Optimal Slot Allocation for European Air Traffic Flow Management

Peter B. M. Vranas

Responding to the increasing congestion of the European airspace, the Central Flow Management Unit (CFMU) of the European Organization for the Safety of Air Navigation (EUROCONTROL) started applying, on 28 April 1995, a centralized slot allocation process by means of a Computer Assisted Slot Allocation (CASA) procedure. Two basic objectives of slot allocation are to ensure the safety of crowded airspace and to minimize the total aircraft delay by optimizing the use of the available capacity. CASA, however, is a heuristic algorithm (based on the "First Planned, First Served" principle) rather than an optimization procedure. The present work continues and builds upon a series of recent attempts by various researchers to examine optimization models for slot allocation. The work consists of two somewhat disparate parts, a theoretical and a computational one. First, a new integer programming model of the European Slot Allocation Problem (SAP) is presented. The model is specifically adapted to the current needs of the CFMU, but some extensions that anticipate possible future needs of the CFMU are also presented. Second, based on a modification of a previous model from the literature, a computational comparison of optimal slot allocation with current slot allocation practice is effected. The computations were performed with real data provided by EUROCONTROL and indicate that, by applying optimal slot allocation, the total delays imposed by CASA can be reduced in some (but not all) cases by about a third, while the maximum delays imposed by CASA can sometimes also be reduced. The comparisons with CASA need to be qualified in several ways; nevertheless, given that the annual costs of delays are estimated to be of the order of billions of U.S. dollars, the computational results suggest that applying optimal slot allocation might result in significant savings.

Peter Vranas was formerly with the German Aerospace Research Establishment.

Received July 18, 1996; accepted August 20, 1996.

Air Traffic Control Quarterly, Vol. 4(4) 249–280 (1997)

CCC 1064-3818/95/030163-20

INTRODUCTION

The Congestion Problem

The annual number of flights in Western Europe[1] has increased from about 2.6 million in 1982 to about 4.5 million in 1992, an increase of about 73 percent (EUROCONTROL, 1994, Ed. 7.0, p. 7). Acute congestion of the air traffic control system has been the result. A similar problem exists in the United States, where each of 33 major airports is expected to exceed 20,000 hours of annual delays by 1997 (Transportation Research Board, 1991, p. 215).

The ground and airborne delays caused by congestion create direct costs to the airlines and indirect (opportunity) costs to the passengers. Direct costs from ground delays include crew, maintenance, and depreciation costs, while direct costs from airborne delays include, in addition, fuel and safety costs (Vranas, 1992, p. 12). Although estimates of congestion costs are not particularly reliable, there seems to be agreement that they amount to billions of U.S. dollars.[2] Given that several European and U.S. airlines are suffering yearly losses that also amount to billions of dollars,[3] the significance of the congestion problem can hardly be denied.

Possible solutions to the congestion problem include the following (Odoni, 1987; Transportation Research Board, 1991; Vranas, 1992): (1) long-term approaches, such as the construction of new airports or of new runways at existing airports; (2) medium-term approaches, such as congestion pricing and use of larger aircraft; and (3) short-term approaches, especially slot allocation policies, the object of the present work.

Reasons for Slot Allocation

A slot allocation policy assigns a time slot for departure to each of a number of flights. If the slot falls later than the scheduled departure time, then allocating the slot is equivalent to imposing a forced ground delay; hence, a slot allocation policy can also be called a ground-holding policy. "Slot allocation" is the preferred term in Eu-

[1] Specifically, in the following 11 European states: Austria, Belgium, France, Germany, Ireland, Luxembourg, Netherlands, Portugal, Spain, Switzeland, and United Kingdom.

[2] For example: (1) the Stanford Research Institute (a consulting company) estimated in 1990 annual costs of $5 billion in Europe due to congestion and ATC inefficiencies (Odoni, 1995; cf. Maugis, 1995); (2) the Federal Aviation Administration estimated in 1988 the annual cost of delays in the U.S. at $1.4 billion (Federal Aviation Administration, 1988); (3) the Air Transport Association of America (an association representing the interests of almost all major U.S. airlines) estimated in 1992 that its members had direct delay costs of $1.5 billion (Odoni, 1995).

[3] In 1995 the U.S. airline industry had a profit of about $2.5 billion, but during the previous 5 years it had accumulated losses of about $13 billion (Gray, 1996). Some airlines in European countries are in a similar plight, but their governments are paying for their losses.

rope, while "ground-holding" is the preferred term in the United States.

A first major objective of slot allocation policies is to *avoid overloads*. An overload occurs when the number of aircraft that are present at an airspace sector exceeds the capacity of the sector. Sector capacities are determined by the maximum numbers of aircraft that air traffic controllers can safely handle during a given time period. Given that overloads sometimes occur, air traffic controllers typically give conservative values for sector capacities; i.e., the number of aircraft that they say they can handle is typically smaller than the number of aircraft that they in fact can handle. If, however, a slot allocation policy is consistently implemented so that overloads almost never occur, it is expected that air traffic controllers will start giving the real values of sector capacities (Dalichampt, 1994). Thus, although the elimination of overloads is meant to guarantee safety, it may also reduce congestion costs by leading indirectly to an increase of capacities and hence to a reduction of delays.

A second major objective of slot allocation policies is to *reduce delay costs* by absorbing airborne delays on the ground. The basic idea is as follows. Suppose that, if an aircraft were to depart on time, it would encounter congestion upon its arrival at its destination airport and thus would have to wait in the air for, say, 30 minutes. Suppose also that, if the aircraft arrived at its destination airport 30 minutes later, it would encounter no congestion and would land immediately. Given that airborne delays are much costlier than ground delays (essentially because of fuel costs), the delay cost (although not the delay itself) can be reduced in the above scenario by holding the aircraft on the ground for 30 minutes.

The above scenario depends on the possibility of an airport being congested during a time period t_1, but not during a later time period t_2. There are two ways in which this can happen. A first way is if poor weather reigns during t_1, but normal weather returns during t_2. This problem is particularly acute in the United States, where instrument meteorological conditions (i.e., poor weather) require the use of instrument flight rules; these require greater spacing of aircraft than do visual flight rules (which apply in normal weather), so that they can reduce airport capacity by as much as 50 percent (Transportation Research Board, 1991, p. 221). In Europe, however, visual flight rules are seldom used (Gainche, 1996) (at least in theory), so that this problem does not appear. However, a second way in which an airport can be congested during t_1 but not during t_2 is possible even with constant capacities and *is* relevant for Europe: it is simply if t_1 is a "peak" traffic period, but t_2 is not.

On the basis of the above illustrative scenario, it may be thought that slot allocation policies cannot reduce the total delay, but can simply transfer airborne delays to the ground. This is not true, as shown by the following simple example. Suppose two flights, f_1 and f_2, are scheduled to depart from a given airport during time period t, while the departure capacity is just one flight during t and very high afterwards. Therefore, either f_1 or f_2 will have to be held on the ground for exactly one time period. Suppose also that (1) whether f_1 is delayed or not, it will encounter no congestion at its destination airport and so will land immediately; and (2) if f_2 departs on time, it will incur an airborne delay of one time period, while if f_2 arrives with a delay of one time period, it will land immediately. Then by holding f_2 rather than f_1 on the ground, the total delay is reduced from two time periods to one, i.e., by 50 percent (and the total delay cost is reduced by an even greater percentage, given that airborne delays are costlier than ground delays).

Current Slot Allocation Practices

The importance of slot allocation policies has been recognized, and steps towards implementing such policies have been taken both in the United States and in Europe. In the United States, the Federal Aviation Administration (FAA) operates the Air Traffic Control System Command Center (ATCSCC) in Washington, D.C. While impressive progress in data collection has been made (Odoni, 1994, pp. 47–48), the ATCSCC is still "plagued with inadequate data" (Booth, 1994, p. 111), particularly because airlines often change their schedules dramatically in bad weather (just when ground-holding is most needed) without informing the ATCSCC. Moreover, although some steps towards applying optimization models have been made, the ATCSCC still assigns ground holds according to heuristic procedures, and assignments affecting each destination are made independently of those affecting other destination airports (Booth, 1994, p. 110).

In Europe EUROCONTROL operates the Central Flow Management Unit (CFMU) in Brussels (Bardin et al., 1992; EUROCONTROL, 1994, Ed. 7.0; EUROCONTROL, 1994, Ed. 1.0; EUROCONTROL, 1993, Ver. 1.0; EUROCONTROL, 1993, Ver. 1.8; Odoni, 1994; Philipp et al., 1994). The CFMU provides air traffic flow management services at three levels: (1) *strategic* activities are executed from several months until 2 days before the day of operation of a flight; (2) *pretactical* activities take place during the 2 days before the day of operation of a flight; and (3) *tactical* activities are carried out on the day of operation of a flight and consist mainly of slot allocation by means of a computerized procedure based on the Computer Assisted Slot Allocation (CASA) algorithm.

The CFMU tactical slot allocation process became operational on 28 April 1995 (Tibichte, 1995). Initially it covered only the French airspace, but it was gradually extended and now covers the whole European airspace (Gainche, 1996). The process works as follows

(EUROCONTROL, 1994, Ed. 1.0). (1) An aircraft operator must file a flight plan at least 3 hours before the intended departure time of the flight. (2) The CFMU sends a Slot Allocation Message to the aircraft operator at least 2 hours before the originally intended departure time of the flight. (3) If the aircraft operator is not happy with the allocated slot, or, e.g., if the CFMU later becomes able to allocate an earlier slot, a new exchange of messages follows (EUROCONTROL, 1994, Ed. 2.0) and may result in the allocation of a new slot to the flight. The tactical activity of the CFMU ends when the flight becomes airborne.

