

# A Critical Survey of Optimization Models for Tactical and Strategic Aspects of Air Traffic Flow Management

Dimitris Bertsimas and Amedeo Odoni  
Massachusetts Institute of Technology

FINAL  
11/28/97  
0017  
000343

## 1. Introduction

This document presents a critical review of the principal existing *optimization models* that have been applied to Air Traffic Flow Management (TFM). Emphasis will be placed on two problems, the Generalized Tactical Flow Management Problem (GTFMP) and the Ground Holding Problem (GHP), as well as on some of their variations.

To perform this task, we have carried out an extensive literature review that has covered more than 40 references, most of them very recent. Based on the review of this emerging field our objectives were to:

- (i) identify the best available models;
- (ii) describe typical contexts for applications of the models;
- (iii) provide illustrative model formulations; and
- (iv) identify the methodologies that can be used to solve the models.

We shall begin our presentation below by providing a brief context for the models that we are reviewing. In Section 3 we shall offer a taxonomy and identify four classes of models for review. In Sections 4, 5, and 6 we shall then review, respectively, models for the Single-Airport Ground Holding Problem, the Generalized Tactical FM P and the Multi-Airport Ground Holding Problem (for the definition of these problems see Section 3 below). In each section, we identify the best available models and discuss briefly their computational performance and applications, if any, to date. Section 7 summarizes our conclusions about the state of the art.

## 2. Context and Terminology

It is important, at the outset to set a context for this review and to introduce some terminology.

It is well known that traffic congestion is a critical problem in the most developed air transportation systems in the world (North America, Western Europe and East Asia). If one combines the various numbers reported, worldwide direct airline costs due to congestion add up to well over \$5 billion per year --an estimate that does not include the cost of lost time to the passengers themselves. This is more than the total profits or losses of *all* the IATA airlines in a typical year. With significant increases in demand expected to continue, congestion is now cited by many experts and organizations as one of the principal impediments to the future growth of the airline industry.

Congestion occurs whenever the capacity of airport runway systems and/or of ATC sectors is exceeded over a period of time. Thus, it is mostly associated with peak traffic hours of the day or peak travel times in the year, as well as with periods of poor weather conditions when airport or en route sector service rates can be significantly reduced. In the absence of the long-term capacity improvements that can be obtained through the construction of additional runways or through advances in ATC, traffic flow management (TFM) is the best available way to reduce the cost of delays. On a day-to-day basis, TFM attempts to “match”, dynamically, air traffic demand with the capacity of airports and airspace sectors of the ATC system. *Ground-holding* and *re-distribution of flows in the airspace* are the two principal devices used by TFM for this purpose.

Ground-holding (or “*gate-holding*” or “*ground-stopping*”) is typically imposed on aircraft flying to congested airports or scheduled to traverse congested airspace. *Ground-holding is the action of delaying take-off beyond a flight's scheduled departure time.* The motivation for doing so is that, as long as a delay is unavoidable, it is safer and less costly for the flight to absorb this delay on the ground before take-off, rather than in the air. Ground holding, as part of TFM, is a relatively recent phenomenon. It was initiated as an *ad hoc* practice in Europe in the 1970's in response to growing air traffic congestion and its use gradually

increased during the 1980s and early 1990s. Similarly, in the United States, ground-holding was used rather sparingly in ATC prior to 1981, i.e., until that year flights were allowed to take-off as soon as ready to do so, except in very unusual circumstances. If delays were encountered, they were absorbed while the aircraft was airborne, typically by circling in the air ("stacking") near the airport of destination. However, widespread use of ground-holding began during the 1981 air traffic controllers' strike in the United States, primarily because this was seen as a way to reduce controller workload by limiting the number of aircraft which were airborne at any given time. When it was realized that ground-holding was also a fuel-saving practice, its use became part of established TFM practice, just as in Europe.

In addition to ground holding, TFM has several other options available to it. As noted above, the most important of these is the *re-distribution of air traffic flows* over the network of airways. This re-distribution can be effected through changes in the routing of flights and can be accomplished in two ways: *strategically*, i.e., by planning in advance the routes of scheduled flights in a region in a way that ensures a desirable distribution of traffic flows; or *tactically*, by re-routing aircraft in "real time", possibly changing an aircraft's flight plan even after that aircraft is already airborne.

Other options beyond ground holding and re-distribution of air traffic flows, include: *speed control* of airborne aircraft; *metering* of air traffic, i.e., controlling the rate at which aircraft go past a given point in airspace; and *airborne holding* en route and, especially, near or inside terminal airspace.

### **3. Model Taxonomy**

We shall classify TFM optimization models next into several distinct classes, according to the type of problem they address. Specifically, we want to review models that address the following TFM problems:

(a) *Ground-Holding Problem (GHP)*: Models in this class are of a tactical nature and attempt to assign ground holding delays to flights, with the objective of minimizing the cost of delays to AOs, while satisfying any existing constraints on ATM capacities at airports or en route. We subdivide the GHP into two subproblems, the *Single-Airport Ground-Holding Problem (SAGHP)* and the *Multi-Airport Ground-Holding Problem (MAGHP)*. As their respective names suggest, the two problems consider, respectively, a single airport at a time (SAGHP) and an entire network of airports simultaneously (MAGHP). In the SAGHP, ground holding times are assigned to flights scheduled to travel to some particular airport A, where scheduled demand is expected to exceed available capacity during some period of time during the day of interest. In the MAGHP, delays are assumed to propagate in the network of airports, as aircraft perform consecutive flights, thus necessitating the examination of an entire set of airports simultaneously.

The GHP can be further subdivided into a “deterministic” version (*deterministic GHP*) and a probabilistic version (*stochastic GHP*). The stochastic version arises because the GHP must often be solved in the presence of considerable uncertainty. In other words, deciding how much ground-holding delay to assign to a flight is complicated by the fact that, in practice, it is often difficult to predict how much delay a flight will actually suffer. The reason is that sector capacities and, especially, airport capacities are oftentimes highly variable and may change dramatically during the course of a day, as weather changes or other events occur. Moreover, small changes in visibility or in the height of the cloud-cover may translate into large differences in airport capacity. It is nearly impossible for meteorologists today to predict these changes to such a high level of accuracy, even over a very short time-horizon of an hour or less. Under such circumstances, ground-holding decisions must be made under uncertainty and must consider the trade-off between “conservative” strategies that may at times assign excessive ground-holds and more “liberal” ones that

may result in more expensive airborne delays. AOs often complain that today's TFM systems may be erring too much on the conservative side, i.e., that there may be too many instances in which aircraft are unnecessarily delayed on the ground by TFM.

(b) *The Generalized Tactical TFM Problem (GTFMP)*: This is the generalized version of the GHP, where in addition to determining release times for aircraft (ground-holds), we also wish to take into consideration the possibility of assigning some airborne delays to flights at specific points on their route. These delays could be absorbed through airborne holding at these points or possibly by exercising speed control or metering of the traffic flow.

Additional problems and types of models could be defined. For example, if we add the possibility of re-routing flights while airborne, we obtain the Traffic Flow Management Rerouting Problem (TFMRP). In this problem, a flight may be re-routed in real time through a different flight path in order to reach its destination, if the current route passes through a region that unexpectedly becomes congested or should be avoided for other reasons, usually related to poor weather conditions.

In this review, we shall be concerned with models addressing only the STFMP, the SAGHP, the GTFMP and the MAGHP. We shall review existing models in each one of these areas in the order just indicated. The MAGHP will be discussed after the GTFMP, because some of the existing models treat the MAGHP as a special case of the TFMP. It is therefore more convenient to familiarize the reader with GTFMP models before proceeding to the MAGHP.

#### **4. The Single-Airport, Ground-Holding Problem (SHGHP)**

The most general version of the SAGHP (dynamic and stochastic) may not be appropriate for many practical situations. Depending on how much information is available and on how this information is updated, alternative versions of the SAGHP may have to be solved. Therefore, several versions of the problem have been addressed by the existing optimization models. For example, deterministic (rather than stochastic) versions will be preferable for locations where either the weather or the airport capacities are stable enough to be approximated as perfectly predictable quantities. Moreover, it should be noted that existing ATC systems never deal explicitly with “probabilities” and thus deterministic models may approximate today’s practice better than stochastic ones. Similarly, static (rather than dynamic) versions may be more appropriate for environments where (i) there are significant lags in updating information concerning weather or capacities at a set of geographically dispersed locations or (ii) an initial ground-holding plan is prepared at a single point in time (typically at the beginning of the day) and that plan is revised only in a marginal way from that point on.

