

Fairness in the Aircraft Landing Problem

M.J. Soomer^{*,a,b}, G.M. Koole^a

^a VU University, De Boelelaan 1081a, 1081 HV Amsterdam, the Netherlands

^b National Aerospace Laboratory NLR, Anthony Fokkerweg 2, 1059 CM Amsterdam, the Netherlands

* email: MJ.Soomer@few.vu.nl

June 9, 2008

Abstract

The last few decades a lot of attention has been paid to Collaborative Decision Making (CDM) in Air Traffic Management. A lot of models are proposed in which information of stakeholders (mainly airlines) is used in some way to improve the decision process. When doing this, fairness becomes an important issue. However, it is not always clear what will be considered fair by the airlines involved.

In this paper we will use the aircraft landing problem to illustrate various definitions of fairness, that stem from the use of airline preferences. In this problem, a landing order and feasible landing times have to be determined for a set of flights at a runway. The airlines cost and the various definitions of fairness are used as objective for the problem. Local search heuristics are introduced to solve these formulations.

Numerical experiments using schedule data from a large European airport are used to evaluate how the fairness definitions and heuristics behave for real life problems. The results show that it is possible to achieve more fairness, while still obtaining considerable cost savings compared to the First Come First Served schedule. The heuristics obtain reductions of up to 26% in the root mean square deviation of the average cost per airline. The different heuristics show that schedules with different trade-offs between efficiency, cost, delay and fairness can be obtained. Hopefully, these results can be a starting point for discussing the fairness issues related to the introduction of CDM processes among air traffic stakeholders.

1 Introduction

The last few decades Collaborative Decision Making (CDM, [11], [12]) has become an important concept in Air Traffic Management. Several models are proposed in which information of stakeholders (mainly airlines) is used to improve the decision process. Air Traffic Management authorities need to make sure that this done in a manner such that the outcome is fair to all airlines. In this paper we will discuss fairness considering the aircraft landing problem.

Airport runway capacity is unpredictable and subject to large changes during operations, mainly because of weather and visibility conditions. The demand for runway capacity at large hub airports is subject to large peaks, because of the hub and spoke network that most airlines use. This often leads to imbalance between capacity and demand on the day of operations. Decisions how to assign scarce landing capacity among competing airlines should be made in a manner that is considered fair by all airlines involved.

Runway landing capacity depends on the weight categories of the aircraft and their landing sequence. This is because the separation that is required between consecutive aircraft depends on the weight categories of the aircraft. A light aircraft flying behind a heavy aircraft requires more separation than the reverse order. However, in practice not much is done actively sequencing flights. Flights approaching the airport are under the guidance of the so-called approach controller from less than 30 minutes before the actual landing. From this moment on, the controller must create a correctly separated flow of aircraft toward the runway. Because of the limited time

available and the high workload of the controllers, hardly any changes to the sequence of the flights can be made at this point.

This leads to the idea to consider runway sequencing at a tactical level (starting a few hours before the actual landings). The problem to determine a landing sequence and landing times for flights at a runway is known as the aircraft landing problem. The landing times are constrained to be within predefined time windows and to allow for the required separation between the flights.

Focusing only on efficient sequences can result in unacceptable delays for individual flights. This could have a large impact on airlines and their passengers, e.g., missed transfers and crew or aircraft scheduling problems. However, flight delay is a poor measure of the impact for airlines and passengers, because the impact of a certain delay might differ from flight to flight depending on, amongst others, the number of (transfer) passengers. This leads to the idea of considering airline preferences for individual flights when determining the runway schedule. By doing this, it is expected that the impact of delays for airlines and their passengers can be reduced and the airlines need less schedule recovery measures.

In another paper [17] we introduced a model for the aircraft landing problem that considers airline preferences. Airlines define a cost function for each arriving flight, which relates the landing time of the flight to costs. The resulting cost functions parameters will be provided to the air traffic control organisation and will be used to generate arrival schedules with minimal total airline cost. Computational experiments show that the total cost savings are tremendous.

It is also the role of air traffic control to assure that air traffic proceeds in an equitable manner. Consequently, taking airlines' operating costs into consideration, must be done in a manner that does not show favor to some airlines at the expense of others. In [17], we propose a scaling mechanism that scales the cost functions in such a way that the average delay cost per minute per flight are equal for all airlines. At the same time the cost ratio between flights of the same airline are preserved. Now the problem is solved minimizing the sum of the scaled cost. The scaling mechanism makes it possible to compare the flights regardless of their airline in a fair way. However, this is no guarantee that the scaled cost and/or delay is shared equally among the airlines. Indeed, our computational experiments showed large differences between the costs and savings of different airlines.

In this paper we aim at modeling fairness more explicitly. It is not trivial to define and measure fairness. We will introduce various definitions of fairness and incorporate these in our formulation of the aircraft landing problem.

These definitions are presented in section 2. In section 3 we will present a MIP formulation of our model. In section 4 an example is given that shows the unequal spread of cost and delay in the minimum total cost schedule and shows schedules that illustrate our definitions of fairness. Local search heuristics focused on finding fair solutions are introduced in section 5.

A large number of problem instances, created from real-life data concerning arrivals during a week at a major European hub, were used to perform computational experiments. The results are presented in section 6.

The resulting schedules will be compared to the the schedule that is obtained without actively sequencing the flights. In the latter schedule the flights are landed (respecting the required separation) in the order in which they approach the airport. Such a schedule resembles current practice. This schedule will be called the First Come First Served (FCFS) Schedule.

The results show that it is possible to improve the fairness compared to both the FCFS reference schedule and the minimum total cost schedule. The heuristics obtain reductions of up to 26% in the root mean square deviation of the average cost per airline. The percentage of airlines that is worse off than in the FCFS schedule is reduced from 13% to 5%. This while still obtaining considerable total cost savings compared to the FCFS schedule (28%), which represents two thirds of the cost savings obtained in the minimum cost solutions. The different heuristics obtain different trade-offs between total cost and fairness. Hopefully, these results can be a starting point for discussing these fairness issues related to the introduction of CDM processes among air traffic stakeholders.

We finish this introduction with an overview of the literature. Beasley et al. [6] gives an extensive literature overview on the aircraft landing problem. The objective usually considered is makespan or total delay. More recent papers have considered different heuristics (e.g., Pinol and

Beasley [8]) and the dynamic version of the problem (e.g., Beasley et al. [7]).