The CASA algorithm (Bardin et al., 1992, pp. 8–11; see also Hughes et al., 1995, pp. 73–84) carries out allocations according to the "First Planned, First Served" rule; i.e., it gives priority not to flights for which flight plans are filed earlier, but rather to flights that, given their intended departure times, have earlier estimated entry times to regulated sectors. (A sector is "regulated" during a time period if its capacity is exceeded by the anticipated demand during that period. A flight is "regulated" if there is at least one sector that is regulated during a time period which includes the estimated entry time of the flight into the sector.[4]) In order to avoid excessive penalization of flights for which flight plans are filed shortly before the intended departure time, CASA reserves some of the available capacity for such "late filers." Moreover, CASA reserves a portion of the available capacity for shorthaul flights in order to avoid discrimination against them. (Such discrimination could arise because longhaul flights, which depart earlier and so are assigned slots earlier, could use up all available capacity by the time flight plans for shorthaul flights are filed.) Finally, CASA automatically reruns every few minutes, in the hope of improving some slot allocations in the light of new data.

The above summary overview of current slot allocation practices shows that, although considerable progress towards implementing such practices has been accomplished, no optimization algorithms are currently in use, so that current slot allocation cannot be described as "optimal." Given the severity of the congestion problem and the significant potential of slot allocation for alleviating this problem (Vranas, 1992), the interest of examining models for optimal slot allocation becomes apparent.

Previous Research on Optimal Slot Allocation

In order to categorize previous research efforts on optimal slot allocation, it is useful to distinguish various versions of the Slot Alloca-

[4] Complications arise for so-called "combined" flights, i.e., flights that go through several regulated sectors. These complications are important, because typically more than 30 percent of regulated flights are combined, and a combined flight typically goes through more than two regulated sectors (Gainche, 1996).

tion Problem (SAP). *Static* versions allocate slots once and for all at the beginning of the time horizon, while *dynamic* versions update slot allocation in the light of new data that become available as time proceeds. *Deterministic* versions consider airport (departure and arrival) and sector capacities as fixed numbers, while *probabilistic* versions consider capacities as random variables. Finally, *decomposed* versions allocate slots for each capacitated element (airport or sector) independently, while *network* versions allocate slots for all capacitated elements simultaneously in order to ensure that the allocations for the various elements are compatible with each other.

A first detailed description of the SAP was given by Odoni (1987). Not surprisingly, research on the SAP first focused on the simpler, though less realistic, decomposed versions of the problem. Terrab (1990) and Terrab et al. (1993) examined the static SAP for a single airport. Andreatta et al. (1987) examined the probabilistic single-airport SAP for a single time period. Richetta (1991) and Richetta et al. (1993, 1995) examined both static and dynamic versions of the probabilistic single-airport SAP. Andreatta et al. (1993) reviewed optimization models for the single-airport SAP.

Some years ago, researchers began attacking network versions of the SAP. Work focused first on versions with several capacitated airports, but no capacitated sector. Vranas (1992) and Vranas et al. (1994a) gave the first models of the static multi-airport SAP. An improvement of those models was suggested by Andreatta et al. (1995), whereas Andreatta et al. (1994) and Brunetta et al. (1995) developed quick and efficient heuristics. Terrab et al. (1995) addressed the probabilistic multi-airport SAP. Finally, Vranas (1992) and Vranas et al. (1994b) gave models for both deterministic and probabilistic versions of the dynamic multi-airport SAP.

The most recent step was the formulation of models for network versions of the SAP which include capacitated sectors in addition to capacitated airports. These models are of special interest for the European case. Lindsay et al. (1993) and Bertsimas et al. (1995) proposed integer programming models, while Helme (1992) presented a multicommodity minimum cost flow model. It seems that the models of Bertsimas et al. (1995) have the best performance in terms of computing time. This is due to the fact that their formulations are "strong," in the sense of including facets of the convex hull of the set of solutions, so that solving the linear programming relaxations of their models usually gives integer solutions. Bertsimas et al. (1995) also proved that the SAP is NP-hard and proposed an integer programming model for rerouting flights.[5]

[5] For more detailed surveys of previous research on optimal slot allocation, see Helme et al. (1992), Lindsay et al. (1993, pp. 256–259), and Tosic and Babic (1995). Küttner (1995) has an extensive bibliography.

Contribution of the Present Work

The present work makes two (only loosely related) contributions. First, a new integer programming model for the static deterministic network SAP is presented. The model is specifically adapted to the current needs of the CFMU and is intended as an alternative to CASA, which is a heuristic rather than an optimization procedure. Given the needs of the CFMU, the emphasis is on capacitated sectors, although an extension of the model incorporating capacitated airports is also presented. Second, and most important, computational results are reported, based not on the new model, but rather on a version of the Bertsimas et al. (1995) model adapted to the needs of the CFMU. The computations were performed with real data provided by EUROCONTROL. Comparisons with the results of CASA indicate that optimal slot allocation can reduce in some (but not all) cases the total delays imposed by CASA by about a third and can sometimes also reduce the maximum delays imposed by CASA. An interesting result is that optimal slot allocation typically delays many fewer flights than CASA does, but the mean delay *per delayed* flight is higher for optimal slot allocation than for CASA. Although the comparisons with CASA need to be qualified in certain ways (see Conclusions), these results can be considered important because they go beyond purely hypothetical assertions of the utility of optimal slot allocation models, and they constitute what seems to be the first concrete evidence to the effect that use of such models might result in a significant reduction of delay costs.

The remainder of the paper consists of two major parts: the next section presents the new integer programming model, whereas the third section reports the computational results based on a version of the Bertsimas et al. (1995) model. The final section evaluates the results and presents suggestions for future research.

AN INTEGER PROGRAMMING MODEL OF THE EUROPEAN SLOT ALLOCATION PROBLEM

This section presents a novel integer programming model of the European SAP. The model is specifically adapted to the current needs of the CFMU, although some extensions that anticipate possible future needs of the CFMU are also presented.

Notation

Consider a set of airports $\mathcal{K} = \{1, \ldots, K\}$, a set of airspace sectors $\mathcal{S} = \{1, \ldots, S\}$, and an ordered set of time periods $\mathcal{T} = \{1, \ldots, T\}$. For instance, $\mathcal{K}$ might be a set of 50 European airports, $\mathcal{S}$ might be a set of 25 sectors of the European airspace regulated by the CFMU during some part of the time horizon, and $\mathcal{T}$ might be a set of 96 time periods of 15 minutes each, amounting to a time horizon of 24 hours. Consider also a set of flights $\mathcal{F} = \{1, \ldots, F\}$. (These are understood as flight *legs*, so that a single aircraft may perform several of them in succession.) $\mathcal{F}$ is the set of all flights of interest, e.g., all flights passing through at least one congested sector. For a typical day, $\mathcal{F}$ may include several thousands of flights.[6]

For each time period $t \in \mathcal{T}$, we are given the departure capacity D_{kt} and the arrival capacity R_{kt} of each airport $k \in \mathcal{K}$, as well as the capacity P_{st} of each sector $s \in \mathcal{S}$. There is a conceptual difference between airport and sector capacities: D_{kt} and R_{kt} are the maximum numbers of aircraft that may *flow* out of or into airport k during time period t, while P_{st} is the maximum number of aircraft that may be *present* at sector s during time period t.[7]

For each flight $f \in \mathcal{F}$, the following data are also given: $k_f^d \in \mathcal{K}$, the airport from which f is scheduled to depart; $k_f^r \in \mathcal{K}$, the airport to which f is scheduled to arrive; $\mathcal{S}_f \subset \mathcal{S}$, the ordered set of sectors that f is scheduled to go through (the sectors are ordered in the sequence in which f is scheduled to go through them); $t_f^d \in \mathcal{T}$, the scheduled departure (take-off) time of f; $t_f^r \in \mathcal{T}$, the scheduled arrival (landing) time of f; $t_f^{s+} \in \mathcal{T}$, the scheduled entry time of f into each sector $s \in \mathcal{S}_f$; $t_f^{s-} \in \mathcal{T}$, the scheduled exit time of f out of each sector $s \in \mathcal{S}_f$; and c_f^g, the cost incurred for each time period that f is delayed on the ground.

Table 1 summarizes the above notation for reference purposes. Table 1 also includes some symbols that will be defined in what follows.

Decision Variables and Sector Capacity Constraints

There is only one decision that needs to be taken about each flight f, namely how many time periods g_f it will be held on the ground. Equivalently, the decision can be described as "allocating a departure slot to each flight." Theoretically, one could also take decisions concerning the speed of each flight *en route*, regulating the time at which a flight will arrive at each sector in its path. This possibility could be accommodated by defining decision variables corresponding to imposed delays (or speed-ups) for each sector in the path of each flight. Such an approach will not be followed here, because it is the policy of the CFMU to avoid imposing airborne delays. The reasoning behind this policy is that imposed airborne delays (or speed-ups) would

[6] Currently more than 20,000 flights cross the European airspace daily (Tibichte, 1995), but usually not more than half of these flights go through regulated sectors.

[7] Sector capacities may also be understood as maximum flows into sectors, and this is in fact how they are understood in CASA. This fact has important implications; see footnote 11.

Table 1. Basic Notation and Acronyms

Symbol	Denotation
Sets	
$\mathcal{K} = \{1, \ldots, K\}$	Set of airports k.
$\mathcal{S} = \{1, \ldots, S\}$	Set of sectors s.
$\mathcal{F} = \{1, \ldots, F\}$	Set of flights f.
$\mathcal{T} = \{1, \ldots, T\}$	Ordered set of time periods t.
Times	
$t_f^d \in \mathcal{T}$	Scheduled departure time of flight f.
$t_f^r \in \mathcal{T}$	Scheduled arrival time of flight f.
$t_f^{s+} \in \mathcal{T}$	Scheduled entry time of flight f into sector s.
$t_f^{s-} \in \mathcal{T}$	Scheduled exit time of flight f out of sector s.
Decision variables	
g_f	Delay decision variables for ground holds.
u_{ft}	Point assignment decision variables for departures.
w_{ft}	Interval assignment decision variables for departures.
d_{kt}^i	Capacity decision variables for departures.
r_{kt}^i	Capacity decision variables for arrivals.
$p_{stt'}^i$	Capacity decision variables for sectors.
Capacities	
D_{kt}	Departure capacity of airport k at period t.
R_{kt}	Arrival capacity of airport k at period t.
P_{st}	Capacity of sector s at period t.
Other symbols	
Δ_{kt}	Number of flights scheduled to depart from k at t.
Γ_{kt}	Number of flights scheduled to arrive at k at t.
$\Pi_{stt'}$	Number of flights scheduled to enter s at t and to exit at t'.
$\mathcal{S}_f \subset \mathcal{S}$	Ordered set of sectors that f is scheduled to go through.
$k_f^d \in \mathcal{K}$	Departure airport of flight f.
$k_f^r \in \mathcal{K}$	Arrival airport of flight f.
c_f^g	Ground delay cost of flight f per time period.
G	Upper bound on the ground delay of a flight.
L_s	Length of sector s.
$\mathcal{F}' \subset \mathcal{F}$	Set of continued flights f'.
$s_{f'}$	Slack of continued flight f'.

be inapplicable in practice, given the short flying times of most flights in the European airspace (Gainche, 1996). This reasoning does not hold for the U.S. airspace, for which a model complete with airborne delays might be appropriate. Such a model has been developed by Bertsimas et al. (1995).

