

**The Multi-Airport Ground-Holding
Problem in Air Traffic Control**

Peter B. Vranas
Dimitris J. Bertsimas
Amedeo R. Odoni

#3385-92-MSA

February 1992

The Multi-Airport Ground-Holding Problem in Air Traffic Control

Peter B. Vranas, Dimitris J. Bertsimas, Amedeo R. Odoni

Massachusetts Institute of Technology

February 4, 1992

Abstract

In 1986, the total profits of the U.S. airline industry were of the order of \$800 million, while its total delay costs due to congestion approached the figure of \$2 billion. Motivated by this important problem of congestion costs and observing that ground delays are far more preferable than airborne delays, we have formulated and studied generic integer programming models in order to assign optimal ground holding delays in a general network of airports, so that the total (ground plus airborne) delay cost of all flights is minimized. All previous research on this problem has been restricted to the single-airport case, which neglects "down-the-road" effects due to transmission of delays between successive flights performed by the same aircraft.

We give three general pure 0-1 integer programming formulations of the problem, one of which also takes into account the possibility of cancelling flights.

We then present a heuristic algorithm which finds a feasible solution to the integer program by rounding the optimal solution of the LP relaxation. Finally, we present extensive computational results with the goal of obtaining qualitative insights on the behaviour of the problem under various combinations of the input parameters. We demonstrate that the problem can be solved in reasonable computation times for networks with at least as many as 6 airports and 3000 flights. Our formulations refer to static deterministic versions of the problem, but they can be easily extended to cover dynamic versions.

Congestion problems are becoming increasingly acute in many major European and American airports. For European airlines, the total yearly delay cost due to congestion (including cost to passengers) was estimated to be \$8 billion in 1989 [1]. For U.S. airlines, the direct delay cost due to congestion was claimed to be \$2 billion in 1986 [2]. Given the fact that, in 1986, the total profits of the U.S. airline industry were about \$800 million [2], it can be seen that congestion problems are a phenomenon of undeniable significance.

Limited capacity is the major cause of congestion. The problem with airport capacity is that it is highly variable, since it is heavily influenced by, among other factors, weather conditions (visibility, wind, precipitation). It is not unusual to encounter 2:1 and even 3:1 ratios between the highest and the lowest capacity of an airport.

Solution approaches to this problem vary according to the contemplated time-horizon. Long-term approaches include construction of additional airports, construction of additional runways at existing airports, and use of larger aircraft. Medium-

term approaches include modification of the temporal pattern of aircraft flow in order to eliminate periods of "peak" demand. Short-term approaches have a planning horizon of 6-12 hours and include, most importantly, "ground-holding" policies. These policies are based on the fundamental fact that airborne delays are much costlier than ground delays, because the former include fuel, maintenance, depreciation, and safety costs. Thus, the motivation underlying ground-holding policies is that one may hold an aircraft on the ground before take-off so that, when the aircraft arrives at its destination, it will not have to wait in the air before landing.

Ground-holding policies have been in effect for several years. The Federal Aviation Administration operates in Washington, DC, an Air Traffic Control System Command Center (ATCSCC, formerly called the Central Flow Control Facility), equipped with outstanding information-gathering capabilities. ATCSCC, however, relies primarily on the judgement of its expert air traffic controllers rather than on any decision-support or optimization models to develop flow management and ground-holding strategies.

The problem of determining how much (if at all) each aircraft must be held on the ground before take-off (and also, possibly, in the air during the flight, e.g., by means of a speed reduction *en route*) in order to minimize the total (ground plus airborne) delay cost will be referred to as the Ground-Holding Problem (GHP). *Static* and *dynamic* versions of the GHP can be distinguished. In the static versions, the ground (and airborne) holds are decided once for all at the beginning of the day, whereas in the dynamic versions they are updated during the course of the day as better weather (and hence capacity) forecasts become available. *Deterministic* and

probabilistic versions of the GHP can also be distinguished, according to whether airport capacities are considered deterministic or probabilistic.

Because each of a large number of aircraft performs more than one flight on any given day, “network” (or “down-the-road”) effects may be important: when a specific aircraft is delayed, in many cases the next flight performed by the same aircraft will also be delayed. Moreover, at a “hub” airport, a late arriving aircraft may delay the departure of several flights, given current airline scheduling practices which emphasize passenger transfers. To the best of our knowledge, previous research on the GHP has neglected network effects, and has been restricted to the single-airport problem. Odoni [2] seems to be the first to have given a systematic description of the problem. Andreatta and Romanin-Jacur [3] proposed a dynamic programming algorithm for the single-airport static probabilistic GHP with one time period. Terrab [1] proposed an efficient algorithm to solve the single-airport static deterministic GHP, as well as several heuristics for the single-airport probabilistic GHP. He also suggested two formulations for the multi-airport static deterministic GHP. Finally, Richetta [4] dealt with the single-airport dynamic probabilistic GHP. It seems that no significant research has been done to date concerning the effects of ground-holding policies on an entire network of airports.

In this paper, the multi-airport GHP is being addressed for the first time. By using a mathematical programming approach, we solve the deterministic network GHP in a quite general setting. We propose several integer programming formulations which have the important advantages of being remarkably simple while capturing the essential aspects of the problem and of being sufficiently flexible to accommodate

various degrees of modeling detail. We present several structural insights on the parameters that influence the problem, based on extensive computational experience. Most importantly, our approach enables one to solve realistic size problems involving, e.g., 6 airports and 3000 flights in very reasonable computation times. (This should be compared with the previous “record” of about 600 flights for the single-airport GHP [1].) Our approach can thus be used to assign ground holds for at least a major part of the network of the most important U.S. or European airports. Although we focus on the static multi-airport GHP, our algorithms could also be used dynamically by solving the problem, say, every two hours, as better capacity estimates become available.

The outline of this paper is as follows. Section 1 defines the problem and gives integer programming formulations of three versions of it. Section 2 proposes a heuristic based on the solution of a linear programming relaxation. Section 3 gives insights on the parameters influencing the behaviour of the problem, based on an extensive series of actual runs. Finally, Section 4 summarizes the results of the paper and points out directions for future research.

