Flight and passenger efficiency-fairness trade-off for ATFM delay assignment

A. Montlaur

Escola d'Enginyeria de Telecomunicació
i Aeronàutica de Castelldefels
Universitat Politècnica de Catalunya, Spain
Email: adeline.de.montlaur@upc.edu

L. Delgado

School of Architecture and Cities
University of Westminster
London, United Kingdom
Email: l.delgado@westminster.ac.uk

Abstract

This paper studies trade-offs between efficiency (performance) and fairness (equity), when assigning ATFM delay pre-tactically (on-ground at origin airport) due to reduced airport capacity at destination. Delay is assigned as the result of the optimisation of a deterministic multi-objective problem considering flight and passenger perspectives when defining objectives of performance and fairness. Two optimisation cases are presented: one where objectives are based on flight metrics, and another one where they are based on passenger metrics. The paper defines and analyses efficiency-fairness trade-offs: the concepts of *price of fairness* for flights and passengers are defined as the percentage of efficiency loss due to the consideration in the optimisation of fairness; whereas the *price of efficiency* is considered as the fairness loss relative to the maximum value of the fairness metric, when considering flight or passenger delay in the optimisation. The optimisation model is based on the ground holding problem and uses various objective functions. For performance, total delay for flights (considering reactionary delay), and total delay for passengers (considering outbound connections) are defined. For fairness, the deviation of flight arrivals from a Ration By Schedule solution, and the deviation of delay experienced by passengers with respect to the one obtained in an RBS situation are used. An illustrative application on traffic at Paris Charles de Gaulle airport, a busy European hub airport, and including realistic values of traffic is modelled. A comprehensive trade-off analysis is presented. Results show, how in some cases, gains on one stakeholder can be achieved without implying any detriment on the other one. Passengers are more sensitive to the optimisation and hence, their consideration when assigning delay is recommended. Further research should explore how to combine flight and passenger indicators in the optimisation and consider how the lack of data availability could be mitigated.

Flight and passenger efficiency-fairness trade-off for ATFM delay assignment

I. Introduction

Airports are limited in capacity by operational constraints (Bazargan et al., 2002; Gilbo, 1993), generating in some cases, a significant imbalance between capacity and demand. Air Traffic Flow and Capacity Management (ATFCM) initiatives are then implemented to smooth traffic arrivals, transferring costly airborne delay, carried out with holdings and/or path stretching, to pre-departure on-ground delay (Carlier et al., 2007).

When a capacity-demand imbalance is detected, a Ration by Schedule (RBS) prioritisation of flights is the current practice for slot assignment (EUROCONTROL, 2015a). RBS policy is considered by the different stakeholders to be the fairest delay assignment, since it respects the original slot planning, but economical optimum cannot be guaranteed and only arrival delay is considered. Nevertheless, while airline reported delay was 14.7 minutes per flight in Europe in 2018, reactionary delay was responsible for 45% of it, heavily influenced by the impact of first-rotation delays (EUROCONTROL, 2019), and should thus not be neglected.

In the current operational environment, the system is optimised considering a flight-centred perspective, however, different stakeholders might experience ATM system performances differently. In particular, passenger-centric metrics might differ from their equivalent flight-centric ones (Cook et al., 2012). In (Montlaur and Delgado, 2017), performances for flight and passenger delays of an extended arrival manager (E-AMAN) were analysed in conjunction with a pre-tactical optimisation of flights. In that work, the assignment of slots was optimised considering either arrival delay for flights, arrival delay for passengers, total delay for flights (considering reactionary delay) or total delay for passengers (considering outbound connections). Results showed that in the scope of an E-AMAN, the distances and possible delays that can be assigned do not justify the application of a more sophisticated strategy than RBS. Nevertheless, when the scope of optimisation was enlarged to include the pre-tactical phase, benefits (and trade-offs) were obtained by optimising the assignment of delay instead of only considering flight schedules. While minimising the total delay for passengers is, as expected, the best strategy from the passengers perspective, it leads to higher reactionary delay for flights with respect to a flight-centric optimisation. Though optimisation carried out in (Montlaur and Delgado, 2017) focused on only one stakeholder at a time and did not include an explicit consideration of equity, it was shown that if focus is given to flight total delay, the benefit per passenger remains similar to the passenger centred optimisation and the variability with respect to the RBS delay assignment was reduced, improving the fairness of the solution.

Equity from a flight perspective was explicitly included in (Montlaur and Delgado, 2018), where flight and passenger delays and equity were conjointly considered in the optimisation scheme. Preliminary conclusions showed how focusing on one objective would affect the two others. The work presented in (Montlaur and Delgado, 2018) focused on the optimisation of three metrics (flight and passenger delays and flight-fairness) at the same time, allowing a better understanding of the trade-offs between performance from a flight delay, passenger delay and flight fairness perspective. The work focused on allowing decision makers to consider informed *a posteriori* articulation of preferences when optimising the delay assignment. To that end, the concept of *price of performance* (understood as *price of trade-off*) and *price of efficiency* as the ones defined in (Bertsimas et al., 2012) were used.

Previous research, focusing on three variables at the same time, lacked some applicability and a definition of fairness for passengers was missing. This paper has three main goals: 1) It considers a classical efficiency-equity trade-off at flight level, which would use easily available data. 2) It then studies the impact on passengers metrics. 3) It finally aims at covering the gap of equity for passengers by suggesting a definition of passenger delay fairness and finally analyses the trade-offs existing between flight and passenger performance and fairness metrics. Note that, during real life operations, the impact of uncertainty would be very relevant as the planned optimisation performed pre-tactically is subject to degradation when tactically realised. This might impact the trade-off between indicators and affect the optimal assignment of slots focusing on the expected outcome, as done for example in (Glover and Ball, 2013). In this paper, however, the assignment of slots is optimised in a deterministic manner since the main objective consists in defining and presenting different metrics. Future work should research the impact of the system's stochasticity on the stability of the solutions.

Section II details the background on fairness on delay assignment with an analysis of literature when allocating resources from a fairness point of view. Section III first explains the optimisation model and objectives used in this work. It then considers how individual objectives are combined in a multi-objective problem and presents the optimisation cases analysed in the paper. Finally, the consideration of efficiency and fairness trade-off with the concepts of *price of fairness* and *price of efficiency* are described. Section IV recalls the main hypothesis of the simulation of traffic that has been used as an example in this paper. The main results (for the multi-objective optimisation and for the performance-fairness trade-offs) and their discussion are reported in Section V. The paper finalises with the conclusions and further work found in Section VI.

TABLE I: Summary of fairness metrics encountered in the literature

	Flight-centered	Airline-Centered	Passenger-centered
(Glover and Ball, 2013; Kuhn, 2013)	deviation from RBS		
(Vossen et al., 2003)	squared deviation from RBS		
(Barnhart et al., 2012)	time-order deviation		
(Manley and Sherry, 2010)		airline's total delay based	
(Bertsimas and Gupta, 2016; Hamdan et al., 2018)	reversal number		
(Jacquillat and Vaze, 2016)		airline disutility	
(Manley and Sherry, 2010)			passenger delay based
(Pilon et al., 2016)		airline total deviation from RBS	

II. BACKGROUND: FAIRNESS ON DELAY ASSIGNMENT

When assigning delay to flights due to capacity-demand imbalances, even though commonly used, RBS is not the only possibility. Extensive research has been conducted to assign the required delay in a most cost-effective manner (Ball et al., 2007; Gilbo, 1993; Dell'Olmo and Lulli, 2003; Vranas et al., 1994; Montlaur and Delgado, 2017). As described in (Bertsimas et al., 2012), this type of resource allocation problems may be viewed as a utility allocation among different parties, which will lead to fairness issue. Note that even the definition of the stakeholders for which the fairness is estimated can be problematic: individual flights, airlines, passengers.

Due to the subjective nature of fairness and different possible interpretations of equity, there is no common definition of fairness allocation. Different proposals have been done such as maximisation of the minimum utility (max-min) (i.e., min-max for minimisation problems) or the α-fairness scheme as the one used in (Bertsimas et al., 2012). See (Young, 1995; Bertsimas et al., 2011) for a more detailed description of different fair metric definitions.

Though no standard of airline equity has been acknowledged in the industry, it is widely accepted by the ATM community that RBS presents a fair allocation of resources, as flights are not prioritised in any specific manner rather than their intended schedule. This has also benefits in terms of data availability and minimisation of gaming opportunity by reporting inaccurate information, for example. Fairness can be considered, with this definition, as being flight-centred and, for example, in (Glover and Ball, 2013) and (Kuhn, 2013), equity is maximised by minimising the deviation from the RBS solution in a stochastic uncertain assignment of slots. Similarly, in the context of a Ground Delay Program (GDP), Vossen et al. (2003) propose a method for mitigating inequities that arise due to exemptions, by minimising the squared deviation between the actual and ideal allocations (obtained from a RBS with no exemptions). Barnhart et al. (2012) also developed a fairness metric to measure deviation from original first-scheduled, first-served planning. They first determine a fair delay threshold for each flight, referred to as the maximum expected delay, and next calculate the time-order deviation for each flight as the amount by which the flight's delay in the controlled schedule exceeds this threshold. They then use an integer programming formulation attempting to directly minimise this metric. They showed promising results, in particular that trade-off between delay and fairness compares favourably. In (Bertsimas and Gupta, 2016), a two-stage approach for network ATFM incorporating fairness followed by airline collaboration was developed. For two flights f and f', a reversal was defined when flight f' arrives before f, when f was originally scheduled to arrive before f'. When a reversal occurs, the number of time periods between the respective arrival times constitutes an overtaking. A discrete optimisation model attempting to incorporate equitable distribution of delays among airlines by minimising the total number of reversals and the total amount of overtaking was proposed, which was seen as a generalisation of the RBS principle to a network setting. Flight fairness was also considered in distributing the delay among flights and in controlling reversals between flights in (Hamdan et al., 2018).

