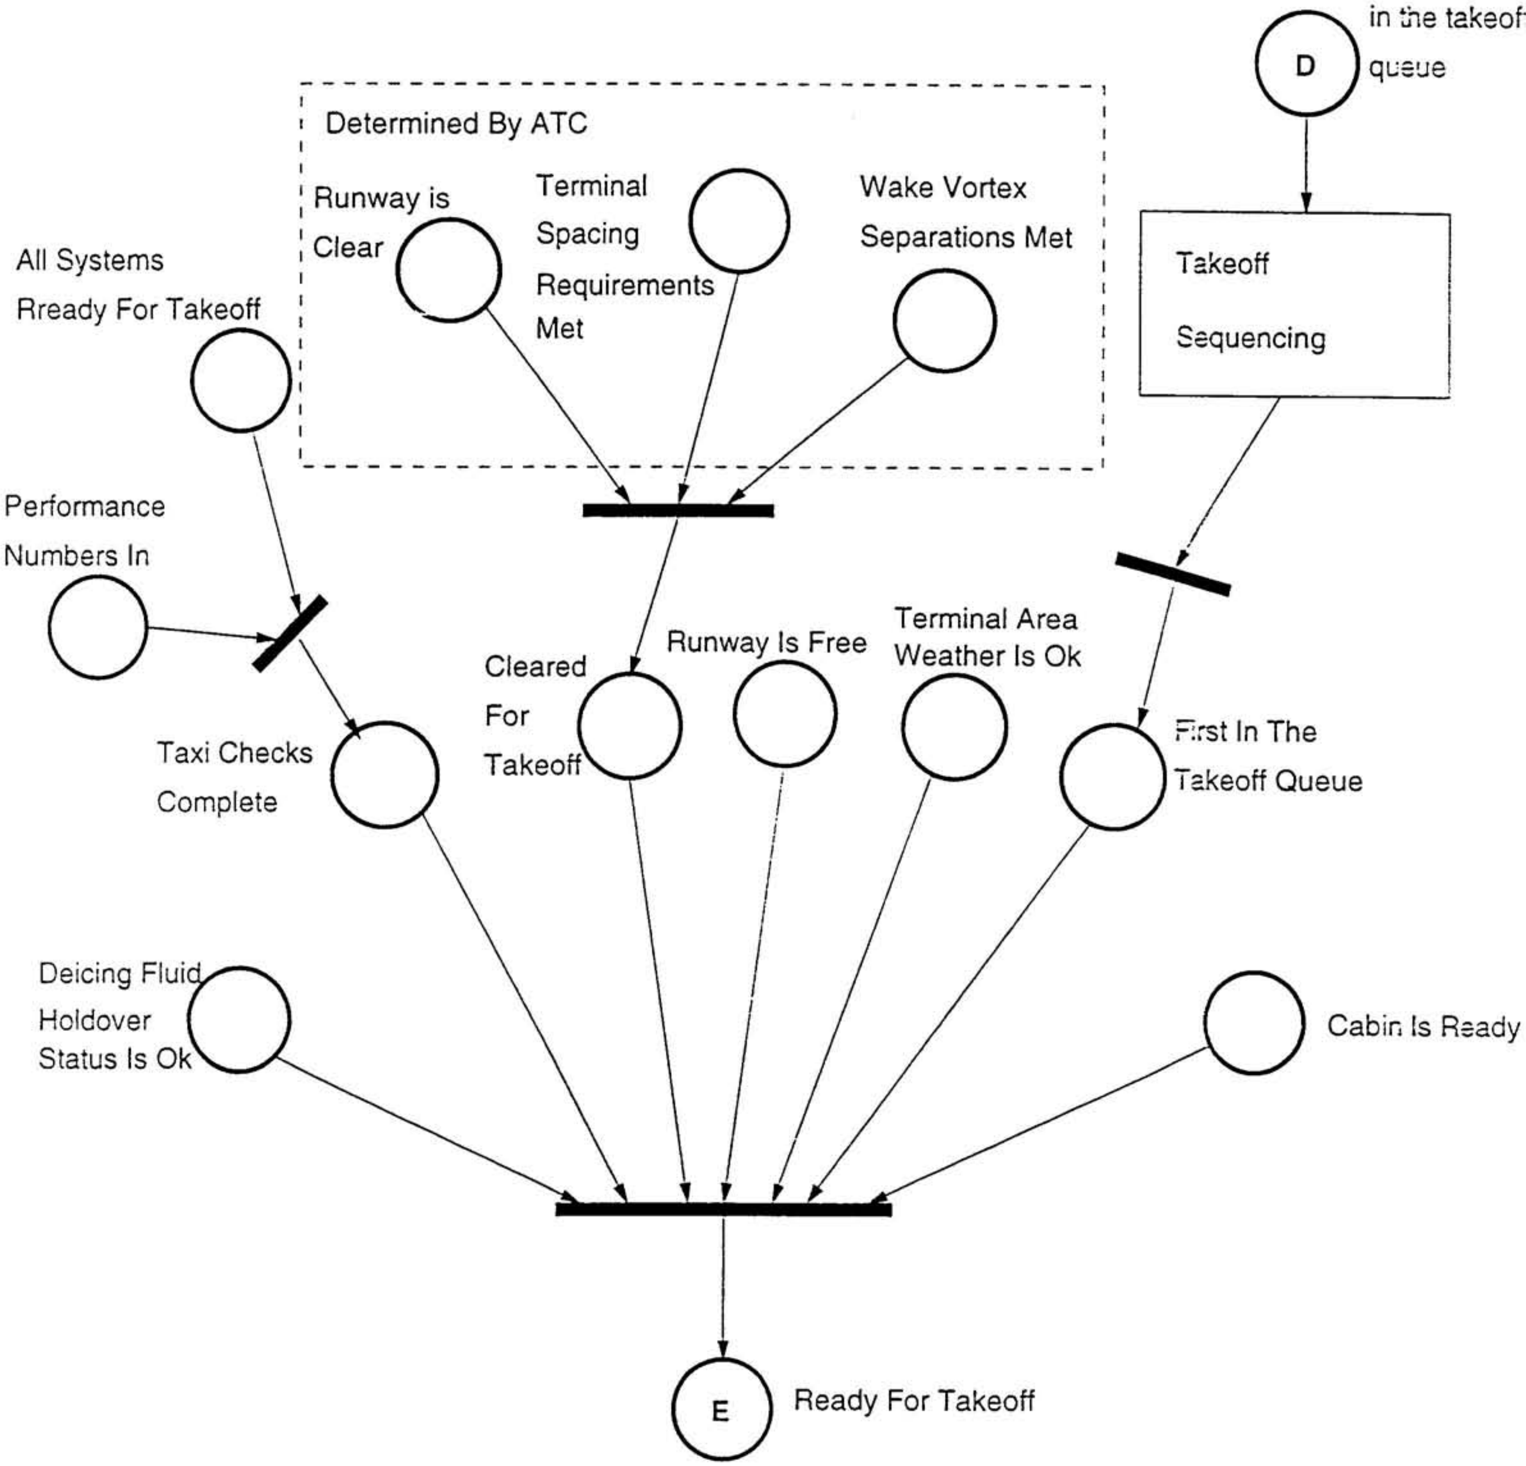

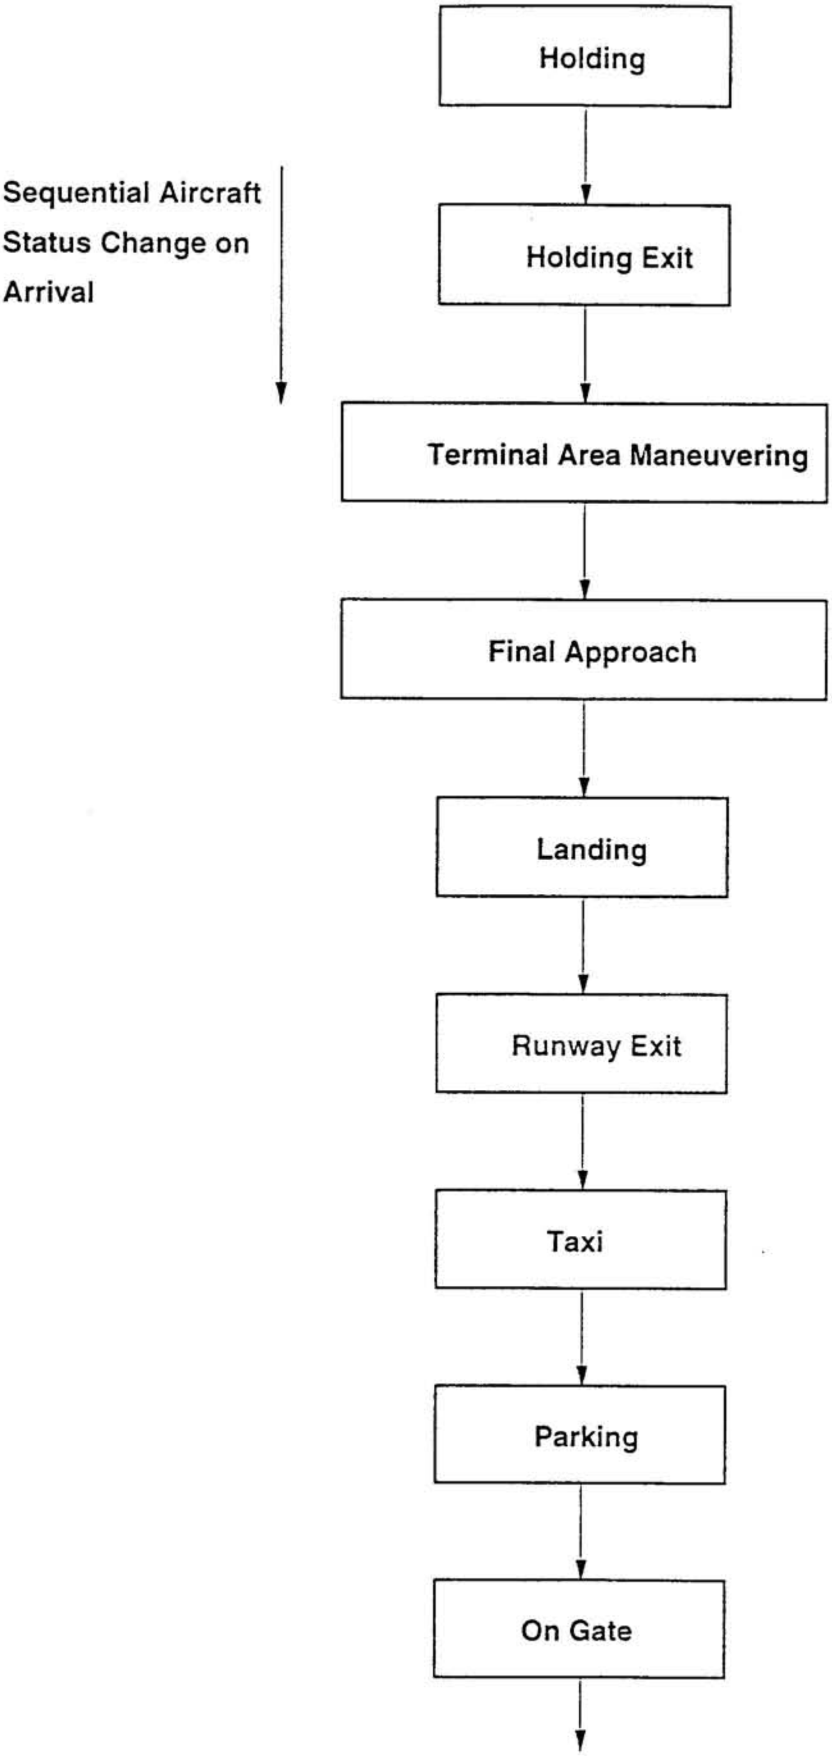
Holding
Sequential Aircraft Status Change on Arrival
Holding Exit
Terminal Area Maneuvering
Final Approach
Landing
Runway Exit
Taxi
Parking
On Gate

Figure B23: Landing Process (Preliminary)

