

Dynamic Ground-Holding Policies for a
Network of Airports

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

WP# 3436-92-MSA June, 1992

Dynamic Ground-Holding Policies for a Network of Airports

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

Massachusetts Institute of Technology

June 8, 1992

Abstract

The yearly congestion costs in the US airline industry are estimated to be of the order of \$2 billion. In [6] we have introduced and studied generic integer programming models for the *static* multi-airport ground holding problem (GHP), the problem of assigning optimal ground holding delays in a general network of airports, so that the total (ground plus airborne) delay cost of all flights is minimized. The present paper is the first attempt to address the multi-airport GHP *in a dynamic environment*. We propose optimal or near-optimal algorithms to update ground-holding decisions as time progresses and more accurate weather (hence capacity) forecasts become available. We propose several pure IP formulations (most of them 0-1), which have the important advantages of being remarkably compact while capturing the essential aspects of the problem and of being sufficiently flexible to accommodate various degrees of modeling detail. For example, one formulation allows the dynamic updating of the mix between departure and arrival capacities by modifying runway use. These formulations enable one to assign and dynamically update ground holds to a sizable portion of the network of the major congested U.S. or European airports. We also present structural insights on the behaviour of the problem by means of computational results, and we find that our methods perform much better than a heuristic which may approximate, to some

extent, current ground-holding practices.

Introduction

In both the United States and Europe, demand for airport use has been increasing quite rapidly during recent years, while airport capacity has been more or less stagnating. Acute congestion in many major airports has been the result. Congestion creates ground and airborne delays at departure and arrival queues. Ground delays entail crew, maintenance, and depreciation costs, while airborne delays entail, in addition, fuel and safety costs.

For U.S. airlines, the total yearly delay costs due to congestion are estimated to be of the order of \$2 billion or more. In order to put this number in perspective, it must be considered that the total *losses* of all U.S. airlines amounted to about \$2 billion in 1991 and \$2.5 billion in 1990. European airlines are in a similar plight [3]. Thus, congestion is a problem of undeniable practical significance.

Not only is the congestion problem severe, it is also expected to get worse. The Federal Aviation Administration (FAA) predicts a 25% increase in demand for airport operations by the year 2000, while no appreciable increase in capacity is expected to be realized. Although capacity could be increased by constructing new airports or new runways at existing airports, significant community opposition to noise makes such new construction unlikely. Other possible solutions to the congestion problem, such as the improvement of Air Traffic Control (ATC) Technologies, the modification of the temporal pattern of aircraft flow in order to eliminate periods of "peak" demand (e.g., by means of congestion pricing), and the use of larger aircraft, are also unlikely to be implemented in the near future [5].

What can be done then?

Ground-holding policies appear to be a promising solution for reducing delay costs.

These short-term policies consider airport capacities and flight schedules as fixed for any given day, and adjust the flow of aircraft on a real-time basis by imposing “ground holds”: delaying the departure of an aircraft by not allowing it to start its engines and leave its gate or parking area even if it is ready to depart [1]. Ground holding makes sense in the following situation. If an aircraft departs on time, it will encounter congestion and will incur an airborne delay upon arrival at its destination; but if its departure is delayed, the aircraft will arrive at its destination at a later time when no congestion will be present and no airborne delay will be incurred. Therefore, the objective of ground-holding policies is to absorb airborne delays on the ground.

The effectiveness of ground-holding policies lies in the following two fundamental facts. First, airborne delays are much costlier than ground delays, because airborne delays entail fuel and safety costs. Second, airport capacity is highly variable, because it depends heavily on weather (visibility, wind, precipitation, cloud ceiling). It is not unusual for the capacity of an airport to be reduced by 50% in inclement weather. Given these two facts, there is significant potential for readjusting aircraft flow when weather (hence airport capacity) forecasts change, and such readjustment can result in a significant cost reduction if ground delays are substituted for the much costlier airborne delays.

The importance of ground-holding policies has long been recognized. The FAA has been operating for several years in Washington, D.C. 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.

Such optimization models are developed in the present paper, which deals with the

dynamic multi-airport Ground-Holding Problem (GHP). The GHP is the problem of determining, for a given network of airports, how much each aircraft must be held on the ground before taking off in order to minimize the total (ground plus airborne) delay cost for all flights. In *dynamic* versions of the problem, ground holds are updated during the course of the day as better weather (and hence capacity) forecasts become available. In *static* versions, on the contrary, ground holds are decided once for all at the beginning of the day. *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 typically performs more than one flight on any given day, “network” (or “down-the-road”) effects may be important: when an aircraft is delayed, the next flight scheduled to be performed by the same aircraft may also have to 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. However, most previous research on the GHP has neglected network effects, and has been restricted to the single-airport problem. Odoni [1] was the first to give a systematic description of the problem. Andreatta and Romanin-Jacur [2] proposed a dynamic programming algorithm for the single-airport static probabilistic GHP with one time period. Terrab [3] 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 a two-airport formulation and a closed-network three-airport formulation for the static deterministic GHP. Richetta [4] dealt with the single-airport dynamic probabilistic GHP. The first attempt to examine network effects was by the authors of the present paper (Vranas, Bertsimas, and Odoni [6]), who dealt with the *static* deterministic multi-airport GHP, proposing several pure 0–1 integer programming (IP) formulations of considerable

generality and compactness. These formulations make possible to solve realistic size problems involving as many as 6 congested airports and 3000 flights.

The present paper is the first attempt to address the multi-airport GHP in a dynamic environment. We propose optimal or near-optimal algorithms to update ground-holding decisions as time progresses and more accurate weather (hence capacity) forecasts become available. We propose several pure IP formulations (most of them 0–1), which have the important advantages of being remarkably compact while capturing the essential aspects of the problem and of being sufficiently flexible to accommodate various degrees of modeling detail. For example, one formulation allows the dynamic updating of the mix between departure and arrival capacities by modifying runway use. These formulations enable one to assign and dynamically update ground holds to a sizable portion of the network of the major congested U.S. or European airports. We also present structural insights on the behaviour of the problem by means of computational results, and we find that our formulations perform much better than a heuristic which may approximate, to some extent, current ground-holding practice.

The outline of this paper is as follows. Section 1 presents two pure 0–1 IP formulations for the dynamic *deterministic* multi-airport GHP. Section 2 gives several interesting extensions of the formulations of Section 1. Section 3 presents formulations for the *probabilistic* multi-airport GHP. Section 4 presents structural insights on the behaviour of the problem on the basis of computational results. Section 4 also proposes and evaluates a heuristic inspired by current ground-holding practice. Finally, Section 5 summarizes the results of the paper and points out directions for future research.

1 Deterministic 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 also 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} or arriving to an airport in \mathcal{K} (or both).

For each flight $f \in \mathcal{F}$, the following data are given: $k_f^d \in \mathcal{K}$, the airport from which f is scheduled to depart; $k_f^a \in \mathcal{K}$, 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 (so the scheduled travel time is $t_f = r_f - d_f$); $c_f^g(\cdot)$, 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).

Decisions are taken at (i.e., just before) various time periods $\tau \in \mathcal{T}' \subset \mathcal{T}$. \mathcal{T}' is the set of decision time periods. (In the static problem, $\mathcal{T}' = \{1\}$, since decisions are taken once for all at period 1.) At each decision time period τ , denote by $\mathcal{F}_\tau \subset \mathcal{F}$ the set of flights for which decisions can be taken; i.e., the set of flights not having yet landed at τ . Partition \mathcal{F}_τ into \mathcal{F}_τ^g , the set of flights not having yet taken off at τ , and \mathcal{F}_τ^a , the set of flights in the air at τ .

The reason why it makes sense to take new decisions at τ is that one has new departure and arrival capacity forecasts $D_k^l(t)$, $R_k^l(t)$, for $t \in \mathcal{T}_\tau := \{\tau, \dots, T\}$. Note that these capacity




Figure 1: Ground and airborne delays as decided at τ .

forecasts extend over the remainder of the time horizon, until time period T , rather than just until the next decision time period $\hat{\tau}$. It will be assumed that these capacity forecasts are perfectly accurate between decision time periods. In other words, the actual airport capacities $D_k(t)$, $R_k(t)$ will be equal to $D_k^\tau(t)$, $R_k^\tau(t)$ for $t \in \{\tau, \dots, \hat{\tau} - 1\}$, $\tau \in T'$.

We introduce now the following decision variables: $g_f^\tau, f \in \mathcal{F}_\tau^g$, is the number of time periods for which one decides at τ to hold flight f on the ground, before allowing f to take off; $a_f^\tau, f \in \mathcal{F}_\tau$, is the number of time periods for which one decides at τ to delay flight f in the air (e.g., by means of an *en route* speed reduction), before allowing f to land. At the next decision time period $\hat{\tau}$, one may change those decisions that will not have been implemented yet. Nevertheless, at τ the situation concerning a flight $f \in \mathcal{F}_\tau^g$ is as depicted in Figure 1.

The actual ground delay of flight $f \in \mathcal{F}$ will be denoted by \hat{G}_f . This is a number rather than a decision variable. Its value becomes determined when flight f departs and remains fixed afterwards. Similar remarks hold for \hat{A}_f , the actual airborne delay of flight f .

Table 1 summarizes the above notation for reference purposes. Table 1 also includes some symbols which will be defined later on.