Recent studies have incorporated inter-airline equity considerations into the objective of ATFM models. In (Manley and Sherry, 2010), an airline equity metric due to flight delay is defined as the negative logarithm of the ratio of an airline's total flight delay over the total GDP flight delay, divided by the ratio of that airline's scheduled flights in the GDP over all GDP flights. Pilon et al. (2016) include equity in a user driven prioritisation process (UDPP) algorithm. They define the baseline delay as the amount of delay a flight would be assigned if no UDPP prioritisation were applied. The UDPP is then constrained such that the total baseline delay of each airspace user should remain the same as, or less than, the level before UDPP. Jacquillat and Vaze (2016) aim to balance the average per-flight displacement of flights from their requested times fairly among the airlines. Their formulation is based on defining each airline's disutility as the weighted average of per-flight displacements, which, ideally, would be the same for all airlines, that is, the weighted sum of displacements experienced by any airline would be proportional to its number of flights scheduled at airport. To maximise inter-airline equity, they first minimise the largest airline disutility, then the second largest, etc.

Finally, Manley and Sherry (2010) defined a passenger equity metric as the negative logarithm of the ratio of passenger delays for a given airport category over the total GDP passenger delay, divided by the ratio of the number of passengers from that airport category over all passengers in the GDP. This implies that the more passengers an airport category has, the more passenger delay it should be assigned. In both cases, perfect equity is represented as 0.

Table I summarises the main approaches and metrics used in literature for fairness models. It can be seen that the most classical choice is flight-centred, though inter-airline equity has been more considered in the past years. Passenger-based fairness metric is more complex to define and to use due to the lack of available data at operational level. In this paper, two definitions

¹ of fairness are considered: a flight-centred approach based on the RBS solution, and a passenger-centred metric based on
² the total delay experienced by passengers with respect to an RBS solution. The flight-centred one will be used to study the
³ classical trade-off between efficiency and equity at flight level. The passenger-centred one helps to fill the gap identified for
⁴ passenger-based fairness metrics, and it will be used to study the trade-off between efficiency and equity, at passenger level.

III. Delay and fairness optimisation models

⁶   This section introduces the mathematical model used in this paper for the optimisation and for the analysis of the trade-off
⁷ between efficiency (delay generated in the system) and fairness (based on the delay experienced by stakeholders). A description
⁸ of *price of fairness* and *price of efficiency* is provided as a tool to understand these trade-offs. The different objective functions
⁹ considered are explored in detail with particular focus on the definition of metrics (and objectives functions) for fairness.

A. Optimisation model

¹¹   As presented in Section I, the assignment of delay to flights in order to deal with capacity-demand imbalances can be seen
¹² as an optimisation problem. Different optimisation techniques can hence be considered. In this paper, delay assignment is
¹³ optimised using a deterministic ground holding problem model (GHP) based on (Ball et al., 2007). In this model, constraints
¹⁴ only apply at the destination to ensure that airport capacities are maintained.

¹⁵   For a given set of time intervals $(t = 1, 2, \ldots, T)$ corresponding to the times of arrival, and a set of aircraft $(f = 1, 2, \ldots, F)$
¹⁶ corresponding to flights that will arrive and then depart from the studied airport, the following inputs are defined:

¹⁷ • b_t is the constrained airport arrival capacity at time interval t and
¹⁸ • $STA(f)$ (scheduled time of arrival) is the earliest time interval at which aircraft f is scheduled to arrive at the constrained
¹⁹   destination airport.

²⁰ To prevent a flight from getting assigned a slot earlier than the earliest time it could arrive, for each flight f, its feasible time
²¹ intervals start at $STA(f)$. The decision variables are defined as:

$$x_{ft} = \quad 1 \text{ if aircraft } f \text{ is assigned to arrive at time interval } t$$
$$0 \text{ otherwise}$$

²² the deterministic ground holding problem can then be formulated as

$$\min \sum_f \sum_t c_{ft} x_{ft} \tag{1}$$

$$\text{subject to} \quad \sum_f x_{ft} \leq b_t, \text{ for all } t \tag{2}$$

$$\sum_t x_{ft} = 1, \text{ for all } f \tag{3}$$

²³ where c_{ft} is the cost of assigning aircraft f to arrive at time t. Note that Equation (2) corresponds to the capacity constraint
²⁴ applied at each time interval t, whereas Equation (3) imposes the fact that a flight must arrive exactly once. More details on
²⁵ this general GHP model can be found for example in (Ball et al., 2007).

²⁶   With the above defined model, it is possible to consider different objective functions by modifying the associated *cost* of
²⁷ assigning a given flight (f) to a given arrival slot (t).

B. Objective functions

²⁹   Optimisation of system performance (total experienced delay) and fairness are modelled. That is, objective functions for
³⁰ these categories are defined for both flights and passengers.

³¹   *1) System performance objective functions (Obj_{Perf}):* As shown in (Cook et al., 2012) and (Montlaur and Delgado, 2017),
³² the delay and cost experienced by passengers differ from the ones obtained with flight-centred metrics. These differences are
³³ partially due to passenger missed connections. For this reason, optimisation functions considering passenger and flight delays
³⁴ are considered in this paper. Given the previously mentioned importance of reactionary delay, total delay (as opposed to just
³⁵ arrival delay) is considered here in both cases. Two objective functions measuring the performance of the system through the
³⁶ total delay are thus modelled as follows:

³⁷ • *Total flight delay ($Obj_{Perf_{Flight}}$):* it includes the reactionary delay, which is defined as the difference between the Actual
³⁸   Time of Departure (ATD) and the subsequent Scheduled Time of Departure (STD). As seen on Figure 1b, it is equivalent
³⁹   to calculate it as the difference between the Actual Arrival Time (ATA), which corresponds to the variable t, and the
⁴⁰   latest time of arrival (LTA) that would not generate delay in the subsequent departure flight of the same aircraft. As
⁴¹   shown on Figure 1a, LTA is calculated as follows: $LTA(f) = STD(f) - MTT(f)$, where $MTT(f)$ is the minimum
⁴²   turnaround time needed for aircraft f; details of how these data were obtained from statistics are found in (Montlaur and

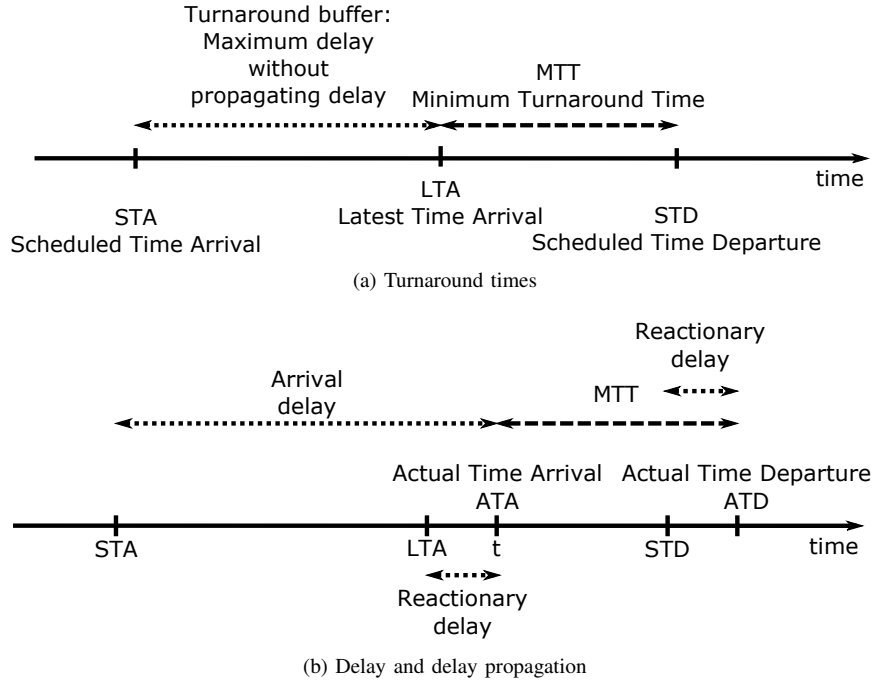(a) Turnaround times

(b) Delay and delay propagation

Fig. 1: Turnaround and delay diagram

Delgado, 2017). Note that in this case, delay is only propagated if the time between $STD(f)$ and $LTA(f)$ is lower than $MTT(f)$, which is a conservative approach since in reality, the turnaround might take longer than $MTT(f)$.

Finally, this objective function is defined as the sum of the arrival delay plus the reactionary delay multiplied by a factor 1.8, corresponding to the extra delay that this reactionary delay will further generate. As reported by (EUROCONTROL, 2015b), in 2014, the ratio of reactionary to primary delay was 0.80, which means that, on average, every minute of primary delay resulted in some additional 0.80 minutes of reactionary delay. Thus, in this model, the total delay per flight is minimised:

$$Obj_{Perf_{Flight}} = \sum_f \sum_t [(t - STA(f)) + 1.8(t - LTA(f))]x_{ft}, \qquad (4)$$

- *Total Pax delay ($Obj_{Perf_{Pax}}$)*: in this objective function, the total delay per passenger to minimise is expressed as follows:

$$Obj_{Perf_{Pax}} = \sum_f \sum_t [Pax_{arr}(f)(t - STA(f)) + Pax_{connec}PropagDelay(f,t) + Pax_{dep}(f)(t - LTA(f))]x_{ft}, \qquad (5)$$

where $Pax_{dep}(f)$ is the number of departure passengers assigned to flight f, and $Pax_{connec}PropagDelay(f,t)$ the propagation delay for each flight, taking into account the number of passengers connecting on the inbound flights and the waiting time at the hub, if the connections are missed, until another flight to their final destination is available. This is a probabilistic model that does not represent individual passenger itineraries. See Section IV and (Montlaur and Delgado, 2017) for more details on this parameter.

2) *Fairness objective functions (Obj_{Fair}):* Two fairness objective functions are considered: one focuses on the deviation of flight arrival time with respect to RBS, whereas the second one focuses on the deviation of the total Pax delay for each flight with respect to the one obtained with an RBS solution.