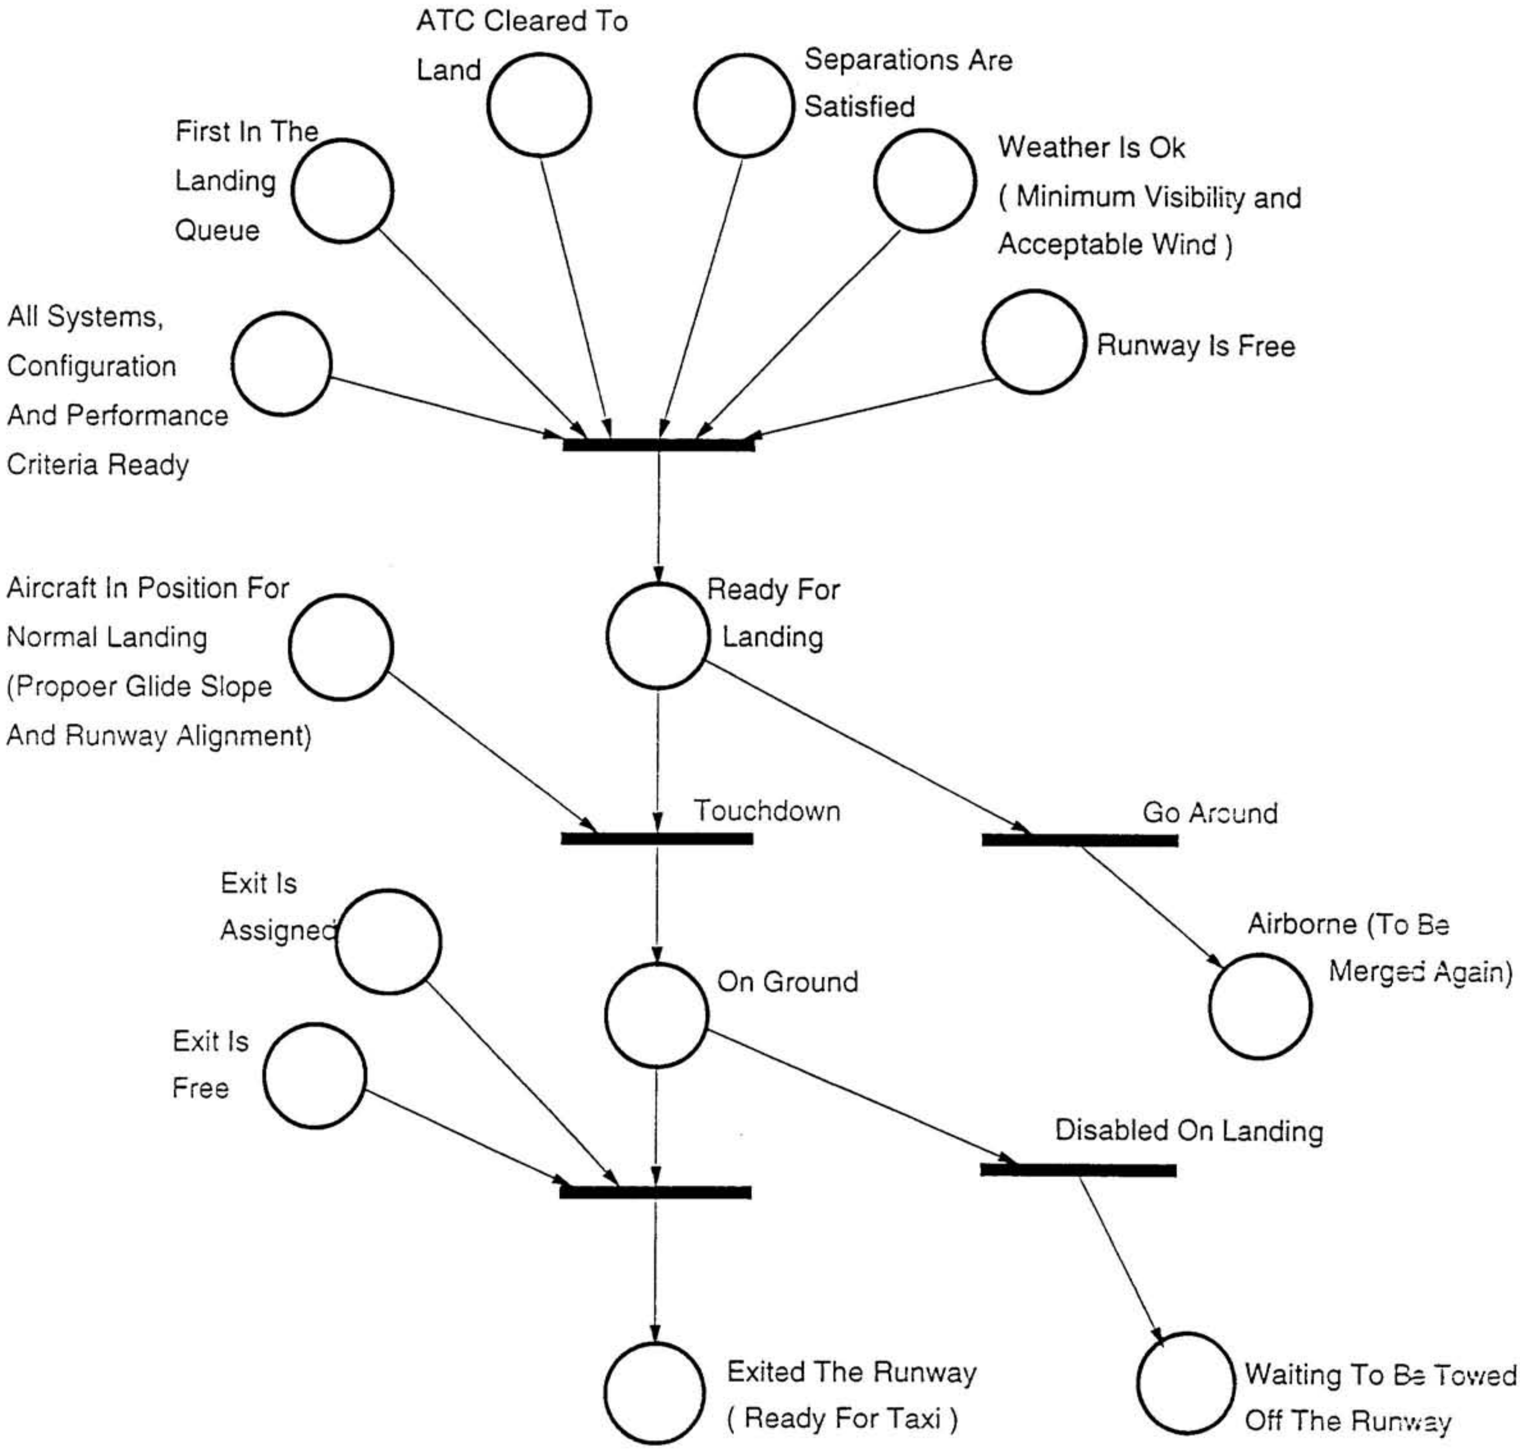

Figure B24: Parking Process (Preliminary)

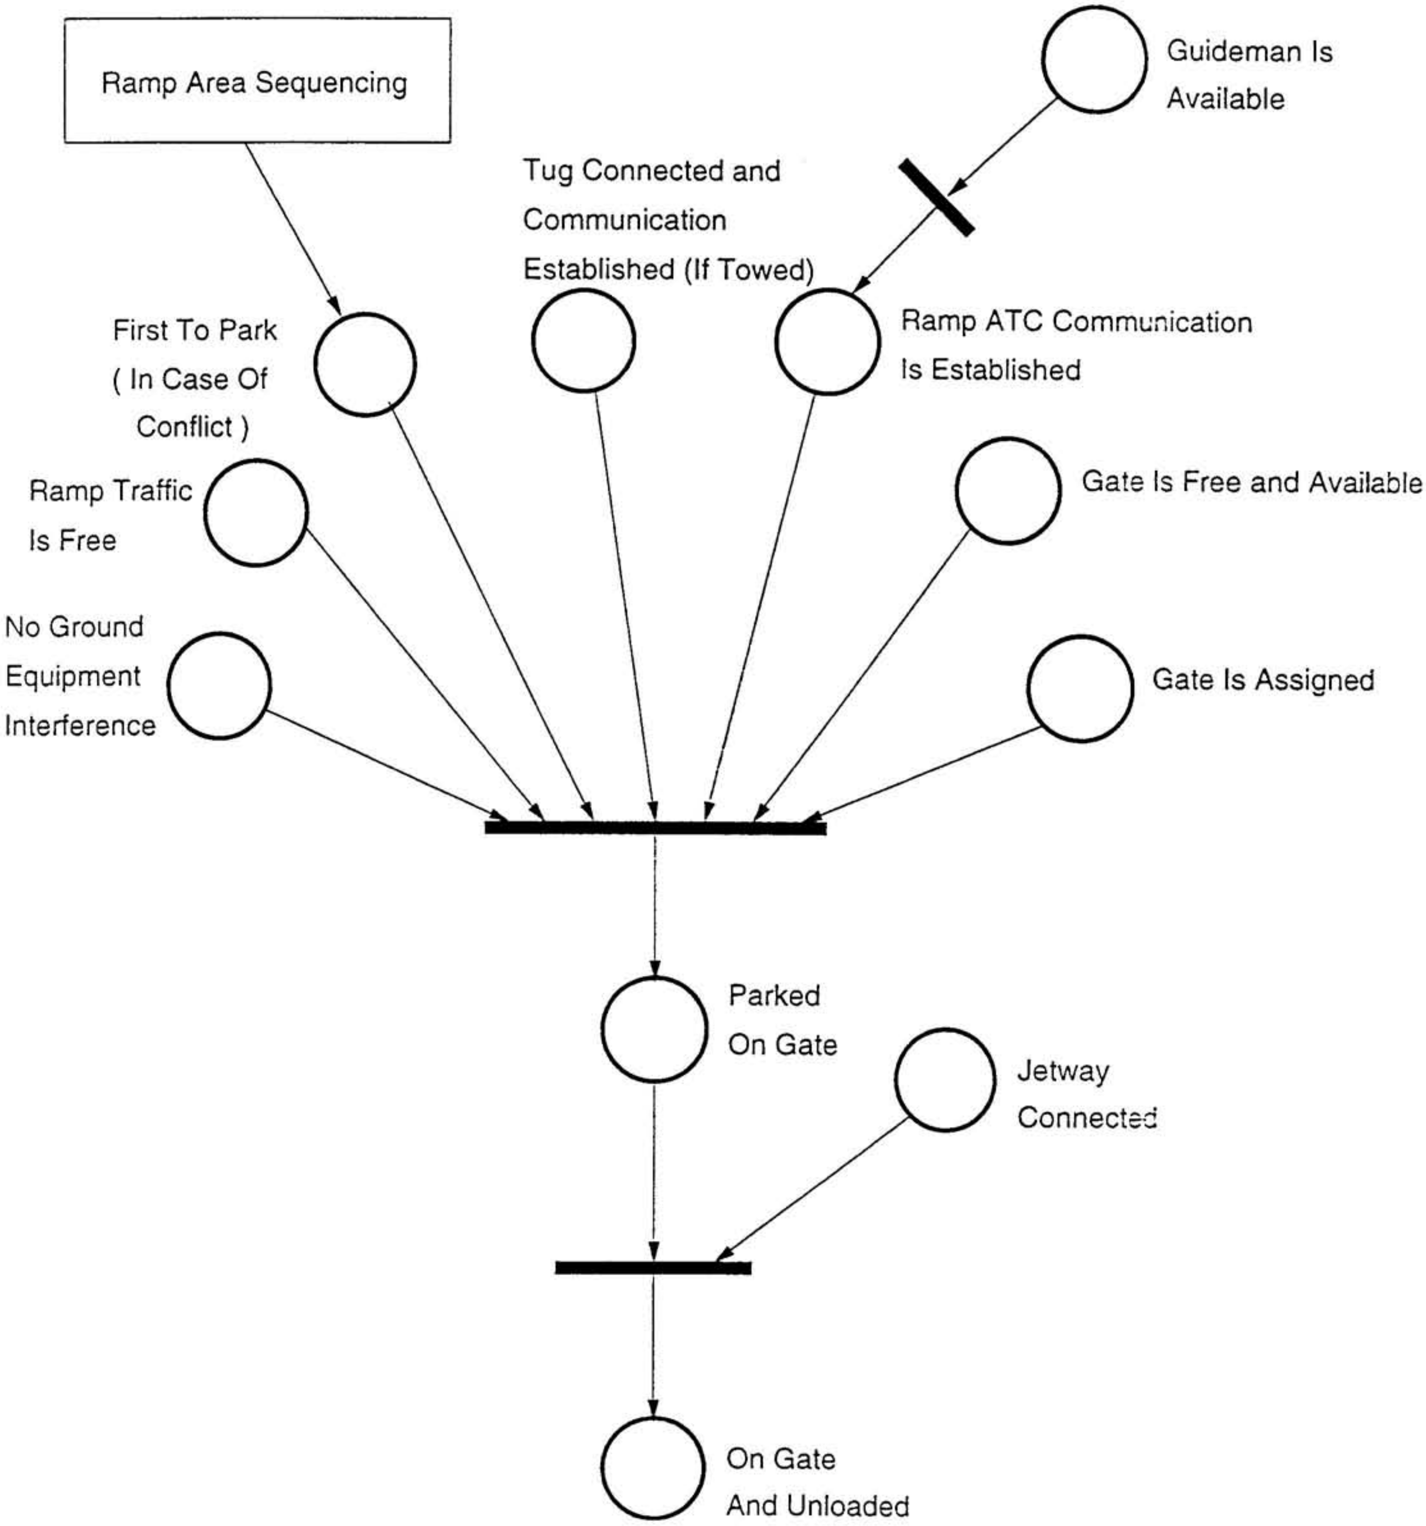

Figure B25: Gate Status Cycle (Preliminary)

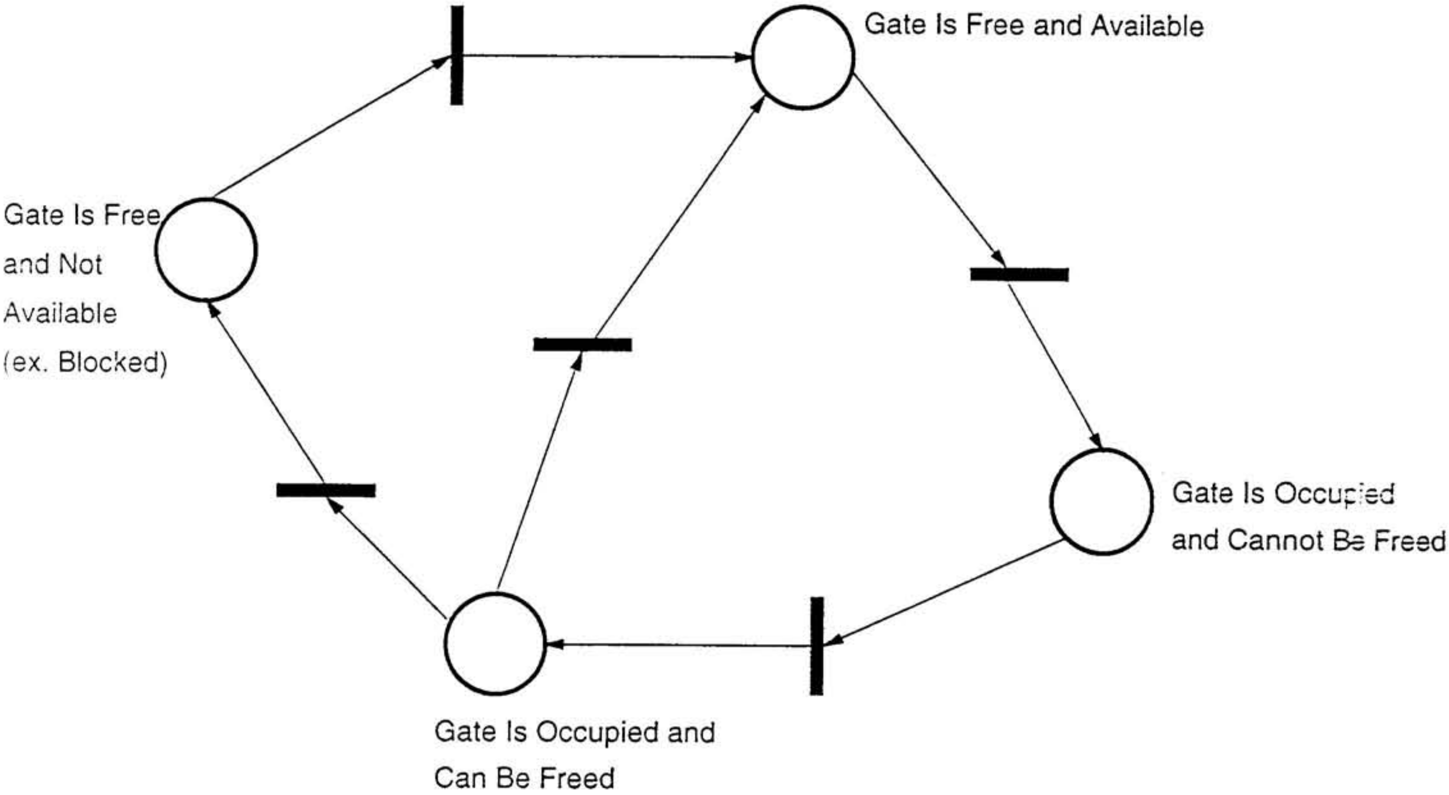

Figure B26: Gate Cycle and Aircraft Gate Processes Combined
(Preliminary)