In this section, we shall present three models for the single-airport GHP that follow exactly the progression outlined above. The first of them, Model 1, is a deterministic model which assumes that the capacity at the destination airport can be forecast with perfect accuracy for the entire period of interest. Model 1 determines how long each aircraft should be kept on the ground before take-off to minimize the total cost of delays. While this model is simple, it will help introduce some fundamental notions and issues, such as the issue of whether delays are distributed equitably among the various aircraft operators (airlines and other airspace users). The question of equitability is often referred to as the “fairness” issue in Traffic Flow Management.

Model 2 will then introduce the problem of uncertainty in the SAGHP by recognizing that future airport capacities (often even over the next hour) cannot be forecast with perfect accuracy. Model 2 therefore treats the GHP as a stochastic decision problem. Model 2, however, does include a simplification by assuming a static environment, i.e., makes ground-holding decisions “once and

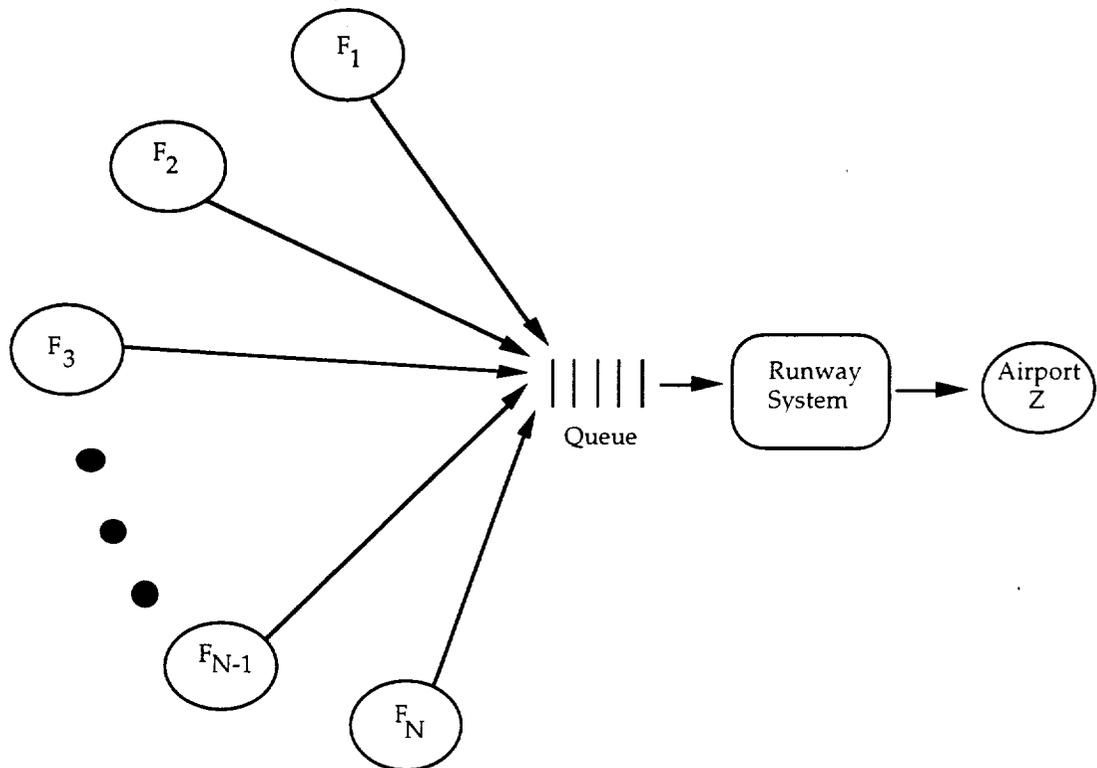
for all" for the entire time period under consideration. By contrast, Model 3 is the most general version of the SAGHP, as it is both stochastic and dynamic.

Models 2 and 3 represent, to our knowledge, the most advanced existing models for the single-airport GHP. They were both developed in the early 1990's and are discussed in far more detail in Andreatta, Odoni and Richetta (1993) --on which this Section largely draws-- and, especially, in Richetta and Odoni (1993, 1994).

#### **4.1 The Simplified Network for the SAGHP Models**

The air traffic network model considered for the single-airport GHP can be described with reference to the single-destination network shown in Figure 1. The model is macroscopic in nature, yet it captures the essential elements needed to solve the single-airport GHP:

- (i)  $N$  aircraft (flights  $F_1, \dots, F_N$ ) are scheduled to arrive at the congested "arrivals" airport  $Z$  from the "departure" airports.
- (ii) Airport  $Z$  is the only capacitated element of the network and thus the only source of delays. All other elements in the network (departure runways, airways, etc.) have unlimited capacity.
- (iii) The departure and travel times of each aircraft are deterministic and known in advance.
- (iv) The time interval of interest is  $[0, B]$ , with the earliest departure for  $Z$  scheduled at 0 and the latest arrival scheduled at  $B$ . The time interval  $[0, B]$  is discretized into  $T$  equal time periods numbered  $1, 2, \dots, T$ . (For example  $B=12$  hours could be subdivided into  $T=72$  periods of 10 minutes each.)
- (v) Ground and air delay cost functions for each flight are known. This means that, for each aircraft, we can estimate the cost of its being delayed for  $x$  minutes in the air and  $y$  minutes on the ground (before takeoff) for any non-negative values of  $x$  and  $y$ .



**Figure 1 "Star" Configuration Network**

Items (ii) and (iii) amount to an assumption of perfectly predictable travel times between the airport of departure of each flight and airport Z (for additional details see Odoni [6]). Thus, this ignores the effect of such tactical actions as speed control, "metering" and path-stretching that may sometimes take place during the "en route" portion of a flight in response to local ATC conditions. It is assumed that the impact of such actions on operating costs is entirely secondary to that of the delays due to congestion at the airport of destination Z - and, thus, to the ground-hold versus air delay trade-off. This is fully justified for the United States ATC system: although no specific statistics are collected on the matter, an overwhelming proportion of delays (certainly more than 95%) are undoubtedly due to airport, not en route, capacity limitations. This is often not the case, however, in Western Europe, where the en route airspace imposes just as severe capacity constraints as the airports - primarily as a result of institutional and political factors. We shall return to this particular point in Part II of this review, when we shall examine more general, capacitated flow models.

## 4.2 Early Models

Andreatta and Romanin-Jacur (1986a, 1986b, 1987, 1988) seem to have been the first to develop optimal approaches to the GHP. They, in fact, addressed a probabilistic version of the SAGHP. They considered a greatly simplified model which highlights some important conceptual aspects of all probabilistic versions of the GHP. This model assumes that congestion may arise only during a *single* given period of time at the specified "arrival" airport Z. The model therefore is neither multi-period nor dynamic, but takes explicitly into account the stochastic nature of airport capacity.

Terrab (1990) presented the first versions of multi-period single-airport models both for a deterministic (see Section 5.3 below) and a stochastic environment. His stochastic formulation was so "fine grain" that it required enormous computational effort to solve, especially as he was using a dynamic programming solution approach. Terrab (1990) also presented a number of interesting heuristic approaches to the stochastic problem. These, however, have been superseded by Models 2 and 3 below.

## 4.3 A Deterministic Model

The simplest multi-period model [Terrab (1990), Odoni and Terrab (1993)] assumes that the capacity of airport Z is a deterministic function of time, known in advance for the entire period of interest. The time horizon consists of T periods in which capacity may be limited (and takes on a different value  $K_j$  in period j) and of an extra period, T+1, in which capacity is large enough so that all the aircraft that are still in the air can land during T+1. This assumption about period T+1 is, of course, true in practice: late at night there is always enough capacity to accommodate all the remaining requests for landings. The model can be formulated mathematically as follows:

### Model 1

$$\text{MIN} \sum_{i=1}^N \sum_{j=t(i)}^{T+1} G_{ij} X_{ij}$$

subject to:

$$\sum_{j=t(i)}^{T+1} X_{ij} = 1 \quad \text{for all } i = 1, \dots, N$$
$$\sum_{i=1}^N X_{ij} \leq K_j \quad \text{for all } j = 1, \dots, T$$
$$X_{ij} \in \{0,1\}$$

where:

$N$  is the number of flights scheduled to land;

$T$  is the number of periods when capacity may be limited;

$K_j$  is the capacity in period  $j$ ;

$X_{ij}$  are the decision variables:  $X_{ij} = 1$  means that aircraft  $i$  will be assigned to land in period  $j$  and  $X_{ij} = 0$  otherwise;

$G_{ij}$  is the cost incurred by aircraft  $i$  when assigned to land in period  $j$ ;

$t(i)$  is the period of time during which aircraft  $i$  was originally scheduled to land.