- *Total flight arrival delay deviation from RBS ($Obj_{RBS_{Flight}}$)*: similarly to (Glover and Ball, 2013) and (Kuhn, 2013), equity is defined here as the total deviation of arrival delay experienced by flights with respect to the RBS solution. Therefore, the deviation with RBS is minimised when minimising:

$$Obj_{Fair_{Flight}} = \sum_f \sum_t |t - RBS(f)|x_{ft} \qquad (6)$$

being $RBS(f)$ the arrival time of flight f in a RBS environment.

- *Total Pax delay deviation from RBS ($Obj_{Fair_{Pax}}$)*: the metric of fairness for passengers defined here considers the total delay deviation that passengers experience due to an assignment of slots different from the RBS solution. Note that, in this case, the deviation is considered based on the delay that would be assigned in the current operations, where RBS is

considered a *fair* assignment, even if this fairness is based on a flight-centric approach. The passenger fairness defined in this paper hence captures the variation with respect to current operations and not necessarily with respect to an even distribution of delay for passengers. The total deviation experienced by Pax is defined as follows:

$$Obj_{Fair_{Pax}} = \sum_f \sum_t |[Pax_{arr}(f)(t - STA(f)) + Pax_{connec}PropagDelay(f,t) + Pax_{dep}(f)(t - LTA(f))] \quad (7)$$

$$-[Pax_{arr}(f)(RBS(f) - STA(f)) + Pax_{connec}PropagDelay(f,RBS(f)) + Pax_{dep}(f)(RBS(f) - LTA(f))]|x_{ft},$$

C. Multi-objective optimisation

The model described in this paper considers trade-offs between system performance and fairness. Therefore, the optimisation problem can be defined as a multi-objective optimisation. This optimisation objective (cost to be minimised) can be described as a single optimisation function defined as the weighted sum of the individual objectives ($F_i(x)$): $\sum_{i=1}^{k} w_i F_i(x)$, see for example (Marler and Arora, 2010). Therefore, here, the single optimisation function considered is defined as:

$$Obj(\alpha) = \alpha Obj_{Fair} + (1 - \alpha)Obj_{Perf} \quad (8)$$

where Obj_{Fair} captures the fairness (as the deviation with respect to RBS for flight and as the deviation on the amount of delay experienced by passenger with respect to RBS solution) and Obj_{Perf} captures the performance (as the total delay).

When considering well-understood preferences, paired comparison methods can be used to define the value of the different weights and unrestricted positive weights should be used (Marler and Arora, 2010). However, when presenting Pareto solutions for an *a posteriori* articulation of preferences, a convex combination of functions can be used, implying that $\sum_{n=1}^{k} w_n = 1$. In this model: $0 \leq \alpha \leq 1$. This approach can be considered as a statement of *fairness* preference as:

- α indicates the preference of fairness, i.e., reduction of difference from the RBS solution,
- $1 - \alpha$ indicates the relevance of flight or passenger delay in the optimisation.

When computing different Pareto solutions to provide an *a posteriori* articulation of preferences, a systematic variation of weights does not necessarily ensure an even distribution of Pareto optimal points and hence an accurate complete representation of the Pareto optimal set (Marler and Arora, 2004). One way to improve this consists in using an objective function transformation method, being the *upper-lower-bound approach* the most robust one (Marler and Arora, 2005). With this method, instead of using $F_i(x)$, single objectives are transformed as:

$$F_i^{trans} = \frac{F_i(x) - F_i^0}{F_i^{max} - F_i^0} \quad (9)$$

This approach, referred to as normalisation, generally leads to F_i^{trans} ranging between zero and one, depending on the accuracy and method with which F_i^{max} and F_i^0 are determined. As described in (Marler and Arora, 2005), the best approach, and the one used here, is to select $F_i^0 = \min\{F_i(x)|x \in X\}$ (being X the feasible design space of the problem), and F_i^{max} as the *Pareto maximum* by defining it such that $F_i^{max} = \max_{1 \leq j \leq k} F_i(x_j^*)$, where x_j^* is the point that minimises the j^{th} objective function.

In the present case, the transformation will be:

$$Obj(\alpha) = \alpha \frac{Obj_{Fair} - Obj_{Fair}^0}{Obj_{Fair}^{max} - Obj_{Fair}^0} + (1 - \alpha) \frac{Obj_{Perf} - Obj_{Perf}^0}{Obj_{Perf}^{max} - Obj_{Perf}^0} \quad (10)$$

In this paper, there are four objective functions defined (two for performance and two for fairness): flight performance (total flight delay), passengers performance (total passenger delay), flight fairness (flight delay deviation with respect to RBS) and passengers fairness (passenger total delay deviation with respect to RBS). We are interested in analysing the system trade-offs between performance and fairness, but instead of considering all four objectives at the same time, we define two case studies:

1) Case 1: Flight-centric optimisation: In this case, the optimisation considers a flight-centric view where performance and fairness are defined according to flight metrics. This case is selected as it is closer to what the ATM system is classically considering. It also has the advantage that the data required to perform this optimisation is available and already used by the community.

$$Obj_{Flight}(\alpha) = \alpha \frac{Obj_{Fair_{Flight}} - Obj_{Fair_{Flight}}^0}{Obj_{Fair_{Flight}}^{max} - Obj_{Fair_{Flight}}^0} + (1 - \alpha) \frac{Obj_{Perf_{Flight}} - Obj_{Perf_{Flight}}^0}{Obj_{Perf_{Flight}}^{max} - Obj_{Perf_{Flight}}^0} \quad (11)$$

2) Case 2: Passenger-centric optimisation: As mentioned trough the paper, passenger experience might differ from flight-centric metrics. For this reason, it is important to capture passenger performance and fairness with their own dedicated metrics. In this second case study, we consider the trade-off between passenger performance and fairness in order to analyse the benefits (or losses) for passengers if they are explicitly considered in a passenger-centric system. Note that in this case, the data required is not fully available to the network manager and hence, the research shows what can be achieved if modelled this way.

$$Obj_{Pax}(\alpha) = \alpha \frac{Obj_{Fair_{Pax}} - Obj^0_{Fair_{Pax}}}{Obj^{max}_{Fair_{Pax}} - Obj^0_{Fair_{Pax}}} + (1 - \alpha) \frac{Obj_{Perf_{Pax}} - Obj^0_{Perf_{Pax}}}{Obj^{max}_{Perf_{Pax}} - Obj^0_{Perf_{Pax}}} \tag{12}$$

Finally, note that results on passenger metrics will be described when optimising the system in a flight-centric approach (Case 1) and vice versa. This will be done in order to understand the impact of the optimisation on both stakeholders and with all four metrics.

D. Price of fairness and efficiency

The model presented in the previous subsection (Section III-C) allows us to compute the Pareto between efficiency and fairness in absolute value. However, specific metrics can be considered to capture those relationships. Trade-off between fairness (or equity) and efficiency has been studied by several researchers in the past years. For example, Bertsimas et al. (2012) aim to balance efficiency and fairness in the context of resource allocation. They identify the notion of α-fairness, which allows the decision maker to trade off efficiency for fairness by means of a single parameter. In (Bertsimas et al., 2011), they introduce the concept of *price of fairness* as the efficiency loss due to the increment of fairness and *price of efficiency* as the fairness loss due to the increment in efficiency in the system. We define similar concepts for our specific problem (considering the use of α as an indicator of preference of fairness).

1) Price of fairness: Bertsimas et al. (2011) introduce the concept of *price of fairness* (POF) as the efficiency loss relative to the maximum system efficiency:

$$POF(U, \alpha) = \frac{SYSTEM(U) - FAIR(U, \alpha)}{SYSTEM(U)} \tag{13}$$

where $SYSTEM(U)$ is the maximum system efficiency, and $FAIR(U, \alpha)$ the efficiency under the fair scheme. POF is thus a number between zero and one, corresponding to the percentage of efficiency loss, compared to the maximum system efficiency, due to the introduction of fairness. In (Montlaur and Delgado, 2018), a similar definition of *price of fairness* was introduced and expanded into *price of trade-off* as the loss in system performance due to applying trade-off in a multi-optimisation environment.

In the present study, we define the *price of fairness* (POF) as the percentage of performance loss due to the consideration in the optimisation of fairness, that is:

$$POF(\alpha) = \frac{Opt_{Perf} - Perf(\alpha)}{Opt_{Perf}} \tag{14}$$

In particular, we distinguish the *price of fairness for flights* (POF_{Flight}) as the percentage of efficiency loss for flights due to the consideration in the optimisation of fairness (by increasing the values of α), and similarly the *price of fairness for passengers* (POF_{Pax}) as the percentage of efficiency loss for passengers due to the consideration in the optimisation of fairness:

$$POF_{Flight}(\alpha) = \frac{Opt_{Perf_{Flight}} - Perf_{Flight}(\alpha)}{Opt_{Perf_{Flight}}} \tag{15}$$

$$POF_{Pax}(\alpha) = \frac{Opt_{PerfPax} - Perf_{Pax}(\alpha)}{Opt_{PerfPax}} \tag{16}$$

where:

- $Perf_{Flight}(\alpha)$ is the value of the total delay of flights for any given value of α, which corresponds to the $Obj_{Perf_{Flight}}$ term of Equation (4). Its optimum value is noted $Opt_{Perf_{Flight}}$, and corresponds in our case to the value of $Obj^0_{Perf_{Flight}}$ (see Equation (11)).
- $Perf_{Pax}(\alpha)$ is the value of the total delay of passengers for any given value of α, which corresponds to the $Obj_{Perf_{Pax}}$ term of Equation (5). Its optimum value is noted $Opt_{Perf_{Pax}}$, and corresponds in our case to the value of $Obj^0_{Perf_{Pax}}$ (see Equation (12)).