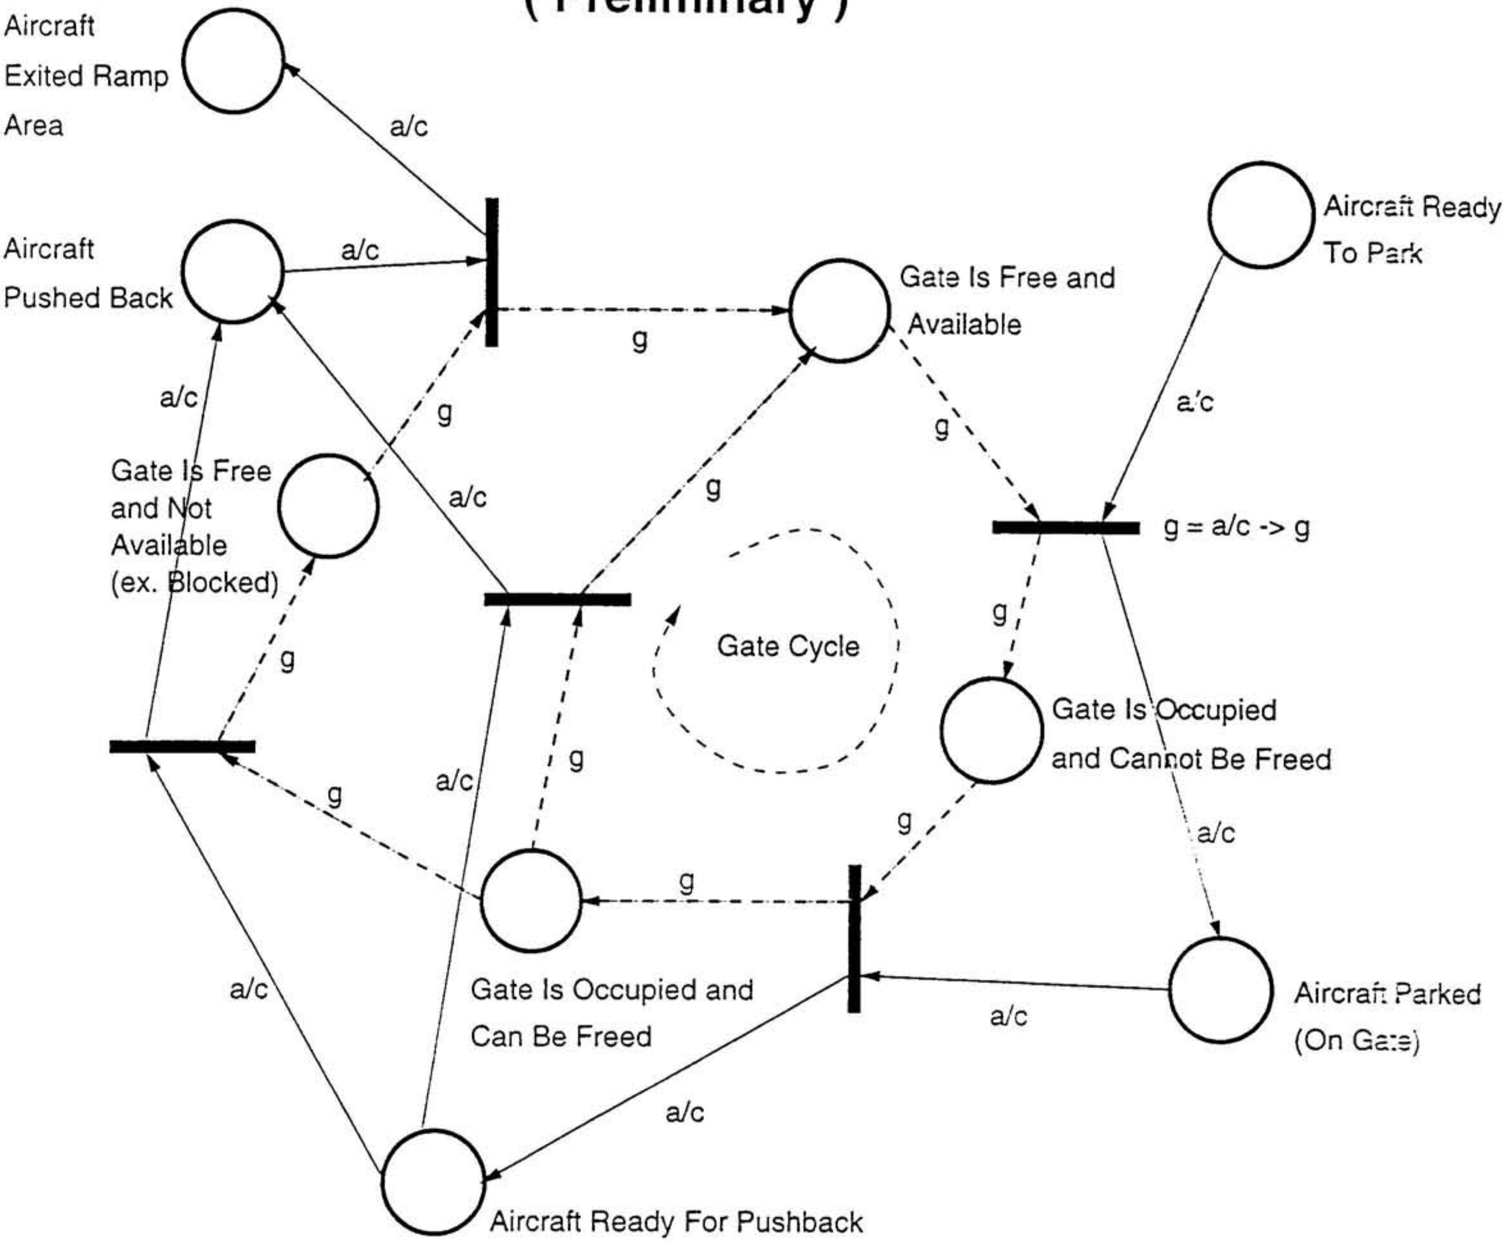

g = gate (the tokens in the gate cycle are gates)

a/c = aircraft (the tokens in the aircraft cycle are aircraft)

Note : This network neglects components other than the gates and the aircraft

Figure B27: Runway Cycles (Preliminary)

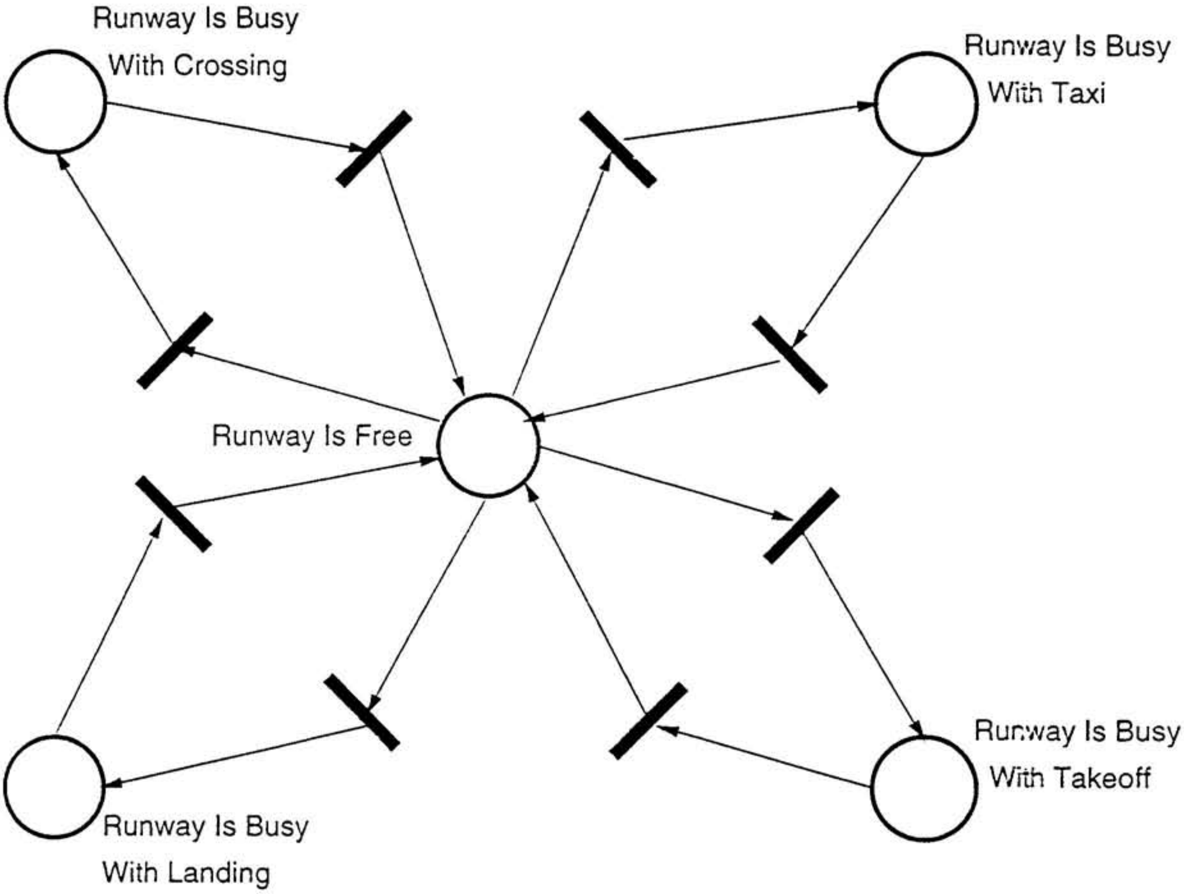

Figure B28: Runway Cycle and Aircraft Landing Combined (Preleminary)

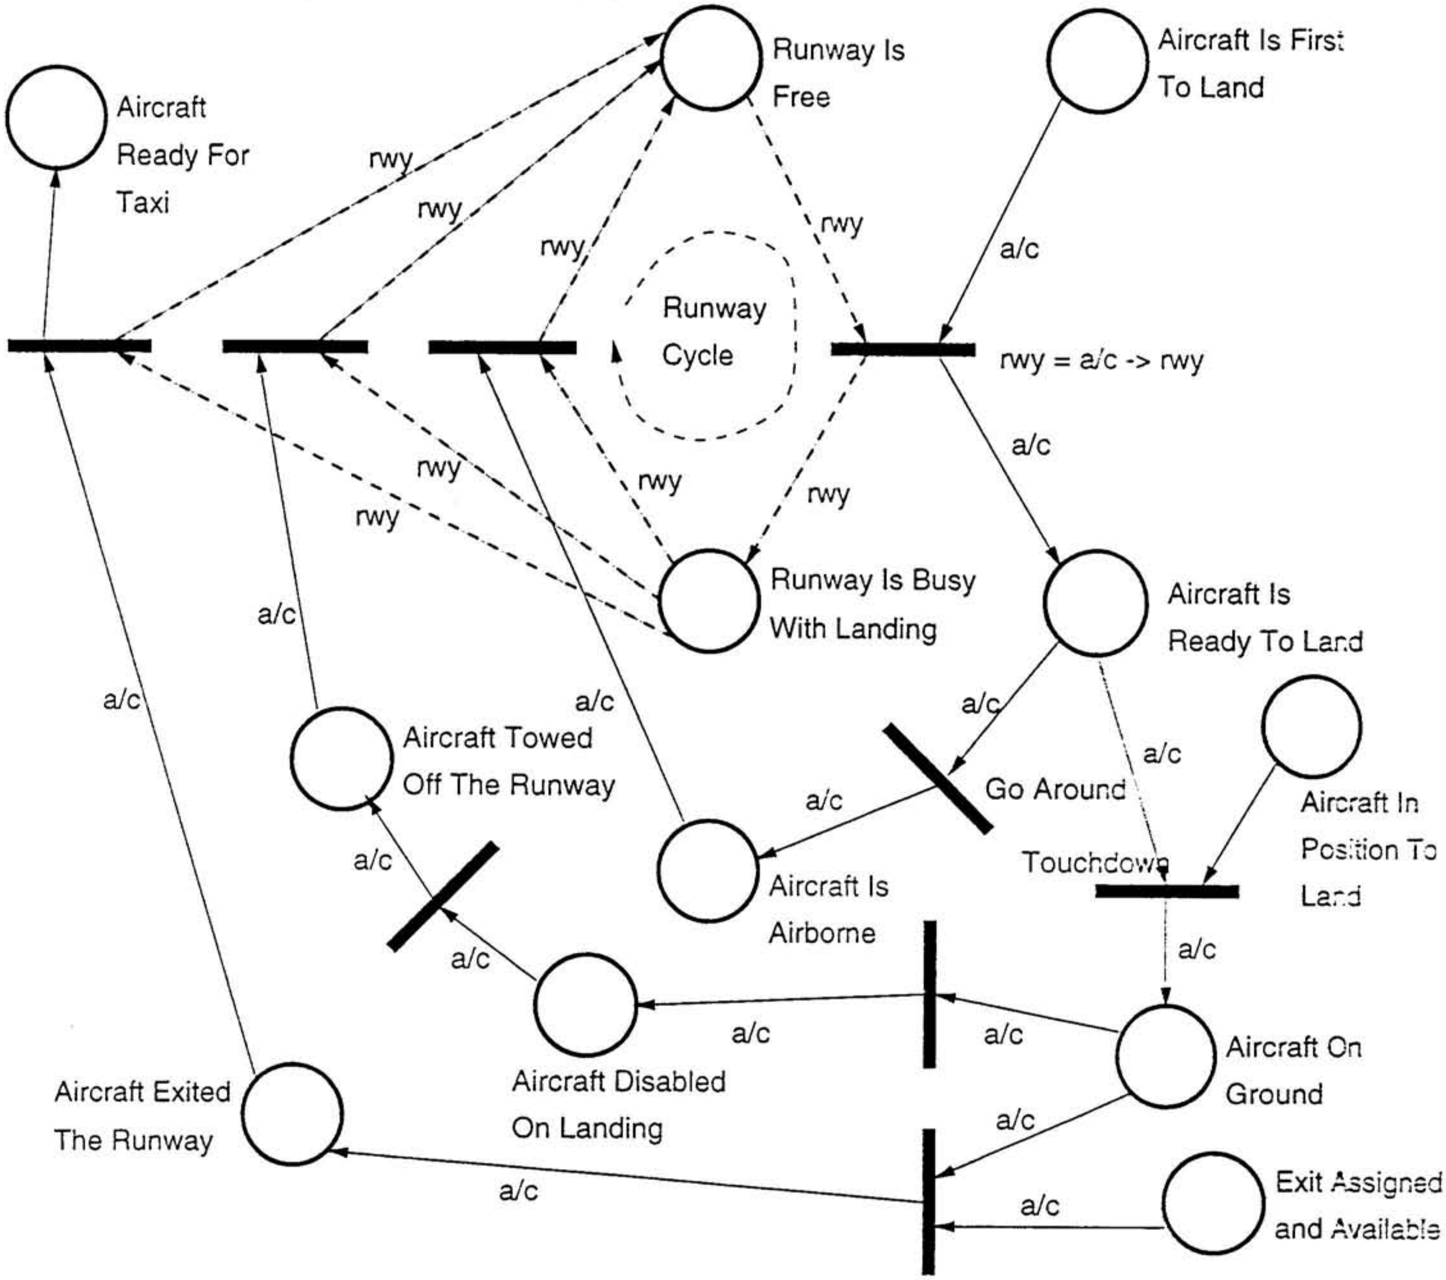

rwy = runway (the tokens in the runway cycle are runways)

a/c = aircraft (the tokens in the aircraft cycle are aircraft)