1 Problem definition and formulations.

1.1 Notation.

Consider a set of airports $\mathcal{K} = \{1, \dots, K\}$ and an ordered set of time periods $\mathcal{T} = \{1, \dots, T\}$. For instance, \mathcal{K} might be the set of the 20 or so busiest U.S. airports, and \mathcal{T} might be a set of 64 time periods of 15 minutes each, amounting to a time

horizon of 16 hours, i.e., the portion of a day from 7am to 11pm (when most flights take place). Consider finally a set of flights $\mathcal{F} = \{1, \dots, F\}$. (Note that a single aircraft may perform several of these flights.) \mathcal{F} is the set of all flights of interest, e.g., all flights departing from an airport in \mathcal{K} and arriving to another airport in \mathcal{K} . This interpretation of \mathcal{F} corresponds to a *closed* network of airports, for which departures from and arrivals to the external world are not considered important. If an *open* network of airports is to be considered, then one of the airports in \mathcal{K} must represent the external world.

For each flight $f \in \mathcal{F}$, the following data are assumed to be known: $k_f^d \in \mathcal{K}$, the airport from which f is scheduled to depart; $k_f^r \in \mathcal{K}$, the airport to which f is scheduled to arrive; $d_f \in \mathcal{T}$, the scheduled departure time of f ; $r_f \in \mathcal{T}$, the scheduled arrival time of f ; $c_f^g(\cdot)$, the ground delay cost function of f (whose argument is the ground delay of f in time periods); and $c_f^a(\cdot)$, the airborne delay cost function of f (whose argument is the airborne delay of f in time periods). For each $(k, t) \in \mathcal{K} \times \mathcal{T}$, the departure capacity $D_k(t)$ and the arrival capacity $R_k(t)$ (in number of aircraft) are also given. Since this paper deals with *deterministic* versions of the GHP, these capacities are considered fixed numbers rather than random variables.

Consider finally the set $\mathcal{F}' \subset \mathcal{F}$ of those flights that are “continued”. A flight is said to be continued if the aircraft which is scheduled to perform it is also scheduled to perform at least one more flight later in the day. For each flight $f' \in \mathcal{F}'$, we assume we know the next flight f scheduled to be performed by the same aircraft, and the “slack” or “absorption” time $s_{f'}$, such that, if f' arrives at its destination at most $s_{f'}$ time periods late, the departure of the next flight f will not be affected. $s_{f'}$ is obviously

equal to the difference between (i) the time interval between the scheduled departure time of f and the scheduled arrival time of f' , and (ii) the minimum “turnaround” time of the aircraft performing both flights.

1.2 Preliminary remarks.

We define the decision variables $g_f, f \in \mathcal{F}$, equal to the number of time periods that flight f is held on the ground before being allowed to take off, and the decision variables $a_f, f \in \mathcal{F}$, equal to the number of time periods that flight f is further held in the air (e.g., by means of a speed reduction *en route*) before being allowed to land. Since this paper deals with *static* versions of the GHP, it is assumed that these ground and airborne holds are decided once for all at the beginning of the day for all flights.

Consider now the following description of the real-world situation. If a flight f is scheduled to depart at period d_f and is delayed on the ground for g_f periods, then it will be *available to depart* at period $d_f + g_f$. Will it actually depart at that period? This will depend on whether the total number of aircraft *available to depart* from airport k_f^d at that time period will exceed or not the available departure capacity. If it does exceed it, then the aircraft performing flight f will have to wait q_f^d time periods in the *departure queue*. q_f^d will depend on the particular *service discipline* adopted for the departure queue. So flight f will actually take off at period $d_f + g_f + q_f^d$. Since flight f will be further delayed in the air for a_f time periods, it will arrive at its destination, airport k_f^r , and will be *available to land* at period $r_f + g_f + q_f^d + a_f$. Will it actually land at that period? This will depend on whether the total number of aircraft *available to land* at airport k_f^r at that period will exceed or not the available landing

capacity. If it does exceed it, then the aircraft performing flight f will have to wait q_f^a time periods in the *arrival queue*, and will actually land at period $r_f + g_f + q_f^d + a_f + q_f^a$. The total cost corresponding to flight f will be the sum of $c_f^g(g_f + q_f^d)$ (the ground delay cost) and $c_f^a(a_f + q_f^a)$ (the airborne delay cost).

Because we are examining the *deterministic* case, the above description can be considerably simplified. In fact, it makes little sense to assign to a flight f a ground hold of g_f time periods such that f will have to further wait q_f^d time periods in the departure queue: one might as well assign to f a total ground hold of $g_f + q_f^d$ time periods such that f will not have to wait at all in the departure queue. Similar remarks hold for airborne delays. Given this simplification, the total ground delay of flight f will be g_f , and its total airborne delay will be a_f , resulting in a cost of $c_f^g(g_f) + c_f^a(a_f)$. This leads us to the following integer programming formulation of the static deterministic multi-airport GHP.

1.3 A pure 0–1 integer programming formulation of the multi-airport GHP.

$$(P_1) \min \quad \sum_{f=1}^F c_f^g g_f + c_f^a a_f$$

$$\text{s.t.} \quad \sum_{f:k_f^d=k} u_{ft} \leq D_k(t), \quad (k, t) \in \mathcal{K} \times \mathcal{T}; \quad (1)$$

$$\sum_{f:k_f^a=k} v_{ft} \leq R_k(t), \quad (k, t) \in \mathcal{K} \times \mathcal{T}; \quad (2)$$

$$\sum_{t \in \mathcal{T}_f^d} u_{ft} = 1, \quad f \in \mathcal{F}; \quad (3)$$

$$\sum_{t \in \mathcal{T}_f^a} v_{ft} = 1, \quad f \in \mathcal{F}; \quad (4)$$

$$g_{f'} + a_{f'} - s_{f'} \leq g_f, \quad f' \in \mathcal{F}'; \quad (5)$$

$$a_f \geq 0, \quad f \in \mathcal{F}; \quad (6)$$

$$u_{ft}, v_{ft} \in \{0, 1\};$$

where u_{ft} is 1 if flight f finally takes off at period t (i.e., if $r_f + g_f = t$) and 0 otherwise, and v_{ft} is 1 if flight f finally lands at period t (i.e., if $r_f + g_f + a_f = t$) and 0 otherwise. For every flight f , exactly one of the variables u_{ft} must be equal to 1 and the others must be equal to zero, and similarly for the variables v_{ft} . (This is ensured by constraints (3) and (4).) Given this fact, the *delay variables* g_f can be expressed in terms of the *assignment variables* u_{ft} in the following way:

$$g_f = \sum_{t \in \mathcal{T}_f^d} tu_{ft} - d_f, \quad f \in \mathcal{F}, \quad (7)$$

where \mathcal{T}_f^d is the set of time periods to which flight f may be assigned to take off, given by: $\mathcal{T}_f^d = \{t \in \mathcal{T} : d_f \leq t \leq \min(d_f + G_f, T)\}$. This simply says that flight f cannot take off before period d_f , its scheduled departure time, nor can it take off after $d_f + G_f$, where G_f is an upper bound on the acceptable ground delay of f .

