An analysis of trajectory centered ATFM with airspace users' preference scores

Yuhan Jiang
Management School
Lancaster University
Lancaster, UK
y.jiang24@lancaster.ac.uk

Luigi De Giovanni
Dip. di Matematica
Università di Padova
Padua, Italy
luigi@math.unipd.it

Carlo Lancia
DISCo
U. di Milano Bicocca (UNIMIB)
Milano, Italy
carlo.lancia@gmaill.com

Guglielmo Lulli
DISCo, U. Milano-Bicocca
guglielmo.lulli@unimib.it,
g.lulli@lancaster.ac.uk

***Abstract*—In this paper, we compare two trajectory-centered Air Traffic Flow Management (ATFM) architectures that assign a trajectory and a departure time slot to each flight. The main feature of these architectures is the use of trajectories that are feasible from the operational point of view and are aligned as much as possible with airspace users' preferences. The first architecture is inspired by the first-come-first-served principle and can be seen as a proxy of the approach currently used in practice. The second one—herein named "preference-aware ATFM architecture"—considers alternative trajectory options and explicitly takes the related preference scores into account: its main objective is to provide better trade-offs between user-preferred trajectories and system efficiency. The analysis herein carried out highlights the benefits of the latter architecture in terms of both system efficiency and satisfaction of airspace users' preferences, thus paving the way to its potential use as a tool for the ATFM collaborative decision-making process.**

Keywords- ATFM; preference scores; trade-off analysis

I. Introduction

The main objective of ATFM is to prevent local demand-capacity imbalances by adjusting the flows of aircraft on a national or regional basis. This function is carried out by the aviation authority (e.g., the Network Manager in Europe or the Air Traffic Control System Command Center in the US) in collaboration with the relevant stakeholders, including airspace users. Dedicated tools, mostly based on the First Come First Served (*FCFS*) rule (also known as the ration-by-schedule principle), are used to support the decision process within the ATFM function. However, due to the "local" nature of the rules underpinning most of these tools, the ATFM solutions implemented to mitigate air traffic system constraints and/or impacting events are often largely suboptimal (e.g., see Ruiz et al. [1]). To overcome this issue, several mathematical models have been developed to identify better and possibly optimal control strategies. These models are designed to provide a holistic view of the whole air traffic system, i.e., at the continental level in Europe and the national level in the US. To tame the related computational challenges, these models are often "macroscopic", see [2] and references therein cited. That is, they do not capture all the tactical details of conditions in every part of the airspace during the day. In the last decade, attention has been progressively turned towards models that are suitable for trajectory-based operations. One of the early models that explicitly considered 4D trajectories for each flight is the model proposed by Sherali et al. [3]. More recently, Dal Sasso et al. [4, 5] explicitly modelled flight-level information in a compact formulation. However, these models have to deal with two major challenges related to the output trajectory. First, it has to be "feasible", i.e., viable from the operational standpoint. Second, it should reflect the airspace user's desiderata, i.e., it should align with the airspace user's business model.

To address these challenges, in this paper, we first review an architecture to compute trajectory options that are acceptable from the airspace users' perspective, meaning both viable and in line with their business model. This architecture includes two main machine-learning modules: one to extract typical *trajectory models*, i.e., trajectories that are usually flown by airspace users; and the second one to compute the preference score for each trajectory model based on key features of the flight and the airspace user. The architecture is completed by a trajectory-centered optimization model that selects a route for each flight, taking explicitly into account the computed preference scores.

The paper analyzes and compares the results of this preference-aware ATFM architecture with an alternative one, whose main engine adopts a *FCFS* discipline. This second architecture, which can be conceived as a proxy of the tool currently used in practice, resolves air traffic congestion by assigning delays to user-requested trajectories. The solutions proposed by the two architectures are compared to each other in terms of system efficiency, measured by total delay and traffic-capacity balance, and adherence to user preferences. The analysis shows the benefits of the preference-aware ATFM, in terms of both efficiency and airspace users' preference satisfaction. Indeed, this architecture allows computing solutions that are reasonably close to the best one in terms of system efficiency without degrading the user's preference satisfaction. On the other hand, our analysis shows the *FCFS* strategy may be largely inefficient. As a by-product, we also evaluate the effectiveness of solutions that minimize delays regardless of preference scores. Although these solutions are, by design, the most efficient, they may suggest trajectories that are far from users' requests and therefore far from being acceptable. The identification of ATFM solutions that can significantly improve the system's efficiency while preserving users' operability is more likely to obtain the consensus of all the stakeholders involved. This promises to facilitate the collaborative decision-making process that leads to agreed flight plans at the pre-tactical stage, somewhat in the same spirit as Xu et al. [6].

The paper is organized as follows. In Section II, we review the analyzed architectures for trajectory-centered. Section III presents the comparative analysis of the solutions output by the two architectures and the potential benefits of the preference-aware ATFM approach. Finally, Section IV draws some conclusions and future research directions.