The following two important observations can be made:

- (i) If the capacity  $K_j$  is known with certainty for all  $j=1, \dots, T$  (as is the case in this deterministic model), then, under any optimal ground holding policy, aircraft may suffer ground holds but never airborne delays. As long as the cost of ground delay is less than the cost of airborne delay per unit of time for any given aircraft, it will always be better to hold an airplane on the ground rather than in the air. For this reason the relevant costs  $G_{ij}$  in Model B are the ground-holding costs for flight  $i$ ; for the deterministic GHP, detailed information on airborne delay costs is not necessary.

- (ii) The above formulation proves that this version of the GHP is of polynomial computational complexity: in fact, Model B is a particular case of the so-called "Transportation Problem" and can be transformed into an "Assignment Problem". Many numerical instances of Model B have been solved using standard Minimum Cost Flow algorithms [Terrab (1990)].

The experimental results reported in Terrab (1990) and in Terrab and Odoni (1993) show that, even with a deterministic knowledge of airport capacity, large savings in total delay costs can be achieved through solutions that assign available capacity optimally. However, the implementability of such optimal strategies is made questionable by the systematic biases they exhibit: typically, they assign most ground-holds to aircraft with low delay costs per unit of time (e.g., general aviation aircraft and regional airlines) while giving priority to aircraft with high delay costs (e.g., wide-body aircraft). In our earlier terminology, such solutions distribute delay "unfairly". It may therefore be necessary to impose additional constraints that force a more equitable distribution of ground-holds. This can be done quite effectively and still yield strategies with significant savings in the cost of delay [Terrab (1990)].

It is also important to anticipate at this point one of the major conclusions of research into TFM problems, namely that it is also possible to achieve large savings, *without discriminating at all* among various types of aircraft. This can be done through good (deterministic or stochastic, static or dynamic) models.

#### **4.4 An Optimal Stochastic and Static Model**

The model to be discussed in this section [Richetta (1991), Richetta and Odoni (1993)] solves, optimally and with very reasonable computational effort, realistic instances of the multi-period, stochastic GHP under only mildly restrictive assumptions. The main feature of this model is that it simplifies the structure of the control mechanism by making ground hold decisions on groups of aircraft (i.e., on aircraft classified according to cost class, and schedule) rather than individual flights. The model will be described here for a static environment, but it can be extended to the dynamic case, as will be explained in the next section.

This model is motivated by the observation that, in practice, the number of alternative capacity profiles (henceforth called scenarios) for airport Z, that can be forecast and dealt with on any given day is small. The model assumes explicitly that there are  $Q$  such scenarios (where  $Q$  is a small number), each having a given probability of materializing. For example, one scenario may call for poor visibility conditions to begin 3 hours from now and last for 4 consecutive hours, whereas another scenario might forecast 4 hours to the beginning of poor visibility conditions and 5 hours for their duration. Model 2 below allows us to make ground-holding decisions in the face of this type of uncertainty. Other important assumptions which permit further reductions in the computational complexity of the model are:

- (i) Aircraft can be classified into a small number of different classes (typically 3 or 4) with aircraft in each class having essentially identical ground-holding delay costs. Let  $C_g^k(i)$  be the function representing the cost of ground holding an aircraft of class  $k$  for  $i$  consecutive time periods
- (ii) The cost of delaying one aircraft in the air for one time period is a constant  $c_a$ , independent of the type of aircraft. Thus,  $c_a$  may be considered as an overall average cost of waiting in the air. This assumption might seem unnatural at first, but is actually based on the following "operational" characteristics of the ATC system:
  - (1) Aircraft which are already airborne are almost always sequenced by ATC in an approximately first-come, first-served (FCFS) way; therefore, there is no need to draw distinctions among different classes of aircraft while airborne.
  - (2) Within reasonable airborne delay levels (i.e., for up to the largest airborne delays observed in practice, which are in the order of one hour) delay cost functions are approximately linear, since "non-linearities", due to factors such as safety, do not yet set in.

Furthermore, computational results have shown that the relative magnitude of average ground and air delay costs affects the ground hold

strategy selected much more significantly than modeling air delay costs in greater detail. This observation provides further support for the “constant  $c_a$ ” assumption.

We shall use the following notation next:

$N_{ki}$  is the number of aircraft of class  $k$  scheduled to arrive at the destination airport during period  $i$  ( $k=1,\dots,K; i=1,\dots,T$ );

$M_{qi}$  denotes the airport capacity in period  $i$  under scenario  $q$  ( $q=1,\dots,Q; i=1,\dots,T$ );

$X_{qkij}$  represents the number of aircraft of class  $k$  originally scheduled to arrive at the destination airport during period  $i$ , and rescheduled to arrive during period  $j$  under capacity scenario  $q$ , due to a ground delay of  $j-i$  time periods ( $q=1,\dots,Q; k=1,\dots,K; i=1,\dots,T; i \leq j \leq T+1$ );

$W_{qi}$  are auxiliary variables representing the number of aircraft unable to land at the destination airport during period  $i$  under capacity scenario  $q$ , i.e., the number of aircraft incurring airborne delay during period  $i$  ( $q=1,\dots,Q; i=1,\dots,T$ );

It should be noted that the decision variables  $X_{qkij}$  defined above are more “aggregate” than the decision variables of the previously discussed Model 1. In the latter model, we were concerned with ground-hold delays at the individual flight level (i.e., Model 1 is a “fine-grain” model) whereas we have now defined somewhat more “aggregate” decision variables: how many flights scheduled to arrive in period  $i$  will instead be rescheduled for period  $j$ .

### Step 1: Deterministic Model.

A first step toward developing a stochastic model is to write a formulation for the deterministic case. Assume that it is known with certainty that, on a given day, a particular capacity scenario  $q$  will materialize, i.e., there is only one scenario to be considered. We then have:

## Model 2a

$$\text{Minimize Cost}(q) = \sum_{k=1}^K \sum_{i=1}^T \sum_{j=i+1}^{T+1} C_g^k (j-i) X_{qkij} + c_a \sum_{i=1}^T W_{qi}$$

subject to:

$$\sum_{j=i}^{T+1} X_{qkij} = N_{ki} \quad k=1, \dots, K; \quad i=1, \dots, T$$

$$W_{qi} \geq \sum_{k=1}^K \sum_{h=1}^i X_{qkhi} + W_{qi-1} - M_{qi} \quad i=1, 2, \dots, T+1$$

$X_{qkij}, W_{qi} \geq 0$  and integer.

The objective function in Model 2a minimizes total (ground plus air) delay costs. The first set of constraints states that all flights scheduled to land during period  $i$  must be rescheduled to arrive at  $i$  or later. The second set represents the flow balance at airport  $Z$  at the end of each time period  $i$ .

One can easily check that the coefficient matrix of Model D is totally unimodular. So one can relax the integrality constraints, since they will be satisfied by any basic feasible solution of the corresponding Linear Programming problem. In fact, it is convenient to model this as a Minimum Cost Flow problem on a network and to solve it through specialized algorithms.

### Step 2: Stochastic Static model.

Suppose now that, in the situation described above, each capacity scenario  $q$  has probability  $\text{Prob}(q)$  of materializing. Clearly, the objective function must be "weighted" over all possible capacity scenarios. Furthermore, since one still has to make ground-hold decisions at  $t=0$ , before knowing which airport capacity scenario will materialize, the decision variables corresponding to each capacity profile must be "coupled" with those corresponding to all the other capacity profiles (see coupling constraints in Model 2b below). The following stochastic

model of the GHP is obtained (see Richetta and Odoni (1993) for additional details):

**Model 2b:**

$$\text{Minimize } \sum_{q=1}^Q \text{Cost}(q) \text{ Prob}(q)$$

subject to:

set of constraints for q=1

set of constraints for q=N

Coupling constraints:  $X_{1kij} = \dots = X_{Qkij} \quad \forall k; i; j$

where, for each value of q from 1 to Q, there is a set of constraints identical to those for the deterministic Model 2a.

Model 2b can be viewed as a Stochastic Programming problem with one stage. It is suitable for application of decomposition techniques and lends itself well to parallel computation. Richetta and Odoni (1993) make the conjecture that the coefficient matrix is unimodular, since, in all their numerical experiments where integrality constraints were relaxed, they obtained integer solutions. The proof of this conjecture, however, is still an open question.