In our case, the maximum system efficiency corresponds to a minimum value of the objective function, thus POF_{Flight} and POF_{Pax} are negative numbers. They, however, also correspond to the percentage of efficiency loss compared to the maximum system efficiency. The best outcome possible for performance ($\alpha = 0$) corresponds to a respective zero price of fairness.

2) Price of efficiency: We now define $Fair(\alpha)$ as the value of the deviation from RBS (for flights or Pax) for any given value of α; its optimum value $Opt_{Fair} = 0$ is obtained for $\alpha = 1$. This corresponds to the Obj_{Fair} term of Equation (8).

The fairness metric adopted in this work is the deviation from Opt_{Fair}. It is minimum for the α-fair allocation corresponding to $\alpha = 1$, and maximum for $\alpha = 0$. As more emphasis is put on fairness (e.g., by selecting a higher value of α), the maximum deviation from RBS is likely to decrease. The fairness gain is now the difference between the fairness metric evaluated at $\alpha = 1$ and the general fairness metric. Bertsimas et al. (2012) call *price of efficiency* the fairness loss (difference between the fairness metric evaluated at the max-min fair allocation and the α-fair allocation) relative to the maximum value of the fairness metric. Here, the fairness loss is also defined as the difference between minimum value of the fairness metric (Opt_{Fair}) and the equivalent of the α-fair allocation $Fair(\alpha)$, but because here $Opt_{Fair} = 0$ the price of efficiency is obtained normalising the fairness loss with respect to its maximum value, that is:

$$POE(\alpha) = \frac{|Opt_{Fair}| - |Fair(\alpha)|}{|Max(Fair)|} = \frac{-|Fair(\alpha)|}{|Max(Fair)|} \tag{17}$$

Again, prices of efficiency are defined both for flights and passengers:

$$POE_{Flight}(\alpha) = \frac{|Opt_{FairFlight}| - |Fair_{Flight}(\alpha)|}{|Max(Fair_{Flight})|} = \frac{-|Fair_{Flight}(\alpha)|}{|Max(Fair_{Flight})|} \tag{18}$$

$$POE_{Pax}(\alpha) = \frac{|Opt_{FairPax}| - |Fair_{Pax}(\alpha)|}{|Max(Fair_{Pax})|} = \frac{-|Fair_{Pax}(\alpha)|}{|Max(Fair_{Pax})|} \tag{19}$$

where:

- $Fair_{Flight}(\alpha)$ is the value of the sum of the absolute deviation of delay with respect to RBS for flights, which corresponds to $Obj_{FairFlight}$ in Equation (6). Its optimum value is noted $Opt_{FairFlight}$, and corresponds in our case to the value of $Obj^0_{FairFlight}$, which is equal to 0, whereas $|Max(Fair_{Flight})|$ corresponds to $Obj^{max}_{FairFlight}$ (see Equation (11)),

- $Fair_{Pax}(\alpha)$ is the value of the deviation of total delay of passengers with respect to RBS, which corresponds to $Obj_{FairPax}$ in Equation (7). Its optimum value is noted $Opt_{FairPax}$, and corresponds in our case to the value of $Obj^0_{FairPax}$, which is equal to 0, whereas $|Max(Fair_{Pax})|$ corresponds to $Obj^{max}_{FairPax}$ (see Equation (12)).

Once again, because the maximum fairness corresponds to a zero value of $|Opt_{Fair}|$, POE is negative and corresponds to the percentage of fairness loss compared to the maximum fairness loss of the system. The most fair solution corresponds to a zero price of efficiency.

IV. PROBLEM DESCRIPTION

In order to test the previously defined metrics and to provide a testing environment for trade-offs between performance and fairness for flights and passengers, a specific traffic scenario is modelled. The arrival demand at Paris Charles de Gaulle (CDG) airport on September 12th, 2014 has been considered for all simulations. It was a busy Friday without any major disruption. Morning traffic, between 5.00 GMT and 11.00 GMT, is analysed. For the traffic scheduled, data from EUROCONTROL Demand Data Repository 2 (DDR2) (EUROCONTROL, 2018) have been used. All details of the problem simulated here can be found in (Montlaur and Delgado, 2017). Next comes a summary of the main hypothesis. This traffic scenario is representative of operations at a hub disrupted by an ATFM regulation.

During the 6 hours of study, the total number of aircraft scheduled to arrive at CDG is 285. The hypothesis that every arriving aircraft will eventually depart is made. Considering the demand data and historic regulations at CDG, an ATFM regulation between 6.00 GMT and 8.00 GMT is modelled. A nominal capacity of 80 arrivals per hour is considered when no regulation is applied, and the regulated capacity is decreased to 40, which is a possible value of capacity during regulated periods at CDG as shown in the DDR2 dataset (EUROCONTROL, 2018). For this pre-tactical optimisation, slot windows of 15 minutes are considered, i.e., 20 (nominal) or 10 (regulated) aircraft every 15 minutes.

For each flight, the type of aircraft has been considered and the number of passengers in each flight has been estimated as a function of the maximum capacity of the aircraft. A triangular distribution has been used to allocate passengers between 60 and 95% of the maximum capacity, with the peak of the distribution at 85%, which is considered as the target average load factor.

As mentioned in Section III-B, the propagation of passenger delay, due to missed connections at the hub, has been modelled by simulating the number of connecting passengers on inbound flights and the waiting time at the hub, if the connections are missed, until another flight to their final destination is available. These data are based on the statistical analysis of a day of itineraries at the hub from individual passengers' itineraries developed in SESAR WP-E ComplexityCost project (Cook et al., 2016; Delgado et al., 2016). The obtained data correspond to the parameter $Pax_{connec}PropagDelay(f,t)$ in Equation (5). This passenger allocation process leads to a total of 39 820 arrival passengers, from which 8 620 are connecting to following flights (21.6% of arrival passengers) and a total of 39 671 departure passengers, during the 6-hour study. Note thus that here, even though it would not be the case in operational systems, all passenger information needed for the optimisation process is considered to be available to the ATC/ATFM at execution level.

TABLE II: Minimum and maximum values of each parameter in [min] for each case of optimisation

		Pareto maximum values		Pareto optimal values	
		$Obj_{Fair_{Flight}}$	$Obj_{Perf_{Flight}}$	$Obj_{Fair_{Pax}}$	$Obj_{Perf_{Pax}}$
Case 1 Flight-centric:	min	0	2553.4	0	429 630
	max	2027	3774.8	422 787	589 102
Case 2 Pax-centric:	min	0	3405.8	0	239 594
	max	2601	3900.8	560 868	589 102

To model the reactionary delay, scheduled turnaround times (STT) and minimum times required to do the turnaround process (MTT) have been computed for each flight based on statistic data at CDG for different types of aircraft.

V. RESULTS

We calculate the Pareto optimal points needed for computation of the Pareto extreme points required in Equations (11) and (12). These optimal points (lightly shaded boxes in Table II) are defined as the points minimising each objective, that is, $\alpha = 1$ for Obj_{Fair}, $\alpha = 0$ for Obj_{Perf}. The darker boxes indicate the maximum of each function. Note that, as shown in Table II, optimal and maximum values are calculated for the chosen optimisation case (Case 1 for flight-based, Case 2 for passenger-based). That is, smaller or larger values of the objective function could be obtained in another optimisation scheme. For example, it can be seen that the maximum value considered for $Obj_{Per_{flight}}$ to define $Obj_{Perf_{Flight}}^{max}$ and $Max(Fair_{Flight})$ will be 3774.9 min, since this is the maximum value obtained in the flight-based optimisation case, but it could reach 3900.8 min in the passenger-based optimisation, i.e., if the system is optimised considering passenger delay and fairness (Case 2). Flight delay could then be larger than the maximum delay observed when the system is optimised considering flight delay and fairness (Case 1). This will be further commented, when studying the obtained results, in the next sections.

A. Optimisation model trade-offs

In this section, we first study the trade-off between fairness and efficiency for the flight-based case, and the effect of flight-based optimisation on passenger metrics. In the second case, we study the contrary.

1) Case 1: Optimising for flight metrics: We first study the trade-off between flight fairness and flight delays. Figure 2a shows that the minimum flight total delay is located at 2553 min for $\alpha = 0$. If some fairness is introduced, by increasing the value of α, deviation from the RBS solution can decrease from 2027 min to 1384 min (32% decrease) at a cost of increasing the flight delay by less than 4%. Once this point is reached, further reducing flight delay deviation comes at higher cost for flight delay.

We then see how optimising for flight objectives affects passenger metrics. In Figure 2b, the point corresponding to fairness only ($\alpha = 1$) can be identified (3774.8 min of flight delay and 589 102 min of passenger delay). When decreasing the value of α and thus increasing the performance importance, it can be seen that flight and passenger delays behave similarly. When improving performance, i.e., decreasing the flight delay, the passenger total delay decreases as well. This highlights the importance of non-connecting passengers on the total delay experienced. Figure 2c shows that, apart from the extreme point corresponding to $\alpha = 1$, increasing the passenger fairness with respect to the RBS solution comes at a relatively low cost for flight delay. Deviation of Pax delay from RBS solution can be reduced by 30% (from 422 787 min to 297 600 min) at a cost of a 3% increase of flight delay (from 2553 min to 2636 min). This can be explained by the fact that though this fairness parameter is taken into account from the passengers point of view, it is related to the RBS solution that is flight-centric. It can be confirmed by Figure 2f that shows how fairness parameters for flights and for passengers vary in the same direction when optimising for flights.

Finally, Figures 2d and 2e, respectively show trade-off between Pax delay and Pax delay deviation, and Pax delay and flight delay deviation when optimising for flight metrics. Figure 2d shows that decreasing Pax delay results in worsening Pax fairness, as it will also be observed when optimising for Pax metrics in Section V-A2. Nevertheless, higher gain can be obtained in Pax fairness (decreasing Pax deviation by 49%, from 422 787 min to 214 500 min) at a low increment of Pax delay (from 459 200 min to 491 700 min, that is, a 7% increase). Similar results are obtained from Figure 2e, higher gain can be obtained in flight fairness, at a relatively low increment of Pax delay.