II. Trajectory-centered ATFM architectures

We here describe the two ATFM architectures that are the object of our comparative study. The key and elementary components of these models are the feasible trajectories, possibly many for each flight. The goal is to assign a trajectory of the given set to each flight and in some cases a delay to meet the airspace capacities. The advantage of using trajectory-centered models—compared to models that generate trajectories by suitably combining network segments—is that they are suitable for reckoning preference scores within the models, since the preference of airspace users for a trajectory is determined by its overall features, which cannot be always decomposed by its segments.

A. *First Come First Served ATFM (Ration-by-schedule)*

The first architecture is inspired by the *FCFS* rule, which currently underpins all the major ATFM tools and initiatives in Europe and the US (see, e.g., [7]): It prioritizes users' preferences so that one trajectory per flight is considered, corresponding to the requested one, even at the cost of larger delays. The architecture assigns the requested trajectory (RT) to each flight and follows an *FCFS* rule to balance demand and airspace capacity (blue blocks of Fig. 1). Flights are processed in order of their scheduled time of departure. The estimated departure time is the result of the minimum delay that is needed to avoid any excess of demand—generated by all the flights that have been already served—with respect to the available capacity at any airport and en-route sector.

B. *ATFM with Preference Scores*

The second approach aims at providing better trade-offs between user-preferred trajectories and system efficiency at the early stage of the collaborative decision-making process, as it selects a feasible trajectory among the ones in a predefined bucket, taking both into account a measure of preference from the airspace user's perspective, and overall requirements on system performance. The preference-aware ATFM architecture is composed of three main components (green blocks of Fig. 1), which are described in the following subsections. Interested readers can find a thorough technical discussion—including a

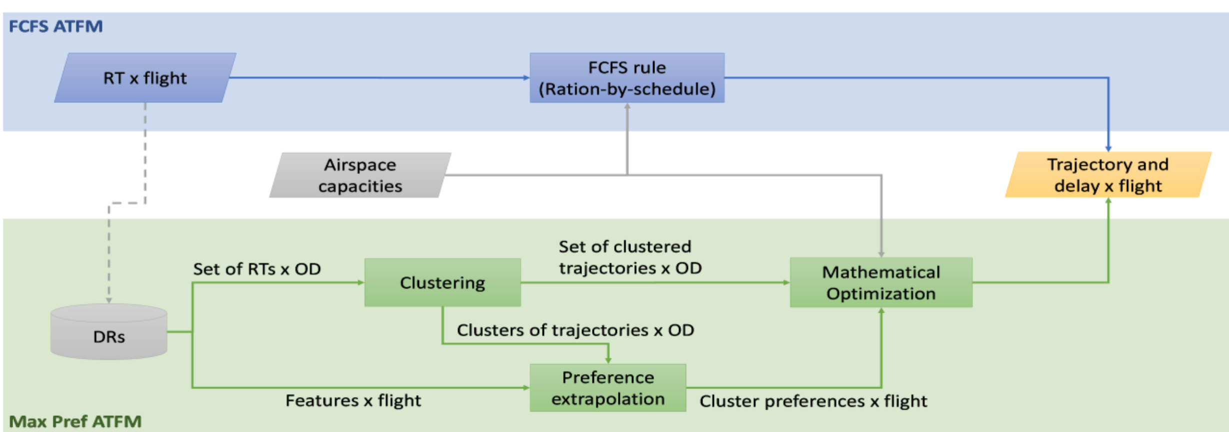

Figure 1. Two trajectory-centered architectures for ATFM.

full description of the Mathematical Optimization module—in [8].

1) *Extraction and Classification of Trajectories*

To obtain, for each flight, a set of trajectories that airspace users are willing to consistently fly, a dataset is extracted from the data repositories (DRs), containing all the stored 4D flight trajectories. Ideally, DRs collect information on the requested trajectories for all flights during a relevant time period. The dataset is then cleaned of outliers, i.e., trajectories that are not consistently flown by airspace users. This can be thought of as trajectories determined by unusual factors (peculiar weather conditions, strikes, airspace closures etc.) and, hence, not likely to be accepted by airspace users in typical situations. To this end, the "Clustering" module of Fig. 1 adopts a density-based technique, namely DBSCAN: based on a measure of the "distance" between trajectories, it clusters together similar trajectories, and it detects outliers by identifying trajectories that are remarkably far from any other trajectory. The clustering step not only allows us to remove outliers but also classifies trajectories into groups of 4D trajectories that are similar to each other in terms of shape (crossed waypoints and related flight levels) and operational speed. An illustrative example of the Rome-Paris origin-destination pair is shown in Fig. 2. We thus associate clusters with distinct trajectory models, in other words, we consider that the elements of a cluster are variants of an underlying trajectory model. We assume that the airspace users' choice depends, at the pre-tactical stage, on the trajectory model,