Similarly, the *delay variables* a_f can be expressed in terms of the *assignment variables* v_{ft} in the following way:

$$a_f = \sum_{t \in \mathcal{T}_f^s} tv_{ft} - r_f - g_f, \quad f \in \mathcal{F}, \quad (8)$$

where \mathcal{T}_f^s is the set of time periods to which flight f may be assigned to land, given

by: $T_f^a = \{t \in \mathcal{T} : r_f \leq t \leq \min(r_f + A_f, T)\}$, where A_f is an upper bound on the acceptable airborne delay of f .

For simplicity of exposition, variables g_f and a_f were kept in formulation (P_1) , but it should be clear that they can be eliminated by mere substitution through (7) and (8), so that u_{ft} and v_{ft} are the only decision variables.

Note that nonnegativity of g_f is guaranteed by (7), whereas nonnegativity of a_f is not guaranteed, this is why constraints (6) are needed.

In the objective function of (P_1) , the cost functions $c_f^g(t), c_f^a(t)$ were replaced by their linear counterparts $c_f^g t, c_f^a t$ (c_f^g, c_f^a being the constant marginal costs). Constraints (1) and (2) are the departure and arrival capacity constraints, respectively. Recall that these have to be satisfied because we choose g_f and a_f such that the queuing delays q_f^d, q_f^a are 0 (and that we can do this because the problem is deterministic). (Strictly speaking, we also need the condition that G_f and A_f be sufficiently large.) Constraints (5) are the *coupling* constraints: they “transfer” any excessive delay of flight f' to its next flight f . In fact, the coupling constraints say that, if flight f' arrives at its destination with a total delay $g_{f'} + a_{f'}$ which is greater than $s_{f'}$ (the “slack” defined above), then the next flight f will have to be delayed on the ground at least $g_{f'} + a_{f'} - s_{f'}$ time periods; otherwise, the departure of the next flight f will not be affected. Note that the existence of these coupling constraints allows us to have a separable objective function: the cost of delaying flight f because of an excessive delay of its previous flight f' is taken into account via the term of the objective function corresponding to f (i.e., $c_f^g g_f$), and so need not be included in the term corresponding to f' . Note also that, if the coupling constraints did not exist, the problem would be

decomposable into K subproblems concerning one airport each, so that one could use the already existing techniques to solve for each of the K airports separately. A final interesting remark concerning the coupling constraints is that they can be interpreted in a more general way than the linking of successive flights scheduled to be performed by the same aircraft; i.e., they can be used to link any pair of flights f' and f such that f cannot be allowed to depart before f' has arrived (possibly because passengers in f' connect to f). In this interpretation, a flight f' may have more than one “next” flights f . This interpretation will *not* be pursued in the sequel.

1.4 A simpler case: infinite departure capacities and zero airborne delays.

Formulation (P_1) is about as general as one could wish for the static deterministic case, but it can be simplified considerably without significant loss of practical applicability. Note, first, that it is usually undesirable to delay aircraft in the air. In fact, the fundamental goal of ground holding policies is to avoid this kind of delay. Therefore, we may eliminate airborne delays *as decision variables*. We will be left then with airborne delays resulting only from arrival queueing (denoted earlier by q_f^a), and our only decision variables will be g_f . (Note that, since the problem is deterministic, q_f^a are determined if g_f and service disciplines for the arrival queues are given.)

Now consider a feasible solution $\{g_f, f \in \mathcal{F}\}$ and the associated arrival queueing delays $\{q_f^a, f \in \mathcal{F}\}$, and compare its cost with the cost of the new solution $\{g_f + q_f^a, f \in \mathcal{F}\}$, in which all airborne delays are incorporated into ground holding delays. Given

that the cost functions are linear, and given that airborne delays are costlier than ground delays (i.e., for any positive t , and for all f , $c_f^a(t) > c_f^g(t)$), it is easy to show that the new solution will have a lower cost than the previous solution. In fact, $c_f^g(g_f + q_f^a) = c_f^g(g_f) + c_f^g(q_f^a) < c_f^g(g_f) + c_f^a(q_f^a)$. It would appear, therefore, that airborne delays need not be taken into consideration, because they will never appear in an optimal solution. The problem, however, is that the new solution $\{g_f + q_f^a, f \in \mathcal{F}\}$ will not necessarily satisfy the departure capacity constraints. This problem disappears if we assume that all departure capacities are infinite, but are we entitled to make this assumption? For practical purposes, this assumption may often be a good approximation, since congestion problems are mostly due to limited landing rather than departure capacities. Moreover, computational experience reported in Section 3 shows that the impact of finite departure capacities is indeed negligible. This *a posteriori* argument justifies the assumption of infinite departure capacities. Note that, in the single-airport case, which is the only case considered so far in the literature, no departure capacities are involved, so that one is in fact rigorously justified to consider only feasible solutions with zero airborne delays (provided the problem is deterministic and the cost functions are linear).

Assuming infinite departure capacities eliminates thus airborne delays and gives the following pure 0-1 integer programming formulation of the static deterministic multi-airport GHP :

$$(P_2) \min \quad \sum_{f=1}^F c_f^g g_f$$

$$\begin{aligned}
\text{s.t. } \sum_{f:k_f^a=k} v_{ft} &\leq R_k(t), \quad (k,t) \in \mathcal{K} \times \mathcal{T}; \\
\sum_{t \in \mathcal{T}_f^a} v_{ft} &= 1, \quad f \in \mathcal{F}; \\
g_{f'} - s_{f'} &\leq g_f, \quad f' \in \mathcal{F}'; \\
v_{ft} &\in \{0,1\}, \quad f \in \mathcal{F}, t \in \mathcal{T}_f^a;
\end{aligned} \tag{9}$$

where g_f is determined by setting $a_f = 0$ in (8):

$$g_f = \sum_{t \in \mathcal{T}_f^a} t v_{ft} - r_f, \quad f \in \mathcal{F}, \tag{10}$$

and variables u_{ft} are now redundant, so that v_{ft} are the only decision variables.

(It can be seen that now G_f and A_f express the same bound, so that they must be taken to be equal.)