The analyses presented in this section indicates that when optimising for flight metrics, it is worth introducing some measure of fairness, because it only slightly increases flight delay, and the penalisation of Pax metrics is also relatively low. However, a solution focused merely on fairness, i.e., RBS, leads to significant sub-optimal assignment of total delay (when reactionary delay and passenger connections are considered) for both flights and passengers.

2) Case 2: Optimising for passengers metrics: We first look at the trade-off between Pax fairness and delay. It can be seen in Figure 3d, that increasing fairness costs more at passenger level than at flight level. Indeed, once the minimum Pax total delay has been identified (239 594 min for $\alpha = 0$), introducing some fairness by increasing α comes at a higher cost than what was previously observed for the flight-based case: a decrease of Pax deviation from the RBS solution of 32% now comes at
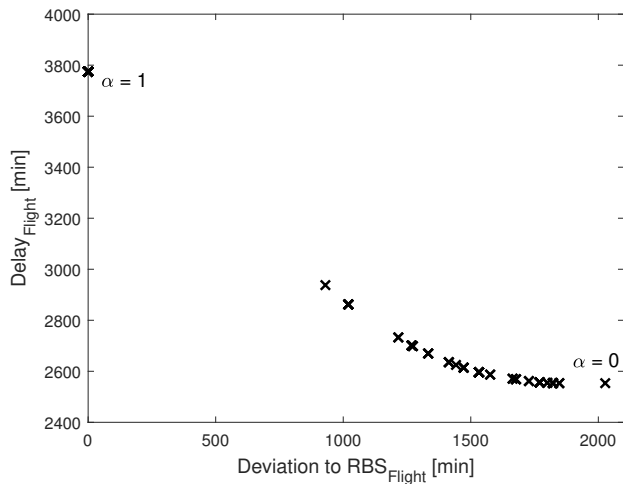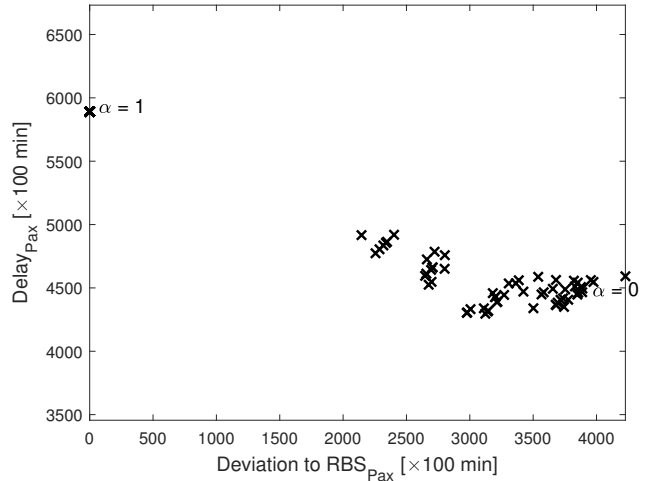
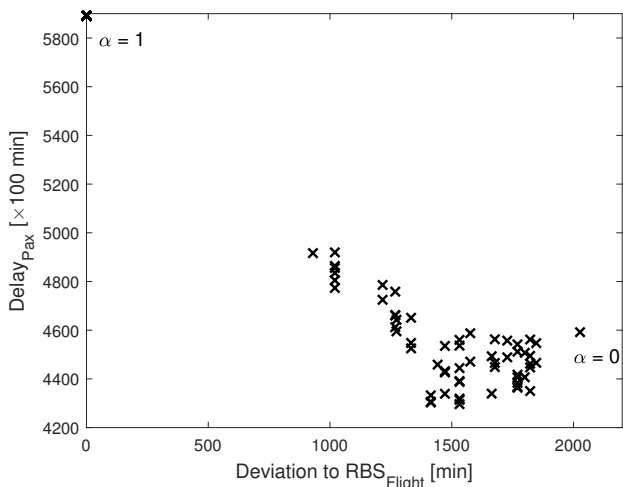(a) Flight delay vs. deviation of flights from RBS
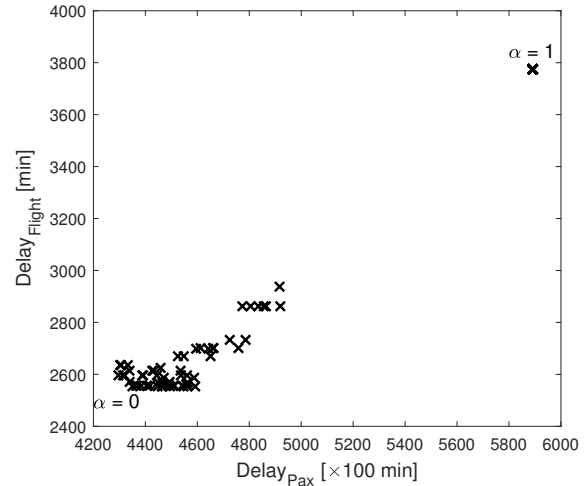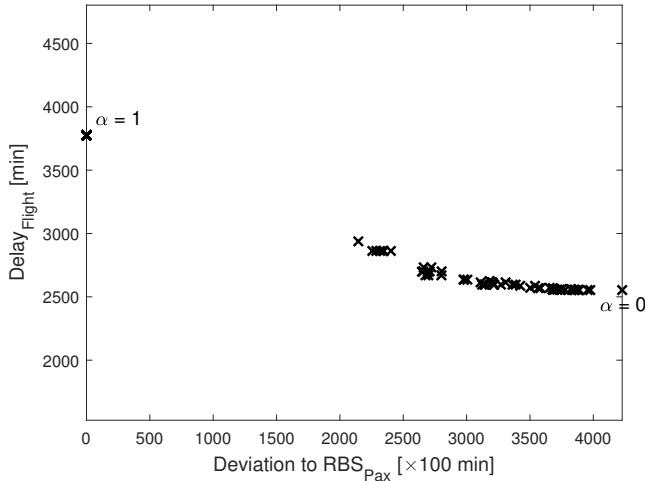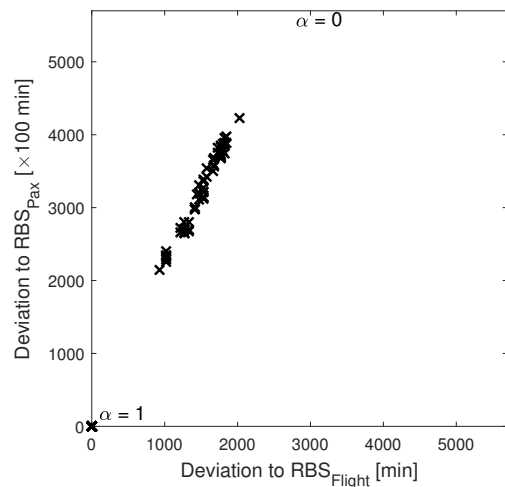
(b) Flight delay vs. Total Pax delay

(c) Flight delay vs. deviation of Total Pax delay from RBS

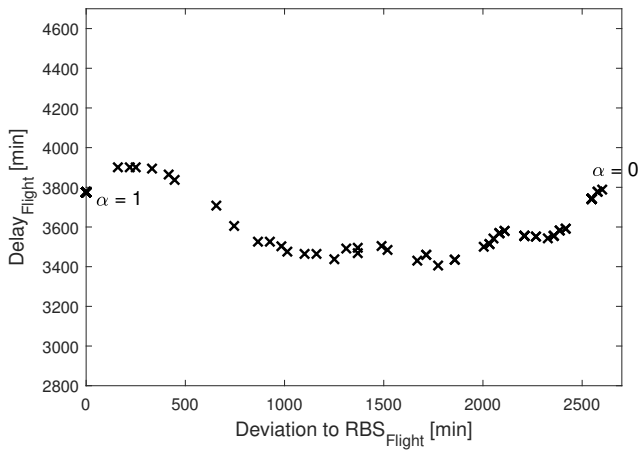(d) Total Pax delay vs. deviation of Total Pax delay from RBS

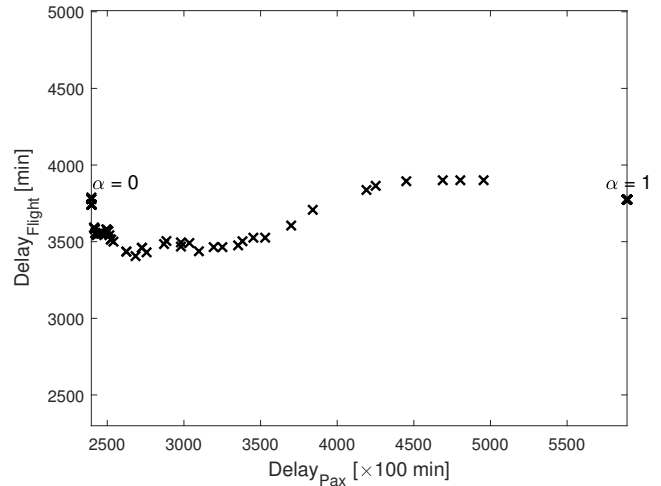(e) Total Pax delay vs. deviation of flights from RBS