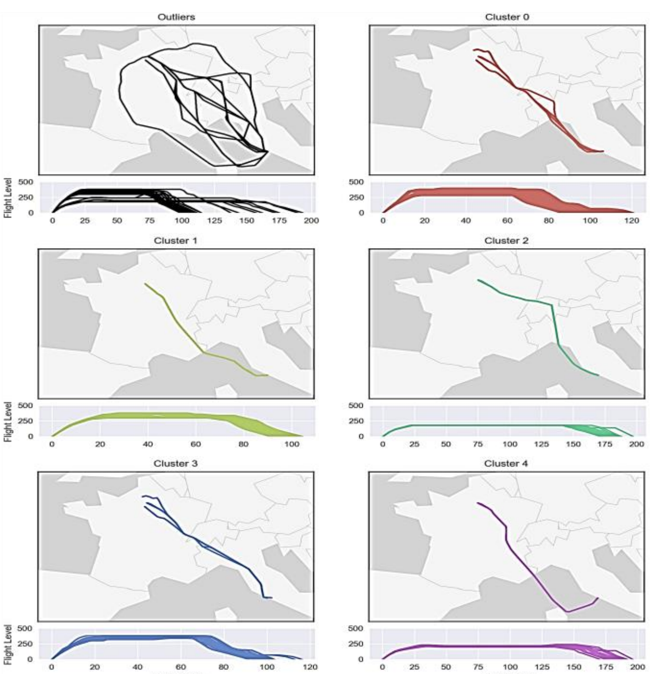

Figure 2. An illustrative example of trajectory-clustering output [10].

rather than on the specific trajectory. Hence, given a flight, the airspace user assigns the same preference score to all the trajectories of the same cluster.

2) *Preference Scores' Computation*

The "Preference extrapolation" module (see Fig. 1) computes the preference scores, i.e., a measure of how much a flight is willing to fly a trajectory option. Preferences are determined by several factors related to both trajectories' and airspace users' features, as well as, possibly, environmental factors (see, e.g., [6]): trajectory length, duration, shape, operational speed, fuel costs, en-route charges, weather conditions, airlines' business model (e.g., legacy air carriers may prefer short routes, whereas some low-cost carriers may lean toward longer routes to avoid high en-route charges) among others. Such determinants, in particular the ones depending on the airspace users, are often only partially known, or even unknown, because of the lack of data sharing, frequently involving confidential business information. However, our purpose is obtaining parameters representing the preference score, i.e., to measure how each flight matches a trajectory as a whole, rather than determining the individual impact of each determinant on preference. Moreover, as observed above, we associate preference scores with the cluster, hence the trajectory model. We thus use cluster membership to approximate the role of all the features depending on the trajectory and their impact on (unknown) airspace users' determinants. This leads to a data-driven approach to preference scoring. As depicted in Fig. 1, the "Preference extrapolation" module receives information on flight features, extracted from DRs, and on trajectory cluster membership. A supervised machine learning classification tool is proposed to discover how flight features are related to cluster membership. In particular, a random forest is trained on the selected features to learn, given a flight, the cluster the trajectory requested for that flight belongs to. A random forest combines different binary tree classifiers in an ensemble (an illustrative example of a binary tree is reported in Fig. 3). Relevant to

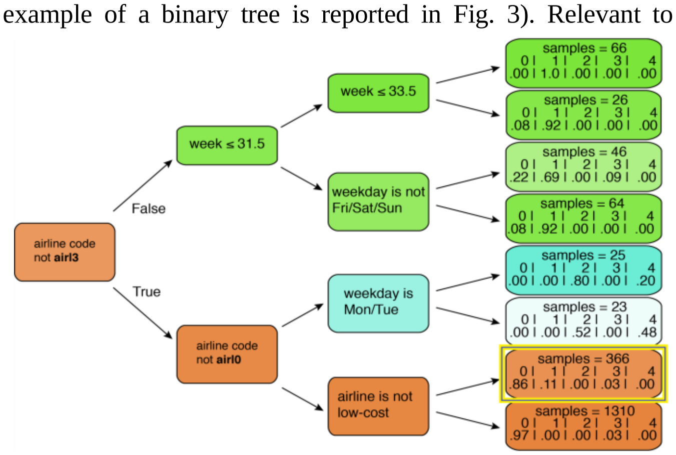

Figure 3. An illustrative binary tree classifier [10].

preference scoring are the figures between 0 and 1 reported in each leaf and representing the association to each cluster. A leaf corresponds to a combination of flight feature values. The association measures the fraction of flights falling in that leaf, and requesting a trajectory in that cluster. We thus use the associations as a measure of the preference assigned by flights to trajectory models. For example, with reference to Fig. 3, 86% of flights that fall in the highlighted leave have requested a trajectory of cluster 0, 11% of flights refer to cluster 1 and the 3% to cluster 3; according to the tree, flights falling in the analyzed leaf will thus have a preference score of 0.86, 0.11 and 0.03 for trajectories of clusters 0, 1 and 3, respectively, and a preference score of 0 for remaining clusters. For each flight, the final preference scores are computed by averaging the associations of the leaves reached in the diverse trees of the random forest.