Symbol	Denotation
$\mathcal{K} = \{1, \dots, K\}$	Set of airports k .
$\mathcal{T} = \{1, \dots, T\}$	Set of time periods t .
$\mathcal{T}_\tau = \{\tau, \dots, T\}$	Set of remaining time periods t .
$\mathcal{T}' \subset \mathcal{T}$	Set of decision time periods τ .
τ'	Decision time period previous to τ .
$\hat{\tau}$	Decision time period next to τ .
$\mathcal{F} = \{1, \dots, F\}$	Set of flights f .
$\mathcal{F}_\tau \subset \mathcal{F}$	Set of flights not having yet landed at τ .
\mathcal{F}_τ^g	Set of flights on the ground at τ .
\mathcal{F}_τ^a	Set of flights in the air at τ .
$k_f^d \in \mathcal{K}$	Departure airport of flight f .
$k_f^a \in \mathcal{K}$	Arrival airport of flight f .
$d_f \in \mathcal{T}$	Scheduled departure time of flight f .
$r_f \in \mathcal{T}$	Scheduled arrival time of flight f .
$c_f^g(\cdot)$	Ground delay cost function of flight f .
$c_f^a(\cdot)$	Airborne delay cost function of flight f .
$D_k^\tau(t)$	Departure capacity forecast at τ for airport k at period t .
$R_k^\tau(t)$	Arrival capacity forecast at τ for airport k at period t .
g_f^τ	Ground delay decision variable at τ for flight f .
a_f^τ	Airborne delay decision variable at τ for flight f .
u_{ft}^τ	Departure assignment decision variable at τ for flight f .
v_{ft}^τ	Arrival assignment decision variable at τ for flight f .
G_f	Actual ground delay of flight f (not a variable).
A_f	Actual airborne delay of flight f (not a variable).
\bar{G}_f	Upper bound on the ground delay of flight f .
\bar{A}_f	Upper bound on the airborne delay of flight f .
\mathcal{T}_f^d	Set of allowable departure periods for flight f .
\mathcal{T}_f^a	Set of allowable arrival periods for flight f .
$\mathcal{F}' \subset \mathcal{F}$	Set of continued flights f' .
$s_{f'}$	Slack time of continued flight f' .

Table 1: Basic notation.




Figure 2: Modeling of coupling constraints.

1.2 Coupling constraints.

Consider 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 on in the day. For each flight $f' \in \mathcal{F}'$, we are given the next flight f scheduled to be performed by the same aircraft.

Network effects will be taken into account in the following way. For each continued flight $f' \in \mathcal{F}'$, we are given the “slack” or “absorption” time $s_{f'}$. The slack is defined as the number of time periods such that, if f' arrives at its destination at most $s_{f'}$ time periods late, the departure of the next flight f is not affected, whereas if f' lands with a delay greater than the slack, the “excess delay” of f' (i.e., the delay minus the slack) is transferred to the next flight f . In the latter case, the next flight f will incur a ground delay at least equal to the excess delay of f' . The situation is depicted in Figure 2, where it can be seen that the slack $s_{f'}$ is 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.

1.3 Assignment decision variables.

The *delay* decision variables g_f^τ and a_f^τ were introduced above. Now we introduce the following *assignment* decision variables: u_{ft}^τ , $f \in \mathcal{F}_\tau^g$, is 1 if one decides at τ that flight f will take off at period t (i.e., if $r_f + g_f^\tau = t$) and 0 otherwise; v_{ft}^τ , $f \in \mathcal{F}_\tau$, is 1 if one decides at τ that flight f will land at period t (i.e., if $r_f + g_f^\tau + a_f^\tau = t$) and 0 otherwise. These new decision variables are introduced because the capacity constraints cannot be expressed in a simple linear way in terms of the more natural delay decision variables.

Moreover, since we do not want to have excessive ground or airborne delays, we also introduce upper bounds on those delays. G_f is the maximum number of time periods that flight f may be held on the ground, and A_f is the maximum number of time periods that flight f may be held in the air. Introduction of these bounds results in no loss of generality, since they can be arbitrarily large. In practice, however, typical values are $G_f = 4-5$ and $A_f = 2-3$, corresponding to maximum ground and airborne delays of about one hour and half an hour, respectively.

Given the above setup, the set \mathcal{T}_f^d of time periods to which flight f may be assigned to take off is given by:

$$\mathcal{T}_f^d = \{t \in \mathcal{T} : d_f \leq t \leq \min(d_f + G_f, T)\}. \quad (1)$$

Similarly, the set \mathcal{T}_f^a of time periods to which flight f may be assigned to land is given by:

$$\mathcal{T}_f^a = \{t \in \mathcal{T} : r_f \leq t \leq \min(r_f + G_f + A_f, T)\}. \quad (2)$$

We are now ready to give a first pure 0–1 integer programming formulation of the static deterministic multi-airport GHP.

1.4 A first pure 0–1 IP formulation.

The dynamic formulation consists in solving, at each decision time period τ , the following pure 0–1 integer program:

$$(I_1^\tau) \min \sum_{f \in \mathcal{F}_\tau^g} c_f^g g_f^\tau + \sum_{f \in \mathcal{F}_\tau} c_f^a a_f^\tau$$

$$\text{s.t.} \quad \sum_{f \in \mathcal{F}_\tau^g: k_f^g = k} u_{ft}^\tau \leq D_k^\tau(t), \quad (k, t) \in \mathcal{K} \times \mathcal{T}_\tau; \quad (3)$$

$$\sum_{f \in \mathcal{F}_\tau: k_f^a = k} v_{ft}^\tau \leq R_k^\tau(t), \quad (k, t) \in \mathcal{K} \times \mathcal{T}_\tau; \quad (4)$$

$$\sum_{t \in \mathcal{T}_f^g \cap \mathcal{T}_\tau} u_{ft}^\tau = 1, \quad f \in \mathcal{F}_\tau^g; \quad (5)$$

$$\sum_{t \in \mathcal{T}_f^a \cap \mathcal{T}_\tau} v_{ft}^\tau = 1, \quad f \in \mathcal{F}_\tau; \quad (6)$$

$$g_{f'}^\tau + a_{f'}^\tau - s_{f'} \leq g_f^\tau, \quad f' \in \mathcal{F}' \cap \mathcal{F}_\tau^g; \quad (7)$$

$$\hat{G}_{f'} + a_{f'}^\tau - s_{f'} \leq g_f^\tau, \quad f' \in \mathcal{F}' \cap \mathcal{F}_\tau^a; \quad (8)$$

$$a_f^\tau \geq 0, \quad f \in \mathcal{F}_\tau; \quad (9)$$

$$u_{ft}^\tau, v_{ft}^\tau \in \{0, 1\}. \quad (10)$$

The formulation presents a dichotomy necessitated by the fact that flights on the ground at τ and flights in the air at τ must be treated differently. The objective function is a sum of two terms, corresponding to the ground delay costs of flights on the ground and to airborne delay costs of all flights in \mathcal{F}_τ . (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).) The departure capacity constraints (3) and the departure assignment constraints (5) refer only to flights on the ground, while the arrival capacity constraints (4) and the arrival assignment constraints (6) refer to all flights in \mathcal{F}_τ . The coupling constraints (7) and (8) (cf. Figure 2) are also divided into two categories because, for the continued flights f' which are already in the air at τ , the ground hold $\hat{G}_{f'}$ has already been determined: it is a number, not a decision variable (constraints (8)).

For simplicity of exposition, variables g_f^τ and a_f^τ were kept in formulation (I_1^τ) , but it should be clear that they can be eliminated by mere substitution in the following way:

$$g_f^\tau = \sum_{t \in \mathcal{T}_f^g \cap \mathcal{T}_\tau} t u_{ft}^\tau - d_f, \quad f \in \mathcal{F}_\tau^g; \quad (11)$$

$$a_f^\tau = \sum_{t \in \mathcal{T}_f^a \cap \mathcal{T}_\tau} t v_{ft}^\tau - r_f - g_f^\tau, \quad f \in \mathcal{F}_\tau^g; \quad (12)$$

$$a_f^\tau = \sum_{t \in \mathcal{T}_f^a \cap \mathcal{T}_\tau} t v_{ft}^\tau - r_f - \hat{G}_f, \quad f \in \mathcal{F}_\tau^a. \quad (13)$$

Therefore, the only decision variables are u_{ft}^τ and v_{ft}^τ . The number of these variables for a given τ is at most $\sum_{f \in \mathcal{F}} (2G_f + A_f + 2)$ (there are at most $G_f + 1$ u_{ft}^τ variables and $G_f + A_f + 1$ v_{ft}^τ variables for a given f). For the typical values $G_f = 4$ and $A_f = 2$, the number of variables becomes at most $12F$, which is a small linear multiple of the number of flights. It can be seen that the number of constraints is at most $3F + F' + 2KT$. Thus formulation (I_1^τ) is quite compact. Note also that the size of the integer program decreases as τ increases.