Note the simplicity of formulation (P_2). The number of constraints is $F + F' + KT$, and the number of variables is at most $\sum_{f \in \mathcal{F}} (G_f + 1)$ which, if all G_f are equal to 4 (corresponding to a maximum ground hold of one hour), becomes $5F$. Therefore, the total number of flights F is the major determinant of the size of the problem. The number of time periods T has almost no influence on the size of the problem, and the same holds for the number of airports K . Of course, the number of airports has an indirect influence on the size of the problem, since it influences the number of flights to be considered. Typically, a major U.S. airport has 1000–2000 operations (landings plus take-offs) each day, corresponding to 500–1000 flights per day. But still, the fact that the problem is insensitive as to how the total number of flights is distributed among airports and time periods is very welcome. This becomes clear in

dynamic versions of the ground-holding problem (not treated in this paper), where the time horizon is limited to a portion of a day, so that fewer flights per airport have to be considered, and it becomes possible to solve the problem for a large number of airports.

Note, finally, that, if the coupling constraints (9) are omitted from the formulation, what is left is exactly the single-airport formulation given, e.g., in [1]. It follows that the coupling constraints (9) are the gist of the model. It is indeed surprising that the network effects can be taken into account in such a simple way without loss of generality. *Simplex sigillum veri* (“the simple is the sign of the true”).

1.5 How to handle infeasibility: Cancelling flights.

If the arrival capacities are low, then formulation (P_2) may become infeasible. Even though the total daily capacity of an airport may be sufficient to accommodate the total number of flights scheduled to arrive at that airport, the problem may still be infeasible if excessive congestion appears during some portion of the day. This is mainly due to the requirement that there be an upper bound, G_f , to the delay of flight f . In order to grasp this point, take the extreme case where the landing capacity of an airport is reduced to zero for $G_f + 1$ successive time periods. Then, if a flight was scheduled to arrive exactly before the zero capacity interval, it will be impossible to reassign this flight and the problem will become infeasible. Similar remarks hold for formulation (P_1).

In situations where delays become excessive, it is common airline practice to cancel some flights. Motivated from this, we developed an alternative formulation which

eliminates infeasibility problems and takes into account the possibility of cancelling flights.

The new formulation will be presented as a generalization of (P_2) . We keep the old decision variables v_{ft} and we define the decision variables $z_f, f \in \mathcal{F}$, to be 1 if flight f is cancelled and 0 otherwise. Denote by M_f the *cancellation cost* of flight f . When a flight in \mathcal{F}' (i.e., a flight that is “continued”) is cancelled, there are two possibilities concerning the next flight initially scheduled to be performed by the same aircraft: either it is performed by a replacement (or a “spare”) aircraft, or it is also cancelled. The first case is more common in practice, but our formulation is general enough to incorporate a combination of both cases. Partition \mathcal{F}' into \mathcal{F}'_1 , the set of those flights in \mathcal{F}' whose cancellation will not affect their next flight, and \mathcal{F}'_2 , the set of those flights in \mathcal{F}' whose cancellation will entail the cancellation of their next flight. We will now first give the new formulation and then comment on it.

$$(P_3) \min \quad \sum_{f=1}^F (c_f^g g_f + (M_f + c_f^g r_f) z_f) \quad (11)$$

$$\text{s.t.} \quad \sum_{f:k_f^g=k} v_{ft} \leq R_k(t), \quad (k, t) \in \mathcal{K} \times \mathcal{T}; \quad (12)$$

$$z_f + \sum_{t \in \mathcal{T}_f^g} v_{ft} = 1, \quad f \in \mathcal{F}; \quad (13)$$

$$g_{f'} - s_{f'} + (s_{f'} + r_{f'} - r_f) z_{f'} \leq g_f, \quad f' \in \mathcal{F}'_1; \quad (14)$$

$$g_{f'} - s_{f'} + (s_{f'} + r_{f'} + G_f + 1) z_{f'} \leq g_f + (r_f + G_f + 1) z_f, \quad f' \in \mathcal{F}'_2; \quad (15)$$

$$v_{ft}, z_f \in \{0, 1\}. \quad (16)$$

The above formulation incorporates some technical tricks which are necessitated by the fact that, when a flight f is cancelled (i.e., $z_f = 1$), then all v_{ft} corresponding to f are 0 (by (13)), so that (10) gives $g_f = -r_f$. Keeping this fact in mind, it can be seen immediately that, when $z_f = 1$, the objective function term corresponding to f is M_f . It is also clear that, when $z_{f'} = 1$, (14) becomes $-r_f \leq g_f$, which holds even if flight f is cancelled (so that cancellation of f' leaves f unaffected). Finally, if $z_{f'} = 1$, (15) becomes $G_f + 1 \leq g_f + (r_f + G_f + 1)z_f$, entailing $z_f = 1$ (since $g_f \leq G_f$ always), which is precisely what we wanted: if f' is cancelled, then f is also cancelled.

The variables g_f were again left in the formulation, but it should be clear that they can be eliminated by mere substitution through (10). It is important to notice that the variables z_f can also be eliminated through (13), provided that (13) is replaced by $\sum_{t \in \mathcal{T}_f^s} v_{ft} \leq 1$. In this case, it can be seen that the objective function becomes: $\min \sum_{f=1}^F [M_f + \sum_{t \in \mathcal{T}_f^s} v_{ft}(c_f^g(t - r_f) - M_f)]$. It follows that the only decision variables left are v_{ft} , so that the new formulation (P_3) has exactly the same number of variables and of constraints as the previous formulation (P_2). (P_3), however, enjoys considerable advantages both in terms of generality (the real-world problem is better approximated) and in terms of flexibility (infeasibility problems are eliminated).

This section has presented three quite general and, we believe, elegant formulations of the static deterministic GHP.

2 A heuristic.

This section presents a heuristic which finds a feasible solution of the integer program (P_3) starting from a feasible solution of the linear programming (LP) relaxation of (P_3) . The next section will show, on the basis of computational experience, that it is easy to solve optimally the LP relaxation of (P_3) and that, when one applies the heuristic to this optimal solution, one gets a "good" feasible solution of the integer program (P_3) .

The heuristic will be presented in rough outline here. An algorithmic presentation is given in the appendix.