Carr, Erzberger and Neumann analyse the effect of using sequence preferences from airlines for their own aircraft ([9]) and of exchanging delays ([10]), using simulations. Their results indicate that both can be done while causing little or no decrease in efficiency.

Another approach to assign landing capacity to flights are so called ground delay programs. These are commonly used in the US. Such a program is issued whenever a sustained period of congestion at an airport is expected. Usually this is caused by a decline in airport arrival capacity (e.g., because of severe weather conditions). This will lead to delays for flights landing at this airport. During a ground delay program these delays are (as much as possible) assigned to flights on the ground prior to departure rather than en-route. According to Inniss and Ball [14] it is about twice as cheap to delay an aircraft on the ground instead of in the air.

The procedure starts with determining the (expected) airport arrival capacity. Then a number of landing slots is defined equal to the determined capacity. This means that landing capacity is considered as input and is not optimized (by sequencing the flights) as in the aircraft landing problem. The landing slots have to be allocated among the flights that were originally scheduled during the period of congestion. There is a lot of attention to use CDM procedures in ground delay programs.

The current practice in the US is to allocate the arrival slots to the airlines by a procedure called *ration by schedule*. This procedure orders the flights according to their scheduled time of arrival and assigns them to the slots in this order. Airlines are now considered the owners of the slots assigned to their flights and can reassign their slots to different flights by slot substitutions and cancellations. After this, a procedure called *compression* is used to fill unassigned slots (caused by cancellations). The idea behind compression is to reward airlines for slots they release, thus encouraging airlines to report cancellations. Vossen and Ball [18] discuss this ground holding procedure that is currently used in practice and formalize it using mathematical models.

A lot of alternative procedures for the compression method are modeled in the literature, involving a shift to decision-making by the airlines. Vossen and Ball ([18] and [19]) propose a model in which airlines can propose requests for slot exchanges which are considered by a mediator (the FAA). According to their results, this improves on-time performance and passenger delay measures.

Abdelghany et al. [2] considers the problem of the airline to assign its flights to the slots it received. It uses a genetic algorithm to solve this such that the overall downline impact resulting from delaying these inbound flights is minimised.

In his PhD thesis, Hall [13] presents two methods to assign arrival slots to airlines using airlines preferences. The first is called the Objective Based Allocation Model. Airlines supply a value for each flight/slot combination. An assignment problem is solved with the objective to maximize the total value. In order to get truthful values, each airline has to pay a fee that is equal to the value lost by the other airlines caused by the presence of the considered airline. The author concludes that although in theory the potential benefit is enormous, it would be very difficult to implement such a method.

The second method is called the Arrival-Departure Capacity Allocation Method. This is an extension to ground delay models where airport arrival and departure capacity (and their relation) are considered at the same time. Airlines get restrictions on the amount of runway capacity they are allowed to use and can decide whether they use this capacity for arrivals or departures (with some restrictions). The results of an extensive simulation study show that the airline objectives increase compared to traditional ground delay programs, especially on low capacity scenarios.

Andreatta et al. [3] consider the multi airport ground holding problem. Their aim is to improve the application of these models, by introducing simple, easy to understand heuristics. Each heuristic is based on a specific "priority rule". The flights are grouped dynamically into a manageable number of classes. Each "priority rule" specifies, for any two flights of different classes, which one has priority over the other. How to assign priority classes to the flights in an equitable manner is not considered.

In our opinion, the work presented in this paper brings the following contributions compared to the existing literature.

When cost functions or weights are used, it is important to consider how these are obtained in practice. If airlines are allowed to supply weights or costs, fairness should be considered. These issues are not addressed in most literature mentioned above.

At the same time, we consider the runway sequencing problem. This gives the possibility to optimise runway capacity, opposed to the mentioned ground delay problems, in which capacity is considered as input. To our opinion, this is essential because the runway capacity is the major bottleneck of the air traffic system.

2 Fairness

Although everybody has a general idea about fairness, it is hard to give a formal definition, especially related to the aircraft landing problem.

There is some literature where fairness is considered in other fields, such as game theory [5], bandwidth sharing problems in computer and telecom networks [15], queuing [4] and job scheduling [16]. These are problem-specific definitions of fairness that are not applicable to the aircraft landing problem.

We want to obtain a safe, efficient and fair schedule with as little costs for the airlines as possible. The cost incurred by an airline for an arriving flight largely depends on its landing time. Each flight has different characteristics, therefore the airline is allowed to provide a different cost function for each individual flight. Such a function relates the arrival time to cost.

We want to allow the airlines as much flexibility as possible in determining the cost functions. At the same time we want to obtain a schedule that does not show favor to some airlines at the expense of others.

In [17] we proposed a restriction on the shape of the functions and the application of a scaling mechanism to obtain this. Thus, the cost of flights of different airlines can be compared in a fair way. However, minimizing the total scaled cost does usually not result in a schedule in which the scaled cost and/or delay is shared equally (or proportionally) among the airlines.

It is complicated to obtain such a schedule because the fairness has to be determined during the optimization process. During this process the total amount of costs and delay is not fixed and changes when it is divided differently over the airlines.

In this paper this issue will be addressed further, starting with the introduction of three definitions of fairness in the next section. In section 2.2 the scaling method is explained.

2.1 Fairness Definitions

The following notation is used throughout this paper:

Let $F = \{1, \dots, N\}$ be the set of all flights to schedule.

Let A be the set of all airlines.

Let $F_a \subset F$ be the set of flights of airline $a \in A$.

Note that $F = \bigcup_{a \in A} F_a$ and $F_a \cap F_b = \emptyset$, for all $a, b \in A, a \neq b$.

Let $f_i(t)$ be the scaled cost function, relating landing time with scaled cost, for flight i .

The restrictions on the scaled cost functions $f_i(t)$, $i \in F$, and the scaling method will be explained in section 2.2.

Absolute Fairness

The most natural way to define fairness is to divide the scaled cost equally (proportional to the number of flights) over the airlines. This means that we will compare the average scaled cost per

flight of the airlines:

$$\bar{c}_a(t_1, \dots, t_N) := \frac{1}{|F_a|} \sum_{i \in F_a} f_i(t_i) \quad a \in A, \quad (1)$$

given a schedule where flight i lands at time t_i .

