Journal of the Operational Research Society (2001) 52, 1338–1349

www.palgrave-journals.com/jors

Optimisation models for re-routing air traffic flows in Europe

P Leal de Matos[1]*, B Chen[2] and RJ Ormerod[2]

[1]Instituto Superior Técnico, SAEG, Lisbon, Portugal; and [2]*Warwick University, Coventry, UK*

This paper focuses on the usability of optimisation models to support the re-routing of air traffic flows in Europe. It discusses different modelling approaches and proposes three integer programming models, with different levels of detail, aimed at re-routing air traffic flows. The models are tested on a set of test data based on the actual traffic crossing the whole French upper airspace. The paper arrives at conclusions on the usability, limitations and extensions of the models.

Keywords: integer programming; optimisation; air transport

Introduction

Congestion in the air transportation system has been plaguing air traffic both in the US and in Europe for nearly 20 y. To protect air traffic control (ATC) from overloads, a planning activity called air traffic flow management (ATFM) emerged during the 1970s. ATFM, by comparing the available capacity with forecast traffic demand some time prior to the flights, tries to anticipate overloads and take control actions to prevent them.

In simple terms, the capacity of an ATC sector is defined as the number of flights that the air traffic control team of that sector is able to supervise per period of time, usually one hour. When the traffic expected to cross the sector exceeds the capacity, traffic delays occur. In the summer of 1998, more than 20% of flights in Europe were delayed due to ATC capacity constraints. On average, each of these flights was delayed by more than 20 min.[1] It is claimed that delays caused by lack of capacity cost European carriers around $5.4 billion in 1998.[2] Moreover, these estimates do not take into account the cost of delays to passengers, or the cost of the heavier burden on the controllers and on other elements of the air transportation system. ATFM tries to limit the extent and impact of those delays.

ATFM control actions range from departure delays to re-routing of flights. Departure delay, or ground delay, means delaying departures of flights heading to congested areas. The idea behind it is that, if delays are unavoidable, it is safer and cheaper to delay the flights on the ground than in the air. Flights can be re-routed to bypass already overloaded elements of the airspace or to prevent overloads occurring.

In continental US there is a single body located near Washington DC which coordinates flow management: the Air Traffic Control System Command Center. Congestion problems in the US are experienced mostly at airports. In Europe, a continent with many countries each with its own airspace, coordinated air traffic control and flow management is more difficult to implement. Many flights in Europe take only one hour or less but have to cross several airspaces. Congestion is felt not only at airports, as is mostly the case in the US, but also in the airspace at the junction points of air routes (also called fixes). Therefore, the thrust of air traffic management and control efforts in Europe has been to integrate and centralise control activities. To this end, the Central Flow Management Unit (CFMU), located in Brussels, was created in 1989 to be the sole provider of air traffic flow management in the 38 countries of the European Civil Aviation Conference (ECAC).

Leal de Matos and Ormerod,[3] in a paper based on fieldwork done at the CFMU, address the differences in timescale and organisation of ATFM between the US and Europe: in the US most of the planning is done a few hours before the flight's depart by the Air Traffic Control System Command Center whereas in Europe the planning starts six months before departure and involves not only flow managers but also different National Administrations, area control centres and aircraft operators' representatives. Accordingly, concepts differ: US authors tend to call all the planning done before the flights take-off 'strategic' and after the flight's take-off 'tactical'. In Europe, there is strategic planning, which goes from 6 months ahead to a few days before departure, pre-tactical planning, which occurs on the two days before departure, and tactical planning, which takes place on the day of departure until take-off. Measures affecting airborne flights are considered strictly in the realm

---

*Correspondence: P Leal De Matos, Centre of Business Studies, Instituto Superior Técnico SAEG, Av. Rovisco Pais, 1049-001 Lisbon, Portugal.
E-mail: Paula.Leal-de-Matos@eurocontrol.be

of ATC rather than ATFM. Accordingly, three levels of re-routing control measures can be identified:

1. Tactical level: re-routing of individual flights on the day of operations, usually flights with considerable ground-delay.
2. Pre-tactical level: re-routing of air traffic flows a few days before the day of operations in order to prevent or alleviate congestion. A flow is defined as a set of flights departing from one airport or an airport area to another airport or airport area at any time.
3. Strategic level: routing of air traffic flows for the season, a few months before the operations, to attain a balanced distribution of traffic and a good use of capacity. Most of these routings are published in a document called the Standard Routing Scheme.

These three levels of re-routing differ in the timeframe considered, the uncertainty involved and the room for manoeuvre. At the strategic level, where the re-routing measures are prepared a few months beforehand to last for a whole summer season, there is scope for significant changes in the routing of flows. However, the information available on traffic demand a few months ahead is very uncertain. As the day of operations gets closer the information available becomes more accurate but the room for introducing changes to the routings diminishes. The distinction between the different levels of re-routing measures can also be blurred, for instance, when re-routing measures are prepared a few weeks before the flights.

Research on ATFM problems started in the late 1980s. Odoni[4,5] defined the air traffic flow management problem area, identified some of the major issues in the field, and suggested the decision support needs, mostly based on the US situation. Leal de Matos and Ormerod[3] provide similar ground-clearing work for European ATFM. Research has concentrated on optimisation models for the allocation of ground delays. Most models are meant for the US case, with congestion limited to airports. They are intended for the tactical level, within the few hours before flight departure.

The progress of these models in terms of the cases they cover has been quite steady. Andreatta and Romanin-Jacur[6] address the case of one airport where congestion lasts for a single period of time. Terrab and Odoni[7] present an exact solution method for a case with one airport, several periods and deterministic capacity. Richetta and Odoni[8] provide a linear programming solution method to a multi-period single airport case where capacity is stochastic, and Vranas et al[9] present integer formulations for a network of airports, taking into account the interdependency between operations at different airports. The multi-airport problem has also been addressed by Navazio and Romanin-Jacur,[10] who provide algorithms to solve it. Andreatta and Brunetta[11] compare existing multiairport ground holding models computationally. Formulations have been developed to deal also with dynamic situations, when information on capacity changes over time.[12,13] Lindsay et al[14] and Tošic and Babic[15] provide a detailed survey of literature on the optimisation of ground delays.