3) *Optimized Assignment of Trajectory and Time Slot*

The assignment of trajectories and departure time slots to flights is obtained by solving a mathematical optimization model, which is at the core of the architecture. The model selects a trajectory within a given set of options, together with possibly a departure delay to satisfy air system capacities. It is formulated by means of binary decision variables. More specifically, a binary decision variable $x_{p,d}^{f}$ is defined for each flight $f \in F$ (the set of flights), trajectory $p \in P_f$ (the predefined set of 4D trajectory options for flight f output by the trajectory extrapolation module), and departure delay $d \in D_f$ (the set of possible delays associated with the time slots available for flight *f*). The variable takes value one if f flies p with a delay d, 0 otherwise. The objective function maximizes the total preference score, obtained by summing up the preference scores associated with each flight f and the selected trajectory p, denoted by G_p^f. In formula,

$$Max_{x \in X} \sum_{f \in F, p \in P_f, d \in D_f} G_p^f \cdot x_{p,d}^f \quad (1)$$

In the model, an ATFM strategy is represented by a vector of decision variables. The set X of feasible solutions includes a number of constraints that guarantee the satisfaction of the following conditions: (*i*) one trajectory and one (possibly null) ground delay is assigned to each flight; (*ii*) system capacities are satisfied for all the elements of the air traffic system (i.e., airports and en-route sectors) and at any period of time; and (*iii*) the total amount of arrival delays is below a given threshold. It is important to clarify that the total delay is the delay observed at the arrival, so in addition to the ground delay associated with the assigned time slot, it also accounts for either possible delay reduction due to air schedule padding or additional delay incurred during the cruising phase due to flying longer routes. Constraints (*iii*) follow from the existing trade-off between total delay and total preference scores, which makes the proposed model bi-objective. However, in this work, we do not have an interest in investigating this trade-off, i.e., computing the Pareto frontier, but rather computing a Pareto efficient solution to the problem, i.e., a solution maximizes the total preference scores given a certain threshold on the maximum total arrival delay that can be assigned. This threshold is established by ATFM authority in collaboration with the relevant stakeholders. The described model, which is a very large-scale optimization problem, is solved using a custom decomposition method, which can compute near-optimal solutions in short computational times (i.e., order of minutes).

III. TRADE-OFFS' ANALYSIS

We present an analysis of the trade-offs between the delay and the preference score that can be achieved by the proposed ATFM architectures. We consider the following settings:

- *Nominal*: the RT is assigned to each flight with no departure delay. Even if the resulting ATFM solution normally exceeds the airspace capacities, we use it as a baseline for comparison;
- *FCFS*: the RT and departure slot are assigned to each flight by the *FCFS* architecture. Using RTs, this setting provides, by definition, a solution with the best possible overall preference score;
- *Min Delay*: a variant of the mathematical optimization model is solved, where preference scores are neglected, and the objective function is turned into total delay minimization. This setting determines, for the set of trajectory options output by the "Clustering" module, the smallest possible total delay;
- *Max Pref 100*: the optimization model maximizes the total preference scores under a delay budget equal to the smallest possible total delay. It computes a "preference-efficient" solution that minimizes the total delay;
- *Max Pref 110*: same as *Max Pref 100* but considering an additional 10% delay budget. It represents one possible trade-off between total preference score and delay.

These settings have been evaluated on 10 instances of the ATFM problem extracted from the Eurocontrol Demand Data Repository DDR2. Instances correspond to the busiest days (i.e., with the largest number of flights) during Summer 2016, with at least one instance per day of the week. Each instance represents a whole day of operations in the European airspace, and the number of flights ranges from 28,508 to 32,128. Data from DDR2 include airspace configurations and capacities, flights and related trajectories. More specifically, for each flight, DDR2 stores, among others, a 4D description of the last Filed Tactical

Flight Model (FTFM, also called "Initial" or "Model 1"). We remark that FTFM may not correspond to the RT, however, we used this model for preference extrapolation since it is the best approximation of airspace users' choices available to us from DDR2. The random forest has been trained on the following flight features: day of the week, week number (to account for seasonal effects), anonymized airline code, airline type ("legacy" or "low cost") and aircraft model (see [8] for details).

In the *FCFS* setting, to obtain a solution with best possible preference score, we use a trajectory with shortest flight time among the ones with largest preference score for each flight. For the departure slot allocation, the '24-hours' time horizon has been discretized in 288 5-minute time intervals. The same discretization is used to check airspace capacity constraints. The *FCFS* setting is intended to approximate the more complex procedures (and tools) used by network managers to allocate slots, e.g., the Computer Aided Slot Allocation tool [7], which is not available to us.

In the following analysis, we focus on the main features of the proposed solutions while avoiding an excessive level of detail. Therefore, we report the average value of the displayed statistics, computed over the set of instances herein considered, unless otherwise specified. Moreover, we may jointly refer to *Min Delay*, *Max Pref* and *Max Pref 110* as optimization approaches or simply optimization.