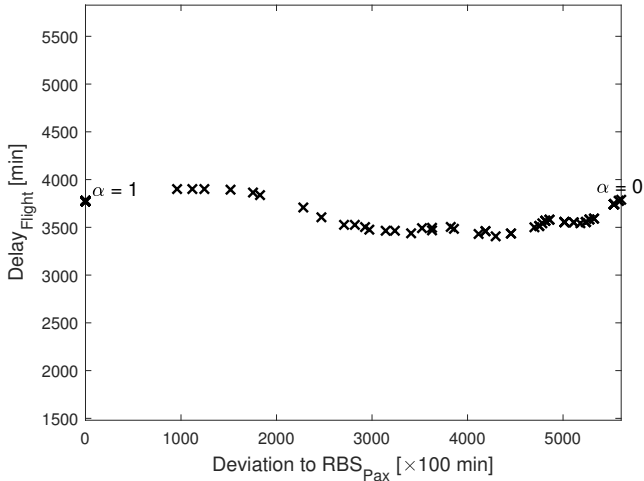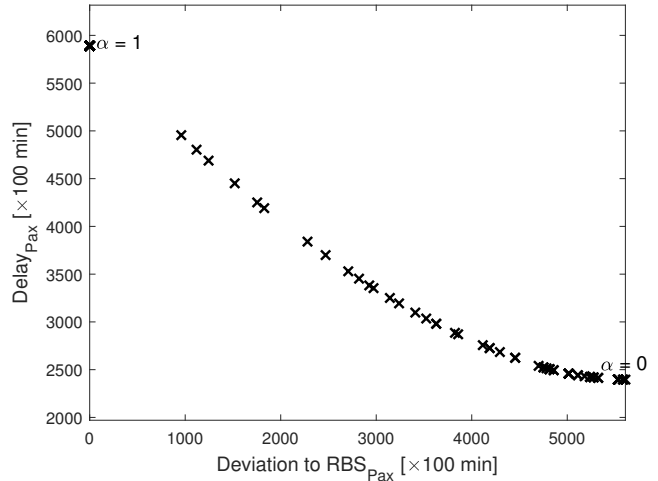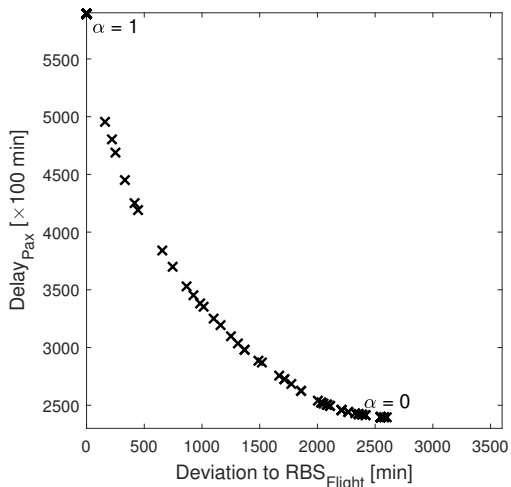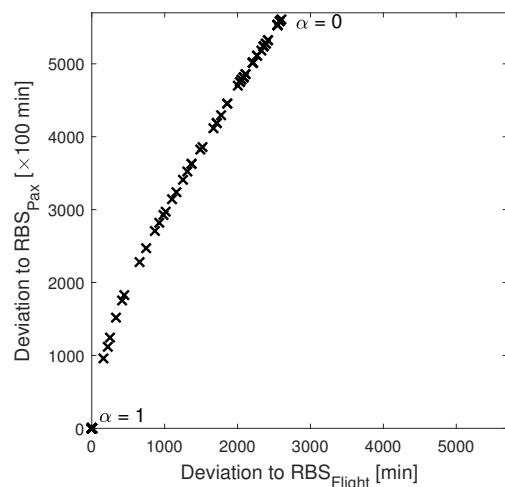(f) Deviation of Total Pax delay from RBS vs. deviation of flights from RBS

Fig. 2: Trade-offs between performance and fairness when minimising flight delay vs. minimising deviation of flights from RBS (case 1)

(a) Flight delay vs. deviation of flights from RBS

(b) Flight delay vs. Total Pax delay

(c) Flight delay vs. deviation of Total Pax delay from RBS

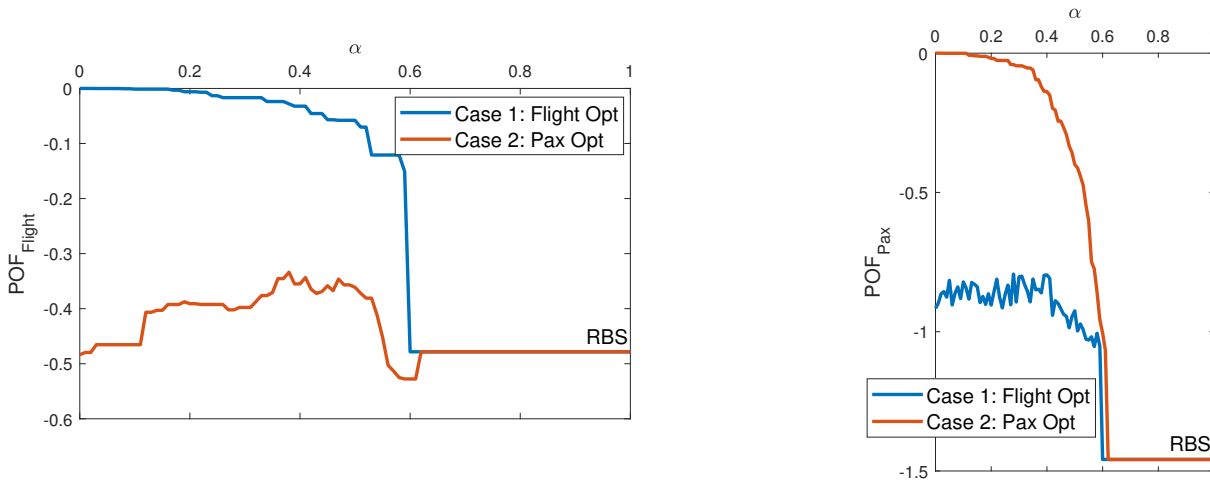(d) Total Pax delay vs. deviation of Total Pax delay from RBS

(e) Total Pax delay vs. deviation of flights from RBS

(f) Deviation of Total Pax delay from RBS vs. deviation of flights from RBS

Fig. 3: Trade-offs between performance and fairness when minimising Total Pax delay vs. minimising deviation of Total Pax delay from RBS (case 2)

(a) Price of fairness for flights as a function of α (reference value of 0 is obtained for $\alpha = 0$, $Opt_{Flight} = 2553.4$ min)

(b) Price of fairness for Pax as a function of α (reference value of 0 is obtained for $\alpha = 0$, $Opt_{Pax} = 239\,594$ min)

Fig. 4: Price of fairness as a function of α

a cost of 20% increase for the Pax total delay. Further reducing the fairness metric after this point comes at an even higher cost as shown in Figure 3d.

Next, we want to study how optimising for Pax-related objectives affects flight metrics. Figures 3a, 3b and 3c respectively show how, when optimising with respect to Pax metrics, changes in deviation of flight delay, total Pax delay and deviation of Pax delay hardly affect values of total flight delay. This result indicates that more focus can be put on Pax metrics without penalising too much total flight delay. A trade-off between total Pax delay and deviation of flight from RBS is nevertheless observed in Figure 3e, similar to the one observed between total Pax delay and deviation of Pax delay from RBS: the reduction of total Pax delay comes at a cost of reduction of fairness for flights. This similarity of results is confirmed by Figure 3f, which shows that deviations to RBS of flight delay and of Pax delay behave similarly when taking into account fairness.

Case 2 shows that when optimising for Pax metrics, introducing some measure of fairness, penalises more Pax delay and flight metrics than what was observed in Case 1.
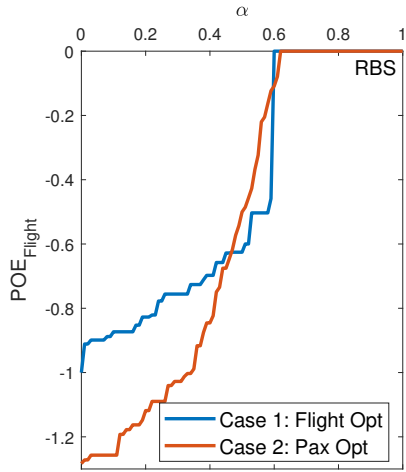
B. Price of fairness and efficiency

As previously discussed, the introduction of *price of fairness* and *price of efficiency* concepts help decision makers to understand the relative changes between efficiency and fairness for the different parameters that can be considered in the optimisation. In our case, we present these indicators for flights and passengers, and the impact of optimising considering only one stakeholder (Case 1 and Case 2) on the performance and fairness of the other.
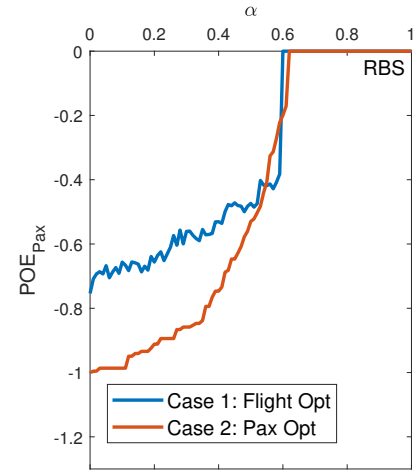
1) Price of fairness: Figure 4 presents the evolution of price of fairness (POF) as a function of α. This representation allows decision makers to quantify the impact of including fairness on delay performance.

Figure 4a shows the price of fairness for flight (POF_{Flight}), when optimising for flight (Case 1) and when optimising for Pax (Case 2). When optimising for flight, for $\alpha = 0$, a maximum POF_{Flight} is achieved and equal to 0, i.e., there is no loss on flight delay efficiency, which corresponds to a total flight delay of 2553.4 minutes. As α, and thus fairness, increases, POF_{Flight} decreases. However, the graph allows us to see that the evolution is at first very flat, pointing to the fact that we can gain in fairness (using higher α) with relatively small loss in efficiency (total flight delays). When α reaches a value close to 0.5, $POF_{Flight} \approx -0.1$, i.e., there is an increment in 10% of the total flight delay. The cost of increasing fairness from that point on is very high, rapidly reducing the performance of total flight delay. To see how much fairness is gained by increasing α, we will then have to check the price of efficiency in Section V-B2.