Research has been reported on optimisation models where congestion also affects en-route sectors.[14,16] Helme[17] describes a multicommodity network flow formulation of this case. Vranas[18] proposes optimisation models for the allocation of ground delays in Europe at the tactical level. The models take into account that flights can cross several congested elements of the airspace. Bertsimas and Stock Patterson[19] present a model for the allocation of ground delays and speed control of airborne traffic. The model takes into account the en-route and airport capacities in the US and was tested on large scale problems of thousands of flights. This model is extended to address situations with stochastic capacity in Alonso et al.[20]

Optimisation models are often regarded with suspicion by users, possibly because they view their mathematical content as rather obscure. They can also be very time-consuming to execute. Andreatta et al[21] describe a heuristic for the allocation of ground delays which is based on priority rules. The heuristic is less time-consuming than exact models and easier to grasp by the users. The use of simulation and artificial intelligence techniques in ATFM is also documented in the literature, as explained in Leal de Matos and Ormerod.[3]

Research on decision support models and systems for re-routing flights at the strategic, pre-tactical and tactical levels of ATFM is still in its infancy. In 1995, a re-routing demonstrator and optimisation models for the re-routing of air traffic flows were presented as part of a project for the CFMU.[22] The re-routing of flights has also been included in optimisation models as a variation or part of a global ATFM problem. In Tošic et al[23] an integer programming model that allocates both ground delays and routes to individual flights is described. A research project aimed at developing decision support aids for re-routing flights, Computer Aided Route Allocation Tools (CARAT), took place between 1995 and 1997 at the Experimental Centre of EUROCONTROL.[24] CARAT provided the prototype of a tool to support the re-routing of individual flights. Bertsimas and Stock Patterson[19] propose variations to their ATFM model to support re-routing of soon to depart or already airborne flights.

In Leal de Matos and Ormerod,[3] it is emphasised that flow managers at the CFMU need tools to support re-routing decisions at the tactical, pre-tactical and strategic levels. The contribution of this paper is to present and evaluate the usability of three optimisation models for pre-tactical re-routings, that is, routings of air traffic flows a few days before the flights take place. The models differ in the level of detail and in the congestion cost functions used. Note that these models are not intended for real-time optimal assignment of aircrafts by air traffic controllers.

The next section discusses different modelling approaches to pre-tactical re-routings and explores different formulations and their impact on the size of the models. The third section describes the models developed and the fourth section focuses on the testing of the models and the results obtained. The fifth section discusses the usability and limitations of this research, and finally the sixth section proposes directions for future research.

Modelling approaches

A key issue in determining the effectiveness of re-routing measures is the degree of authority that flow managers at the CFMU can exercise.[3] At present, only some of the routings at the strategic level, or those in contingencies or in severely congested situations, are mandatory. All other re-routings tend to be advisory. Mandatory re-routing measures apply to flows, during certain periods and are usually negotiated beforehand with airline representatives and the area control centres involved; they cannot be imposed on an individual flight basis. The choice of air route for a particular flight is seen as a commercial decision to be taken by the airline, given air traffic control constraints.

However, there is an on-going debate on the adequacy of the present situation, and whether there should be more or less regulation.[3] Some stakeholders in flow management argue in favour of a firmer regulatory control, where responsibility for the provision of flight plans, including the flight route, lies with ATFM. In this case, airlines just file the airports of departure and destination, type of aircraft, number of passengers and state their preferences. The nascent research on optimisation models for re-routing measures is meant for situations when flow management does have the authority to route individual flights.[19,23]

In practice, at present, flow managers, when considering pre-tactical re-routing measures, group flights into flows, according to origin and destination airport areas. They then identify alternative routes for the flows and compare capacity with demand for ATC sectors, in an iterative way. The alternative routes have to be acceptable to airlines, that is, they cannot be too long or too costly. The modelling approach taken in this paper is based on this practice and assumes that flow managers have authority to issue re-routing measures applying to flows during a very well defined period, from a few hours to a day. Routes are not allocated on an individual flight basis.

Flights are grouped into flows according to their origin-destination, and the problem of re-routing air traffic flows is solved in two stages: (1) routes problem: identify acceptable and alternative routes for each flow; (2) assignment problem: given a set of flows, a set of acceptable routes and a set of capacity constrained sectors, assign a route to each flow so that the total cost of re-routings and congestion is minimised. The models presented in this paper address the second stage. This approach results in smaller, easier-to-solve models as it requires grouping the flights into flows before reaching the optimisation phase. However, if needed, flows can be further divided into smaller groups of flights. In fact, if the flow variables are replaced by flight variables the models here presented can also be formulated in terms of individual flights.

Formulations and size

When modelling the above mentioned assignment problem, there is a question of how to represent congestion. At least, two possibilities can be considered: (1) use penalties whenever traffic demand exceeds the capacity of an ATC sector, or (2) use ground delays to keep the demand within the capacity. Possibility (2) is justified by the fact that congestion results in ground delays, but it can lead to large-size integer problems. It should be stressed that, at this level of planning, ground delays are included in the problem just to support the decision on the re-routing of flows. The actual allocation of ground delays will be done by the CFMU computer system, TACT, on the day of the flights. Possibility (1) reduces substantially the size of the problems, but because it does not take into account the cumulative effect of capacity/demand imbalances over time it may underestimate congestion. Both possibilities are explored in this paper.

Models with ground delays have two types of decision variables: (1) variables assigning one route to each flow, (2) variables assigning ground delays (or departure time intervals) to flights. The first type of variables depends on the number of flows and the choice of routes available. The definition of the ground delay variables, given the large number of flights involved, is not immediate. If a binary variable is defined for each flight in a flow, on each route and time period, the number of variables easily reaches 100 000. This is an unmanageable number of variables for an integer programming model. For instance, suppose the following decision variables are defined:

$$y_{ijzt} = \begin{cases} 1 & \text{if flight z of flow i is delayed on route j at t} \\ 0 & \text{otherwise} \end{cases}$$

To illustrate the difference in scale, consider a simplified scenario, which involves only part of the French airspace, of 20 flows with 60 flights and 2 alternative routes each, and a period of 50 time intervals, the number of ground delay variables could total 120 000.

Another possibility is to model ground-delay variables in terms of ‘number of flights delayed’. For example, consider the number d_{ijt} of flights of flow i departing on route j in t and the number y_{it} of flights of flow i ground-delayed in t:

d_{ijt} number of flights of flow i departing on route j at t

y_{it} number of flights of flow i ground delayed at t.

The number of variables for the above scenario is reduced from 120 000 to 3000. A drawback of this approach is that the length of delays affecting the flights is not taken into account.

To overcome this drawback, variables can be defined in a more detailed way, using the number d_{ijtt}, of flights of flow i on route j that are scheduled to depart in t and are departing in t:

$d_{ijtt'}$ number of flights of flow i, on route j that are scheduled to depart at t and are departing at t'.