A. Analysis of Delays

The bar chart in Fig. 4 displays the average total delay of *FCFS*, *Max Pref 100*, *Max Pref 110* and *Min Delay* settings. We observe that the optimization approach, whether the preference scores are considered (*Max Pref* settings) or not (*Min Delay*), allows for reducing the total arrival delay substantially. On average, the optimization method reduces the total delay by 66% compared to *FCFS*. Indeed, the average value of the total arrival delay drops from almost 1700 hours to 500 hours. It is evident that the optimization approach is able to exploit the schedule

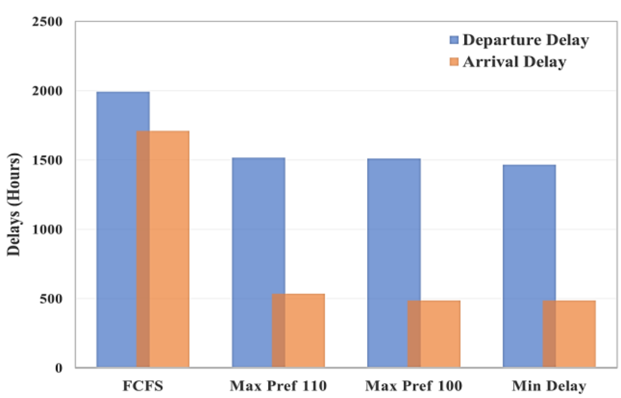

Figure 4. Total average delay comparison.

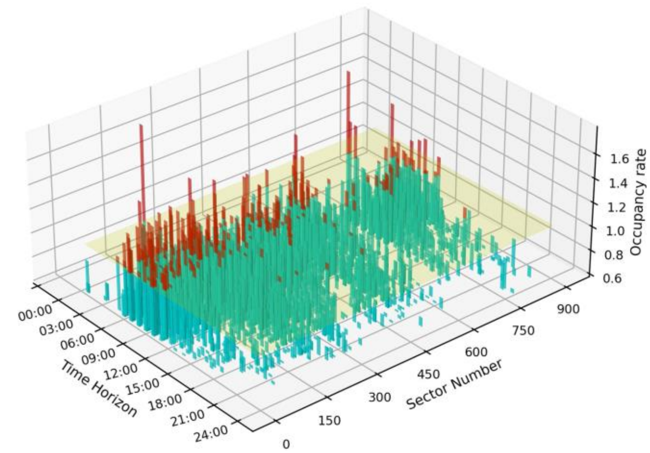

Figure 5. Sectors' occupation rate for the August 28, 2016 instance.

padding associated with some of the available trajectories to provide better punctuality performances. It assigns departure delays—which is in any case 15% smaller than the one assigned by *FCFS*—to flights that have more buffer to recover departure delays. The observed values of total delay are not abnormal (see also Ruiz et al. [1]), and they are obtained with a moderate level of congestion of the European airspace. As an example, the heatmap of Fig. 5 highlights the *Nominal* level of congestion for each sector (on the *y*-axis) and period of the day (on the *x*-axis) for the August 28, 2016 instance. Approximately, 10% (91/920) of the en-route sectors faced congestion, with air traffic demand exceeding capacity by 10% or more. Congestion is happening mostly in the time window between 6 and 10 am (GMT).

The system delay reduction obtained by the optimized approaches compared to *FCFS* is due to a judicious assignment of delay as well as some fine-tuned re-routing of flights. Fig. 6 illustrates both these aspects for sector LECPGXX—which is in the Balearic airspace—. More specifically, Fig. 6 depicts the *Nominal* air traffic demand (orange line), the *FCFS* (blue line)

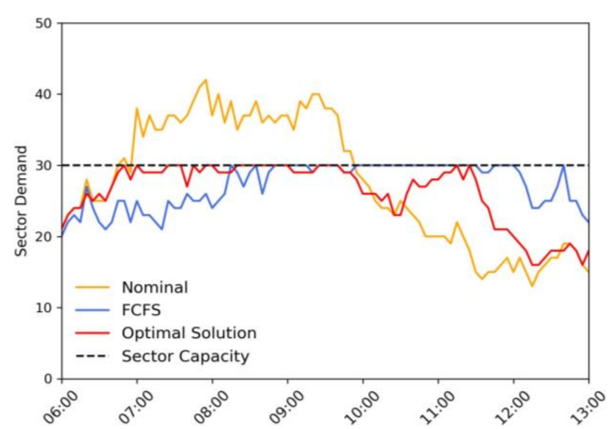

Figure 6. Sectors' occupation rate for the August 28, 2016 instance.

and the *Min Delay* (red line) solution demand for the considered sector between 6 am and 1 pm (GMT) on August 28, 2016. We observe that the *Nominal* sector demand exceeds the capacity limit for an extended period of time. The *Min Delay* solution does not delay any flight in the time periods before the occurrence of congestion. In the time periods of congestion, the number of flights scheduled to fly in the sector is very close to the available capacity. Some flights—originally scheduled to cross the en-route sector during the period of congestion—are rerouted to reduce the (nominal) demand and consequently the amount of assigned delays. On the other hand, *FCFS* shows some perverse behavior. Due to the presence of other congested sectors, some of the demand is delayed even before the congestion takes place. Moreover, during the period of congestion, not all the available capacity is fully exploited thus exacerbating the congestion phenomenon and increasing substantially the amount of delay assigned.