As an alternative to decomposition, one can substitute directly the coupling constraints into the rest of model, hence reducing both the number of variables and the number of constraints.

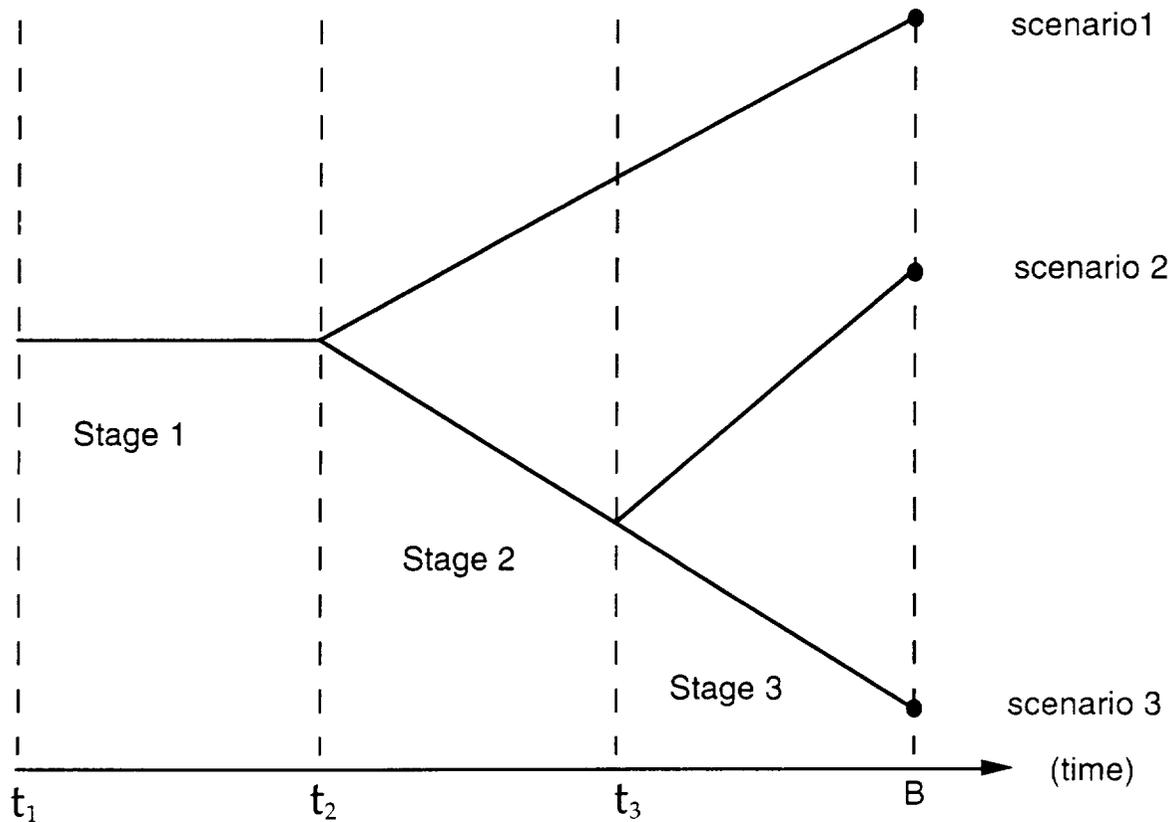
## 4.5 An Optimal Stochastic and Dynamic Model

Before describing the extension of Model 2b to the dynamic, stochastic SAGHP, an observation about a fundamental difference between dynamic and static models must be made. In the dynamic problem, ground-holding strategies are revised over time as capacity forecasts are updated. Strategy revisions take into consideration the “current state” of the ATC system, including any earlier decisions regarding ground-holds. Thus, the expected cost of ground plus air delays is minimized by deciding, *at the beginning of each time period*, whether eligible flights will be allowed to depart or will be held on the ground. By contrast, the static solution imposes “once and for all” ground-holds at time zero (i.e., at the beginning of the first time period of the day). In terms of modeling, this means that, in addition to scheduled arrival times, *scheduled departure times* must also be considered explicitly in developing dynamic strategies.

The dynamic version of Model 2b [Richetta and Odoni (1994)] can now be described with reference to Figure 2. The dynamic evolution of the capacity forecasts and the implicit updating of the associated probabilities is modeled through a probability tree. Taking a forecast consisting of three capacity scenarios for airport Z as an example, Figure 2 shows that the forecast is updated three times during the interval  $[0, B]$ . These three instants (denoted  $t_1$ ,  $t_2$  and  $t_3$  in Figure 2) define three stages comprising the time intervals  $[t_1, t_2)$ ,  $[t_2, t_3)$ , and  $[t_3, B]$  (a capacity forecast consisting of  $Q$  capacity scenarios would consist of at most  $Q$  stages). Within each stage, the probability of each of the scenarios for future airport landing capacity does not change. Therefore, an optimal dynamic solution to the GHPP assigns and/or revises ground-holds at the beginning of each stage. The time at which stage  $s$  starts will be referred to as  $t_s$  below.

As in the static Model 2b, an implicit assumption here is that the number of alternative scenarios at the beginning of each day,  $Q$ , is quite small - probably 4 or less. This assumption is important for obtaining quick numerical solutions. A small value of  $Q$  is, once again consistent with current weather forecasting technology which has advanced to the point where the type of weather conditions in a specific geographic area can be predicted with reasonable accuracy, but the exact timing of weather fronts and their local severity are uncertain. A typical example of the type of situation that can be addressed

through this approach is one in which, at the beginning of a day, there is an expectation of some deterioration in weather conditions in early afternoon which may result in severe loss of landing capacity (scenario 1) limited loss (scenario 2) or no loss at all (scenario 3), each with a certain probability.



**Figure 2**  
**Number of Stages Defined by the Times of Possible Changes in Airport Capacities**

In the static solution to the GHP, the time interval  $[0, B]$  comprises a single stage, resulting in a "here-and-now" solution which assigns ground-holds at  $t_1$ . In the dynamic case there are up to  $Q$  stages at which we make ground hold decisions.

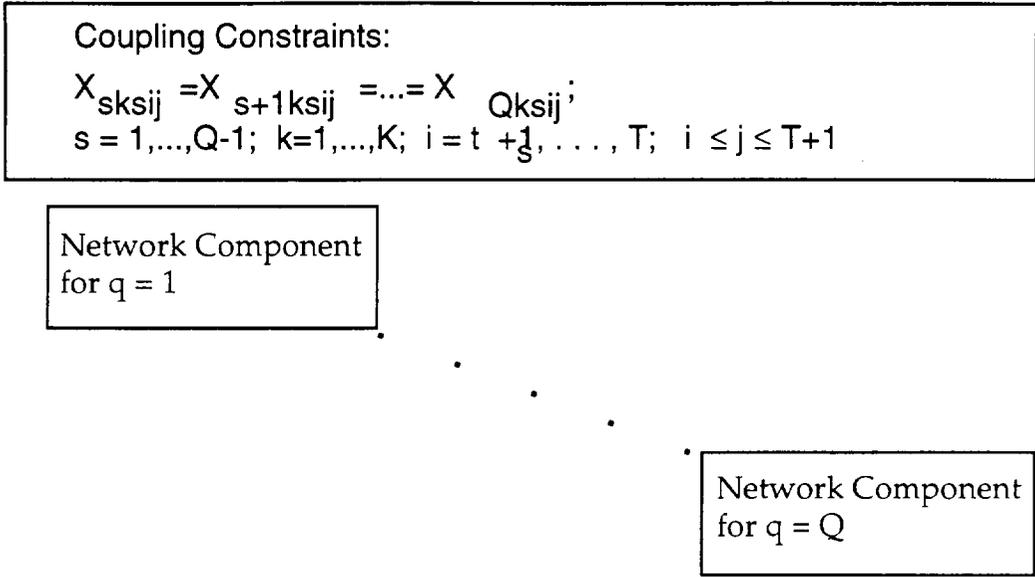
### Model 3

Due to the role of scheduled departure times, the dynamic Model 3 requires a modification in some of the notation defined in the previous section. Specifically:

$N_{ksi}$  is the number of aircraft of class  $k$  scheduled to depart during stage  $s$  and arrive at airport  $Z$  during period  $i$  ( $k=1,\dots,K$ ;  $s = 1,\dots,Q$ ;  $i = t_s+1, t_s+2,\dots, T$ )

$X_{qksij}$  is the number of aircraft of class  $k$  scheduled to depart during stage  $s$  and arrive at airport  $Z$  during time period  $i$  which are rescheduled to arrive during period  $j$ , under capacity scenario  $q$  ( $k=1,\dots,K$ ;  $q, s = 1, \dots,Q$ ;  $i = t_s+1, t_s+2,\dots, T$ ;  $i \leq j \leq T+1$ ).