It seems thus that the *delay decision variables* g_f are all one needs. Unfortunately, it is not possible to express the departure, arrival, and sector capacity constraints as linear functions of these decision variables. A previous approach to this problem (Vranas, 1992; Vranas et al., 1994; Vranas et al., 1994a, 1994b; see also Bertsimas et al.,

1995) has been to define *(point) assignment decision variables*: u_{ft} was defined to be 1 if flight f is allocated a departure slot at time period t (i.e., if $t_f^d + g_f = t$), and 0 otherwise. Introducing these assignment decision variables allows one to express the departure and arrival capacity constraints in a simple way, but has the unfortunate effect of greatly increasing the size of the model. In fact, denoting by G the maximum number of time periods that a flight may be delayed on the ground (e.g., $G = 10$ corrresponds to a typical maximum ground hold of 2.5 hours), the number of u_{ft} decision variables is about $(G + 1)F$.[8] If the number of flights is, e.g., $F = 10{,}000$, then there are 110,000 decision variables, a very large number for an *integer* programming model.

In order to keep the number of decision variables as small as possible so that the model may be solvable within a reasonably short amount of time, use of assignment decision variables will be avoided. Define instead *(sector) capacity decision variables* as follows: $p^i_{stt'}$ is the number of aircraft that are scheduled to enter sector s at time period t and to exit sector s at time period t', but are held on the ground for i time periods ($i \in \{0, \ldots, G\}$). These variables allow one to express the sector capacity constraints in the following way:

$$\sum_{n=0}^{L_s} \sum_{x=0}^{L_s - n} \sum_{i=0}^{G} p^i_{s,t-n-i,t+1+x-i} \leq P_{st}, \quad (s,t) \in \mathcal{S} \times \mathcal{T}, \tag{1}$$

where L_s is the length of sector s, i.e., the maximum number of time periods that an aircraft needs to go through s.

The reasoning behind (1) is as follows. The left-hand side of the inequality in (1) must be the number of aircraft that are present at sector s during time period t. The aircraft present at s at t are the aircraft that have entered s at a time $t^+ \leq t$ and that will exit s at a time $t^+ \geq t + 1$.[9] The possible entry times t^+ of these aircraft at s range from $t - L_s$ to t and, for each given t^+, the possible exit times t^- from s range from $t + 1$ to $t^+ + 1 + L_s$. Therefore, the sector capacity constraints become:

[8] It may not be *exactly* $(G + 1)F$ because the scheduled departure time of some flights may be so close to the end of the time horizon that fewer than $G + 1$ departure slots will be available for each of those flights. This insignificant complication will be ignored in what follows.

[9] Aircraft are assumed to enter or exit sectors only at time instants coinciding with the beginning of time periods. The real entry and exit times of flights to and from sectors are rounded off so that, if an aircraft is present at a sector during part of a time period, then it is present at the sector during the whole time period. Therefore, $p^i_{stt} = 0$; if $p^i_{stt'} > 0$, then $t' \geq t + 1$.

$$\sum_{t^+=t-L_s}^{t} \sum_{t^-=t+1}^{t^++1+L_s} p_{st^+t^-} \leq P_{st}, \quad (s,t) \in \mathcal{S} \times \mathcal{T}, \tag{2}$$

where $p_{st^+t^-}$ is the number of aircraft that have entered s at t^+ and will exit at t^-. $p^0_{st^+t^-}$ of these aircraft were scheduled to enter s at t^+ and to exit at t^- and took off on time; p^1_{s,t^+-1,t^--1} of these aircraft were scheduled to enter s at $t^+ - 1$ and to exit at $t^- - 1$ and took off after a ground hold of one time period; and so on up to p^G_{s,t^+-G,t^--G}. Therefore, $p_{st^+t^-} = \sum_{i=0}^{G} p^i_{s,t^+-i,t^--i}$. Substitution into (2) and change of variables gives (1).

The Basic Integer Programming Model

We are now in a position to give the basic integer programming model of the European SAP, a model specifically adapted to the current needs of the CFMU.

$$(P_1) \ \min \ \textstyle\sum_{f=1}^{F} c^g_f g_f$$

$$\text{s.t.} \quad \textstyle\sum_{n=0}^{L_s} \sum_{x=0}^{L_s-n} \sum_{i=0}^{G} p^i_{s,t-n-i,t+1+x-i} \leq P_{st}, \quad (s,t) \in \mathcal{S} \times \mathcal{T}; \tag{3}$$

$$\textstyle\sum_{i=0}^{G} p^i_{stt'} = \Pi_{stt'}, \quad (s,t) \in \mathcal{S} \times \mathcal{T}, t+1 \leq t' \leq t+1+L_s; \tag{4}$$

$$g_f \leq G, \quad f \in \mathcal{F}; \tag{5}$$

$$\text{Delay} - \text{capacity constraints}; \tag{6}$$

$$g_f, p^i_{stt'} \ \text{nonnegative integers}. \tag{7}$$

The objective function, to be minimized, is the total cost of delays in the airport network. Although in theory one could assign different delay costs per time period, c^g_f, to different flights f (e.g., larger aircraft could have higher delay costs), in practice this would violate the principle of equity among all users of the airport network, a principle of high importance for the CFMU (EUROCONTROL, 1994, Ed. 1.0, p. III-03; see also Bardin et al., 1992, p. 16). Therefore, all c^g_f are taken to be equal; this amounts to minimizing the sum of all delays rather than the total cost of delays.

Constraints (3) are the sector capacity constraints explained in detail earlier. Constraints (4) express the fact that, among the aircraft that are scheduled to enter sector s at t and to exit at t' (the number of these aircraft being denoted by $\Pi_{stt'}$), $p^0_{stt'}$ aircraft will not be delayed at all, $p^1_{stt'}$ will be delayed one time period, and so on up to $p^G_{stt'}$. It follows that, for each combination $\langle s,t,t' \rangle$ for which $\Pi_{stt'} \neq 0$, one of the decision variables $p^0_{stt'}, p^1_{stt'}, \ldots, p^G_{stt'}$ can be eliminated by means of the equality constraints (4).[10] Therefore, there are G decision variables left for each such combination <s,t,t'>. Denoting by $\mathcal{S}_t \subset \mathcal{S}$ the set of sectors that are regulated at time t, the total number of $p^i_{stt'}$ decision variables can be seen to be $G \sum_{t \in \mathcal{T}} \sum_{s \in \mathcal{S}t}(L_s + 1)(L_s + 2)/2$. For the typical values of about 350 sector-hours (1,400 sector-periods) regulated each day by the CFMU, $G = 10$, and $L_s = 0$,[11] the number of $p^i_{stt'}$ decision variables becomes 14,000. Adding to these the F decision variables g_f (e.g., $F = 10{,}000$), the total number of decision variables in model (P_1) becomes 24,000. What seems to be the best previously available model of the SAP (Bertsimas et al., 1995), suitably adapted to meet the current needs of the CFMU (see model (P_3) below), has $GF = 100{,}000$ decision variables. Therefore, model (P_1) achieves a reduction of more than 75 percent in the number of decision variables for the above typical values of the parameters.

Constraints (5) simply express the fact that the maximum allowable ground hold is G time periods. Constraints (6) indicate that model (P_1) is, as it stands, incomplete, because there is no guarantee that, in the optimal solution, the values of the delay decision variables g_f and those of the capacity decision variables $p^i_{stt'}$ will be compatible (as they should be, given that, in the real-world problem, determining g_f suffices to determine $p^i_{stt'}$). Spelling out constraints (6) is relegated to the appendix, because it requires a rather lengthy treatment that would complicate the exposition here.

Extensions of the Basic Model

This subsection presents two extensions of the basic model (P_1). These extensions anticipate possible future needs of the CFMU. The first extension is straightforward and shows how to incorporate into (P_1) constraints for departure and arrival capacities at airports. The second extension comes from previous work (Vranas, 1992; Vranas et al., 1994; Vranas et al., 1994a, 1994b) and shows how to take into account the propagation of delays that results from connections between flights.

Departure and Arrival Capacity Constraints. Define *departure capacity decision variables* as follows: d^i_{kt} is the number of flights that

[10] If, for instance, variables $p^0_{stt'}$ are eliminated, then the equality constraints (4) must be replaced by the following inequality constraints:

$$\sum_{i=1}^{G} p^i_{stt'} \leq \Pi_{stt'}. \tag{8}$$

[11] The value $L_s = 0$ corresponds to the current understanding by CASA of sector capacities as maximum flows into sectors (see footnote 7). Given this understanding of sector capacities, the sector capacity constraints could have been expressed like the departure and arrival capacity constraints (9) and (10) (see the subsection on Departure and Arrival Capacity Constraints), but the more general way (1) of expressing them was preferred.