B. Analysis of preferences

A key aspect of the proposed approach is the inclusion of airspace users' preferences. We here show their implications and impact on the ATFM solutions. The stacked bar chart of Fig. 7 describes the distribution of preference levels of the assigned routes for each of the proposed solutions. The stacked bar of the *FCFS* solution represents the distribution of the preference level of the most preferred route. The majority of flights (67.2%) have a preferred route with a preference level greater than 0.8. Only a relatively small percentage of flights (i.e., 6.5%) have a preferred route with a preference level in the interval [0.2, 0.4]. The *Min Delay* stacked bar provides the distribution of the preference level for the routes assigned by optimization model that minimizes the total delay. Comparing the *FCFS* and the *Min Delay* stacked bars, we observe that about 10% of the flights fly a route that is not the preferred one. This corresponds to the amount of rerouting which is needed to minimize the total delay. However, not all the rerouting suggested by the optimal solution is necessary. Indeed, for each instance, there are several and possibly many optimal solutions —all with the same amount of total delay but with a different level of rerouting—. Indeed, the efficient solution denoted as *Max Pref 100* has almost half of the rerouting with respect to the *Min Delay* solution. The result has been obtained without assigning any additional delay. If we accept an increment of the total delay equal to 10%, the amount of rerouting of the corresponding solution (*Max Pref 110*) is further reduced. In this scenario, the number of rerouted flights is 3.3% of the total.

In Fig. 8, we give a more accurate account of the preference level difference between the preferred route and the assigned one for each of the following solutions: *Min Delay*, *Max Pref 100* and *Max Pref 110*. For the *FCFS* solution, there is no such difference because all the flights fly the most preferred route. For each flight, the preference level difference (%) is computed using the following formula:

$$\frac{P_f^* - P_f}{P_f^*} \cdot 100 \qquad (2)$$

where P_f^* is the value of flight *f*'s maximum preference, and P_f is the preference of the trajectory assigned to flight *f*. From the stacked bar chart in Fig. 8, we observe that the majority of rerouted flights in the *Min Delay* solution are assigned to a route with a preference level that is at least 80% smaller than the maximum preference level (i.e., the preference level of the most preferred route). This percentage of flights drastically reduces for the *Max Pref* solutions. Indeed, for the *Max Pref 110*, the percentage of flights that are assigned to a route with a sensibly reduced preference level is slightly higher than 2%.

C. Analysis by Airports and Airlines

We now analyze the distribution of delays across airports and airlines. The objective is to verify if the better performances of the proposed preference-aware architecture are achieved at cost of an unfair distribution of delays, meaning that the proposed

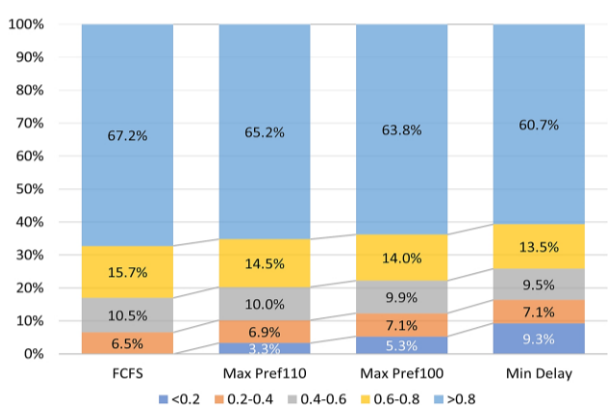

Figure 7. Distribution of the preference levels.

Figure 8. Distribution of the preference level difference (%) between the preferred and the assigned route.

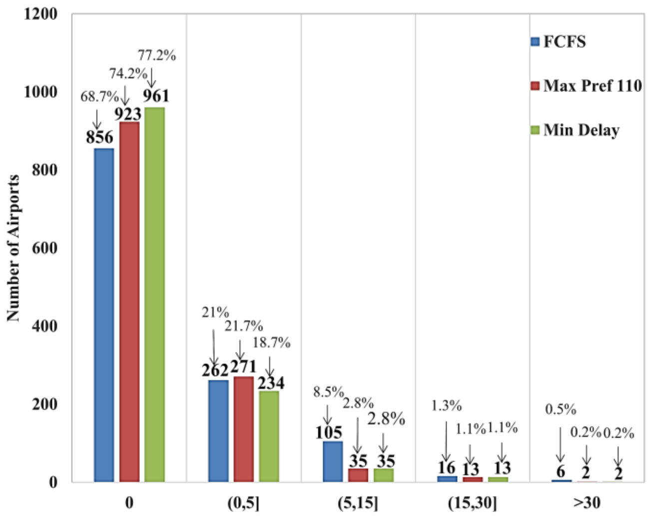

Figure 9. Average arrival delay (in minutes) distribution per airport.

solutions end up in consistently discriminating against certain airlines and/or flights originating from or terminating in certain airports. For the sake of readability of the charts, we display statistics for only three of the settings herein analyzed, i.e., *FCFS*, *Min Delay* and *Max Pref* 110.