These average scaled cost per flight should be equal or almost the same for all airlines. To measure how fair a schedule is we can use the root mean square deviation:

$$\sigma_{\bar{c}}(t_1, \dots, t_N) := \sqrt{\frac{1}{|A|} \sum_{a \in A} \left(\bar{c}_a(t_1, \dots, t_N) - \frac{1}{|A|} \sum_{b \in A} \bar{c}_b(t_1, \dots, t_N) \right)^2} \quad (2)$$

A problem is that the total scaled cost are not fixed but depend on the order and times the flights are scheduled. This means that there will always be a trade-off between efficiency (total scaled cost) and fairness (the division of these total cost over the airlines).

Therefore, it would make no sense to make a change (to improve fairness) in a schedule that results in the same cost for all airlines except one, which costs increase. On the other hand we don't want to improve the cost of an airline that has already low average cost on the expense of an airline with high average cost. It is also hard to use the above measure directly in the optimization process. Therefore, during the optimization process the scaled cost of the airline that is the worst off is minimised:

$$\min_{t_1, \dots, t_N} \max_{a \in A} \{ \bar{c}_a(t_1, \dots, t_N) \} \quad (3)$$

This objective will be iteratively improved during the optimization process. In the obtained schedule every airline will have scaled cost that are equal or smaller than the obtained maximum.

Relative Fairness

It is natural to compare the current schedule to a reference schedule. Usually this will be an earlier schedule. Fairness can be assessed by comparing the airlines improvement compared to this reference schedule. In this research we will use the schedule that would be obtained without actively sequencing the flights as reference schedule. In this schedule it is assumed that the flights are landed (respecting the required separation) in the order in which they approach the airport. This schedule is called the First Come First Served (FCFS) schedule.

Let \hat{t}_i be the landing time of flight i in the reference schedule. In this case we will compare the ratio of the airline cost in the considered schedule and the reference schedule:

$$\Delta_a(t_1, \dots, t_N) := \frac{\sum_{i \in F_a} f_i(t_i)}{\sum_{i \in F_a} f_i(\hat{t}_i)} \quad a \in A. \quad (4)$$

Ideally we want every airline to have some minimal improvement (or maximum deterioration) in the new schedule. Therefore we can measure this by the percentage of airlines that is worse off than in the reference schedule:

$$\frac{1}{|A|} \sum_{a \in A} 1_{\{\Delta_a(t_1, \dots, t_N) > 1\}} \quad (5)$$

Or equivalently to the percentage of airlines that has an improvement less than a certain fixed ratio.

Considering the trade-off between total cost and fairness, we can obtain this by maximizing the improvement of the airline, that is the worst off (lowest improvement), during the optimization process:

$$\min_{t_1, \dots, t_N} \max_{a \in A} \{ \Delta_a(t_1, \dots, t_N) \} \quad (6)$$

Fairness Measured by Delay

Delay is used as measure for fairness. Thus, fairness is not measured in terms of airline cost. A trade-off between those two has to be found. This trade-off might be more efficient than with absolute or relative fairness because the fairness is assessed by a different measure.

To measure fairness we will compare the average airline delay per flight compared to the reference schedule:

$$\bar{d}_a(t_1, \dots, t_N) := \frac{1}{|F_a|} \sum_{i \in F_a} (t_i - \hat{t}_i)^+ \quad a \in A. \quad (7)$$

Similarly as with the scaled cost, the fairness measured by delay can be evaluated by the root mean square deviation of the average airline delays:

$$\sigma_{\bar{d}}(t_1, \dots, t_N) := \sqrt{\frac{1}{|A|} \sum_{a \in A} \left(\bar{d}_a(t_1, \dots, t_N) - \frac{1}{|A|} \sum_{b \in A} \bar{d}_b(t_1, \dots, t_N) \right)^2} \quad (8)$$

In the optimization process we will minimize the average delay of the airline that is the worst off:

$$\min_{t_1, \dots, t_N} \max_{a \in A} \{ \bar{d}_a(t_1, \dots, t_N) \} \quad (9)$$

2.2 Airline Cost Scaling Mechanism

We will shortly explain the cost function scaling mechanism, as introduced in [17], in this section again.

The cost functions supplied by the airlines are required to be convex and piecewise linear and to have a minimal cost of zero at a time within the interval they are defined on.

The scaling mechanism ensures that the average cost per time unit per flight are approximately the same for all airlines. This is done by introducing a single scaling factor for each airline. All cost functions for flights from this airline are multiplied with this scaling factor. In this way the original cost ratio between flights of the airline is preserved.

Let us make this more precise. Consider airline $a \in A$ with $|F_a|$ arriving flights, with convex piecewise linear cost functions $\kappa_i(t)$ for $i \in F_a$. Let E_i and L_i be the earliest and latest possible landing times of aircraft i , respectively. So, $\kappa_i(t)$ is defined on the interval $[E_i, L_i]$.

To obtain equity, these cost functions will be scaled to new cost functions $f_i(t) = \alpha_a \kappa_i(t)$ ($i \in F_a$). The scaling factors α_a are determined per airline. This ensures that the ratio between costs of their own flights are preserved in the scaled objective functions. α_a is defined such that:

$$\frac{1}{|F_a|} \sum_{i \in F_a} \frac{\int_{E_i}^{L_i} \alpha_a \kappa_i(t) dt}{(L_i - E_i)^p} = 1.$$

So,

$$\alpha_a = |F_a| \left(\sum_{i \in F_a} \frac{\int_{E_i}^{L_i} \kappa_i(t) dt}{(L_i - E_i)^p} \right)^{-1},$$

where p is a parameter to minimise the effect of differences in the length of the landing intervals. It is preferable to choose p just over 2. This gives a small flexibility reward for flights with a larger time interval.

3 Model

In this section we will formulate a model for the airport landing model, which provides landing times for a given set of flights at a runway. The time between landings should be equal or larger than the minimum separation time required for safe operations. The separation times required

between every pair of flights are given. The landing time of each flight should be within the given landing time interval of the flight. We will use different objectives to compare (feasible) schedules.

The MIP formulation of the problem is given in section 3.1. Various objectives related to fairness for the problem are introduced in section 3.2.

3.1 MIP Formulation

In this section a Mixed Integer Programming (MIP) formulation of the model is given. The basic notation and constraints are similar to those of Beasley et al. [6].