Having solved the program (I_1^τ) , one solves the program $(I_1^{\hat{\tau}})$ corresponding to the next decision period $\hat{\tau}$ by taking as inputs $D_k^{\hat{\tau}}(t)$, $R_k^{\hat{\tau}}(t)$, $\mathcal{T}_{\hat{\tau}}$, and by updating the flight sets as follows:

$$\mathcal{F}_{\hat{\tau}}^g = \mathcal{F}_\tau^g \setminus \{f \in \mathcal{F}_\tau^g : d_f + \hat{g}_f^\tau < \hat{\tau}\}; \quad (14)$$

$$\mathcal{F}_{\hat{\tau}}^a = (\mathcal{F}_\tau^a \setminus \{f \in \mathcal{F}_\tau^a : r_f + \hat{G}_f + \hat{a}_f^\tau < \hat{\tau}\}) \cup \{f \in \mathcal{F}_\tau^g : (d_f + \hat{g}_f^\tau < \hat{\tau})(r_f + \hat{g}_f^\tau + \hat{a}_f^\tau \geq \hat{\tau})\}, \quad (15)$$

where \hat{g}_f^τ and \hat{a}_f^τ are the optimal values returned by (I_1^τ) . In words, (14) simply says that the new set of flights on the ground is equal to the previous set of flights on the ground minus the flights that were assigned by (I_1^τ) to leave before the new decision period $\hat{\tau}$. Similarly, (15) says that the new set of flights in the air is equal to the old set of flights in the air minus any flights in that set that were assigned to land before the new decision period, plus any flights that were previously on the ground, were assigned to depart before

the new decision period, and were assigned to land at or after it.

The final cost resulting from the above dynamic formulation is $\sum_{f \in \mathcal{F}} (c_f^g \hat{G}_f + c_f^a \hat{A}_f)$. At $\hat{\tau}$ one updates the cost by adding to it the sum of $c_f^g \hat{g}_f^\tau$ for $f \in \{f \in \mathcal{F}_\tau^g : d_f + \hat{g}_f^\tau < \hat{\tau}\}$ (i.e., flights which were on the ground at τ but left before $\hat{\tau}$), and the sum of $c_f^a \hat{a}_f^\tau$ for $f \in \{f \in \mathcal{F}_\tau^a : r_f + \hat{G}_f + \hat{a}_f^\tau < \hat{\tau}\} \cup \{f \in \mathcal{F}_\tau^g : r_f + \hat{g}_f^\tau + \hat{a}_f^\tau < \hat{\tau}\}$ (i.e., flights that either were in the air or were on the ground at τ and landed before $\hat{\tau}$). One also sets $\hat{G}_f = \hat{g}_f^\tau$ and $\hat{A}_f = \hat{a}_f^\tau$ for flights in the above sets.

An interesting feature of formulation (I_1^T) is that decisions can be updated or overruled as long as they have not been implemented. For instance, aircraft can be assigned ground or airborne delays smaller or larger than what they had been previously assigned, and aircraft not having yet taken off can be given priority over aircraft already in the air. The last possibility is not expected to appear often in the optimal solution, given that c_f^a is higher than c_f^g . Nevertheless, in practice one would almost never want to deal with this possibility, so formulation (I_1^T) may be too general. The second dynamic formulation, presented below, always gives aircraft in the air priority over aircraft not having yet taken off.

1.5 A second pure 0–1 IP formulation.

The second formulation assumes infinite departure capacities. For the static case, it has been shown [6,7] that this assumption results in no significant loss of generality if the real departure capacities are somewhat higher than arrival capacities, as is the case in practice. It has also been shown that, when departure capacities are infinite, if formulation (I_1^1) has an optimal solution, then it has an optimal solution in which no flight incurs an airborne delay. In the dynamic case, however, airborne delays cannot be completely eliminated even if departure capacities are infinite. This is because, at a given decision time, the new arrival

capacity forecasts may be significantly reduced with respect to the forecasts of the previous decision time. Then it can happen that even the number of aircraft already in the air exceeds the new capacity forecasts, so that some of these aircraft may have to wait in the air when they arrive at their destination.

For a given decision time period τ , define the *excess* at airport k and time period t , denoted by E_{kt}^τ , as the number of aircraft currently in the air which will arrive at k at t minus the currently forecasted capacity of k at t . At each decision time period τ , one needs to do the following preliminary calculations in order to find what the excesses E_{kt}^τ are:

BEGIN

FOR $k = 1$ TO K DO:

$$E_{k,\tau-1}^\tau = E_{k,\tau-1}^{\tau'}$$

FOR $t = \tau$ TO T DO:

$$E_{kt}^\tau = \max(E_{k,t-1}^\tau, 0) + |\{f \in \mathcal{F}_\tau^a : (k_f^a = k)(r_f + \hat{g}_f^{\tau'} = t)\}| - R_k^\tau(t).$$

CONTINUE t

CONTINUE k

END

$E_{k,\tau-1}^{\tau'}$ are the excesses calculated at τ' , the decision time previous to τ . (At decision time 1, one begins with $E_{k0}^1 = 0$.) These previously calculated excesses are the actually realized excesses, since capacity forecasts at τ' are accurate until $\tau - 1$. Any positive excess at a time period $t - 1$ is transferred to the next time period t . As a result of the above preliminary calculations, if $E_{kt}^\tau < 0$, then $-E_{kt}^\tau$ is the *available arrival capacity* of airport k at time period t , i.e., the currently forecasted capacity minus the number of aircraft that

are already in the air and will arrive at k at t . On the other hand, if $E_{kt}^T \geq 0$, then there is no available capacity at k at t , and the best one can do, supposing that aircraft in the air have priority over aircraft on the ground, is to not assign any of the aircraft currently on the ground to arrive at k at t .

Another preliminary calculation which needs to be carried out at each decision time period concerns the unavoidable airborne delays of flights already in the air, arising when the new capacity forecasts are not sufficient to accommodate these flights. These delays, which will be denoted by α_f^T (numbers, not variables), will be needed for the coupling constraints and for the calculation of the total cost. They are calculated by means of the excesses E_{kt}^T in the following way:

BEGIN

FOR $f \in \mathcal{F}_\tau^a$ DO:

$$\alpha_f^T = \alpha_f^{\tau'} .^1$$

CONTINUE f

FOR $k = 1$ TO K DO:

FOR $t = \tau$ TO T DO:

IF $E_{kt}^T > 0$ THEN

Select E_{kt}^T flights in $\{f \in \mathcal{F}_\tau^a : (k_f^a = k)(r_f + \hat{g}_f^{\tau'} + \alpha_f^T = t)\}$ and

set $\alpha_f^T = \alpha_f^T + 1$ for them.

END IF

CONTINUE t

¹At $\tau = 1$, one begins with $\alpha_f^1 = 0$.

CONTINUE k

END

The selection of the flights to be delayed depends on the arrival queueing discipline, usually FCFS.

We are now in the position to give the second dynamic deterministic formulation, with infinite departure capacities:

$$(I_2^\tau) \min \quad \sum_{f \in \mathcal{F}_\tau^g} c_f^g g_f^\tau$$

$$\text{s.t.} \quad \sum_{f \in \mathcal{F}_\tau^g: k_f^a = k} v_{ft}^\tau \leq \max(-E_{kt}^\tau, 0), \quad (k, t) \in \mathcal{K} \times T_\tau; \quad (16)$$

$$\sum_{t \in T_f^a \cap T_\tau} v_{ft}^\tau = 1, \quad f \in \mathcal{F}_\tau^g; \quad (17)$$

$$g_{f'}^\tau - s_{f'} \leq g_f^\tau, \quad f' \in \mathcal{F}' \cap \mathcal{F}_\tau^g; \quad (18)$$

$$\hat{G}_{f'} + \alpha_{f'}^\tau - s_{f'} \leq g_f^\tau, \quad f' \in \mathcal{F}' \cap \mathcal{F}_\tau^a; \quad (19)$$

$$v_{ft}^\tau \in \{0, 1\}, \quad (20)$$

where:

$$g_f^\tau = \sum_{t \in T_f^a \cap T_\tau} t v_{ft}^\tau - r_f, \quad f \in \mathcal{F}_\tau^g. \quad (21)$$

Constraints (16) say that, if $E_{kt}^\tau < 0$, then the number of aircraft assigned to arrive at airport k at time period t must not exceed the available excess capacity $-E_{kt}^\tau$; whereas, if $E_{kt}^\tau \geq 0$ (i.e., no available capacity at k at t), then v_{ft}^τ will be 0 for all f that could arrive at k at t , so that no new aircraft will be assigned to arrive at k at t . Note that decisions are taken only for flights on the ground at τ . Flights in the air at τ influence the decisions by means of (16) (by determining the excesses) and of (19) (their airborne delays enter into the coupling constraints).

The total cost is calculated and updated at each decision time period in a way analogous to that explained in the previous subsection for formulation (I_1^τ) .

2 Extensions of the deterministic formulations.

2.1 Flight cancellations.

In situations where delays become excessive, it is common airline practice to cancel some flights, especially at hub airports. Motivated from this fact, we developed formulations which take into account the possibility of cancelling flights. These formulations have the additional advantage that they escape infeasibility problems which might arise with formulations (I_1^T) and (I_2^T) . Infeasibility occurs when airport capacities are low: even though the total daily capacity of an airport may be sufficient to accommodate the total number of flights scheduled to depart from or arrive at that airport, the problem may still be infeasible if excessive congestion appears during some portion of the day.