After substitution of the decision variables  $X_{qksij}$  and of the demands  $N_{ksi}$  into the objective function and corresponding network model (see Model 2a above), Model 3 is entirely analogous to Model 2b, except that the coupling constraints are as shown in Figure 3. Note that there is one set of coupling constraints for each stage  $s$  of the problem. This reflects the fact that, at the beginning of each stage, we must assign ground-holds without knowing which of the possible capacity forecasts will actually materialize. This reasoning is again similar to that for Model 2b. The reader is referred to Richetta (1991) and to Richetta and Odoni (1994) for a detailed description of Model 3.



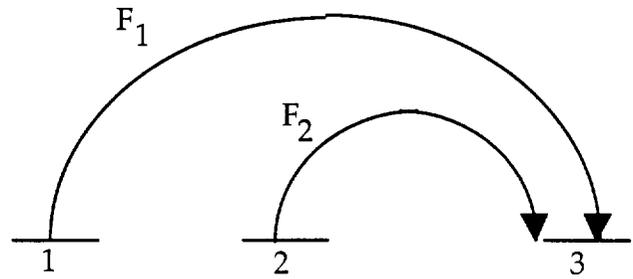
**Figure 3**  
**Constraint Matrix Structure for the Multistage Problem**

The resulting optimization problem may be solved by using standard techniques of multistage Stochastic Programming, with the stages corresponding to the time instants,  $t_s$ , when new or updated information may become available.

Finally, we note that additional constraints, such as placing limits on the maximum acceptable ground-holds and/or airborne delays, can be introduced into Models E and F easily.

A Simple Example

The difference between dynamic and static strategies is illustrated by the following idealized example, which involves only two flights. Figure 4 shows a diagram of the flight "schedule". Flight  $F_1$  is scheduled to depart at time 1,  $F_2$  is scheduled to depart at time 2, and both flights are scheduled to arrive at an airport Z at period 3.



**Figure 4**  
**Flight Schedule**

Landing capacity,  $M$ , at  $Z$  during the arrival period, time 3, is limited to one or two flights according to the probability tree shown in Figure 5. We assume that capacity during the next time period, 4, is unrestricted. The probability tree of Figure 5 shows the possible evolution of capacity at  $Z$  over time, as perceived at time 1. As happens in practice, airport capacity during time period 3 is partially correlated to capacity available during time 2. If time 2 capacity is 2, then there is a greater chance of having a high capacity during time 3; while if time 2 capacity is 1, there is a greater probability of having limited capacity during time 3.

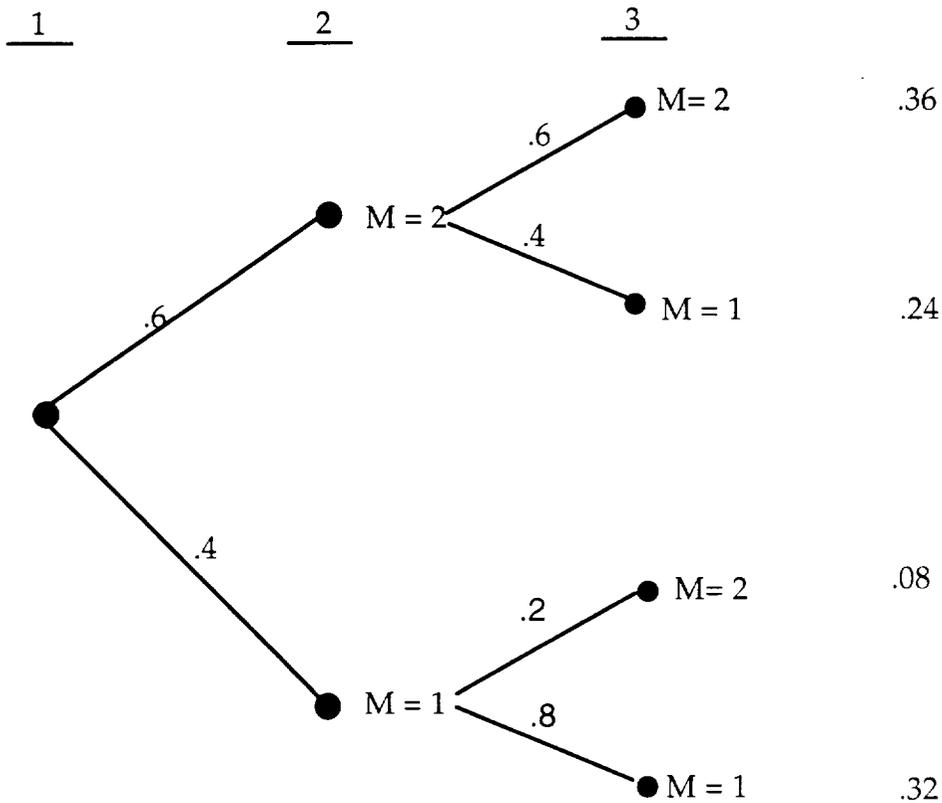
Next we specify the ground and air delay costs for  $F_1$  and  $F_2$ . Since  $F_1$  and  $F_2$  are both scheduled to arrive during time 3, and time 4 capacity is unrestricted, we only need to consider the cost of one period of delay:

| <u>Flight</u> | <u>Ground Delay Cost</u> | <u>Air Delay Cost</u> |
|---------------|--------------------------|-----------------------|
| $F_1$         | $c_{g1} = \$1000$        | $c_{a1} = \$2000$     |
| $F_2$         | $c_{g2} = \$1100$        | $c_{a2} = \$2200$     |

In line with what we would expect in a real situation, the cost of air delay is higher than that of ground delay reflecting the higher operational cost of airborne aircraft. The two aircraft have different costs reflecting factors such as aircraft type, passenger load, fuel efficiency, connection schedules, etc.

Time Period:

Probability



**Figure 5**  
**PMF of Airport Landing Capacities**

The static solution assigns ground-holds to both flights,  $F_1$  and  $F_2$ , at the beginning of time 1 based only on the information available then. Since the probability of having capacity limited to one landing during time 3 is .56, the optimal static strategy is to let the more costly flight,  $F_2$ , depart according to schedule and delay  $F_1$  one time period for an optimal static cost of \$1000 (if flight  $F_1$  were also allowed to depart on time, the expected cost would be  $(.56)(2000) + (.44)(0) = \$1,120$ ).

In the dynamic problem we make ground-hold decisions on a period by period basis, using the history of airport capacities to produce an updated capacity forecast regarding future airport landing capacities. Consider the following dynamic strategy:

Let  $F_1$  depart at time 1 (i.e., according to schedule). At time 2, if time 2 capacity is 1, delay the departure of  $F_2$  one time period; otherwise, let  $F_2$  depart according to schedule. By conditioning on the value of airport capacity at time 2 we see that the expected delay cost of this strategy is  $(.4)(1,100) + (.6)(.4)(2,000) = \$920$ , representing a significant cost improvement vs. the optimal static strategy. The reader can verify that this is indeed the optimal dynamic strategy. Note that if both aircraft are allowed to depart on time under this strategy (the probability of this event is .6), and the capacity at time 3 turns out to be equal to 1, then the less expensive flight,  $F_1$ , is the one which is held in the air for one time period.

The example also points to what is indeed a systematic bias in dynamic decision-making for the GHP: optimal strategies favor long-range flights over short-range ones, in the sense that long-range flights are more likely to be allowed to take-off on time (i.e., with no or little ground-holding). Intuitively, good dynamic strategies would tend to be more "active" with short-range flights (i.e., impose more ground-holds on them) in order to take advantage of the improved state of knowledge at the time when short-range flights are scheduled to depart. Current practice partly reflects this tendency: for example, flights to the United States from Europe or non-stop coast-to-coast flights in the United States are typically exempt from ground-holding.