Let $F = \{1, \dots, N\}$ be the set of all flights to schedule. Let

- E_i : Earliest possible runway time for flight i $i \in F$
- L_i : Latest possible runway time for flight i $i \in F$
- S_{ij} : Required separation time when flight i uses the same runway before flight j $i, j \in F, i \neq j$

The main decision variables are the landing times of the flights. Further the formulation requires some additional decision variables to represent the sequence of the flights:

- t_i : landing time for flight i $i \in F$
- $\delta_{ij} = \begin{cases} 1 & \text{if flight } i \text{ lands before flight } j \\ 0 & \text{otherwise} \end{cases}$ $i, j \in F, i \neq j$

To make sure the variables act as described, the following constraints are introduced:

$$E_i \leq t_i \leq L_i \quad i \in F \quad (10)$$

$$\delta_{ij} + \delta_{ji} = 1 \quad i, j \in F, j > i \quad (11)$$

Constraint (11) ensures that either flight i lands before flight j or the reverse. These variables are needed in the separation constraints.

We must ensure that proper separation is maintained between pair of flights using the same runway. To obtain this, we introduce the following sets of pair of flights, determined by their possible runway time intervals :

- U : the set of pairs (i, j) of flights for which it is undetermined whether flight i lands before flight j or the other way around
- V : the set of pairs (i, j) of flights for which flight i definitely lands before flight j , but for which the separation is not automatically satisfied
- W : the set of pairs (i, j) of flights for which aircraft i definitely lands before flight j , and the separation is automatically satisfied

More formally:

- $U = \{(i, j) | E_j \leq E_i \leq L_j \text{ or } E_j \leq L_i \leq L_j \text{ or } E_i \leq E_j \leq L_i \text{ or } E_i \leq L_j \leq L_i, i, j \in F, i \neq j\}$
- $V = \{(i, j) | L_i < E_j \text{ and } L_i + S_{ij} > E_j, i, j \in F, i \neq j\}$
- $W = \{(i, j) | L_i < E_j \text{ and } L_i + S_{ij} \leq E_j, i, j \in F, i \neq j\}$.

The following constraints will ensure the proper separation:

$$\delta_{ij} = 1 \quad (i, j) \in V \cup W \quad (12)$$

$$t_j \geq t_i + S_{ij} \quad (i, j) \in V \quad (13)$$

$$t_j \geq t_i + S_{ij}\delta_{ij} - (L_i - E_j)\delta_{ji} \quad (i, j) \in U \quad (14)$$

If flight i definitely precedes flight j then we can fix δ_{ij} (constraint (12)). For $(i, j) \in V$ the order is known but the separation still needs to be ensured (constraint (13)). This must also be done for the pairs in U . This is done by constraint 14 for the pair (i, j) if flight i lands before flight j ($\delta_{ij} = 1, \delta_{ji} = 0$). If this is not the case this constraint is superfluous. Note that if $(i, j) \in U$ then $(j, i) \in U$ and constraint (14) ensures the separation for both orders.

As explained in section 2 each airline provides a convex piecewise linear cost function for each of its flights, that is then scaled to obtain equity among airlines. Consider such a cost function $f_i(t)$ for flight i . This function can be written as a set of linear functions on a number of connected intervals. Let X_{i1}, \dots, X_{iK_i} be the K_i breakpoints of $f_i(t)$. The function consists of $(K_i + 1)$ linear functions with slopes A_{i0}, \dots, A_{iK_i} and intercepts B_{i0}, \dots, B_{iK_i} :

$$f_i(t) = \begin{cases} A_{i0}t + B_{i0} & 0 \leq t \leq X_{i1} \\ A_{i1}t + B_{i1} & X_{i1} \leq t \leq X_{i2} \\ \vdots & \vdots \\ A_{iK_i}t + B_{iK_i} & X_{iK_i} \leq t \end{cases},$$

And because of the convexity the following holds:

$$f_i(t) = \max_{k=0, \dots, K_i} \{A_{ik}t + B_{ik}\}.$$

These functions are not linear in the current decision variables t_i , and therefore new decision variables c_i are introduced:

c_i : cost for landing flight i , $i = 1, \dots, N$

We need constraints to ensure that c_i represents the cost function correctly:

$$c_i \geq A_{ik}t_i + B_{ik} \quad i = 1, \dots, N; k = 0, \dots, K_i \quad (15)$$

These constraints will make sure that c_i is equal or greater than the (scaled) cost for flight i . The c_i variables will be used in the different objectives that are introduced in the next section.

In [17] the minimum total scaled cost are used as objective:

$$\min \sum_{i=1}^N c_i. \quad (16)$$

3.2 Fairness

In this section, we will define objectives which are related to our fairness definitions.

Absolute Fairness

The objective to minimize the maximum average airline cost per flight as defined in equation (3) is easily formulated in terms of our MIP model.

Let c_{\max} be the decision variable that represents the maximum airline average cost. We can model this using the following objective and constraints.

$$\min c_{\max} \quad (17)$$

$$c_{\max} \geq \frac{1}{|F_a|} \sum_{i \in F_a} c_i \quad \forall a \in A \quad (18)$$

As we mentioned in section 2 there is always a trade-off between total cost and fairness, therefore we adapt the objective to represent this trade-off:

$$\min c_{\max} + \frac{\epsilon}{N} \sum_{i=1}^N c_i \quad (19)$$

ϵ should be chosen small ($0 < \epsilon \ll 1$) to focus on absolute fairness.

Relative Fairness

Equation (6) can be formulated in terms of the MIP model as follows.

Let \hat{c}_f be the cost involved with landing flight f using the reference schedule. Let Δ_{\max} be the maximum airline ratio of the cost in the current schedule and the reference schedule. So $(1 - \Delta_{\max})$ represents the minimum improvement. We can now model this using the following objective and constraints.

$$\min \Delta_{\max} \quad (20)$$

$$\Delta_{\max} \geq \frac{\sum_{i \in F_a} c_i}{\sum_{i \in F_a} \hat{c}_i} \quad \forall a \in A \quad (21)$$

Again we want also to consider the trade-off between total cost and fairness, by adapting the objective in the following way:

$$\min \Delta_{\max} + \frac{\epsilon}{\sum_{i=1}^N \hat{c}_i} \sum_{i=1}^N c_i \quad (22)$$

Fairness measured by delay

The minimum of the maximum average airline delay per flight (equation (9)) can also be incorporated in our MIP formulation. Let \hat{t}_i be the timetable landing time of flight i . We can now model this using the following objective and constraints.

$$\min d_{\max} + \frac{\epsilon}{N} \sum_i c_i \quad (23)$$

$$d_{\max} \geq \frac{1}{|F_a|} \sum_{i \in F_a} d_i \quad \forall a \in A \quad (24)$$

$$d_i \geq t_i - \hat{t}_i \quad \forall i \quad (25)$$

$$d_i \geq 0 \quad \forall i \quad (26)$$

In this case we included the trade-off between total cost and fairness straight away, because they are assessed by different measures.