Consider a feasible solution $\{v_{ft} : f \in \mathcal{F}, t \in \mathcal{T}_f\} \cup \{z_f : f \in \mathcal{F}\}$ of the LP relaxation of (P_3) and denote by Φ the set of "problematic" flights $f \in \mathcal{F}$, i.e., the set of flights for which some integrality constraint is violated. The heuristic gives a "rounding" scheme for flights in Φ which leaves undisturbed, as far as possible, the remaining flights (which already satisfy integrality). The basic idea of the heuristic is to treat each flight in Φ once for all.

The heuristic starts by partitioning Φ into classes, each class corresponding to an aircraft and containing all and only the flights of Φ scheduled to be performed by that aircraft. The heuristic treats each class separately; the order in which the classes are treated is more or less arbitrary.

Each class is treated in the following way. The flights in the class are examined one at a time, in the order in which they are scheduled to be performed by the aircraft defining the class. For each specific flight ϕ , the heuristic takes the following actions.

For each time period t at which ϕ can be allowed to land, it computes the available “capacity slacks” $R_{k_\phi^a}(t) - \sum_{f:k_f^a=k_\phi^a} v_{ft}$, which will be denoted by $S_\phi(t)$. (If some v_{ft} have already been updated by new values, then the new values are used in the computation of the capacity slacks.) It can be seen that, if $S_\phi(t) \geq 1 - v_{\phi t}$, then it is possible to assign flight ϕ to period t without violating the corresponding capacity constraint. If this is possible for no t , then flight ϕ is cancelled and we are done with it. Otherwise, i.e., when there are time periods to which it is possible to assign flight ϕ without violating the corresponding capacity constraint, flight ϕ is assigned to the earliest such period, τ . (Recall that this assignment is made once for all.) After this assignment, all constraints involving flight ϕ are satisfied, with the possible exception of the coupling constraints.

In order to deal with the coupling constraint linking flight ϕ with its next flight $\hat{\phi}$ (if such a next flight exists), the heuristic removes certain time periods from the set of time periods at which $\hat{\phi}$ can be allowed to land, and proceeds to examine $\hat{\phi}$. The removed time periods are those which would violate the coupling constraint in question if $\hat{\phi}$ were assigned to them (given that ϕ has already been assigned to τ). It can be seen that, if flight ϕ has a *previous* flight ϕ' , the coupling constraint linking ϕ' and ϕ need not be dealt with while examining flight ϕ , because it has already been dealt with when examining flight ϕ' (since ϕ is the *next* flight to ϕ' .)

As pointed out above, this is only a rough outline, and a more rigorous formulation is given in the appendix.

3 Structural insights.

This section investigates the behaviour of the GHP on the basis of extensive computational experience. The investigation is conducted in three parts, each part dealing with one of the formulations (P_1) , (P_2) , and (P_3) . For each formulation, we examine the variation, as a function of the input parameters, of the optimal objective function values of the following three mathematical optimization problems: the integer program (denoted by I), the corresponding linear programming relaxation (denoted by L), and the “decomposed” program (denoted by D), defined as the integer program without the coupling constraints.

It is important to understand the role of D in the comparison. The “decomposed GHP” corresponding, e.g., to formulation (P_2) is simply (P_2) without the coupling constraints (9). Solving the decomposed GHP is equivalent to solving the GHP for each airport separately, and then adding the optimal objective function values corresponding to the various airports. Note that the optimal objective function value of the decomposed GHP is equal to the optimal objective function value of the LP relaxation of the decomposed GHP, because the constraint matrix of any single-airport GHP is totally unimodular [1]. Therefore, D can be defined as a linear rather than an integer program.

Denote the optimal values of I , L , and D by v_I , v_L , and v_D , respectively. Now the greater the gap between v_D and v_I (and, *a fortiori*, the greater the gap between v_D and v_L), the greater the impact of the network effects. A large gap between v_D and v_I presumably justifies one in pursuing the application of algorithms pertaining

to the multi-airport (coupled) GHP rather than solving for each airport separately by means of the existing methods for the single-airport GHP. This much is clear. What is less clear is how a very *small* gap between v_D and v_I should be interpreted. A small gap would not necessarily mean that the multi-airport GHP is valueless. Consider the extreme case where $v_D = v_I$. The zero gap means that we could ignore the coupling constraints without any change in the optimal value of I . But if D has multiple optimal solutions, then solving it will not necessarily give a solution satisfying the coupling constraints, i.e., a solution *feasible* for I .

It should be noted that the objective of this section is to investigate the behaviour of the problem under various combinations of the input parameters, *not* to demonstrate the efficiency of any particular algorithm. In fact, we solved the various instances of the problem by using the well-known commercial package *MPSX*, rather than any custom-tailored algorithm. We give CPU times simply in order to indicate whether the problem can be solved in reasonable time, rather than in order to provide any "good" bounds on computation times.

This section is divided into three subsections. The basic conclusions are reached in the first subsection, which deals with formulation (P_2). The second subsection, which deals with formulation (P_1), verifies that the impact of finite departure capacities would be negligible in many practical cases. Finally, the third subsection deals with formulation (P_3) (with flight cancellations) and with the performance of the heuristic.

3.1 The model without flight cancellations.

This subsection deals with formulation (P_2). We consider first a test case with $K = 3$ airports, $T = 100$ time periods, $F = 1800$ flights (600 flights per airport), and $F' = 600$ flights. With the exception of capacities, all parameters were kept fixed in this test case: the cost function slopes were 50, the slacks were 0, the upper bounds on the delays were 4 time periods, and the scheduled arrival times were arbitrarily chosen.

As mentioned in Section 1, if arrival capacities are very low, the problem becomes infeasible. Let us consider only cases in which the arrival capacity of any given airport is constant over the whole time horizon: $R_k(t) = R_k$. Then it is found that, for the particular test case under consideration, for $(R_1, R_2, R_3) = (10, 10, 10)$ the problem is feasible, while for $(9, 9, 9)$ the problem is infeasible. Furthermore, for $(9, 10, 10)$, $(10, 10, 9)$, $(9, 10, 9)$, and $(10, 10, 8)$ the problem is feasible, while for $(10, 9, 10)$, $(8, 10, 10)$, and $(10, 10, 7)$ the problem is infeasible. These results give us a fairly good picture of the border between capacity regions corresponding to feasibility and to infeasibility for the test case under consideration. Delimitation of this border is important because it is there that the greatest delays are expected to occur: if capacities are very high, then there is little need to delay aircraft.