This formulation, which draws on a ground-delay model presented in Vranas,[18] results in 51 000 variables for the above simplified scenario, a large number, but still substantially smaller than the formulation with binary variables. The models with ground delays that are to be described are based on the last two sets of variables discussed above.

Models

Our models assign a route to each traffic flow in order to minimise the cost of congestion and re-routings. Three integer programming models resulting from different ways of measuring congestion and congestion costs are presented: BALDIST, DELINT1 and DELINT2.

- BALDIST. Congestion is measured in terms of penalty variables that are activated whenever traffic demand is above the capacity of an ATC sector. The model minimises the sum of the estimated cost of congestion and the cost of re-routing subject to capacity constraints and constraints on the assignment of routes to flows.
- DELINT1. Congestion is measured in terms of the number of ground delayed flights. The ground delay variables are included in the model to support the decision on re-routing and not to allocate ground delays to individual flights. Therefore, unlike in BALDIST, flights ground delayed can build up over time. The model minimises the sum of the estimated cost of ground-delay and the cost of re-routing subject to capacity and assignment constraints plus constraints defining and relating the two types of variables: assignment and ground delay variables.
- DELINT2. Congestion is measured in terms of more detailed ground delay variables than in DELINT1. This model takes into account not only the number of flights ground delayed, but also the length of the ground delay.

To formulate the three models the following is assumed:

1. All flights have identical cost functions. This assumption ensures that there is equity between flights, but means that the model does not represent actual flight costs. Ground delay and re-routing cost functions are in the model to account for different trade-offs between re-routing and ground delaying flights and to compare various re-routing scenarios.
2. All flights in a flow, that is, flights with the same origin destination, fly the same route at the same speed. The limitations of this assumption are attenuated by the fact that airlines tend to follow the same (cheapest) route and use the same type of aircraft for the same city-pairs. It should also be noted that the time intervals considered are long and the models are not detailed to the point of providing exact times for individual flights. Finally, at the pre-tactical level, ATFM does not have accurate information about individual flights.
3. The period of time for which flow re-routings are being considered is divided into p identical time intervals. These time intervals work as time units: the events ‘flight departure,’ ‘flight arrival’, ‘flight entry in sector’ are assumed to take place at the beginning of the corresponding time interval. Parameters like ‘the time it takes to get to a certain sector on a certain route’ are measured in ‘number of time intervals’. If a flight crosses two sectors in the same time interval, then the number of time intervals it takes to get to both sectors is the same and the crossing time for these sectors is 0 time intervals. This assumption is consistent with the way capacity of an air traffic control sector is defined for air traffic flow management: ‘number of flights per time interval.’

The following is a summary of notation that we are to use in the models.

m	total number of traffic flows
n	total number of flights
l	total number of sectors
p	total number of time intervals
i	index for traffic flows $(i = 1, \ldots, m)$
z	index for flights above sector capacity $(z = 1, \ldots, \overline{Z})$
$\overline{Z}$	maximum number of flights allowed to exceed capacity $(\overline{Z} \leqslant n)$
N_i	set of flights in flow i $\left(\sum_{i=1}^{m} \mid N_i \mid = n \right)$
k	index for sectors $(k = 1, \ldots, l)$
R_i	set of routes acceptable to flow i
j	index for routes ($j \in \bigcup_i R_i$); each route is defined to be a sequence of sectors $k, k', \ldots$ and a sequence of corresponding entry times $t_{jk}, t_{jk'}, \ldots$
L_k	set of routes that cross sector k
t	index for time intervals $(t = 1, \ldots, p + 1)$
t_{jk}	number of time intervals it takes to get from departure point to sector k on route j, which include the entry interval
τ_{jk}	number of time intervals it takes to cross sector k on route j, which excludes the entry interval $(\tau_{jk} = t_{jk'} - t_{jk}$, where k' is the sector that immediately follows sector k in route j)
$\overline{Y}$	limit on the number of flights ground delayed
$\overline{q}$	maximum number of ground delayed intervals allocated to a flight
u_{kt}	capacity of sector k during interval t

f_{it}	number of flights of flow i originally scheduled to depart in interval t $\left(\sum_{i=1}^{m} f_{it} = \lvert N_i \rvert, i = 1, \ldots, m\right)$
M_{it}	upper bound on the number of flights of flow i that eventually depart in interval t on a route, for example, in DELINT1 $M_{it} = \min\{\lvert N_i \rvert, f_{it} + \overline{Y}\}$
c_j	additional cost of route j for each flight
h_{zk}	marginal cost of the zth flight above capacity in sector k
g_{it}	cost of ground delay per time interval and per flow
$g(t)$	cost of ground delay for t time intervals

Variables

$$x_{ij} = \begin{cases} 1 & \text{if flow i is assigned to route j} \\ 0 & \text{otherwise} \end{cases}$$

$$o_{ztk} = \begin{cases} 1 & \text{if there is a zth flight above capacity in} \\ & \text{interval t in sector k} \\ 0 & \text{otherwise} \end{cases}$$

y_{it}	number of flights of flow i ground-delayed in interval t
d_{ijt}	number of flights of flow i on route j departing in interval t
$d_{ijtt'}$	number of flights of flow i on route j that originally were scheduled to depart in interval t and are ground delayed to depart in t′.

Model BALDIST

In this model we implicitly assume that the cost of the nth flight exceeding the capacity of an ATC sector is at least as big as that of the (n − 1)th flight.

$$\min w = \sum_{i=1}^{m} \sum_{j \in R_i} \sum_{t=1}^{p} c_j f_{it} x_{ij} + \sum_{z=1}^{\overline{Z}} \sum_{k=1}^{l} \sum_{t=1}^{p} h_{zk} o_{ztk} \qquad (1)$$

subject to

$$\sum_{i=1}^{m} \sum_{j \in (R_i \cap L_k)} \sum_{r=0}^{\tau_{jk}} f_{i,t-t_{jk}-r+1}\, x_{ij} - \sum_{z=1}^{\overline{Z}} o_{ztk} \leqslant u_{kt} \qquad (k = 1, \ldots, l; t = 1, \ldots, p) \quad (2)$$

$$\sum_{j \in R_i} x_{ij} = 1 \quad (i = 1, \ldots, m) \qquad (3)$$

$$x_{ij} \in \{0, 1\} \quad (j \in R_i; i = 1, \ldots, m) \qquad (4)$$

$$O_{ztk} \in \{0, 1\} \quad (z = 1, \ldots, \overline{Z};\ t = 1, \ldots, p;\ k = 1 \ldots, l) \qquad (5)$$

The objective (1) of the model is to minimise the total cost of re-routing and congestion. Expressions (2) are the capacity constraints for each ATC sector in each time interval and expressions (3) make sure that a flow is assigned to one and only one route.