4 Example

To illustrate the difference between our definitions of fairness and the trade-off with total cost, a small example is presented in this section.

Consider the following flights to plan:

id	airline	delay cost	Timetable	FCFS	Min cost	Absolute	Relative	Delay
1	A	2	00:00	00:00	00:00	00:00	00:00	00:00
2	A	6	00:01	00:02	00:02	00:02	00:02	00:02
3	B	2	00:02	00:04	00:10	00:10	00:10	00:08
4	B	4	00:03	00:06	00:08	00:04	00:04	00:04
5	A	5	00:04	00:08	00:04	00:08	00:06	00:10
6	B	7	00:05	00:10	00:06	00:06	00:08	00:06

We assume that the flights cannot land earlier than their timetable time. Between all flights a separation of 2 minutes is required. We assume that the cost are linear in the amount of delay.

The cost per minute delay differs per flight and is listed in the table. Note that the average cost per minute delay are the same for both airlines, so scaling is not necessary.

In the table above optimal schedules considering different objectives are shown. In the table below the delay and costs and their division over the airlines are shown.

The FCFS schedule is the minimum total cost schedule when the flights are ordered according to their timetable time. This schedule is used as reference schedule. In this case the total cost are 77 and the total delay is 15 minutes. Airline B has (almost) twice the delay and cost of airline A.

In the minimum total cost schedule, the total cost are only 49. However, almost all the delay and cost are at the expense of airline B.

If we focus on absolute fairness, the costs can be shared almost equally. The total cost are only slightly higher than in the minimum total cost schedule. Airline B has large savings compared to the FCFS schedule, while airline A's cost remain equal.

Considering relative fairness, it is possible to accomplish savings compared to the FCFS schedule for both airlines. Airline A saves 38% and airline B 20%. However, since the total cost are higher than in the minimum total cost schedule, there is a clear trade-off between relative fairness and total cost. In this particular case, the minimum cost schedule could be preferable when considering relative fairness, because this schedule gives almost the same savings for airline B (18%) and larger savings for airline A.

In the four schedules considered, airline B receives much more delay than airline A. When the objective is fairness measured by delay, we obtain a schedule with almost the same delay for both airlines. Again, there is a trade-off between fairness and cost. This schedule has the highest total cost. However, Airline B has the lowest cost of all considered schedules. But airline A has larger cost and delay than in the any of the other schedules.

Airline	FCFS		Min cost		Absolute		Relative		Delay	
	delay	cost	delay	cost	delay	cost	delay	cost	delay	cost
A	5	26	1	6	5	26	3	16	7	36
B	10	51	14	42	10	27	12	41	8	23
Total	15	77	15	49	15	53	15	57	15	59

This example shows that the minimum total cost schedule can have an unequal spread of cost and delay over the airlines. By considering fairness explicitly, we can obtain schedules with a fairer spread of cost and delay. However, there is a trade-off because the total cost in these schedules are often larger. Schedules with different trade-offs between total cost, delay and fairness can be obtained.

5 Local Search Heuristics

In [17] a heuristic to solve the aircraft landing problem was introduced. This heuristic uses local search (to select a sequence for the flights) and LP (to determine the optimal landing times for this sequence) repeatedly. An estimate of the objective improvement of every member of the current neighborhood is calculated to evaluate promising neighbors first. An outline of the algorithm is given below.

LOCAL SEARCH()

- 1 $\pi = \text{FIND INITIAL SOLUTION}()$
- 2 $N(\pi) := \text{Set of neighbors of } \pi$
- 3 Estimate objective improvement for all members of $N(\pi)$
- 4 **while** $N(\pi) \neq \emptyset$
- 5 **do** $\pi' = \text{neighbor with maximum estimated improvement in } N(\pi)$
- 6 **if** $LP(\pi')$ is feasible and $z_{LP(\pi')} \leq z_{LP(\pi)}$
- 7 **then** $\pi = \pi'$
- 8 $N(\pi) = \text{Set of neighbors of } \pi$

```

9           Estimate objective improvement for all members of  $N(\pi)$ 
10          else  $N(\pi) = N(\pi) \setminus \pi'$ 
11 return  $\pi$  and the optimal landing times for this sequence

```

As initial solution, we can use the approach currently used in practice. The flights are usually landed in the order they approach the airport. Therefore this order is called the First Come First Served (FCFS) sequence. When scheduling, this means the sequence is formed by the flights, sorted according to their expected arrival times. Solving the LP will give the optimal landing times (respecting separation requirements) given this order. This initial solution is also used as reference schedule when considering relative fairness and to calculate cost savings.

In all our heuristics we will use a shift neighborhood. A neighbor is obtained by removing a flight and inserting it at a different position in the sequence.

In the next section we will introduce restrictions on this neighborhood and different methods to determine the order in which the neighbors are evaluated.

5.1 Neighborhood Selection

To reduce the computation time, it is preferable to select a neighbor in a manner that finds an improvement by evaluating as few neighbors as possible. Therefore, we will evaluate promising neighbors first. In this section various methods to identify promising neighbors are presented. It is essential that the methods are fast (compared to the neighbor evaluations).

Additionally, we could reduce the size of the neighborhood. Of course this brings the risk of deteriorating the final solution. However, if by choosing a smart criteria (e.g., related to fairness), it could also help to find solutions with certain features.

In this section we will introduce restrictions on the neighborhood and different methods to determine the order in which the neighbors will be evaluated.

Estimated Cost Improvement

This method estimates the improvement in scaled cost between the current solution and the neighbor. The estimation uses the landing times in the current solution, to estimate the landing times and involved scaled cost in the neighbor. This can be done by estimating the cost for the flight at the i -th position in the neighbor by using this flights' cost function with the optimal time for the flight at the i -th position in the current solution.

Now we can sort the neighbors according to their estimated cost improvement. We will evaluate the neighbor with the largest estimated cost improvement first. If this does not actually improve the objective, we will continue to evaluate the neighbor with the second largest estimated improvement, etc.