Table 1 gives the optimal objective function values of L , D , and I for the various capacity cases. It is seen that these values always turn out to be very close. An examination of the optimal solution of D , however, reveals that usually about 180 to 200 of the 600 coupling constraints are violated. It follows that solving the decomposed

Capacities	v_D	v_L	v_I	# of coupling constr. that D violates	% of $f \in \mathcal{F}'$ delayed in I	% of $f \in \mathcal{F} - \mathcal{F}'$ delayed in I
(10,10,10)	43,550	43,550	43,550	179	12%	30%
(9,10,10)	51,900	52,800	52,900	204	18%	36%
(10,10, 9)	48,500	49,000	50,600	183	17%	34%
(9,10, 9)	56,850	57,450	57,950	238	20%	40%
(10,10, 8)	55,650	56,700	58,000	235	19%	37%

Table 1: Behaviour of the test case around the capacity border between feasibility and infeasibility.

problem is probably of little use as far as getting a feasible solution to the coupled problem is concerned. Nevertheless, solving the decomposed problem provides a good indicator of what the optimal value of the coupled problem will be.

The proximity of v_D and v_I needs an explanation, but we must first ascertain that it is a common phenomenon rather than a peculiar feature of the particular test case under consideration. To this end, we examined a systematic series of test cases. In all these cases, T is kept fixed and equal to 64 (corresponding to a 16-hour time horizon with 15-minute periods), and K is determined by F via the assumption that 500 flights are scheduled to land at each airport during the time horizon. Three cases for F are examined: 1000, 2000, and 3000 flights (corresponding, respectively, to 2, 4, and 6 airports). For each particular F , four values of F' are examined, corresponding to a ratio F'/F equal to 0.20, 0.40, 0.60, and 0.80. The results are summarized in Table 2. The capacities appearing in the table for any particular case are at the infeasibility borders (and were found by trial and error). The cost function slopes are always 50, all slacks are 1, and all upper bounds on delays are 4.

These results lead to the following conclusions. First, the gap between v_D and

F	F'/F	Capac.	v_D	t_D	v_L	t_L	v_I	t_I	
1000	0.20	(12,14)	71,000	218	71,000	258	71,000	371	63
1000	0.40	(10,10)	56,000	235	56,000	327	56,000	894	84
1000	0.60	(11,11)	84,200	242	84,300	377	84,700	6958	168
1000	0.80	(10,10)	65,000	235	65,000	453	65,500	9512	128
2000	0.20	all 14	96,300	664	96,300	731	99,000	5126	117
2000	0.40	all 14	88,400	652	89,933	973	93,200	9522	195
2000	0.60	all 12	71,600	644	71,600	1148	71,800	13607	252
2000	0.80	all 17	53,250	617	57,387	1603	65,500	18093	355
3000	0.20	all 12	128,000	1188	129,200	1453	129,400	11360	110
3000	0.40	all 18	55,800	1208	55,800	1808	57,300	13291	119
3000	0.60	all 17	90,200	1166	96,550	2547	99,687	17980	414
3000	0.80	all 18	80,500	1180	84,250	3072	87,012	25021	232

Table 2: Results for various cases at the infeasibility border.

v_I is always small. Second, the computation times (given in CPU seconds) t_D and t_I are quite reasonable, but t_I can become excessive. Third, as one would expect, the computation times increase as F increases, because the number of constraints and the number of variables increase. Fourth, for any given F , t_D does not vary significantly with F' , while t_L and t_I increase as F' increases. This is due to the fact that an increase in F' increases the number of constraints of L and I (which have $KT + F + F'$ constraints), while it leaves unaffected the number of constraints of D ($KT + F$). Finally, the last column in Table 2 gives the number of flights for which the optimal solution of L had noninteger values. It can be seen that this number is usually small, around 10% of F . This observation provided the motivation for the development of the heuristic given in Section 2.

Now we must explain the fact that v_D and v_I are typically very close. Two possible explanations suggest themselves. The first is that the capacities at the border between feasibility and infeasibility, although they cannot be lowered in the context of the

present model, are still too high for network effects to have a severe impact. This explanation, if true, would undermine the utility of the model as a representation of the real-world situation. This explanation, however, is not true. First, v_D and v_I are very close even with quite low capacities (see the second and the fourth rows of Table 2). Second, in Subsection 3.3, where formulation (P_3), which is immune to infeasibility, is examined, it will be seen (cf. fifth row of Table 3) that v_L and v_D are very close even with capacities as low as 256 aircraft per airport per day (4 per period) (again, with 500 aircraft scheduled to land, so that the remaining flights are cancelled).

The second possible explanation for the proximity of v_D and v_I is that, if there is a choice (in I) between delaying a continued flight and a noncontinued flight, it will usually be preferable to delay the latter, since delaying the former would probably result in a greater total cost (because the next flight might also have to be delayed). If this is the case, then, in the optimal solution of I , few flights in \mathcal{F}' will be delayed. This effect would be particularly noticeable for small slacks. A look at the last two columns of Table 1 corroborates this hypothesis. A second way to confirm this hypothesis is by varying the cost function slopes so as to disadvantage continued flights. If continued flights have much lower marginal costs than noncontinued flights, then it may often be preferable to delay a continued rather than a noncontinued flight when a choice is available, with the consequence that network effects may be significant. The test case with 1800 flights was run with capacities equal to 10 and with cost function slopes equal to 10 for the continued flights and equal to 100 for the noncontinued flights, and the results were: $v_D = 13,950$ and $v_L = 22,811$, a very

significant gap. (Cf. also Table 6 in Subsection 3.3.)

This subsection reached the following conclusions. First, the gap between v_D and v_I is usually very small. Second, the explanation seems to be that, given equal marginal costs, continued flights are typically not delayed in optimal solutions of I , so that network effects are not very significant. Third, although network effects do not have a very significant impact on the optimal objective function values, solving the decomposed GHP appears to be of little use for the network problem, because optimal solutions of D typically violate a large percentage of the coupling constraints of I .

3.2 The negligible impact of finite departure capacities.

In order to check the impact of finite departure capacities and to demonstrate that formulation (P_1), which has about twice as many variables and constraints as formulation (P_2), can be also solved in reasonable computation times, we examined the problems of the first two rows of Table 2 with various departure capacities. In order to make meaningful comparisons, we kept the scheduled arrival times unchanged. The new data, besides the departure capacities, were the scheduled departure times or, equivalently, the flight times. Table 3 gives results for various combinations of departure capacities and flight times. Airborne marginal delay costs were taken to be 75 versus ground marginal delay costs of 50.