We will give a formulation (I_3^T) generalizing (I_2^T) . Keep the old decision variables u_{ft}^T and v_{ft}^T , and define the new decision variables $z_f^T, f \in \mathcal{F}_\tau^j$, to be 1 if flight f is cancelled at τ 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, especially in hub airports where most cancellations take place, but the 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. At τ solve:

$$\begin{aligned}
 (I_3^T) \quad \min \quad & \sum_{f \in \mathcal{F}_\tau^g} (c_f^g g_f^T + (M_f + c_f^g r_f) z_f^T) \\
 \text{s.t.} \quad & \sum_{f \in \mathcal{F}_\tau^g: k_f^g = k} v_{ft}^T \leq \max(-E_{kt}^T, 0), \quad (k, t) \in \mathcal{K} \times \mathcal{T}_\tau;
 \end{aligned}$$

$$\begin{aligned}
z_f^\tau + \sum_{t \in \mathcal{T}_f^g \cap \mathcal{T}_\tau} v_{ft}^\tau &= 1, & f \in \mathcal{F}_\tau^g; & \quad (22) \\
g_{f'}^\tau - s_{f'} + (s_{f'} + r_{f'} - r_f)z_{f'}^\tau &\leq g_f^\tau, & f' \in \mathcal{F}'_1 \cap \mathcal{F}_\tau^g; \\
g_{f'}^\tau - s_{f'} + (s_{f'} + r_{f'} + G_f + 1)z_{f'}^\tau &\leq g_f^\tau + (r_f + G_f + 1)z_f^\tau, & f' \in \mathcal{F}'_2 \cap \mathcal{F}_\tau^g; \\
\hat{G}_{f'} + \alpha_{f'}^\tau - s_{f'} &\leq g_f^\tau, & f' \in \mathcal{F}' \cap \mathcal{F}_\tau^a; \\
v_{ft}^\tau, z_f^\tau &\in \{0, 1\}.
\end{aligned}$$

where g_f^τ are given by (21).

The above formulation incorporates some technical tricks which are necessitated by the fact that, when a flight f is cancelled (i.e., $z_f^\tau = 1$), then all v_{ft}^τ corresponding to f are 0 (by (22)), so that (21) gives $g_f^\tau = -r_f$. Keeping this fact in mind, it can be seen immediately that, when $z_f^\tau = 1$, the objective function term corresponding to f is M_f . One can similarly simplify the coupling constraints when $z_{f'}^\tau = 1$.

The variables g_f^τ were again left in the formulation, but it should be clear that they can be eliminated by mere substitution through (21). Now an important point is that the variables z_f^τ can also be eliminated through (22), provided that (22) is replaced by $\sum_{t \in \mathcal{T}_f^g \cap \mathcal{T}_\tau} v_{ft}^\tau \leq 1$.

The fact that the new formulation (I_3^τ) has exactly the same number of variables and of constraints as (I_2^τ) is particularly significant, since (I_3^τ) enjoys considerable advantages both in terms of generality (the real-world problem is better approximated) and in terms of flexibility (infeasibility problems are eliminated).

2.2 Interdependent departure and arrival capacities.

The departure and arrival capacities of a given airport at a given time period are often not independent, because they are determined by the way in which runway use is allocated among departing or arriving aircraft. If all runways are exclusively used for landings,




Figure 3: An example of a region of possible combinations between the departure capacity d_{kt}^τ and the arrival capacity r_{kt}^τ of airport k at time period t , as forecasted at τ .

arrival capacity reaches a maximum value $R_{k,max}(t)$ determined by the minimum separation between successive landings, while departure capacity is 0. If all runways are exclusively used for take-offs, departure capacity reaches a maximum value $D_{k,max}(t)$ determined by the minimum separation between successive take-offs (which is less than the minimum separation between landings, so that $D_{k,max}(t) > R_{k,max}(t)$), while arrival capacity is 0. Intermediate cases give departure and arrival capacities belonging to a region with the general shape of a two-dimensional convex polytope (Figure 3). Note that this region differs among airports and, for a given airport, it can change with time (because weather can change).

The above situation can be easily taken into account by generalizing formulation (I_1^τ) . This is achieved by introducing the new integer (not 0–1) decision variables d_{kt}^τ and r_{kt}^τ standing, respectively, for the departure and the arrival capacities of airport k at time t , as forecasted at τ . These new decision variables will replace the constants $D_k^\tau(t)$ and $R_k^\tau(t)$

in the right-hand sides of the capacity constraints. In addition, new sets of constraints will be needed, one set for each time period, ensuring that d_{kt}^τ and r_{kt}^τ fall within the region of their possible combinations. These constraints will be of the following general form:

$$\alpha_{kt}^{i\tau} d_{kt} + \beta_{kt}^{i\tau} r_{kt} \leq \gamma_{kt}^{i\tau}, \quad (k, t) \in \mathcal{K} \times \mathcal{T}_\tau, i \in \{1, \dots, I_{kt}^\tau\}, \quad (23)$$

where $\alpha_{kt}^{i\tau}, \beta_{kt}^{i\tau}, \gamma_{kt}^{i\tau}$ are constants and I_{kt}^τ is the number of linear constraints describing the departure-arrival capacity region of airport k at time period t , as forecasted at τ . For instance, for the region shown in Figure 3, the following two constraints are needed for period t :

$$3r_{kt}^\tau + d_{kt}^\tau \leq 3R_{k,max}(t); \quad (24)$$

$$r_{kt}^\tau + 3d_{kt}^\tau \leq 3D_{k,max}(t). \quad (25)$$

Variables d_{kt}^τ and r_{kt}^τ can be eliminated. In fact, the departure and arrival capacity constraints (3) and (4) can be replaced by equalities since constraints (23) are inequalities. Then d_{kt}^τ and r_{kt}^τ will represent the *used* capacities of airport k at time period t , as forecasted at τ . Thus constraints (3) and (4) will be removed from (I_1^τ) and will be replaced by:

$$\alpha_{kt}^{i\tau} \left(\sum_{f:k_d^f=k} u_{ft}^\tau \right) + \beta_{kt}^{i\tau} \left(\sum_{f:k_a^f=k} v_{ft}^\tau \right) \leq \gamma_{kt}^{i\tau}, \quad (k, t) \in \mathcal{K} \times \mathcal{T}_\tau, i \in \{1, \dots, I_{kt}^\tau\}$$

. It is remarkable that this generalization of (I_1^τ) has exactly the same (number of) variables as (I_1^τ) and only slightly more constraints.

2.3 Further extensions.

2.3.1 Hub airports: more than one “next” flights.

In hub airports, an arriving flight typically has passengers connecting to several departing flights. This can be taken into account in any of the above formulations by means of an

easy extension. It suffices to *reinterpret* the coupling constraints as linking not only a pair of flights scheduled to be performed by the same aircraft, but also any pair of flights f' and f such that f is not allowed to leave before f' lands, because passengers in f' connect to f . With this reinterpretation of the coupling constraints, a continued flight may have more than one next flights. Therefore, the formulations remain unchanged and a number of new coupling constraints is simply added to them.

Note that the slack in a coupling constraint of the new kind (connecting flights) will typically be different from the slack in a coupling constraint of the old kind (continued flights). This is because the turnaround time involved in connecting flights is restricted to moving passengers and their luggage, while the turnaround time involved in continued flights also includes cleaning and refuelling the aircraft.

2.3.2 En route speeding.

Sometimes there is the possibility to speed up an aircraft *en route*, so that the aircraft may arrive even before its scheduled arrival time. This possibility can be easily taken into account in any of the formulations with airborne delays presented in this chapter. It suffices to take r_f to be not the scheduled arrival time, but the earliest possible arrival time. If, for instance, an aircraft is scheduled to arrive at time period 28, but may be speeded up so as to arrive up to two periods earlier, r_f for the corresponding f will be equal to 26. An airborne “delay” \hat{A}_f equal, e.g., to 1 will correspond to a speeding up of one time period, whereas an \hat{A}_f equal to 3 will correspond to a slowing down of one time period. The actual arrival time of flight f will of course also depend on its ground delay: if f departs with a ground delay \hat{G}_f equal to 3 and is speeded up by one time period, it will arrive with a total delay of two time periods, i.e., at period 30 (which is $r_f + \hat{G}_f + \hat{A}_f = 26 + 3 + 1$). Note that

the upper bound A_f will have to be increased by 2 in this example.

3 Probabilistic formulations.

One way to model the case of probabilistic capacities is by considering that capacity forecasts take the forms of various scenarios, each scenario coming with a given probability of realization. In symbols, there are L possible capacity scenarios, and scenario l , having probability p_l ($\sum_{l=1}^L p_l = 1$), is denoted by $R_k^l(t)$, $(k, t) \in \mathcal{K} \times \mathcal{T}$. Note that a capacity scenario involves capacity forecasts for all airports of the network. In other words, p_l is the probability that airport 1 will have capacities $R_1^l(t)$, $t \in \mathcal{T}$, and that airport 2 will have capacities $R_2^l(t)$, $t \in \mathcal{T}$, and so on. This is because capacities at various airports may not be independent, especially for airports close enough to have similar weather.