If the optimisation considers passenger delay and fairness (Case 2), it is observed in Figure 4a, that when focus is given to passenger delay (i.e., $\alpha = 0$), the flight performance suffers nearly a 50% increment on flight delay with respect to the minimum delay possible for flights obtained in Case 1. This decrease in total delay performance is expected as the optimisation focuses on passenger delay. Note however, that this decrease is even larger than the worst performance obtained in Case 1 with $\alpha = 1$, i.e., for maximum fairness. Another interesting fact is that as fairness for passenger is increased, the performance for flights does not get further reduced but instead improves. This indicates that focusing on increasing passenger fairness (e.g., with $\alpha = 0.5$) achieves a reduction in flight performance that is lower than if focus is given to passenger performance. From that point on, increasing the passenger fairness impacts negatively the total flight delay.

(a) Price of efficiency for flights as a function of α (b) Price of efficiency for Pax as a function of α

Fig. 5: Price of efficiency as a function of α

Figure 4b shows the price of fairness from a passenger total delay perspective (POF_{Pax}), for both cases of optimisation. When optimising for Pax, for $\alpha = 0$, a maximum POF_{Pax} corresponding to zero loss on passenger delay efficiency is achieved. This corresponds to a total passenger delay of $239\,594$ min. But contrary to the study of price of fairness for flights, as soon as α increases, passenger performances are significantly affected. This points out the sensitivity of the optimisation on the performance experienced by passengers. The total delay that passengers experience is affected by the relative importance given to fairness. Finally, the performance from a passenger perspective can deteriorate up to nearly 150%, while for flights, the total delay would only increase up to around an extra 50% with respect to the best optimisation for flight. This means that greater savings can be done for passengers, but also, that they are more susceptible to experience worse performances. We could expect this, since fairness maintains flights close to their RBS, which is centred on minimising flight arrival delay.

Note also how, if the optimisation is performed focusing on flights, the performance for passengers is at a maximum around 80% worse than if the optimisation is done considering passengers. Hence, by not considering passengers in the optimisation, there is a worsening on the total delay of at least 80% with respect to the best values achieved with passengers. However, the optimisation also shows that the delay experienced by passengers is more stable when the optimisation focuses on flights. The performance of passengers remains flat up to around $\alpha = 0.5$ when it decreases. Therefore, if the optimisation is done considering only flight performance and fairness, one can increase the fairness up to around $\alpha = 0.5$ with a relatively small loss on flight performance and without significantly affecting passenger performance.

Another approach to the price of performance results is to consider, instead of the loss of performance with respect to the lowest amount of total delay (i.e., $\alpha = 0$), the gain in performance by increasing the relative importance of performance with respect to fairness. This is particularly interesting when considering Case 1. Note that in current operations, where RBS is the used practice, the maximum fairness would be equivalent to an optimisation of Case 1 with $\alpha = 1$, i.e., minimum deviation with respect to the delay assigned in RBS. We can then observe that flight performance (i.e., total delay experience by flights) is nearly 50% higher in delay than the best that can be achieved (see Figure 4a), and this is also producing the worst performance for passenger total delay (see Figure 4b). Hence, introducing the consideration of total flight delay (i.e., reducing α), one can increase the performance significantly. Fairness between $\alpha = 0.5$ and $\alpha = 1$ has a high impact on performance. This will have, however, an effect on fairness, which will be explored in the next section.

2) Price of efficiency: In the previous section, we have discussed how changes in fairness in the objective function impact the performance, i.e., the amount of delay that is generated. The price of efficiency helps the decision maker to understand the impact of α in the changes of fairness of the solution. Remember that, as seen in Equation (17), POE has been defined as the ratio of the deviation of the fairness obtained by the optimisation, with respect to the optimum value of fairness, and the maximum deviation from fairness that can be expected (i.e., with $\alpha = 0$). With this definition, $POE = 0$ means that the solution provides a deviation of zero with respect to RBS. Note that a $POE = -1$ indicates the maximum possible deviation with respect to RBS for each of the cases independently. However, as presented in Table II, when optimising in one case, it is possible to obtain higher deviations to RBS than in the other one (e.g., the maximum delay deviation with respect to RBS in Case 1 for flight fairness is 2027 min, however, when optimising in Case 2 (passenger-centric), the deviation for flights increases to 2601 min). This means that the price of efficiency could lead to a degradation larger than 100% in those cross evaluations.

Figure 5 shows that this worst performance, i.e., the largest deviation with respect to RBS and hence the least fair solution,
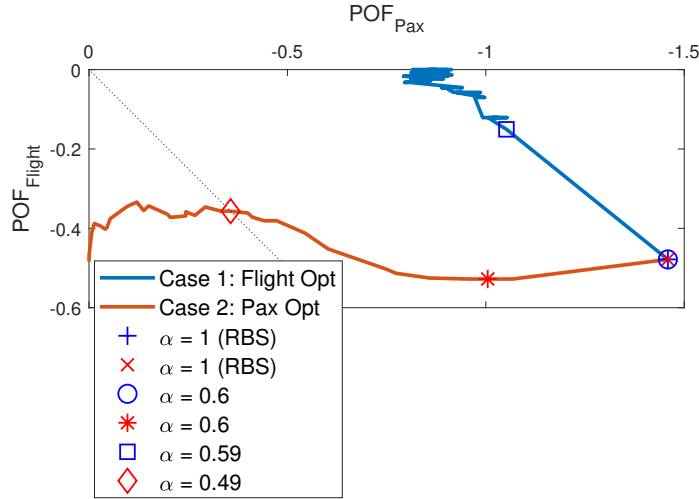
Fig. 6: Price of fairness for Flights vs Price of fairness for Pax

is obtained, as expected, when $\alpha = 0$ as the optimisation is only focusing on delay. As α increases, the deviation from RBS decreases and POE increases. The deviation with respect to RBS is initially relatively large, but is quickly reduced once α gets close to 0.6. This is consistent with the decrease in performance observed in the results presented in the previous section (Section V-B1).

As shown in Figure 5a, when optimising considering passenger delay and fairness (Case 2), the deviation of fairness for flight can be up to over 120%. Then flight fairness shows the same degradation when optimising for flights or for Pax when $\alpha \geq 0.5$. This degradation is removed once $\alpha = 0.6$. Therefore, it does not seem to be a benefit, from a fairness degradation point of view, to use values of α greater than 0.6. In Figure 5b, the maximum degradation for Pax fairness when optimising for flight (Case 1) is only around 70%, and as previously observed for POE_{Flight}, once α gets close to 0.6, Pax fairness deviations are similar for flight and for passenger optimisation, and no improvement is observed for $\alpha > 0.6$.

C. Discussion of results

In the previous sections, POF and POE for flights and passengers have been presented. In this section, the relationships between POF_{Flight}, POF_{Pax} and POE are jointly further explored.

Figure 6 shows the trade-off between POF_{Flight} and POF_{Pax} for each possible value of α, i.e., it presents the degradation on performance (total delay) for flights and passengers when fairness is considered. See how, most of the time (always in Case 1 and for $\alpha > 0.49$ in Case 2), better POF_{Flight} than POF_{Pax} is obtained. This shows once again how sensitive passenger delay is to the selection of parameters in the optimisation. The fact that fairness has been defined with respect to a flight-centric metric might improve the robustness of flight performances as α changes, whereas performance for passengers is worse when a solution closer to RBS is selected. Optimising considering only passenger delays (Case 2) provides similar reduction in performance for flights (fluctuating between -0.35 and -0.5) but can improve significantly passenger performances. That is, we can achieve almost the same total delay for flights while reducing significantly the total delay experienced by passengers. Whereas if only total delay for flights is optimised, passenger delay also improves, since reactionary delay is minimised, which affects passenger delay too. However, note that, if flights are optimised (Case 1) without considering passengers, then the maximum improvement possible for passengers is limited to a worsening of around 90% with respect to the best possible passenger delay achievable and that by default, if one does not want to compromise fairness, i.e., using $\alpha \geq 0.5$, then the passenger performance will be between -100% and -150% of the best achievable value. This highlights the importance of considering passenger delay when optimising the system.

It is worth noticing that current operations (RBS) provide a sub-optimal performance (the worst) for both flights and passengers. This points out the importance of considering an optimisation that might compromise some of the fairness allocation of resources but that can improve significantly the performance of the system. This is further explored next.

Figure 6 describes the impact of the optimisation parameters on the performance, but in order to quantify the impact on fairness, we need to refer to Figure 7. This figure shows trade-offs between gain in fairness (POE) and gain in efficiency (POF) for flights and passengers. Figure 7a shows that for most values of α ($\alpha \leq 0.59$ in Case 1, and $\alpha \leq 0.53$ in Case 2), $POE_{Flight} \leq POF_{Flight}$, that is, fairness performance is usually worse than performance of flight delay. Once again, if the
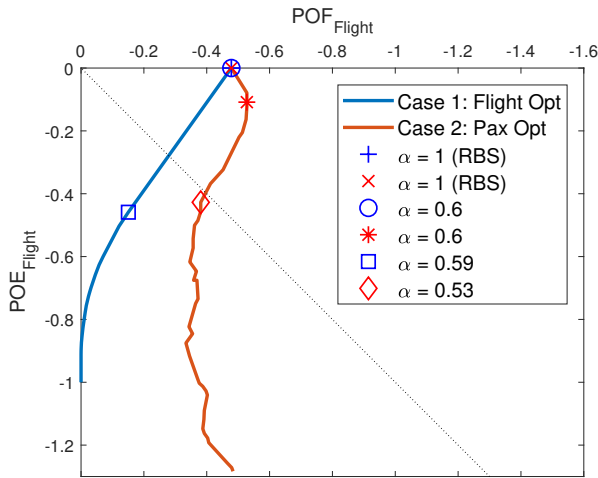
(a) Price of efficiency for flights vs. price of fairness for flight
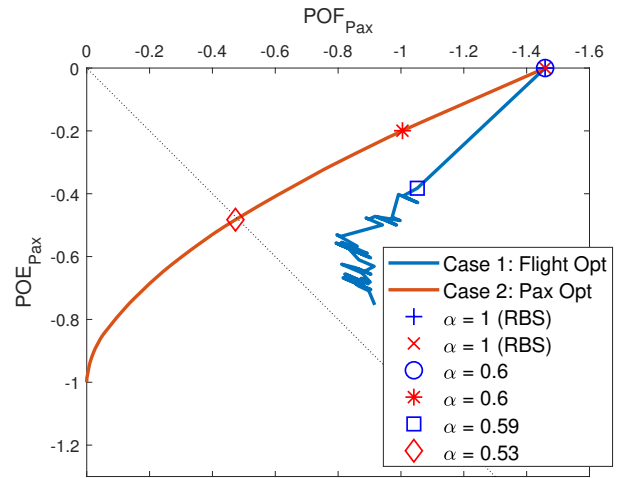
(b) Price of efficiency for Pax vs. price of fairness for Pax

Fig. 7: Price of efficiency vs Price of fairness

optimisation focuses on passengers (Case 2), a significant improvement on POE_{Fligh} can be achieved with a small impact on the total delay experienced by flights (POF_{Flight}).