#### **4.6 Computational Experience**

A large number of computational experiments have been carried out with Models 2 and 3 using data from Boston's Logan International Airport, which is the 21st busiest airport in the world in terms of number of passengers served. Detailed information regarding all aspects of these experiments can be found in Richetta (1991) and in Richetta and Odoni (1993, 1994). The experiments suggest, if properly adapted, Models 2 and 3 may be able to indicate ground-holding strategies that could result in significant reductions of delay costs. These same experiments were also encouraging with regard to the feasibility of employing the models as real-time decision support tools for ATC flow managers. However, no real-time experiments that may include professional air traffic controllers in a simulated environment have been carried out to date. Such an effort would undoubtedly spur further research on (i) the desirable type of "interfacing" between controllers and the decision support tools and (ii)

variations on the models described here that might provide additional flexibility and options to controllers.

## **5. The Generalized Tactical Traffic Flow Management Problem (GTFMP)**

Much work has been performed on model development for the generalized tactical TFM problem (GTFMP). Several exact optimization algorithms have been proposed. Moreover, the well-known Computer Assisted Slot Allocation (CASA) algorithm that EUROCONTROL uses to support its tactical level planning is essentially a heuristic algorithm for solving the GTFMP.

One can identify five distinct types of existing models for the GTFMP, the first three of which attempt to solve the problem optimally, while the other two, including CASA, aim at an approximate solution. A common characteristic of all five types of models is that they view the ATM infrastructure (airports, terminal areas, navigation fixes, en route sectors, airways) as a multiple origin-destination network on which traffic flows have to be assigned. In all cases these traffic flows can vary over time. Another common characteristic is that, in all five approaches, the models are deterministic, i.e., it is assumed that available capacity as well as flight demand is known in advance. In four of the five cases, the models deal with uncertainty by essentially assuming that the GTFMP will be solved anew every time conditions change, i.e., whenever capacity and/or demand change sufficiently to warrant such re-solving of the GTFMP (e.g., every few hours). However, the CASA heuristic recognizes uncertainty more explicitly by setting aside a number of flight slots for last minute exigencies.

The first optimal approach -- and the most successful computationally to date -- is the one due to Bertsimas and Stock (1994). They have formulated the GTFMP as a 0-1 Integer Programming Problem with six sets of constraints. The objective is to minimize the total cost of delaying aircraft on the ground (through ground holding) and in the air (through airborne holding). The constraints ensure that the traffic flows recommended by the model will not exceed available capacities and will satisfy certain "connectivity" relationships, so that the recommended solutions are meaningful physically. Specifically, the first three sets of constraints ensure that the traffic flows, in any discrete time interval during the period of interest, will not exceed the departure capacity of any airport in the

network, the arrival capacity of any airport and the sector capacity of any sector, respectively. The other three sets of constraints represent the three types of connectivity in the problem: connectivity between sectors, between airports and in time. All these relationships lead to a quite complex optimization model. As noted above, it is believed that this model has the best performance among existing models at this time.

The second major existing GTFMP model is the so-called Time Assignment Model or TAM [Lindsay, Boyd and Burlingame (1993), Boyd, Burlingame and Lindsay (1994), Burlingame, Boud and Lindsay (1995)]. This is also a 0-1 IP optimization model that consists of an objective function and five sets of constraints. The objective function, as in the case of the Bertsimas and Stock model, is to minimize the total cost of delaying flights. The first set of constraints specifies that capacities cannot be exceeded; the second defines a lower bound for the flight time from one node ("fix") of the underlying network to the next; the third specifies that each flight can pass over any particular fix only once; the fourth gives the earliest time interval during which a flight may depart; and the last set of constraints specifies a minimum "turn-around" time on the ground for each aircraft between flight legs. (It should be noted that the Bertsimas and Stock (1994) model also defines, implicitly, some of the same restrictions, through its connectivity constraints.)

The third optimization model is known as STN, Space-Time Network [Helme (1992, 1994)]. This is a model formulated as a multicommodity minimum-cost flow on a network. It deals with aggregate flows, not individual flights, attempts to minimize total delay costs and includes departure, en route and arrival capacity constraints. While the formulation of this model is straightforward and easy to understand, its computational performance has been quite disappointing and efforts toward its further development have apparently been abandoned.

The first of the two heuristic approaches, the Multiple Airport Scheduler (MAS) is a hybrid of optimization routines and heuristics [Epstein, Futer and Medvedovsky (1992)]. Information about this model, which was developed in connection with the Advanced Traffic Management System of the FAA, is limited, because little has been published about it. The goal of the MAS is

apparently to maximize flow in the ATM network subject to the usual capacity constraints (departure, en route, arrival), flight connectivity constraints and "fairness-to-user" constraints. A decomposition approach (the model is eventually decomposed into a number of single-airport problems) is used to speed up the solution. While the model apparently runs quickly, little is known about the quality of the solutions it produces or its actual computational performance.

Finally, the CASA heuristic has been described very clearly in Philipp and Gainche (1994) [see also EUROCONTROL (1993a, 1993b)] and a detailed description will not be repeated here. The heuristic gives priority to flights on a "First Planned, First Served" (FPFS) basis, meaning that the flights with the earliest departure times are considered first. CASA thus considers departing flights sequentially and, when necessary, assigns to each a ground-holding delay consistent with the most restrictive capacity constraint that the flight will encounter between its origin and its destination. As already noted, and in the interest of fairness, CASA also reserves a portion of the available capacity for short-haul flights and/or for flights that may, for some reason, file a flight plan shortly before their intended departure time. The rationale here is that, in the absence of such practice, all available slots might be consumed by long-haul flights, that file flight plans early and have early departure times. CASA also automatically updates its "solutions" every few minutes, in the hope that, as conditions change, the algorithm can discover ever-improving slot allocations. CASA is undoubtedly a very efficient tool computationally, typically returning solutions within about half a minute for problems that would be considered large for an optimization model. What is not yet known about CASA is whether or not the TFM slot allocations/ground-holding assignments it returns are of good quality. Some very preliminary evidence (see Section 6.1) may suggest that, in some cases, there may be considerable room for improvement in this respect.

### **5.1 Computational Results for the GTFMP**

We next discuss some computational results for the best of the existing optimization models, the Bertsimas and Stock (1994) model, which has been extensively tested using real data from both the US and the European networks.

As an example, two data sets of realistic size, obtained directly from the Official Airline Guide (OAG) flight schedules, were provided by the FAA. The first consisted of 278 flights, 10 airports and 178 sectors, tested over a 7 hour time frame with 5 minute intervals. The second of these data sets consisted of 1002 flights, 18 airports, and 305 sectors tested over an 8 hour time frame with 5 minute intervals. The sector crossing times, sector and airport capacities, and required turnaround times were all provided by the FAA. These data sets are comparable to those in the problems being solved daily by the FAA.

For the first problem, consisting of 43226 constraints and 18733 variables, an optimal solution of the linear programming relaxation was found in approximately 30 minutes on a SUN SPARC 20 workstation using CPLEX 3.0 as the optimization solver and GAMS 2.25 as the modeling language. Furthermore, the solution obtained was completely integral. In other words, there is no need to use any integer programming methods. The second and larger data set consisting of 151662 constraints and 69497 variables, was solved to optimality in approximately 2 hours, again achieving completely integral solutions.

Similar results, with essentially the same model, were obtained by Vranas (1995) for the European network. For a data set provided by EUROCONTROL, a problem involving 2293 flights and 25 sectors and with all costs equal to one (i.e., the objective is to minimize the total delay), an optimal, completely integral solution was found in approximately one hour in a SUN 10 workstation. The total delay in the optimal solution was roughly 40% lower than the delay assigned by CASA with its FPFS algorithm. This illustrates the significant impact that a linear optimization to the problem might have in practice, in certain cases. On the other hand, this evidence is very preliminary. Out of a total of five cases, tried by Vranas (1995) the optimization algorithm's solution had about 40% less delay than the CASA solution in 3 cases, but only about 5% savings, on average in the other two. It is also not clear how carefully, the solutions of the optimization algorithm were checked, at the individual flight level, for practical feasibility. Finally, it should be noted that CASA obtained in a matter of about 30 seconds a solution whereas, as noted, computation times were of the order of 30 minutes to 2 hours for the variant of the Bertsimas and Stock model used by Vranas.

Another observation that is very important for achieving short computational times is that in the absence of capacity constraints, the remaining inequalities define a network flow problem, for which we know that an optimal integral solution exists and can be found by the simplex method. As a result, the model is first ran as a network flow problem by ignoring the capacity constraints and a basis is found. Then, the capacity constraints are introduced and the model is solved as a linear program using the dual simplex method.