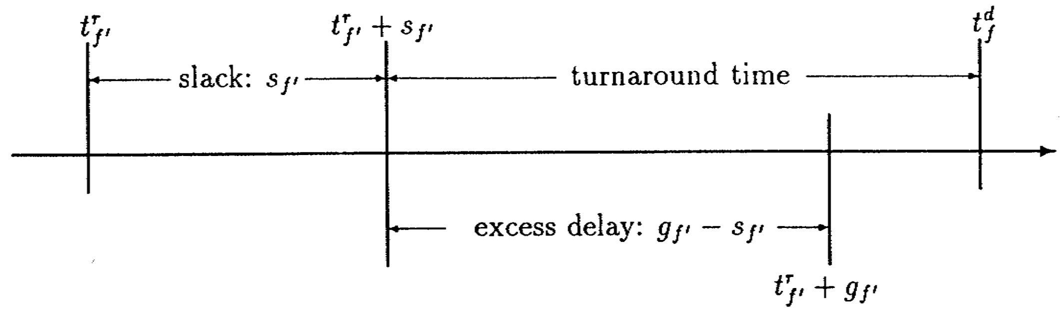

Figure 1. Modeling of connections between flights.

are scheduled to depart from airport k at time period t, but depart with a delay (ground hold) of i time periods ($i \in \{0, \ldots, G\}$). Similarly, define r^i_{kt} (*arrival capacity decision variables*) as the number of flights that are scheduled to arrive at airport k at time period t, but arrive with a delay (due to their ground hold) of i time periods. The departure and the arrival capacity constraints are then expressed in the following way:

$$\sum_{i=0}^{G} d^i_{k,t-i} \leq D_{kt}, \quad (k, t) \in \mathcal{K} \times \mathcal{T}; \tag{9}$$

$$\sum_{i=0}^{G} r^i_{k,t-i} \leq R_{kt}, \quad (k, t) \in \mathcal{K} \times \mathcal{T}. \tag{10}$$

Connections Between Flights. Some aircraft perform successive flights. Denote by $\mathcal{F}'$ the subset of $\mathcal{F}$ that contains the "continued" flights. 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.

For each continued flight f', we are also given its "slack" $s_{f'}$. The slack is defined as the number of time periods such that, if f' lands 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 1, where it can be seen that the slack $s_{f'}$ is equal to the difference between (1) the time interval between the scheduled departure time of f and the scheduled arrival time of f', and (2) the minimum "turnaround" time of the aircraft performing both flights.

Given the above setup, connections between flights are taken into account by the following constraints:

$$g_{f'} - s_{f'} \leq g_f, \ f' \in F'. \tag{11}$$

It should be noted that f' and f need not be physically performed by the same aircraft. Constraints (11) also hold for any pair of flights (f', f) such that f is not allowed to leave before f' lands, because some passengers from f' must transfer to f.

Although previous work (Vranas, 1992; Vranas et al., 1994a) has demonstrated that the propagation of delays has important effects, constraints (11) were not included in the basic model (P_1). The reason is that the CFMU has no data showing which successive flights are performed by the same aircraft or which flights are connected. It is hoped, however, that such data will become available in the future.

The Extended Integer Programming Model. Combining model (P_1) presented earlier with the departure (9) and arrival (10) capacity constraints, as well as with the constraints for connections (11), we get the following extended model:

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

$$\text{s.t.} \quad \textstyle\sum_{n=0}^{L_s} \sum_{x=0}^{L_s-n} \sum_{i=0}^{G} p^i_{s,t-n-i,t+1+x-i} \leq P_{st}, \quad (s,t) \in \mathcal{S} \times \mathcal{T}; \tag{12}$$

$$\textstyle\sum_{i=0}^{G} d^i_{k,t-i} \leq D_{kt}, \quad (k,t) \in \mathcal{K} \times \mathcal{T}; \tag{13}$$

$$\textstyle\sum_{i=0}^{G} r^i_{k,t-i} \leq R_{kt}, \quad (k,t) \in \mathcal{K} \times \mathcal{T}; \tag{14}$$

$$\textstyle\sum_{i=0}^{G} p^i_{stt'} = \Pi_{stt'}, \quad (s,t) \in \mathcal{S} \times \mathcal{T}, t + 1 \leq t' \leq t + 1 + L_s; \tag{15}$$

$$\textstyle\sum_{i=0}^{G} d^i_{kt} = \Delta_{kt}, \quad (k,t) \in \mathcal{K} \times \mathcal{T}; \tag{16}$$

$$\textstyle\sum_{i=0}^{G} r^i_{kt} = \Gamma_{kt}, \quad (k,t) \in \mathcal{K} \times \mathcal{T}; \tag{17}$$

$$g_{f'} - s_{f'} \leq g_f, \quad f' \in \mathcal{F}'; \tag{18}$$

$$g_f \leq G, \quad f \in \mathcal{F}; \tag{19}$$

$$\text{Delay} - \text{capacity constraints}; \tag{20}$$

$$g_f, p^i_{stt'} \text{ nonnegative integers.} \tag{21}$$

Δ_{kt} denotes the number of aircraft that are scheduled to depart from airport k at time period t; similarly for Γ_{kt} and arrivals. Constraints (16) and (17) are straightforward analogs of constraints (15). The total number of decision variables can be seen to be F (delay variables) + $2GKT$ (departure and arrival capacity variables) + G $\sum_{t \in T} \sum_{s \in S_t} (L_s + 1)(L_s + 2)/2$ (sector capacity variables).

Constraints (20) indicate again that model (P_2) is, as it stands, incomplete. Spelling out constraints that relate the delay and the capacity decision variables is the object of the appendix.

COMPUTATIONAL COMPARISON OF OPTIMAL SLOT ALLOCATION WITH CURRENT SLOT ALLOCATION PRACTICE

The purpose of this section is to compare optimal slot allocation with slot allocation as currently practised by the CFMU. There are several measures on which such a comparison might be based, but the most important one is arguably the sum of the delays imposed on all flights. The minimization of this total imposed delay is the objective of optimal slot allocation models. A secondary measure of comparison is the maximum imposed delay. Reducing this maximum delay is one way of approaching the goal of equitable treatment of all flights. Other measures of comparison, such as the percentage of flights that are delayed and the mean delay per delayed flight, will also be considered.

The computations were performed not with model (P_1), but with the following model (P_3), which is an adaptation to the current needs of the CFMU of a model developed by Bertsimas and Stock (1995)—itself a development of models presented by Vranas (1992) and Vranas et al., (1994a, 1994b).

$$(P_3) \min \sum_{f=1}^{F} c_f^g(t_f^d(w_{ft_f^d} - 1) + \sum_{t=t_f^d+1}^{t_f^d+G} t(w_{ft} - w_{f,t-1}))$$

$$\text{s.t.} \sum_{f:s\in\mathcal{S}_f} (w_{f,t-(t_f^{s+}-t_f^d)} - w_{f,t-(t_f^{s-}-t_f^d)}) \leq P_{st}, \quad (s,t)\in\mathcal{S}\times\mathcal{T}; \quad (22)$$

$$w_{f,t+1} \geq w_{ft}, \quad f\in\mathcal{F}, t_f^d \leq t \leq t_f^d + G - 1; \quad (23)$$

$$w_{f,t+G} = 1, \quad f\in\mathcal{F}; \quad (24)$$

$$w_{ft} \in \{0,1\}. \quad (25)$$

The (*interval*) *assignment decision variables* w_{ft} are defined as follows: w_{ft} is 1 if flight f takes off by time period t (i.e., during t or an earlier period) and 0 otherwise (i.e., if f takes off after t). Therefore, if f takes off at t, $w_{ft'} = 1$ for $t' \geq t$, and $w_{ft''} = 0$ for $t'' \leq t - 1$. It can be shown that what multiplies c_f^g in the objective function is in fact g_f. Constraints (22) are the sector capacity constraints, and constraints (24) ensure that f departs by $t + G$, i.e., that $g_f \leq G$. For further details see Bertsimas et al. (1995). (Constraints (23) are briefly explained in footnote 13.)

Computational experiments carried out by Bertsimas et al. (1995) and by Andreatta et al. (1996) indicate that the Bertsimas-Stock model is quite efficient. A reason for preferring model (P_3) to model (P_1) for the computations was that, as Bertsimas and Stock have found and explained (Bertsimas et al., 1995), the linear programming (LP) relaxation of their model normally has completely integral solutions. By solving the LP relaxation rather than the integer program, one achieves of course substantial savings in computation time. It will be seen later, however, that the solutions of the LP relaxation of model (P_3) turned out *not* to be completely integral. It should also be noted that the computations were performed with datasets corresponding to the initial period of operation of the CFMU, when only flights crossing the French airspace were regulated. Solving model (P_1) might be computationally advantageous now that the CFMU operations have expanded to cover the whole European airspace: as explained earlier, for $F = 10{,}000$, model (P_1) has about 75 percent fewer decision variables than model (P_3).

The Scenarios

The computations were based on data kindly provided by EUROCONTROL.[12] The datasets correspond to the approximately 5,800 flights that cross the French airspace daily. When the demand for a certain sector is expected to exceed the sector capacity for a portion of the day, a corresponding "regulation" is activated (corresponding to a sector capacity constraint in models (P_1) – (P_3)), and all flights scheduled to cross the given sector during the given portion of the day are subject to this regulation. Clearly, a flight may be subject to multiple regulations. Moreover, only a portion of the approximately 5,800 flights (usually about 1,500–2,000) will be regulated on any given day. Now that the tactical operations of the CFMU have expanded to cover the whole European airspace, the total number of flights that are regulated on any given day typically exceeds 6,000.

As shown in Table 2, five scenarios were considered, involving from 16 to 44 regulations and from 1,154 to 2,293 flights. Very roughly, scenarios $S1$, $S2$, and $S3$ were related as follows (Tibichte, 1995). $S1$ included all flights that were scheduled to cross a sector at a time when the demand for that sector exceeded its capacity. After CASA turned out the imposed delays for each of these flights, scenario $S2$ was derived from $S1$ by means of horizontal and vertical reroutings designed to ensure the feasibility of the CASA solution. After CASA turned out the imposed delays for $S2$, scenario $S3$ was derived from $S2$ by eliminating some redundant regulations. Given this indirect "operational optimization" process, it can be expected that, as one proceeds from $S1$ to $S3$, the CASA solution will somehow improve. This prediction was indeed confirmed, as discussed later.