Note : This network neglects the components othe than the runways and the aircraft

Figure B29: Runway Cycle and Aircraft Takeoff Combined
(Preliminary)

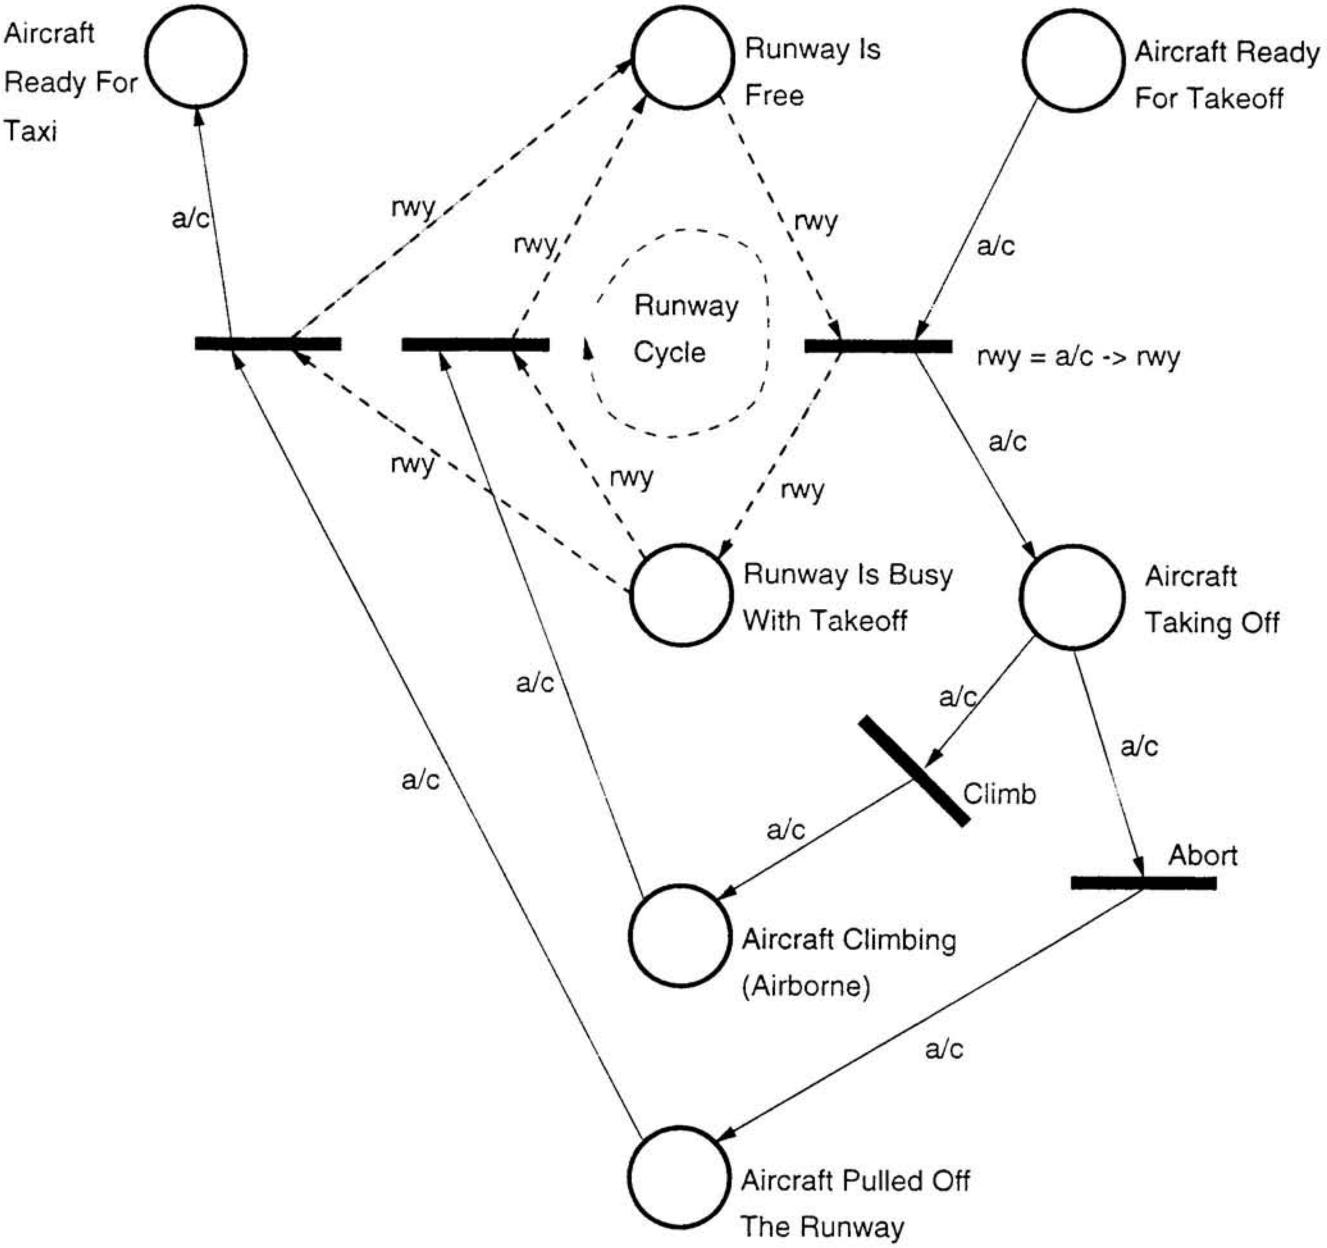

rwy = runway (the tokens in the runway cycle are runways)

a/c = aircraft (the tokens in the aircraft cycle are aircraft)

Note : This network neglects the components other than the runways and the aircraft

APPENDIX C: A HEURISTIC ALGORITHM FOR SCHEDULING AND SEQUENCING DEPARTURES ON A RUNWAY

by Ruben Barocio Cots

The objective of this paper is to present an heuristic algorithm limited to departures only that will determine a position shifted departure sequence such that the throughput time measured from the start of the takeoff roll of the first aircraft and ending with the start of the takeoff roll of the last scheduled aircraft will be minimized.

Airport geometry.

For the purpose of a first approach to the problem, a simple airport geometry was assumed. The airport is assumed to consist of a single runway and two aprons (gate areas) designated D and E at each side of the runway (see Figure 1). For departure purposes only, and for each apron, there is just one taxiway connecting that apron to the two runway thresholds. It is also assumed that the geometry of the apron area is such that once an aircraft has started its push-back procedure, no other aircraft from the same apron area behind the former will be able to overtake it. For runway clearing practices, it will also be assumed that other taxiways exist such that arriving aircraft will be able to conform to standard runway occupancy times.

Algorithm for departures.

Step 1: Identifying departing aircraft.

Let $\{TDAS_i\}$=Time of Departure from Apron Stand i.

Apron stand i is defined to be the ith position in the terminal ramp (be it apron D or apron E) from which the aircraft will start its taxi to the active runway. Additionally, for apron D, i=1, 2,..., m, and for apron E, i=m+1, m+2, ..., n. The TDASi is then defined to be a time at which a certain aircraft will be ready to start its taxi to the runway as reported by the airline some time before the push-back procedure. In this instance, we assume a static case, i.e., the TDAS, once they are reported by the airlines will be deterministic.

Thus, the set of $TDAS_i$ is defined to contain:

Apron D	Apron E
$TDAS_i(D_1)$	$TDAS_i(E_1)$
$TDAS_i(D_2)$	$TDAS_i(E_2)$
.	.
.	.
.	.
$TDAS_i(D_o)$	$TDAS_i(E_p)$
i=1, 2, ..., m	i=m+1, m+2, ..., n

As can be seen, there are a total of o+p departures to be scheduled.

For the purpose of identifying aircraft, D1 will be the aircraft with the earliest TDAS from its own apron stand in apron D. Aircraft D2 will be the aircraft with the next earliest TDAS, etc. The same applies to aircraft originating from Apron E. Thus:

TDAS(D1)<TDAS(D2)<TDAS(D3)<... for aircraft originating from Apron D, and

TDAS(E1)<TDAS(E2)<TDAS(E3)<... for aircraft originating from Apron E.

Note that there is no relationship between the order of the aircraft and the apron stand that it occupies. That is, aircraft E4 may occupy apron stand 8.

Example:

Consider the next apron areas and TDASs:

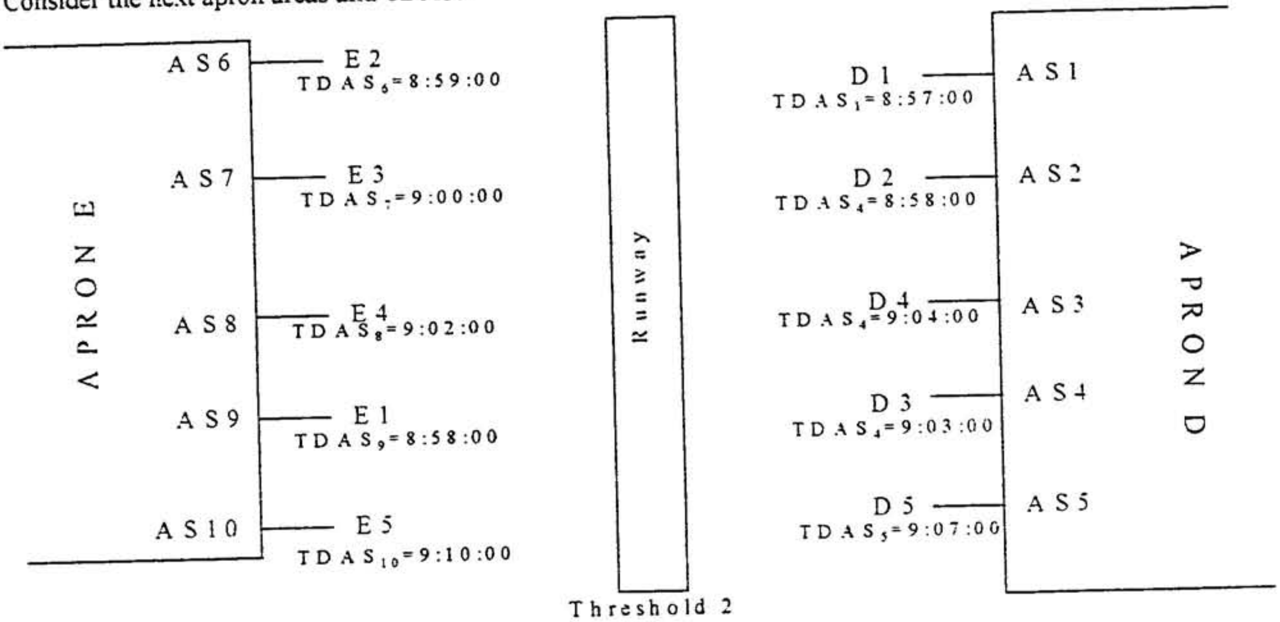

As can be seen, the aircraft on Apron Stand (AS) 1 has the earliest TDAS amongst all aircraft departing from Apron D, and thus, will be called D1. The next earliest TDAS is that of the aircraft in AS2, and thus will be called D2. The third earliest TDAS corresponds to the aircraft on AS4, and will be called D3. Note that there is no correlation between the order of the aircraft according to their TDAS and their apron stands. So, for each apron we have:

Aircraft (D)	TDAS(D)	Apron Stand	Aircraft (E)	TDAS(E)	Apron Stand
D1	TDAS$_1$(D1)=8:57:00	AS1	E1	TDAS$_9$(E1)=8:58:00	AS9
D2	TDAS$_2$(D2)=8:58:00	AS2	E2	TDAS$_6$(E2)=8:59:00	AS6
D3	TDAS$_4$(D3)=9:03:00	AS4	E3	TDAS$_7$(E3)=9:00:00	AS7
D4	TDAS$_3$(D4)=9:04:00	AS3	E4	TDAS$_8$(E4)=9:02:00	AS8
D5	TDAS$_5$(D5)=9:07:00	AS5	E5	TDAS$_{10}$(E5)=9:10:00	AS10

Step 2: Determining a FCFS sequence.

Let $\{TX_{ij}\}$=Taxi time from the ith apron stand to the jth takeoff point, including push-back time.

The jth takeoff point is defined to be a point on the runway from which the aircraft will start its takeoff roll. In this instance, the jth takeoff point can only be either of the two runway thresholds, e.g., j=1 for runway 36 and j=2 for runway 18. As stated under Airport Geometry, it is assumed that the taxi routes from point i to point j are fixed and that no overtaking is possible amongst aircraft originating from the same apron.

The set of TX_{ij} will then contain:

Apron D	Apron E
$TX_{ij}(D_1)$	$TX_{ij}(E_1)$
$TX_{ij}(D_2)$	$TX_{ij}(E_2)$
.	.
.	.
.	.
$TX_{ij}(D_o)$	$TX_{ij}(E_p)$

Let {ETD}=Estimated Time of Departure from the takeoff point, defined to be the time at which the aircraft will start its takeoff roll.