In the probabilistic GHP, static policies are subject to a “paradox” which does not appear in the deterministic GHP. At the beginning of the day, one knows the possible capacity scenarios and their probabilities of realization, but of course one does not know which scenario will be realized. If the problem is static, one must make irrevocable decisions concerning ground holds at the beginning of the day. But sooner or later some scenario will be realized, and at that point one should normally take into account the new information and update ground holds. So the paradox of the static probabilistic GHP is that new information will inevitably become available but will not be taken into account. The static *deterministic* GHP encounters no similar problem because, by assumption, capacity forecasts are perfectly accurate and no new information will become available.

The above considerations show that static probabilistic formulations may be of somewhat limited practical interest in themselves. Nevertheless, they can be used as building blocks for dynamic probabilistic formulations. This is entirely analogous to the way in

which the static deterministic formulations of [6] were used as building blocks for the dynamic deterministic formulations of Section 1. As an example, we will present below, for the case of infinite departure capacities, a static probabilistic formulation and its dynamic extension.

3.1 A static probabilistic formulation.

Define the decision variables g_f , equal to the ground delay of flight f , and v_{ft}^l , equal to 1 if scenario l is realized and flight f lands at t , and equal to 0 otherwise. Denote also by a_f^l the airborne delay of flight f if scenario l is realized. Under scenario l , the total delay of flight f , $g_f + a_f^l$, is equal to the difference between the actual arrival time, $\sum_{t \in \mathcal{T}_f^a} tv_{ft}^l$, and the scheduled arrival time, r_f , so that:

$$a_f^l = \sum_{t \in \mathcal{T}_f^a} tv_{ft}^l - r_f - g_f, \quad f \in \mathcal{F}, l \in \{1, \dots, L\}. \quad (26)$$

Assuming infinite departure capacities, we have now the following static probabilistic IP formulation:

$$(IP) \quad \min \quad \sum_{f \in \mathcal{F}} (c_f^g g_f + c_f^a \sum_{l=1}^L p_l a_f^l)$$

$$\text{s.t.} \quad \sum_{f: k_f^a = k} v_{ft}^l \leq R_k^l(t), \quad (k, t) \in \mathcal{K} \times \mathcal{T}, l \in \{1, \dots, L\}; \quad (27)$$

$$\sum_{t \in \mathcal{T}_f^a} v_{ft}^l = 1, \quad f \in \mathcal{F}, l \in \{1, \dots, L\}; \quad (28)$$

$$g_{f'} + a_{f'}^l - s_{f'} \leq g_f, \quad f' \in \mathcal{F}', l \in \{1, \dots, L\}; \quad (29)$$

$$g_f \in \{0, 1, \dots, G_f\}, \quad f \in \mathcal{F}; \quad (30)$$

$$v_{ft}^l \in \{0, 1\}. \quad (31)$$

Although formulation (IP) looks superficially similar to previously presented formulations, it has several peculiarities which need mentioning. By solving (IP), one gets values for g_f and v_{ft}^l . The values for g_f are the ground holds which will be implemented no matter which

capacity scenario is ultimately realized, since we are examining the static case. On the contrary, which values v_{ft}^l will be implemented will depend on the capacity scenario that will be realized. If, for instance, scenario 3 is realized, then flight f will land at the time period t for which v_{ft}^3 is equal to 1. Therefore, it can be seen that g_f *cannot* be expressed in terms of v_{ft}^l : there are two independent sets of decision variables. Moreover, (IP) is not a *pure* 0–1 IP formulation, since g_f are not binary variables.

Another important comment concerns the coupling constraints (29). These must ensure that the ground holds g_f , which are irrevocably decided at the beginning of the day, will be feasible no matter which capacity scenario is ultimately realized. But the capacity scenario which will be realized may affect the airborne delays of continued flights, which may in turn affect the ground delays of their next flights. This problem is solved, in (29), by having the ground delay of a continuing flight be at least equal to the *maximum* excess delay of its previous flight over all possible capacity scenarios.²

A final remark concerns the size of formulation (IP) . There are $LKT + (L + 1)F + LF'$ constraints, and the number of decision variables is at most $F + \sum_{l=1}^L \sum_{f \in \mathcal{F}} (A_f^l + 1)$, where A_f^l is an upper bound on the airborne delay of flight f under scenario l . Such an upper bound cannot be arbitrarily imposed, but can be calculated for given arrival profiles and capacity scenario. In the worst case, $A_f^l = T - r_f$ will do. Therefore, the number of constraints and the number of variables can become excessive, especially if L is large, but may remain manageable for small L .

Extensions of formulation (IP) like those presented in Section 1 (e.g., with flight cancellations) are straightforward.

² $g_f \geq g_{f'} + a_{f'}^l - s_{f'}, l \in \{1, \dots, L\}$, is equivalent to: $g_f \geq \max(g_{f'} + a_{f'}^l - s_{f'}, l \in \{1, \dots, L\})$.

3.2 A dynamic probabilistic formulation.

Extending the static probabilistic formulation of the previous subsection to the dynamic probabilistic case is rather straightforward, but there is a minor complication. The complication concerns the way of modeling the additional information that emerges as time goes on. In an extreme case, one of the possible capacity scenaria is realized at a time δ and all uncertainty is eliminated after that time. In a more realistic case, at various *realization time periods* $\delta \in \Delta \subset T$ the probabilities of the various scenaria are updated to p_i^δ . In this case the reasonable thing to do is to identify the set of decision time periods T' with the set of realization time periods Δ , since it is exactly at realization time periods that new information becomes available.

The situation can be explained with the help of Figure 4, which refers to a case with $L = 3$ capacity scenaria and is also the basis of the computational results presented in Chapter 5.³ At time $\tau_0 = 1$, one knows that scenaria $S1, S2, S3$ will be ultimately realized with probabilities $p_1^{\tau_0} = p_1, p_2^{\tau_0} = (1 - p_1)p_2, p_3^{\tau_0} = (1 - p_1)(1 - p_2)$, respectively. (Moreover, one knows the capacities with certainty until time τ_1 , since all three scenaria coincide until that time.) So at τ_0 one solves formulation (IP) with $p_i^{\tau_0}$ as above. Now at time τ_1 new information is obtained: either scenario $S1$ is realized or it is not. This new information gives $p_i^{\tau_1}$. If $S1$ is realized, then $p_1^{\tau_1} = 1$ and $p_2^{\tau_1} = p_3^{\tau_1} = 0$. If $S1$ is not realized, then $p_1^{\tau_1} = 0, p_2^{\tau_1} = p_2$, and $p_3^{\tau_1} = 1 - p_2$. So at τ_1 one solves formulation (IP) with $p_i^{\tau_1}$ as above.⁴ Similarly at τ_2 .

We give now the dynamic probabilistic formulation corresponding to (IP) . The notation

³As was pointed out above, a capacity scenario includes capacity forecasts for all airports of the network. Figure 4 gives only the parts of scenaria $S1, S2$, and $S3$ that correspond to a given airport.

⁴Of course at τ_1 the new scenaria extend from τ_1 to T (rather than from τ_0 to T as they did at τ_0).




Figure 4: Modeling additional information over time in the dynamic probabilistic GHP. There are three possible capacity scenarios: $S1$, $S2$, and $S3$. Overall, $S1$ has lower capacities than $S2$, and $S2$ has lower capacities than $S3$. All three scenarios coincide between time periods τ_0 and τ_1 , and scenarios $S2$ and $S3$ coincide between time periods $\tau_1 + 1$ and τ_2 . At τ_1 , $S1$ is realized with probability p_1 and all uncertainty is eliminated. Otherwise (with probability $1 - p_1$), at τ_2 , $S2$ is realized with probability p_2 or $S3$ is realized with probability $1 - p_2$.

generalizes that of Section 1 (cf. Table 1).

$$(I_P^\tau) \min \sum_{f \in \mathcal{F}_\tau^g} c_f^g g_f^\tau + \sum_{f \in \mathcal{F}_\tau} c_f^a \sum_{l=1}^L p_l^\tau a_f^{l\tau} \quad (32)$$

$$\text{s.t.} \quad \sum_{f \in \mathcal{F}_\tau: k_f^a = k} v_{ft}^{l\tau} \leq R_k^{l\tau}(t), \quad (k, t) \in \mathcal{K} \times \mathcal{T}_\tau, l \in \{1, \dots, L\}; \quad (33)$$

$$\sum_{t \in \mathcal{T}_f^a \cap \mathcal{T}_\tau} v_{ft}^{l\tau} = 1, \quad f \in \mathcal{F}_\tau, l \in \{1, \dots, L\}; \quad (34)$$

$$g_{f'}^\tau + a_{f'}^{l\tau} - s_{f'} \leq g_f^\tau, \quad f' \in \mathcal{F}' \cap \mathcal{F}_\tau^g, l \in \{1, \dots, L\}; \quad (35)$$

$$\hat{G}_{f'} + a_{f'}^{l\tau} - s_{f'} \leq g_f^\tau, \quad f' \in \mathcal{F}' \cap \mathcal{F}_\tau^a, l \in \{1, \dots, L\}; \quad (36)$$

$$g_f^\tau \in \{0, 1, \dots, G_f\}, \quad f \in \mathcal{F}; \quad (37)$$

$$v_{ft}^{l\tau} \in \{0, 1\}, \quad (38)$$

where:

$$a_f^{l\tau} = \sum_{t \in \mathcal{T}_f^a \cap \mathcal{T}_\tau} t v_{ft}^{l\tau} - r_f - g_f^\tau, \quad f \in \mathcal{F}_\tau^g, l \in \{1, \dots, L\}; \quad (39)$$

$$a_f^{l\tau} = \sum_{t \in \mathcal{T}_f^a \cap \mathcal{T}_\tau} t v_{ft}^{l\tau} - r_f - \hat{G}_f, \quad f \in \mathcal{F}_\tau^a, l \in \{1, \dots, L\}. \quad (40)$$

This concludes the presentation of the formulations. The next section will present computational results based on the above framework.