Models with ground delays

In the two models with ground delays, we consider that the capacity of any sector in time interval p + 1, one just after the end of the period during which the re-routing measures apply, is infinite. In practice, this means that, in the time interval after the end of the re-routing, the difference between capacity and demand will be sufficiently large to allow the backlog of ground-delayed flights to depart.

Model DELINT1

$$\min w = \sum_{i=1}^{m} \sum_{j \in R_i} \sum_{t=1}^{p} c_j f_{it} x_{ij} + \sum_{i=1}^{m} \sum_{t=1}^{p} g_{it} y_{it} \qquad (6)$$

subject to

$$\sum_{i=1}^{m} \sum_{j \in (R_i \cap L_k)} \sum_{r=0}^{\tau_{jk}} d_{ij,t-t_{jk}-r+1} \leqslant u_{kt} \qquad (k = 1, \ldots, l; t = 1, \ldots, p) \quad (7)$$

$$\sum_{j \in R_i} x_{ij} = 1 \quad (i = 1, \ldots, m) \qquad (8)$$

$$d_{ijt} \leqslant M_{it}\, x_{ij} \quad (j \in R_i; i = 1, \ldots, m; t = 1, \ldots, p + 1) \qquad (9)$$

$$\sum_{j \in R_i} d_{ijt} = f_{it} + y_{i,t-1} - y_{it} \qquad (i = 1, \ldots, m;\ t = 1, \ldots, p + 1) \quad (10)$$

$$x_{ij} \in \{0, 1\} \quad (j \in R_i; i = 1, \ldots, m) \qquad (11)$$

$$0 \leqslant y_{it} \leqslant \overline{Y} \text{ and integer} \quad (i = 1, \ldots, m; t = 1, \ldots, p) \qquad (12)$$

$$y_{i0} = 0, \quad y_{i,p+1} = 0 \quad (i = 1, \ldots, m) \qquad (13)$$

$$d_{ijt} \geqslant 0 \text{ and integer} \qquad (j \in R_i; i = 1, \ldots, m; t = 1, \ldots, p + 1) \quad (14)$$

The objective (6) of the model is to minimise the total cost incurred in re-routing plus the aggregated cost of ground delays. The cost of ground delays is considered to be independent of the length of ground delay. Expressions (7) make sure that all the flights present in a sector in a certain time interval do not exceed the capacity of that sector. Expressions (9) ensure that flights do not depart on routes that have not been assigned to their flow. Expressions

(10) state that the total flights of a flow departing in a time interval t is equal to the total flights of that flow scheduled to depart in t plus the flights ground delayed in (t − 1) minus the flights to be ground delayed in t. Expressions (12) define the ground delay variables as integer and impose an upper limit on the number of flights of a flow ground delayed. Expressions (13) say that there is no ground delay at the beginning of the decision period and that all possible ground delay backlog is cleared up right after the end of the decision period.

Model DELINT2

$$\min\ w = \sum_{i=1}^{m} \sum_{j \in R_i} \sum_{t=1}^{p} c_j f_{it}\, x_{ij} + \sum_{i=1}^{m} \sum_{j \in R_i} \sum_{t=1}^{p} \sum_{t'=t+1}^{t+\bar{q}} g(t'-t) d_{ijtt'} \tag{15}$$

subject to

$$\sum_{i=1}^{m} \sum_{j \in (R_i \cap L_k)} \sum_{r=0}^{\tau_{jk}} \sum_{t=1}^{t'-t_{jk}-r+1} d_{ijt,t'-t_{jk}-r+1} \leqslant u_{kt'} \quad (k = 1, \ldots, l;\ t' = 1, \ldots, p) \tag{16}$$

$$\sum_{j \in R_i} x_{ij} = 1 \quad (i = 1, \ldots, m) \tag{17}$$

$$\sum_{t'=t}^{t+\bar{q}} d_{ijtt'} \leqslant M_{it}\, x_{ij} \quad (j \in R_i;\ i = 1, \ldots, m; t = 1, \ldots p) \tag{18}$$

$$\sum_{j \in R_i} \sum_{t'=t}^{t+\bar{q}} d_{ijtt'} = f_{it} \quad (i = 1, \ldots, m;\ t = 1, \ldots, p) \tag{19}$$

$$x_{ij} \in \{0, 1\} \quad (j \in R_i;\ i = 1, \ldots, m) \tag{20}$$

$$d_{ijtt'} \geqslant 0 \text{ and integer} \quad (j \in R_i; i = 1, \ldots, m; t = 1, \ldots, p; t' = t, \ldots, t + \bar{q} \leqslant p + 1). \tag{21}$$

The objective of the model is again to minimise the cost incurred in re-routing and the estimated cost of ground delays, taking into account the length of delays. As in the previous model, expressions (16) and (17) are respectively the capacity constraints on ATC sectors and the constraints on the assignment of routes to flows. Expressions (18) ensure that no flights depart on routes which their flows do not use and expressions (19) ensure that the number of flights departing equals the number of flights scheduled.

Testing the models

The set of data used to test the models is based on the actual traffic crossing the whole French upper airspace on 25/04/96, from 03:00 h to 22:00 h, totalling 3582 flights. The French airspace was chosen because it is at the crossroads of the European airspace, with approximately 25% of the whole of ECAC traffic, and many of its sectors are often congested. The period from 03:00 h to 22:00 h is similar to the periods to which some re-routing control measures apply.

The French-controlled airspace is divided into five regions under the responsibility of five air traffic control centres: Aix, Bordeaux, Brest, Paris and Reims. Each region is, in turn, broken up into sectors. Sectors can have different configurations, for instance, two contiguous sectors can be merged for a certain period of the day if the traffic is expected to decrease or if there are fewer air traffic controllers. The set of possible configurations is pre-defined at the strategic level. For this test, the configuration chosen is the one with the largest capacity that was available on 25/04/96 with 41 sectors.

At present, ATC capacity for ATFM is defined hourly. Therefore, hourly time intervals were used in this test. The period from 03:00 h to 22:00 h is divided into 19 hourly time intervals and in the case of DELINT1 and DELINT2 with an additional 20th interval to allow delayed flights to depart.

To limit the number of overcapacity variables in BALDIST, the number of flights above the capacity of a certain sector during a time interval cannot exceed 19, the lowest number for which the problem becomes feasible. Note that only one of the 41 sectors (sector AO) is constrained by this limit. Overcapacity per time interval in all other sectors is significantly below this value.