In Fig. 9, we display the average arrival delay distribution across airports. Each bar of the chart represents the number of airports with an average arrival delay (measure in minutes) equal to zero, or in one of the following intervals (0, 5], (1, 15], (15, 30], and greater than 30 minutes. We can observe that the *FCFS* distribution has a fatter tail, with a larger number of airports showing larger average arrival delays.

Fig. 10 shows the average arrival delay per flight at the ten busiest airports: the boxplots display the distribution among instances. For all the considered airports, *Max Pref 110* and *Mind Delay* show very similar values. The *FCFS* average arrival delay is larger—and in some cases at a significant level—than the one assigned with *Max Pref 110* and *Min Delay*, with the only exception of Airport 2. At Airport 2, the average arrival delay is similar in all three settings here considered, even though in some instances *FCFS* shows better punctuality performance. In Fig. 11, we report the percentage of delayed flights at each of the 10 busiest airports. We use a stacked bar to highlight both the fractions of flights with small (up to 15 minutes, in blue) and large (greater than 15 minutes, in orange) delays. With the only exception of Airport 2, the optimization approaches provide a smaller percentage of delayed flights, especially in the *Min Delay* setting. Moreover, it is important to observe that the optimization approach consistently reduces the number of flights with large delays, even for Airport 2. Both Fig. 10 and Fig.11 show that, overall, optimization smooths differences between airports and improves fairness compared to *FCFS*.

As far as the distribution across airlines, we observe similar trends. Fig. 12 reports the average arrival delay per flight for the 10 largest (in terms of number of movements recorded in our dataset) airlines. As expected, the optimization provides lower values of the average arrival delay than *FCFS*, and none of the airlines would be better off with the *FCFS* solution. Indeed, some airlines have a large improvement in their punctuality. Moreover, as depicted in Fig. 13—showing the percentage of (slightly and sensibly) delayed flights—the *FCFS* setting not only leads to a greater number of delayed flights, but also with larger delays. Indeed, the orange bars are uniformly and, in some cases, substantially bigger in the *FCFS* setting. Again, optimization seems to improve fairness with respect to *FCFS*, showing reduced differences between airlines.

IV. CONCLUSIONS

To improve the ATM system's performance, the importance of sharing and using trajectory preference information at an early stage of the decision process is widely recognized. This is seen as a key enabler for the full implementation of innovative ATM initiatives, like the TBO (trajectory-based operations) concept. In this paper, we have shown the potential benefits of using such information, and we have analyzed the existing trade-off between system efficiency and airspace users' preferences.

To overcome the information sharing issue, we have considered an ATFM architecture made of both machine-learning and optimization modules. Machine learning is used to approximate the set of feasible trajectories and related airspace users' preference scores. The optimization model aims at finding the overall best preferred ATFM decisions while guaranteeing a minimum system performance level.

Our study, on a set of realistic instances, shows that the proposed architecture computes (delays) efficient solutions with just 6% of flights diverted from the most preferred trajectory. This fraction can be further reduced if some additional delay is allowed. Moreover, the performance of the preference-aware architecture compares very favorably with the two ends of the ATFM interventions spectrum, i.e., *Min Delay* and *FCFS* respectively. Indeed, the preference-aware ATFM solutions retain most of the system efficiency, i.e., with total delay at the level of the *Min Delay* solutions, and achieve large preference scores close to the *FCFS* architecture. Finally, a preliminary fairness analysis suggests that the optimization approach reduces bias in the computed solutions. From the practical point of view, this means that the preference-aware ATFM is able to satisfy the air traffic demand by assigning a moderate amount of delays and/or rerouting. In the case of rerouting, the aim is to assign a new route that is acceptable from the airspace user's perspective. Therefore, the preference-aware ATFM architecture has the potential to facilitate the ATFM

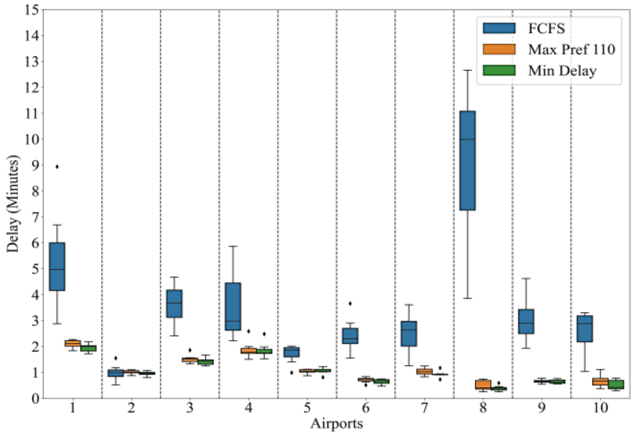

Figure 10. Average arrival delay (in minutes) for the 10 busiest airports.