Scenario $S4$ corresponds to a case where a breakdown occurs in the CASA system. To face this situation, sector capacities are reduced initially by 75 percent and then gradually restored to their original values after the breakdown is repaired. Finally, scenario $S5$ is a case

[12] The invaluable assistance of M. A. Tibichte of the EUROCONTROL Experimental Centre in Brétigny, France, is most gratefully acknowledged.

Table 2. Computational Results

Scenario	S1		S2		S3			S4		S5
1 # of flights	2,293		2,199		1,777			1,154		1,808
2 # of regulated sectors	25		25		22			44		16
CASA results										
3 CASA total delay (min)	60,349		44,234		41,581			52,779		50,288
4 CASA max. delay (min)	172		146		178			327		184
Optimization results										
5 Max. delay (min)	150	135	150	120	180	150	135	360	345	180
6 (# of 15-min periods)	(10)	(9)	(10)	(8)	(12)	(10)	(9)	(24)	(23)	(12)
LP relaxation optimum										
7 LP total delay (min)	34,740	35,295	28,440	29,505	27,945	28,155	28,425	51,000	51,000	45,285
8 LP/CASA total delay	58%	58%	64%	67%	67%	68%	68%	97%	97%	90%
9 CPU time Sun Sparc 2 (s)	3,514	3,856	5,405	1,912	2,608	1,257		1,336		7,380
10 CPU time Sun Sparc 5 (s)	1,100	1,200	1,200	780	700	500	300	250	250	
11 # of iterations CPLEX 3	11,320	11,835	16,604	7,642	12,192	6,264		8,778		22,098
12 # of iterations CPLEX 4	11,516	11,835	12,344	9,270	11,593	7,453	5,903	6,702	6,679	
13 # of noninteger variables	641	662	239	575	350	78	329	48	60	
Integer solution										
14 IP total delay (min)		35,565	28,515		28,005	28,185	28,995	51,075	51,075	
15 IP/CASA total delay		59%	64%		67%	68%	70%	97%	97%	
16 LP/IP total delay		99%	100%		100%	100%	98%	100%	100%	
17 CPU time Sun Sparc 5 (s)		5,046	12,006		4,631	2,271	5,902	336	429	
18 # of iterations CPLEX 4		50,297	36,446		21,194	13,400	19,270	7,144	7,960	
19 # of b&b nodes		643	866		519	483	834	45	78	

where more than 50 percent of the flights are subject to multiple regulations.

Computation Times

The fundamental advantage of CASA over an optimization model is that CASA, being a heuristic that requires no iterations, runs in very short computation times. For example, CASA ran in 35 seconds for scenario *S*3 and in 25 seconds for scenario *S*4. An optimization model requires thousands of iterations (lines 11, 12, and 18 of Table 2) and is thus bound to be much slower than CASA. The important question however is: how much slower? Slot allocations need to be updated dynamically at short intervals because new data keep coming in during the course of the day. (The CFMU calls this dynamic updating the "True Revision Process" and uses a series of message exchanges for it [EUROCONTROL 1994, Ed. 2.0].) Currently CASA reruns automatically every, e.g., 3 minutes, as well as whenever some modification for a flight is received. The CFMU has indicated a willingness to consider the possibility of a slot allocation process that might rerun as infrequently as, say, every 10 or at most 15 minutes (Gainche, 1996). Can this bound be achieved by optimization models?

It will be seen that it can, subject to some qualifications. The LP relaxations for most scenarios were solved twice: once on a Sun Sparc 2 workstation (40 MHz, 64 MB RAM, 4.2 MFlop) using CPLEX 3.0 as the solver, and once on a Sun Sparc 5 workstation (110 MHz, 192 MB RAM, 14.9 MFlop) using CPLEX 4.0 as the solver. (The modeling tool was GAMS 2.25.) A comparison of lines 9 and 10 of Table 2 shows that the computation times were reduced by about 60 percent to 80 percent. Line 17 of Table 2 gives the CPU times required to find a first integer solution on the Sun Sparc 5 with CPLEX 4.0. These times range from 5.6 to 200.1 minutes. The upper end of this range clearly exceeds the 15-minute bound, but it must be taken into account that, until the recent release of the Sun Ultra computer, the Sun computers were the slowest workstations on the market (Lowe, 1996). CPLEX Optimization, Inc., has most kindly agreed to solve, for comparison purposes, scenario *S*3 with an upper bound on delays of 150 minutes (Lowe, 1996). They used an SGI Power Challenge (4 CPU, 75 MHz) and found a first integer solution in 121 seconds, as opposed to the 2,271 seconds reported in Table 2. The optimal solution (with a total delay of 28,170 minutes) was found in a total of 339 seconds. Thus the 15-minute bound seems to be achievable.

It might be objected that the SGI Power Challenge is in fact a supercomputer and is thus prohibitively expensive. But the computation cost might be worthwhile, since it might result in a significant reduction of delays (as discussed later) and thus in annual savings of hundreds of millions of U.S. dollars in congestion costs. Moreover, operations research departments of airline companies use similar computers to solve routinely optimization problems of structure similar to that of model (P_3) (e.g., American Airlines uses an SGI Power Challenge, 2 CPU, 90 MHz).

Four further remarks might be made concerning computation times.

(1) Computation times depend not only on the hardware, but also on the algorithm. The computations reported in Table 2 were performed with the "netopt" command of CPLEX (extraction of a network and then dual Simplex method), which was found to have much better performance than either the "optimize" (primal Simplex) or the "tranopt" (dual Simplex) command. The Barrier algorithm of CPLEX was not available, but was used by CPLEX Optimization, Inc., in their solution of scenario $S3$, and might have contributed significantly to the reduction of the computation time. This is indicated by comparisons of the performance of the Barrier algorithm with that of the other algorithms of CPLEX, in computations performed by Andreatta and Brunetta (1996).

(2) In most computations performed in previous research with the Bertsimas-Stock model (Andreatta and Brunetta, 1996; Bertsimas, 1995), the solutions of the LP relaxation were completely integral, so that one did not need to engage in the time-consuming branch-and-bound process in order to find an integer solution. Line 13 of Table 2 indicates that, although the percentage of noninteger variables in the optimal solution of the LP relaxation was always small, it was never zero. In fact, a comparison of lines 10 and 17 shows that branching on these noninteger variables until an integer solution was found increased the total computation time quite significantly. Thus the hope that one can get away just by solving the LP relaxation is shattered.[13]

(3) Recall that CPLEX Optimization, Inc., required 121 seconds to find a first integer solution but 339 seconds to find the IP optimum. The IP optimum was 28,170 minutes (1,878 periods), whereas the LP optimum was 28,155 minutes (1,877 periods), and the value of the first integer solution reported in Table 2 was 28,185 minutes (1,879 periods). It seems thus that, once a first integer solution is found,

[13] A plausible theoretical explanation can be offered for this result. In the original Bertsimas et al. (1995) model, three kinds of connectivity constraints are facets of the convex hull of the set of solutions: sector connectivity constraints (an aircraft cannot enter a sector before leaving the previous one), flight connectivity constraints (an aircraft cannot depart for a next flight before it has arrived and stayed some time on the ground), and time connectivity constraints (if a flight has departed by time t, then it has departed by time $t + 1$). The capacity constraints for airports and sectors are *not* facets. Now model (P_3) is a special case of the Bertsimas-Stock model, a case in which no airborne delays are imposed, so that the sector and the time connectivity constraints drop out, and only the time connectivity constraints (23) remain; thus model (P_3) may not be as strong as the original Bertsimas-Stock model.

continuing the computation in order to find the IP optimum results in only a negligible improvement of the objective function that is not worth the significant increase in computation time. There is considerable evidence to indicate that this is always the case: the IP and LP optimal values are always very close (Vranas, 1992; Vranas et al., 1994a; cf. line 16 of Table 2).

(4) A further decrease in computation time might be achieved by developing heuristic algorithms to find a near-optimal solution. Very successful such heuristics have been developed for optimization models without sector capacities by Andreatta et al. (1994) and by Brunetta et al. (1995); it is reasonable to hope that these heuristics can be extended to the case in which sector capacities exist. (See Maugis, 1995, for further ways to find approximate solutions.)

Comparing CASA Solutions with Optimization Solutions

The comparisons will be based on the total delays, the maximum delays, and the distributions of delays. Some words will also be said about the delays of simple versus combined flights.

Comparing Total Imposed Delays. Several computations were performed for each of the five scenarios, in order to find by trial-and-error the lowest upper bound G on delays that resulted in a feasible problem. The values of G are given in lines 5 (in minutes) and 6 (in number of 15-minute periods) of Table 2. For instance, scenario $S1$ was solved with $G = 10$ and with $G = 9$; it was infeasible for $G = 8$. Similarly, $S2$ was infeasible for $G = 7$, $S3$ for $G = 8$, $S4$ for $G = 22$, and $S5$ for $G = 11$.

Line 15 (cf. line 8) of Table 2 shows that, for scenarios $S1$, $S2$, and $S3$, the optimal total delay was about two-thirds the CASA total delay or less. Moreover, the ratio of the optimal total delay over the CASA total delay increased from 59 percent for $S1$ to about 68 percent for $S3$, confirming thus the prediction that the CASA solution improves as one proceeds from $S1$ to $S3$ by means of the indirect "operational optimization" process roughly described earlier.

For scenario $S4$, the optimal total delay was almost the same as the CASA total delay. It may be the case that for the extremely restricted sector capacities of this "breakdown" scenario (several capacities had a value of zero or one flight per hour after the breakdown), a first-come-first-served procedure like that applied by CASA gives near-optimal solutions.

Finally, scenario $S5$ gives a ratio of LP optimal over CASA total delay of 90 percent.