While the formulation is not always integral, it was integral for these and other test cases. A partial explanation for this is that the three sets of constraints that express the three types of connectivity in the problem are facets of the convex hull of the set of feasible solutions.

An alternative approach to the one we have presented is due to Lindsay et. al. (1993). They propose integer programming formulations for a version of TFMP that tracks a flight as it passes from fix to fix in the airspace. As the linear programming relaxations of these formulations are not very strong, branch and bound is needed to generate integral solutions. However, by developing a wide array of novel formulation-strengthening techniques, the dependence on "pure" branch and bound, as well as the computation times, are actually reduced.

## **6. The Multi-Airport, Ground Holding Problem (MAGHP)**

As noted earlier, models for the MAGHP differ from those for the SAGHP in that they consider many, rather than a single, airports simultaneously. However, the best available models in this area are not extensions or derivatives of models for the SAGHP. For example, the important research on the stochastic SAGHP, described in Section 4, has not been extended to the MAGHP. In fact, the only research on optimization models for the MAGHP which has partly addressed some stochastic issues is that of Vranas (1992), but has done so only in an implicit way. Thus, all models for the MAGHP are deterministic ones.

In general, existing optimization models for the MAGHP are related in approach and in certain cases are direct extensions of models for the generalized tactical FMP described in Section 5. The common characteristic of these models is that they assume that capacity constraints exist only at arrival airports and

departure airports, not in en route sectors. This assumption is quite true in the United States where more than 95% of serious delays can be attributed to capacity constraints at either the arrival airport or the departure airport, but is certainly not true in Europe, where the en route sectors currently place as much of a capacity constraint as the terminal areas. It should be noted, however, that the European situation is partly due to the fact that most major airports in Europe are operating with strict limits on how many operations can be scheduled there ("slots"). This means that demand at these airports is restricted *in advance*, i.e., when the airline schedules are made, to be approximately in balance with airport capacity under all weather conditions. In contrast, the absence of any limits on scheduling of operations at airports in the United States (only four airports have any such limits and even those may be rescinded soon) has resulted in a situation where demand far exceeds capacity at many airports whenever weather conditions are less than good.

For the time being, therefore, models for the MAGHP may be more directly applicable to the United States environment than to the European one. On the other hand, it would seem that some of the heuristic approaches that have been devised for the MAGHP could be extended to the case where the en route sectors also pose capacity restrictions and thus could be made more applicable to the European environment.

The optimization models that have been proposed for the MAGHP are essentially four, although a number of variations exist. They are described in: (1) Vranas (1992) and Vranas, Bertsimas and Odoni (1994a, 1994b); (2) Andreatta and Brunetta (1995); (3) Bertsimas and Stock (1994); and (4) Terrab and Paulose (1993). All four models have similar fundamental rationale: they all attempt to minimize the total cost of ground-holding delays and airborne delays, subject to capacity constraints on arrival and on departure and to connectivity constraints that ensure that solutions will be meaningful. The differences among these models lie primarily in the way the decision variables to be optimized are defined, and can thus be characterized as mostly technical, rather than substantive.

Andreatta and Brunetta (1995) have performed a careful comparison of the computational performance of the first three of the models mentioned above

and have concluded that the Bertsimas and Stock (1994) model performs best in most cases. The model is essentially a special case of the Bertsimas and Stock model for the TFMP which is described in the beginning of Section 6. The MAGHP model does not include any constraints related to the en route sectors and, primarily for this reason, can be solved much more efficiently than the model of the TFMP model.

The MAGHP can also be extended in an important way by including constraints that capture the interdependence between arrival capacity and departure capacity that exists at many airports (see Bertsimas and Stock (1994) for details). With this extension the model can be used not only to assign ground-holding and airborne delays to aircraft, but also to determine the optimal allocation of capacity between arrivals and departures at an airport during any given time period.

Two heuristics based on "priority rules" have also been proposed recently for the MAGHP by Andreatta, Brunetta and Guastalla (1994) and by Navazio and Romanin-Jacur (1995). These priority rules attempt to capture simple ways to assign ground-holding delays to sequences of aircraft and thus to develop locally optimal ground-holding strategies. The heuristics have performed very well in a limited number of tests to date, giving very good quality solutions (compared to the optimal) quickly. They are currently under active investigation and may offer a good alternative to exact optimization models in the future.

### **6.1 Computational Results for the MAGHP**

To illustrate the computational performance of MAGHP models, we review in this section computational results reported in Bertsimas and Stock (1993) and Vranas et. al. (1994a). (The reader is also referred to Andreatta and Brunetta (1995) for many additional test cases.) The datasets consist of 2 and 6 airports with 500 flights per airport, 1000 and 3000 flights respectively. Four levels of flight connectivity were considered. These levels give the ratios of continued flight to total flights,  $|C|/|F|$ , as 0.20, 0.40, 0.60, and 0.80 (see Appendices A and B for further explanation.) A 16 hour day was divided into 15 minute time intervals. All experiments were performed on a Sun SPARCstation 10 model 41. GAMS was used as the modeling tool and CPLEXMIP 2.1 was used

as the solver. The results obtained using the above datasets and the (MAGHP) formulation are summarized in Tables 1 and 2.

| F    | C / F | Dep Capacity | Arr Capacity | Time | % Nonint |
|------|-------|--------------|--------------|------|----------|
| 1000 | 0.20  | 32           | 15           | 262  | 0        |
| 1000 | 0.40  | 17           | 10           | 741  | 4        |
| 1000 | 0.60  | 20           | 14           | 359  | 0        |
| 1000 | 0.80  | 20           | 20           | 283  | 0        |

**Table 1. Results at the infeasibility border for 3000 Flights**

| F    | C / F | Dep Capacity | Arr Capacity | Time | % Nonint |
|------|-------|--------------|--------------|------|----------|
| 3000 | 0.20  | 20           | 20           | 5475 | 0        |
| 3000 | 0.40  | 20           | 20           | 4703 | 0        |
| 3000 | 0.60  | 20           | 20           | 5407 | 0        |
| 3000 | 0.80  | 20           | 20           | 9411 | 0        |

**Table 2: Results at the infeasibility border for 3000 Flights**

Tables 1 and 2 give results at the infeasibility border for each case. The infeasibility border is the set of critical values for the departure and arrival capacities, in units of flights per time interval, under which the problem becomes infeasible. We expect that it is in this region that the problem is very relevant practically and is hardest to solve. The critical capacity levels were found by a series of trial and error tests. The times reported are in CPU seconds and the "% Nonint" column is the percentage of total flights whose solution was noninteger.

## 7. Conclusions

In this section we list the principal conclusions we have reached as a result of the review presented in this document.

1. The Single-Airport Ground-Holding Problem (SAGHP) has been investigated in depth and provides a fine paradigm of optimization models that

include consideration of uncertainty and of the dynamic nature of TFM decision-making. However, the existing SAGHP models generally assume that the en route airspace does not impose any major constraints on the flow of air traffic and that the real capacity constraints exist at airports and terminal areas. This assumption, while quite valid in the United States, is not realistic for the European ATM system today.

2. The most advanced existing applications of optimization modeling to TFM deal with the Generalized Tactical Traffic Flow Management Problem (TFMP). Optimization models in this area hold serious promise of providing a viable alternative to heuristic algorithms, such as CASA, potentially leading to significant improvements in the performance of the TFM system. The computational requirements for obtaining optimal solutions and the difficulty of gaining TFM manager acceptance of such approaches would appear to be the main obstacles to adopting such optimization models at this time. It is our recommendation that further research in this area be strongly encouraged and pursued.

3. Multi-Airport Ground-Holding Problems (MAGHP) can be solved quite efficiently through existing optimization models. Most of these models can be viewed as essentially modified (and simplified) versions of models for the TFMP.

4. Regarding the issue of computational performance and feasibility the principal conclusions regarding existing solution approaches to the large-scale optimization problems, TFMP and MAGHP, are:

- a. In all but one instance of the MAGHP test cases that have been attempted with the Bertsimas and Stock (1994) models and in all instances of TFMP the relaxations of MAGHP and TFMP yielded integral solutions.
- b. The integrality of solutions was not affected by problem parameters, nor the size of the problem.
- c. The computational time required to obtain an optimal solution increases with the degree of connectivity as well as with the size of the problem.
- d. These models represent the strongest formulations proposed to date for this class of problems. Combined with state-of-the-art optimization libraries,

they allow the solution of truly large, realistic problems in reasonable computational times.