$ETD(D_\alpha)=TDAS_i(D_\alpha)+TX_{ij}(D_\alpha)$ for aircraft departing from apron D, and

$ETD(E_\alpha)=TDAS_i(E_\alpha)+TX_{ij}(E_\alpha)$ for aircraft departing from apron E.

It is at the takeoff point that the FCFS sequence for each apron is defined. That is, the FCFS is defined by the order in which aircraft (within the stream of departures originating at each individual apron) will get to the jth takeoff point according to their respective ETDs.

Let FCFS(D) be the FCFS sequence for aircraft departing from apron D, and FCFS(E) be the FCFS sequence for aircraft departing from apron E The FCFS sequence of departures for each apron is then defined by ordering the ETD(•), such that:

FCFS(D): $ETD(D_\alpha)<ETD(D_\sigma)<ETD(D_\nu)<\ldots$ for aircraft departing from apron D, and

FCFS(E): $ETD(E_\tau)<ETD(E_\omega)<ETD(E_\varepsilon)<\ldots$ for aircraft departing from apron E.

For example, consider that the ordering of TDAS(D) are originally sequenced as follows: $TDAS_i(D_\alpha)<TDAS_i(D_\beta)<TDAS_i(D_\gamma)$. Due to different taxi times, TX_{ij}, for each airplane, the original sequence will be modified to a FCFS measured at the takeoff point. Suppose that $ETD(D_\beta)<ETD(D_\gamma)<ETD(D_\alpha)$. That is, according to the ETD, aircraft β will be assigned position x in the FCFS(D) departure sequence, aircraft γ will be assigned position x+1, aircraft α will be assigned position x+2, and so forth.

Let us continue with the example. We now depict the taxi times from each apron stand to the runway threshold.

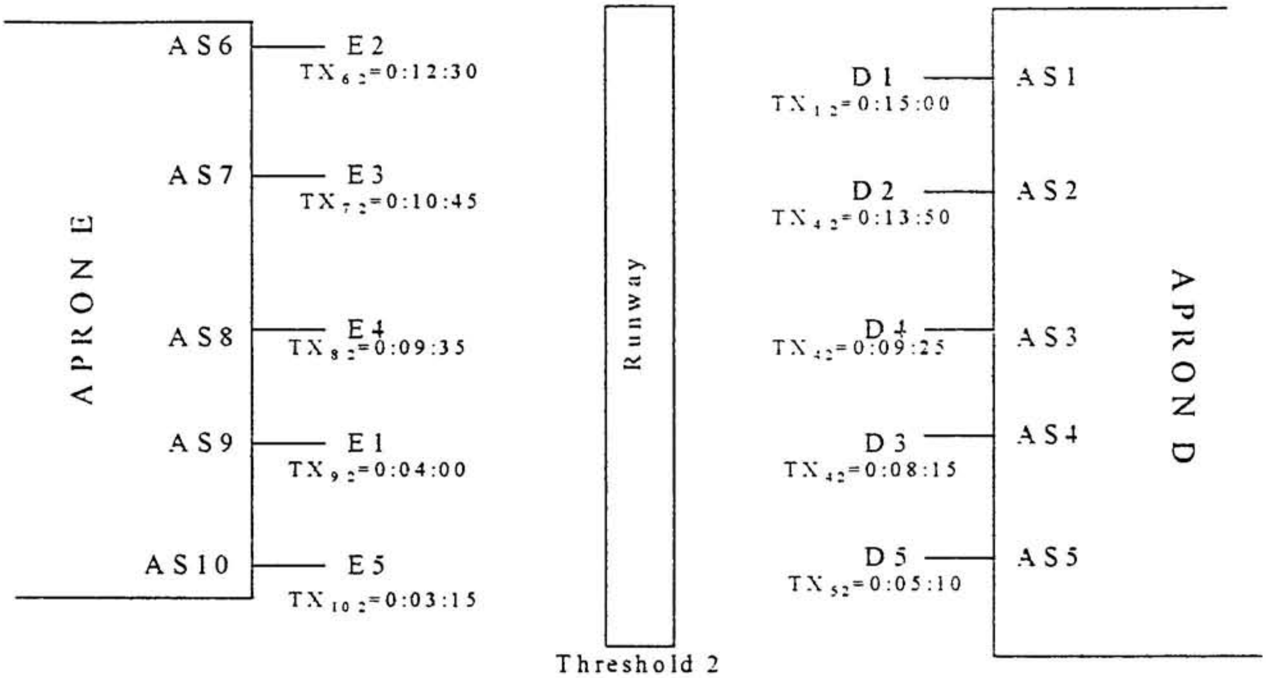

Knowing the TDAS for each aircraft and the TX from each apron stand to the runway threshold, we can now calculate the ETD for each aircraft:

Apron D $ETD(D)=TDAS_i(D)-TX_{ij}(D)$	Apron E $ETD(D)=TDAS_i(E)-TX_{ij}(E)$
$ETD(D_1)=TDAS_1(D_1)-TX_{1,2}(D1)$=8:57:00−0:15:00=9:12:00	$ETD(E1)=TDAS_9(E1)-TX_{9,2}$=8:58:00−0:04:00=9:02:00
$ETD(D2)=TDAS_2(D2)+TX_{2,2}(D2)$=8:58:00−0:13:50=9:12:50	$ETD(E2)=TDAS_6(E2)-TX_{6,2}$=8:59:00−0:12:30=9:11:30
$ETD(D3)=TDAS_4(D3)-TX_{4,2}(D3)$=9:03:00−0:08:15=9:11:15	$ETD(E3)=TDAS_7(E3)-TX_{7,2}$=9:00:00−0:10:45=9:10:45
$ETD(D4)=TDAS_3(D4)-TX_{3,2}(D4)$=9:04:00−0:09:25=9:13:25	$ETD(E4)=TDAS_8(E4)-TX_{8,2}$=9:02:00−0:09:35=9:11:35
$ETD(D5)=TDAS_5(D5)-TX_{5,2}(D5)$=9:07:00−0:05:10=9:12:10	$ETD(E5)=TDAS_{10}(E5)-TX_{10,2}$=9:10:00−0:03:15=9:13:15

Now, we order each aircraft within each individual apron according to their ETD to obtain the FCFS sequences for each apron, FCFS(D) and FCFS(E):

FCFS(D)		FCFS(E)	
Position in FCFS(D)	Aircraft	Position in FCFS(E)	Aircraft
1	D3: ETD(D3) =9:11:15	1	E1: ETD(E1) =9:02:00
2	D1: ETD(D1)=9:12:00	2	E3: ETD(E3)= 9:10:45
3	D5: ETD(D5) =9:12:10	3	E2: ETD(E2)= 9:11:30
4	D2: ETD(D2)=9:12:50	4	E4: ETD(E4) =9:11:35
5	D4: ETD(D4)= 9:13:25	5	E5: ETD(E5)=9:13:15

Once the FCFS sequence for each apron is defined, the two sequences are merged into a single FCFS sequence by ordering the aircraft from each apron according to their ETD. The first aircraft in the merged sequence will be called F_1, the second F_2, and so forth. Thus, the merged sequence will be composed by aircraft F_1, F_2, F_3, …, F_{a+b}. Note that the FCFS sequence is not optimal in the sense that it will minimize throughput time as defined above. For example, suppose that the FCFS sequence is:

FCFS sequence	Type of aircraft	Scheduled ETD	Time in trail
1	D1=Heavy	00:00:00	-
2	D2=Large/Medium	00:00:10	120 seconds
3	D3=Small	00:00:20	60 seconds

The throughput time for this sequence will be 120+60=180 seconds (note that the differences in ETD are not binding).

Now consider that a new sequence can be constructed:

FCFS sequence	Type of aircraft	Scheduled ETD	Time in trail
2	Large/Medium	00:00:10	-
3	Small	00:00:20	60 seconds
1	Heavy	00:00:00	60 seconds

In this case, the throughput time will be 10+60+60=130 seconds. Note that this sequence achieves a savings of 50 seconds when compared to the original FCFS sequence, despite the fact that departures will start 10 seconds later. If throughput times are to be reduced in order to better use an airport's capacity, it is then clear that altering the FCFS sequence is a good means to achieve that objective. However, the problem must be approached carefully, since if we followed the same train of thought, there could be cases in which huge delays could accumulate for certain types of aircraft. The above procedure, per se, could systematically place heavy aircraft at the end of the line, time and time again. Imagine a situation in which a Heavy aircraft occupies the first spot in the FCFS sequence, followed by 5 small aircraft. In order to optimize the sequence, the heavy aircraft would have to be moved to the 5th spot, and if non-heavy aircraft were to continue to line-up in the sequence, it could very well be the case that the heavy aircraft would continually be put at the end of the line. Clearly, this is unacceptable, and a constraint must be introduced to avoid this potential problem. Clearly, such a constraint must limit the number of spots that an aircraft may be shifted backward relative to the FCFS sequence. For illustrative purposes, this paper assumes that any aircraft may not be moved backward (pushed to the back of the sequence) more than two spots relative to the FCFS sequence. However, this not be the case, and the constraint may be treated as a decision variable, set to any number of backward shifts which the decision maker deems appropriate.

In the previous box, the FCFS sequences for each apron, FCFS(D) and FCFS(E) had been calculated. Now, we merge those two sequences into one sequence by ordering the aircraft according to their ETD to obtain the FCFS sequence.

FCFS(D)		FCFS(E)		Merging of FCFS(D) and FCFS(E) to produce FCFS	
Position in FCFS(D)	Aircraft	Position in FCFS(E)	Aircraft	Position in FCFS	Aircraft
		1	E1: ETD(E1) =9:02:00	1	E1=F1: ETD(F1) =9:02:00
		2	E3: ETD(E3)= 9:10:45	2	E3=F2: ETD(F2)= 9:10:45
1	D3: ETD(D3) =9:11:15			3	D3=F3: ETD(F3) =9:11:15
		3	E2: ETD(E2)= 9:11:30	4	E2=F4: ETD(F4)= 9:11:30
		4	E4: ETD(E4) =9:11:35	5	E4=F5: ETD(F5) =9:11:35
2	D1: ETD(D1)=9:12:00			6	D1=F6: ETD(F6)=9:12:00
3	D5: ETD(D5) =9:12:10			7	D5=F7: ETD(F7) =9:12:10
4	D2: ETD(D2)=9:12:50			8	D2=F8: ETD(F8)=9:12:50
		5	E5: ETD(E5)=9:13:15	9	E5=F9: ETD(F9)=9:13:15
5	D4: ETD(D4)= 9:13:25			10	D4=F10: ETD(F10)= 9:13:25

Step 3: Construction of a throughput-time minimizing position shifted sequence (PSS).

The Position Function (PF).

To choose a PSS, we look at the throughput times achieved by each possible combination of departure sequences. Since there is a constraint which does not allow any single aircraft to be position shifted backward (pushed to the back of the sequence) more than 2 slots when compared to the FCFS sequence, it suffices to explore the throughput times achieved by the combinations of departure sequences of three aircraft (taken according to the FCFS sequence) at a time. Based on the minimum throughput-time combination, each aircraft, within the sequence being considered is assigned a position within that same sequence. That is, within the three-aircraft sequence that minimizes throughput time, one aircraft will be assigned position 1, another will be assigned position 2 and yet another will be assigned position 3. That aircraft, whose position is 1, will leave the set of three aircraft being considered, and will be assigned to the Position Shifted Sequence (except if any of the two other aircraft has already been shifted backward two slots relative to the FCFS sequence, in which case, that aircraft will be the one assigned to the PSS and will thus, leave the set). Then, the next aircraft in the FCFS sequence will enter the three-aircraft set and the process will be repeated. Every time the process is repeated will now be called an iteration. Note that this single process does not guarantee that an aircraft will not be pushed to the back of the line (when compared to the FCFS sequence) more than two slots. It is therefore necessary to introduce another process which keeps track of the number of slots that each aircraft has been position shifted backward in order to constrain that shifting to no more that two slots.

Formally: let $P(F_{(\bullet)}/k) \in [1, 2, 3]$ be the Position Function (PF) for aircraft number (•) in the FCFS sequence given the kth iteration, i.e., solely within the three aircraft being considered in each iteration.

By looking at the throughput times achieved by each possible combination of three available aircraft taken sequentially from the FCFS sequence, we choose the optimal sequence. In this case, available aircraft refers to the fact that the aircraft being considered have not been assigned to the PSS and thus, have not yet been removed from the set that conforms the FCFS sequence. Based on that, the Position Function for each aircraft considered will assign a value to each aircraft which represents the position of that particular aircraft in the sequence of the total of three aircraft. For each change of iteration, the values assigned by the Position Function are reset to 0.

The Counter Function (CF).

The Counter Function is, as it name implies, a counter. Its purpose is to count the number of times that an aircraft has been considered throughout the iterations of the Position Function process in an effort to identify which aircraft, amongst the set of three aircraft being considered at a time, has been position shifted backward by two slots relative to the FCFS sequence.

Formally, $CF(F_{(\bullet)}/k)=k-(\bullet)$. If $CF(F_{(\bullet)}/k)=2$, then aircraft $F_{(\bullet)}$ has been position shifted backward by two slots and must therefore be assigned to the PSS in iteration k.

The Position Shifting Function (PSF).

The Position Shifting Function is part of a process that operates in parallel with the Position Function process and which keeps track of the number of slots that each aircraft has been position shifted relative to the FCFS sequence, but solely within the set of three aircraft considered by the Position Function. Its objective will be clarified in a few lines.

Let $PS(F_{(\bullet)}/k)\in[-2, -1, 0, 1, 2]$ be the position shifting function which denotes the number of slots that aircraft number $(\bullet)$ has been position shifted relative to the FCFS sequence in the kth iteration. If the aircraft is position-shifted backward (pushed to the back of the sequence), the PS Function will take positive values and vice versa. Its objective is to revise the ETD of aircraft that have been position shifted backwards: if an aircraft is shifted backwards, i.e., $PS(F_{(\bullet)}/k)=1$, or $PS(F_{(\bullet)}/k)=2$, its ETD must be revised, since it would be impossible for it to leave before an aircraft that has a later ETD!

Therefore:

$ETD(F_{(\bullet)}/k)=ETD(G_{(k-1)})+1s\ \forall\ F_{(\bullet)}$ such that $PS(F_{(\bullet)}/k-1)=1$ or $PS(F_{(\bullet)}/k-1)=2$ if $ETD(G_{(k-1)})-ETD(F_{(\bullet)}/k)>0$

The 1 second that we add to the revised ETD is to avoid any potential confusions that could arise if the algorithm were to implemented using a computer.