4 Structural insights.

4.1 Comparing dynamic scenaria.

The goal of this section is to gain insight on the behaviour of the dynamic multi-airport GHP by examining a relatively realistic case with three capacity scenaria and two realization times (identical with the two decision times), on the model of Figure 4. The computations will be carried out with formulation (I_2^τ) . In the most general case (see Figure 4), one solves first $(I_2^{\tau_0})$ with capacity forecasts equal to $S1$ or $S2$ or $S3$, for $t \in \{1, \dots, T\}$. Then one

solves $(I_2^{\tau_1})$ with new capacity forecasts (e.g., $S2$ or $S3$ if the previous forecast was $S1$), but now for $t \in \{\tau_1, \dots, T\}$. Finally one solves $(I_2^{\tau_2})$ with yet new capacity forecasts and for $t \in \{\tau_2, \dots, T\}$. In special cases, one may not need to solve $(I_2^{\tau_1})$, $(I_2^{\tau_2})$, or both, if the capacity forecast does not change. Suppose, for instance, that the forecast at τ_0 is $S2$ and that $S1$ is realized at τ_1 . Then one needs to solve $(I_2^{\tau_1})$ with forecast $S1$, but no new problem needs to be solved at τ_2 , since the forecast will inevitably remain $S1$.

A particular combination of (at most three) problems to be solved in the dynamic case will be referred to as a *(dynamic) scenario* (not to be confused with a *capacity scenario*, which is one of $S1, S2, S3$). In the example at the end of the last paragraph, the dynamic scenario will be referred to as 2-1, since one solves $(I_2^{\tau_0})$ with forecasts $S2$ and then $(I_2^{\tau_1})$ with forecasts $S1$. Some reflection should convince the reader that, assuming all branches in Figure 4 have nonzero probabilities, there are 15 possible dynamic scenarios that one may have to solve, depending on which capacity scenarios are forecasted and which are realized. These 15 dynamic scenarios are given in Table 2.

To be quite explicit, the relationship between the “forecast” and the “realization” columns of Table 2 is the following: only a capacity scenario which can be realized may be forecasted, and *any* capacity scenario which can be realized may be forecasted. At τ_0 , for instance, all three capacity scenarios are possible, so any of them may be forecasted. At τ_1 , if $S1$ is realized, then only $S1$ can be forecasted, whereas, if $S1$ is not realized, then either $S2$ or $S3$ may be forecasted. In practice, of course, which capacity scenario will be forecasted will depend on the probabilities of the branches in Figure 4 and on the forecasting method (e.g., most-probable, worst-case, etc). Probabilities and forecasting methods will be introduced in the next section; for the moment we are just examining all possible cases.

We want to examine the possible dynamic scenarios before introducing probabilities of

Dynamic Scenario	Forecast at τ_0	Realization at τ_1	Forecast at τ_1	Realization at τ_2	Forecast at τ_2
1	$S1$	$S1$	$S1$	—	$S1$
1-2	$S1$	not $S1$	$S2$	$S2$	$S2$
1-3	$S1$	not $S1$	$S3$	$S3$	$S3$
1-2-3	$S1$	not $S1$	$S2$	$S3$	$S3$
1-3-2	$S1$	not $S1$	$S3$	$S2$	$S2$
2	$S2$	not $S1$	$S2$	$S2$	$S2$
2-1	$S2$	$S1$	$S1$	—	$S1$
2-3	$S2$	not $S1$	$S3$	$S3$	$S3$
2-2-3	$S2$	not $S1$	$S2$	$S3$	$S3$
2-3-2	$S2$	not $S1$	$S3$	$S2$	$S2$
3	$S3$	not $S1$	$S3$	$S3$	$S2$
3-1	$S3$	$S1$	$S1$	—	$S1$
3-2	$S3$	not $S1$	$S2$	$S2$	$S2$
3-2-3	$S3$	not $S1$	$S2$	$S3$	$S3$
3-3-2	$S3$	not $S1$	$S3$	$S2$	$S2$

Table 2: The 15 possible dynamic scenaria (cf. Figure 4-1).

realization because there are interesting and insightful comparisons to be made. As an example, the cost of scenario 3 must be lower than (or equal to) the cost of scenario 3-3-2, because in 3-3-2 the new capacity forecasts at τ_2 are lower than the previous forecasts, so some already departed aircraft may have to wait in the air. In other cases the comparison is less clear. Compare, for instance, the costs of scenaria 1-2-3 and 1-3-2. At τ_2 some departed aircraft in 1-3-2 may have to wait in the air, while this is not the case in 1-2-3. But the departed aircraft in 1-3-2 have probably had lower ground holds than in 1-2-3, since the previous capacity forecasts were more optimistic in 1-3-2 than in 1-2-3. There is thus a trade-off between ground and airborne delay costs, and it is interesting to pursue the investigation further.

Computational experiments were performed for four cases with $K = 3$ airports and $F = 1500$ flights. All cases have the same profile of scheduled arrivals but different percentages

of continued flights, F'/F , ranging from 0.20 to 0.80. The four cases are comparable, in the sense that cases with lower F'/F are obtained from cases with higher F'/F by eliminating some connections between flights. All four cases have slacks equal to 0 and identical arrival capacity profiles in the spirit of Figure 4, with the difference that $S2$ has a positive slope and $S1$ is constant (rather than $S1$ and $S2$ having negative slopes, as they have in Figure 4). $S1$ is at the infeasibility limit, and is equal to 11, 10, and 10 aircraft per time period for airports 1, 2, and 3, respectively. The two realization and decision times are $\tau_1 = 21$ and $\tau_2 = 41$; the time horizon has $T = 64$ time periods. Ground delay costs are 50 and airborne delay costs are 75.

The results of the computations are shown in Tables 3 and 4. It should be noted that the computations were performed with the LP relaxation (L_2^T) rather than (I_2^T), because the values of (L_2^T) and (I_2^T) are very close for identical cost functions [6,7]. Nevertheless, rounding mistakes are expected.

After some reflection, one can make the following observations on the basis of Tables 3 and 4. First, the 15 dynamic scenaria fall into two groups, call them group A and group B . Group A has three subgroups of three scenaria each. Subgroup $A1$ consists of scenaria 1, 2-1, and 3-1; subgroup $A2$ consists of scenaria 2, 1-2, and 3-2; and subgroup $A3$ consists of scenaria 3, 1-3, and 2-3. It can be seen that, for any of the four values of F'/F , all three scenaria within any of the above three subgroups have equal values. In symbols, we always have:

$$v_1 = v_{2-1} = v_{3-1} > v_2 = v_{1-2} = v_{3-2} > v_3 = v_{1-3} = v_{2-3}. \quad (41)$$

Why should this be so? The reason is probably that all scenaria within the above subgroups (with the exception of the static scenaria 1, 2, 3) have capacity updates only once, and at a rather early time period ($\tau_1 = 21$ with $T = 64$). Recall that all three

Dynamic Scenario	No capacity updates	Cap. updated at $\tau_1 = 21$ only	Cap. updated at $\tau_2 = 41$ only	Cap. updated at both τ_1 and τ_2
1 1-2 1-3 1-2-3 1-3-2	50,550	35,700 32,250		32,500 35,850
2 2-1 2-3 2-2-3 2-3-2	35,700	50,550 32,250	32,500	35,850
3 3-1 3-2 3-2-3 3-3-2	32,250	50,550 35,700	35,850	32,500
1 1-2 1-3 1-2-3 1-3-2	50,550	36,100 32,800		32,400 35,200
2 2-1 2-3 2-2-3 2-3-2	36,100	50,550 32,800	32,667	35,775
3 3-1 3-2 3-2-3 3-3-2	32,800	50,550 36,100	34,950	32,869

Table 3: Values of the 15 dynamic scenaria for $F'/F = 0.20$ (upper part of the table) and for $F'/F = 0.40$ (lower part).