It can be seen from Table 3 that, when flight times are uniform (e.g., all equal to 2 time periods) or slightly nonuniform, the differences between finite and infinite departure capacities are negligible. It is only with wildly nonuniform flight times

F	F'/F	Arr. Cap.	Dep. Cap.	Flight times	v_L
1000	0.20	(12,14)	∞	—	71,000
1000	0.20	(12,14)	(12,14)	Uniform: 2	71,000
1000	0.20	(12,14)	(15,17)	Nonuniform: 1 or 2	71,500
1000	0.40	(10,10)	∞	—	56,000
1000	0.40	(10,10)	(10,10)	Uniform: 2	56,000
1000	0.40	(10,10)	(15,15)	Nonuniform: 1 to 30	64,965
1000	0.40	(10,10)	(16,16)	Nonuniform: 1 to 30	57,633

Table 3: Results for various cases with finite departure capacities.

that some minor differences appear. These results justify us in pursuing our investigation with the more manageable formulation (P_2), although formulation (P_1) is also manageable (running times for the cases of Table 3 were about 2000 CPU seconds).

3.3 The model with flight cancellations.

Table 4 gives results for selected cases from Table 2, but for formulation (P_3), and for various capacities and cancellation costs M . The rows with "infinite" cancellation costs correspond to formulation (P_2) and are taken from Table 2. All marginal delay costs were equal to 50.

These results strongly support the conclusion that, for cancellation costs greater than 100 times the marginal delay cost (i.e., here, $M > 5000$), no flight is ever cancelled, so that models (P_2) and (P_3) give the same results. For cancellation costs greater than 20 times the marginal delay costs ($M > 1000$), few flights are cancelled, so that the optimal values of (P_2) and (P_3) are very close. Finally, for cancellation costs less than 10 times the marginal delay cost ($M < 500$), more flights are cancelled and significant differences between (P_2) and (P_3) emerge. Note also that, in that

F	F'/F	Cap.	M	v_D	t_D	v_L	t_L	v_H
1000	0.60	11	1000	70,300	297	70,300	479	78,500
1000	0.60	10	1000	117,000	286	117,000	475	125,450
1000	0.60	08	1000	240,700	280	241,805	524	253,750
1000	0.60	06	1000	402,600	274	403,476	513	411,500
1000	0.60	04	1000	582,300	272	583,417	484	586,700
1000	0.60	11	100	28,700	283	28,700	498	30,250
1000	0.60	11	1000	70,300	297	70,300	473	78,500
1000	0.60	11	10000	84,200	276	84,300	444	240,700
1000	0.60	11	∞	84,200	242	84,300	377	—
2000	0.20	14	500	77,500	652	77,500	803	82,000
2000	0.20	14	1000	94,000	691	94,000	922	103,300
2000	0.20	14	5000	96,300	717	96,300	931	165,200
2000	0.20	14	∞	96,300	664	96,300	731	—
2000	0.40	14	500	73,100	815	74,983	1020	75,800
2000	0.40	14	1000	86,100	690	86,372	1102	93,650
2000	0.40	14	5000	88,400	675	89,933	1176	168,900
2000	0.40	14	∞	88,400	652	89,933	973	—
3000	0.60	17	100	38,250	1119	38,693	1911	42,350
3000	0.60	17	500	71,800	1128	72,240	1708	84,600
3000	0.60	17	750	81,000	1148	81,338	1931	95,000
3000	0.60	17	1000	87,000	1187	87,156	2114	130,300
3000	0.60	17	10000	90,200	1248	96,550	3767	667,750
3000	0.60	17	∞	90,200	1166	96,550	2547	—
3000	0.80	18	100	36,600	1114	38,042	1846	58,900
3000	0.80	18	500	71,500	1140	71,559	2320	83,350
3000	0.80	18	750	78,700	1128	78,707	2693	106,800
3000	0.80	18	1000	80,500	1235	82,214	2900	111,350
3000	0.80	18	10000	80,500	1230	84,250	3227	509,900
3000	0.80	18	∞	80,500	1180	84,250	3072	—

Table 4: Results for various cases with flight cancellations.

F	F'/F	Cap.	M	v_D	t_D	v_L	t_L
3000	0.80	nonun.	500	232,800	1142	252,045	1973
3000	0.80	nonun.	750	302,700	1200	330,040	2217
3000	0.80	nonun.	1000	366,200	1215	403,127	2228

Table 5: Results for various cases with flight cancellations and nonuniform capacities.

last region of cancellation costs, the slope of the optimal value as a function of the cancellation cost becomes quite abrupt.

The last column of Table 4 shows the value v_H of the objective function corresponding to the feasible solution found by the heuristic. It can be seen that v_H is quite close to v_L (hence to v_I) for small cancellation costs. For large cancellation costs, however, the heuristic performs poorly. This was to be expected, because the heuristic will inevitably cancel some flights, and these will inflate the objective function value if the cancellation cost is excessive. This is not worrisome, however, since, as was pointed out above, for cancellation costs above 1,000 few flights are cancelled, so that for such high cancellation costs neither formulation (P_3) nor the heuristic have much practical use.

Table 5 gives results concerning cases with nonuniform arrival capacities. It can be seen that gaps between v_D and v_I are quite significant.

Finally, another reason why the optimal values of L and D are so close might be the assumption that the costs of all flights are equal. In order to check this, we ran some cases with three classes of costs: 40% of all flights had cost 100, 40% had cost 50, and 20% had cost 20, corresponding to the relative direct operating costs of large, medium-sized, and small aircraft, respectively. Aircraft performing continued flights

F	F'/F	Cap.	M	v_D	t_D	v_L	t_L
3000	0.60	nonun.	500	305,690	1236	373,271	2099
3000	0.60	nonun.	750	385,830	1253	491,791	2219
3000	0.60	nonun.	1000	460,230	1305	601,331	2332

Table 6: Results for various cases with flight cancellations and three cost classes.

were generally assigned to the large- or medium-cost category. Results are shown in Table 6. It can be seen that the differences are significant (22-27%).

The last column of Table 4 shows the value v_H of the objective function corresponding to the feasible solution found by the heuristic. It can be seen that v_H is quite close to v_L (hence to v_I) for small cancellation costs. For large cancellation costs, however, the heuristic performs poorly. This was to be expected, because the heuristic will inevitably cancel some flights, and these will inflate the objective function value if the cancellation cost is excessive. This is not worrisome, however, since, as was pointed out above, for cancellation costs above 1,000 few flights are cancelled, so that for such high cancellation costs neither formulation (P_3) nor the heuristic have much practical use.