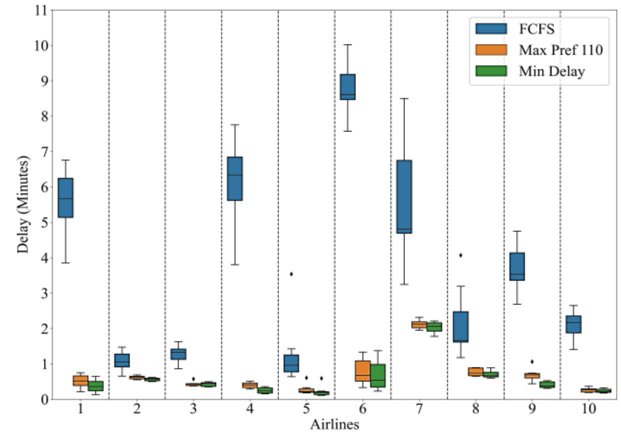

Figure 12. Average arrival delay (in minutes) for the 10 largest airlines.

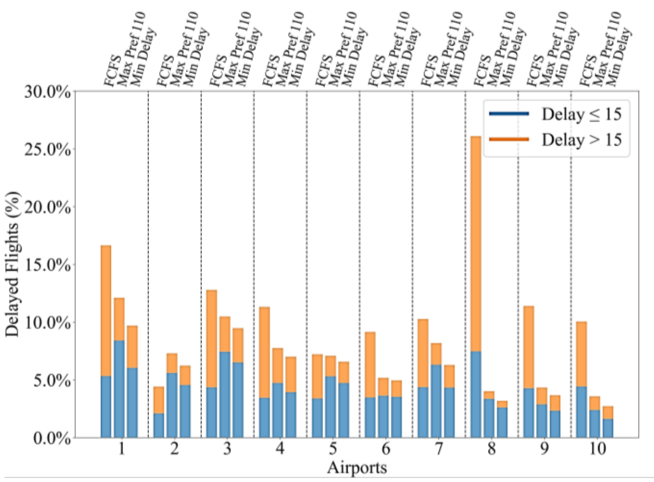

Figure 11. Percentage of delayed flights at the 10 busiest airports.

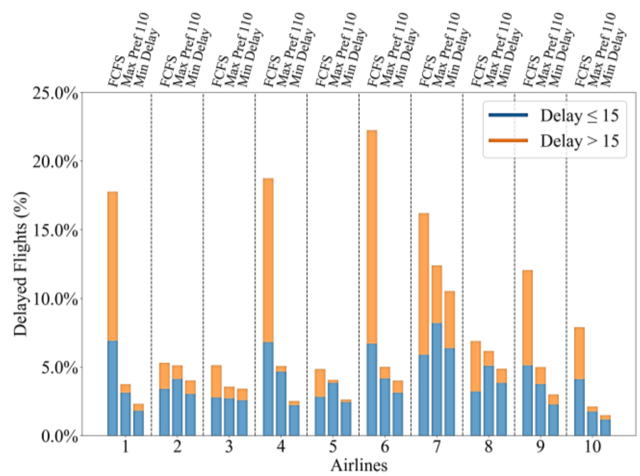

Figure 13. Percentage of delayed flights of 10 largest airlines.

collaborative decision-making process for the identification of agreed flight plans at the pre-tactical stage.

As future research, further assessment of the proposed analysis may be achieved by additional experiments and comparison to the more complex tools currently used by network managers. Moreover, we plan to explore in more detail the equity/fairness aspects of the model's solutions and potentially modify the model accordingly.

REFERENCES

[1] S. Ruiz, H. Kadour, P. Choroba, "A novel air traffic flow management model to optimise network delay", 13th USA/Europe Air Traffic Management Research and Development Seminar, pp. 1–10, 2019.

[2] H. Balakrishnan, B. Chandran, "Optimal large-scale air traffic flow management, unpublished (available at http://tinyurl.com/4um89w6r).

[3] H. D. Sherali, J.C. Smith, A.A. Trani, "An airspace planning model for selecting flight-plans under workload, safety, and equity considerations", Transp. Scie., 36 (4), pp. 387–397, 2002.

[4] V. Dal Sasso, et al, "Incorporating stakeholders' priorities and preferences in 4D trajectory optimization", Trans. Res. Part B, 117, pp. 594–609, 2018.

[5] V. Dal Sasso, et al, "Planning efficient 4D trajectories in air traffic flow management", European J. of Op. Research, 276(2), pp. 676–687, 2019.

[6] Y. Xu, et al, "A framework for collaborative air traffic flow management minimizing costs for airspace users: Enabling trajectory options and flexible pre-tactical delay management", Transp. Res. Part B, 134, pp. 229-255, 2020.

[7] A. Tibichte, M. Dalichamp, "ATFM Modelling Capability – AMOC". Technical report, Eurocontrol, EEC Note No 28/97, 1997 (available at http://tinyurl.com/39j6zsxm).

[8] Lancia, C., De Giovanni, L., & Lulli, G. (2024). Data-driven optimization for air traffic flow management with trajectory preferences. Transp. Scie., 58(2), 540-556, 2024.

[9] L. Delgado, "European route choice determinants", Eleventh USA/Europe ATM Research and Development Seminar, 2015.

[10] C. Lancia, L. De Giovanni, G. Lulli, "Data Analytics for Trajectory Selection and Preference-Model Extrapolation in the European Airspace", In: B. Fortz, M. Labbé (eds.), OR Proceedings 2018, pp. 563–570, 2019.