Airline Selection

Incorporating fairness criteria can be done by only considering cost improvements of certain airlines.

When absolute fairness is considered this could be the airlines of which the average scaled cost in the current solution are close enough to the maximum airline average scaled cost. That is, for airlines $a \in A$ for which

$$\frac{1}{|F_a|} \sum_{f \in F_a} c_f \geq \beta \max_{a \in A} \left\{ \frac{1}{|F_a|} \sum_{f \in F_a} c_f \right\} \quad 0 \leq \beta \leq 1. \quad (27)$$

When relative fairness is considered this could be the airlines of which the cost ratio in the current solution are close enough to the maximum cost ratio:

$$\frac{\sum_{f \in F_a} c_f}{\sum_{f \in F_a} \hat{c}_f} \geq \beta \max_{a \in A} \left\{ \frac{\sum_{f \in F_a} c_f}{\sum_{f \in F_a} \hat{c}_f} \right\} \quad 0 \leq \beta \leq 1. \quad (28)$$

	Objective	Airline Selection
(1)	total scaled cost (FCFS)	
(2)	total scaled cost	no
(3)	absolute fairness	yes
(4)	relative fairness	yes
(5)	absolute fairness	adaptive threshold
(6)	absolute fairness	after cost minimization
(7)	fairness measured by delay with scaled cost	no
(8)	fairness measured by delay with real cost	no

Table 1: Heuristics

The airline threshold β determines how sensitive this control is. Choosing $\beta = 1$ will only consider changes that involve flights of the maximum cost airline(s). This might be too restrictive and might lead to bad local minima (w.r.t. total scaled cost). The same can be said of β close to 0 w.r.t. fairness criteria.

β can be dynamically updated during the algorithm. Thus, first a solution that has little total cost can be found and then the fairness of this solution can be improved. This can be done by choosing $\beta = 0$ initially and increasing β with small steps with each improvement in objective.

5.2 Heuristics

In table 1, eight heuristics are listed, combining different objective functions, neighborhoods and neighborhood selection criteria. These heuristics are aimed at finding different trade-offs between cost, delay and fairness.

Heuristic (1) is the reference schedule. This is the minimum cost schedule given that the flights will land in timetable order.

Heuristic (2) is aimed at finding the minimum scaled cost schedule.

Heuristic (3) uses the absolute fairness objective (19). The neighborhood is restricted by airline selection (with a fixed threshold $\beta = 0.9$).

Heuristic (4) uses the relative fairness objective (22). The neighborhood is restricted by airline selection (with a fixed threshold $\beta = 0.9$).

Heuristic (5) uses the absolute fairness objective (19) again. The neighborhood is restricted by airline selection with an adaptive threshold starting at 0 and increasing with 0.1 with every improvement of the objective until a maximum value of 0.9.

Heuristic (6) uses the absolute fairness objective (19) again. To avoid fair solutions with large total cost, the solution from heuristic (2) is used as initial solution. The neighborhood is restricted by airline selection (with a fixed threshold $\beta = 0.9$).

Both heuristic (7) and (8) measures fairness by delay using (LP) objective (23). The neighborhood is not restricted by airline selection. Heuristic (7) uses the total scaled cost as the efficiency measure. Heuristic (8) uses the original cost functions as supplied by the airlines. The cost scaling is not performed in this case because the fairness is measured by delay.

6 Computational experiments

In this section, we assess the performance of the local search heuristics. A large number of instances, created using schedule data from a major European hub, were tested. These data contain all arrivals from a week in September 2004. The data included airline, flight number, aircraft type, arrival runway and scheduled arrival times.

It is assumed that the cost function of every flight has a minimum of zero cost at the scheduled time of arrival of the flight, according to the timetable.

The structure of the cost functions for the home carrier and its partners, was determined in cooperation with specialists from this airline. Its perceived delay costs are strongly related to the number of missed transfers. This is quite natural, since this airline uses the airport as hub, and consequently has a lot of passengers transferring at the airport. Exact passenger flows and related costs were not provided by the airline for reasons of confidentiality. Instead the number of missed transfers per 15 minutes of delay were generated randomly. These were translated into convex piecewise linear cost functions by using the cumulative number of missed transfers up to time t as the slope of the cost function at time t .

The main objective for the other airlines is punctuality. The latest landing time before other flights from the airline are affected by the delay of the flight considered, is also taken into account. This is represented in the cost function by choosing a positive slope between the scheduled time and this time and a steeper slope hereafter. This point in time and the exact slopes are generated randomly for each flight.

Our planning horizon is several hours before the flights will land. That means that flights from within Europe have not departed when planning. These flights are assumed to have a maximum departure delay and a possible landing interval of 3 hours. Intercontinental flights are en-route and are able to arrive between 25 minutes before and 30 minutes after schedule.

In Table 2 the arrival separation distances under good visibility conditions, according to international regulations, are listed.

		Following aircraft		
		Light	Medium	Heavy
Leading aircraft	Light	3	3	3
	Medium	5	3	3
	Heavy	6	5	4

Table 2: Wake vortex separation in nautical miles for different weight categories

In low-visibility conditions, the separation distances must be larger.

The required separation distance under low-visibility conditions is the maximum of the required wake-vortex and low-visibility separations. The actual required separation *time* between two flights is calculated using the weight categories and approach speeds of the aircraft. The aircraft types used for the flights are available from the data.

In the experiments, low-visibility conditions requiring 6 nautical miles separation were used. This separation causes a decreased arrival capacity which will often result in large delays. It is interesting to assess the (scaled) costs and their spread among the airlines, resulting from those delays.

The data contained 3978 flights of 121 airlines. 139 instances were created by dividing the arrivals by runway and time. A runway is only used continuously for at most a few hours, depending on demand and weather conditions. The flights landing in such a period on a runway, are considered as a single instance. These instances contained between 1 and 117 flights. All the instances were solved using the local search heuristics introduced in section 5.2. The heuristics are implemented in C++. The LP problems are solved using the COIN-OR LP solver [1].

6.1 Results

All results will be presented relative to the FCFS schedule, resulting from heuristic (1). This is the schedule that is obtained without actively sequencing the flights. It is assumed that the flights are landed (respecting the required separation) in the order in which they approach the airport. As mentioned before this schedule resembles current practice. For our data it means that the FCFS order is determined by the scheduled times in the timetable. The landing times for the FCFS schedule are obtained by solving the LP for the FCFS order. Thus, the resulting landing times respect the operational separation requirements and are within the feasible landing time intervals. The FCFS schedule is also used as initial solution by the local search heuristic (see section 5).