Dynamic Scenario	No capacity updates	Cap. updated at $\tau_1 = 21$ only	Cap. updated at $\tau_2 = 41$ only	Cap. updated at both τ_1 and τ_2
1	63,188			
1-2		40,250		
1-3		35,750		
1-2-3				37,895
1-3-2				38,574
2	40,250			
2-1		63,188		
2-3		35,750		
2-2-3			36,020	
2-3-2				40,092
3	35,750			
3-1		63,188		
3-2		40,250		
3-2-3				37,013
3-3-2			38,447	
1	70,500			
1-2		46,100		
1-3		40,150		
1-2-3				42,974
1-3-2				45,922
2	46,100			
2-1		70,500		
2-3		40,150		
2-2-3			41,315	
2-3-2				44,755
3	40,150			
3-1		70,500		
3-2		46,100		
3-2-3				41,337
3-3-2			46,700	

Table 4: Values of the 15 dynamic scenarios for $F'/F = 0.60$ (upper part of the table) and for $F'/F = 0.80$ (lower part).

capacity scenaria coincide until τ_1 . In conclusion, (41) seems to say that, *if one can get the correct capacity forecasts early enough in the day, one can almost completely compensate for incorrect capacity forecasts made at the beginning of the day.* It should be noted that one does not always expect the equalities (41) to hold exactly. For instance, scenario 3-2 can sometimes have a higher value than scenario 2, because some aircraft in the air at τ_1 may have to incur airborne delays when the forecast shifts from $S3$ to $S2$.⁵

The second group, group B , consists of two subgroups of three scenaria each. Subgroup $B2$ consists of scenaria 1-3-2, 2-3-2, and 3-3-2; subgroup $B3$ consists of scenaria 1-2-3, 2-2-3, and 3-2-3.⁶ These subgroups are detected by means of the case $F'/F = 0.20$ (upper part of Table 3), where all scenaria within each of these two subgroups have equal values. In symbols:

$$v_{1-2-3} = v_{2-2-3} = v_{3-2-3} < v_{1-3-2} = v_{2-3-2} = v_{3-3-2}, \text{ for } F'/F = 0.20. \quad (42)$$

Examination of the other three cases for F'/F supports (42), although the equalities become approximate.

The equalities in group B have presumably the same origin as those in group A : incorrect forecasts are corrected early enough, so that their influence is minimized. Scenaria within each subgroup of group B (e.g., scenaria 1-2-3 and 3-2-3) differ only up to time $\tau_1 = 21$.

One can also compare scenario 3 with 2-2-3, or scenario 2 with 3-3-2. Here we have incorrect predictions which are corrected late in the day (at $\tau_2 = 41$ with $T = 64$). It is seen that only minor differences appear. In other words, the values of scenaria within subgroup $B2$ are quite close to the values of scenaria within subgroup $A2$ (similarly for subgroups $B3$

⁵For an example of a case where (41) do not hold exactly, see Table 7.

⁶The reason why the two subgroups of group B were named $B2$ and $B3$, rather than $B1$ and $B2$, is that subgroups $B2$ and $A2$ (and, similarly, subgroups $B3$ and $A3$) share an important feature: all scenaria within $B2$ and $A2$ end in 2. In other words, all six scenaria in these two subgroups correspond to the case where capacity scenario $S2$ is ultimately realized.

and A3). Therefore, it seems that obtaining the correct capacity forecasts even relatively late in the day suffices to minimize the impact of incorrect capacity forecasts made at the beginning of the day.

The main result of this subsection is that, if one obtains the correct capacity forecasts before the end of the day, then one can, for the most part, compensate for incorrect capacity forecasts made earlier on. Referring again to Figure 4, one can understand the reason. The difference between different dynamic scenarios is mainly due to aircraft which may have to wait in the air when the new forecast is lower than the previous one. Such aircraft are in the air at the current decision time period, so they will arrive at their destination soon. So the only differences between capacity scenarios that really matter are the differences in the vicinity of decision time periods. But such differences are usually small, because the capacity scenarios usually diverge smoothly rather than abruptly. This is why dynamic scenarios ending in i will have values quite close to v_i . What mostly matters is to obtain the correct forecasts, not *when* to obtain them.

4.2 Comparing forecasting methods.

Two forecasting methods will be examined: the most-probable forecast and the worst-case forecast. Referring again to Figure 4, at time τ_0 the *most-probable* forecast is $S1$ if $p_1 > 1/2$, while it is $S2$ if $p_1 < 1/3$ and $p_2 > 1/2$. All possible cases are given in Table 5. Similarly, at τ_1 , supposing $S1$ is not realized, $S2$ is most probable when $p_2 \geq 1/2$. The *worst-case* forecast, on the other hand, is obviously $S1$ at τ_0 and, if $S1$ is not realized, $S2$ at τ_1 .

For given values of the probabilities p_1 and p_2 , the expected value of the policy which uses the most-probable forecast is:

$$MP = p_1 v_{MP|1} + (1 - p_1)p_2 v_{MP|2} + (1 - p_1)(1 - p_2)v_{MP|3}, \quad (43)$$

Range of p_1	Range of p_2	Most probable scenario at τ_0
$p_1 > 1/2$	$0 \leq p_2 \leq 1$	$S1$
$1/3 < p_1 < 1/2$	$p_2 > p_1/(1 - p_1)$	$S2$
	$p_1/(1 - p_1) > p_2 > (1 - 2p_1)/(1 - p_1)$	$S1$
	$p_2 < (1 - 2p_1)/(1 - p_1)$	$S3$
$p_1 < 1/3$	$p_2 > 1/2$	$S2$
	$p_2 < 1/2$	$S3$

Table 5: Most probable capacity scenario at τ_0 for various probability combinations (cf. Figure 4).

where $v(MP|i)$ is the value of the dynamic policy given that capacity scenario i is ultimately realized. For instance, if $S2$ is realized and $p_1 > 1/2$, $p_2 < 1/2$, the most-probable forecasts will be $S1$ at τ_0 and $S3$ at τ_1 , so that $v_{MP|2} = v_{1-3-2}$. It is seen that $v_{MP|i}$ depends not only on i but also on p_1 and p_2 .

For given values of the probabilities p_1 and p_2 , the expected value of the policy which uses the worst-case forecast is:

$$WC = p_1 v_{WC|1} + (1 - p_1)p_2 v_{WC|2} + (1 - p_1)(1 - p_2)v_{WC|3}. \quad (44)$$

It is easily seen that, for the case of Figure 4, $v_{WC|1} = v_1$, $v_{WC|2} = v_{1-2}$, and $v_{WC|3} = v_{1-2-3}$, regardless of the values of p_1 and p_2 .

In order to compare the dynamic with the static GHP, we also included in the comparison the expected value of a *random-selection* static policy:

$$RS = p_1 v_1 + (1 - p_1)p_2 v_2 + (1 - p_1)(1 - p_2)v_3, \quad (45)$$

where v_i is the value of scenario i . The random-selection static forecast is defined by means of a probabilistic event: at τ_0 , one performs an experiment which yields outcomes 1, 2, and 3, with probabilities p_1 , $(1 - p_1)p_2$, and $(1 - p_1)(1 - p_2)$, respectively. If outcome $i \in \{1, 2, 3\}$ occurs, then the random-selection forecast is Si , and is not updated at later decision times.

Table 6 gives the expected values of the two dynamic policies under consideration and of the static random-selection policy for typical values of p_1 and p_2 corresponding to the combinations of Table 5. It is seen that the expected values of both dynamic policies are always very close to the expected value of the random-selection policy and to each other. It seems, therefore, that both forecasting methods perform equally well. On the other hand, the dynamic policies seem to result in no significant cost savings over the static random-selection policy.

These results were expected, given the conclusion of Subsection 4.1. In fact, $v_{MP|i}$ and $v_{WC|i}$ are always close to v_i , so (43), (44), and (45) entail the approximate equality of MP , WC , and RS .

4.3 A dynamic FCFS heuristic.

The following dynamic heuristic is inspired by current ground-holding practice. It gives priority to aircraft in the air over aircraft on the ground, and it assigns available capacity on an FCFS basis.

BEGIN

Initialize: $g_f = 0$.

FOR $t = \tau$ TO T DO:

FOR $k = 1$ TO K DO:

FOR $f \in \mathcal{F}_t^a$ DO:

IF $k_f^a = k$ AND $r_f + \hat{G}_f + \alpha_f = t$ AND f has a next flight \hat{f}

THEN $g_{\hat{f}} = \max(\hat{G}_f + \alpha_f - s_f, 0)$. Similarly if \hat{f} has a next

flight and so on.

CONTINUE f

F'/F	p_1	p_2	Random-selection	Most-probable	Worst-case
0.20	0.8	0.8	47,442	47,452	47,452
0.20	0.8	0.2	47,028	47,586	47,068
0.20	0.4	0.8	41,226	41,256	41,256
0.20	0.4	0.6	40,812	40,872	40,872
0.20	0.4	0.4	40,398	40,434	40,488
0.20	0.4	0.2	39,984	40,002	40,104
0.20	0.2	0.8	38,118	38,158	38,158
0.20	0.2	0.2	36,462	36,486	36,622
0.40	0.8	0.8	47,528	47,512	47,512
0.40	0.8	0.2	47,132	47,096	47,068
0.40	0.4	0.8	41,484	41,276	41,436
0.40	0.4	0.6	41,088	40,992	40,992
0.40	0.4	0.4	40,692	40,476	40,548
0.40	0.4	0.2	40,296	40,158	40,104
0.40	0.2	0.8	38,462	38,441	38,398
0.40	0.2	0.2	36,878	36,694	36,622
0.60	0.8	0.8	58,420	58,506	58,506
0.60	0.8	0.2	57,880	57,813	58,224
0.60	0.4	0.8	48,885	48,918	49,143
0.60	0.4	0.6	48,345	48,860	48,860
0.60	0.4	0.4	47,805	47,403	48,577
0.60	0.4	0.2	47,265	47,049	48,295
0.60	0.2	0.8	44,117	44,161	44,461
0.60	0.2	0.2	41,958	41,669	43,330
0.80	0.8	0.8	65,382	65,495	65,495
0.80	0.8	0.2	64,668	64,661	65,160
0.80	0.4	0.8	55,146	55,286	55,485
0.80	0.4	0.6	54,432	55,110	55,110
0.80	0.4	0.4	53,718	53,675	54,735
0.80	0.4	0.2	53,004	53,076	54,360
0.80	0.2	0.8	50,028	50,214	50,480
0.80	0.2	0.2	47,172	47,268	48,979

Table 6: Expected values of dynamic policies and of static random-selection policies for various probability combinations.