Let Tit $_{\alpha,\beta}$ be the time in trail between the departures of aircraft α and aircraft β. Tit $_{\alpha,\beta}$ will take the following values:

Leading Aircraft	Following Aircraft: Heavy	Large/Medium	Small
Heavy	90s	120s	120s
Large/Medium	60s	60s	60s
Small	60s	60s	60s

We now proceed trough an entire iteration in our efforts to produce a PSS.

First, we proceed to choose the optimal sequence and to assign values to the Position Functions and Position Shifting Functions for each aircraft considered. The following table describes in detail the construction of the possible sequences and the corresponding values produced by the Position Function and the Position Shifting Function for each of the aircraft being considered.

Throughput time for a sequence of 3 aircraft Choose:		Position Function in the kth iteration			Position Shifting Function in the kth iteration		
		$P(F_\alpha/k)$	$P(F_\beta/k)$	$P(F_\delta/k)$	$PS(F_\alpha/k)$	$PS(F_\beta/k)$	$PS(F_\delta/k)$
min	$ETD(F_\alpha)+Tit_{\alpha,\beta}+Tit_{\beta,\delta}$	1	2	3	0	0	0
	$ETD(F_\alpha)+Tit_{\alpha,\delta}+Tit_{\delta,\beta}$	1	3	2	0	1	-1
	$ETD(F_\beta)+Tit_{\beta,\alpha}+Tit_{\alpha,\delta}$	2	1	3	1	-1	0
	$ETD(F_\beta)+Tit_{\beta,\delta}+Tit_{\delta,\alpha}$	3	1	2	2	-1	-1
	$ETD(F_\delta)+Tit_{\delta,\alpha}+Tit_{\alpha,\beta}$	2	3	1	1	1	-2
	$ETD(F_\delta)+Tit_{\delta,\beta}+Tit_{\beta,\alpha}$	3	2	1	2	0	-2

Where $\alpha<\beta<\delta$, e.g.. $\alpha=2$, $\beta=4$, $\delta=5$.

Given that the minimum throughput-time sequence has been chosen, the final positioning of aircraft in the departing sequence is calculated as follows:

Let $PSS(G_{(k)}/k)$ define the aircraft (amongst the set of aircraft that conform the FCFS sequence, i.e.. $F_{(\bullet)}$) that will enter the Position Shifted Sequence in iteration k.

As each iteration takes place, one of the aircraft in the FCFS sequence will be assigned a position in the PSS. Obviously, in iteration k=1, the aircraft assigned to the PSS will be in position 1 of that PSS and will then be called G1. In iteration k=2, another aircraft will be assigned to the PSS, will occupy position 2 in the PSS and will be called G2 and so on.

Choose $F_{(\bullet)}$ to satisfy:

$$PSS(G_{(k)}/k)=\left|\begin{array}{ll} \min\{P(F_\alpha/k), P(F_\beta/k), P(F_\delta/k)\} & \text{if } CF(F_{(\bullet)}/k)\neq 2, \\ PS(F_{(\bullet)}/k) & \text{if } CF(F_{(\bullet)}/k)=2 \end{array}\right|$$

As can be seen, the aircraft that will enter the PSS in iteration k is the first aircraft in the optimal sequence, i.e.. that aircraft whose Position Function has a value of 1, given that the optimal sequence has been chosen (this can be seen on the first line of the decision rule above), except if there is an aircraft whose Counter Function in the kth iteration equals a value of 2, in which case, the constraint is binding, and that aircraft must be the one to enter the PSS in iteration k (second line of the decision rule above).

The above process is repeated until all aircraft are assigned to the PSS. The total number of iterations for this to take place will be of o+p-2: 1≤k≤o+p-2, since the last iteration will assign the remaining three aircraft to the PSS. The final product is a Position Shifted Sequence which is optimal given the constraint.

In this box, we continue with the example.

First, suppose that the type of aircraft for each departing airplane is as follows:

Aircraft (FCFS)	F1	F2	F3	F4	F5	F6	F7	F8	F9	F10
Aircraft type	Heavy	Heavy	Small	Medium	Medium	Heavy	Heavy	Small	Heavy	Medium
ETD	9:02:00	9:10:45	9:11:15	9:11:30	9:11:35	9:12:00	9:12:10	9:12:50	9:13:15	9:13:25

Note: the Throughput time minimizing sequence is depicted in the box.

Iteration k=1					P(F1/1)	P(F2/1)	P(F3/1)	PS(F1/1)	PS(F2/1)	PS(F3/1)
Min	9:02:00 +	0:01:30 +	0:02:00 =	9:05:30	1	2	3	0	0	0
	9:02:00 +	0:02:00 +	0:01:00 =	9:05:00	1	3	2	0	1	-1
	9:10:45 +	0:01:30 +	0:02:00 =	9:14:15	2	1	3	1	-1	0
	9:10:45 +	0:02:00 +	0:01:00 =	9:13:45	3	1	2	2	-1	-1
	9:11:15 +	0:01:00 +	0:01:30 =	9:13:45	2	3	1	1	1	-2
	9:11:15 +	0:01:00 +	0:01:30 =	9:13:45	3	2	1	2	0	-2

CF(F1/1)=1-1=0 PS(F1/1)=0

CF(F2/1)=1-2=-1 PS(F2/1)=1

CF(F3/1)=1-3=-2 PS(F3/1)=-1

PSS(G1/1)=P(F1/1) and G1=F1.

ETD(G1)=ETD(F1)=9:02:00

Iteration k=2					P(F2/2)	P(F3/2)	P(F4/2)	PS(F2/2)	PS(F3/2)	PS(F4/2)
Min	9:10:45 +	0:02:00 +	0:01:00 =	9:13:45	1	2	3	0	0	0
	9:10:45 +	0:02:00 +	0:01:00 =	9:13:45	1	3	2	0	1	-1
	9:11:15 +	0:01:00 +	0:02:00 =	9:14:15	2	1	3	1	-1	0
	9:11:15 +	0:01:00 +	0:01:00 =	9:13:15	3	1	2	2	-1	-1
	9:11:30 +	0:01:00 +	0:02:00 =	9:14:30	2	3	1	1	1	-2
	9:11:30 +	0:01:00 +	0:01:00 =	9:13:30	3	2	1	2	0	-2

CF(F2/2)=2-2=0 PS(F2/2)=2

CF(F3/2)=2-3=-1 PS(F3/2)=-1

CF(F4/2)=2-4=-2 PS(F4/2)=-1

PSS(G2/2)=P(F3/2) and G2=F3.

ETD(G2)=ETD(F3)=9:11:15

For the next iteration (k=3), note that ETD(F2), which originally was 9:10:45, must now be revised and set to be 9:11:15, according to:

$ETD(F_{i,\bullet}/k)=ETD(G_{(k-1)})+1s\ \forall F_{i,\bullet}$ such that $PS(F_{i,\bullet}/k-1)=1$ or $PS(F_{i,\bullet}/k-1)=2$ if $ETD(G_{(k-1)})-ETD(F_{i,\bullet}/k)>0$. To emphasize this, the revision has been highlighted. The same will occur every time a revision takes place.

Iteration k=3

					P(F2/3)	P(F4/3)	P(F5/3)	PS(F2/3)	PS(F4/3)	PS(F5/3)
Min	9:11:16 +	0:02:00 +	0:01:00 =	9:14:16	1	2	3	0	0	0
	9:11:16 +	0:02:00 +	0:01:00 =	9:14:16	1	3	2	0	1	-1
	9:11:30 +	0:01:00 +	0:02:00 =	9:14:30	2	1	3	1	-1	0
	9:11:30 +	0:01:00 +	0:01:00 =	9:13:30	3	1	2	2	-1	-1
	9:11:35 +	0:01:00 +	0:02:00 =	9:14:35	2	3	1	1	1	-2
	9:11:35 +	0:01:00 +	0:01:00 =	9:13:35	3	2	1	2	0	-2

CF(F2/3)=3-2=1 PS(F2/3)=2

CF(F4/3)=3-3=0 PS(F4/3)=-1

CF(F5/3)=3-4=-1 PS(F5/3)=-1

PSS(G3/3)=P(F4/3) and G3=F4.

ETD(G3)=ETD(F4)=9:11:30

Iteration k=4

					P(F2/4)	P(F5/4)	P(F6/4)	PS(F2/4)	PS(F5/4)	PS(F6/4)
Min	9:11:31 +	0:02:00 +	0:01:00 =	9:14:31	1	2	3	0	0	0
	9:11:31 +	0:01:30 +	0:02:00 =	9:15:01	1	3	2	0	1	-1
	9:11:35 +	0:01:00 +	0:01:30 =	9:14:05	2	1	3	1	-1	0
	9:11:35 +	0:01:00 +	0:01:30 =	9:14:05	2	3	1	2	-1	-1
	9:11:50 +	0:01:30 +	0:02:00 =	9:15:20	3	1	2	1	1	-2
	9:11:50 +	0:02:00 +	0:01:00 =	9:14:50	3	2	1	2	0	-2

CF(F2/4)=4-2=2 PS(F2/4)=1

CF(F5/4)=4-5=-1 PS(F5/4)=-1

CF(F6/4)=4-6=-2 PS(F6/4)=0

Note that in this iteration, the optimal sequence would be given by F5, F2, F6. However, CF(f2/4)=2, which indicates that F2 has been position shifted 2 slots backward. That is, in this case, the constrain is binding, and thus F2 must enter the PSS at this point.

PSS(G4/4)=P(F2/4) and G4=F2.

ETD(G4)=ETD(F2)=9:11:31

Iteration k=5

					P(F5/5)	P(F6/5)	P(F7/5)	PS(F5/5)	PS(F6/5)	PS(F7/5)
Min	9:11:35 +	0:01:00 +	0:01:30 =	9:14:05	1	2	3	0	0	0
	9:11:35 +	0:01:00 +	0:01:30 =	9:14:05	1	3	2	0	1	-1
	9:11:50 +	0:02:00 +	0:01:00 =	9:14:50	2	1	3	1	-1	0
	9:11:50 +	0:01:30 +	0:02:00 =	9:15:20	2	3	1	2	-1	-1
	9:12:30 +	0:02:00 +	0:01:00 =	9:15:30	3	1	2	1	1	-2
	9:12:30 +	0:01:30 +	0:02:00 =	9:16:00	3	2	1	2	0	-2

CF(F5/5)=5-5=0 PS(F5/5)=0

CF(F6/5)=5-6=-1 PS(F6/5)=0

CF(F7/5)=5-7=-2 PS(F7/5)=0

PSS(G5/5)=P(F5/5) and G5=F5.

ETD(G5)=ETD(F5)=9:11:35

teration k=6					P(F6/6)	P(F7/6)	P(F8/6)	PS(F6/6)	PS(F7/6)	PS(F8/6)
Min	9:11:50 +	0:01:30 +	0:02:00 =	9:15:20	1	2	3	0	0	0
	9:11:50 +	0:02:00 +	0:01:00 =	9:14:50	1	3	2	0	1	-1
	9:12:00 +	0:01:30 +	0:02:00 =	9:15:30	2	1	3	1	-1	0
	9:12:00 +	0:02:00 +	0:01:00 =	9:15:00	3	1	2	2	-1	-1
	9:12:10 +	0:01:00 +	0:01:00 =	9:14:10	2	3	1	1	1	-2
	9:12:10 +	0:01:00 +	0:01:00 =	9:14:10	3	2	1	2	0	-2

CF(F6/6)=6-6=0 PS(F6/6)=1

CF(F7/6)=6-7=-1 PS(F7/6)=1

CF(F8/6)=6-8=-2 PS(F8/6)=-2

PSS(G6/6)=P(F8/6) and G6=F8.

ETD(G6)=ETD(F8)=9:12:10

teration k=7					P(F6/7)	P(F7/7)	P(F9/7)	PS(F6/7)	PS(F7/7)	PS(F9/7)
Min	9:12:11 +	0:01:30 +	0:01:30 =	9:15:11	1	2	3	0	0	0
	9:12:11 +	0:01:30 +	0:01:30 =	9:15:11	1	3	2	0	1	-1
	9:12:12 +	0:01:30 +	0:01:30 =	9:15:12	2	1	3	1	-1	0
	9:12:12 +	0:01:30 +	0:01:30 =	9:15:12	2	3	1	2	-1	-1
	9:13:15 +	0:01:30 +	0:01:30 =	9:16:15	3	1	2	1	1	-2
	9:13:15 +	0:01:30 +	0:01:30 =	9:16:15	3	2	1	2	0	-2

CF(F6/7)=7-6=1 PS(F6/7)=0

CF(F7/7)=7-7=0 PS(F7/7)=0

CF(F9/7)=7-9=-2 PS(F9/7)=0

PSS(G7/7)=P(F6/7) and G7=F6.

ETD(G7)=ETD(F6)=9:12:11

eration k=8					P(F7/8)	P(F9/8)	P(F10/8)	PS(F7/8)	PS(F9/8)	PS(F10/8)
in	9:12:12	0:01:30	0:02:00	9:15:42	1	2	3	0	0	0
	9:12:12	0:02:00	0:01:00	9:15:12	1	3	2	0	1	-1
	9:13:15	0:01:30	0:02:00	9:16:45	2	1	3	1	-1	0
	9:13:15	0:02:00	0:01:00	9:16:15	2	3	1	2	-1	-1
	9:13:25	0:01:00	0:01:30	9:15:55	3	1	2	1	1	-2
	9:13:25	0:01:00	0:01:30	9:15:55	3	2	1	2	0	-2

CF(F7/8)=8-7=1 PS(F7/8)=0

CF(F9/8)=8-9=-1 PS(F9/8)=1

CF(F10/8)=8-10=-2 PS(F10/8)=-1

PSS(G8/8)=P(F7/8) and G8=F7.

ETD(G8)=ETD(F7)=9:12:12

Note that since there are only three airplanes left, there is no need for further iterations. and the position of each aircraft will be given by its Position Function:

PSS(G9/8)=P(F10/8) and G9=F10

ETD(G9)=ETD(F10)=9:13:25

PSS(G10/8)=P(F9/8) and G10=F9

ETD(G10)=ETD(F9)=9:13:26, where the ETD has been revised.

Finally, the resulting PSS is as follows:

Aircraft (PSS)	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10
Aircraft (FCFS)	F1	F3	F4	F2	F5	F8	F6	F7	F10	F9
Original aircraft	E1	D3	E2	E3	E4	D2	D1	D5	D4	E5
ETD (revised)	9:02:00	9:11:15	9:11:30	9:11:31	9:11:35	9:12:10	9:12:11	9:12:12	9:13:25	9:13:26
Aircraft type	Heavy	Heavy	Small	Medium	Medium	Heavy	Heavy	Small	Heavy	Medium

Step 4: Revising the Time of Departure from Apron Stand.