Figure 1: Total scaled cost and the root mean square deviation of the average airline scaled cost




Figure 2: Total delay and the root mean square deviation of the average airline delay

In figure 1 the total scaled cost resulting from the schedules generated by the eight heuristics is shown by the black bars. The heuristics obtain savings from 19% to 42% in terms of scaled cost compared to the FCFS schedule. As expected, heuristic (2) obtains the largest total cost savings.

In order to evaluate the absolute fairness the root mean square deviation of the average airline scaled cost $\sigma_{\bar{c}}$, as defined by equation (2), is depicted by the gray bars.

Heuristic (5) does only marginally lower the airline root mean square deviation compared to heuristic (2). Since also the total scaled cost savings are not competitive, the conclusion is that this heuristic often gives sub optimal solutions. Heuristic (3) lowers the root mean square deviation more substantially, but is still dominated by heuristic (4), where both the total scaled cost and root mean square deviation are lower.

In figure 2 the total delay resulting from the schedules generated by the eight heuristics is shown by the black bars. The reductions in delay are small. This is expected, since the FCFS schedule inherently will not have large delays. The reduction of delay can only be obtained by sequences that reduce the total separation time required. This can usually only be done by increasing the delay of at least one flight (that is moved to a position further in the sequence). As expected, heuristic (7) and (8) obtain the largest delay reduction (9%).

The gray bars depict the root mean square deviation of the average airline delay $\sigma_{\bar{d}}$, as defined by equation (8). This measures the fairness measured by delay. Heuristic (3), (4), (7) and (8) reduce this root mean square deviation substantially.

In figure 3 the total real cost resulting from the schedules generated by the eight heuristics is shown by the black bars. The gray bars depict the root mean square deviation of these cost. The results are similar to that of the scaled cost. There is only a large difference for heuristic (8). This follows from the explicit focus of this heuristic on the real cost (as trade off to the fairness measured by delay). This heuristic gives the largest cost savings (38%) compared to the FCFS schedule. This is more than the minimum total scaled cost schedule which obtains 35% real cost savings.

In figure 4 the percentage of airlines that have larger (real) cost compared to the FCFS schedule




Figure 3: Total real cost and the root mean square deviation of the average airline real cost




Figure 4: Percent of airlines that have no real cost improvement compared to the FCFS schedule

is shown, as defined in equation (5). This is a measure for the relative fairness. Heuristic (4) has the lowest percentage of airlines without any improvement (5%), as expected. Although heuristic (8) reduces the spread in the average delays of the airlines and has the lowest total real cost, it has the highest percentage of airlines that have larger (real) cost compared to the FCFS schedule (34%). So the more equal spread of delay does not necessarily lead to cost savings for all airlines.

Heuristic (6) indeed seem to improve the absolute fairness of its initial solution (the minimum total cost schedule resulting from heuristic (2)) slightly. It is still obtaining 40% total cost savings. This heuristic is suitable if total cost is the most important consideration and fairness only a secondary issue. It is dominated by the other heuristics on most fairness measures.

Heuristic (7) and (8) provide a large delay reduction and as expected a good fairness measured by delay. Surprisingly these heuristics provides also considerable cost savings. Heuristic (8) even obtains the lowest total real cost. However, these heuristics score the worst on relative fairness.

Heuristic (4) performs well on absolute fairness, relative fairness and fairness measured by delay. It is better than heuristic (1), (2), (5) and (6) on all these fairness measures. At the same time it obtains 30% savings in total scaled cost compared to the FCFS schedule. It can be concluded that this heuristic gives a good balance between the various criteria and may therefore be preferred over the others.

6.2 Analysis by airline category

To look into more detail to the fairness of the results, we defined 4 categories of airlines, based on the number of flights of the airline in the dataset. The first category consists of a single airline, the hub airline. This airline has almost half of the total number of flights. The large airlines have at least 70 flights per week (on average at least 10 per day). The medium airlines have at least 21 flights per week (on average at least 3 per day).

For each airline a the average cost per flight $\bar{c}_a(t_1, \dots, t_N)$ can be calculated for a given schedule. These numbers are used to calculate airline average scaled cost for each category. Let A^m be the

Category	Total Flights	Number of Airlines
Hub Airline	1880	1
Large Airlines	941	7
Medium Airlines	539	15
Small Airlines	618	98

Table 3: Airline categories

set of airlines in category m . The category average \tilde{c}_m is now calculated by

$$\tilde{c}_m := \frac{1}{|A^m|} \sum_{a \in A^m} \bar{c}_a(t_1, \dots, t_N).$$

In table 4 the normalised average category scaled cost are shown. For category m these are obtained by

$$\frac{\tilde{c}_m}{\frac{1}{4} \sum_{r=1}^4 \tilde{c}_r}.$$

Heuristic	Hub airline	Large airlines	Medium airlines	Small airlines
(1)	0.68	0.95	1.19	1.19
(2)	0.79	0.85	1.19	1.16
(3)	1.07	1.02	0.87	1.04
(4)	0.85	1.07	1.03	1.04
(5)	0.92	0.90	1.05	1.12
(6)	0.85	0.82	1.14	1.19
(7)	0.80	0.88	1.14	1.17
(8)	0.38	1.04	1.24	1.33

Table 4: Normalised average scaled cost per airline category

It is clear that using FCFS schedule the medium and small airlines have larger scaled cost per flight on average. The costs are 16% and 19% larger than the average over the categories. This is because these airlines have a relatively large share of their flights during peak periods, in which delays are much more likely to occur.

In table 5 the real cost for the airlines in each category divided by the real cost in the category using the FCFS schedule as obtained by heuristic (1) are shown. For category m this is calculated by

$$\frac{\sum_{a \in A^m} \sum_{i \in F_a} \kappa_i(t_i)}{\sum_{a \in A^m} \sum_{i \in F_a} \kappa_i(\hat{t}_i)}.$$

Heuristic	Hub airline	Large airlines	Medium airlines	Small airlines
(1)	1.00	1.00	1.00	1.00
(2)	0.67	0.53	0.65	0.83
(3)	1.13	0.78	0.54	0.72
(4)	0.85	0.77	0.64	0.76
(5)	1.00	0.69	0.72	0.87
(6)	0.73	0.50	0.65	0.82
(7)	0.80	0.64	0.68	0.77
(8)	0.52	0.90	0.83	0.72

Table 5: Relative real cost compared to heuristic (1) (FCFS schedule)

Heuristic (2) focuses only on total scaled cost minimization. The average improvement (in real cost) per category lies between 17% for the small airlines and 47% for the large airlines.