Define $S_k(t) := \{f \in \mathcal{F}_\tau^g : (k_f^a = k)(r_f + g_f = t)\}$.

Define $S_k(t) := |S_k(t)|$.

IF $S_k(t) > -E_k(t)$ THEN:

 Choose $Q_k(t) := S_k(t) + \min(E_k(t), 0)$ flights in $S_k(t)$.

 FOR $f = 1$ TO $Q_k(t)$ DO:

 Set $g_f = g_f + 1$.

 IF f has a next flight \hat{f} THEN:

 IF $g_f > s_f$ THEN:

 Set $g_{\hat{f}} = g_{\hat{f}} + 1$, and similarly if \hat{f} has a next flight and so on.

 END IF

 END IF

 CONTINUE f

END IF

CONTINUE k

CONTINUE t

END

Computations were performed for the same case as in Subsection 4.1 (3 airports and 1500 flights), but only for $F'/F = 0.80$, and with slacks equal to 1 instead of 0. The new capacity scenaria are lower, since the infeasibility limit is lower (due to the increase of the slack). The new $S1$ is equal to 10, 9, and 9 aircraft per time period for airports 1, 2, and 3, respectively.

The values of the 15 dynamic scenarios, both for formulation (I_2^T) and for the dynamic FCFS heuristic, are given in Table 7. It is seen that, as was the case in the static problem, the FCFS heuristic always results in a cost much higher than what one can achieve by solving the multi-airport GHP optimally. This is a significant result, since the dynamic FCFS heuristic is, to some extent, a reasonable approximation of the essential aspects of current ground-holding practice.

Finally, Table 8 gives the expected values of all four dynamic policies and of the corresponding two static random-selection policies for the case under consideration. It is seen that the two dynamic policies which use the FCFS heuristic perform about equally well, but perform much more poorly than the two policies which use formulation (I_2^T). Moreover, in accordance with the results of the previous subsection, it is seen that the dynamic policies result in no significant cost savings over the static random-selection policies.

4.4 Summary of results.

This section has reached the following main conclusions on the dynamic GHP.

(1) If incorrect capacity forecasts made at the beginning of the day are corrected early enough, then their influence on the total cost of the dynamic problem can be minimized.

(2) The most-probable and the worst-case methods of forecasting capacities given probabilistic information perform about equally well, if their performance is measured by their expected values.

(3) The dynamic policies using either a most-probable or a worst-case forecasting method perform, in terms of their expected values, about as well as a static random-selection policy. (Note, however, that the static random-selection policy will have a greater variance than the dynamic policies, since the latter are designed to minimize the effect of initial incorrect

Dynamic Scenario	No capacity updates	Cap. updated at $\tau_1 = 21$ only	Cap. updated at $\tau_2 = 41$ only	Cap. updated at both τ_1 and τ_2
1	151,850 (97,163; 56.3%)			
1-2		108,600 (62,814; 72.9%)		
1-3		104,500 (60,765; 72.0%)		
1-2-3				99,750 (63,805; 56.34%)
1-3-2				103,825 (64,373; 61.3%)
2	108,550 (63,050; 72.2%)			
2-1		151,900 (98,034; 54.9%)		
2-3		104,450 (61,975; 68.5%)		
2-2-3			99,700 (65,194; 52.9%)	
2-3-2				103,775 (66,550; 55.9%)
3	104,450 (61,000; 71.2%)			
3-1		151,900 (97,347; 56.0%)		
3-2		108,550 (63,287; 71.5%)		
3-2-3				99,700 (63,307; 57.5%)
3-3-2			103,775 (66,650; 55.7%)	

Table 7: Values of (a) the dynamic FCFS heuristic, and (b) (in parentheses) the corresponding exact optimum and the percentage of cost overestimation resulting from applying the heuristic (overestimation = (value of heuristic/exact optimum)-1).

F'/F	p_1	p_2	RS	MP	WC	RS_{FCFS}	MP_{FCFS}	WC_{FCFS}
0.80	0.8	0.8	90,258	90,333	90,333	143,026	142,846	142,846
0.80	0.8	0.2	90,012	90,028	90,452	142,534	142,353	141,784
0.80	0.4	0.8	76,449	77,301	76,673	125,378	124,828	124,838
0.80	0.4	0.6	76,203	76,791	76,791	124,886	123,776	123,776
0.80	0.4	0.4	75,957	76,928	76,910	124,394	123,278	122,714
0.80	0.4	0.2	75,711	76,217	77,029	123,902	123,349	121,652
0.80	0.2	0.8	69,545	70,390	69,842	116,554	115,804	115,834
0.80	0.2	0.2	68,561	69,173	70,318	114,586	113,832	111,586

Table 8: Expected values of dynamic policies and of static random-selection policies for various probability combinations.

predictions and succeed in so doing, given (1) above.)

(4) Finally, and most importantly, assigning ground holds by means of the dynamic FCFS heuristic is highly inefficient compared with optimal ground holds based on the formulations of Chapter 4. This result is particularly important because the dynamic FCFS heuristic may approximate, to some extent, current ground-holding practice.

5 Conclusions.

This paper presented an efficient method for modeling various versions of the dynamic multi-airport GHP. It should be clear by now that the formulations of this paper constitute a *method* rather than a mere bag of tricks. By this we mean that they are sufficiently flexible to accommodate various degrees of modeling detail, rather than being rigidly limited to some particular cases.

In order to appreciate the flexibility of the method, recall that, starting with the simplest formulations (I_1^T) and (I_2^T) for the static deterministic case, the following extensions were presented:

(1) Flights may be cancelled. When a continued flight is cancelled, the next flight scheduled to be performed by the same aircraft can be either also cancelled or not affected at all (because, e.g., it is performed by a spare aircraft).

(2) Continued flights may have more than one “next” flights, e.g., because passengers in the continued flight connect to several other flights (especially in hub airports).

(3) Aircraft in the air may be speeded up, even to the point of arriving at their destination before their scheduled landing time.

(4) Departure and arrival capacities may be interdependent; their mix may be modified from time to time by controlling runway use.

Two points are worth making here. First, the above list is an open one. New generalizations may be devised according to the demands of the real-world situation. Second, and most important, these generalizations are not mutually exclusive: they can be combined to yield even more general formulations.

As far as the computational results are concerned, it was found that large-scale GHPs involving as many as 3 congested airports and 1500 flights are solvable. It was shown that optimal ground holds can result in significant cost savings with respect to a heuristic approximating, to some extent, current ground-holding practice. It was also found that what matters is to eventually obtain the correct capacity forecast, not *when* to obtain it.

The following directions for future research on the dynamic multi-airport GHP seem worth pursuing.

(1) The real-world GHP should be further investigated in order to find out whether any interesting aspects have been left out in the formulations of this paper. If such aspects are found, it is possible that they could be incorporated in the formulations by means of additional constraints or other modifications. As an example, one could try to extend the

formulations in order to take into account that delays sometimes occur not only because of limited airport capacities but also because of congestion in terminal area and *en route* airspace.

(2) In order to further reduce running time and to make it feasible to assign ground holds on a real-time basis even for a very large network of airports, special-purpose algorithms for solving the formulations should be sought. Such algorithms could take into account the special structure of the constraint matrix.

(3) Reduction of computation times could also be achieved by looking for heuristics that would provide good feasible solutions of the formulations, rather than insisting on optimal solutions.

(4) Finally, the most important task for the future is to apply in practice the ground holds arising from the formulations, at first on a limited scale or in a simulation environment, in order to obtain realistic feedback which would then spur further research.

We conclude by expressing the earnest wish that the work presented in this paper will ultimately find its way into actual practice and will result in a more efficient way of operating the air travel network.

Acknowledgments

We would like to thank the Charles S. Draper Laboratory for supporting our 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.

References

- [1] 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.
- [2] Andreatta, G. & G. Romanin-Jacur. "Aircraft Flow Management under Congestion". *Transportation Science*, v.21 (1987): pp.249–253.
- [3] Mostafa Terrab. "Ground Holding Strategies For Air Traffic Control". Ph.D. thesis, Massachusetts Institute of Technology, February 1990.
- [4] Octavio Richetta. "Ground Holding Strategies for Air Traffic Control under Uncertainty". Ph.D. thesis, Massachusetts Institute of Technology, June 1991.
- [5] Transportation Research Board. *Winds of Change: Domestic Air Transport since Deregulation*. National Research Council, Washington, D.C., 1991.
- [6] Peter B. Vranas, Dimitris J. Bertsimas, & Amedeo R. Odoni. "The Multi-Airport Ground-Holding Problem in Air Traffic Control". Submitted for publication to *Operations Research*, February 1992.
- [7] Peter B. Vranas. "The Multi-Airport Ground-Holding Problem in Air Traffic Control". Ph.D. thesis, Massachusetts Institute of Technology, June 1992.