In DELINT2, flights can be delayed a maximum of 4 time intervals. Note that the time intervals could be shorter. However, in ATFM practice, at the moment, there is practically no data available on capacity for intervals finer than one hour, especially at the pre-tactical and strategic levels.

As a result of considering part of the airspace, it is implicitly assumed that all other airspace elements not considered in the test (airports, lower airspace, neighbouring airports, etc) do not have capacity constraints. The effect of this assumption is attenuated by the fact that the French airspace is one of the main bottlenecks of the European airspace.

The 138 flows identified are sorted by flow group. Flows belong to the same group if the portion of their routes crossing the French upper airspace is the same. The flights considered for re-routings totalled 920, which is approximately 26% of the traffic on that date. The best route for each flow was obtained from the flight plans, selecting the most frequently filed route on that date. The alternative routes, depending on the choice of routes filed, were either obtained from the flight plans or determined by calculating the distance. The flying times were also obtained from the flight plans. Only those routes acceptable to airlines, that is, routes whose flying time is not significantly larger than the flying time of the best routes (in this test, extra flying time less than or equal to 30 minutes), were chosen.

The cost of the alternative routes is an estimate of the fuel cost incurred with the re-routing, by flying longer or at a lower altitude. The cost c_j of a route j is calculated in the following way:

$$c_j = 10\alpha_j + \mu\Delta_j$$

where μ represents the cost of a minute of extra flying time and is given by

$$\mu = \begin{cases} 10 & \text{if } \Delta_j \leqslant 15 \\ 100 & \text{if } 15 < \Delta_j \leqslant 30. \end{cases}$$

α_j denotes the number of minutes flying in sectors lower than the sectors in the best route and Δ_j is the additional flying time of route j.

The routes chosen for each flow are described in Leal de Matos.[25]

The cost of congestion is represented differently depending on the model. For BALDIST the following cost function is used:

$$h_{zk} = 400\, z^2 (\forall k)$$

In DELINT1, the cost of ground delay varies with the time interval and flow and is calculated in the following way:

$$g(i, t) = a(1 + m(t)/\max(1, f_{it}))$$

where a is a constant representing the basic cost of delay, which is set at $a = 147$, approximately the same cost of 15 minutes of extra flying time. In addition, m(t) is obtained as follows:

$$m(t) = \sum_{k=1}^{l} \max\left(0, \sum_{i=1}^{m} \sum_{j\in(R_i\cap L_k)} \sum_{r=0}^{\tau_{jk}} f_{i,t-t_{jk}-r} - u(k, t)\right) \qquad (t = 1, \ldots, p)$$

The idea behind it is that, in a congested situation, the heavier a flow is, the more acceptable it becomes to delay it.

The cost of ground delay in DELINT2 varies with the number of delay time intervals in the following way:

$$g(t' - t) = a(t' - t)^2$$

where a is a basic cost of ground delay, and is set at $a = 2500$, which is equal to the cost of 25 minutes of extra flying time, t' is the departure time interval and t is the time interval in which the flight was scheduled to depart with $t < t'$.

Table 1 shows the resulting size for each of the models, BALDIST, DELINT1 and DELINT2. The congestion variables in BALDIST depend on the upper limit on the number of flights that can be above capacity during a certain interval (19 in this case) and are determined once the flows are assigned to routes. Some variables are dependent, for example, in DELINT1: $d_{ijt} = (f_{it} + y_{i,t-1} - y_{it})x_{ij}$.

Results

The models were solved using the modelling system GAMS 2.25 coupled with the integer programming package LAMPS 1.66 on a SUN/SPARC workstation. Table 2 shows a summary of the results. For each model, two situations are compared: one where all flights take the best routes, and another where flows are re-routed. The difference in cost between the solutions obtained for each situation is substantial.

The gap between the optimal value of the linear relaxation and the optimum (integer) value is very small in BALDIST. This might be due to the format of the constraints matrix, with most variables having coefficients 1 or -1 (for an explanation see Williams[26]). Consequently, the branch-and-bound (B&B) search (measured in number of B&B nodes) and the execution time are very short (see Table 2). The first optimisation runs of DELINT1 showed that it is hard to solve to optimality, therefore to limit the execution time, the best solution obtained after 10 000 iterations, corresponding to approximately 8 minutes of CPU time, was taken. However, the difference between the value of this solution and the lower bound provided by the linear relaxation of the model is very small: 0.71%, meaning that the value of this solution is, at most, 0.71% away from the optimum value. DELINT2, despite being a larger model than DELINT1, appears to be less hard to solve to optimality: the solution presented is optimum and was obtained in very little time.

Table 1 BALDIST, DELINT1 and DELINT2-problem size

	BALDIST	DELINT1	DELINT2
Capacity constraints	779	779	779
Assignment constraints	138	138	138
Flights have to be routed onto flow route	—	6880	6536
Relation flights scheduled/departing	—	3580	3401
Total constraints	917	11377	10854
Assignment variables	303	303	303
Congestion variables	14801	—	—
Ground-delay variables	—	3401	—
Departure variables	—	6880	30616
Total variables	15104	10584	30919

Table 2 Summary of results

	BALDIST	DELINT1	DELINT2
(1) Cost if all flights take best routes	7834 400	12 323 675	9685 000
(2) Best value after 10 000 iterations		7181 207	
(3) Optimum value	4086 660	—	3522 140
Variation between (1) and (2) or (3)	−47.84%	−41.73%	−63.63%
(4) Linear relaxation-optimum value	4086 036	7130 246	3515 978
Variation between (4) and (2) or (3)	0.02%	0.71%	0.17%
Flows re-routed	36	34	38
Flights re-routed	346	312	351
(5) Flights ground delayed if all take best routes	—	2154	1273
(6) Flights ground delayed with re-routings	—	1011	560
Variation between (5) and (6)	—	−53.06%	−56.01%
Total ground delay if all flights take best routes (minutes)	—	129 240	116 760
Total ground delay with re-routings	—	60 660	46 620
CPU time (s)	33	466	281
Number of B&B nodes	20	368	82