Comparing Maximum Imposed Delays. By comparing lines 4 and 5 of Table 2 one can reach the following two conclusions. First, it is sometimes feasible to reduce the maximum delay of CASA; the reduction is about 20 to 25 percent for scenarios $S1$, $S2$, and $S3$. Second, this reduction is not always possible, as indicated by scenario $S5$. Notice that, for scenario $S4$, $G = 22$ (i.e., 330 minutes) results in an infeasible solution, whereas CASA found a feasible solution with a maximum delay of 327 minutes. This result indicates that the two solutions are not directly comparable, presumably because the EUROCONTROL data were transformed in such a way that, if a flight was present at a sector during some part of a 15-minute period, it was taken to be present at the sector during the whole 15-minute period (cf. footnote 9). This incomparability deserves further investigation.

Lines 7 and 14 of Table 2 indicate that, in the cases in which it is possible to reduce the maximum imposed delay, one can do so without appreciably increasing the total imposed delay. For instance, a 25 percent reduction of G from 12 to 9 periods for scenario $S3$ results in an increase of less than 4 percent (from 28,005 to 28,995 minutes) in the total delay of the first integer solution.

Comparing Distributions of Delays. Table 3 gives some data on the distributions of delays for the CASA solutions of scenarios $S1$–$S4$, as well as the complete integer solutions for the seven problems for which such solutions were found (cf. Table 2). A comparison of lines 10 and 2 (cf. line 12) of Table 3 shows that the (near-)optimal solution reduces drastically the number of delayed flights, in some cases by about a third and in other cases by almost two-thirds.[14] Given that this reduction is more than the reduction in total delay (line 15 of Table 2), *the mean delay per delayed flight increases significantly in the optimal solution* (lines 1, 8, and 9 of Table 3).

It can be seen also from lines 13–37 of Table 3 that the optimal solutions for scenarios $S1$–$S3$ have a concentration of delays at the highest possible delay G. Such a concentration is not observed for scenario $S4$, presumably because for $S4$ G has a value about twice the value it has for $S1$–$S3$. It would be interesting to see whether the CASA solutions exhibit similar effects, but unfortunately the complete CASA solutions were not available.

Simple Versus Combined Flights. Given that combined flights (by definition) are subject to more than one regulation, they are expected to incur a greater mean delay than simple (i.e., noncombined) flights. Relevant data were available for the CASA solution of scenario $S3$. Line 4 of Table 4 shows that 36 percent of delayed flights

[14] This effect is so large that it survives a possible distortion due to time discretization: some of the flights that CASA delays less than 15 minutes might be considered to be nondelayed for purposes of comparison with the optimal solution.

Table 3. Distributions of Delays

	Scenario	S1	S2	S3			S4	
	CASA results							
1	Mean delay/delayed (min)	40	34		37		56	
2	# of delayed flights	1,501	1,302		1,133		942	
3	% of delayed flights	65%	59%		64%		82%	
4	# flights delayed < 15 min				318		179	
5	# flights delayed 30–60 min	364	393		340		305	
6	# flights delayed > 60 min	366	181		170		263	
	Optimization results							
7	G (# of 15-min periods)	9	10	12	10	9	24	23
8	Mean delay/dealyed (min)	61	57	71	61	65	86	85
9	(1)/(8) [CASA/Integer]	66%	60%	52%	61%	57%	65%	66%
10	# delayed flights	583	501	394	462	447	597	603
11	% of delayed flights	25%	23%	22%	26%	25%	52%	52%
12	(10)/(2) [Integer/CASA]	39%	38%	35%	41%	39%	63%	64%
13	# flights with delay 1	180	166	112	148	134	146	171
14	# flights with delay 2	85	84	67	68	64	101	85
15	# flights with delay 3	63	47	35	46	35	58	59
16	# flights with delay 4	37	38	25	34	28	49	43
17	# flights with delay 5	22	34	20	23	21	26	32
18	# flights with delay 6	32	25	16	24	21	25	34
19	# flights with delay 7	32	18	19	23	17	34	25
20	# flights with delay 8	30	16	10	19	39	25	20
21	# flights with delay 9	102	30	10	22	88	9	12
22	# flights with delay 10	0	43	10	55	0	15	12
23	# flights with delay 11	0	0	23	0	0	13	16
24	# flights with delay 12	0	0	47	0	0	9	5
25	# flights with delay 13	0	0	0	0	0	18	9
26	# flights with delay 14	0	0	0	0	0	9	12
27	# flights with delay 15	0	0	0	0	0	8	12
28	# flights with delay 16	0	0	0	0	0	9	4
29	# flights with delay 17	0	0	0	0	0	8	10
30	# flights with delay 18	0	0	0	0	0	7	9
31	# flights with delay 19	0	0	0	0	0	5	8
32	# flights with delay 20	0	0	0	0	0	5	6
33	# flights with delay 21	0	0	0	0	0	4	7
34	# flights with delay 22	0	0	0	0	0	5	4
35	# flights with delay 23	0	0	0	0	0	4	8
36	# flights with delay 24	0	0	0	0	0	5	0
37	# flights with delay > 24	0	0	0	0	0	0	0

were combined in the CASA solution, whereas 61 percent of delayed flights were combined in the integer (near-)optimal solution (with $G = 10$). Line 8 of Table 4 shows that combined flights incurred 55 percent of the total delay in the CASA solution, but 78 percent in the optimal solution. One might think then that the optimal solution discriminates heavily against combined flights. But a better measure of possible discrimination might be the ratio *mean delay per combined*

Table 4. Simple Versus Combined Flights

Scenario $S3$	CASA	Optimization
1 # of delayed flights	1,133	462
2 # of simple delayed flights	730	179
3 # of combined delayed flights	403	283
4 % of delayed flights that are combined	36%	61%
5 Total delay (min)	41,581	28,185
6 Total delay of simple flights (min)	18,789	6,255
7 Total delay of combined flights (min)	22,792	21,930
8 % of total delay incurred by combined	55%	78%
9 Mean delay per delayed flight (min)	37	61
10 Mean delay per simple delayed flight (min)	26	35
11 Mean delay per combined delayed flight (min)	57	77
12 Ratio of line 11 over line 10	2.2	2.2

delayed flight over mean delay per simple delayed flight, and this ratio has the same value (line 12 of Table 4) for both the CASA and the optimal solutions. Therefore, there is a sense in which the optimal solution does not discriminate against combined flights more than CASA does.

Sensitivity Analysis

Upper Bounds on Delays. By comparing the two columns of Table 3 corresponding to the integer solutions for scenario $S4$, it is seen that the same total delay can be achieved with different values of the upper bound G on delays. The mean delays (per delayed flight) are about the same for both solutions (86 minutes for $G = 24$ versus 85 minutes for $G = 23$) and the same holds for the standard deviations (84 minutes for $G = 24$ versus 86 minutes for $G = 23$). A comparison of the three columns of Table 3 corresponding to the integer solutions for scenario $S3$ leads to similar conclusions. It seems thus that it is profitable to find, by trial and error, the lowest value of G for which the problem remains feasible.

Time Discretization. The LP relaxation of scenario $S3$ was solved on the Sun Sparc 5 by discretizing time into periods of 5 minutes and with $G = 30$ periods (150 minutes). The LP optimum was found in about 10,200 CPU seconds after 46,226 iterations, and the total delay was 28,920 minutes. This value is only 2.7 percent higher than the corresponding value for periods of 15 minutes. This result indicates that the effects of time discretization may not be significant, although further investigation is certainly called for.

Summary of Results

To summarize, the computational results reported in this section constitute (defensible) evidence for the following conclusions.

(1) It is possible to solve two optimality integer programs corresponding to the portion of the European air traffic that crosses the French airspace in less than 10 minutes by using supercomputers and the Barrier algorithm of CPLEX. The computation time can be significantly reduced if one settles for a first integer solution, which in any case usually differs only negligibly from the optimal solution. However, one cannot settle for just the optimal solution of the LP relaxation.

(2) In some (but not all) cases, the optimal solution can reduce significantly the total delay imposed by CASA. In some (but not all) cases, this reduction in the total delay can be achieved concurrently with a reduction of the maximum delay imposed by CASA.

(3) The optimal solution delays many fewer flights than CASA does, but the mean delay per delayed flight increases appreciably in the optimal solution.

(4) In one sense, the optimal solution discriminates against combined flights more than CASA does, because it increases the percentage of delayed flights that are combined and the percentage of the total delay that is incurred by combined flights. In another sense, however, the optimal solution does not discriminate against combined flights more than CASA does, because it leaves unchanged the ratio of the mean delay per combined flight over the mean delay per simple flight.

(5) The length of the time periods into which time is discretized does not seem to affect appreciably the total delay in the optimal solution.

CONCLUSIONS

This paper has presented what seems to be the first concrete evidence for the importance of incorporating optimization models into the European slot allocation process. The computational results of the previous section suggest that optimal slot allocation can sometimes reduce both the total and the maximum imposed delays, resulting in a utilization of the European airspace that is at once more efficient and more equitable.

The above conclusion needs two kinds of qualifications.

(1) Besides ensuring safety and optimizing the use of the available capacity, an important objective of the CFMU is to guarantee the equitable treatment of all users of the European airspace. As long as this *principle of equity* is understood as implying the "First Planned, First Served" principle applied by CASA, equity will conflict with the minimization of delays, and an optimization procedure will seem out of place. This reasoning, however, assumes that the objective of equity is more important than the objective of optimizing the use of

the available capacity, and it is not clear why this should be the case. Moreover, it is not clear why the principle of equity should be understood as implying the First Planned, First Served principle; another possible way to ensure equity is to include in the optimization models constraints guaranteeing that the mean delay per flight is the same for all airlines (Vranas, 1992; cf. Maugis, 1995).