Remember that given that the PSS has been constructed, some airplanes had their ETD revised. Due to this, it is now necessary to revise the TDAS. Also, you will remember that there was no change of notation from the original ETD to the revised ETD. This was so because the implementation of the algorithm into a computer program can deal with that fact, and also because it will reduce the number of variables that must be defined. By the same token, the revised TDAS will not be changed notation-wise.

So, with the revised ETD in hand (obtained from step 3), it is an easy matter to deduct the applicable TX_{ij} to find the revised TDAS:

$TDAS(G_{(\bullet,j)})=ETD(G_{(\bullet,j)})-TX_{ij}(G_{(\bullet,j)})$

Example:

Aircraft (PSS)	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10
Aircraft (FCFS)	F1	F3	F4	F2	F5	F8	F6	F7	F10	F9
Original aircraft	E1	D3	E2	E3	E4	D2	D1	D5	D4	E5
ETD (revised) minus	9:02:00	9:11:15	9:11:30	9:11:31	9:11:35	9:12:10	9:12:11	9:12:12	9:13:25	9:13:26
TX	0:04:00	0:08:15	0:12:30	0:10:45	0:09:35	0:13:50	0:15:00	0:05:10	0:09:25	0:03.15
Revised TDAS	8:58:00	9:03:00	8:59:00	9:00:46	9:02:00	8:58:20	8:57:11	9:07:02	9:04:00	9:10:11

The results are presented graphically using the airport map:

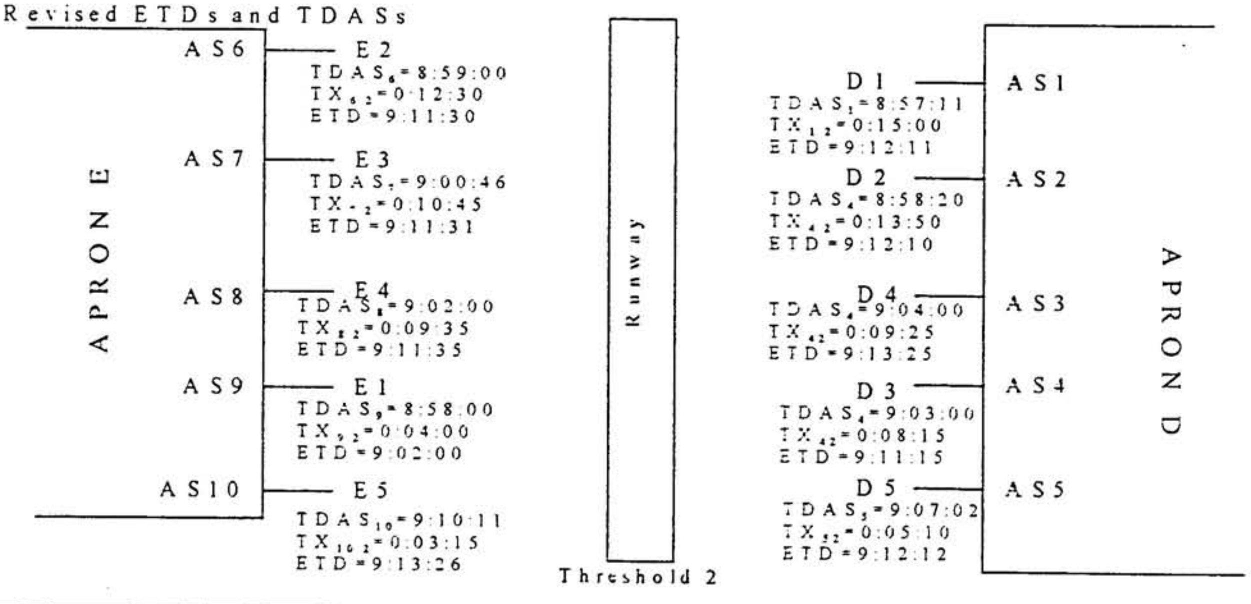

Step 5: Defining Release Times.

The Release Time from Apron stand i is defined to be the time at which ATC, now with an optimal aircraft sequence, will release aircraft from their respective apron stands (gates) in order to avoid potential "cut-offs" in the apron area, which would in turn avoid the implementation of such a sequence.

Obviously, aircraft originating from each apron will have to be released in the order prescribed by the PSS. However, if $TX_{\alpha j}(D_\delta)$ is greater than $TX_{\beta j}(D_\varepsilon)$, there is a possibility that if aircraft were to be released from their apron stands according their respective TDAS, the PSS sequence could not be maintained and the final result would be an unfeasible PSS.

From the airport geometry assumed in this paper, it is clear that the time it takes an aircraft to travel from its apron stand and cross another apron stand in its way to the runway will be given by the difference of the taxi times of such apron stands to the same runway. If we refer to the airport map presented above, it can be seen that the time that it would take an aircraft to travel from, say Apron Stand 6 (AS6 in apron E) across Apron Stand 9 (AS9 in apron E) on its way to runway threshold 2, would be given by $TX_{62}-TX_{92}$=0:12:30-0:04:00=0:08:30, i.e., 8 minutes and 30 seconds.

So, to avoid any potential cutoffs, the following restrictions must then be introduced:

Let {RTASi}=Release Time from Apron stand i.

$$RTASi(G_\beta)= \left| \begin{array}{ll} TDAS(G_\alpha)+[TX_j(G_\alpha)-TX_j(G_\beta)] & \text{if } TX_j(G_\alpha)-TX_j(G_\beta)>0 \text{ and } RTASi(G_\beta)> TDAS(G_\beta)\forall\ G\in D \\ TDAS(G_\beta) & \text{if } TX_j(G_\alpha)-TX_j(D_\beta)<0.\ \forall\ G\in D \end{array} \right|$$

and

$$RTASi(G_\beta)= \left| \begin{array}{ll} TDAS(G_\alpha)+[TX_j(G_\alpha)-TX_j(G_\beta)] & \text{if } TX_j(G_\alpha)-TX_j(G_\beta)>0 \text{ and } RTASi(G_\beta)> TDAS(G_\beta)\forall\ G\in E \\ TDAS(G_\beta) & \text{if } TX_j(G_\alpha)-TX_j(D_\beta)<0.\ \forall\ G\in E \end{array} \right|$$

where $\alpha<\beta$.

Note that the i subscripts have been dropped to indicate that each TX applies according to the gate that G_α and G_β are occupying. Finally, note that there are two distinct RTAS: one for those airplanes G (belonging to the PSS) that originate from apron D, and one for airplanes G that originate from apron E.

What these restriction say is that for airplanes originating from the same apron, the airplane that is ready to initiate its pushback procedure must wait unit the preceding aircraft passes behind it, if the preceding aircraft is originating from an apron stand that is further away from the takeoff point than the apron stand from which the present aircraft is originating.

Example.

First, we focus on the airplanes departing from Apron D. Remember that the order of aircraft G that originate from Apron D is as follows:

Aircraft (PSS)	G2	G6	G7	G8	G9
Original aircraft	D3	D2	D1	D5	D4
Revised TDAS	9:03:00	8:58:20	8:57:11	9:07:02	9:04:00
Apron Stand	4	2	1	5	3

$$RTASi(G_\beta)=\left|\begin{array}{ll} TDAS(G_\alpha)+[TX_j(G_\alpha)-TX_j(G_\beta)] & \text{if } TX_j(G_\alpha)-TX_j(G_\beta)>0 \text{ and } RTAS_i(G_\beta)> TDAS(G_\beta)\forall\ G\in D \\ TDAS(G_\beta) & \text{if } TX_j(G_\alpha)-TX_j(G_\beta)<0,\ \forall\ G\in D \end{array}\right|$$

Aircraft G2=D3, being the first of the sequence, is released at 9:03:00. On its way to the runway, this aircraft will cross only AS5, and $TX_{42}(G_2)-Tx_{i4}(G_\beta)<0$ for i=1, 2, 3. Then, we have $TX_{42}(G_2)-TX_{52}(G_8)$=0:08:15-0:05:10=0:03:05>0 So, initially, the first line of the above constraint applies: $RTAS_5(G_5)$=9:03:00+0:03:05=9:06:05. See now that $RTAS_5(G_8)<TDAS_5(G_8)$, which does not comply with the second condition of the first line. Therefore $RTAS_5(G_8)=TDAS_5(G_8)$=9:07:02 for the time being (since the comparison must be made for each aircraft that crosses behind AS5 and that departs prior to G_8.

The same process is done for every departing aircraft.

In this simulation, no such conflicting situations were encountered.

Step 6. Defining Revised Times of Departure (RTD).

Now we must define the Revised time of Departure (RTD) for each airplane in the PSS.

Let $RTD(G_k)$ be the time at which aircraft k in the PSS will be ready to start its takeoff roll. Then,

$$RTD(G_k)=\left|\begin{array}{ll} ETD(G_k) & \text{if } k=1 \text{ or } ETD(G_k)-ETD(G_{k-1})>Tit_{k-1,k} \\ ETD(G_k)+Tit_{k-1,k} & \text{if } k\neq 1 \text{ and } ETD(G_k)-ETD(G_{k-1})<Tit_{k-1,k} \end{array}\right|$$

Obviously, the RTD for the first aircraft of the day will be its ETD. The same will occur if the Time in trail (Tit) between the G_{k-1} and G_k aircraft is not binding, that is, the difference between the ETDs of the kth and k-1 aircraft is longer than the applicable Time in trail. On the other hand, if such a difference is shorter than the Time in trail, then the RTD will be the ETD plus the applicable Time in trail.

Example.

Aircraft (PSS)	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10
Aircraft type	Heavy	Small	Medium	Heavy	Medium	Small	Heavy	Heavy	Medium	Heavy
Time in trail (s)	NA	120s	60s	60s	120s	60s	60s	90s	120s	90s
ETD	9:02:00	9:11:15	9:11:30	9:11:31	9:11:35	9:12:10	9:12:11	9:12:12	9:13:25	9:13:26

The first airplane will depart at 9:02:00. For the second airplane (G_2), we calculate: $ETD(G_2)-ETD(G_1)$=9:11:15-9:02:00=0:09:15>0:02:00, and thus, the first line of the restriction applies ($ETD(G_k)-ETD(G_{k-1})<Tit_{k-1,k}$) and $RTD(G_2)$=9:11:15. For the rest of the aircraft (G3 through G10), all $ETD(G_k)-ETD(G_{k-1})<Tit_{k-1,k}$, and thus, $RTD(G_k)=ETD(G_k)+Tit_{k-1,k}$

The final results are as follows:

Aircraft (PSS)	G1	G2	G3	G4	G5	G6	G7	G8	G9	G10
Aircraft (FCFS)	F1	F3	F4	F2	F5	F8	F6	F7	F10	F9
Original aircraft	E1	D3	E2	E3	E4	D2	D1	D5	D4	E5
Position shifted	0	-1	-1	2	0	-2	1	1	-1	1
RTAS	8:58:00	9:03:00	8:59:00	9:00:46	9:02:00	8:58:20	8:57:11	9:07:02	9:04:00	9:10:11
ETD	9:02:00	9:11:15	9:11:30	9:11:31	9:11:35	9:12:10	9:12:11	9:12:12	9:13:25	9:13:26
RTD	9:02:00	9:11:15	9:12:15	9:13:15	9:15:15	9:16:15	9:17:15	9:18:45	9:20:45	9:21:45

We now proceed to show the throughput times for three sequences: the FCFS sequence. the PSS sequence which we have developed, and an unconstrained PSS, PSS(U), where the heavy aircraft were pushed to the back of the line without any consideration of a maximum of position shifts.

Aircraft FCFS	RTD	Aircraft PSS	RTD	Aircraft PSS(U)	RTD
F1	9:02:00	F1	9:02:00	F1	9:02:00
F2	9:10:45	F3	9:11:15	F3	9:11:15
F3	9:12:45	F4	9:12:15	F4	9:12:15
F4	9:13:45	F2	9:13:15	F5	9:13:15
F5	9:14:45	F5	9:15:15	F8	9:14:15
F6	9:15:45	F8	9:16:15	F10	9:15:15
F7	9:17:15	F6	9:17:15	F2	9:16:15
F8	9:19:15	F7	9:18:45	F6	9:17:45
F9	9:20:15	F10	9:20:45	F7	9:19:15
F10	9:22:15	F9	9:21:45	F9	9:20:45
Throughput time (last-first)	0:20:15	Throughput time (last-first)	0:19:45	Throughput time (last-first)	0:18:45
Throughput time (last-second)	0:11:30	Throughput time (last-second)	0:10:30	Throughput time (last-second)	0:09:30

As can be seen, if throughput time were to be measured from the first aircraft to the last, the PSS sequence developed in throughout the example provides a time savings of 30 seconds over the FCFS sequence. Using the same measure, the unconstrained PSS provides a time savings of 1 minute and 30 seconds. However, note that this way of measuring the throughput time might not be representative of reality, since there was a period of time (which will be referred to as free period) in which no departures were scheduled. It therefore seems better to measure throughput time considering only those periods in which there were constant departures. In our example, this means that the RTD of the second and last aircraft of the sequence must be compared to calculate the throughput time of the sequence. As shown, the PSS in this case provides a time savings of 1 minute over the FCFS sequence, and the unconstrained PSS provides a time savings of 2 minutes. It could be said that the results are not impressive, but that is not true. Certainly, the ability of the algorithm to reduce throughput time is a function not only of the Heavy-Nonheavy aircraft mix amongst the departures, but also of the relation in time of those two types of aircraft, and most importantly, of the existence of free periods in the departure schedule.

APPENDIX D

An Introduction to the Aircraft Sequencing Problem with Arrivals and Departures

Alp Muharremoglu

1 Introduction

With the increasing use of air transportation, the control of arriving and departing aircraft in an airport's near terminal area became a very complex task, especially during peak hours. The air traffic controller has to make several decisions in a relatively short time, such as deciding which aircraft should use which runways, in what sequence should the aircraft depart or land, what maneuvers should be executed, etc. The problem as a whole is very difficult to attack from a mathematical perspective, and what previous researchers have done in the past has been to focus on some subproblem and make simplifying assumptions about them so that the problem becomes tractable.

The approach we are going to take in this research is similar to the one described above. We shall focus on the problem of scheduling of aircraft that are waiting to land or to depart, on a single runway, in real time. Depending on the progress, the multiple runway case can be considered later. Because of differences in types of aircraft and characteristics, the duration for which an aircraft occupies the runway can be different among different classes of planes. In addition, according to safety regulations, any two coaltitudinal aircraft must maintain a minimum horizontal separation.. This horizontal separation depends on the type of the aircraft that is preceding as well as the type of the aircraft that is following. As a result of these variations, theoretically it is possible to find a schedule that optimizes a function that is thought to be related to the airport workload while keeping passenger comfort and safety and possibly operating costs in mind.