Table 3 Flights re-routed

Flow group	BALDIST	DELINT1	DELINT2	Flights re-routed by all models
UK to Balearics, Barcelona and Alicante	0 (0)	6 (1)	7 (1)	0 (0)
Germany (exc., West) and Switzerland To Balearics and Barcelona	2 (1)	2 (1)	2 (1)	2 (1)
W. Germany to Balearics and Barcelona	5 (1)	5 (1)	7 (2)	5 (1)
Barcelona and Balerics to West Germany	0 (0)	2 (1)	2 (1)	0 (0)
Barcelona and Alicante to Brussels and Amsterdam	8 (2)	9 (2)	0 (0)	0 (0)
Barcelona, Balearics and Alicante to Germany and Switzerland	3 (1)	0 (0)	3 (1)	0 (0)
Barcelona, Balearics and Alicante to UK	12 (4)	6 (2)	19 (4)	6 (2)
Madrid to Frankfurt and Stuttgart	8 (2)	2 (1)	2 (1)	2 (1)
Madrid to Southeast Germany and Switzerland	0 (0)	0 (0)	0 (0)	0 (0)
Madrid to West Germany	0 (0)	0 (0)	0 (0)	0 (0)
Athens and Rome to Lisbon and Madrid	0 (0)	0 (0)	0 (0)	0 (0)
North Italy to Lisbon and Madrid	0 (0)	0 (0)	0 (0)	0 (0)
Lisbon and Madrid to Athens and Rome	0 (0)	0 (0)	0 (0)	0 (0)
Lisbon and Madrid to North Italy	0 (0)	0 (0)	0 (0)	0 (0)
UK (exc. London), Brussels and Amsterdam to Switzerlamd	15 (1)	0 (0)	0 (0)	0 (0)
London to Switzerland	21 (1)	21 (1)	21 (1)	21 (1)
Switzerland to Brussels and Amsterdam	0 (0)	6 (1)	6 (1)	0 (0)
Geneva to UK	18 (2)	18 (2)	18 (2)	18 (2)
Zurich to UK	21 (1)	0 (0)	21 (1)	0 (0)
UK to Italy	4 (2)	1 (1)	3 (3)	0 (0)
Italy to UK	4 (2)	4 (2)	2 (1)	2 (1)
Paris to Italy	57 (5)	67 (7)	67 (7)	57 (5)
Italy to Paris	7 (1)	0 (0)	0 (0)	0 (0)
Paris to Toulouse	0 (0)	0 (0)	0 (0)	0 (0)
Paris to Marseilles and Nice	82 (2)	82 (2)	82 (2)	82 (2)
Toulouse to Paris	30 (1)	30 (1)	30 (1)	30 (1)
Brussels, Amsterdam and West Germany to Madrid and Malaga	4 (1)	6 (2)	4 (1)	4 (1)
Germany (exc. West) to Madrid and Malaga	16 (4)	16 (4)	16 (4)	16 (4)
UK to Madrid and Malaga	29 (2)	29 (2)	29 (2)	29 (2)
South Germany to Canary Islands	0 (0)	0 (0)	0 (0)	0 (0)
Germany (exc. South) to Canary Islands	0 (0)	0 (0)	10 (1)	0 (0)
UK to Canary Islands	0 (0)	0 (0)	0 (0)	0 (0)
Total flights re-routed	346 (36)	312 (34)	351 (38)	274 (24)

Table 3 shows the number of flights re-routed by each model. The number of flows re-routed is also shown between brackets. Note that the decision variables are flows not individual flights. There are 24 flows (274 flights) that are re-routed whichever of the three models is used. These re-routings move flows from more congested sectors to less congested ones. For instance, the flows Italy to UK are repeatedly re-routed from the best route, a route that crosses a congested sector, UFXF, to a route crossing the less congested sectors UH and UE. Some flows are not re-routed whichever model is used, such as the North Italy to Lisbon and Madrid flows, because they cross sectors that clearly are not congested on this date.

Figure 1 shows expected congestion in terms of number of flights above capacity as a percentage of capacity by each sector. Four situations are depicted, if all flights take the best routes, and if flights are re-routed using BALDIST, DELINT1 or DELINT2, respectively. In the cases of DELINT1 and 2, all flights above capacity are ground delayed. Sectors for which there are no flights above capacity expected are not shown in the figure.

The number of flights above capacity per time interval is shown in Figure 2. Time intervals for which there are not flights above capacity such as 3 h and 4 h are not included. As expected, flow re-routings smooth congestion peaks and, in general, there is not much difference between DELINT2, DELINT1 and BALDIST in smoothing congestion peaks both by sector and time interval. Sectors such as Z1N1H1 and AO are still significantly congested after the flow re-routings.

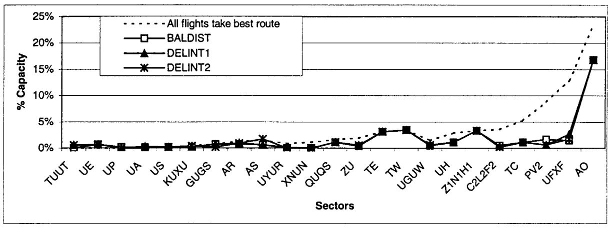

Figure 1 Flights overcapacity/ground delayed by sector.

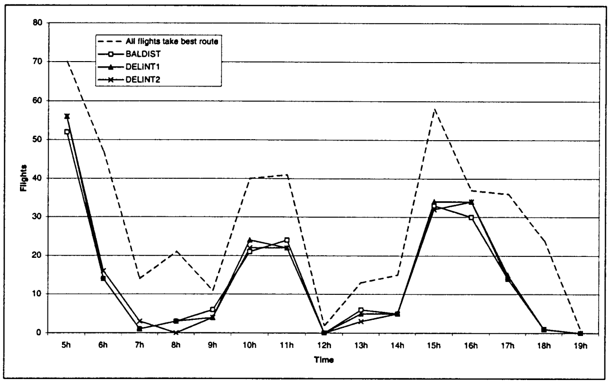

Figure 2 Flights overcapacity/ground delayed by time interval.

To assess the quality of BALDIST routings in terms of ground delay, DELINT2 was run using the flow re-routings obtained from BALDIST. DELINT2 provides significantly less ground delay than DELINT1 and distributes the ground delay more evenly between the flows, as shown in Table 4 and Figure 3. It turns out that DELINT2 with BALDIST routing scheme results in ground delay almost as small as DELINT2.

Discussion

The analysis of results leads to conclusions on two inter-related levels: on the usefulness of re-routing control measures and on the usefulness of the optimisation models. Re-routing control measures appear to be useful if there are imbalances in the distribution of congestion, and if the range of flows considered for re-routing is adequate to re-distribute congestion. Analysing the results of this test, it is clear that, after a certain point, given the flows considered for re-routing and the reductions in congestion already made, re-routing control measures have very little effect in reducing congestion. It is also apparent that there are some sectors whose very severe congestion peaks can be attenuated but not eliminated by applying re-routing control measures. This suggests the need to increase sector capacity during the peak hours to address the more serious and persistent congestion problems. The effect of increasing capacity during the peak hours could be tested and its cost compared with the savings in congestion costs.