## REFERENCES

- G. Andreatta and L. Brunetta (1995), *Multi-Airport Ground Holding Problem: A Computational Evaluation of Exact Algorithms*, Technical Report, Department of Pure and Applied Mathematics, University of Padova, Padova, Italy.
- G. Andreatta, L. Brunetta and G. Guastalla (1994), *Multi-Airport Ground Holding Problem: A Heuristic Approach Based on Priority Rules*, Technical Report, Department of Pure and Applied Mathematics, University of Padova, Padova, Italy.
- G. Andreatta, A. R. Odoni and O. Richetta (1993), "Models for the Ground-Holding Problem", in *Computation and Information Processing in Air Traffic Control*, L. Bianco and A. R. Odoni, editors, Springer-Verlag, Berlin, 1993, pp. 125-168.
- G. Andreatta and G. Romanin-Jacur (1986a), *The Flow Management Problem*, LADSEB-CNR, Report 03/86, University of Padova, Padova, Italy.
- G. Andreatta and G. Romanin-Jacur (1986b) *Flow Management Problem and Optimal Priority Rule*, LADSEB-CNR, Report 04/86, University of Padova, Padova, Italy.
- G. Andreatta and G. Romanin-Jacur (1987), "Aircraft Flow Management under Congestion", *Transportation Science*, 21, No. 4, pp. 249-253.
- G. Andreatta and G. Romanin-Jacur (1988) "Aircraft Flow Management Problem", in *Operational Research '87* (G.K. Rand, editor) Elsevier, pp. 626-634.
- D. Bertsimas and S. Stock (1994), *The<sup>1</sup>air traffic flow management problem with enroute capacities*, Operations Research Center Working Paper # 3726-94, Massachusetts Institute of Technology.
- G. R. Booth (1994), "Flow Management in the United States", in *Advanced Technologies for Air Traffic Flow Management*, H. Winter and H.G. Nüsser, editors, Springer-Verlag, Berlin, pp. 107-120.
- E. A. Boyd, R. Burlingame and K. Lindsay (1994), *An Algorithm for the Management of National Air Traffic*, Report INEN/OR/WP/10/5-94, Texas A&M University, College Station, Texas, USA.

R. Burlingame, E. A. Boyd and K. Lindsay (1995), "Solving Large Integer Programs Arising from Air Traffic Flow Problems," in *Intelligent Scheduling Systems*, Donald E. Brown and W. T. Scherer (eds.), Kluwer Academic Publishers, Boston, MA, pp. 215-226.

J. S. De Armon and A. R. Lacher (1996), *Aggregate Flow Directives as a Ground Delay Strategy: Concept Analysis Using Discrete Event Simulation*, The MITRE Corporation, McLean, VA, USA.

D. Delahaye, J-M Alliot, M. Schoenauer and J-L Farges (1994), "Genetic Algorithms for Air Traffic Assignment", *Proceedings of 11th European Conference on Artificial Intelligence* (A. Cohn, editor), John Wiley & Sons, Ltd., pp. 473-476.

D. Duytschaever (1993), "The Development and Implementation of the Eurocontrol Central Air Traffic Flow Management Unit", *Journal of Navigation*, 46, No. 3, pp. 343-352.

L. Epstein, A. Futer and L. Medvedovsky (1992), *Multiple Airport Scheduling*, ATMS Interim Report, Volpe National Transportation Systems Center, Cambridge, MA, USA.

EUROCONTROL, Central Flow Management Unit (1993a), *Guide to Computer Assisted Slot Allocation (CASA)*, EUROCONTROL document TACT/URD/CASAGUID, Version 1.0.

EUROCONTROL, Central Flow Management Unit (1993b), *Slot Allocation Calculation*, EUROCONTROL document TACT/URD/CASAURD, Version 1.8.

E. Gilbo (1993), "Airport Capacity: Representation, Estimation, Optimization", *IEEE Transactions on Control Systems Technology*, 1, No. 3, September 1993, pp. 144-154.

G. D. Glockner (1995), *Effects of Air Traffic Congestion Delays Under Several Flow Management Policies*, Working Paper, Georgia Institute of Technology, School of Industrial and Systems Engineering, Atlanta, Georgia, USA.

M. P. Helme (1992), "Reducing Air Traffic Delay in a Space-Time Network", *Proceedings of the 1992 IEEE International Conference on Systems, Man and Cybernetics*, Chicago, pp. 236-242.

M. P. Helme (1994), *The Potential of Network Flow Models for Managing Air Traffic*, Report MTR 94W0000058, The MITRE Corporation, McLean, VA, USA.

K. S. Lindsay, E. A. Boyd and R. Burlingame (1993), "Traffic Flow Management Modeling with the Time Assignment Model", *Air Traffic Control Quarterly*, 1, no. 3, pp. 255-276.

- K. S. Lindsay (1993), *Traffic Flow Management Modeling with the Time Assignment Model*, Report MTR 93W0000119, The MITRE Corporation, McLean, VA, USA.
- P. Matos and R. Ormerod (1995), *The application of operational research to European air traffic flow management - understanding the context*, Technical Report, Warwick Business School Business Bureau, No 200, Coventry, UK.
- J. M. Mulvey and S. A. Zenios (1987), "Real-time operational Planning for the U. S. Air Traffic System", *Applied Numerical Mathematics*, 3, pp. 427-441.
- L. Navazio and G. Romanin-Jacur (1995), *Multiple Connections Multi-Airport Ground Holding Problem: Models and Algorithms*, Working Paper, Department of Electronics and Informatics, University of Padova, Padova, Italy.
- A. R. Odoni (1987), "The Flow Management Problem in Air Traffic Control" in *Flow Control of Congested Networks*, A. R. Odoni, L. Bianco and G. Szego (editors) Springer-Verlag, Berlin, pp. 269-288.
- A. R. Odoni (1994), "Issues in Air Traffic Flow Management", in *Advanced Technologies for Air Traffic Flow Management*, H. Winter and H.G. Nüsser, editors, Springer-Verlag, Berlin, pp. 43-63.
- W. Philipp and F. Gainche (1994), "Air Traffic Flow Management in Europe", in *Advanced Technologies for Air Traffic Flow Management*, H. Winter and H.-G. Nüsser, editors, Springer-Verlag, Berlin, pp. 64-106.
- O. Richetta and A. R. Odoni (1993), "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control", *Transportation Science*, 27, pp. 228-238.
- O. Richetta and A. R. Odoni (1994), "Dynamic Solution to the Ground-Holding Policy Problem in Air Traffic Control", *Transportation Research A*, 28, pp. 169-185.
- O. Richetta (1991), *Ground Holding Strategies for Air Traffic Control Under Uncertainty*, Operations Research Center Technical Report No. 198, MIT, Cambridge, MA.
- M. Terrab (1990), *Ground Holding Strategies in Air Traffic Control*, Operations Research Center Technical Report No. 196, MIT, Cambridge, MA.
- M. Terrab and A. R. Odoni (1993), "Strategic Flow Control on an Air Traffic Network", *Operations Research*, 41, pp. 138-152.
- M. Terrab and S. Paulose (1993), *Dynamic Strategic and Tactical Air Traffic Flow Control*, RPI Technical Report.

- V. Tasic, and O. Babic (1995), "Air Route Flow Management - Problems and Research Efforts", *Transportation Planning and Technology*, 19, pp. 63-72.
- P. B. Vranas (1992), *The Multi-Airport Ground Holding Problem in Air Traffic Control*, PhD Thesis, Operations Research Center, MIT, Cambridge, MA.
- P. B. Vranas, D. Bertsimas and A. R. Odoni (1994a), "The Multi-Airport Ground-Holding Problem in Air Traffic Control", *Operations Research*, 42, No. 2, pp. 249-261.
- P. B. Vranas, D. Bertsimas and A. R. Odoni (1994b), "Dynamic Ground-Holding Policies for a Network of Airports", *Transportation Science*, 28, No. 4, pp. 275-291.
- P. Vranas (1995), *Optimal Slot Allocation for European Air Traffic Flow Management*, Working Paper, German Aerospace Research Establishment (DLR), Braunschweig, Germany.
- M. Wald , "Austere Future Looms for F.A.A.", *New York Times*, October 15, 1995, p. A1.
- M. C. Wambsganss (1996), *Collaborative Traffic Flow Management*, Working Paper, Metron Inc., Washington, DC, USA.
- S. Zenios (1991), "Network Based Models for Air Traffic Control", *European Journal of Operational Research*, 50, pp. 166-178.