Heuristic (3) especially reduces the cost 46% for the medium and 28% for the small airlines. This however comes with a cost increase for the hub airline. Overall it results in the category average scaled cost being much closer together.

Using heuristic 4, the average percentage improvements for the categories are closer together. The improvement is between 15% for the hub airline and 36% for the medium airlines.

Heuristic (8) focuses on real cost instead of scaled cost. The different cost structure of the hub airline has a large impact on the results. The real cost of delaying a flight (with a considerable number of transfer passengers) of the hub airline are higher than of other airlines, because of the cost related to the missed transfers. Without scaling this leads to larger absolute and relative cost savings for the hub airline. However, the average real cost per flight are still (a little) above average for the hub airline. The average delay is almost equal over the categories.

The airlines that are worse off than in the FCFS schedule fall mostly in the category of small airlines. This is directly related to their small number of flights. Twenty airlines in this category have only one flight in the dataset. Some of these flights are scheduled on time in the FCFS schedule and have a (small) delay in some of the schedules obtained using the other heuristics. The percentage of airlines that are worse off in this category flights is between 6% with heuristic (4) and 22% with heuristic (7). Since this still represents a minority of the airlines in this category, we expect that all airlines would achieve cost savings over longer periods of times (where these airlines will have a larger total number of flights). That is due to the fact that the schedules will be different from week to week and it is unlikely that the same flights will always receive delays.

7 Conclusions

In this paper, we evaluated a model that incorporates airline cost in the aircraft landing problem and explicitly considers fairness between airlines. We introduced three different measures of fairness: absolute fairness, relative fairness and fairness measured by delay. Several local search heuristics focused on obtaining schedules that perform well on the different fairness measures were introduced.

A large number of instances, created using schedule data from a major European hub, were tested. The results show that it is possible to obtain schedules with considerable cost savings compared to the FCFS schedule, and to achieve more fairness compared to the minimum total cost schedule. Different heuristics drastically reduce the root mean square deviation of the airline cost and delay per flight, compared to the heuristic that only focuses on total cost minimization. The cost savings compared to the FCFS schedule are still considerable.

The different heuristics show that schedules with different trade-offs between efficiency, cost, delay, absolute and relative fairness can be obtained. The heuristic that focus on relative fairness does also good on several other criteria considered, such as total cost savings and absolute fairness. Although this heuristic does not dominate all other heuristics on all criteria, it gives a good balance between the criteria. Therefore it might be preferred over the other heuristics.

Hopefully, these results can be a starting point for discussing these fairness issues related to the introduction of CDM processes among air traffic stakeholders.

References

- [1] Coin-or. <http://www.coin-or.org>.
- [2] A. Abdelghany, K. Abdelghany, and G. Ekollu. A genetic algorithm approach for ground delay program management: The airlines' side of the problem. *Air Traffic Control Quarterly*, 12(1):53-74, 2004.
- [3] G. Andreatta, L Brunetta, and G. Guastalla. Multi-airport ground holding problem: a heuristic approach based on priority rules. In L. Bianco, P. Dell'Olmo, and A.R. Odoni,

- editors, *Modeling and Simulation in Air Traffic Management*, pages 71–89. Springer-Verlag, Berlin, 1997.
- [4] B. Avi-Itzhak and H. Levy. On measuring fairness in queues. *Advances in Applied Probability*, 36:919–936, 2004.
- [5] J.B. Barbanel and S.J. Brams. Cake division with minimal cuts: envy-free procedures for three persons, four persons, and beyond. *Mathematical Social Sciences*, 48(3):251–269, 2004.
- [6] J.E. Beasley, M. Krishnamoorthy, Y.M. Sharaiha, and D. Abramson. Scheduling aircraft landings - the static case. *Transportation Science*, 34(2):180–197, 2000.
- [7] J.E. Beasley, M. Krishnamoorthy, Y.M. Sharaiha, and D. Abramson. Displacement problem and dynamically scheduling aircraft landings. *Journal of the Operational Research Society*, 55(1):54–64, 2004.
- [8] J.E. Beasley and H. Pinol. Scatter search and bionomic algorithms for the aircraft landing problem. *European Journal of Operational Research*, 171(2):439–462, 2006.
- [9] G.C. Carr, H. Erzberger, and F. Neuman. Airline arrival prioritization in sequencing and scheduling. In *2nd USA/EUROPE Air Traffic Management R&D Seminar*, 1998.
- [10] G.C. Carr, H. Erzberger, and F. Neuman. Delay exchanges in arrival sequencing and scheduling. *Journal of Aircraft*, 36:785–791, 1999.
- [11] Eurocontrol. European cdm portal. <http://www.euro-cdm.org>.
- [12] E.P. Gilbo and K.W. Howard. Collaborative optimization of airport arrival and departure traffic flow management strategies for cdm. In *3rd USA/EUROPE Air Traffic Management R&D Seminar*, 2000.
- [13] W.D. Hall. *Efficient Capacity Allocation in a Collaborative Air Transportation System*. PhD thesis, MIT, 1999.
- [14] T. Inniss and M.O. Ball. Estimating one-parameter airport arrival capacity distributions for air traffic flow management. *Air Traffic Control Quarterly*, 12(3):223–251, 2004.
- [15] L. Massoulié and J. Roberts. Bandwidth sharing: Objectives and algorithms. *IEEE/ACM Trans. on Networking*, 10(3):320–328, 2004.
- [16] G. Sabin, G. Kochhar, and P. Sadayappan. Job fairness in non-preemptive job scheduling. In *International Conference on Parallel Processing, ICPP*, pages 186–194, 2004.
- [17] M.J. Soomer and G.J. Franx. Scheduling aircraft landings using airlines’ preferences. *European Journal of Operational Research*, 190(1):277–291, 2008.
- [18] T. Vossen and M. Ball. Optimization and mediated bartering models for ground delay programs. *Naval Research Logistics*, 53(1):75–90, 2006.
- [19] T. Vossen and M. Ball. Slot trading opportunities in collaborative ground delay programs. *Transportation Science*, 40(1):29–43, 2006.