The test indicates that the optimisation models can be of use in re-routing flows and can provide savings in ground delays. In the cases studied, re-routings reduced total ground delay by more than 50% and produced cost savings (cost of congestion + cost of re-routings) of more than 40%. However, these results should be seen in context: they compare with an extreme situation where all flights take the best routes, irrespective of the congestion situation. In a real environment, both airlines and flow managers will take action to limit the extent of ground delays. For instance, some airlines will re-route flights to bypass congested elements of the airspace. Further evaluation of the models is needed in a dynamic environment to assess more fully their impact on congestion. An important question is whether these re-routings reduce the need for re-routing individual flights or for slot allocation regulations at the tactical level.

There are 274 flights that are re-routed in every model. There are also flow groups that have, at least, one flow re-routed whichever model is used. In a context where these models are used daily, these results raise concerns over equity between airspace users, even though the traffic will vary according to the day of the week. Possibilities of addressing this inequity are to include the number of

Table 4 Ground-delay (minutes) of BALDIST, DELINT1 and DELINT2

	BALDIST routing scheme	DELINT1	DELINT2
North Italy to Lisbon and Madrid	60	0	60
West Germany to Balearics and Barcelona	60	0	0
Germany (exc. West) and Switzerland To Balearics and Barcelona	60	0	0
Barcelona and Alicante to Brussels and Amsterdam	120	0	0
Barcelona, Balearics and Alicante to UK	60	0	0
London to Switzerland	0	0	60
Switzerland to Brussels and Amsterdam	180	0	60
Geneva to UK	60	0	60
UK to Balearics, Barcelona and Alicante	360	0	120
UK (exc. London), Brussels and Amsterdam to Switzerland	180	0	120
Zurich to UK	0	0	120
Brussels, Amsterdam and West Germany to Madrid and Malaga	120	0	120
Germany (exc. South) to Canary Islands	300	0	240
UK to Canary Islands	180	0	240
Paris to Italy	180	0	0
Paris to Marseilles and Nice	0	0	300
Paris to Toulouse	240	0	360
Italy to UK	120	0	420
Toulouse to Paris	300	0	600
Madrid to Frankfurt and Stuttgart	0	60	60
Italy to Paris	10080	31440	10380
Others	34740	29160	33300
Total	47400	60660	46620

times a flow has been re-routed in the cost of re-routing a flow or to exclude a priori those flows that have been re-routed previously.

As shown in Table 2, BALDIST is undoubtedly the smallest and fastest model. It provided the optimum solutions in approximately 30 s. DELINT1 is the hardest model to solve of the three, since it was not possible to obtain the optimum solution. However, it provided a feasible solution whose value was less than 0.8% away from the optimum value. Despite being substantially larger than DELINT1, DELINT2 appears to be easier to solve: the optimum solution was reached in less than 10 min. Considering the impact on congestion, the execution time and the size, BALDIST appears to be the most efficient model of the three. The comparisons in the previous section show that BALDIST results in an alleviation of congestion, both in terms of capacity-demand imbalances and ground delay, which is almost the same as DELINT2. DELINT2 has the advantages of taking into account the length of ground delay affecting flights and providing more detailed information on the impact of the re-routing control measures. DELINT1 is out-performed by both the other models in terms of impact on congestion and execution time.

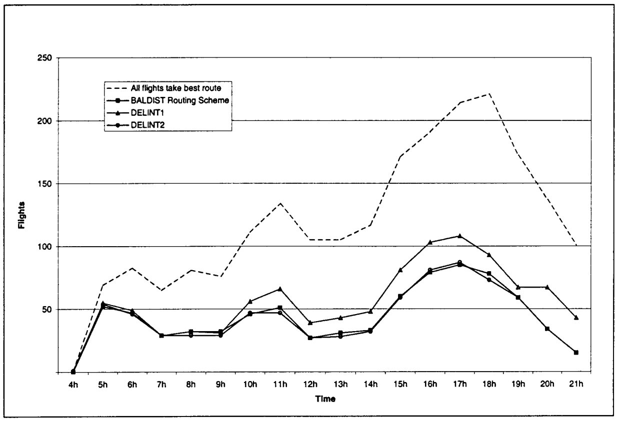

Figure 3 Flights ground delayed by time interval.

As suggested by Odoni,[27] the optimisation models here presented can also be used as ‘benchmarks’, against which various re-routing possibilities are evaluated. It is often not feasible to re-route all the flows in an optimum solution. It becomes important, therefore, to evaluate the impact of re-routing fewer flows and to identify subsets of flows that provide good approximate solutions.

Usability and limitations

This section discusses the usability and limitations of the models and their testing. The testing of the models indicates that they can reduce air traffic control delays and help deciding which flows to re-route. The models were designed to meet the queries likely to be made in pre-tactical ATFM: the flows to be re-routed. They can also allow for changes to features such as traffic demand, capacity of sectors, constraints on which flows to re-route, costs of delay, costs of re-routings and number of time intervals. However, other features, such as the definition of sectors, routes and flows to be included in the re-routing control measures are not easily changed. In fact, some of these operations require expertise provided either by a flow manager and/or a system with some degree of 'intelligence'.

The execution time of the models is not critical because they are aimed at pre-tactical ATFM that is, planning that takes place in the two days before the flights. Another possible use of the models is in contingency situations (eg radar breakdown) to suggest re-routings so that flights bypass a certain part of the airspace, in which case execution time becomes important. The three models provided either optimal solutions or solutions whose value was less than 0.8% away from the optimum value in less than 8 min. Other research in scheduling of aircraft has highlighted the value of using solution methods that provide good approximate solutions (see, for instance, Beasley et al[28]). Comparing the models, the test suggests that BALDIST, the simplest model, is the most efficient of the three in terms of size, execution time and quality of the solutions. DELINT1, in spite of being smaller than DELINT2, proved harder to solve to optimality and its use is not recommended. DELINT2 offers the best results in alleviating congestion, and provides more information than the other models, but it is a large model thus requiring much more space than DELINT1 and BALDIST and more time than BALDIST to be solved. The importance and relevance of these differences in performance will depend on the optimisers and computer resources available and on how close to the optimum value the solutions need to be. Note that the models were run on a SPARC station; if they ran on a more powerful computer the execution time would be shorter.

The models, in their present form, re-route flows not taking into account how many times they have been re-routed before and penalising the flows which comprise more flights (however, the cost function used in the model can reflect how many times a flow or flight has been re-routed).