4 Conclusions.

The multi-airport GHP was shown to be tractable. Our formulations capture the essential aspects of the problem, for the static deterministic case at least, and do so in a very simple way. It is this simplicity, reflected in the small numbers of constraints and variables, that is responsible for the tractability of large-scale GHPs.

It is not yet clear how large a network one can deal with by means of our formu-

lations. We went up to 6 airports and 3,000 flights, but one could probably go far beyond this if one were willing to use supercomputers. This would not be unrealistic, given the importance of the practical problem.

A principal conclusion arising from the computational experience we presented is that network effects are not very significant if all cost functions are assumed identical. This assumption is incorrect, since the delay of large aircraft is more costly than the delay of small aircraft. However, the practice of implicitly considering all cost functions identical seems to be a well-entrenched practice of the FAA, which avoids "discriminating" in any way among users.

A second basic insight was that the relative insignificance of the network effects does not imply that network formulations are useless. In fact, the opposite is true, because solutions of the decomposed problem typically violate a large number of the coupling constraints.

A third important result was that departure capacities can be taken to be infinite in most practical cases, so that airborne delays can be avoided and the size of the formulations can be significantly reduced.

A direction for future research is to extend our formulations to the dynamic deterministic case. We have already performed this extension, which is relatively straightforward, although it needs to incorporate some subtleties due to the fact that airborne delays cannot be totally avoided in the dynamic case.

Another interesting direction of research, referring to the dynamic case, would be to run our formulations for limited time horizons of, say 2 hours. This would dramatically decrease the size of the problem for a given number of airports, and so

would enable one to tackle much larger networks of airports.

The most challenging direction for future research is, of course, the case of probabilistic airport capacities (see [4] for the single-airport problem). This case may require a totally new approach.

Appendix: Algorithmic description of the heuristic.

The heuristic takes as input a solution $\{v_{ft} : f \in \mathcal{F}, t \in \mathcal{T}_f\} \cup \{z_f : f \in \mathcal{F}\}$ which is feasible for the LP relaxation of (P_3) , and gives as output a solution which is feasible for (P_3) . The heuristic is presented here for the case in which, when a flight is cancelled, the next flight scheduled to be performed by the same aircraft is not affected. The other case, in which the next flight is also cancelled, can be treated *mutatis mutandis*.

BEGIN

Define $\Phi := \{\phi \in \mathcal{F} : (z_\phi \notin \{0, 1\}) \vee (\exists t)(v_{\phi t} \notin \{0, 1\})\}$.

Partition Φ into its equivalence classes corresponding to the equivalence relation "is performed by the same aircraft as": $\Phi = \bigcup_{\psi=1}^{\Psi} \Phi_\psi$.

Order each class according to the order in which the flights in the class are scheduled to be performed by the aircraft defining the class: $\Phi_\psi = \{\phi_{\psi 1}, \dots, \phi_{\psi, \Xi(\psi)}\}$.

Order the classes, e.g., in decreasing order of the cost of their first flight, and break ties, e.g., according to the increasing order of scheduled arrival times for first flights.

FOR $\psi = 1$ TO Ψ DO:

FOR $\xi = 1$ TO $\Xi(\psi)$ DO:

Put $\phi = \phi_{\psi\xi}$.

IF $\xi = 1$ THEN:

Define $\bar{\mathcal{T}}_\phi := \mathcal{T}_\phi$.

IF ϕ has a previous flight ϕ' THEN:

Remove from $\bar{\mathcal{T}}_\phi$ those t that are smaller than $r_\phi + g_{\phi'} - s_{\phi'}$ (because, if ϕ were assigned to such a t , then the coupling constraint linking ϕ and ϕ' would be violated).

END IF

END IF

Define the Capacity slacks $S_\phi(t) := R_{k_\phi^a(t)} - \sum_{f: k_f^a = k_\phi^a} v_{ft}$, $t \in \bar{\mathcal{T}}_\phi$.

Define $\hat{\mathcal{T}}_\phi := \{t \in \bar{\mathcal{T}}_\phi : S_\phi(t) \geq 1 - v_{\phi t}\}$.

IF $\hat{\mathcal{T}}_\phi = \emptyset$ THEN

Cancel ϕ : Put $z_\phi = 1$, $v_{\phi t} = 0$, $t \in \mathcal{T}_\phi$.

CONTINUE ξ

END IF

Assign current flight to τ , the smallest element of $\hat{\mathcal{T}}_\phi$: put $z_\phi =$

0 , $v_{\phi\tau} = 1$, $v_{\phi t} = 0$, $t \in \mathcal{T}_\phi \setminus \{\tau\}$.

IF ϕ has a next flight $\hat{\phi}$ THEN:

IF $\tau - r_\phi - s_\phi > g_{\hat{\phi}}$ AND $\hat{\phi} \notin \Phi$ THEN

Include $\hat{\phi}$ in Φ_ψ as $\phi_{\psi, \xi+1}$ and modify subsequent indices ξ accordingly.

END IF

Define $\bar{\mathcal{T}}_{\hat{\phi}} = \{t \in \mathcal{T}_{\hat{\phi}} : t - r_{\hat{\phi}} \geq \tau - r_\phi - s_\phi\}$.

END IF

CONTINUE ξ

CONTINUE ψ

END

Acknowledgments

We would like to thank the Charles S. Draper Laboratory for supporting work on this project. The research of D. Bertsimas was also partially supported by the National Science Foundation with a Presidential Young Investigator award DDM-9158118. We would also like to thank the undergraduate students Mercury Schroepel and Beryl Castello for performing some of the computational work reported.

References

- [1] Mostafa Terrab, "Ground Holding Strategies For Air Traffic Control", Ph.D. thesis, Massachusetts Institute of Technology, February 1990.

- [2] Amedeo R. Odoni, "The Flow Management Problem in Air Traffic Control", in A.R.Odoni, L.Bianco, G.Szego (eds), *Flow Control of Congested Networks*, Springer-Verlag, Berlin, 1987: pp.269-288.
- [3] Andreatta, G. & G. Romanin-Jacur, "Aircraft Flow Management under Congestion", *Transportation Science*, v.21 (1987): pp.249-253.
- [4] Octavio Richetta, "Ground Holding Strategies for Air Traffic Control under Uncertainty", Ph.D. thesis, Massachusetts Institute of Technology, June 1991.