If focus is given to the trade-offs between POF_{Pax} and POE_{Pax}, one can see in Figure 7b that total passenger delay is more sensitive to the parameters of the optimisation. Note how in Case 2, there is a significant continuous trade-off between the two. It is possible to improve the efficiency (POF_{Pax}) but always at the expense of decreasing the fairness (POE_{Pax}). This is also the case in the optimisation when flights are considered (Case 1). Improving passenger total delay comes at a cost on the fairness of the solution (as flights with more passengers or more connecting passengers will tend to be prioritised).

In current operations, flights are prioritised considering the RBS principle. As presented in Figure 6, the results on performance are compromised for both flight and passenger delay. If the optimisation is done considering passenger delay and fairness (Case 2), it is possible to use a lower value of α (e.g., $\alpha = 0.6$), which provides improvements for passenger performance without any negative impact on the flight total delay. RBS optimisation ensures the most fair assignment of delay, but as presented in Figure 7a, in Case 1, when the weighted optimisation has an impact on the outcome (i.e., $\alpha < 0.6$), it produces a significant change of flight fairness ($\alpha = 0.59$ already produces a reduction over 40% on the fairness), but also a significant impact on delay ($\alpha = 0.59$ improves the total flight delay performance from 48% of degradation of flight performance to 15% only). In Case 2, when passengers are optimised, the changes are smoother.

Finally, it is worth noticing how there is a significant difference in the performance experienced by flights and passengers when only one of the stakeholders is considered (see for example the trade-offs between Case 1 and Case 2 on performance in Figure 4 or the different impacts on fairness in Figure 5). This indicates that intermediate solutions might be achieved by considering the four parameters (performance and fairness for flights and passengers simultaneously) leading to operationally more acceptable trade-offs.

VI. Conclusions and further work

In this paper, arrival delay due to ATFM regulations has been assigned as the optimisation of a multi-objective problem considering system performance (total delay) and fairness (deviation with respect to RBS solution). This has been done considering the perspective of flights and passengers. Two optimisation cases have been discussed: one where flights metrics are considered and another one focused on passengers indicators. A new definition for passenger fair assignment of delay has been suggested as the deviation with respect to the delay experienced in an RBS situation.

The concept of *price of fairness*, as the relative loss of performance due to the consideration of fairness in the optimisation process, and *price of efficiency*, as the loss of fairness due to the weight of efficiency on the optimisation, have been defined for both cases.

A specific traffic scenario has been modelled in order to test these concepts, consisting in considering operations at a hub including reactionary delay and passenger connections. The optimisation trade-off analysis between performance and fairness shows that unless fairness is not reduced significantly in the optimisation preference, the solution obtained is close to an RBS one. Then trade-offs arise. The optimisation presented indicates how passenger efficiency, i.e., the delay experienced by passengers, can be significantly improved if they are considered in the optimisation with relatively small impact on flight delay. However, if total flight delay wants to be reduced with respect to RBS, then the reduction of passenger delay is limited.

Passenger metrics have a larger range in terms of fairness than flight metrics, i.e., changes to improve passenger performance lead to quick changes in fairness. But, it is also possible to improve some of the passenger metrics (e.g., delay) without having a negative impact on flight ones with respect to current operations. Therefore it might be worth considering an optimisation where both stakeholders are simultaneously taken into account. When only considering flight metrics, it might be difficult to assess all the underlying trade-offs. Data availability for passenger metrics and optimisation is however an issue.

In this work, fairness has been defined as the total deviation from RBS, but information on how this deviation is distributed among flights, airlines or type of passengers, could also be considered and studied in the future. Additional airports and regulations should also be modelled to analyse trade-offs between flight and passenger delay and fairness in a more general manner. Collaborative mechanisms to obtain the required information for the optimisation process (passengers number, connection times, etc.) should be considered in order to facilitate the model operational implementation. Finally, the impact of uncertainty on the different identified trade-offs should be analysed.

REFERENCES

Ball, M., Barnhart, C., Nemhauser, G., and Odoni, A. (2007). *Chapter 1 Air Transportation: Irregular Operations and Control*, volume 14 of *Handbooks in Operations Research and Management Science*, pages 1–67. Elsevier.

Barnhart, C., Bertsimas, D., Caramanis, C., and Fearing, D. (2012). Equitable and efficient coordination in traffic flow management. *Transportation Science*, 46(2):262–280.

Bazargan, M., Fleming, K., and Subramanian, P. (2002). A simulation study to investigate runway capacity using TAAM. In *2002 Winter Simulation Conference*.

Bertsimas, D., Farias, V. F., and Trichakis, N. (2011). The price of fairness. *Operations Research*, 59(1):17–31.

Bertsimas, D., Farias, V. F., and Trichakis, N. (2012). On the efficiency-fairness trade-off. *Management Science*, 58(12):2234–2250.

Bertsimas, D. and Gupta, S. (2016). Fairness and collaboration in network air traffic flow management: An optimization approach. *Transportation Science*, 50(1):57–76.

Carlier, S., de Lépinay, I., Hustache, J., and Jelinke, F. (2007). Environmental impact of air traffic flow management delays. In 7th *USA/Europe air traffic management research and development seminar (ATM2007)*.

Cook, A., Delgado, L., Tanner, G., and Cristóbal, S. (2016). Measuring the cost of resilience. *Journal of Air Transport Management*, 56 Part A:38–47.

Cook, A., Tanner, G., Cristóbal, S., and Zanin, M. (2012). Passenger-oriented enhanced metrics. In 2nd *SESAR Innovation Days*.

Delgado, L., Cook, A., Tanner, G., and Cristóbal, S. (2016). Quantifying resilience in ATM, contrasting the impacts of four mechanism during disturbance. In 6th *SESAR Innovation Days*.

Dell'Olmo, P. and Lulli, G. (2003). A dynamic programming approach for the airport capacity allocation problem. *IMA Journal of Management Mathematics*, 14(3):235–249.

EUROCONTROL (2015a). ATFCM operations manual – Network manager. Technical report, EUROCONTROL.

EUROCONTROL (2015b). Performance Review Report - An Assessment of Air Traffic Management in Europe during the Calendar Year 2014. Technical report, EUROCONTROL.

EUROCONTROL (2018). DDR2 Reference manual for generic users. Technical report, EUROCONTROL.

EUROCONTROL (2019). Network operations report 2018. Technical report, EUROCONTROL.

Gilbo, E. P. (1993). Airport capacity: representation, estimation, optimization. *IEEE Transactions on Control Systems Technology*, 1(3):144–154.

Glover, C. N. and Ball, M. O. (2013). Stochastic optimization models for ground delay program planning with equityefficiency tradeoffs. *Transportation Research Part C: Emerging Technologies*, 33:196 – 202.

Hamdan, S., Cheaitou, A., Jouini, O., Jemai, Z., Alsyouf, I., and Bettayeb, M. (2018). On fairness in the network air traffic flow management with rerouting. In *9th International Conference on Mechanical and Aerospace Engineering*.

Jacquillat, A. and Vaze, V. (2016). Interairline equity in airport scheduling interventions. *Transportation Science*, 52(4):941964.

Kuhn, K. D. (2013). Ground delay program planning: Delay, equity, and computational complexity. *Transportation Research Part C: Emerging Technologies*, 35:193 – 203.

Manley, B. and Sherry, L. (2010). Analysis of performance and equity in ground delay programs. *Transportation Research Part C: Emerging Technologies*, 18(6):910 – 920. Special issue on Transportation Simulation Advances in Air Transportation Research.

Marler, R. and Arora, J. (2004). Survey of multi-objective optimization methods for engineering. *Structural and Multidisciplinary Optimization*, 26(6):369–395.

Marler, R. T. and Arora, J. S. (2005). Function-transformation methods for multi-objective optimization. *Engineering Optimization*, 37(6):551–570.

Marler, R. T. and Arora, J. S. (2010). The weighted sum method for multi-objective optimization: new insights. *Structural and Multidisciplinary Optimization*, 41(6):853–862.

Montlaur, A. and Delgado, L. (2017). Flight and passenger delay assignment optimization strategies. *Transportation Research Part C: Emerging Technologies*, 81:99 – 117.

Montlaur, A. and Delgado, L. (2018). Arrival trade-offs considering total flight and passenger delays and fairness. In 8^{th} *SESAR Innovation Days*.

Pilon, N., Cook, A., Ruiz, S., Bujor, A., and Castelli, L. (2016). Improved flexibility and equity for airspace users during demand-capacity imbalance - an introduction to the user-driven prioritisation process. In 6^{th} *SESAR Innovation Days*.

Vossen, T., Ball, M., Hoffman, R., and Wambsganss, M. (2003). A general approach to equity in traffic flow management and its application to mitigating exemption bias in ground delay programs. *Air Traffic Control Quarterly*, 11(4):1–11.

Vranas, P. B., Bertsimas, D. J., and Odoni, A. R. (1994). Dynamic ground-holding policies for a network of airports. *Transportation Science*, 28(4):275–291.

Young, H. P. (1995). *Equity In Theory and Practice*. Princeton University Press.