The following limitations can be identified in the test of the models:

- The models were tested on the French upper airspace. It was implicitly assumed that there was no congestion in the remainder of the ECAC airspace. The test was limited to one, albeit typical, day of traffic.
- In preparing the data for testing the models, decisions were made in defining the scope of the optimisation problem, for example in grouping flights into flows, selecting the routes or in setting limits to ground delays and sectors overcapacity. These decisions affected the results obtained.
- The cost functions used in the models were based on the literature, on published operational costs of aircraft and on the interviews with four airlines. The input from airlines consisted of examples of trade-offs between ground delays and re-routings, and information on overall costs. To ensure equity between flights, it was assumed that all flights had identical ground delay and re-routing cost functions.

Ultimately, an underlying limitation of the test is that the models were tested using a simplified and static representation of 'the real world'.

Directions for future research

The following directions for future research are suggested:

- To test the models presented here on an even larger airspace that includes several European countries, or even the whole ECAC airspace.
- To investigate further how to define the scope of the optimisation problem regarding what sectors are to be considered and what flows are relevant and how they should be defined. Sensitivity analysis, which addresses the impact of decisions made in the definition of the scope of the optimisation problem on the results, should be carried out. This includes sensitivity to upper limits on ground delays and sector overcapacity.
- To fine-tune the cost functions used in the optimisation models in consultation with aircraft operators and ATC costing centres. Inclusion of priority indices in the cost function of optimisation models to improve equity between airspace users should be investigated. The impact of different cost functions on equity should be examined.
- To investigate other definitions of equity. This research applied the definition of equity in use at the CFMU, equity between flights. Other definitions of equity should be investigated, such as equity between aircraft operators or between passengers. The question of how acceptable these are to the various stakeholders in ATFM should be addressed.
- To forecast traffic demand and capacity of en-route ATC sectors as stochastic variables. At present, they are considered deterministic for planning purposes.

Abbreviations

ATC	Air Traffic Control
ATFM	Air Traffic Flow Management
CARAT	Computer Aided Route Allocation Tools
CFMU	EUROCONTROL Central Flow Management Unit
DSS	Decision Support System
ECAC	European Civil Aviation Conference
EUROCONTROL	European Organisation for the Safety of Air Navigation

Acknowledgements—We would like to thank everyone at the CFMU for their contribution to this research and particularly to Francis Gainche who supervised part of the research. Marcel Richard and all other Users' Requirements Section and Central Executive Unit staff provided much needed information, data and comments. This research was partly supported by Programa Praxis XXI/JNICT (Portugal).

References

1 EUROCONTROL (1999). Delays to Air Transport in Europe. Annual Report 1998.

2 Flight International (1999). Euro ATC delays could rise by 70%. 13–19 October: 10.

3 Leal de Matos P and Ormerod R (2000). The application of operational research to European air traffic flow management. Eur J Opl Res 123: 125–144.

4 Odoni AR (1987). The flow management problem in air traffic control. In: Odoni AR, Bianco L, Szegö G (eds). Flow Control of Congested Networks. NATO ASI Series Series F: Computer and Systems Science 38. Springer-Verlag: Berlin, pp 269–288.

5 Odoni AR (1994). Issues in air traffic flow management. In: Proceedings of Conference on Advanced Technologies for Air Traffic Flow Management. Deutsche Forschungsanstalt für Luft- und Raumfahrt e.V. (DLR): Bonn.

6 Andreatta G and Romanin-Jacur G (1987). Aircraft flow management under congestion. Trans Sci 21: 249–253.

7 Terrab M and Odoni AR (1993). Strategic flow management for air traffic control. Opns Res 41: 138–152.

8 Richetta O and Odoni AR (1993). Solving optimally the static ground-holding policy problem in air traffic control. Trans Sci 27: 228–238.

9 Vranas PBM, Bertsimas DJ and Odoni AR (1994). The multi-airport ground-holding problem in air traffic control. Opns Res 42: 249–261.

10 Navazio L and Romanin-Jacur G (1998). The multiple connections, multi-airport ground holding problem: models and algorithms. Trans Sci 32: 268–276.

11 Andreatta G and Brunetta L (1998). Multiairport ground holding problem: a computational evaluation of exact algorithms. Opns Res 46: 57–64.

12 Richetta O (1995). Optimal algorithms and a remarkably efficient heuristic for the ground-holding problem in air traffic control. Opns Res 43: 758–770.

13 Vranas PBM, Bertsimas DJ and Odoni AR (1994). Dynamic ground-holding policies for a network of airports. Trans Sci 28: 275–291.

14 Lindsay K, Boyd EA and Burlingame R (1993). Traffic flow management modeling with the time assignment model. Air Traffic Control Q 1: 255–276.

15 Tošic V and Babic O (1995). Air route flow management—problems and research efforts. Transp Plan Technol 19: 63–72.

16 Tošic V, Babic O, Cangalovic M and Hohlacov D (1995). Some models and algorithms for en route air traffic flow management. Trans Plan Technol 19: 147–164.

17 Helme M (1992). Reducing air traffic delay in a space-time network. IEEE Int Conf Syst, Man Cybern 1: 236–242.

18 Vranas PBM (1996). Optimal slot allocation for European air traffic flow management. Air Traffic Control Q 4: 249–280.

19 Bertsimas D and Stock Patterson S (1998). The air traffic flow management problem with enroute capacities. Opns Res 46: 406–422.

20 Alonso A, Escudero LF and Ortuño MT (2000). A stochastic 0-1 program based approach for the air traffic flow management problem. Eur J Opl Res 120: 47–62.

21 Andreatta G, Urunetta L and Guastalla G (1998). Flow management problem: recent computational algorithms. Control Engng Prac 6: 727–733.

22 Matos P (1995). Re-routing Study. Internal Report, Central Flow Management Unit-EUROCONTROL.

23 Tošic V, Babic O, Cangalovic M, and Hohlacov D (1995). A Model to Solve En Route Air Traffic Flow Management Problem: A Temporal and Spatial Case. Internal report, University of Belgrade.

24 Loubieres P (1996). Computer Aided Route Allocation Tools (CARAT)—Phase 1: Modelisation and Algorithms. EUROCONTROL Experimental Centre Internal Report.

25 Leal de Matos P (1998) The Development of Decision Support Models for European Air Traffic Flow Management. PhD Thesis, University of Warwick, Great Britain.

26 Williams HP (1993). Model Solving in Mathematical Programming. Wiley: Chichester.

27 Odoni AR (1996). Flow management in transition. Air Traffic Control Q 4: 225–227.

28 Beasley JE, Sonander J, and Havelock P (2001). Scheduling aircraft landings at London Heathrow using a population heuristic. J Opl Res Soc 52: 483–493.

Received December 1999;
accepted May 2001 after two revisions