(2) The optimal solutions turned out by the integer programming models are not *exactly* comparable to the CASA solutions, mainly for the following reasons. First, CASA goes through an "operational optimization" process which ensures the geometric feasibility of the slot allocation pattern; this feasibility is not ensured by the integer programs. Second CASA does not discretize time, whereas the integer programs do. Although evidence was adduced earlier that the discretization effects may not be important in terms of the total delay, these effects may still be important in terms of the feasibility of solutions, as indicated by the fact that, as discussed earlier, in one case CASA found a feasible solution, but the corresponding integer program was infeasible.[15] In spite of these qualifications, however, some of the differences between the CASA and the optimal solutions reported in the previous section are of sufficient magnitude to indicate that they would survive an exact comparison.

Although, at the current state of the art, optimization models cannot be used for the dynamic updating of slots (unless they are implemented on supercomputers), such models could very well be used even now for the *pre*tactical operations of the CFMU, in which there is not such a high time pressure as in the tactical activities. Moreover, if quick and efficient heuristics, such as those presented in Andreatta et al. (1994) and Brunetta et al. (1995), were extended to the case with sector capacities, near-optimal solutions could be found even with ordinary workstations sufficiently quickly to allow dynamic updating of slots.

Following are some suggestions for future research.

(1) More insight is needed into the conditions under which optimal slot allocation improves the CASA solutions. For example, it was seen earlier that the CASA solution was near optimal for scenario $S4$ and close to optimal for scenario $S5$. More computational comparisons are needed to clarify such phenomena.

(2) The effect of discretizing time should be further investigated. In particular, it should be clarified why sometimes the integer program is infeasible, whereas CASA finds a feasible solution.

(3) Comparisons between optimal slot allocation and the CASA solution should also be effected for bigger problems with, say, 10,000

[15] See Vranas (1992), and Vranas et al. (1994a) for another case in which two optimization problems had almost the same optimal values, but the optimal solution of the one was highly infeasible for the other.

regulated flights, since the CFMU operations currently cover such big problems. It cannot just be assumed that the advantage of optimization models over CASA will remain the same when the size of the problem changes; the advantage may increase or decrease. Moreover, it is not clear whether such big problems can be solved within 10 minutes, even by using supercomputers. Model (P_1) should also be investigated for big problems, because its smaller size may allow it to perform better than model (P_3).

(4) Finally, the structure of optimal slot allocation and that of the CASA solutions should be compared in more detail, with respect, e.g., to the variations from the First Planned, First Served principle induced by the optimal solutions. Moreover, alternative ways to define equity should be considered and incorporated as constraints in the optimization models, and their effect on the optimal solutions should be investigated.

We conclude by expressing the hope that the present work will spur further research on this eminently practical problem.

APPENDIX: RELATING DELAY AND CAPACITY DECISION VARIABLES

In models (P_1) and (P_2) of the second section, if the delay-capacity constraints are omitted, the optimal solution has all g_f equal to 0—an absurd result. It should be thus emphasized that the discussion of the delay-capacity constraints was omitted from the second section not because of any lack of importance, but because it requires a rather lengthy exposition.

The following subsection formulates the basic delay-capacity constraints. The second subsection gives a method for generating a large number of additional delay-capacity constraints and discusses which of these constraints it is advisable to include in models (P_1) and (P_2).

The Basic Delay-Capacity Constraints

Denote by $\mathcal{F}_{stt'}$ the set of flights that are scheduled to enter sector s at time period t and to exit at t'. (The number of such flights was denoted above by $\Pi_{stt'}$—see Table 1.) Similarly, denote by $\mathcal{F}^d_{kt}$ the set of flights that are scheduled to depart from airport k at time period t and by $\mathcal{F}^a_{kt}$ the set of flights that are scheduled to arrive at k at t. Then the following constraints hold:

$$\sum_{f\in\mathcal{F}_{stt'}} g_f = \sum_{i=1}^{G} ip^i_{stt'}, \ (s,t) \in S \times T, \ t+1 \leq t' \leq t+1+L_s; \quad (26)$$

$$\sum_{f\in\mathcal{F}^d_{kt}} g_f = \sum_{i=1}^{G} i d^i_{kt}, \quad (k,t) \in \mathcal{K} \times \mathcal{T}; \tag{27}$$

$$\sum_{f\in\mathcal{F}^r_{kt}} g_f = \sum_{i=1}^{G} i r^i_{kt}, \quad (k,t) \in \mathcal{K} \times \mathcal{T}. \tag{28}$$

To understand, e.g., constraints (26), notice that $p^{G-1}_{stt'}$ of the flights in $\mathcal{F}_{sst'}$ have $g_f = G$ and so contribute a total of $Gp^G_{stt'}$ to the sum at the left-hand side of (26); $p^G_{stt'}$ of the flights in $\mathcal{F}_{stt'}$ have $g_f = G - 1$ and so contribute a total of $(G - 1)p^{G-1}_{stt'}$ to the sum at the left-hand side of (26); and so on. Similarly for constraints (27) and (28).

A given variable g_f appears in exactly one of constraints (27) and in exactly one of (28), but it appears in as many of constraints (26) as there are regulated sectors that flight f goes through.

Additional Delay-Capacity Constraints

It will be shown now how to generate additional constraints relating the delay decision variables g_f with the sector capacity decision variables $p^i_{stt'}$. Constraints relating g_f with d^i_{kt} and with r^i_{kt} can be similarly generated.

Consider any set of flights $\emptyset$ which are scheduled to enter sector s at t and to exit at t' (i.e., consider any subset of $\mathcal{F}_{stt'}$). Denote by $\Pi^i_{stt'}$ the number of flights in $\emptyset$ that have $g_f = i$. Clearly, $0 \le \Pi^i_{stt'} \le p^i_{stt'}$. Consider now the sum $\Sigma^G_{i=0}(i - a)\Pi^i_{stt'}$, where a is any integer from 1 to G. We will relate this sum to $\Sigma_{f\in\emptyset}(g_f - a)$.

By a reasoning similar to that establishing constraints (26), it can be seen that $\Sigma_{f\in\emptyset}(g_f - a) = \Sigma^G_{i=0}(i - a)\Pi^i_{stt'}$. By neglecting the terms with nonpositive coefficients $i - a$ in the right-hand side of the previous equality, we get $\Sigma_{f\in\emptyset}(g_f - a) \le \Sigma^G_{i=a+1}(i - a)\Pi^i_{stt'}$. But given that $\Pi^i_{stt'}, \le p^i_{stt'}$, we finally get:

$$\sum_{f\in\emptyset} (g_f - a) \le \sum_{i=a+1}^{G} (i - a)p^i_{stt'} \ \forall\ \emptyset \subset \mathcal{F}_{stt'}. \tag{29}$$

A similar reasoning, neglecting the nonnegative rather than the nonpositive coefficients $i - a$ in the sum $\Sigma^G_{i=0}(i - a)\Pi^i_{stt'}$, establishes the following additional constraints:

$$\sum_{f\in\emptyset} (a - g_f) \le \sum_{i=0}^{a-1} (a - i)p^i_{stt'} \ \forall\ \emptyset \subset \mathcal{F}_{stt'}. \tag{30}$$

Substituting the values of a from 1 to G in (29) and (30) gives, e.g., for $G = 4$ the following seven sets of constraints[16]:

$$\sum_{f\in\phi} (g_f - 3) \leq p^4_{stt'}; \tag{31}$$

$$\sum_{f\in\phi} (g_f - 2) \leq 2p^4_{stt'} + p^3_{stt'}; \tag{32}$$

$$\sum_{f\in\phi} (g_f - 1) \leq 3p^4_{stt'} + 2p^3_{stt'} + p^2_{stt'}; \tag{33}$$

$$\sum_{f\in\phi} (1 - g_f) \leq p^0_{stt'}; \tag{34}$$

$$\sum_{f\in\phi} (2 - g_f) \leq 2p^0_{stt'} + p^1_{stt'}; \tag{35}$$

$$\sum_{f\in\phi} (3 - g_f) \leq 3p^0_{stt'} + 2p^1_{stt'} + p^2_{stt'}; \tag{36}$$

$$\sum_{f\in\phi} (4 - g_f) \leq 4p^0_{stt'} + 3p^1_{stt'} + 2p^2_{stt'} + p^3_{stt'}, \ \forall\ \phi \subset \mathcal{F}_{stt'}. \tag{37}$$

In order to understand the meaning of the above sets of constraints, consider, e.g., constraints (31). Consider first the case in which ϕ has just one element. Then (31) become: $(\forall f\in\mathcal{F}_{stt'})\ g_f - 3 \leq p^4_{stt'}$. Therefore, if there is at least one flight in $\mathcal{F}_{stt'}$ with $g_f = 4$, the corresponding constraint will give: $p^4_{stt'} \geq 1$. Consider next the case in which ϕ has two elements. Then, if there are at least two flights in $\mathcal{F}_{stt'}$ with $g_f = 4$, the corresponding constraint will give: $p^4_{stt'} \geq 2$. In conclusion, if all of constraints (31) are included in the models, then $p^4_{stt'}$ will be no less than the value it should have (the value it should have being, of course, the number of flights in $\mathcal{F}_{stt'}$ with $g_f = 4$).

Similar reasoning shows that, *if* $p^4_{stt'}$ has the value it should have, then including all of constraints (32) in the models ensures that $p^3_{stt'}$ will be no less than the value it should have. This means that including constraints (32) makes sense only if *all* of constraints (31) are also included. However, even if all of constraints (31) are included,

[16] The eighth set of constraints is $\Sigma_{f\in\phi} g_f \leq \Sigma^4_{i=1}\, ip^i_{stt'}$, which is entailed by (26).

constraints (32) might still not be effective; this is because (31) do not exclude the case in which $p^4_{stt'}$ is *greater* than the value it should have. This case, however, is not to be expected, given constraints (26) and the fact that the sum of all g_f is to be *minimized.*

Two conclusions arise from the above discussion. First, (32) should be included in the models only if all of (31) are also included, (33) should be included only if all of both (31) and (32) are also included, and similarly for (34) to (37). Second, (32) are less useful than (31), (33) are less useful than (32), and similarly for (34) to (37).