In reality, the arrival and departure requests are generated in a dynamic fashion over time. In this research, initially we are going to try to come up with a way to schedule the arrivals and departures of a given set of aircraft, namely to solve a static version of the problem. The controller can at any time apply this algorithm to the set of aircraft waiting to be serviced by the runway. Whenever a new plane enters the queue or after a specified time period, the algorithm can be applied again. Nevertheless, some concerns regarding the fairness and the applicability of the control have to be considered (e.g., the quoted schedule of an aircraft should not be changed very frequently as new arrivals occur).

Dear [1] has considered a sequencing problem for arriving aircraft. He proposed the CPS rule, which eliminated some of the above mentioned concerns and yielded relatively good performance values. We are planning to use the CPS rule (or a variation of it) in our case as well. The difference between Dear's problem and our problem is that now the problem includes departures as well. This adds a complication, because the controller might want to give priority to arrivals (or departures) if many planes are waiting in the air (waiting to depart). Therefore, we are going to include a parameter in our formulation that can be chosen by the controller, which will reflect this choice of priority.

2 The Objective

The problem at hand is to sequence n arrivals and m departures on a single runway. We are going to calculate the separation matrix among all types of operations (arrivals and departures). This matrix gives the minimum time needed after one operation is completed until another operation can be performed on the runway. (For example, how much time is needed before a large aircraft can land after a medium size aircraft takes off).

For each operation i there is a required runway occupancy time p_i. A minimum time of t_{ij} has to pass before operation j can be performed after operation i. Let's define Π to be a sequence of operations, $\Pi = (\Pi(0), \Pi(1), \ldots \Pi(n+m))$, where $\Pi(i)$ is the index of the i^{th} aircraft in the sequence and define $\Pi(0) = 0$. For example, $\Pi(1) = 2$ means that the first aircraft to use the runway is an aircraft of type 2. Define d_i^{Π} as the delay absorbed by the i^{th} aircraft in the sequence when Π is applied (delay ends when aircraft completes landing or take-off). Initially, we are assuming that all the planes are available for landing or take off at time 0. Then:

$$d_i^{\Pi} = \sum_{k=1}^{i} \left[t_{\Pi(k-1)\Pi(k)} + p_{\Pi(k)} \right]$$

There are going to be two different objectives.

Objective 1: Minimize the total weighted delay of all airplanes. This objective can be written as:

$$\min_{\Pi} \sum_{i \in (A \cup D)} \omega_i \cdot d_i^{\Pi} \tag{1}$$

where ω_i := The weight associated with aircraft i

A := The set of arriving aircraft

$D :=$ The set of departing aircraft

This objective is very general, because there is a weight associated with each aircraft. In our case, we are going to have ω, a constant weight for all the arrivals and (1-ω), a constant weight for all the departures (alternatively, one could use a constant weight for each [aircraft type, operation] pair). So the objective reduces to:

$$\min_{\Pi} \omega \cdot \sum_{i \in A} d_i^{\Pi} + (1-\omega) \cdot \sum_{i \in D} d_i^{\Pi} \tag{2}$$

Objective 2: Minimize the weighted makespan

$$\min_{\Pi} \omega \cdot d_{LA}^{\Pi} + (1-\omega) \cdot d_{LD}^{\Pi} \tag{3}$$

where d_{LA}^{Π} and d_{LD}^{Π} are the delays of the last landing aircraft and the last departing aircraft respectively.

3 The Constraints

The first constraint is the resource constraint, namely the available runway time and available airspace in the final approach route. In fact, this research is relevant because of these constraints. This constraint is captured in the separation matrix that is to be calculated and in the times that the aircraft need to occupy the runway.

Since keeping a particular aircraft in the air for a long time is not desirable due to safety reasons, the proposed algorithm should not assign excessive delays to individual aircraft. In addition, any significant discrimination among the different classes of aircraft is not desirable. So, in some sense, the algorithm should be fair.

Another important problem might arise when the algorithm is being applied in a real world setting. Since the static optimization problem is going to be solved every time a new aircraft enters the terminal area, if the algorithm allows global changes in every iteration, the schedule of an aircraft might change several times before the aircraft is actually serviced. So, the algorithm should be designed so that after the schedule is quoted to the pilot of the aircraft, it should either be frozen or be allowed to change only a small number of times.

A last but equally important consideration to be taken into account is that the algorithm should be very fast to be executed in real time. During a congested period, planes might enter the terminal area as frequently as every 30 seconds. The controller has to run the algorithm and after getting the results he/she should be given some time to finalize the schedule.

4 The Non-Constrained Case

4.1 Minimizing Total Weighted Delay

The problem with the objective function as in 1 with the runway and terminal area airspace constraints turns out to be a special kind of a machine scheduling problem. This is the single machine scheduling problem with sequence dependent setup times and a weighted completion time objective. In the aircraft scheduling setting, the runway corresponds to the machine, the minimum separation times correspond to the sequence dependent setup times and the completion time is the delay that the aircraft absorbs. Previous work on the single machine scheduling problem has focused either on sequence independent setup times, or in the case where setup times are sequence dependent, on minimizing the makespan. The makespan objective reduces the problem to the well known TSP, for which an extensive literature and fairly good exact and approximate solution techniques exist. We have not been able to find any work on the problem with the weighted completion time as the objective. The main difference between the TSP and the weighted completion time problem is that a job's absolute position in the sequence, in addition to its relative position, is important in the weighted completion time problem. Ragatz [2] has a study on "A branch and bound method for minimum tardiness sequencing on a single processor with sequence dependent setup times" and Rubin and Ragatz [3] have a genetic search algorithm for scheduling in a sequence dependent setup environment to minimize tardiness. Tan, Narasimhan, Rubin and Ragatz [4] propose another heuristic approach for the problem. The objective in all these papers is to minimize the total tardiness. Minimizing the total completion time is a special case of the tardiness problem with all due dates set equal to zero. The problem with this approach might be that since the algorithms were designed for an environment with due dates, the performance of the heuristics (time and/or objective function value) could be negatively affected when all due dates are zero. But even if this approach yielded satisfactory results, the objective would be to minimize total completion time, rather than weighted completion time.

Tamimi [5] and Smith [6] have worked on minimizing total weighted tardiness on uniform machines with sequence dependent setup times. The problem they are working on has multiple machines and positive due dates. Tamimi proposes a genetic algorithm for this problem. The performance of this approach in our setting is not tested but in the paper, Tamimi tests his algorithm on problems with at most 20 jobs. In the aircraft scheduling case, a typical number of aircraft waiting to be serviced during peak hours at one time could be around 45-50.

Psaraftis [7] has a dynamic programming approach to the aircraft sequencing problem with only arrivals. His work is not directly applicable to our problem, because we are dealing with arrivals and departures at the same time and we want to include a parameter in the formulation that will allow the controller to assign different priorities to arrivals and departures depending on the specific situation.

This review shows that no immediate answer to the problem at hand can be found in the existing literature. One could try to use the existing algorithms (probably by modifying them a little) and see if the results are satisfactory or not. On the other hand, one could also try to start from scratch and try to come up with a way to sequence the aircraft. Psaraftis observed that the aircraft sequencing problem was highly structured and exploited this structure to reduce the runtime significantly and came up with a DP algorithm that can be executed in real time. The general case of the problem (Objective 1a and no classification of jobs) is a very difficult problem to solve. Nevertheless, we believe that a similar approach to Psaraftis' approach can be taken in our problem as well (by exploiting the structure), so that an exact algorithm using DP can be found that will execute very fast. If this attempt does not prove to be very successful, a heuristic procedure for the problem can be proposed.

4.2 Minimizing the Weighted Makespan

If the objective was to minimize the makespan, meaning to minimize the time at which all the aircraft have been serviced, the problem would be reducible to the TSP and as mentioned earlier, an extensive literature exists on this famous problem. The most popular approach has been branch and bound. and a summary of these branch and bound methods can be found in Balas and Toth [8]. A number of construction and improvement heuristics have been proposed on the problem, and some of these heuristics can solve problems with about 6,000 cities within 3% of the best existing lower bound on the problem. A summary of computational solutions for the TSP can be found in Reinelt [9].

Our problem differs from the original TSP in the sense that we want to minimize the total weighted makespan of the arrivals and the departures. The TSP is known to be NP-hard. In our problem, we are adding a new complication, namely we are saying that there are two different types of cities. The makespans (the time when the last city of a particular type is visited) of both types of cities is then weighted and the sum is minimized. So, it is very likely that this problem is NP-Hard as well.

The observation that was made in the case of minimizing the weighted completion time, namely that the absolute position of the city affects the cost as well as its relative position, applies to this case too. For example, suppose that all arriving planes land before all the planes waiting to depart are allowed to do so. In this case, scheduling a particular departure i before all the arrivals are serviced has a different contribution to the objective function than scheduling the same departure after the same [aircraft type, operation] pair j, after all the arrivals have landed. (The contribution of $(t_{ji} + p_i)$adds to both the makespan of arrivals and departures in the first case, whereas in the second case, it only adds to the makespan of the departures).

Let's introduce a dummy arrival r and a dummy departure t with $t_{ir} = t_{it} = p_r = p_t = 0, \forall i \in (A \cup D)$. Set $\omega_r = \omega$, $\omega_t = (1 - \omega)$ and $\omega_i = 0, \forall i \in (A \cup D)$. Finally introduce two precedence constraints such that arrival r cannot land before all other arrivals have landed and departure t cannot take-off before all other departures have taken-off. Now, our problem with the objective function as in 3 is transformed into one with the objective as in 1 with two additional precedence constraints. This problem is not necessarily easier than the makespan problem, but this relation can be useful in the future, since we are planning to work on both problems simultaneously.

So, similar to the other objective, there exists no past research that can directly be applied to attack the problem. We believe that the general case of the problem is very difficult like the previous one, but again, we hope that an algorithm can be developed for this special case by exploiting the structure. One could also try to come up with a heuristic for the general case which could then be applied to the aircraft sequencing problem, but that requires a different focus than the aircraft sequencing problem. One is forced to look for heuristics in the general case and no exact algorithm can be found that executes in polynomial time, since the problem is NP-hard. (If $NP \neq P$). If a heuristic performs well for the general problem, it would be a very successful work, and a rather significant one, but there is no guarantee that it will perform well in the aircraft sequencing problem. Especially if an optimal solution can be found in real time in non-polynomial ways, this will still outperform any heuristic.

5 The Constrained Case

As mentioned earlier, an optimization of the scheduling problem with only the resource constraints could lead to undesired situations such as excessive delays, discrimination among aircraft types, too frequent changes in the schedule etc. Dear introduced the concept of Constrained Position Shifting, referred to as CPS, which simply constrained the positions that an aircraft can take in the schedule, depending on the initial sequence. The initial sequence is the FCFS schedule of the planes. A number called Maximum Position Shifts (MPS) was used. So, if MPS is 2, a plane that was 5th in the initial sequence can only land 3rd, 4th, 5th, 6th or 7th. (5 ± 2). This approach eliminates almost all concerns mentioned above. Even though the schedule is not optimal, the deviation from the optimal schedule was shown to be not too significant, especially for large MPS values (for the cases studied). Psaraftis proposed a dynamic programming algorithm for CPS scheduling when the problem was to schedule an ordered list of aircraft waiting to land on a single runway.

The CPS methodology could be used in our setting as well. Since we are considering arrivals and departures at the same time, the number of operations and the number of classes (in our case a class is an [aircraft type, operation] pair) are larger than when only arrivals are considered. So, an enumeration of all the permutations could possibly not be feasible in real time in some cases, especially when MPS is a high value. A dynamic programming approach could be more efficient.

Both in Dear's work and Psaraftis' work, there is a single MPS value for all the aircraft. In our case, it might be desirable to have different MPS values for different operations (or even different classes). This would change the state space in the DP formulation, but our prediction is that the work associated with this change is not significant.

6 Conclusion

This report introduces the aircraft sequencing problem on a single runway when there are arrivals and departures present. It gives a brief summary of the related past literature and points out that there is no previous research that is directly applicable to our problem. It also makes the observation that the problem is highly structured and that a relatively efficient method could possibly be developed. In the real world, most airports operate with more than one runway, but the inclusion of departures in the analysis will be a step towards a potential future research, where multiple runways and later multiple airports could be considered. So, analyzing the aircraft sequencing problem with arrivals and departures seems to be promising as far as a hope for a solution is concerned aas well as relevant to real world needs.

References

[1] Dear, R.G., "*The Dynamic Scheduling of Aircraft in the Near Terminal Area,*" Ph.D. Thesis, Flight Transportation Laboratory Report R76-9. M.I.T., Cambridge, MA (1976)

[2] Ragatz, G. L., "*A Branch and Bound Method for Minimum Tardiness Sequencing on a Single Processor with Sequence Dependent Setup Times,*" Proceedings: 24th Annual Meeting of the Decision Sciences Institue, 1375-1377 (1993)

[3] Rubin, P. A. and Ragatz, G. L., "*Scheduling in a Sequence Depenent Setup Environment with Genetic Search,*" Computers & Operations Research, 22(1), 85-99 (1995)

[4] Tan, K. C., Narasimhan, R., Rubin, P. A. and Ragatz, G. L., "*A Simulated Annealing Approach to Minimize Tardiness,*" 26th Annual Meeting of the Decision Sciences Institue, 1514-1516 (1996)

[5] Tamimi, S. A., "*A Genetic Algorithm for Minimizing the Total Weighted Tardiness on Uniform Machines with Sequence Dependent Setup Times,*" Masters Thesis, IMfgE Dept., Wichita State University, Wichita, KS. (1997)

[6] Smith, B. K., "*Minimizing Total Weighted Tardiness on Uniform Parallel Machine with Sequence Dependent Setup Times*," Masters Project, IMfgE Dept., Wichita State University, Wichita, KS. (1994)

[7] Psaraftis, H. N., "*A Dynamic Programming Approach to the Aircraft Sequencing Problem*," Flight Transportation Laboratory Report R78-4, M.I.T., Cambridge, MA (1978)

[8] Balas, E. and Toth, P., "Branch and Bound Methods. In *The Traveling Salesman Problem* (Edited by E.L. Lawler, J.K. Lenstra, A.H.G. Rinooy Kan and D.B. Shmoys)," pp. 361-401 (1985)

[9] Reinelt, G., "*The Traveling Salesman Problem. Computational Solutions for TSP Applications*," Springer-Verlag, Berlin Heidelberg (1994)