A problem may arise because the number of any of the above sets of constraints looks quite high: it is equal to the number of subsets of $\mathcal{F}_{stt'}$ minus 2,[17] namely to $2^{\Pi_{stt'}} - 2$. On the other hand, $\Pi_{stt'}$ is normally fairly small, rarely exceeding 10. So the number of these constraints may not be inordinately high after all. But a more serious problem might be that, even if all of the above seven sets of constraints were included in the models, it might still be the case that in the optimal solution, the values of g_f and $p^i_{stt'}$ would not be fully compatible.

The easiest course of action in such a case would be to keep the values of g_f turned out in the optimal solution and to modify the values of $p^i_{stt'}$ in order to make them compatible with g_f. This may require the increase of some $p^i_{stt'}$, resulting thus possibly in the temporary violation of some sector capacity constraints; this violation would have to be then heuristically compensated for. It should be noted that such a way of proceeding should not be unacceptable to the CFMU. Indeed, currently the CFMU accepts the concept of "slot forcing" (EUROCONTROL, 1993, Ver. 1.8, pp. 14–15; see also EUROCONTROL, 1993, Ver. 1.0, p. 10), which is a temporary violation of the sector capacity constraints (a violation subsequently compensated for by CASA (Gainche, 1996)), arising from the fact that some flights are subject to "multiple regulations" (i.e., pass through several congested sectors): when the time comes to allocate a slot to such a flight, all available capacity in some of the sectors through which the flight passes may have already been used up.

In conclusion, the best way to relate the delay and the capacity decision variables cannot be fully determined *a priori*, but needs to take into account computational experiments. Such experiments were not undertaken in the context of the present work, because it was considered more urgent to compare current slot allocation practice with optimal slot allocation based on an established model.

[17] It can be seen that, for $\emptyset = \mathcal{F}_{stt'}$, all of (31)–(37) are entailed by (26) and (4).

ACKNOWLEDGMENTS

The present work was carried out at the German Aerospace Research Establishment (DLR) in Braunschweig, Germany, for the most part while I held a research fellowship from the Alexander von Humboldt Foundation. I feel deeply indebted towards the Alexander von Humboldt Foundation for its support. I am also most grateful to DLR, especially to Prof. Dr. H. Winter, Dr. U. Völckers, and Dr. U. Brokof, for securing *ideal* working conditions. I would also like to express my gratitude to EUROCONTROL, especially to MM. F. Gainche, A. Vidal, M. Dalichampt, and A. Tibichte, for providing the data as well as much useful information. The help of M. A. Tibichte was crucial for the timely completion of the present work. Finally, I would like to thank the following people for various kinds of assistance: D. Bertsimas, R. Küttner, G. Linde, T. Lowe, K. Muth, A. Odoni, M. Schiele, M. Schubert, S. Stock, and especially J. Theuerkauf.

ACRONYMS

ATCSCC	Air Traffic Control System Command Center
CASA	Computer Assisted Slot Allocation
CFMU	Central Flow Management Unit
CPU	central processing unit
EUROCONTROL	European Organization for the Safety of Air Navigation
FAA	Federal Aviation Administration
SAP	Slot Allocation Problem

REFERENCES

Andreatta, G., and G. Romanin-Jacur (1987), "Aircraft Flow Management under Congestion," *Transportation Science,* 21, 249–253.

Andreatta, G., A. R. Odoni, and O. Richetta (1993), "Models for the Ground-Holding Problem," in *Large-Scale Computation and Information Processing in Air Traffic Control,* L. Bianco and A. R. Odoni (eds.). Springer-Verlag, Berlin, 125–168.

Andreatta, G., L. Brunetta, and G. Guastalla (1994), "Multi-Airport Ground Holding Problem: A Heuristic Approach Based on Priority Rules," Internal Report 9/94, Dip. di Matematica Pura e Applicata, Università di Padova.

Andreatta, G., and G. Tidona (1995), "A Note on the Multi-Airport Ground-Holding Problem in Air Traffic Control," *Operations Research* (to appear).

Andreatta, G., and L. Brunetta (1996), "Multi-Airport Ground Holding Problem: A Computational Evaluation of Exact Algorithms," *Operations Research* (to appear).

Bardin, E., and D. Dufraisse (1992), "CASA Study Simulations Report," CAP Gemini International Support.

Bertsimas, D. J., and S. Stock (1995), "The Air Traffic Flow Management Problem with Enroute Capacities," *Operations Research* (to appear).

Booth, G. R. (1994), "Flow Management in the United States," in *Advanced Technologies for Air Traffic Flow Management*, pp. 107–120.

Brunetta, L., G. Guastalla, and L. Navazio (1995), "A New Approach for Solving the Multi-Airport Ground Holding Problem," Internal Report 25/95, Dip. di Matematica Pura e Applicata, Università di Padova.

Dalichampt, M. (1994), EUROCONTROL Experimental Centre, Personal communication.

EUROCONTROL: Central Flow Management Unit (1994), "Guide to ATFM Message Exchange," EUROCONTROL document ref. TACT/USD/MSGGUID. Edition 2.0, 22 June 1994.

EUROCONTROL: Central Flow Management Unit (1994), CFMU Project: Executive Summary, Edition No. 7.0.

EUROCONTROL: Central Flow Management Unit (1994), Overview of CFMU Operations, Edition No. 1.0.

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

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

EUROCONTROL: Experimental Centre (1995), "ATFM Simulation of CFMU Phase 1," EEC Report No. 291.

Federal Aviation Administration System Capacity and Requirements Office (1988), *1990–91 Aviation System Capacity Plan*, DOT/FAA/SC-90-1, Department of Transportation, Washington, D.C., pp. 1–11.

Gainche, F. (1996), (1994), EUROCONTROL: Central Flow Management Unit, Personal communications.

Gray, P. (April 1996), "How High Can They Fly?" *Time Magazine*.

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

Helme, M. P., K. S. Linday, and S. V. Massimini (1992), "Optimization Models for Air Traffic Flow Control Management: New Formulations," MTR-92W0000152, The MITRE Corporation, McLean, VA.

Hughes, A. J., and J. L. Mann (1995), Final Report for WPR2.3 of Project MUFTIS, Document ref.: MFTS/DRA/WPR2.3/1.

Küttner, R. (1995), "Verkehrsstromtheorie Luftverkehr: Systemanalytische Untersuchungen im Luftverkehr," Deutsche Forschungsanstalt für Luft- und Raumfahrt (DLR) Braunschweig. Unpublished.

Lindsay, K. S., E. A. Boyd, and R. Burlingame (1993), "Traffic Flow Management Modeling with the Time Assignment Model," *Air Traffic Control Quarterly*, 1, 255–276.

Lowe, T. A. (1996), CPLEX Optimization, Inc., Personal communications.

Maugis, Lionnel (1995), "Mathematical Programming for the Air Traffic Flow Management Problem with En-Route Capacities," To be presented at the XIV World Conference on Operational Research, Vancouver, Canada. (First version published as technical report CENA/R95-022, June 1995.)

Odoni, A. R. (1987), "The Flow Management Problem in Air Traffic Control," in *Flow Control of Congested Networks*, A. R. Odoni, L. Bianco, and G. Szego (eds.), Springer-Verlag, Berlin, 269–288.

Odoni, A. R. (1994), "Issues in Air Traffic Flow Management," in *Advanced Technologies for Air Traffic Flow Management*, pp. 43–63.

Odoni, A. R. (1995), personal communication.

Philipp, W., and F. Gainche (1994), "Air Traffic Flow Management in Europe," *Advanced Technologies for Air Traffic Flow Management*, pp. 64–106.

Richetta, O. (1991), "Ground Holding Strategies for Air Traffic Control Under Uncertainty," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Richetta, O., and A. R. Odoni (1993), "Solving Optimally the Static Ground-Holding Policy Problem in Air Traffic Control," *Transportation Science*, 27, 228–238.

Richetta, O., and A. R. Odoni (1995), "Dynamic Solution to the Ground-Holding Policy Problem in Air Traffic Control," *Transportation Research* (to appear).

Terrab, M. (1990), "Ground Holding Strategies for Air Traffic Control," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Terrab, M., and A. R. Odoni (1993), "Strategic Flow Control on an Air Traffic Network," *Operations Research*, 41, 138–152.

Terrab, M., and S. Paulose (1995), "Dynamic Strategic and Tactical Air Traffic Flow Control," RPI Technical Report.

Tibichte, A. (1995), EUROCONTROL Experimental Centre, Personal communications.

Tosic, V., and O. Babic (1995), "Air Route Flow Management—Problems and Research Efforts," *Transportation Planning and Technology*, 19, 63–72.

Transportation Research Board (1991), *Winds of Change: Domestic Air Transport since Deregulation*, National Research Council, Washington, D.C.

Vranas, P. B. (1992), "The Multi-Airport Ground-Holding Problem in Air Traffic Control," Ph.D. Thesis, Massachusetts Institute of Technology, Cambridge, MA.

Vranas, P. B., D. J. Bertsimas, and A. R. Odoni (1994a), "The Multi-Airport Ground-Holding Problem in Air Traffic Control," *Operations Research*, 42, 249–261.

Vranas, P. B., D. J. Bertsimas, and A. R. Odoni (1994b), "Dynamic Ground-Holding Policies for a Network of Airports," *Transportation Science*, 28, 275–291.

Winter, H., and H.-G. Nüßer (eds.) (1994), *Advanced Technologies for Air Traffic Flow Management*, Proceedings of an International Seminar Organized by Deutsche Forschungsanstalt für Luft-und Raumfahrt (DLR), Bonn, Germany, April 1994. Lecture Notes in Control and Information Sciences 198. London: Springer.

BIOGRAPHY

Peter B. Vranas's undergraduate degrees were in the fields of naval architecture and marine engineering in Greece, and economics in France. In 1987 he went to MIT, where he received a masters degree in ocean systems management and a doctorate in operations research. After military service in the Greek Navy, he was an Alexander von Humboldt Research Fellow at the German Aerospace Research Establishment in Germany, where most of the work reported here was performed. He is now a Ph.D. student in the Department of Philosophy at the University of Michigan.